

- Navigate linked references
- Download citations
- Explore related articles
- Search keywords

Atomistic Simulations of Activated Processes in Materials

G. Henkelman

Department of Chemistry and the Institute for Computational and Engineering Sciences, University of Texas, Austin, Texas 78712; email: henkelman@utexas.edu

Annu. Rev. Mater. Res. 2017. 47:199-216

First published as a Review in Advance on February 27, 2017

The Annual Review of Materials Research is online at matsci.annualreviews.org

https://doi.org/10.1146/annurev-matsci-071312-121616

Copyright © 2017 by Annual Reviews. All rights reserved

Keywords

activated processes, rare-event dynamics, transition state theory, surface chemistry, diffusion, solid-solid phase transitions

Abstract

Activated processes in materials are important for many of the properties that make them function. Batteries and catalysts are examples for which understanding how the component materials function on a timescale of milliseconds to seconds is critical to making improvements in a rational way. Modeling materials over these long timescales, relative to the timescale of atomic vibrations, is one of the grand challenges of the field. Transition state theory is central to bridging this timescale gap, and in the materials community, the harmonic approximation and the determination of saddle points to quantify reaction rates are ubiquitous. In this review, single- and doubleended methods for saddle point finding are discussed, as well as how finding saddle points can be used in the adaptive kinetic Monte Carlo method to model materials properties on the timescale of activated processes. Applications of surface diffusion and chemistry, phase boundary migration, and solid-solid phase transitions are presented.

1. INTRODUCTION

MD: molecular dynamics

PES: potential energy surface

State: an associated set of points on the potential energy surface

Minimum energy

path: a path that follows the gradient from minima across saddle points; also termed an intrinsic reaction coordinate Many of the atomic processes that govern the interesting properties of materials are activated and slow with respect to the timescale of atomic vibrations. One reason for this is that many materials properties that we care about occur on the human timescale of seconds or longer. A suitable battery material, for example, should be able to charge and discharge on a timescale of minutes. Catalysts are needed to accelerate reaction rates that have turnover frequencies slower than 1/s because that is the timescale that we are willing to wait for.

A central challenge for materials modeling is that a direct simulation of atomic-scale dynamics is limited by a time step discretization of femtoseconds. Even with a billion time steps, a molecular dynamics (MD) simulation is limited to microseconds. When the electronic structure is needed to accurately describe interatomic interactions, simulating even nanoseconds of dynamics is an effort. **Figure 1** shows the gap between what we can simulate directly with MD and the activated processes of interest for materials.

1.1. Classes of Potential Energy Surfaces

An alternative strategy for understanding the mechanism and timescale of activated processes in materials is to explore the potential energy surface (PES) directly. There are many computational methods for this in part because there are different classes of energy landscapes and different types of questions that can be asked about kinetic properties. Examples from two dimensions in this space of variables are illustrated in **Figure 2**, specifically (*i*) how smooth or rough the PES surface is and (*ii*) whether the final state of the activated process is known or whether it should be determined over the course of the calculation.

In the simplest class of system (shown in **Figure 2***a*), the energy landscape is smooth, and the transition between a modest number of (perhaps even two) states is desired. When both the initial and final states are known, methods such as the nudged elastic band (NEB) (1–3) or the closely related string method (4) can be used to determine the minimum energy path between reactant and product(s). For increasingly rough potentials, in the examples of a small peptide (**Figure 2***b*) and a model two-dimensional potential (**Figure 2***c*) (5), no single saddle captures the mechanism or rate of the reaction. These cases require double-ended sampling methods, such as path sampling (6, 7), milestoning (8), forward flux sampling (9, 10), transition interface sampling (11), the finite temperature string (5), and discrete path sampling (12–14). However, in the case of protein folding (such as a known crystal structure) (**Figure 2***d*), the many local minima and long transition timescales can be an overwhelming obstacle, even for the best double-ended methods available.

For the systems shown in **Figure** *2e***–***b*, there is a different quantity of interest. Instead of one being concerned about the mechanism and rate of a reaction to a known final state, the starting point is an initial state, and the question is what happens to the system over time. **Figure** *2e* shows



Figure 1

Timescale gap between what can be simulated directly with molecular dynamics and the activated processes that govern many materials properties. MD denotes molecular dynamics.



Classification of rare-event systems based upon the roughness of the potential energy surface and whether the final state is known. When the final states are known (a-d), double-ended methods are appropriate for determining the reaction mechanism and rate (e.g., the nudged elastic band in panels *a* and *b* and transition path sampling in panels *c* and *d*). When the final state is not known (e-b), single-ended methods are appropriate and can be used to determine the final state of the system as time evolves. In panels *e* and *f*, accelerated dynamics methods and adaptive kinetic Monte Carlo are effective tools. Methods for finding rates to unknown final states for rough landscapes as in panels *g* and *b* are still an open area of research.

that atoms are deposited on a surface, and the questions are how they diffuse and ripen and what islands might form. The kinetically stable island shapes are not known at the outset of the simulation but are what one might be interested in learning. **Figure 2f** shows that there are a set of reactant molecules on a surface, and the question is what products will form. The mechanism and rate of reactions are also interesting, but the products and the active site(s) are particularly important. For surface reactions in the presence of a solvent (**Figure 2g**), the problems of rough landscapes become increasingly important. Finally, **Figure 2b** shows the grand challenge [e.g., the Critical Assessment of protein Structure Prediction (CASP) competitions; **http://predictioncenter.org/**] of starting with a protein sequence and trying to predict how it will fold over long times and on a rough potential surface.

The challenges associated with rough landscapes are particularly difficult in soft matter, in which there may not be a clear separation between enthalpy barriers of activated processes and the thermal energy. This review focuses on hard materials and surface chemistry, in which bondbreaking events typically have barriers significantly larger than the thermal energy. For example, a catalytic reaction with an energy barrier of 0.5 eV, which is $20 \times k_BT$ at room temperature, will be active on a timescale of milliseconds. A separation of timescales between that of molecular vibrations and the rare events of interest allows for the use of methods based on transition state theory (TST). Established methods are routinely used for systems for which initial and final states of reactions are known, as in **Figure 2***a*. More recent methods that can be used to map out diffusion and reaction pathways for which the appropriate final state is not known, as in **Figure 2***e*–*g*, are also discussed.

TST: transition state theory



Model two-dimensional potential showing a transition state (*dashed line*) and the set of saddle points on the boundary of the initial state that can be used to estimate escape rates within the harmonic approximation to transition state theory.

1.2. Transition State Theory

The clearly defined potential well in **Figure 2***a* is shown again as a contour plot in **Figure 3**. If the thermal energy is on the order of the contour spacing, then the rate of escape from the initial state will be an activated process. The mechanisms and rate of escape are therefore not easily modeled directly with MD, because the trajectory would spend the majority of time in the red region at the bottom of the well. TST (15, 16) provides a solution to this problem by describing the rate of escape from the initial state as an equilibrium flux through a transition state (TS) surface. If one defines the initial state as all points in configuration space that minimize to the initial state minimum, then a good TS would be the boundary of that initial state. In practice, it is not usually possible to find an analytic expression for such a TS, and so approximations are necessary. **Figure 3** illustrates a suitable surface at $x = x^{\dagger}$. Note that only part of the TS surface is shown and that formally the TS should enclose the initial state so that any reactive trajectory must cross the TS. In practice, however, high-energy regions can be omitted from the TS if their statistical weight is insignificant. Given a TS that bounds the initial state, one can use statistical mechanics to calculate the equilibrium flux through the surface as

$$k_{\rm TST} = \frac{1}{2} \langle \delta(x - x^{\dagger}) | v_{\perp} | \rangle_{\rm IS}, \qquad 1.$$

where $\langle ... \rangle_{IS}$ is an ensemble average over the initial state (IS), including the TS boundary; the $\delta(x - x^{\dagger})$ terms gives the ratio of the TS partition function to the reactant; and $|v_{\perp}|$ is the magnitude

of the velocity normal to the TS surface. The factor of 1/2 is to consider only the outgoing flux through the surface in the rate of escape from the initial state.

An assumption of TST is that every trajectory crossing the TS surface corresponds to a single reactive event from the initial state to a final state. In practice, primarily due to limitations in the construction of good TS surfaces, trajectories will recross the TS surface and either fail to be reactive or recross before a reaction occurs. Because all reactive trajectories must cross the TS surface at least once, and TST counts each crossing as reactive, k_{TST} is an upper bound on the true rate. This is a powerful principle because it means that the TS can be variationally optimized to lower the TST rate to approach the true rate (17). Additionally, one can compute a transmission coefficient κ by using short trajectories from the TS and counting the ratio of crossing points to reactive trajectories. The true escape rate is then $k = \kappa k_{\text{TST}}$ (18, 19).

In the fields of computational surface chemistry and material sciences, the vast majority of calculations do not rely on Equation 1 to evaluate reaction rates both because determining a good TS can be challenging and because sampling is computationally expensive, especially when based upon electronic structure methods. Instead, an enormous simplification to TST-the harmonic approximation—is most commonly used. The idea behind the harmonic approximation to TST (HTST) is that there are bottleneck regions on the TS surface that dominate the rate in Equation 1. These bottlenecks, or reaction channels, are characterized by first-order saddle points on the PES (Figure 3). The saddle points are the lowest-energy points that need to be crossed (locally) to escape the initial state. The single unstable mode at each saddle defines the local normal to the TS, and the TS is approximated as a hyperplane with a harmonic potential given by the N-1 remaining stable modes (indicated by a purple line at each saddle in Figure 3). It is important to understand that the HTST is used to calculate the escape rate via each local channel characterized by a saddle point. The total rate of escape from the initial state can be approximated by summing over all escape channels. In this way, HTST can be viewed as approximating the TS by a set of hyperplanes passing through each saddle on the boundary of the initial state. Also, although it is common to hear the terms saddle point and transition state used interchangeably, this association is appropriate only in one dimension. In high-dimensional systems, the harmonic TS is the hyperplane that passes through the saddle with a normal along the unstable mode, and more generally in TST, the TS is an N-1-dimensional surface that divides the initial state from product states.

The HTST approximation allows for an analytic evaluation of Equation 1,

$$k_{\rm HTST} = \frac{\prod_{i=1}^{N} v_i}{\prod_{i=i}^{N-1} v_i} e^{-\Delta E/k_{\rm B}T},$$
2.

where v_i and v_j are the real normal mode frequencies at the minimum and saddle point, respectively, and ΔE is the potential difference between the saddle and minimum (20). There is one fewer mode at the saddle because the unstable mode does not contribute to the HTST rate. Equation 2 is particularly attractive because the HTST calculation is based upon only local properties of minima and saddle points and does not require additional sampling of the PES. The exponential term, which contains the energy barrier, typically dominates reaction rates. The vibrational prefactor is often approximated as 10^{13} /s for elementary reactions because this is a characteristic atomic vibrational frequency. The prefactor can also be calculated explicitly to capture entropic contributions to the rate within the harmonic approximation.

A common misconception associated with HTST is that it is a zero-temperature approximation, perhaps because no finite-temperature MD calculations are required to evaluate Equation 2. Clearly, though, Equation 2 is temperature dependent and can provide highly accurate rates for temperatures where the harmonic approximation holds. More precisely, HTST will be accurate

for reactions that follow an Arrhenius rate dependence in which the entropic contributions to the rate are temperature independent.

Within the HTST approximation, the most challenging aspect of calculating a reaction rate is not evaluating Equation 2 but, especially in high-dimensional systems, finding the saddle points that characterize reaction pathways from the initial state of interest.

2. SADDLE POINT FINDING METHODS

Saddle point finding methods can be classified as single- or double-ended algorithms. The doubleended methods, which are most widely used in material science, require specification of both the initial and final states and aim to find a saddle connecting them through a minimum energy path. Single-ended methods require only an initial geometry. Single-ended methods are typically started either within an initial state to find a mechanism of escape or near a TS, with the aim of rapid convergence to a saddle. Single- and double-ended methods have both strengths and weaknesses, and both tools are valuable when a PES is being explored.

2.1. Double-Ended Methods

An example of a double-ended method is the NEB (1–3, 23), which is illustrated in **Figure 4***a*. An NEB calculation starts with an initial path; this could be a linear interpolation between the initial and final states or a path that is constructed to be closer to the minimum energy path (24). The tangent to the path at the *i*th image, \hat{t}_i , is defined with an up-winding scheme as the direction to the higher-energy neighbor (2). There are two components in the NEB force on each image, F_i^{NEB} : the projection of the real force perpendicular to the tangent, F_i^{\perp} , and a spring force $F_i^{\text{S}\parallel}$, parallel to the tangent. The spring force maintains equal spacing between the images while they relax according to the real force perpendicular to the path. As the NEB force goes to zero, the images lie equally spaced along the minimum energy path, where the perpendicular force is zero.

In the climbing image NEB (CINEB) method (3), the highest-energy image climbs up the minimum energy path to the saddle point. This is accomplished by inverting the component of the true force on the climbing image along the tangent, $F^{\text{CI}} = F - 2F \cdot \hat{\tau}$. Note that the force on



Figure 4

(*a*) The double-ended nudged elastic band (NEB) method and (b,c) the single-ended minimum-mode following dimer saddle point finding method. Abbreviation: MEP, minimum energy path. Adapted with permission from References 21 and 22.

the climbing image highlights a common element of most saddle point finding methods. First, a direction is determined to approximate the reaction coordinate, or the unstable mode at the saddle. Then, the methods move the system up the potential in the reaction coordinate direction while minimizing in all other directions. When a stationary point is reached, it will be at a maximum in the one direction and a minimum in the others—the definition of a first-order saddle. The primary difference between saddle point finding methods is the definition of the direction used to approximate the reaction coordinate at the saddle. In the CINEB method, this reaction coordinate is the local tangent along the path, $\hat{\tau}$. Because the climbing image does not feel any spring forces, it can be viewed as a single-ended method in which the reaction coordinate is defined by the neighboring images. Because the entire path is relaxed collectively, the local tangent converges parallel to the minimum energy path, and the climbing image converges to the highest saddle along the path. Multiple climbing images can be used to resolve paths with multiple saddles. In fact, all critical points along a path can be found by converging local maxima along the path to saddles and local minima along the path to minima. In practice, however, it is generally more efficient to break up multistep paths into separate CINEB calculations.

There are a number of double-ended methods related to the NEB, as well as modifications to it. A closely related approach is the string method, in which images are redistributed along the path via interpolation so that an equal-distance constraint is imposed at every iteration. An advantage of the string approach is that no spring constant is required. In practice, however, the spring forces in the NEB do not limit convergence rates, and the force projection allows for a minimal implementation into codes that have their own optimization functions. A comparison of the performance of the string and NEB methods shows no significant difference in convergence rates for typical calculations in material science (21).

Several modifications to the double-ended methods aim to improve performance in specific systems. In the adaptive NEB (25), first a loosely converged path is calculated, and then images are redistributed with higher density in the region of a saddle. The adaptive approach improves resolution near the saddle without distributing many images along the entire path. One should note, however, that a climbing image will converge precisely to a saddle with a sufficiently defined tangent; a subsequent increase in resolution near the saddle will not improve precision. Additionally, the accurate calculation of a minimum energy path is most efficiently accomplished by following the steepest descent paths from a converged saddle (21).

Another approach is to loosely converge a double-ended method to determine the neighborhood of a saddle point and then switch to a single-ended method to converge efficiently to the saddle with the precision desired (2). The Wales group has shown that using a doubly NEB (26) efficiently locates saddle regions. A subsequent eigenvector following (27) calculation is then an efficient way of finding saddles between the local minima (28). The doubly NEB is particularly helpful for systems with high initial forces because the extra spring forces, in perpendicular directions other than the direction of the true force, favor a short and smooth path during convergence.

Finally, a method to avoid high forces at the start of an NEB calculation is the growing string (29), in which images initiated in the initial and final state minima are pulled together along the minimum energy path during convergence.

This is by no means an exhaustive list of double-ended saddle point finding methods; a review of other methods can be found in the literature (30).

2.2. Single-Ended Methods

A diligent researcher can often determine the mechanism of chemical reactions with a combination of intuition, optimization, and double-ended saddle searches. However, numerous examples show

Minumum-mode (min-mode) following methods:

a class of single-ended saddle point finding methods in which the potential is maximized along the lowest curvature mode and minimized in all other degrees of freedom

Accelerated molecular dynamics (AMD): include parallel replica dynamics, hyperdynamics, and temperatureaccelerated dynamics that unexpected reaction mechanisms can be important to the dynamics of interest (31–38). Relying on intuition is dangerous in these cases because there may not be an indication that important reaction mechanisms are missing. A powerful alternative, in these cases, is single-ended saddle point finding methods that do not require the specification of a final state.

An example of a single-ended saddle point finding method is illustrated in **Figure 4***b*,*c*. In **Figure 4***b*, minimum-mode (min-mode) following searches are started near an initial state minimum. As in the case of the NEB climbing image, a modified force, F^{\dagger} , is used to direct configurations to a saddle point. This specific flavor of min-mode following method, termed the dimer method (22), uses an iterative finite difference optimization of the local minimum curvature direction, $\hat{\tau}$, as shown in **Figure 4***c*. The term dimer refers to a pair of images separated from the current configuration by a finite displacement. Minimization of the dimer energy, while holding the center point fixed, orients the dimer along the lowest curvature mode. The dimer force is defined as the true force with the component along $\hat{\tau}$ inverted. Following the dimer force corresponds to a maximization along the lowest curvature mode and to a minimization in all other directions. Upon convergence to a zero force, the center of the dimer will be at a saddle, and the dimer will be oriented along the negative curvature mode.

As with the NEB and the double-ended methods, there are a number of improvements and alternative implementations of the dimer min-mode following algorithm. Improvements include the use of a single image in the dimer and forward difference (39), increasing the dimer rotation efficiency (40), and the use of superlinear optimizers for both dimer rotation and translation (41). Developed concurrently with the dimer method, the hybrid-eigenvector method following uses a Rayleigh-Ritz approach to approximate the lowest curvature (27). A recent comparison between these two methods shows that they are in fact equivalent and have no significant difference in performance (28). A related approach, the activation-relaxation technique nouveau, uses the Lanczos method to identify the lowest curvature mode along which the system is driven up the potential while relaxing in perpendicular directions (42).

The single-ended methods are appropriate for converging from an initial configuration to a nearby saddle. Often, however, one is interested in saddle points located on the boundary of an initial state minimum to determine the rate of escape and the possible reaction mechanisms. The specific goal of targeting unique saddles that are on the boundary of an initial state has led to specialized methods. In the κ -dimer method, the isocontour curvature is used to restrain a min-mode following search to the initial state as it converges to a saddle (43). Unique saddles can be targeted by separating initial points and avoiding previous search paths (44). In the biased gradient squared descent method, stationary points are systematically found at increasingly higher potential energy from an initial state minimum (45). The use of high-temperature trajectories with periodic quenching, as in the temperature-accelerated dynamics method (46), is another technique for finding saddles on the boundary of an initial state, which provides a confidence measure in the HTST rate of escape (47). These and related techniques facilitate the exploration of reaction events available to a system from an initial state minimum and, as discussed next, the dynamical evolution of the system.

3. ATOMISTIC SIMULATIONS OVER LONG TIMESCALES

An important class of methods for modeling the dynamics of materials over long timescales is accelerated molecular dynamics (AMD). Voter and coworkers (48) contributed an excellent review of these methods. The central idea of AMD methods is to run one or more dynamical trajectories and use the principles of statistical mechanics to accelerate the frequency of rare events. There are three primary methods for accelerating rare-event dynamics. In parallel replica dynamics, multiple

trajectories are used to sample an initial state so that the escape time can be accelerated roughly by the number of trajectories that are run in parallel (49). In hyperdynamics, a bias potential is introduced to destabilize a reactant state, but not the TS, so that escape times are accelerated by a boost factor determined as the trajectory samples the bias potential (50). Finally, in temperatureaccelerated dynamics, a high-temperature trajectory is used to accelerate possible escape events; the event that occurs and the escape time are extrapolated to the low temperature of interest from the barrier of the escape pathway (46). Many extensions of these methods have been developed over the past 15 years, and these three ideas remain central to efforts to accelerate MD.

An alternative approach to modeling rare events in materials is to explore the potential surface, for example, with saddle point searches, and then use TST as a basis for following the dynamics of the system. An example of this approach is the adaptive kinetic Monte Carlo (AKMC) method.

3.1. Adaptive Kinetic Monte Carlo

Finding the saddle points at the boundary of a reactant state provides sufficient information, for example, using HTST, to determine the escape rate and possible reaction mechanisms. The KMC algorithm can then be used to follow the state-to-state kinetics (51–54).

In KMC, a list of all possible reaction mechanisms and rates must be known at the start of a simulation. From each state visited, a table of possible events is constructed, and one event is selected with a probability proportional to its rate $p_i = r_i / \sum r_j$, where r_i is the rate of event *i*. A single event is chosen stochastically, the system is moved to the final state of that event, and the system clock is advanced by a time step $\Delta t = -\ln \mu / \sum r_j$ corresponding to a time drawn from the exponential distribution of escape times. Here, μ is a uniform random number between 0 and 1. The KMC algorithm is repeated in each state visited to give a state-to-state trajectory.

In AKMC (55), as illustrated in **Figure 5**, the possible reactive events are found with saddle point searches in each new state visited along a trajectory. The method is described as adaptive because the event table can change, or adapt, to the state of the simulation rather than being fixed as in standard KMC. The adaptive aspect of AKMC is important because it allows the simulation to follow unexpected reaction mechanisms; in KMC, no unexpected elementary reaction is possible.



Figure 5

(*a*) In the adaptive kinetic Monte Carlo (AKMC) method, saddle searches are used to find reactive events (P_1 – P_4); the saddles are used to construct an event table from each new state visited, and the KMC algorithm is used to evolve the system from state to state. (*b*) With AKMC, the reactive events can be complex and long range and can involve multiple atoms; they do not need to be anticipated as in standard KMC. (*c*) The efficiency of AKMC simulations can be improved by storing local information about events into a kinetic database from which subsequent saddle geometries can be predicted. KDB denotes kinetic database. Adapted with permission from Reference 33.

AKMC: adaptive kinetic Monte Carlo

Additionally, AKMC facilitates an off-lattice description of atomic positions and can include longrange collective mechanisms that would be difficult to include in a standard KMC event table (33).

KDB: kinetic database

The advantages of AKMC are clear, but there is also a considerable computational cost. A KMC step requires only a couple of random numbers and some bookkeeping to keep track of changes to the state of the system and the event table. Typically, one can do millions or billions of KMC steps in a simulation. In AKMC, one must do saddle searches on the PES at each new state visited. Particularly when atomic interactions are evaluated with density functional theory (DFT), saddle point searches can be much more expensive than an entire KMC simulation. It is critical, therefore, to have an efficient parallel implementation of AKMC, such as in the EON code (56). It is also valuable to save information about saddles that are found so that they can be used later in the simulation or in subsequent simulations. At least two strategies are effective for reusing saddle information. In the first, termed recycling, every saddle in the event table at the start of a KMC step is updated to incorporate the selected event (37). The resulting set of saddle configurations are reconverged for the next KMC step. Events in the table that are spatially separate from the one selected converge rapidly so that in the large-system limit, new saddle searches are required only in the location of the selected event (37). The second approach is to store events in a database. In the kinetic database (KDB) (57), local geometries of initial, saddle, and final states for each mechanism found are saved so that they can be matched to any subsequent state visited.

The use of AKMC has increased in recent years, facilitated by the increase in the number of computational cores available in modern computers. There are now a few different implementations and variations on the algorithm. In the kinetic activation-relaxation technique (58, 59), local events are stored as graphs, rather than as geometries. The strength of this approach is that a known event can be efficiently matched to a new geometry by using existing subgraph matching algorithms (60). Alternatively, a grid-based approach to pattern recognition of local environments avoids the explicit definition of bonds and retains the fast matching property (61). A weakness of these methods is that graph edges or grid entries are discrete, so there is a trade-off between the size of the environment used to describe an event and the number of graphs in the database. A method described as on-the-fly KMC is an AKMC algorithm for solids, which focuses on defects in the crystal as objects (62). The off-lattice positions of the atoms in each defect become objects for which kinetic events are learned. The self-learning KMC is an AKMC method focusing on surface diffusion (63, 64). In self-evolving atomistic KMC (65), emphasis is placed upon activated volumes in which local events are discovered. Given the relative simplicity of the AKMC method and the range of systems that it can be applied to, it is not surprising that variations are being developed and optimized.

3.2. Coarse Graining

The AKMC method extends the timescale of atomic simulations from that of atomic vibrations to that of state-to-state kinetics. An ongoing challenge for AKMC and AMD methods is another timescale gap between the fast events available to a system and the slower events of interest. In catalysis, for example, diffusive events are typically much faster than the bond-breaking events that are most important for the function of the catalysts. In terms of the PES, this situation is described as the small-barrier problem, in which events with low barriers occur exponentially more frequently than those with high barriers. The separation of timescales between the fast and slow events often limits the simulation time that can be reached with AKMC or AMD. In some cases, these small barriers can be smoothed over (66) or integrated over in terms of a transition network (67, 68) or the barriers can be safely increased (69), but in general the low-barrier problem is an outstanding research challenge.

An approach that effectively deals with low barriers in an AKMC simulation is Monte Carlo with adsorbing Markov chains (MCAMC) (70). The inputs to the MCAMC method are a set of connected recurring states, one of which is the current state, and the set of escape rates to neighboring adsorbing states. By solving the kinetic master equation, MCAMC provides an escape time to a specific adsorbing state. When there is a set of states with fast rates between them, sometimes described as a superbasin on the PES, MCAMC allows AKMC to escape the superbasin in a single step. Remarkably, there are no additional approximations made by the MCAMC method for any chosen set of recurring states.

Although MCAMC is effective for escaping superbasins containing a modest number of states, it does not entirely solve the small-barrier problem. Going back to the example of catalysis with fast diffusion and slow bond-breaking reactions, one can recognize that there will be a combinatorially large number of states that can be reached by reactant molecules diffusing on the surface. Furthermore, even if these superbasin states could be enumerated, and the time for a chemical reaction reached with MCAMC, one would then have to start all over again in a new large superbasin. Hence, the small-barrier problem remains one of the most significant research challenges for AMD/AKMC methods.

4. APPLICATIONS

4.1. Surface Chemistry and Diffusion in Materials

The AKMC method has been applied to a wide range of material systems. Examples include the dynamics of Cu atoms in a grain boundary showing the presence of local icosahedral motifs in the disorder interface (72), the diffusion of hydrogen in the grain boundaries of Al (73), and the diffusion of water molecules on the surface of ice (74). **Figure 6** shows a calculation to determine the mechanism and timescale by which the A15 phase of Mo transforms into the more stable BCC phase, using an embedded atom method potential (75). In **Figure 6***a*, the propagation of a boundary between the two phases in Mo is calculated over a timescale of 200 µs. The interface



Figure 6

(*a*) AKMC simulations of a transition between the A15 (*red*) and BCC (*blue*) phases in Mo with interatomic forces described using an empirical potential. (*b*) Visualizing the energy landscape as a disconnectivity graph shows that disorder at the interface (*gray*) results in a rough landscape and in a superbasin consisting of numerous states connected by low barriers. (*c*) The number of states in a typical superbasin of interface motion grows combinatorially with the size of the system, showing how significant the low-barrier problem can be for AKMC/AMD methods. Adapted with permission from Reference 71.



AKMC simulations in which interatomic forces are described with density functional theory. (*a*) Pd cluster formation on MgO(100). (*b*) Decomposition of methanol to CO on a Cu(100) surface. Adapted with permission from Reference 77.

between the A15 and BCC phases remains disordered. It is remarkably difficult to understand the elementary steps associated with the phase transition because these disordered atoms move in highly collective and concerted ways. Only after many such elementary steps do some disordered atoms slide into BCC positions while others in the A15 phase become disordered. This is a clear example of a kinetic process that could not be modeled using an on-lattice KMC approach. **Figure 6***b* shows a disconnectivity graph (76) of a layer transition where each point corresponds to the energy of a local minimum and the vertices indicate the TS energies between them. Clusters of states with similar energies and low diffusion barriers between them correspond to superbasins on the potential surface. When the system size is increased, as shown in **Figure 6***c*, the number of superbasin states grows combinatorially, making the simulation nearly intractable.

AKMC simulations can also be based upon forces and energies from DFT (37). **Figure 7** shows two examples. In **Figure 7***a*, isolated Pd atoms on MgO(100) diffuse and form a small cluster on a timescale of 9 ms (35). Remarkably, this tetramer is able to diffuse by rolling along the surface with a barrier lower than that of any smaller cluster (78). In **Figure 7***b*, the decomposition mechanism of a methanol molecule is calculated and found to occur on a timescale of seconds on the Cu(100) surface (77).

4.2. Solid-Solid Phase Transitions

A recent extension to saddle point finding methods has made it possible to study solid-solid phase transitions and reactions in which atomic motion is coupled to changes in the supercell of a material. Earlier work in this area extended the NEB method to allow either the cell degrees of freedom to relax along an atomic transition (79) or the atoms to relax along a transition described by the cell degrees of freedom (80). These methods were unified by the generalized solid-state NEB (G-SSNEB) method (81), which combines both cell and atomic degrees of freedom into a single vector. A Jacobian was introduced to scale the cell degrees of freedom so that the calculated minimum energy path for any concerted solid-solid phase transition would be independent of the supercell chosen. This same generalized configuration vector, containing both the atomic positions and the cell degrees of freedom, provided the basis of the solid-state dimer method (82).



(*a*) Concerted mechanisms for the A15-to-BCC phase transition in bulk Mo found with the solid-state dimer method. (*b*) The solid-state nudged elastic band method reveals that this phase transformation switches to a nucleation-and-growth mechanism in a larger supercell. Adapted with permission from Reference 82.

Figure 8 shows an example of calculations that are made possible with these solid-state saddle point finding tools. The system is the same A15-to-BCC phase transition in Mo that is discussed above, except that, whereas in **Figure 6** AKMC was used to determine atomic transitions that led to phase boundary migration, **Figure 8** shows a phase transition of the bulk material. In **Figure 8***a*, the solid-state dimer was used to find concerted reaction paths between the A15 and BCC phases in a modestly sized supercell. Interestingly, solid-solid phase transitions are an example in which the single-ended dimer method has a distinct advantage over the double-ended G-SSNEB method. A problem with using a double-ended method for solid-solid phase transitions is that the final state, as well as the corresponding identity of every atom between the initial and final states, must be known. Finding the mechanisms in **Figure 8***a* was possible because a relaxation from the saddles leading out of the A15 initial state shows the appropriate atomic ordering in the BCC final state.

Concerted mechanisms for a solid-solid phase transition have an activation energy that scales with the size of the system. For large systems, a nucleation-and-growth mechanism will dominate because the barrier for a local mechanism does not increase with system size. In the case of the A15-to-BCC phase transformation in Mo, the concerted mechanism found by the solid-state dimer was represented in a larger supercell, and the G-SSNEB method relaxed to the lower-energy nucleation-and-growth mechanism in **Figure 8***b*. This complicated pathway would be extremely difficult to find with the solid-state dimer method, in which the product state is not specified. These calculations are an example of how the single- and double-ended saddle point finding methods are complementary tools for determining reaction pathways.

5. DISCUSSION

There is a concerted effort in the material community to strive for accuracy in our description of interatomic interactions and electronic structure. The failure of DFT to accurately predict important material properties, such as band gaps, have pushed the community toward higher levels of theory. Strongly correlated electronic systems, including those found in transition metal oxides,



With limited computational resources, there will necessarily be a trade-off between the accuracy of interatomic interactions and the computational effort that can be used to understand the atomic structure and dynamics of materials. Although there is a considerable effort toward more precise electronic structure methods, there remains an opportunity for tools that improve sampling and dynamics to better understand how materials function. Abbreviations: AKMC, adaptive kinetic Monte Carlo; AMD, accelerated molecular dynamics; COMB, charge-optimized many-body potential; DFT, density functional theory; MD, molecular dynamics; NEB, nudged elastic band; QMC, quantum Monte Carlo; RPA, random phase approximation.

have provided additional pressure to go beyond Kohn-Sham DFT toward methods—including those based on the random phase approximation (RPA), Green's function methods (GW), or quantum Monte Carlo—that can correct for electronic self-interaction and better describe electronic correlation. Alternatively, by including explicit orbital density dependence, many of the problems of Kohn-Sham DFT can be mitigated at a lower computational cost relative to RPA or GW (83). The effort to improve the accuracy of electronic structure calculations is important and necessary for material modeling.

What is less appreciated is that a focus on accuracy of the electronic structure will necessarily result in a reduction in sampling of atomic structure and dynamics. Figure 9 illustrates this point in a qualitative way. Given a fixed amount of computational resources, one can choose to focus on accuracy of the electronic structure or on sampling of atomic structure or try to find some middle ground between these two extremes. Clearly there is no single correct answer in this space; there are problems that require the highest levels of electronic structure, such as the example of band gaps, even if only a single calculation is affordable. In other cases, including the examples mentioned in this review, interesting and unexpected aspects of material properties require sampling of the PES. This exploration may require using DFT with modest system sizes and settings, reactive potentials such as the reactive force field (ReaxFF) (84) and the charge-optimized many-body potential (COMB) (85), or empirical potentials, when available. My feeling is that there is pressure from the community toward improved accuracy and less focus on improved sampling and dynamics. Reviewers will readily point out limitations in basis sets, functionals, system size, k-point sampling, and convergence criterion. A focus on precision, however, comes at the expense of sampling. A precise rate calculated for an assumed reaction mechanism is misleading if there is a lower-energy pathway that was not discovered because precision was emphasized over exploration of the PES. I think there is a challenge, and an opportunity, to have a range of models of varying precision to facilitate a range of sampling methods that allow for the possibility of discovering interesting and important material properties that may not be anticipated.

SUMMARY POINTS

- 1. Computational tools based upon transition state theory are important for understanding the kinetic properties and function of materials.
- 2. Double-ended saddle point finding methods are robust and widely used to find reaction pathways to known product states.
- 3. Single-ended saddle point finding methods provide rapid convergence to nearby saddles and can find reaction mechanisms when the product state is not known.
- 4. Adaptive kinetic Monte Carlo couples saddle point finding methods with kinetic Monte Carlo to model state-to-state dynamics.
- 5. With adaptive kinetic Monte Carlo, reaction mechanisms can be complex, can be off lattice, can involve many atoms, and can be an interesting result of the simulation.
- 6. Material configurations that combine both atomic and cell degrees of freedom allow saddle point finding methods to find the mechanism of solid-solid phase transitions.

FUTURE ISSUES

- 1. Rough potential energy landscapes, found in some complex materials, pose a significant challenge for existing rare-event simulation methods based upon stationary points.
- 2. Better computational tools, coupled with electronic structure codes, would help the community explore potential energy landscapes and model the function of materials.

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

The National Science Foundation has provided ongoing support for the methodology work through grants CHE-0645497, CHE-1152342, and CHE-1534177. The Welch Foundation has funded applications under grant F-1841. Sustained computational resources have been provided by the Texas Advanced Computing Center. I would like to thank Hannes Jónsson, Arthur Voter, and David Wales for their helpful suggestions and comments.

LITERATURE CITED

- Jónsson H, Mills G, Jacobsen KW. 1998. Nudged elastic band method for finding minimum energy paths of transitions. In *Classical and Quantum Dynamics in Condensed Phase Simulations*, ed. BJ Berne, G Ciccotti, DF Coker, pp. 385–404. Singapore: World Sci.
- Henkelman G, Jónsson H. 2000. Improved tangent estimate in the nudged elastic band method for finding minimum energy paths and saddle points. *J. Chem. Phys.* 113:9978–85
- 3. Henkelman G, Uberuaga BP, Jónsson H. 2000. A climbing image nudged elastic band method for finding saddle points and minimum energy paths. *J. Chem. Phys.* 113:9901–4

- Weinan E, Ren W, Vanden-Eijnden E. 2002. String method for the study of rare events. Phys. Rev. B 66:052301
- Weinan E, Ren W, Vanden-Eijnden E. 2005. Finite temperature string method for the study of rare events. *J. Phys. Chem. B* 109:6688–93
- Dellago C, Bolhuis P, Csajka FS, Chandler D. 1998. Transition path sampling and the calculation of rate constants. *7. Chem. Phys.* 108:1964–77
- Bolhuis PG, Dellago C, Chandler D. 2000. Reaction coordinates of biomolecular isomerization. PNAS 97:5877–82
- Faradjian AK, Elber R. 2004. Computing time scales from reaction coordinates by milestoning. J. Chem. Phys. 120:10880–89
- Allen RJ, Frenkel D, ten Wolde PR. 2006. Simulating rare events in equilibrium or nonequilibrium stochastic systems. *J. Chem. Phys.* 124:024102
- Allen RJ, Frenkel D, ten Wolde PR. 2006. Forward flux sampling-type schemes for simulating rare events: efficiency analysis. 7. Chem. Phys. 124:194111
- van Erp TS, Moroni D, Bolhuis PG. 2003. A novel path sampling method for the calculation of rate constants. J. Chem. Phys. 118:7762
- 12. Wales DJ. 2002. Discrete path sampling. Mol. Phys. 100:3285-305
- Wales DJ. 2004. Some further applications of discrete path sampling to cluster isomerization. *Mol. Phys.* 102:891–908
- 14. Wales DJ. 2006. Energy landscapes: calculating pathways and rates. Int. Rev. Phys. Chem. 25:237-82
- 15. Wigner E. 1938. The transition state method. Trans. Faraday Soc. 34:29-41
- 16. Eyring H. 1938. The theory of absolute reaction rates. Trans. Faraday Soc. 34:41-48
- 17. Keck JC. 1967. Variational theory of reaction rates. Adv. Chem. Phys. 13:85
- Chandler D. 1978. Statistical mechanics of isomerization dynamics in liquids and the transition state approximation. *J. Chem. Phys.* 68:2959–70
- Voter AF, Doll D. 1985. Dynamical corrections to transition state theory for multistate systems: surface self-diffusion in the rare-event regime. *J. Chem. Phys.* 82:80–92
- Vineyard GH. 1957. Frequency factors and isotope effects in solid state rate processes. J. Phys. Chem. Solids 3:121–27
- Sheppard D, Terrell R, Henkelman G. 2008. Optimization methods for finding minimum energy paths. *J. Phys. Chem.* 128:134106
- Henkelman G, Jónsson H. 1999. A dimer method for finding saddle points on high dimensional potential surfaces using only first derivatives. *J. Chem. Phys.* 111:7010–22
- Mills G, Jónsson H, Schenter GK. 1995. Reversible work transition state theory: application to dissociative adsorption of hydrogen. Surf. Sci. 324:305–37
- Smidstrup S, Pedersen A, Stokbro K, Jónsson H. 2014. Improved initial guess for minimum energy path calculations. J. Chem. Phys. 140:214106
- Maragakis P, Andreev SA, Brumer Y, Reichman DR, Kaxiras E. 2002. Adaptive nudged elastic band approach for transition state calculation. *J. Chem. Phys.* 117:4651–58
- Trygubenko SA, Wales DJ. 2004. A doubly nudged elastic band method for finding transition states. *J. Chem. Phys.* 120:2082–94
- 27. Munro LJ, Wales DJ. 1999. Defect migration in crystalline silicon. Phys. Rev. B 59:3969
- Chill ST, Stevenson J, Ruhle V, Shang C, Xiao P, et al. 2014. Benchmarks for characterization of minima, transition states and pathways in atomic systems. *J. Chem. Theory Comput.* 10:5476–82
- Peters B, Heyden A, Bell AT, Chakraborty A. 2004. A growing string method for determining transition states: comparison to the nudged elastic band and string methods. *J. Chem. Phys.* 120:7877–86
- Henkelman G, Jóhannesson G, Jónsson H. 2000. Methods for finding saddle points and minimum energy paths. In Progress on Theoretical Chemistry and Physics, ed. S Schwartz, pp. 269–99. New York: Kluwer Acad.
- 31. Feibelman PJ. 1990. Diffusion path for an Al adatom on Al(001). Phys. Rev. Lett. 65:729-32
- Trushin OS, Salo P, Ala-Nissila T. 2000. Energetics and many-particle mechanisms of two-dimensional cluster diffusion on Cu(100) surfaces. *Phys. Rev. B* 62:1611–14
- Henkelman G, Jónsson H. 2003. Multiple time scale simulations of metal crystal growth reveal importance of multi-atom surface processes. *Phys. Rev. Lett.* 90:116101

- Uberuaga BP, Smith R, Cleave AR, Montalenti F, Henkelman G, et al. 2004. Structure and mobility of defects formed from collision cascades in MgO. *Phys. Rev. Lett.* 92:115505
- Xu L, Campbell CT, Jónsson H, Henkelman G. 2007. Kinetic Monte Carlo simulations of Pd deposition and island growth on MgO(100). *Surf. Sci.* 601:3133–42
- Uberuaga BP, Hoagland RG, Voter AF, Valone SM. 2007. Direct transformation of vacancy voids to stacking fault tetrahedra. *Phys. Rev. Lett.* 99:135501
- Xu L, Henkelman G. 2008. Adaptive kinetic Monte Carlo for first-principles accelerated dynamics. *J. Chem. Phys.* 129:114104
- Xu L, Henkelman G. 2008. Calculations of Ca adsorption on MgO(100) surface: determination of the binding sites and growth mode. *Phys. Rev. B* 77:205404
- Olsen RA, Kroes GJ, Henkelman G, Arnaldsson A, Jónsson H. 2004. Comparison of methods for finding saddle points without knowledge of the final states. J. Chem. Phys. 121:9776–92
- Heyden A, Bell AT, Keil FJ. 2005. Efficient methods for finding transition states in chemical reactions: comparison of improved dimer method and partitioned rational function optimization method. *J. Chem. Phys.* 123:224101
- Kästner J, Sherwood P. 2008. Superlinearly converging dimer method for transition state search. J. Chem. Phys. 128:014106
- Malek R, Mousseau N. 2000. Dynamics of Lennard-Jones clusters: a characterization of the activationrelaxation technique. *Phys. Rev. E* 62:7723–28
- Xiao P, Wu Q, Henkelman G. 2014. Basin constrained κ-dimer method for saddle point finding. *J. Chem. Phys.* 141:164111
- Pedersen A, Hafstein SF, Jónsson H. 2011. Efficient sampling of saddle points with the minimum-mode following method. SIAM J. Sci. Comput. 33:633–52
- Duncan J, Wu Q, Promislow K, Henkelman G. 2014. Biased gradient squared descent saddle point finding method. *J. Chem. Phys.* 140:194102
- Sørensen MR, Voter AF. 2000. Temperature-accelerated dynamics for simulation of infrequent events. *7. Chem. Phys.* 112:9599–606
- Chill ST, Henkelman G. 2014. Molecular dynamics saddle search adaptive kinetic Monte Carlo. J. Chem. Phys. 140:214110
- Voter AF, Montalenti F, Germann TC. 2002. Extending the time scale in atomistic simulation of materials. Annu. Rev. Mater. Res. 32:321–46
- 49. Voter AF. 1998. Parallel replica method for dynamics of infrequent events. Phys. Rev. B 57:R13985-88
- 50. Voter AF. 1997. Hyperdynamics: accelerated molecular dynamics of infrequent events. *Phys. Rev. Lett.* 78:3908–11
- 51. Bortz AB, Kalos MH, Lebowitz JL. 1975. A new algorithm for Monte Carlo simulation of Ising spin systems. *J. Comput. Phys.* 17:10
- Gillespie DT. 1976. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *J. Comp. Phys.* 22:403–34
- Voter AF. 1986. Classically exact overlayer dynamics: diffusion of rhodium clusters on Rh(100). *Phys. Rev.* B 34:6819–29
- Voter AF. 2007. Introduction to the kinetic Monte Carlo method. In *Radiation Effects in Solids*, NATO Science Series II: Mathematics, Physics and Chemistry, ed. KE Sickafus, EA Kotomin, BP Uberuaga, pp. 1–23. Dordrecht, Neth: Springer
- Henkelman G, Jónsson H. 2001. Long time scale kinetic Monte Carlo simulations without lattice approximation and predefined event table. *J. Chem. Phys.* 115:9657–66
- Chill ST, Welborn M, Terrell R, Zhang L, Berthet JC, et al. 2014. Eon: software for long time simulations of atomic scale systems. *Model. Simul. Mater. Sci. Eng.* 22:055002
- 57. Terrell R, Welborn M, Chill ST, Henkelman G. 2012. Database of atomistic reaction mechanisms with application to kinetic Monte Carlo. *J. Chem. Phys.* 137:014105
- El-Mellouhi F, Mousseau N, Lewis LJ. 2008. Kinetic activation-relaxation technique: an off-lattice selflearning kinetic Monte Carlo algorithm. *Phys. Rev. B* 78:153202
- Béland LK, Brommer P, El-Mellouhi F, Joly JF, Mousseau N. 2011. Kinetic activation-relaxation technique. *Phys. Rev. E* 84:046704

- 60. McKay BD. 1981. Practical graph isomorphism. Congr. Numer. 30:45-87
- Nandipati G, Kara A, Shah SI, Rahman TS. 2012. Off-lattice pattern recognition scheme for kinetic Monte Carlo simulations. *J. Comput. Phys.* 231:3548–60
- Vernon LJ, Kenny SD, Smith R. 2011. Growth mechanisms for TiO₂at its rutile (110) surface. *Phys. Rev.* B 83:075412
- Trushin OS, Karim A, Kara A, Rahman TS. 2005. Self-learning kinetic Monte Carlo method: application to Cu(111). *Phys. Rev. B* 72:115401
- Konwar D, Bhute VJ, Chatterjee A. 2011. An off-lattice, self-learning kinetic Monte Carlo method using local environments. J. Chem. Phys. 135:17103
- Xu H, Osetsky YN, Stoller RE. 2011. Simulating complex atomistic processes: on-the-fly kinetic Monte Carlo scheme with selective active volumes. *Phys. Rev. B* 84:142103
- Miron RA, Fichthorn KA. 2004. Multiple-time scale accelerated molecular dynamics: addressing the small-barrier problem. *Phys. Rev. Lett.* 93:128301
- Boulougouris GC, Theodorou DN. 2007. Dynamical integration of a Markovian web: a first passage time approach. J. Chem. Phys. 127:084903
- Wales DJ. 2009. Calculating rate constants and committor probabilities for transition networks by graph transformation. *J. Chem. Phys.* 130:204111
- Chatterjee A, Voter AF. 2010. Accurate acceleration of kinetic Monte Carlo simulations through the modification of rate constants. *J. Chem. Phys.* 132:194101
- Novotny MA. 1995. Monte Carlo algorithms with absorbing Markov chains: fast local algorithms for slow dynamics. *Phys. Rev. Lett.* 74:1–5
- Duncan J, Harjunmaa A, Terrell R, Drautz R, Henkelman G, Rogal J. 2016. Collective atomic displacements during complex phase boundary migration in solid-solid phase transformations. *Phys. Rev. Lett.* 116:035701
- Pedersen A, Henkelman G, Schiøtz J, Jónsson H. 2009. Long time scale simulation of a grain boundary in copper. New J. Phys. 11:073034
- Pedersen A, Jónsson H. 2009. Simulations of hydrogen diffusion at grain boundaries in aluminum. Acta Met. 57:4036–45
- Pedersen A, Karssemeijer LJ, Cuppen HM, Jónsson H. 2015. Long-timescale simulations of H₂O admolecule diffusion on ice Ih(0001) surfaces. *J. Phys. Chem. C* 119:16528
- Zhou XW, Wadley HNG, Johnson RA, Larson DJ, Tabat N, et al. 2001. Atomic scale structure of sputtered metal multilayers. *Acta Mater*. 49:4005–15
- Becker OM, Karplus M. 1997. The topology of multidimensional potential energy surfaces: theory and application to peptide structure and kinetics. J. Chem. Phys. 106:1495–517
- Xu L, Mei DH, Henkelman G. 2009. Adaptive kinetic Monte Carlo simulation of methanol decomposition on Cu(100). *J. Chem. Phys.* 131:244520
- Xu L, Henkelman G, Campbell CT, Jónsson H. 2005. Small Pd clusters, up to the tetramer at least, are highly mobile on the MgO(100) surface. *Phys. Rev. Lett.* 95:146103
- 79. Trinkle DR, Hennig RG, Srinivasan SG, Hatch DM, Jones MD, et al. 2003. New mechanism for the α to ω martensitic transformation in pure titanium. *Phys. Rev. Lett.* 91:025701
- Caspersen KJ, Carter EA. 2005. Finding transition states for crystalline solid-solid phase transformations. PNAS 102:6738–43
- Sheppard D, Xiao P, Chemelewski W, Johnson DD, Henkelman G. 2012. A generalized solid-state nudged elastic band method. *J. Chem. Phys.* 136:074103
- Xiao P, Sheppard D, Rogal J, Henkelman G. 2014. Solid-state dimer method for calculating solid-solid phase transitions. *J. Chem. Phys.* 140:174104
- 83. Jónsson H. 2011. Simulation of surface processes. PNAS 108:944-49
- van Duin ACT, Dasgupta S, Lorant F, Goddard WA. 2001. ReaxFF: a reactive force field for hydrocarbons. *J. Phys. Chem. A* 105:9396–409
- Yu J, Sinnott SB, Phillpot SR. 2007. Charge optimized many-body potential for the SiSiO₂ system. *Phys. Rev. B* 75:085311

$\mathbf{\hat{R}}$

v

Annual Review of Materials Research

Contents

Novel Functionality Through Metamaterials (Venkatraman Gopalan, Don Lipkin & Simon Phillpot, Editors)

Control of Localized Surface Plasmon Resonances in Metal Oxide Nanocrystals Ankit Agrawal, Robert W. Johns, and Delia J. Milliron	1
DNA-Driven Assembly: From Polyhedral Nanoparticles to Proteins Martin Girard, Jaime A. Millan, and Monica Olvera de la Cruz	33
Harnessing Instabilities to Design Tunable Architected Cellular Materials Katia Bertoldi	51
Negative-Poisson's-Ratio Materials: Auxetic Solids Roderic S. Lakes	63
Sound Absorption Structures: From Porous Media to Acoustic Metamaterials <i>Min Yang and Ping Sheng</i>	83
Structured X-Ray Optics for Laboratory-Based Materials Analysis Carolyn A. MacDonald	115
Synchrotron X-Ray Optics Albert T. Macrander and XianRong Huang	135

Current Interest

Active Crystal Growth Techniques for Quantum Materials <i>Julian I. Schmehr and Stethen D. Wilson</i>	153
Atomic-Scale Structure-Property Relationships in Lithium Ion Battery	155
Electrode Materials Zhenzhong Yang, Lin Gu, Yong-Sheng Hu, and Hong Li	175
Atomistic Simulations of Activated Processes in Materials	
G. Henkelman	199

 Deformation of Crystals: Connections with Statistical Physics James P. Sethna, Matthew K. Bierbaum, Karin A. Dahmen, Carl P. Goodrich, Julia R. Greer, Lorien X. Hayden, Jaron P. Kent-Dobias, Edward D. Lee, Danilo B. Liarte, Xiaoyue Ni, Katherine N. Quinn, Archishman Raju, D. Zeb Rocklin, Ashivni Shekhawat, and Stefano Zapperi
Heusler 4.0: Tunable Materials Lukas Wollmann, Ajaya K. Nayak, Stuart S.P. Parkin, and Claudia Felser
Physical Dynamics of Ice Crystal Growth Kenneth G. Libbrecht
Silicate Deposit Degradation of Engineered Coatings in Gas Turbines: Progress Toward Models and Materials Solutions David L. Poerschke, R. Wesley Jackson, and Carlos G. Levi
Structural and Functional Fibers Huibin Chang, Jeffrey Luo, Prabhakar V. Gulgunje, and Satish Kumar
Synthetic Two-Dimensional Polymers Marco Servalli and A. Dieter Schlüter 361
Transparent Perovskite Barium Stannate with High Electron Mobility and Thermal Stability Woong-Jhae Lee, Hyung Joon Kim, Jeonghun Kang, Dong Hyun Jang, Tai Hoon Kim, Jeong Hyuk Lee, and Kee Hoon Kim
Visualization of Atomic-Scale Motions in Materials via Femtosecond X-Ray Scattering Techniques <i>Aaron M. Lindenberg, Steven L. Johnson, and David A. Reis</i>
X-Ray Tomography for Lithium Ion Battery Research: A Practical Guide Patrick Pietsch and Vanessa Wood

Indexes

Cumulative Index of Contributing Authors,	Volumes 43–47	
---	---------------	--

Errata

An online log of corrections to *Annual Review of Materials Research* articles may be found at http://www.annualreviews.org/errata/matsci