Accelerated Molecular Dynamics

Timothy C. Germann
Materials Science (X-7)
Arthur F. Voter
Theoretical Chemistry and Molecular Physics (T-12)
Los Alamos National Laboratory

Slide 1

"Summer School on Computational Materials Science,"
University of Illinois at Urbana-Champaign, June 2001

Collaborators: Mads R. Sørensen, Francesco Montalenti, Sriram Swaminarayan
James Sprague (NRL)

ILLINOIS Los Alamos NATIONAL LABORATORY

Outline of the Lectures

- Motivation: Why Accelerated MD?
- Infrequent events and transition state theory
- Hyperdynamics
  - Basic concept
  - Mathematical justification
  - Trial applications
- Parallel Replica Dynamics
- Temperature Accelerated Dynamics
- Application to thin film growth at experimental deposition rates

Slide 2

Conventional Molecular Dynamics Capabilities

Slide 3

- The maximum integration timestep in MD simulations is determined by the fastest motion in the system, typically vibrations.
- Typical molecular vibration frequencies range up to 3000 cm^{-1} (period = 10 fs), and optical phonon frequencies are on the order of 10 THz (period = 100 fs).
- A typical MD timestep is therefore on the order of 1 fs.
- Millions of timesteps only reach ns of simulation time.
- High-frequency stretch or bend modes can be frozen in molecular systems (e.g. SHAKE/RATTLE) or multiple timesteps used (e.g. r-RESPA), but the gain is typically only a factor of 2-3.

Need for Accelerated Molecular Dynamics Methods

Slide 4

Many important problems lie beyond the timescales accessible to conventional MD, due to the infrequent surface or bulk diffusion mechanisms involved, including:

- Vapor-deposited thin film growth (metals and semiconductors) *especially heteroepitaxial layers*
- Stress-assisted diffusion of point defects, defect clusters, and dislocations
- Ion implantation or radiation damage annealing *especially for complex crystal phases such as $\alpha-Pu$*

Vapor-deposited thin film growth

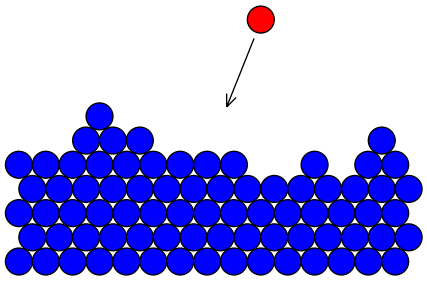

- Each deposition event lasts less than 2 ps – ideally suited for traditional MD.
- Experimental growth conditions are typically 1 monolayer per second or slower.

Slide 5

- A delicate interplay of various diffusion barriers causes a nonmonotonic temperature dependence of the final surface roughness, from smooth layer-by-layer growth to rough 3D growth and back to layer-by-layer.
- These activated processes, which determine the final morphology, occur on timescales far longer than that of traditional MD.

ILLINOIS Los Alamos NATIONAL LABORATORY

Neutron/ion bombardment simulations

- Radiation damage or ion implantation simulations involve a rapid (ps) dissipation of the primary knock-on atom (PKA) kinetic energy, which is again well-suited to standard molecular dynamics.
- However, the typical α decay rate or ion flux is many orders of magnitude slower; for instance, ^{239}Pu undergoes α decay with a half-life of 24,400 years. On average, each Pu atom is displaced from its lattice site once every 10 years, and during the same time period, 411 α-decays will occur for every million Pu atoms.

Slide 6

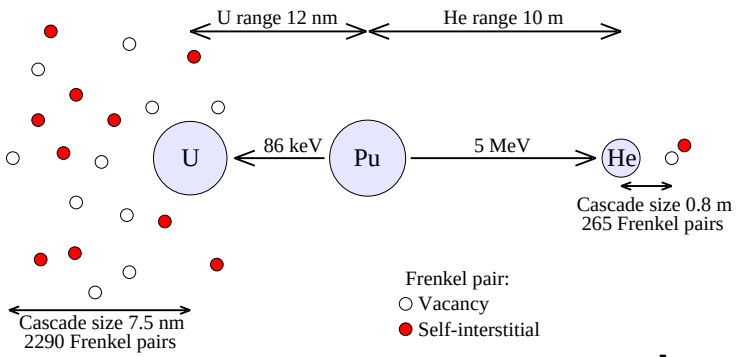

ILLINOIS Los Alamos NATIONAL LABORATORY

One possible solution: kinetic Monte Carlo

Slide 7

- kMC is extremely useful and has been applied to each of these problems, sometimes in combination with direct MD:
  - *Simulations of energetic beam deposition: From picoseconds to seconds*, J. Jacobsen, B.H. Cooper, and J.P. Sethna, Phys. Rev. B **58**, 15847 (1998).
  - *Self-decay-induced damage production and micro-structure evolution in fcc metals: An atomic-scale computer simulation approach*, T. Diaz de la Rubia *et al.*, J. Comp.-Aided Mat. Design **5**, 243 (1998).
- kMC requires a prior knowledge (catalog) of all possible events.
- Even for one of the simplest systems – submonolayer diffusion on a fcc(100) surface – new diffusion mechanisms discovered during the 1990s raise questions about the completeness of kMC rate catalogs…

ILLINOIS Los Alamos NATIONAL LABORATORY

Unexpected events

- Before 1990, most kMC catalogs for fcc(100) surface diffusion only contained a hop rate dependent on the local environment (occupancy of sites 1-10):

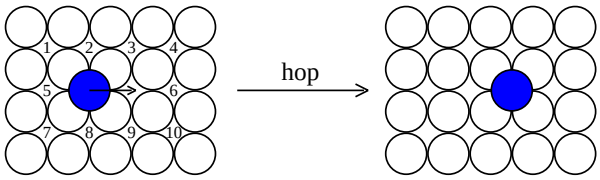

Slide 8

- A concerted exchange mechanism, only discovered in 1990, is *dominant* for some metals:

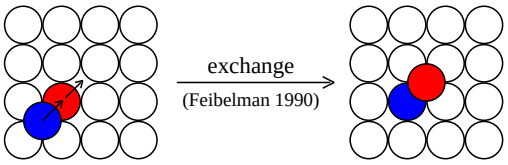

- The situation becomes nearly hopeless if we are interested in grain boundaries, strained systems (as in lattice-mismatched heteroepitaxial growth), or complex crystal structures…

ILLINOIS Los Alamos NATIONAL LABORATORY

Complex crystal phases: α-Pu

The low-temperature phase of plutonium (stable below 388 K at $P = 0$) resembles a distorted hcp lattice:

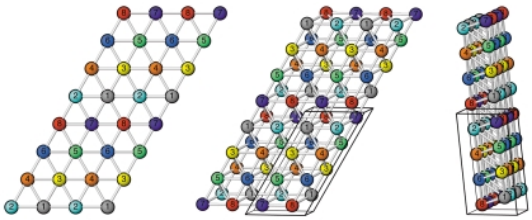

hexagonal lattice

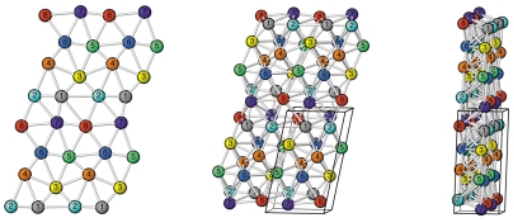

monoclinic α phase

- Predicting diffusion mechanisms for the α-phase structure is out of the question; even if they could all be enumerated, the resulting rate catalog may be prohibitively large.

ILLINOIS — Los Alamos NATIONAL LABORATORY

Slide 9

kMC with on-the-fly saddle searching

One possibility is to compute a kMC rate catalog whenever a new state is entered, using one of several saddle-finding methods:

- Mode-following techniques; particularly promising in this class is the dimer method, since no Hessian matrix is required:

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

- High-temperature sampling of possible escape routes

Can one ever be sure that all relevant mechanisms have been seen?

An ideal computational method for studying these phenomena would be MD-like, in that nothing about the possible escape paths needs to be specified in advance, only the interatomic potential (and not even that, in the case of *ab initio* MD).

ILLINOIS — Los Alamos NATIONAL LABORATORY

Slide 10

Activation-Relaxation Technique (ART)

Event-Based Relaxation of Continuous Disordered Systems,

G.T. Barkema and N. Mousseau,
Phys. Rev. Lett. **77**, 4358 (1996).

The activation-relaxation technique: an efficient algorithm for sampling energy landscapes,

G.T. Barkema and N. Mousseau,
Comput. Mater. Sci. **20**, 285 (2001).

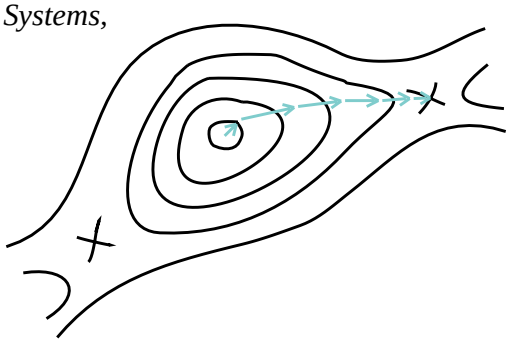

Slide 11

- The activation step pushes the configuration up a valley towards a saddle point using the modified force

$$\mathbf{G} = \mathbf{F} - (1+\alpha)(\mathbf{F}\cdot\boldsymbol{\Delta}\mathbf{x})\boldsymbol{\Delta}\mathbf{x}$$

- Relaxation locates the basin on the other side of the saddle point.
- A Metropolis Monte Carlo probability may be used to accept or reject this event, or it may simply be added to a kMC rate catalog.

Infrequent Event Systems

- All of these are examples of a more general class of problems, that of infrequent (or activated) event systems:

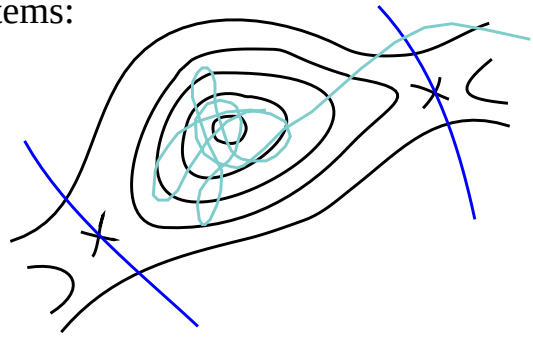

Slide 12

- The system vibrates (and vibrates and vibrates...) in a many-dimensional potential basin the vast majority of the time.
- Occasionally – but only very rarely – it escapes through a dividing surface to a new potential basin.
- The objective is to accelerate these escapes *without* corrupting the relative escape probabilities for various exit pathways.

Transition State Theory (TST)

- For infrequent event systems, assuming no correlated crossing events

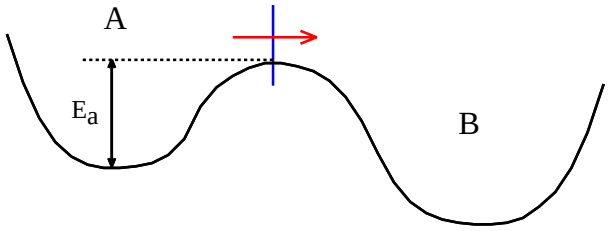

- $k_{A\to B}^{TST}$ = equilibrium flux through a particular dividing surface
  $= \langle |v| \delta(x) \rangle_A$
  $= \nu_0 \ \exp\left(-E_a/k_B T\right)$ ⟵ harmonic approximation

Slide 13

- No dynamics required (although minimal dynamics in the dividing surface region can be used to improve the classical TST rate, or in various "quantum transition state theory" methods)
- Can we exploit TST to develop methods which do not require any knowledge of where dividing surfaces are?

Three Recent Methods for Infrequent Events

Hyperdynamics

A method for accelerating the molecular dynamics simulation of infrequent events, A.F. Voter, J. Chem. Phys. **106**, 4665 (1997).

Hyperdynamics: Accelerated Molecular Dynamics of Infrequent Events, A.F. Voter, Phys. Rev. Lett. **78**, 3908 (1997).

- Assumptions: infrequent events, TST

Slide 14

Parallel Replica Dynamics

Parallel replica method for dynamics of infrequent events, A.F. Voter, Phys. Rev. B **57**, R13985 (1998).

- Assumption: infrequent events

Temperature Accelerated Dynamics

Temperature-acclerated dynamics for simulation of infrequent events, M.R. Sørensen and A.F. Voter, J. Chem. Phys. **112**, 9599 (2000).

- Assumptions: infrequent events, *harmonic* TST, lower bound on preexponential factor

Hyperdynamics: Basic Approach

- Consider the Vineyard (harmonic TST) expression for the rate:

$$k_{A\to B}^{hTST}(T) = \frac{\prod_{i}^{3N} \nu_i^{min,A}}{\prod_{i}^{3N-1} \nu_i^{sad,AB}} \; e^{-E_a^B/k_B T}$$

Slide 15

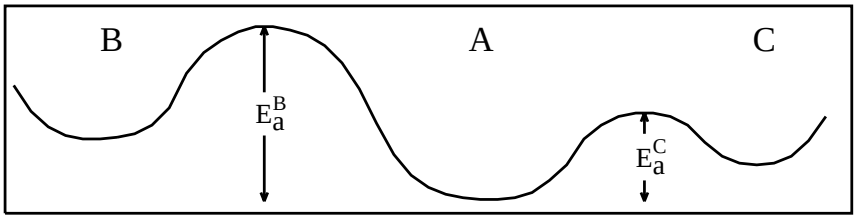

- The ratio of any two escape rates from state A only depends on the local properties at the saddles:

$$\frac{k_{A\to B}^{hTST}(T)}{k_{A\to C}^{hTST}(T)} = \prod_{i}^{3N-1} \left(\frac{\nu_i^{sad,AC}}{\nu_i^{sad,AB}} \right) e^{(E_a^C - E_a^B)/k_B T}$$

Hyperdynamics: Basic Approach

- Therefore we are free to modify the basins however we like, as long as the saddle properties are unaffected. Add a positive bias potential $\Delta V_b(\mathbf{r})$ to the potential surface $V(\mathbf{r})$ in the wells:

Slide 16

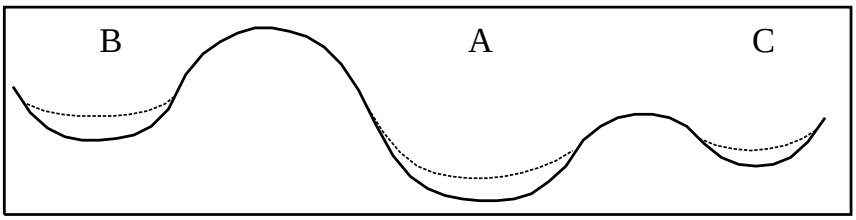

- The TST rates increase by $e^{\Delta V_b(\mathbf{r}_A)/k_B T}$ (neglecting the prefactor)
- Since the escape rates are uniformly raised, a valid state-to-state sequence is obtained, but at an accelerated pace.
- In fact, harmonic TST is not required, only to assume TST is exact (implying no correlated crossing events), as we now show…

Hyperdynamics: Mathematical Foundations

The transition state theory (TST) thermal rate constant is given by the canonical phase-space average $\langle\rangle_A$ of the flux through a dividing surface:

$$k_{A\rightarrow}^{TST} = \langle |v_A| \delta_A(\mathbf{r}) \rangle_A$$

On a biased potential $V(\mathbf{r}) + \Delta V_b(\mathbf{r})$, using $\beta \equiv 1/k_B T$,

Slide 17

$$k_{A\rightarrow}^{TST} = \frac{\langle |v_A| \delta_A(\mathbf{r}) e^{\beta \Delta V_b(\mathbf{r})} \rangle_{A_b}}{\langle e^{\beta \Delta V_b(\mathbf{r})} \rangle_{A_b}}$$

But since $\Delta V_b(\mathbf{r}) = 0$ whenever $\delta_A(\mathbf{r}) \neq 0$,

$$k_{A\rightarrow}^{TST} = \frac{\langle |v_A| \delta_A(\mathbf{r}) \rangle_{A_b}}{\langle e^{\beta \Delta V_b(\mathbf{r})} \rangle_{A_b}}$$

So the TST escape rate on the biased surface is related to that on the original PES by

$$k_{A_b\rightarrow}^{TST} = k_{A\rightarrow}^{TST} \ \langle e^{\beta \Delta V_b(\mathbf{r})} \rangle_{A_b}$$

Hyperdynamics: Mathematical Foundations

The average escape time $\tau_{esc}^A = 1/k_{A\rightarrow}^{TST}$, so inverting the previous expression gives

$$\tau_{esc}^A = \tau_{esc}^{A_b} \ \langle e^{\beta \Delta V_b(\mathbf{r})} \rangle_{A_b}$$

Evaluating $\langle\rangle_{A_b}$ from a n_{tot}-step MD trajectory on the biased surface,

Slide 18

$$\tau_{esc}^A = \tau_{esc}^{A_b} \ \frac{1}{n_{tot}} \sum_i^{n_{tot}} e^{\beta \Delta V_b(\mathbf{r}(t_i))}$$

Events are thus accelerated by the average "boost" factor

$$\frac{\tau_{esc}^A}{\tau_{esc}^{A_b}} = \frac{1}{n_{tot}} \sum_i^{n_{tot}} e^{\beta \Delta V_b(\mathbf{r}(t_i))}$$

and we can simply run a hyperclock which advances by $\Delta t_{hyper} = \Delta t_{MD} \ e^{\beta \Delta V_b(\mathbf{r}_i)}$ each timestep.

Defining a bias potential

Consider the simplest possible V_b:

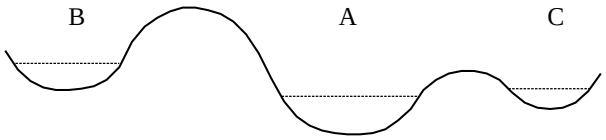

Simple bias potential for boosting molecular dynamics with the hyperdynamics scheme, M.M. Steiner, P.-A. Genilloud, and J.W. Wilkins, Phys. Rev. B **57**, 10236 (1998).

Slide 19

- Very cost effective (zero or negative overhead)
- Works well for low-dimensional systems (few atoms)
- Vanishing boost for large systems

Conclusion: may be useful for small first principles applications, but not for general many-dimensional systems (tens of atoms)

Recognizing Saddle Points

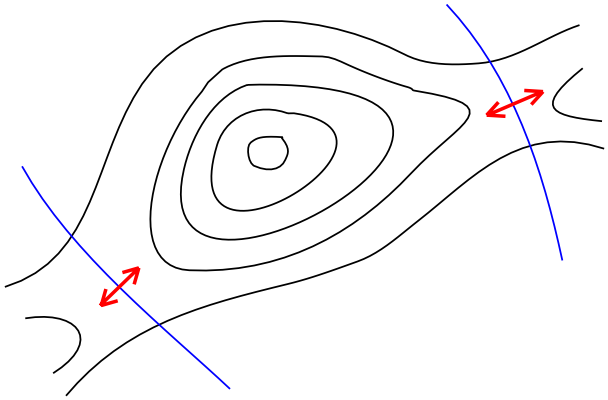

The dividing (hyper-) surface has an intrinsically nonlocal definition. However, an *approximate* local criterion may be obtained by noting that:

Slide 20

- At a first-order saddle point, the gradient $\mathbf{g} = \mathbf{0}$ and the Hessian $\mathbf{H}$ has exactly one negative eigenvalue.
- The corresponding eigenvector $\mathbf{C}_1$ points along the reaction coordinate, and is orthogonal to the dividing surface.
- Similarly, the ridge line has at least one negative eigenvalue, $\epsilon_1 < 0$, and the gradient along $\mathbf{C}_1$ is zero.

Defining a bias potential

- Along a ridge line, the lowest eigenvalue (ϵ_1) of the Hessian is negative, and along its eigenvector ($\mathbf{C}_1$), the gradient is zero.
- Therefore, a possible bias potential is one that turns on smoothly when either of these conditions is violated:

$$\Delta V_b(\mathbf{r}) = a\theta(\epsilon_1)\epsilon_1^2 + c(\mathbf{C}_1 \cdot \mathbf{g})^2,$$

Slide 21

where $\theta(x)$ is a step function (1 when $x > 0$ and 0 when $x < 0$) and g is the potential gradient, $g_i \equiv \partial V(\mathbf{r}) \; / \; \partial x_i$.

- Alternatively, a bias potential which exactly "fills in" a cosine potential $V(x) = (h/2)\cos(x/d)$ with barrier height h and period $2\pi d$ may be used:

$$\Delta V_b^{cos}(x) = \frac{h}{2}\left[1 + \frac{\epsilon_1}{(\epsilon_1^2 + (\mathbf{C}_1 \cdot \mathbf{g})^2/d^2)^{1/2}}\right]$$

Extracting a relevant "boosted" time

- Each timestep we advance the t_{hyper} clock by an amount

$$\Delta t_{hyper} = \Delta t_{MD} \exp\left(\Delta V_b(\mathbf{r}_i)/k_B T\right)$$

- The resulting t_{hyper} evolution is highly nonlinear:

Slide 22

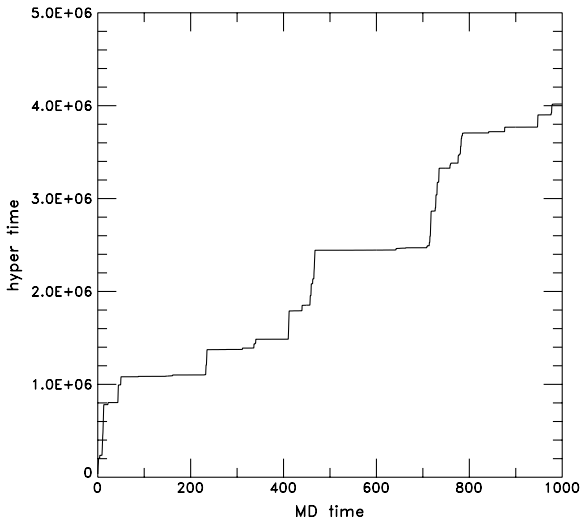

Example: Two-dimensional model potential†

Slide 23

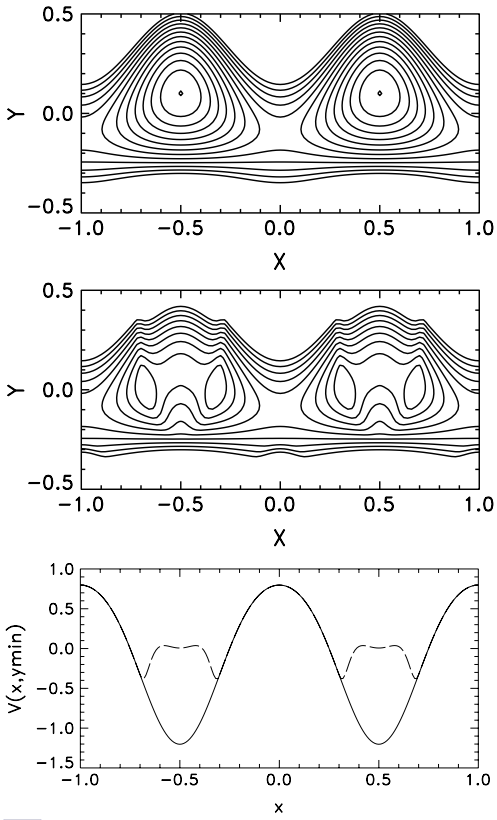

† A. F. Voter, J. Chem. Phys. **106**, 4665 (1997).

Arrhenius plot with boost factors:

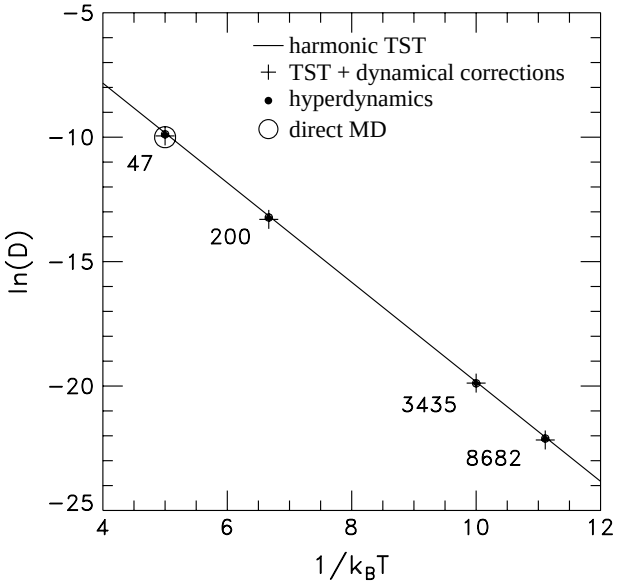

2-D model with various bias potentials

Slide 24

- Exponential boost factor is apparent:
- In the lab, you can experiment with different bias potential choices
- More complex variants by adding another cos term and/or additional "y-like" degrees of freedom

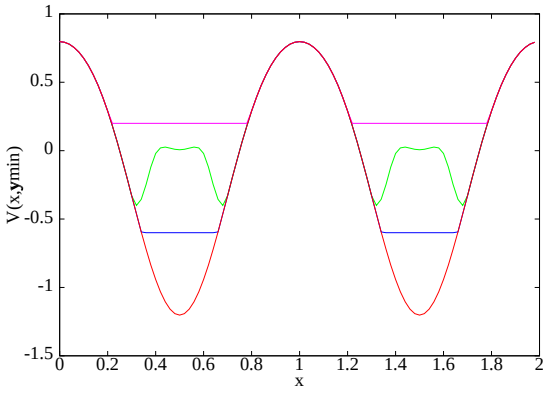

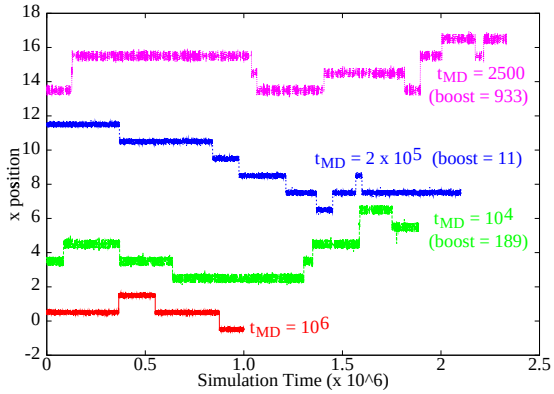

Embedded Atom Method (EAM) Potentials

- Pair potentials (e.g. Lennard-Jones) are unable to capture many-body effects in covalent species and metals
- EAM is based on density functional theory
- Used for metals, especially those with empty or filled d bands

Slide 25

$$E_{tot} = \sum_i F_i(\rho_i) + \frac{1}{2}\sum_{i \neq j} \phi_{ij}(R_{ij})$$

where

$$\rho_i = \sum_{j \neq i} \rho_j^a(R_{ij})$$

- F_i represents the embedding energy for each atom into a host electron gas due to neighboring atoms
- ϕ_{ij} is a core-core electrostatic (repulsive) pair interaction

Hessian-free bias potentials†

For large systems (100 or more atoms), setting up and diagonalizing the Hessian $\mathbf{H}$ becomes a serious bottleneck. Since only the lowest eigenvalue (or lowest few eigenvalues) is needed, iterative Lanczos-type methods can be used. Better yet is a bias potential which does not even require an explicit Hessian construction:

Slide 26

- The lowest eigenvalue ϵ_1 corresponds to the direction $\mathbf{s}_1$ which minimizes the second derivative $\partial^2 V(\mathbf{r})/\partial \mathbf{s}_1^2$.
- Express ϵ_1 as a numerical derivative,

$$\epsilon_1(\mathbf{s}_1) \simeq \frac{V(\mathbf{r}+\eta\mathbf{s}_1) + V(\mathbf{r}-\eta\mathbf{s}_1) - 2V(\mathbf{r})}{\eta^2}$$

- Using steepest descent or conjugate gradient methods, vary the direction $\mathbf{s}_1$ to minimize ϵ_1.

† A. F. Voter, Phys. Rev. Lett. **78**, 3908 (1997).

"λ method" for $\mathbf{C}_1 \cdot \mathbf{g}$

In addition to ϵ_1, we also need the gradient projection on its eigenvector, $g_{1p} \equiv \mathbf{C}_1 \cdot \mathbf{g}$ and its derivative. This is done in the same spirit:

- Numerically determine the lowest eigenvalues $\epsilon^{\pm}$ for the modified Hessian $\mathbf{H} \pm \lambda \mathbf{g}\mathbf{g}^{\dagger}$ by minimizing

Slide 27

$$\epsilon^{\pm}(\mathbf{s}) \simeq \frac{V(\mathbf{r}+\eta\mathbf{s}) + V(\mathbf{r}-\eta\mathbf{s}) - 2V(\mathbf{r})}{\eta^2} \pm \lambda \left[\frac{V(\mathbf{r}+\eta\mathbf{s}) + V(\mathbf{r}-\eta\mathbf{s})}{2\eta} \right]^2$$

- Combine to approximate g_{1p} as $\quad (g_{1p})^2 \simeq \dfrac{\epsilon^+ - \epsilon^-}{2\lambda}$
- Only a few dimer iterations are needed to converge $\epsilon^{\pm}$ using s_1 as a first guess to s
- Gradient of g_{1p} has a simple form
- But numerical instabilities for small λ need to be investigated

More on bias potentials...

This iterative N-scaling approach:

- only requires first derivatives of $V(\mathbf{r})$
- typically requires 20-50 force calls for EAM potentials

For numerical stability and other problems (e.g., $g_{1p} \to 0$) when the lowest two eigenvalues cross, an added eigenvalue repulsion term is useful:

Slide 28

$$\Delta V_b^{\Delta\epsilon} = \begin{cases} a[1 - 3q^2 + 2q^3] & \text{if } \epsilon_2 - \epsilon_1 \le \Delta\epsilon_c \\ 0 & \text{otherwise} \end{cases}$$

where $q \equiv (\epsilon_2 - \epsilon_1)/\Delta\epsilon_c$. By maintaining orthogonality of $\mathbf{s}_2$ to the numerically determined $\mathbf{s}_1$, ϵ_2 can be computed in a similar manner as ϵ_1.

The final bias potential is $\Delta V_b = \Delta V_b^{cos} + \Delta V_b^{\Delta\epsilon}$, with the ΔV_b^{cos} height h set to be somewhat less than the lowest anticipated barrier.

Ramping up the bias potential

An automatic procedure may also be used to set the bias strength h dynamically during a simulation in a "safe" manner:

- If no event has occured after running dynamics for a time t, then we can put a lower bound on the lowest barrier E_a for escape from this state, using a typical prefactor $\nu_0 \simeq 10^{12}$:

  $t < 1/$ (fastest rate)

  fastest rate $\simeq 10^{12} e^{-E_a^{min}/k_B T}$

- Therefore, $E_a^{min} > k_B T \ln\left(10^{12} t\right)$ and the bias potential strength h can be set to a value slightly below this value.

- Reset $h = 0$ (or some safe minimum value) and restart this procedure after each transition to safely ramp up the bias potential.

Slide 29

Example: Ag_{10}/Ag(111) diffusion†

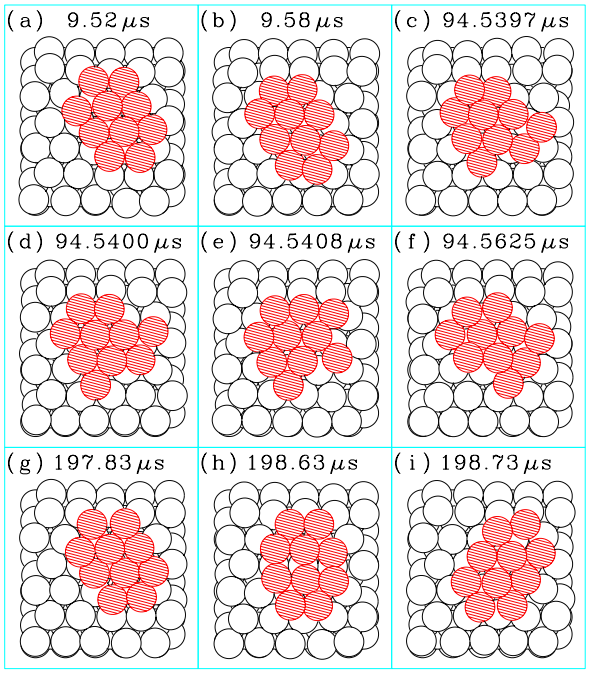

- average boost = 8310
- overhead per step $\simeq 20$
- computational boost = 400

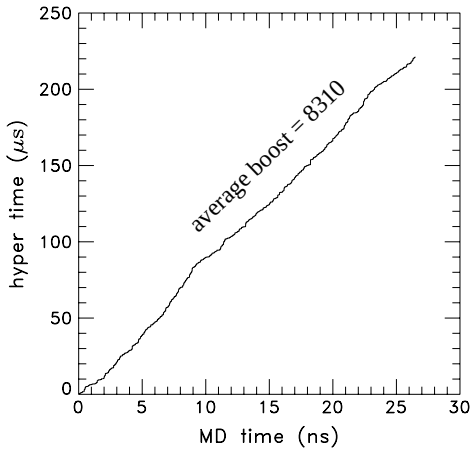

Slide 30

† A. F. Voter, Phys. Rev. Lett. **78**, 3908 (1997).

System size effects

- Increasing the system size by adding spectator atoms has no effect on the performance (only making each MD step more expensive).
- The achievable boost drops as the number of "active" atoms in the system (i.e., the number of low-frequency eigenmodes) is increased, and it takes longer to converge the eigenvalues:

Slide 31

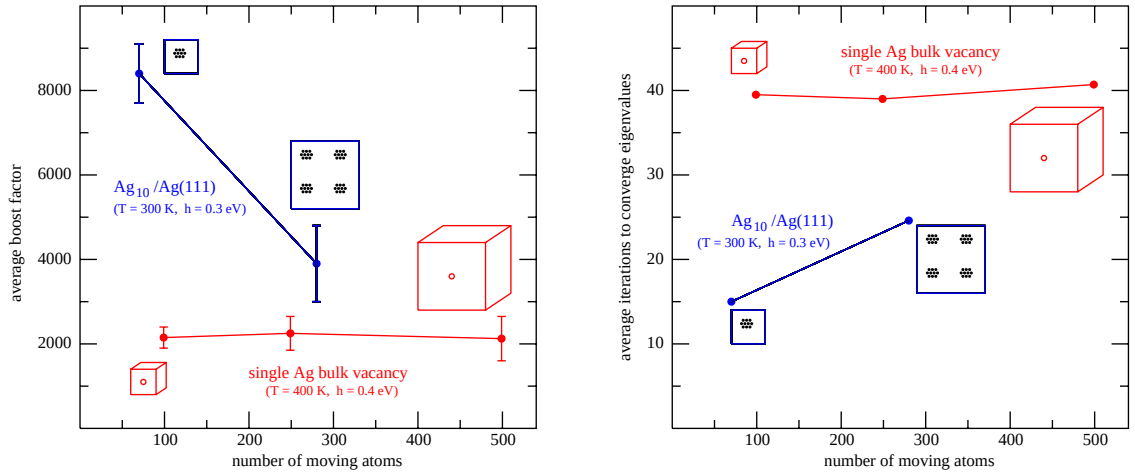

The Low-Barrier Problem

Often, a set of states may be connected by transitions with small activation energies, limiting the boost which could be achieved once a more stable configuration is reached via a higher-E_a mechanism:

Slide 32

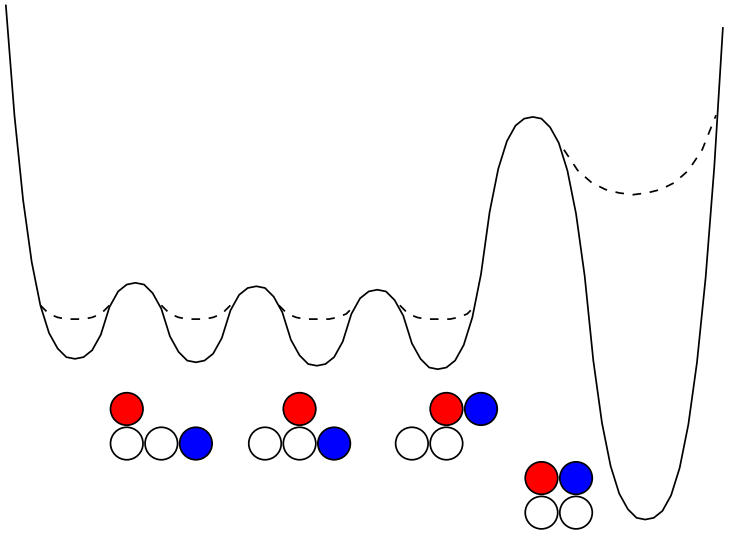

This is still an unsolved problem for hyperdynamics!

Parallel Replica Dynamics: Basic Concept

Slide 33

- Harness parallel or distributed processors to extend simulation time for infrequent event systems by independently exploring phase space

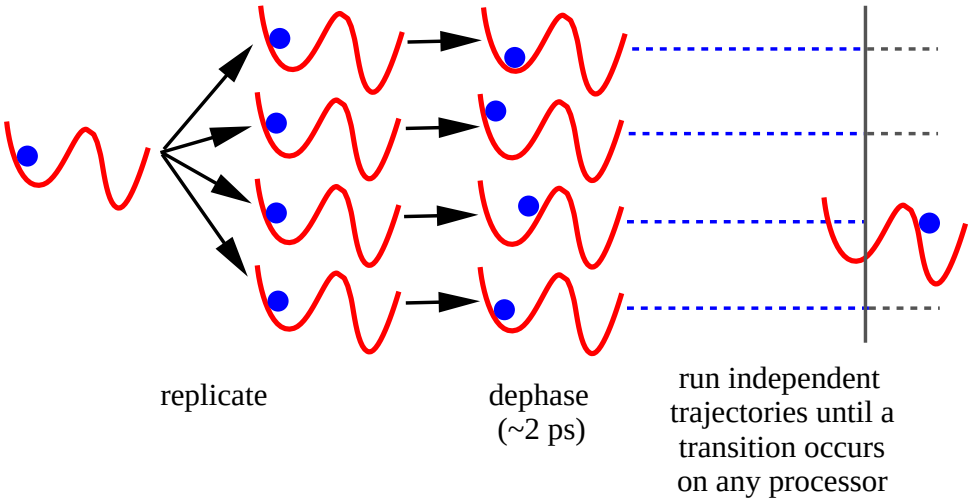

- When any processor detects a transition, sum the trajectory times from all processors, and restart all processors in the new state.

Detecting Transitions

Slide 34

- Do a periodic (say, every 2 ps) steepest descent or conjugate gradient minimization just long enough to ascertain whether the configuration has crossed to a different basin.

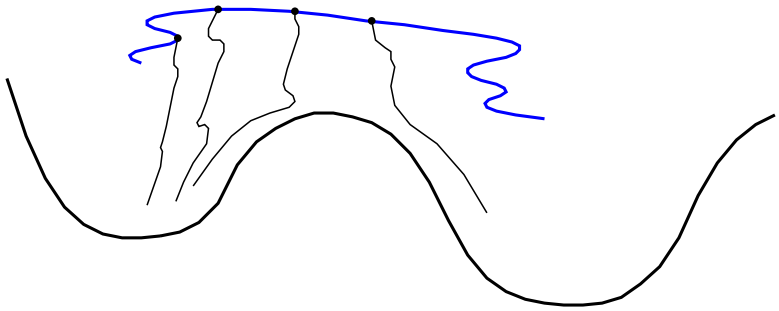

- Usually, only a few steps are required to say no transition has occured.
- Near saddle points, 100 steps may be needed.

Parallel Replica Dynamics: Mathematical Foundations

- Assume only infrequent events having a first-order rate process, with a total escape rate equal to the sum of all n individual escape path rates: $k_{tot} = \sum_i^n k_i$.
- The escape-time probability distribution is $p(t) = k_{tot} e^{-k_{tot} t}$.
- If we run M replicas, we have Mn escape paths with a total escape rate Mk_{tot}, whose probability distribution function is

Slide 35

$$p(t_1) = M k_{tot} e^{-M k_{tot} t_1}$$

where t_1 is the simulation time accumulated on any one processor.

- Rewrite this using $t_{sum} \equiv M t_1$:

$$p(t_{sum}) = \frac{1}{M} \; p\left(\frac{t_{sum}}{M}\right) = k_{tot} \; e^{-k_{tot} t_{sum}}$$

where the left equality comes from $p(t/M)dt/M = p(t)dt$.

Parallel Replica Dynamics: Mathematical Foundations

- We therefore obtain the correct escape-time distribution,

$$p(t_{sum}) = k_{tot} \; e^{-k_{tot} t_{sum}}$$

- The same analysis is valid for inequivalent processors; M is replaced by a real-valued quantity (the sum of relative processor speeds), and the summed time becomes $t_{sum} = \sum_i^M t_i$.

Slide 36

- Only infrequent event dynamics was assumed; TST is not required.
- Recrossing or other correlated events are allowed to occur by continuing the dynamics on the processor which detected a transition for an additional time $\Delta t_{corr} \simeq 2$ ps.

Further Developments

Multiplicative speedups may be obtained by running accelerated MD trajectories on each replica, e.g. using hyperdynamics:

Slide 37

- Ag bulk vacancy diffusion
- 16 Pentium processors (200 MHz)
- Ramped bias potential
- 1.9 ms simulation time in 65 hours (10 events)

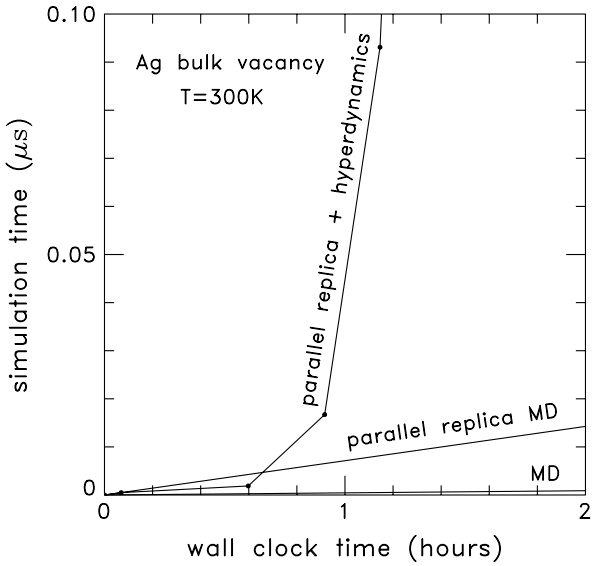

ILLINOIS — Los Alamos NATIONAL LABORATORY

Cu/Cu(100) epitaxial growth

The combined parallel replica hyperdynamics has been demonstrated on epitaxial surface growth at experimental deposition rates:

Slide 38

- 100 K and 120 K
- 10 monolayers (ML) per second
- 1200 ASCI Red (333 MHz Pentium II) processors
- 24 hours wall-clock time
- 0.3 sec boosted simulation time (3 ML)

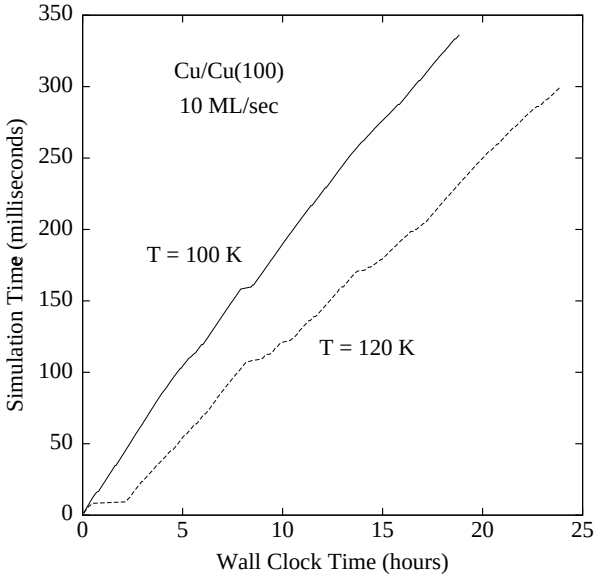

Ramping of bias potential may have been overly aggressive…

ILLINOIS — Los Alamos NATIONAL LABORATORY

Other Applications of Parallel Replica Dynamics

- Since the only assumption is infrequent events, the parallel replica method has potential uses beyond those considered here (entropic bottlenecks, non-atomistic systems, . . .)
- Folding@home: V.S. Pande Group (Stanford)
  - Some proteins are known to exhibit an exponential distribution of time-to-folding.
  - The fastest proteins fold on a μs timescale
  - Use free energy monitoring to detect transition
  - Widely distributed and heterogeneous processors (like SETI@home)
- SurfaceDiffusion@home, anyone?

Slide 39

Temperature Accelerated Dynamics

Again, the basic concept is rather simple:

- Run a trajectory at an elevated temperature T_{high}, but reject all attempted transitions (more on how this is done in a bit. . .).
- Extrapolate the high-temperature escape times $\{t_{i,high}\}$ to times $\{t_{i,low}\}$ at the lower temperature T_{low}, using the Arrhenius relationship and measured activation energies $\{E_a^i\}$.
- Accept the transition with the shortest $t_{i,low}$.

Slide 40

The key question is how long one needs to run the T_{high} simulation. . .

High-Temperature Sampling

- Vineyard (harmonic TST):

$$k^{hTST}(T) = n_{path} \frac{\prod_i^{3N} \nu_i^{min}}{\prod_i^{3N-1} \nu_i^{sad}} e^{-E_a/k_B T}$$

Slide 41

- Competition between entropic prefactor and activation energy E_a
- High T will tend to favor pathways with soft ν_i^{sad} (e.g. A $\rightarrow$ B) over those with small barriers (e.g. A $\rightarrow$ C):

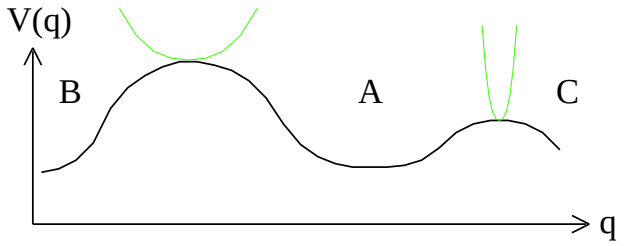

ILLINOIS Los Alamos NATIONAL LABORATORY

Basin-Confined Trajectories

Slide 42

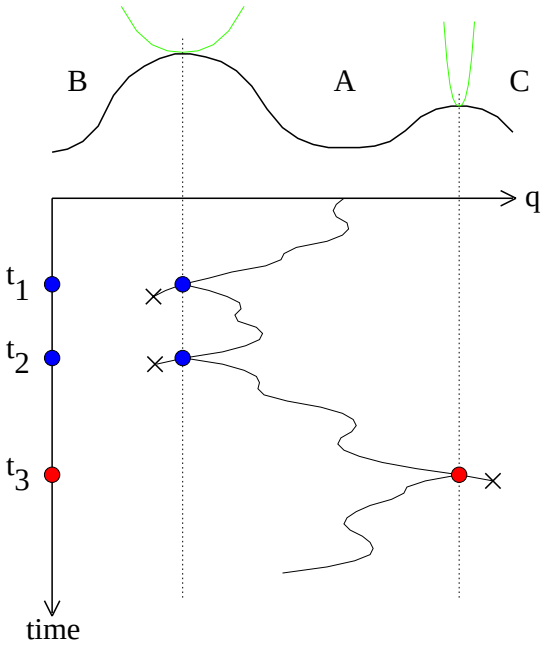

- Monitor trajectory for transitions (using periodic quench)
- When a transition occurs for the first time, carefully find its saddle point to calculate the activation energy E_a.
- Record transition time $t_{i,high}$ and E_a^i.
- Reflect trajectory back into basin A and continue.
- This gives $\{t_{i,high}\}$ properly sampled from an exponential distribution for each pathway.

ILLINOIS Los Alamos NATIONAL LABORATORY

Finding the Saddle Point

Slide 43

- Periodically save configurations as the trajectory proceeds
- When a transition is detected, find the two saved configurations which bracket the dividing surface

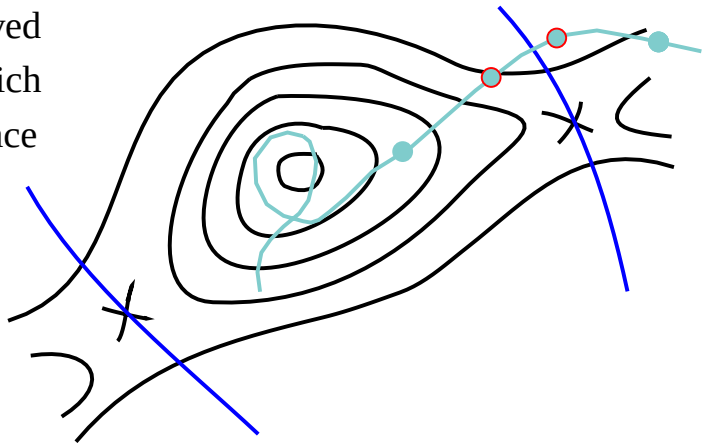

- Use the nudged elastic band method (H. Jónsson *et al.*, 1998) to optimize the minimum energy path and locate the saddle point.
- Check that steepest descent in each direction leads to the right basins.

TAD: Mathematical Foundations

Slide 44

- Assume infrequent events with first-order kinetics, i.e. a probability distribution for each escape path i given by $p(t_i) = k_i e^{-k_i t}$.
- Assume harmonic TST, with an Arrhenius form $k_i = \nu_i e^{-E_a^i / k_B T}$.
- Finally, assume that the prefactors ν_i have a lower bound, ν_{min}.
- Run basin-constrained simulation at a high temperature T_{high} to obtain a sequence of waiting times $\{t_{i,high}\}$.
- Extrapolate these to a sequence of low-temperature waiting times $\{t_{i,low}\}$ using the Arrhenius relationship:

$$\{t_{i,low}\} = \{t_{i,high}\} \exp\left[E_a^i(\beta_{low} - \beta_{high})\right]$$

  where $\beta \equiv 1/k_B T$.

- The transition with the shortest $t_{i,low}$ is accepted and the T_{low} clock is advanced by this time.

The Arrhenius view

$$k_i = \nu_i e^{-E_a^i / k_B T}$$

Slide 45

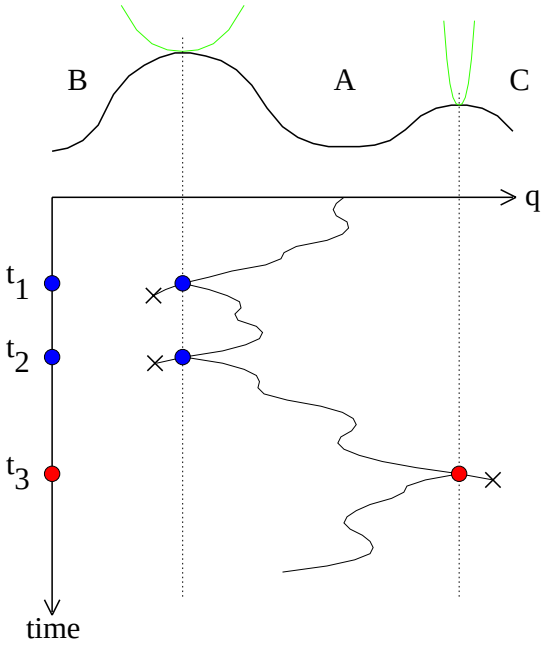

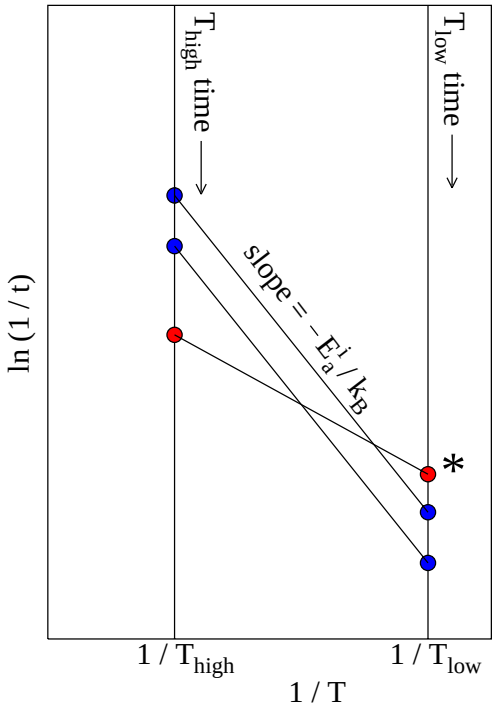

When can we stop?

Slide 46

- The average first occurrence time for event i is $\tau_i \equiv 1/k_i$
- There is probability δ that event i will first occur after $t = \tau_i \ln(1/\delta)$
- We assumed ν_{min}
- With confidence $1 - \delta$, all transitions with slope (i.e., E_a) less than the line with intercept $\nu_{min}^{eff} = \nu_{min} / \ln(1/\delta)$ have occured.

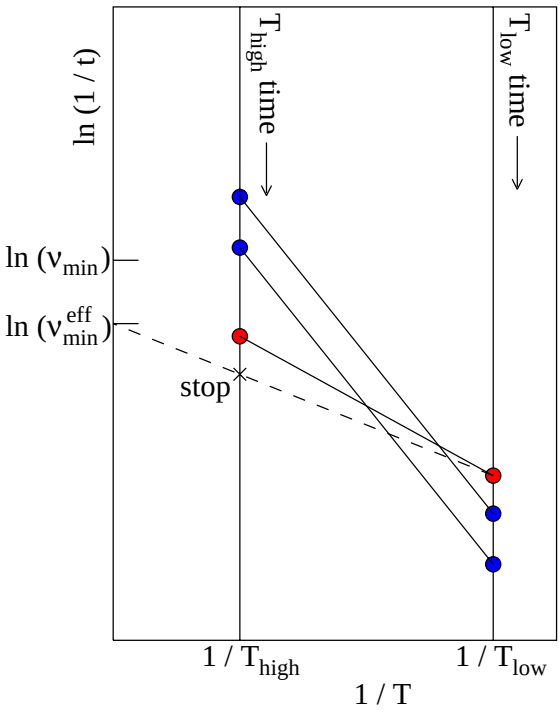

$$t_{high}^{stop} = (\nu_{min}^{eff} t_{low}^{short})^{T_{low}/T_{high}} / \nu_{min}^{eff}$$

Synthetic kMC Mode for the Low Barrier Problem

Slide 47

- Keep cumulative history of states.
- When a fast process has been attempted a number of times, estimate the rate. Put the process in the *synthetic* class.
- Exclude synthetic transitions from the time-extrapolation scheme.
- Generate random escape times for synthetic processes from exponential distribution using estimated rate.

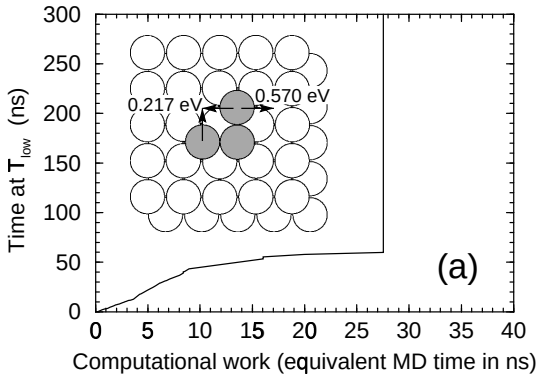

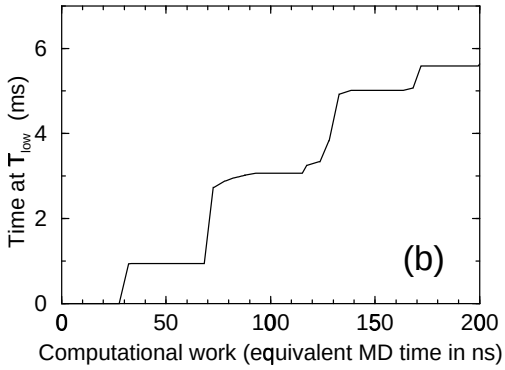

TAD Example: 1/4 ML Cu/Cu(111) at 150 K†

Slide 48

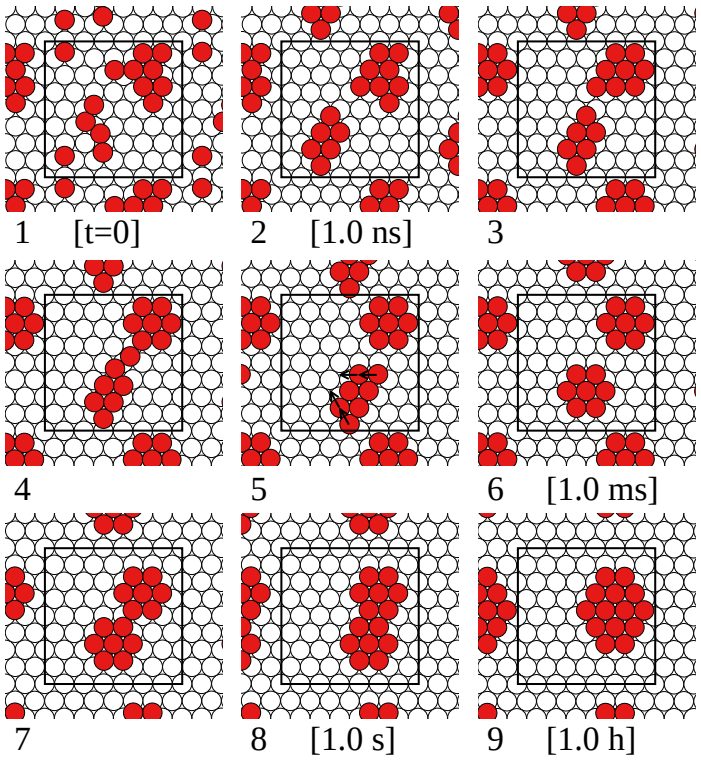

† M.R. Sørensen and A.F. Voter, J. Chem. Phys. **112**, 9599 (2000).

Recap: Comparison of Different Methods

Hyperdynamics

- assumes infrequent events, TST
- requires construction of a valid bias potential
- low barriers are a problem

Parallel Replica Dynamics

Slide 49

- only assumes infrequent events
- can be combined with hyperMD (done) or TAD (harder)
- full parallel efficiency if $\tau_{rxn}/n_{proc} \gg \tau_{dephase}, \tau_{corr}$

Temperature Accelerated Dynamics

- assumes infrequent events, harmonic TST, minimum prefactor
- synthetic (kMC) mode reduces low-barrier problem
- probably the most powerful of the three, given its assumptions

Application to Thin Film Growth

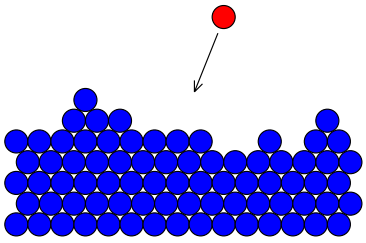

Slide 50

- Use embedded atom method (EAM) potentials for Cu, Ag, ...
- Periodic boundary conditions in lateral directions
- For each deposition event, select a random position above the surface and run ordinary MD for 2 ps
- Use accelerated dynamics (hyper + parallel, TAD, ...) to follow activated processes between depositions (μsec to sec)
- Study morphology, mechanisms, etc. as function of the deposition parameters (rate, temperature, impact energy, incidence angle, ...)

Application to Thin Film Growth

- Homoepitaxial growth at temperatures up to 100 K have been demonstrated using hyper + parallel replica dynamics (TCG and AFV) at rates up to 10 monolayers per second, and using on-the-fly kinetic Monte Carlo with the dimer method by Henkelman and Jónsson at 1 monolayer per millisecond.

Slide 51

- Work is in progress with Francesco Montalenti (LANL) and James Sprague (NRL) using TAD at *experimental* deposition rates (e.g., one monolayer every 15 seconds) for Cu/Cu(100), Ag/Ag(100), and Ag/Cu(100).
- All of these simulations exhibit a surprisingly large number of concerted, multiatom events which would be omitted from most (if not all) kinetic Monte Carlo rate catalogs.
- Stay tuned. . .

Concerted Smoothing Events

Several multiatom concerted events seen during Cu/Cu(100) growth:

Slide 52

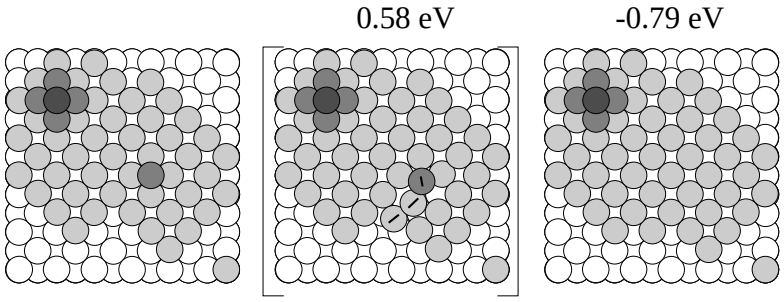

3-atom exchange step-down

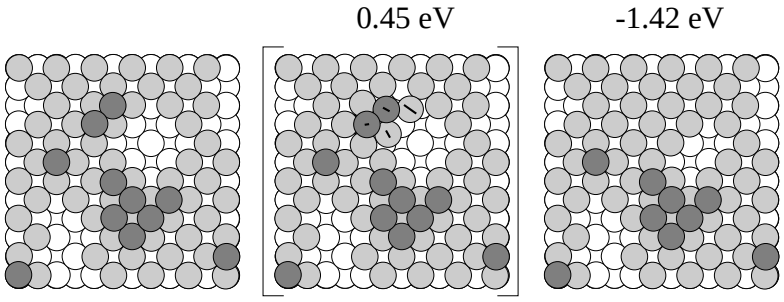

Dimer exchange step-down (4-atom)