Introduction to Accelerated Molecular Dynamics Methods

Danny Perez and Arthur F. Voter
Theoretical Division, T-12
Los Alamos National Laboratory

Uncertainty Quantification Workshop
April 25-26, 2008
Tucson, AZ

Acknowledgments

Blas P. Uberuaga (LANL, MST-8)
Francesco Montalenti (U. Milano-Bicocca)
Graeme Henkelman (U. Texas at Austin)
James A. Sprague (NRL)
Mads Sorensen (Novo Nordisk A/S, Copenhagen)
Sriram Swaminarayan (LANL, MST-8)
Steve Stuart (Clemson)
Roger Smith (U. Loughborough)
Robin Grimes (Imperial College)
Kurt Sickafus (LANL, MST-8)
Jacques Amar (U. Toledo)
Yunsic Shim (U. Toledo)
Yuri Mishin (George Mason U.)
Soo Young Kim (LANL postdoc, T-12)
Abhijit Chatterjee (LANL postdoc, T-12)

DOE Office of Basic Energy Sciences
LDRD (LANL Internal)
SCIDAC (DOE)

Outline

- The molecular dynamics (MD) timescale problem
- Infrequent-event systems
- Transition State Theory
- Accelerated molecular dynamics methods
  - Hyperdynamics
  - Parallel-replica dynamics
  - Temperature accelerated dynamics (TAD)
- Ongoing challenges and recent advances
- Conclusion

The MD time-scale problem

For many systems, we need to simulate with full atomistic detail.

Molecular dynamics (MD) (the integration of the atomistic equations of motion) can only reach nanoseconds to microseconds due to the stiffness of the equations of motion (timestep is limited to fs).

Processes we want to study often take much longer:

- vapor-deposited film growth (s)
- STM/AFM surface manipulation, nanoindentation (ms - s)
- bulk and surface diffusion processes
- radiation damage annealing (ns, μs, ms, s, …, years)
- protein folding (μs - s)

Such slowly evolving systems share a common feature: their long-time dynamics consists of infrequent jumps between different states (i.e., activated processes). The problem is that these systems are way too complex to map out completely. We thus cannot use Kinetic Monte Carlo to generate long-time trajectories.

Infrequent Event System

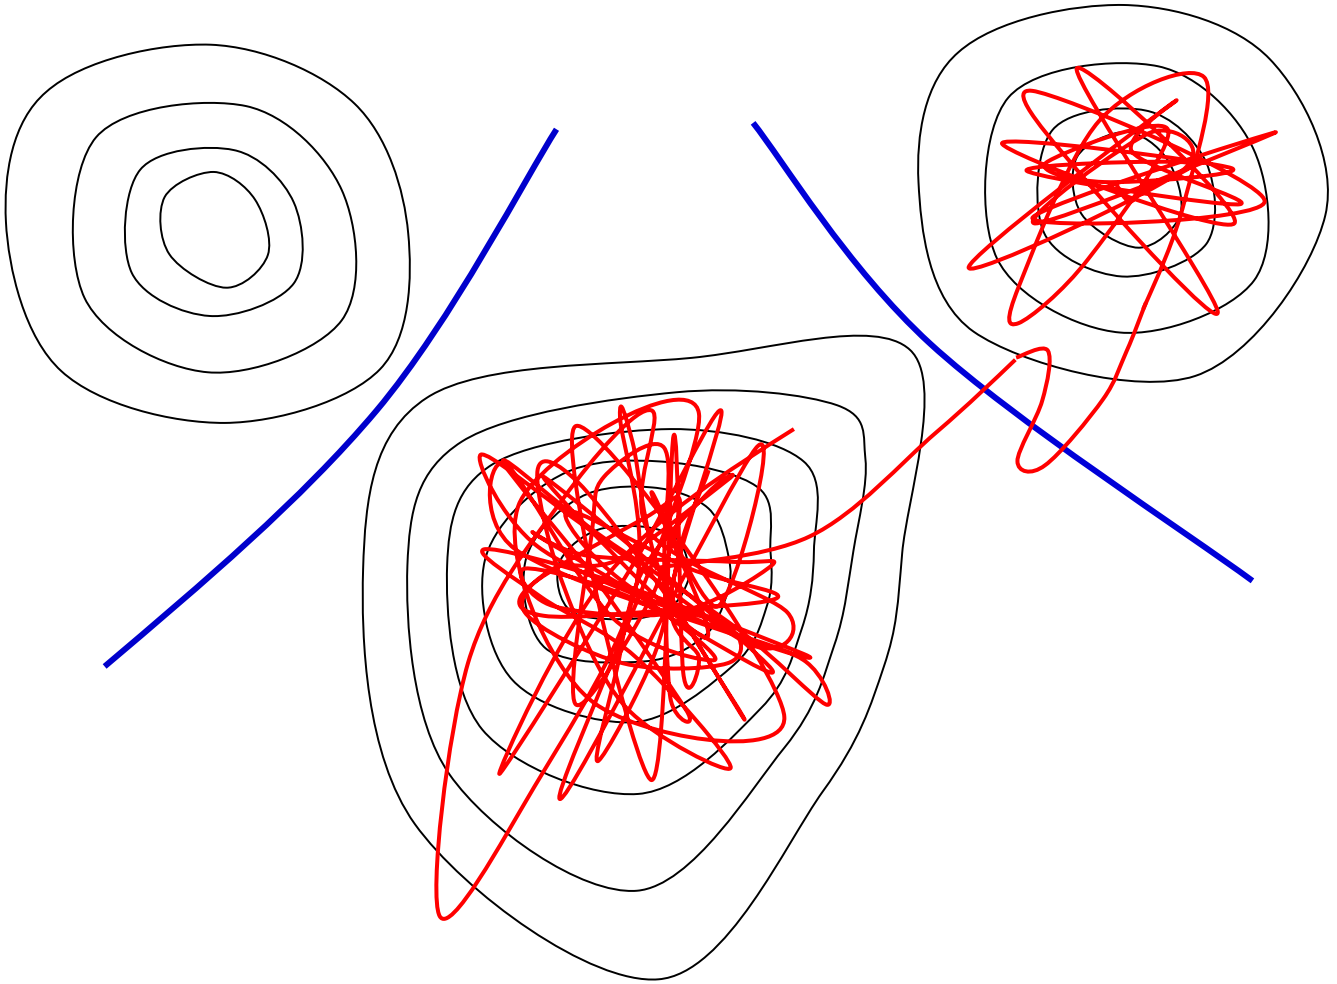

Indeed, the system vibrates in one of the 3N dimensional basins many times before finding an escape path. The trajectory finds an appropriate way out (i.e., proportional to the rate constant) without knowing about any of the escape paths except the one it first sees. Can we exploit this?

Transition State Theory (TST)

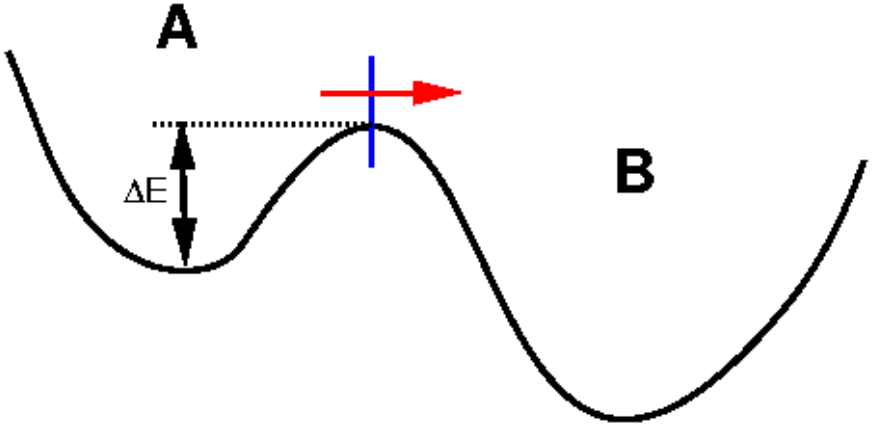

Marcelin (1915)
Eyring, Wigner,…

TST escape rate = equilibrium flux through dividing surface at x=q

$$k_{A\to B}^{TST} = \langle \delta(x-q)|\dot{x}|\rangle_A = Z^q / Z^A$$ (exact flux)

$$k_{A\to B}^{HTST} = \upsilon_0 e^{-\Delta E/k_B T}$$ (harmonic approx.)

- classically exact rate if no recrossings or correlated events
- no dynamics required
- excellent approximation for materials diffusion
- entails an exponential distribution of escape times

Accelerated molecular dynamics (AMD) concept

Let the trajectory, which is smarter than we are, find an appropriate way out of each state. The key is to coax it into doing so more quickly, using statistical mechanical concepts (primarily transition state theory).

With these AMD methods, we can follow a system from state to state, reaching time scales that we can't achieve with molecular dynamics. However, we have to sacrifice the short time dynamics to do so.

AMD methods are not sampling methods as they generate a single long state-to-state trajectory at the time. Often, even just one of these long trajectories can reveal key system behavior. If desired, we can go back through the trajectory to determine rates and properties in more detail, using conventional methods, and/or we can run more long trajectories to gather statistics.

Accelerated Molecular Dynamics Methods

Hyperdynamics (1997)

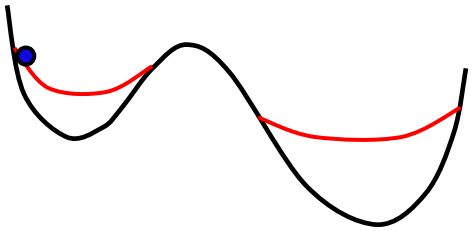

- Design bias potential that fills basins.
- MD on biased surface evolves correctly from state to state.
- Accelerated time is statistical quantity.

Parallel Replica Dynamics (1998)

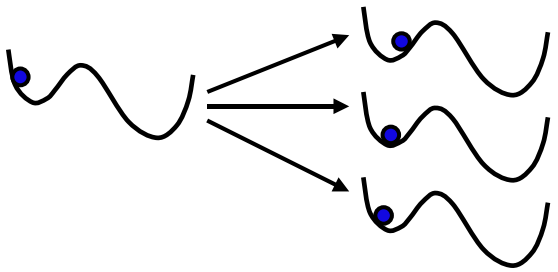

- Parallelizes time.
- Very general -- any exponential process.
- Gives exact dynamics.
- Boost requires multiple processors

Temperature Accelerated Dynamics (2000)

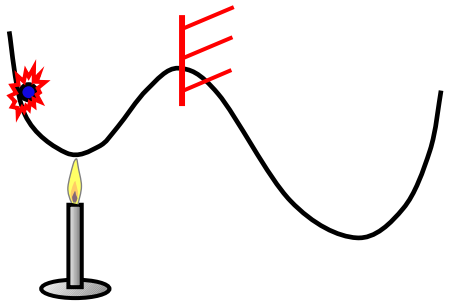

- Raise temperature of MD in this basin.
- Intercept and block every attempted escape.
- Accept event that would have occurred first at the low temperature.
- More approximate; good boost.

Characteristics of the AMD methods

- All three methods can give very large boost factors when events are very infrequent

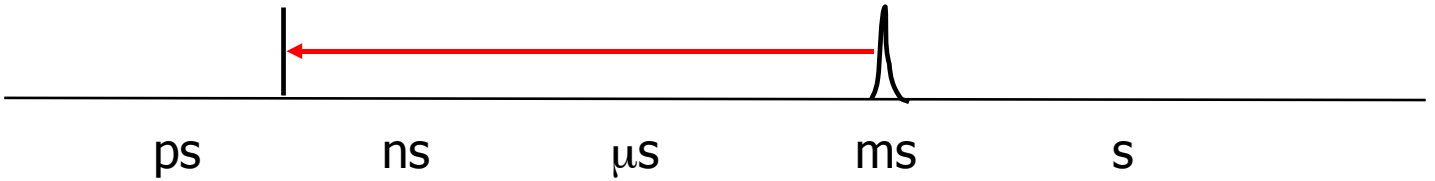

- Hyperdynamics
  - requires designing a valid and effective bias potential
  - assumes TST holds (no recrossings)
  - no need to detect transitions
- Parallel Replica Dynamics
  - requires M processors for boost of M
  - most general, most accurate
  - can treat entropic bottlenecks
- Temperature Accelerated Dynamics
  - most approximations, but still fairly accurate
  - assumes harmonic transition state theory

Parallel-Replica Dynamics

Parallel Replica Dynamics

Concept: Follow many replicas of the system on a parallel computer to parallelizes time evolution

Assumptions:

- exponential distribution of first-escape times

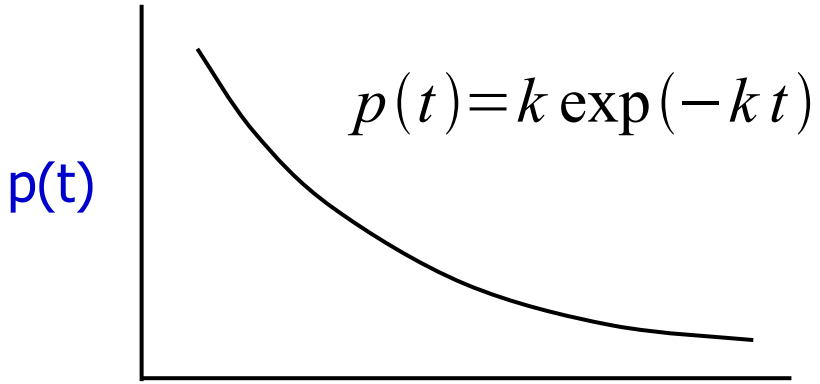

Must know:

- how to detect transitions
- correlation time

AFV, Phys. Rev. B, 57, R13985 (1998)

Parallel Replica Dynamics Procedure

Replicate entire system on each of M processors.

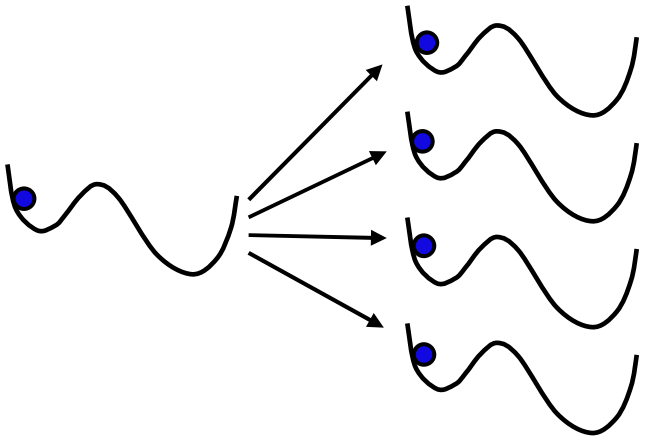

Parallel Replica Dynamics Procedure

Randomize momenta independently on each processor.

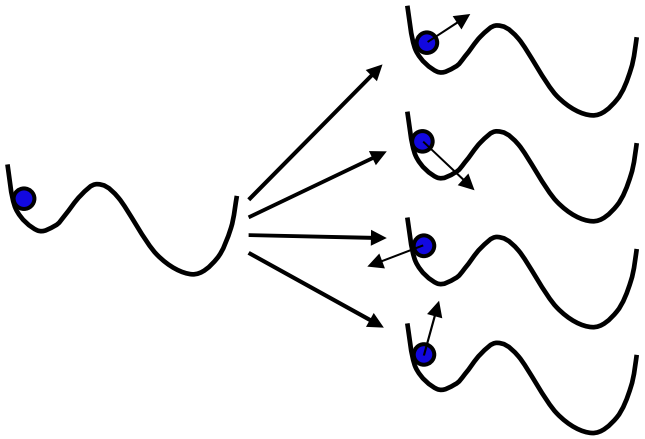

Parallel Replica Dynamics Procedure

Run MD for short time ($\tau_{dephase}$) to dephase the replicas.

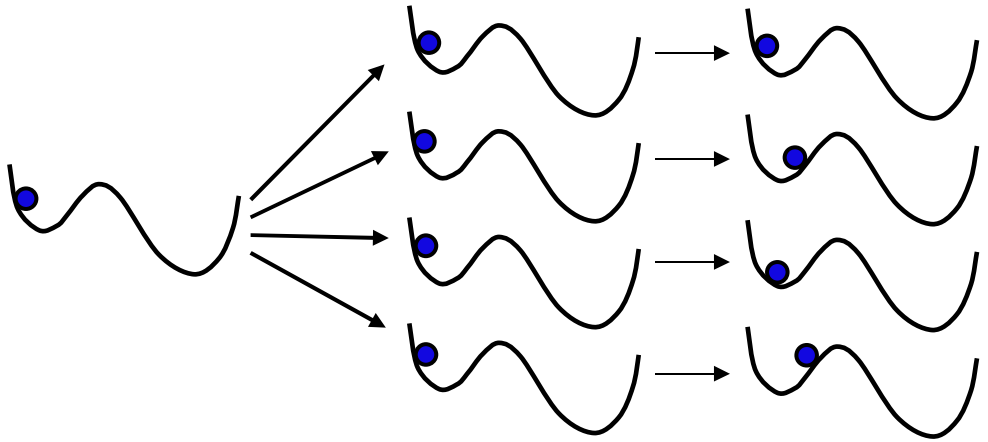

Parallel Replica Dynamics Procedure

Start clock and run thermostatted MD on each processor.
Watch for transition…

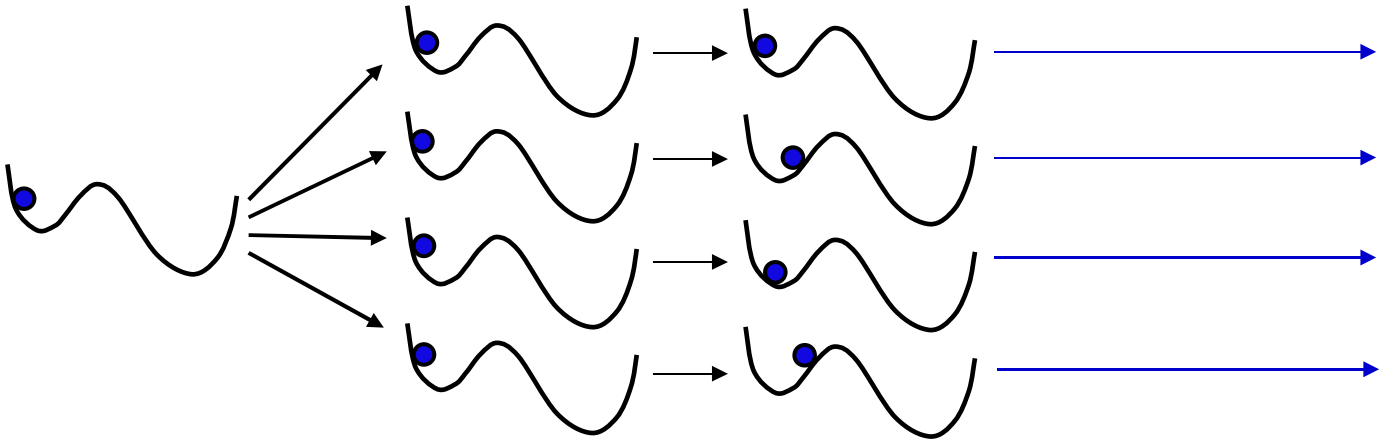

Parallel Replica Dynamics Procedure

Stop all trajectories when first transition occurs on *any* processor.

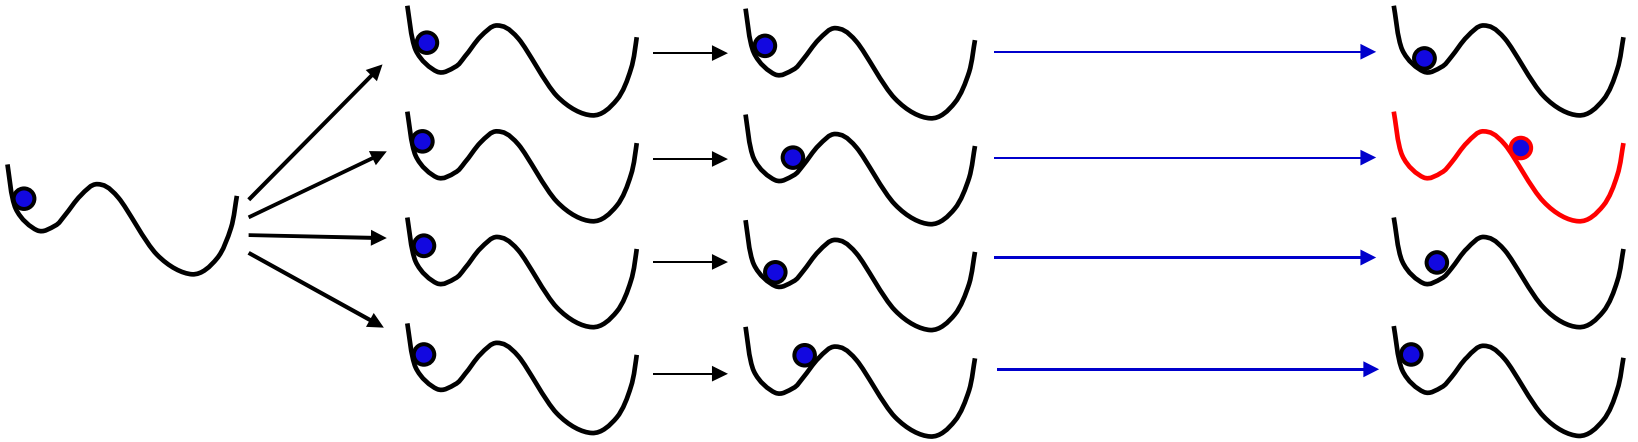

Parallel Replica Dynamics Procedure

Sum the trajectory times over all M processors. Advance simulation clock by this t_{sum}

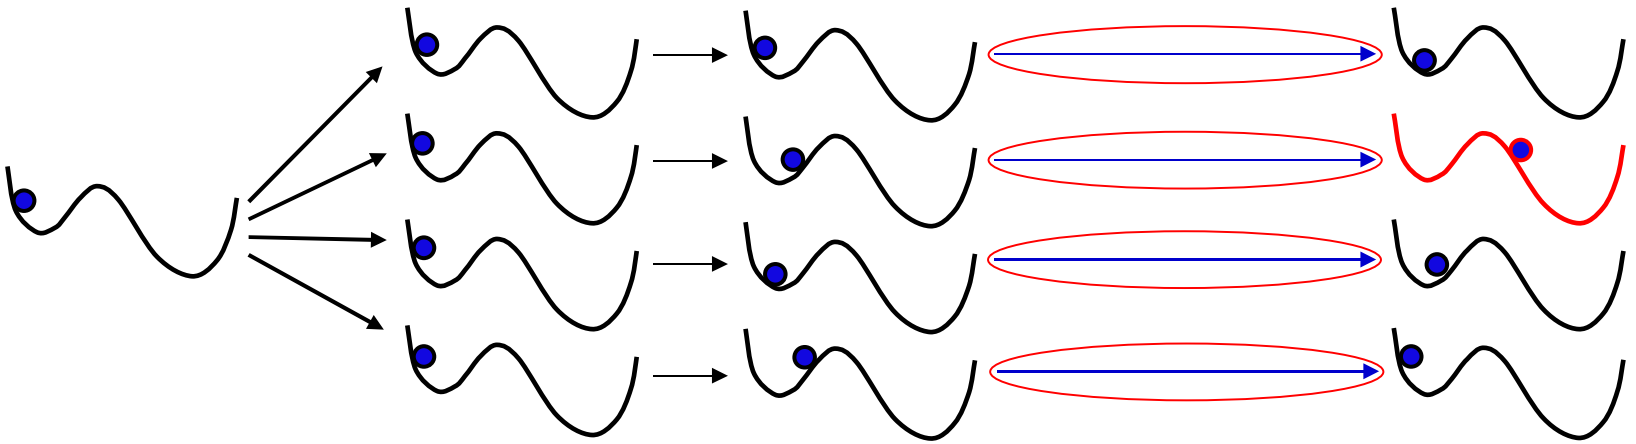

Parallel Replica Dynamics Procedure

On the processor where a transition occurred, continue trajectory for a time τ_{corr} to allow correlated dynamical events.

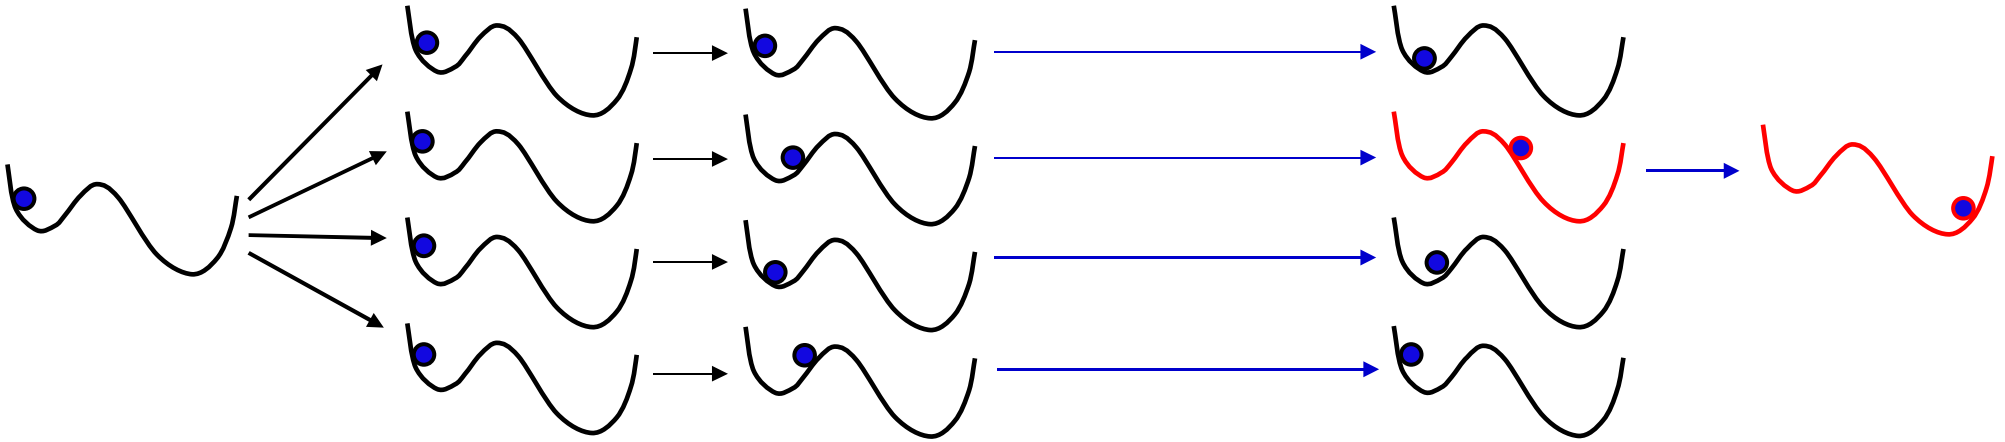

Parallel Replica Dynamics Procedure

Advance simulation clock by τ_{corr}.

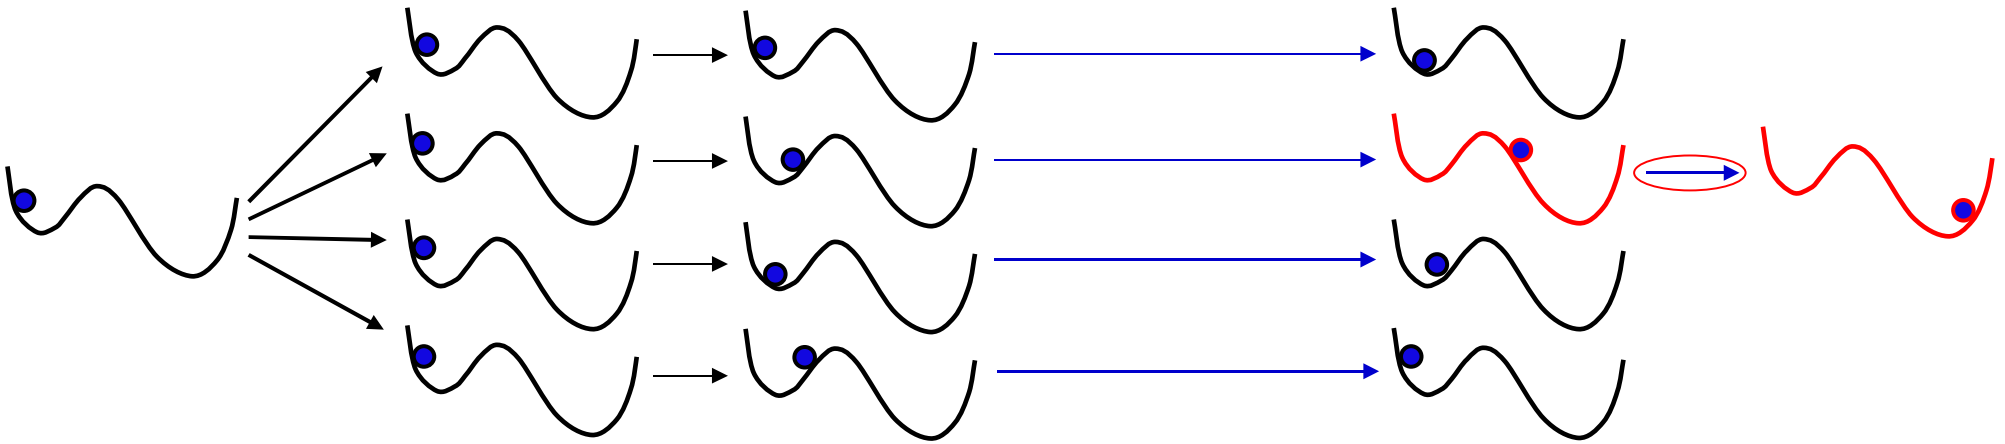

Parallel Replica Dynamics Procedure

Replicate the new state and begin procedure again.

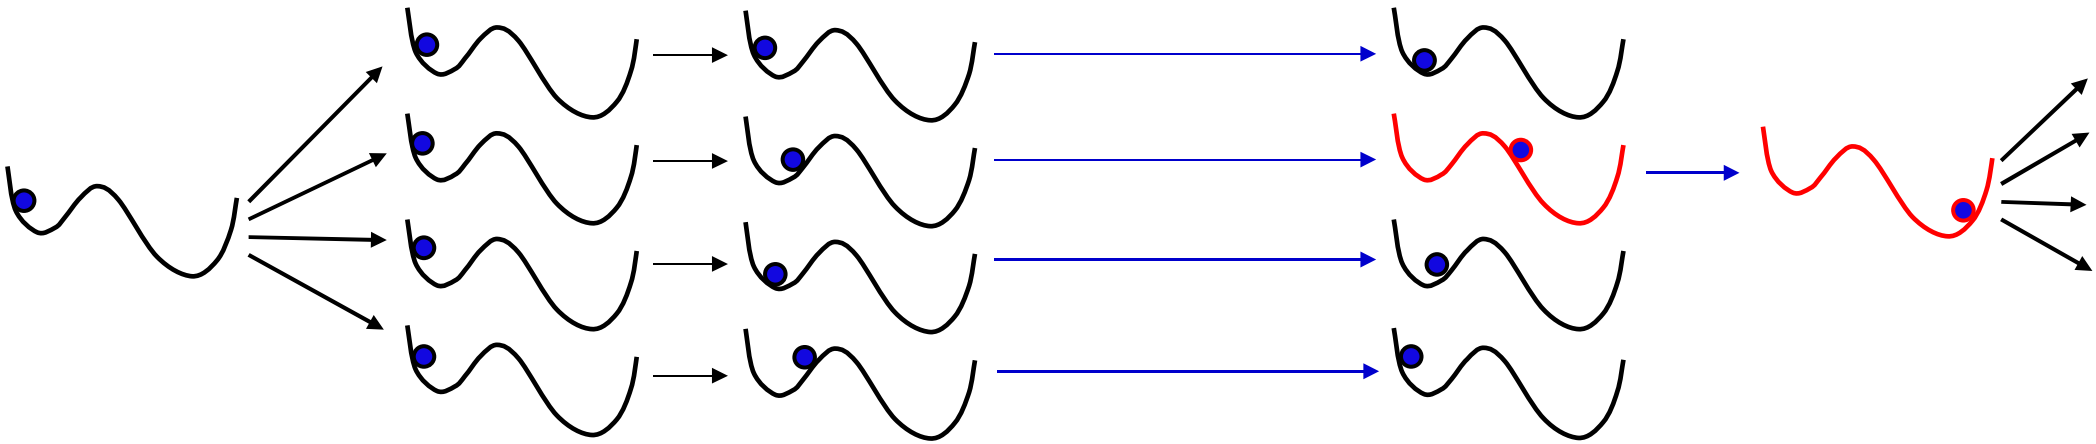

Long time annealing of 20 vacancy void in Cu

- EAM Copper
- Parallel-replica simulation of 20-vacancy void annealing at T=400 K
  - 20 vacancies is one too many for "perfect" void
- Total simulation is 7.82 μs
- At 1.69 μs, void transforms to SFT
- Equivalent single processor time: 1.3 years
- Very complex transition pathway

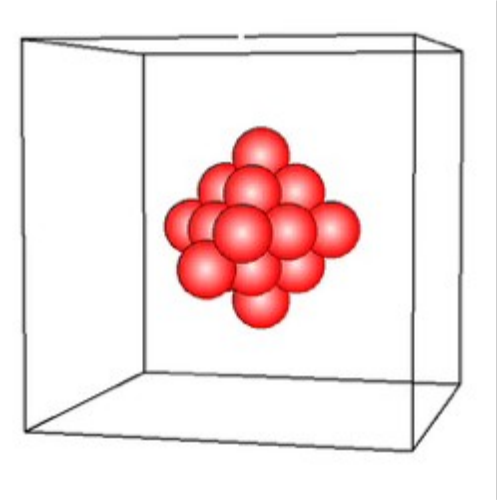

Red atoms=vacancies
Blue atoms=interstitials
Bulk atoms not shown

Completely new transformation pathway for the formation of stacking fault tetrahdera (SFT)

Uberuaga, Hoagland, Voter, Valone, PRL **99**, 135501 (2007)

Transformation pathway for 20 vacancy void

- Full path for transformation to SFT calculated with NEB
- Initial barrier is >2 eV
- Vineyard prefactor for first step is 2×10^{36} Hz !
- Driven by entropy increase as extra volume is made available to system as void collapses

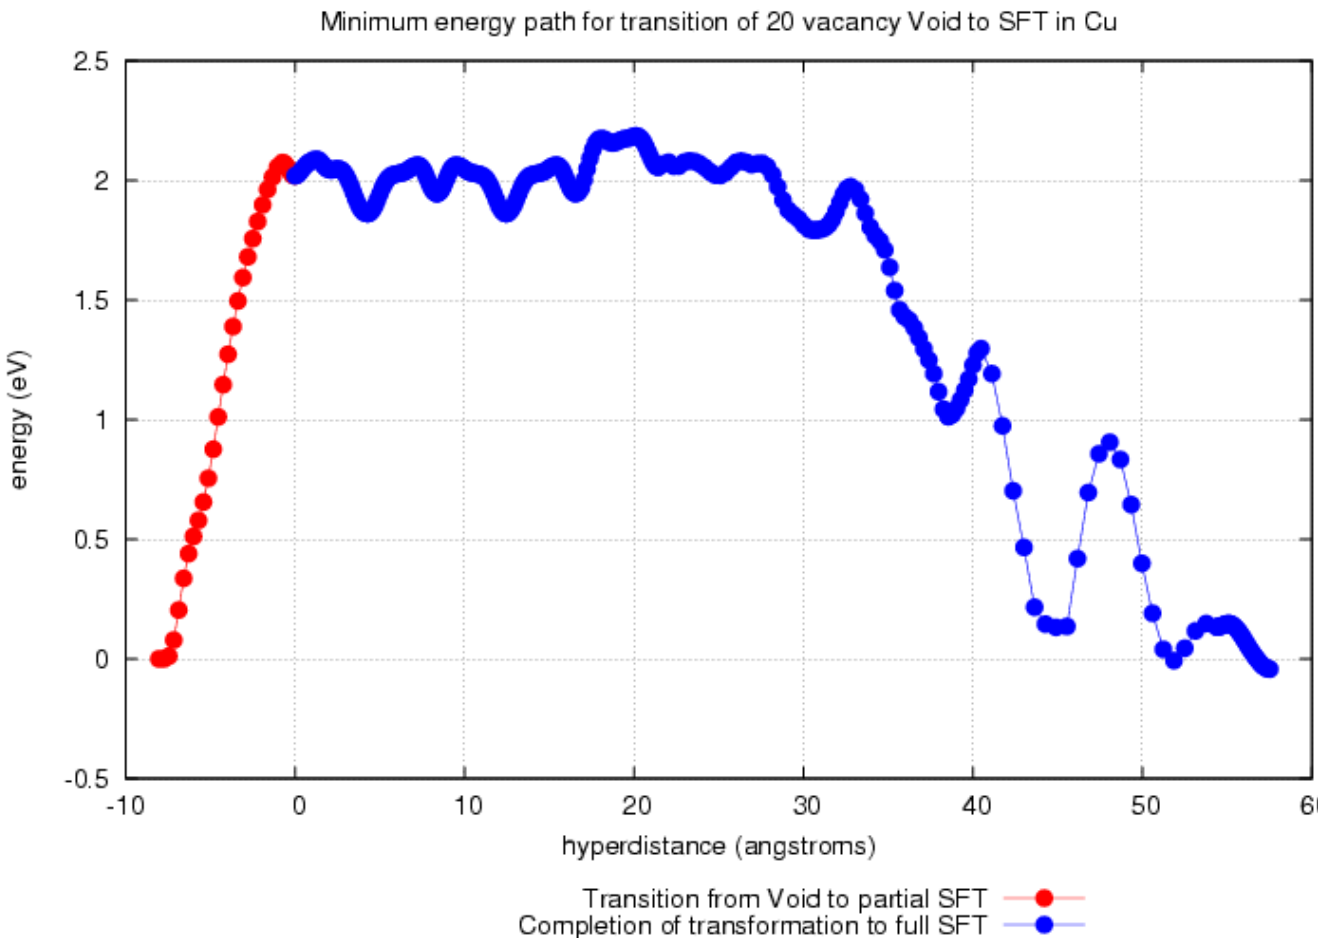

Uberuaga, Hoagland, Voter, Valone, PRL **99**, 135501 (2007)

Summary: Parallel Replica Dynamics

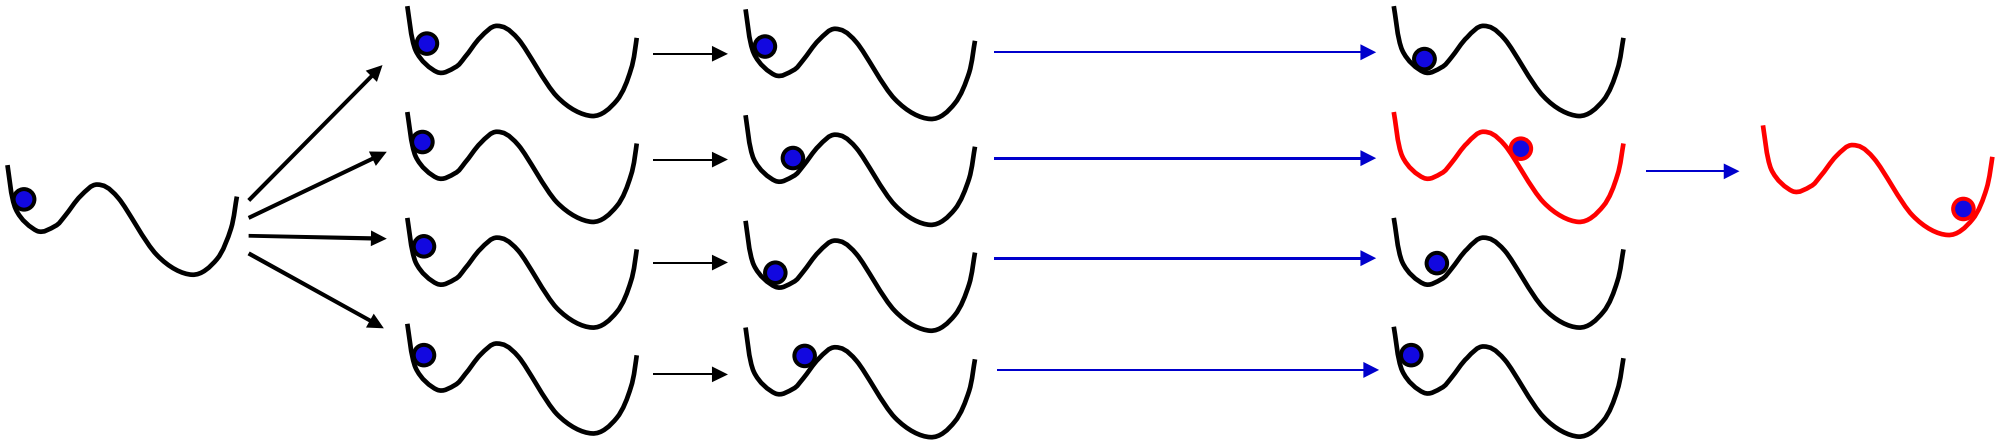

The summed time (t_{sum}) obeys the correct exponential distribution, and the system escapes to an appropriate state.

State-to-state dynamics are thus correct; τ_{corr} stage even releases the TST assumption [AFV, Phys. Rev. B, 57, R13985 (1998)].

Maximal boost is equal to M

Good parallel efficiency if $\tau_{rxn} / M >> \tau_{dephase} + \tau_{corr}$

Applicable to any system with exponential first-event statistics

Hyperdynamics

Hyperdynamics

Concept: Fill the basins with a bias potential to increase the rate of escape and renormalize the time accordingly.

Assumptions:

- transition state theory (no recrossings)

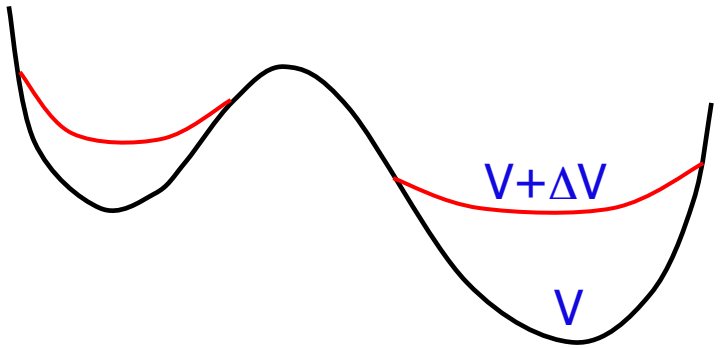

Procedure:

- design bias potential ΔV which is zero at **all** dividing surfaces so as not to bias rates along different pathways.
- run thermostatted trajectory on the biased surface ($V+\Delta V$)
- accumulate hypertime as

$$t_{hyper} = \Sigma \Delta t_{MD} \exp[\Delta V(R(t))/k_B T]$$

Result:

- state-to-state sequence correct
- time converges on correct value in long-time limit (vanishing relative error)

AFV, J. Chem. Phys. 106, 4665 (1997)

The hypertime clock

System coordinate

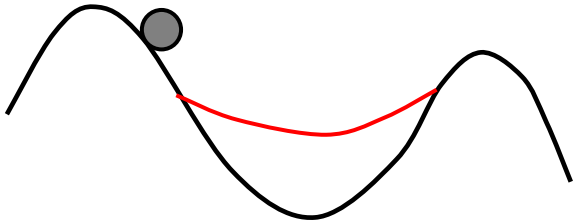

MD clock

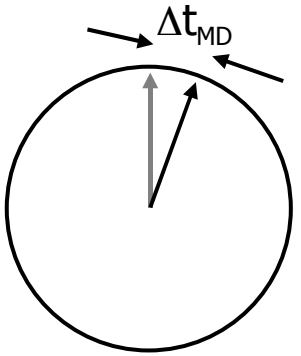

hypertime clock

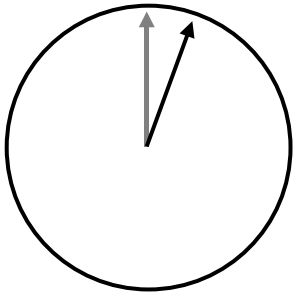

The hypertime clock

System coordinate

MD clock

hypertime clock

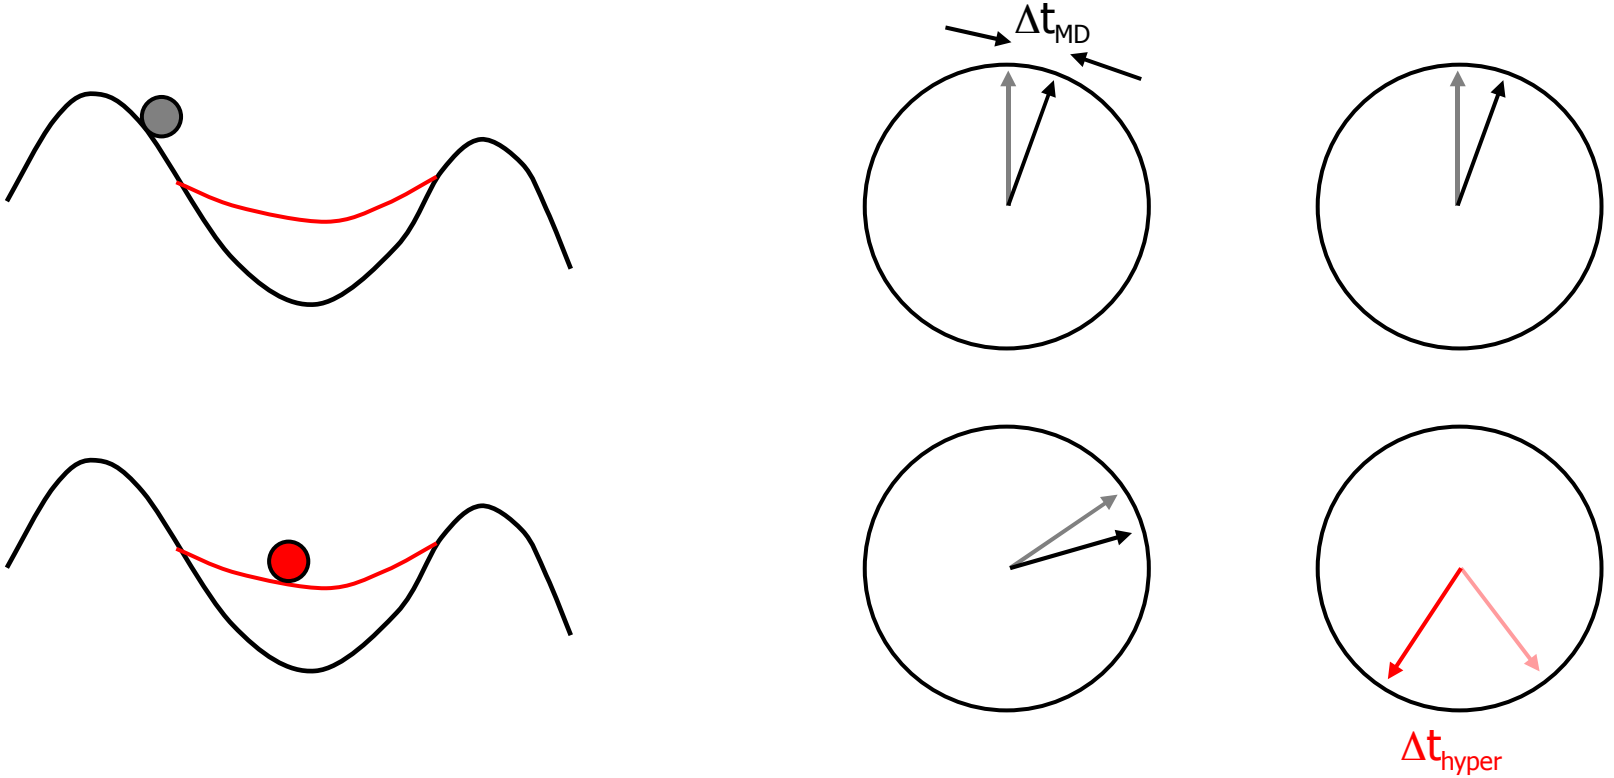

The hypertime clock

System coordinate

MD clock

hypertime clock

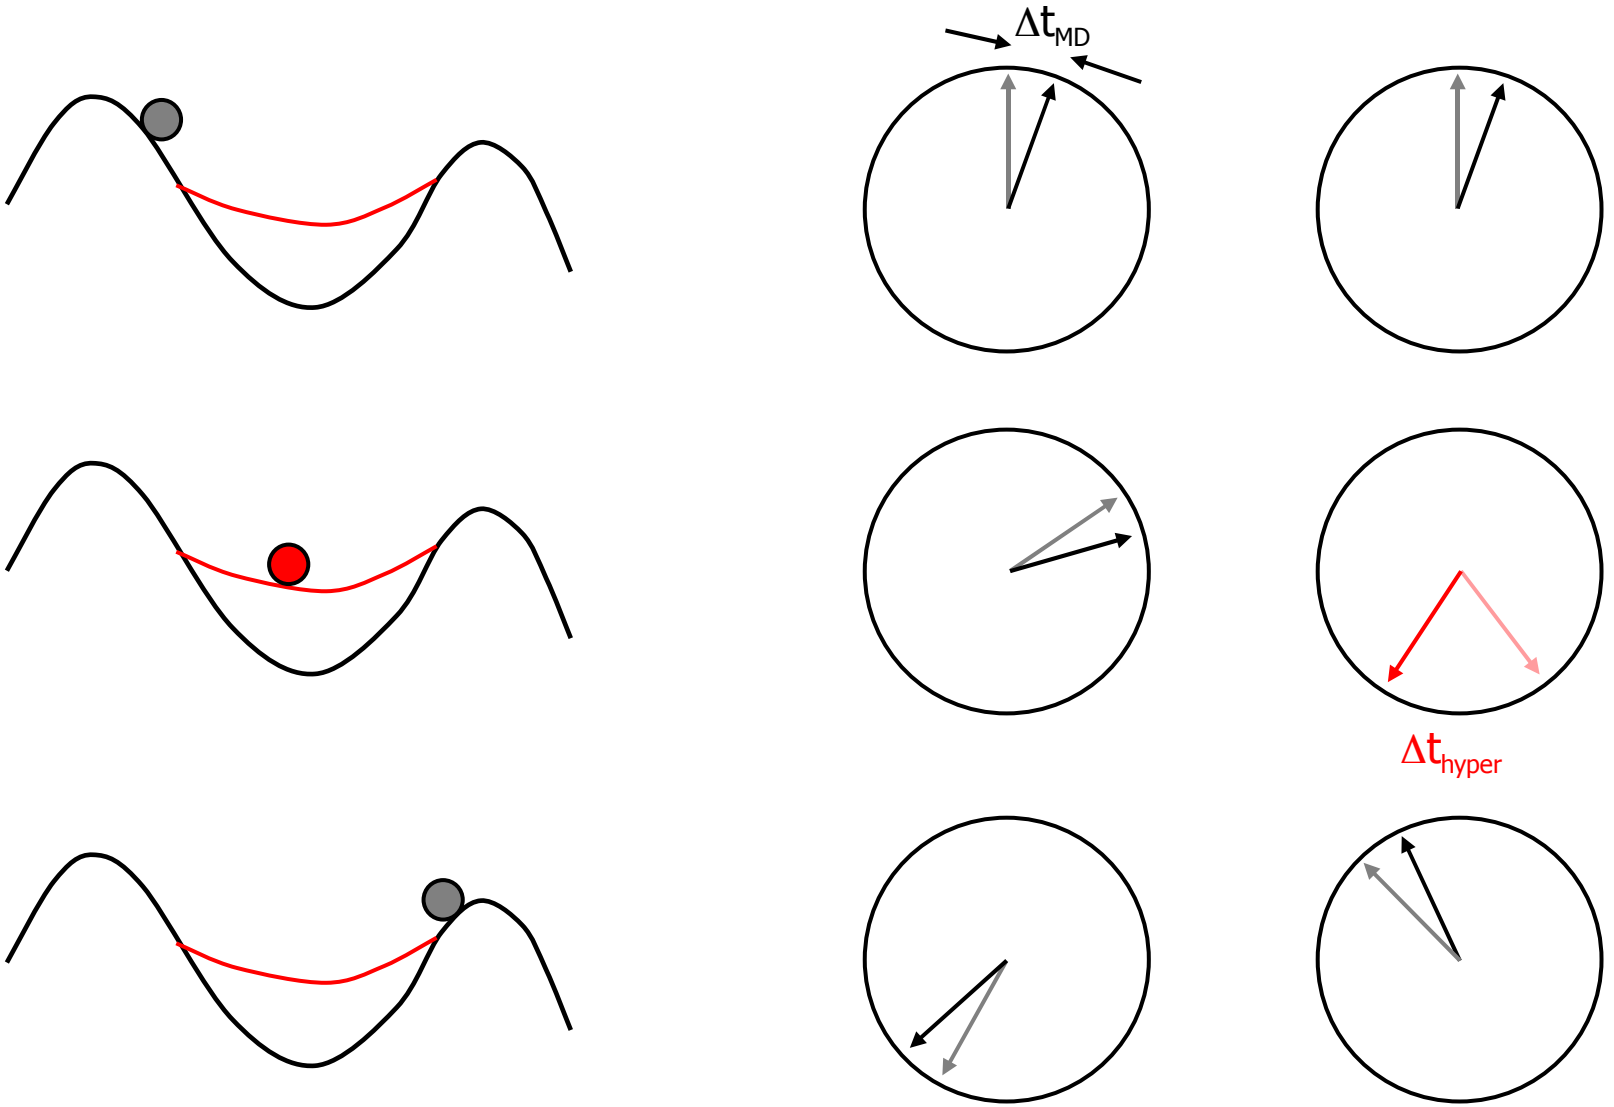

Boost = hypertime/(MD clock time)

Hyperdynamics

Key challenge is designing a bias potential that meets the requirements of the derivation and is computationally efficient. This is very difficult since we do not have any a priori information about neighboring states nor about the dividing surfaces in between them. Futher, we have to work in very high dimension.

A few forms have been proposed and tested. Still a subject of ongoing research…

We recently proposed a self-learning version of the Bond-Boost potential of Miron and Fichthorn that automatically adapts to the system at hand, thus requiring no a priori parametrization.

For discussion, see

Voter, Montalenti, and Germann, Ann. Rev. Mater. Res. 32, 321 (2002)

Hyperdynamics bias potential

An extremely simple form: flat bias potential

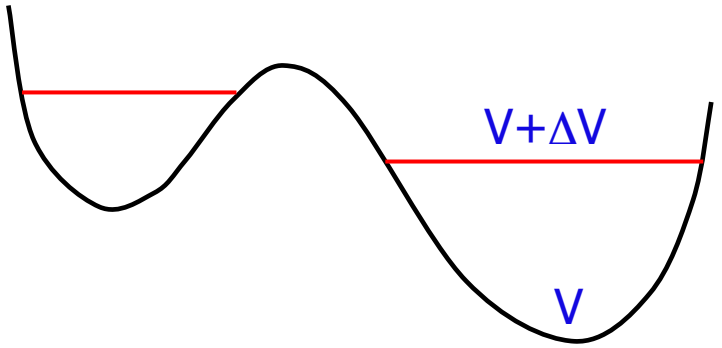

M. M. Steiner, P.-A. Genilloud, and J. W. Wilkins, Phys. Rev. B **57**, 10236 (1998).

- no more expensive than normal MD (negative overhead(!))
- very effective for low-dimensional systems
- diminishing boost factor for more than a few atoms.

Hessian-based bias potential

Detect ridgetop using local approximation of Sevick, Bell and Theodorou (1993),

$$\varepsilon_1 < 0 \quad \text{and} \quad C_1 g = 0$$

(ε_1, C_1 = lowest eigenvalue, eigenvector of Hessian; g = gradient)

Design bias potential that turns off smoothly in proximity of ridgetop

Iterative method for finding ε_1 using only first derivatives of potential

Iterative method for finding $C_1 g$ and its derivative using only first derivatives of potential

Good boost, but very tight convergence required for accurate forces

AFV, Phys. Rev. Lett., **78**, 3908 (1997)

Bond-boost bias potential

R.A. Miron and K.A. Fichthorn J. Chem. Phys. **119**, 6210 (2003)

Assumes any transition will signal itself by significant changes in bond lengths

ΔV = sum of contributions from every bond

Envelope function forces ΔV --> 0 when any bond is stretched beyond some threshold value

Very promising:

- fairly general
- very low overhead
- for metal surface diffusion, boost factors up to 10^6

Simple bond-boost bias potential

(Danny Perez and AFV, to be published)

Simple:
Only the "most distorted" bond contributes to the bias potential at any time (captures the essence of the Miron-Fichthorn approach)

Self-learning:
Increases bias strength parameter for each bond on the fly in a way that ensures the hyperdynamics requirements are maintained -- maximum safe boost.

Hypertime increases exponentially with MD time during first few vibration periods; system quickly reaches maximum boost.

Bond-boost bias potential

Ag monomer on Ag (100) at T=300K: long time behavior

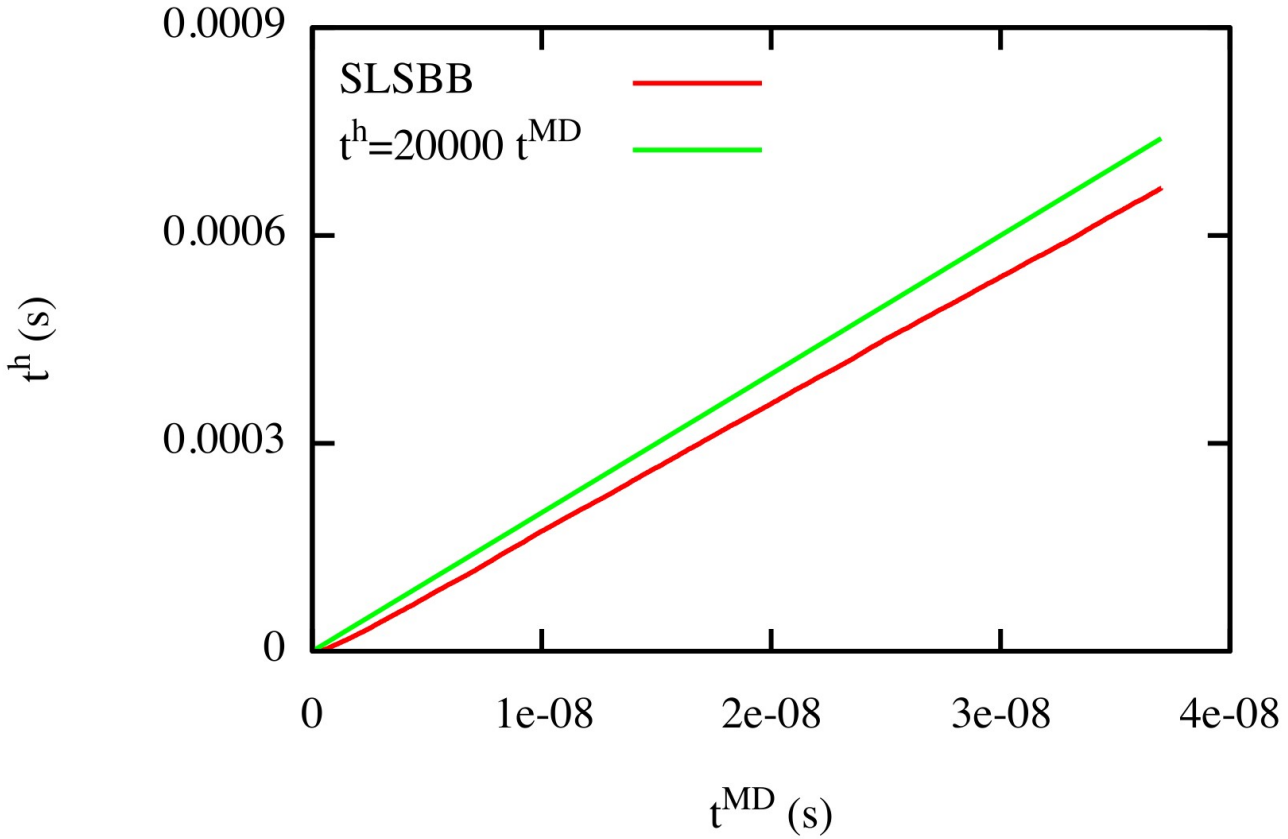

Bond-boost bias potential

Ag monomer on Ag (100) at T=300K: learning phase

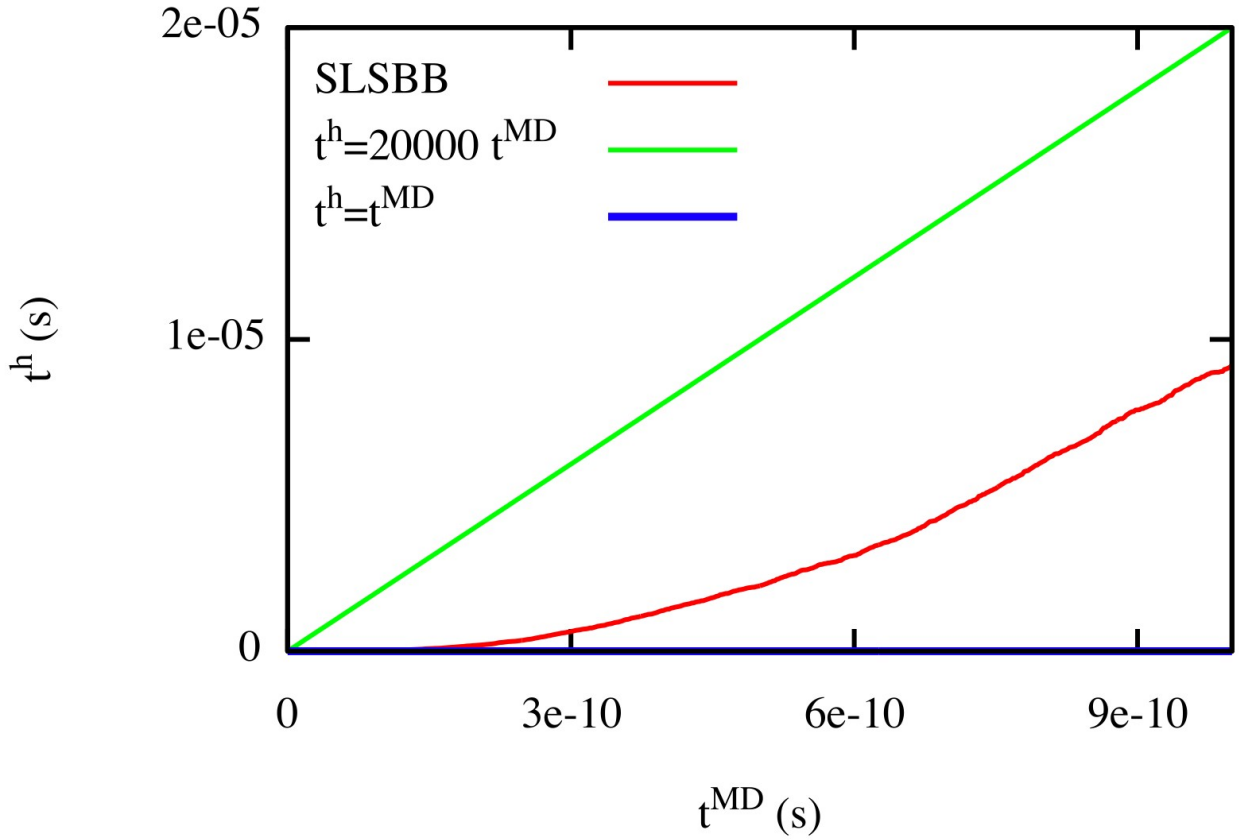

Summary - Hyperdynamics

Powerful if an effective bias potential can be constructed

Need not detect transitions

Boost factors climbs exponentially with inverse temperature (can reach thousands or even millions)

Especially effective if barriers high relative to T

Lots of possibilities for future development of advanced bias potential forms

TAD

Temperature Accelerated Dynamics (TAD)

Concept:

Raise temperature of system to make events occur more frequently.
Filter out the events that should not have occurred at the lower temperature.

Assumptions:

- infrequent-event system
- transition state theory (no correlated events)
- harmonic transition state theory (gives Arrhenius behavior)

$$k = \nu_0 \exp[-\Delta E/k_B T]$$

- all preexponentials (ν_0) are greater than ν_{min}

[Sørensen and Voter, J. Chem. Phys. 112, 9599 (2000)]

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

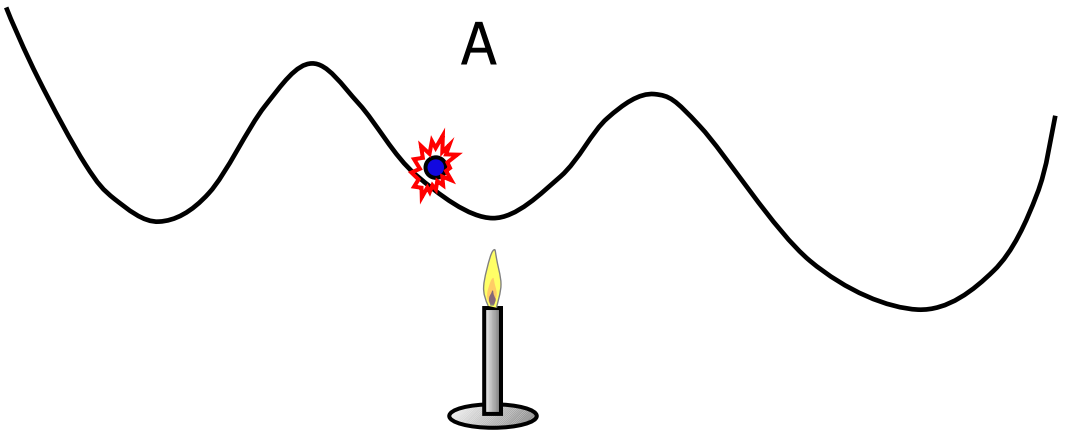

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

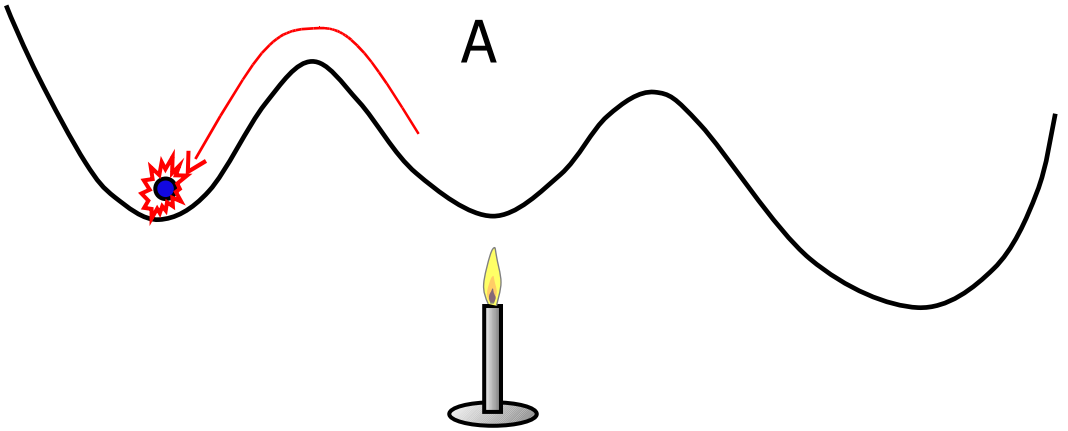

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

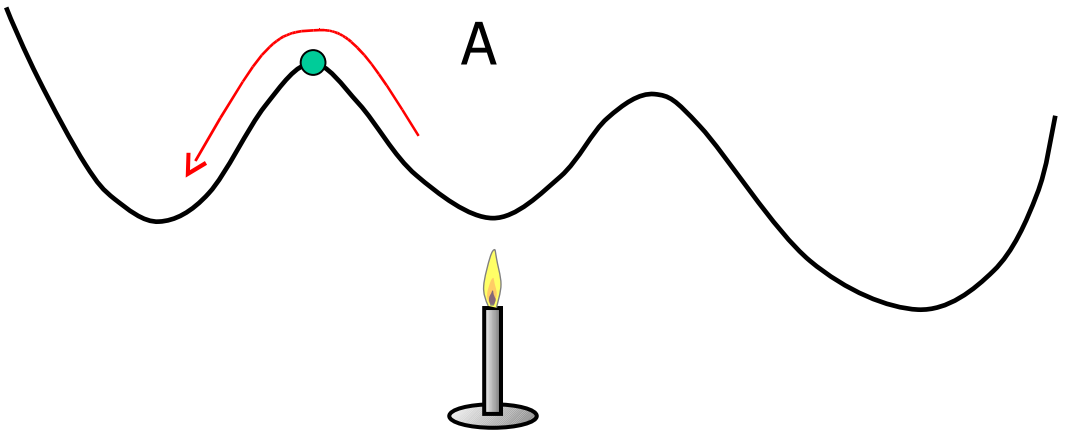

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

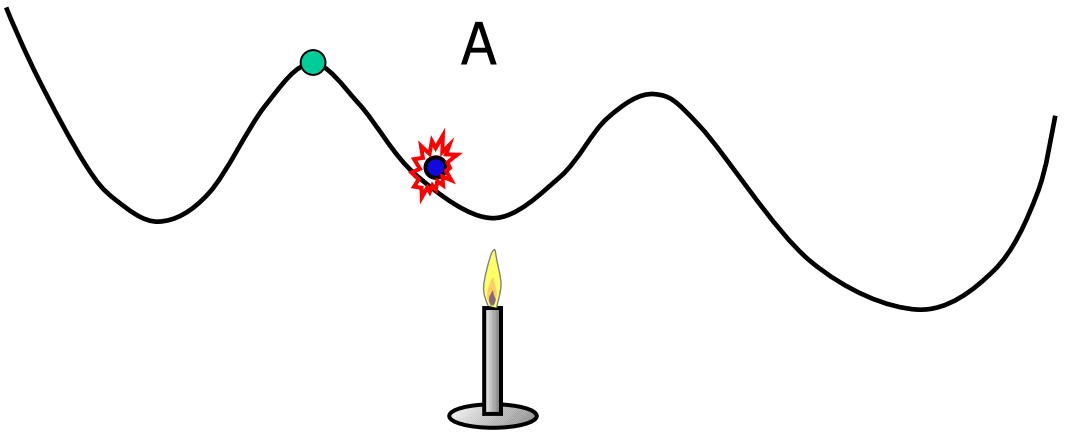

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

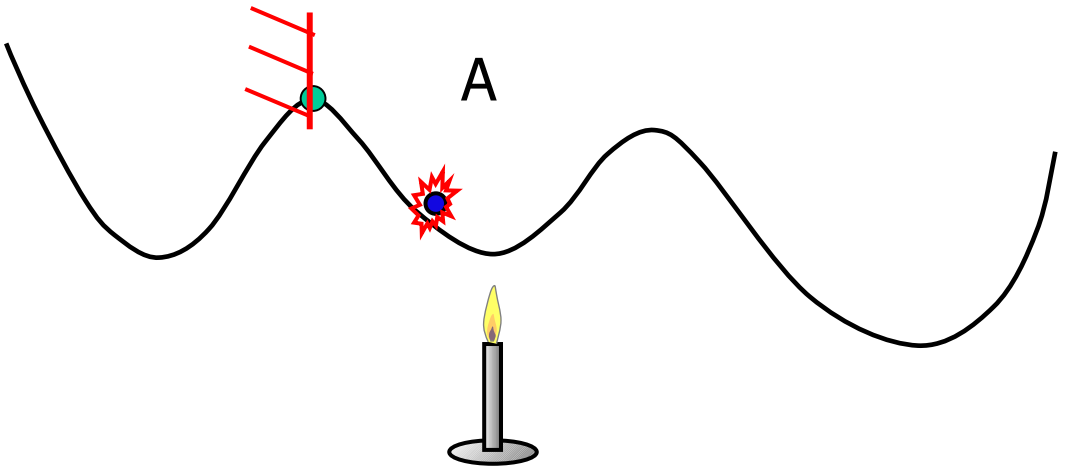

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

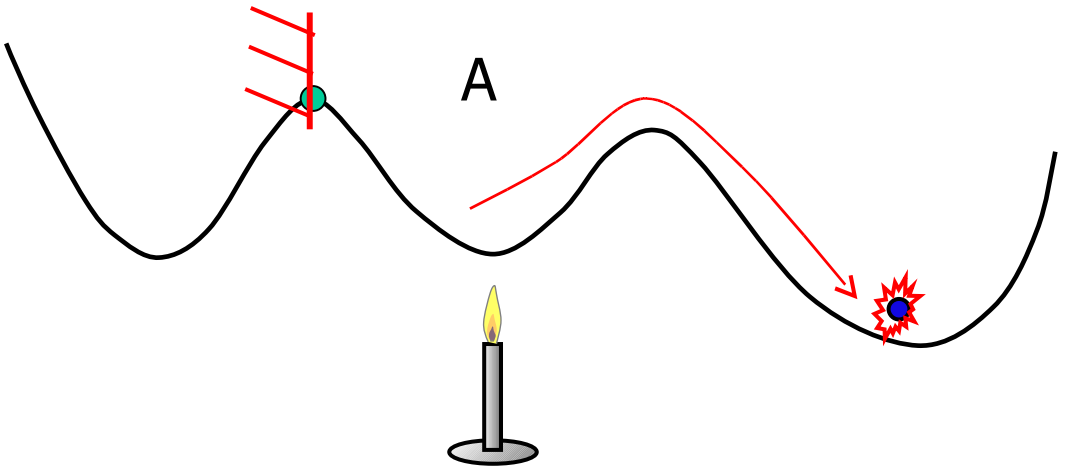

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

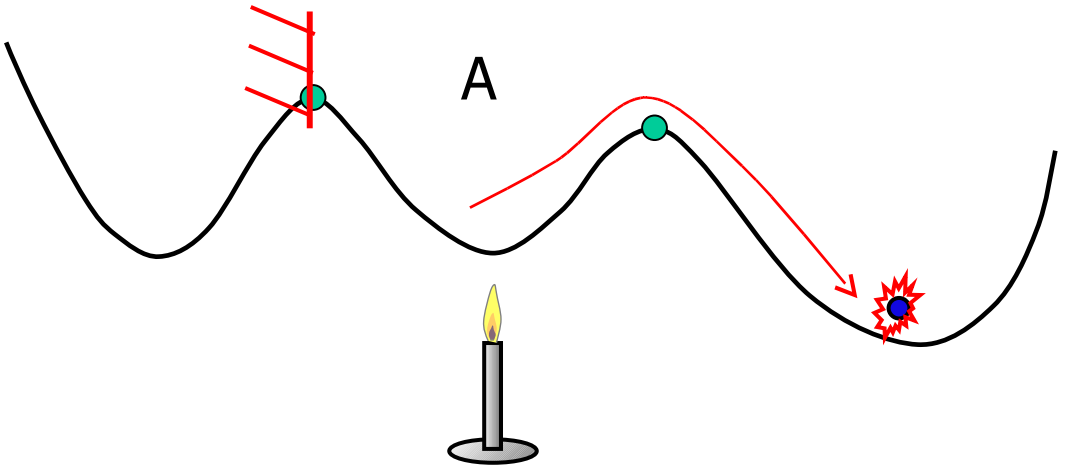

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

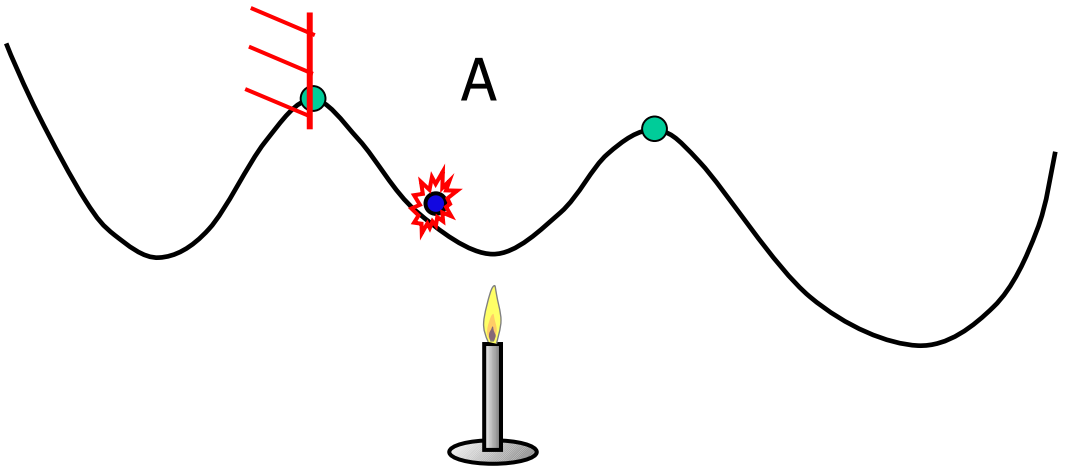

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

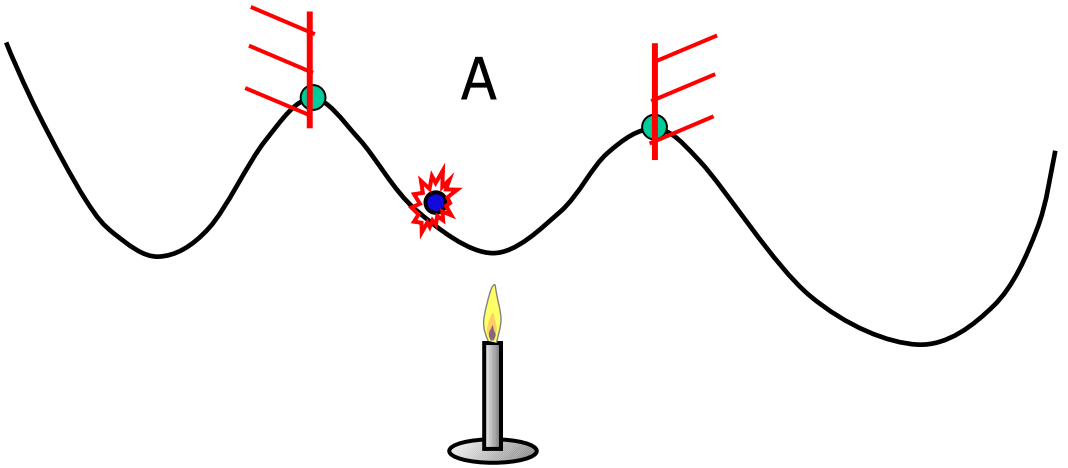

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

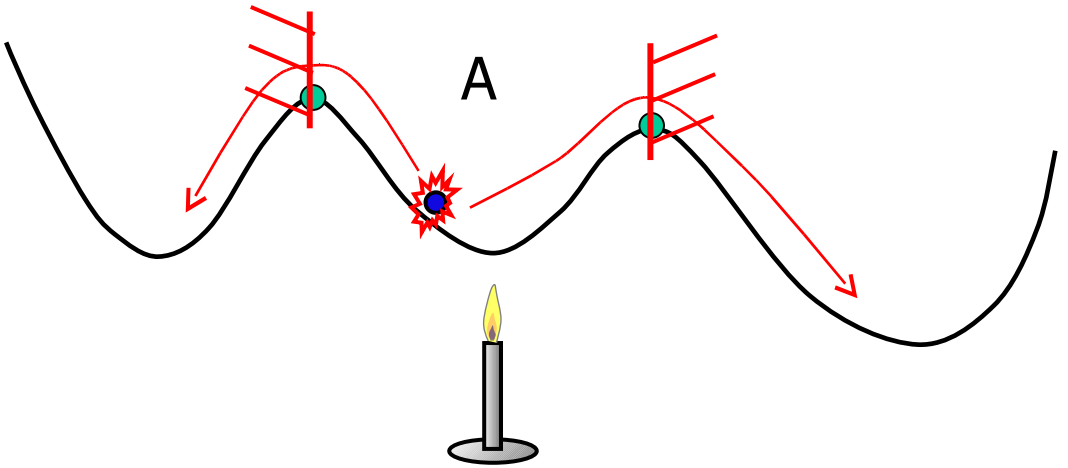

TAD temperature-extrapolated time

Because each rate is assumed to be Arrhenius,

$$k = \nu_0 \exp[-\Delta E/k_B T] ,$$

the time for each particular event at high T can be extrapolated to low T:

$$t_{low} = t_{high} \exp[\Delta E(1/k_B T_{low} - 1/k_B T_{high})] .$$

This time is sampled correctly from the exponential distribution at low T, mapped from the high T sample:

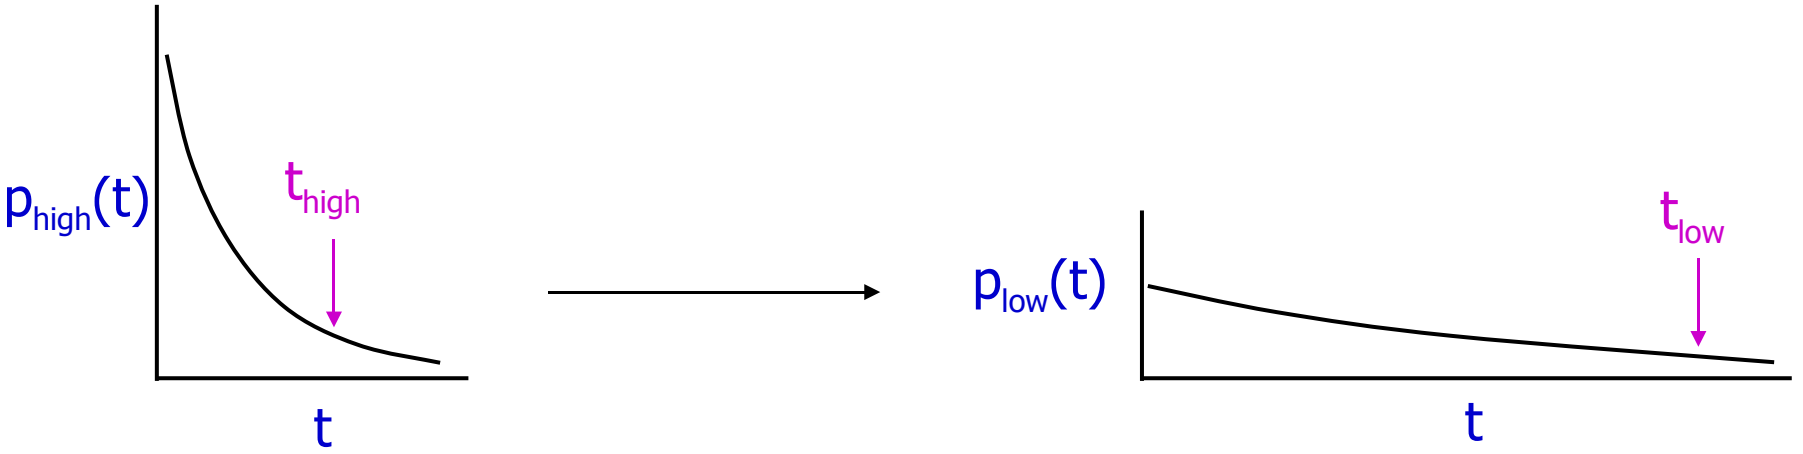

The Arrhenius view

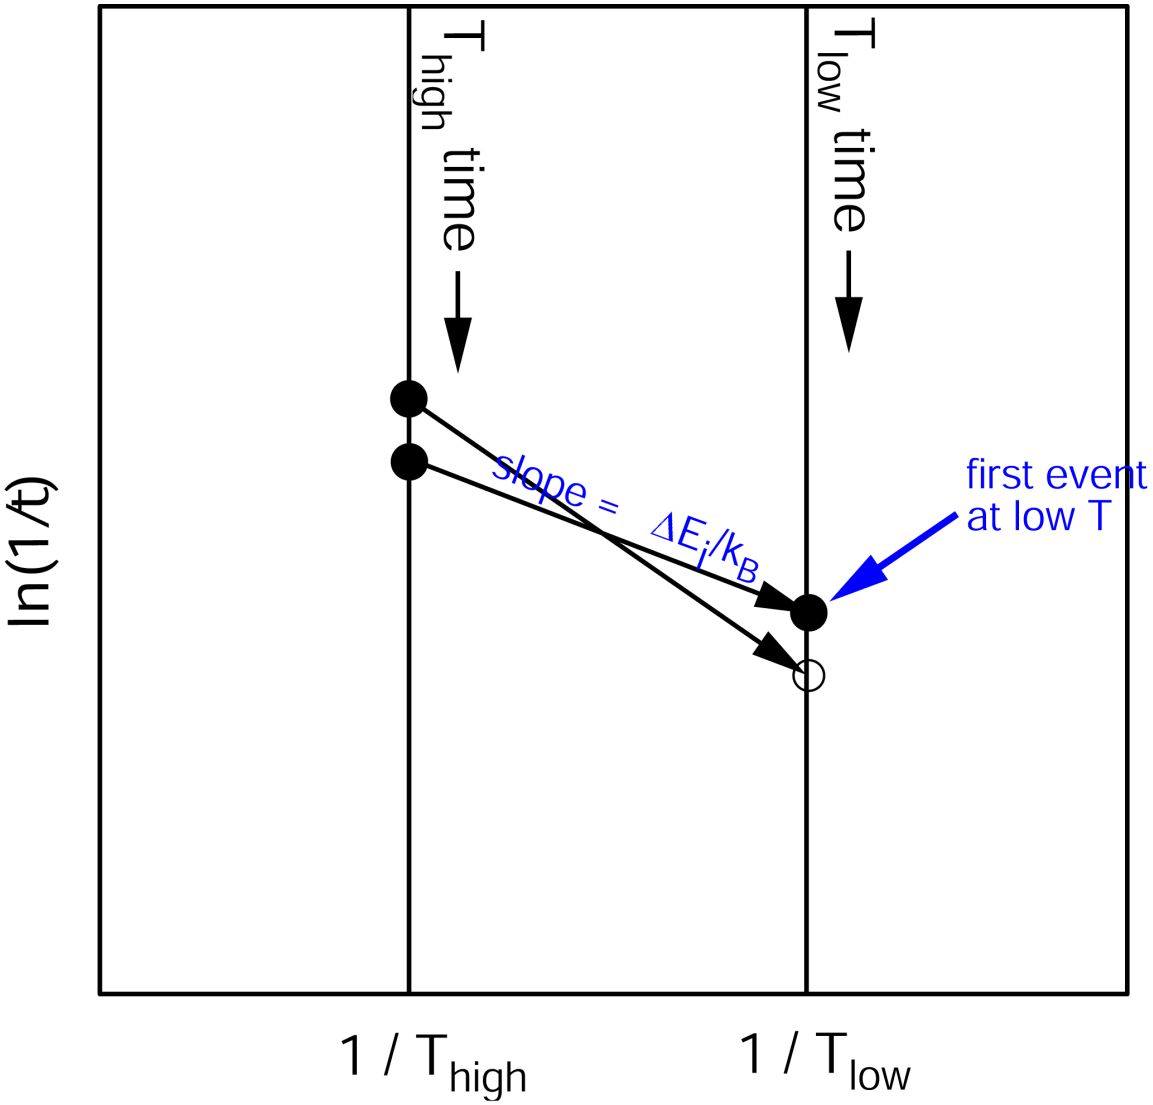

when can we stop?

The confidence line

For a pathway with rate k, the time τ required to be certain with confidence $1-\delta$ that at least one escape will occur is given by

$$\tau = (1/k)\ \ln(1/\delta)$$

For an Arrhenius rate, $k = \nu_0 \exp(-E_a/k_B T)$, all but fraction δ of the first escapes will occur above the line with slope E_a and intercept $\ln[\nu_0/\ln(1/\delta)]$

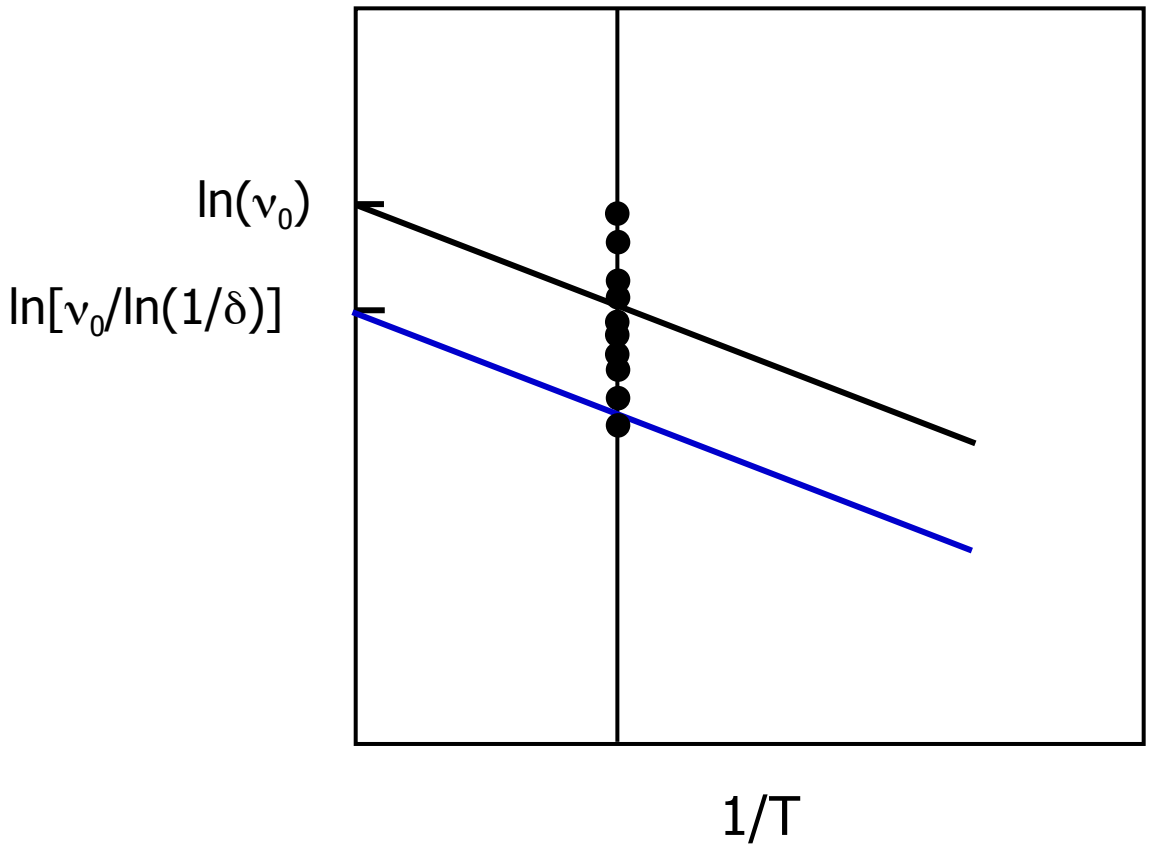

The confidence line

For a pathway with rate k, the time τ required to be certain with confidence $1-\delta$ that at least one escape will occur is given by

$$\tau = (1/k)\ \ln(1/\delta)$$

For an Arrhenius rate, $k = \nu_0 \exp(-E_a/k_B T)$, all but fraction δ of the first escapes will occur above the line with slope E_a and intercept $\ln[\nu_0/\ln(1/\delta)]$

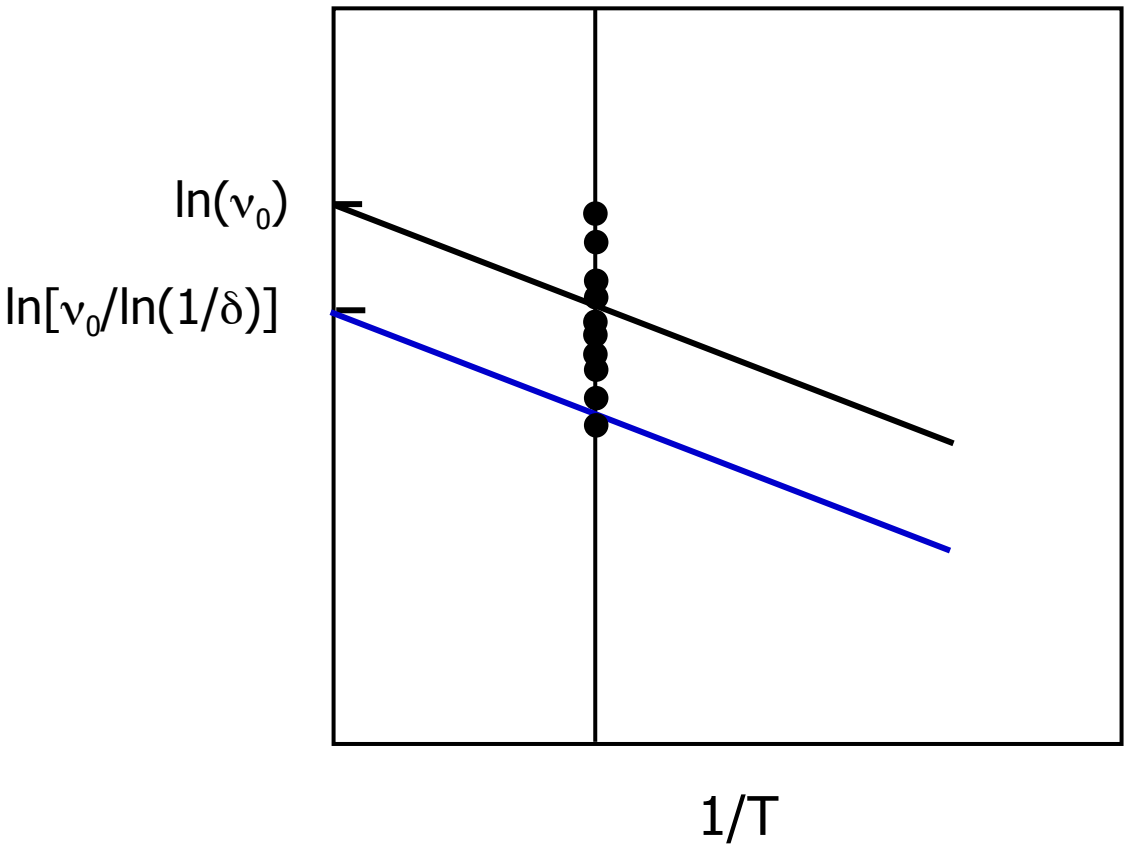

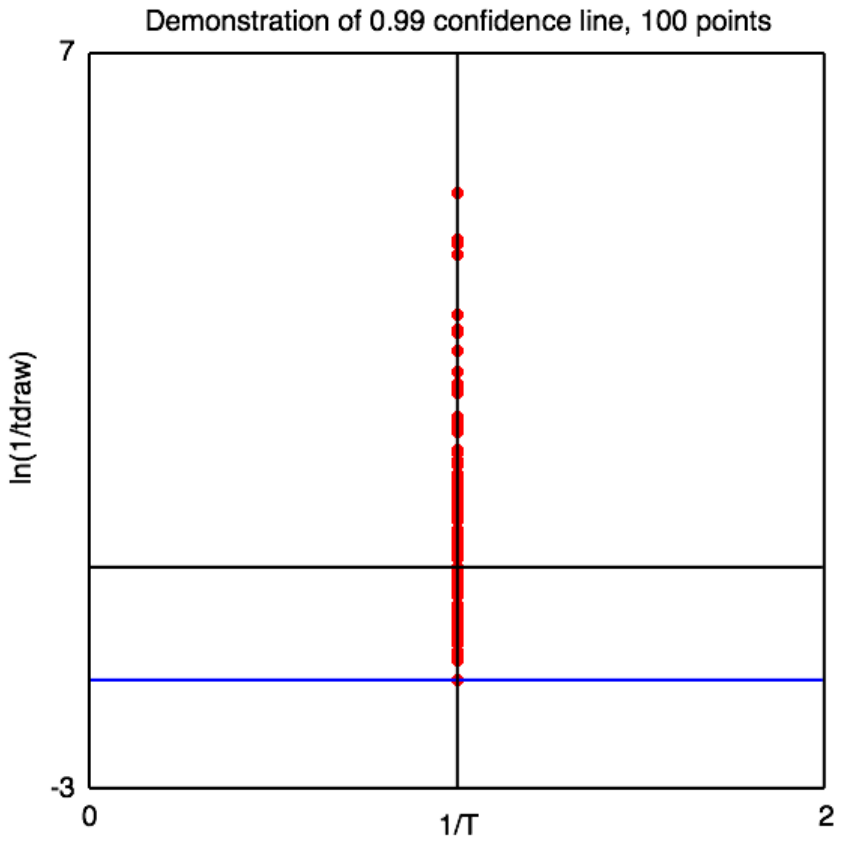

TAD - when can we stop the MD and accept an event?

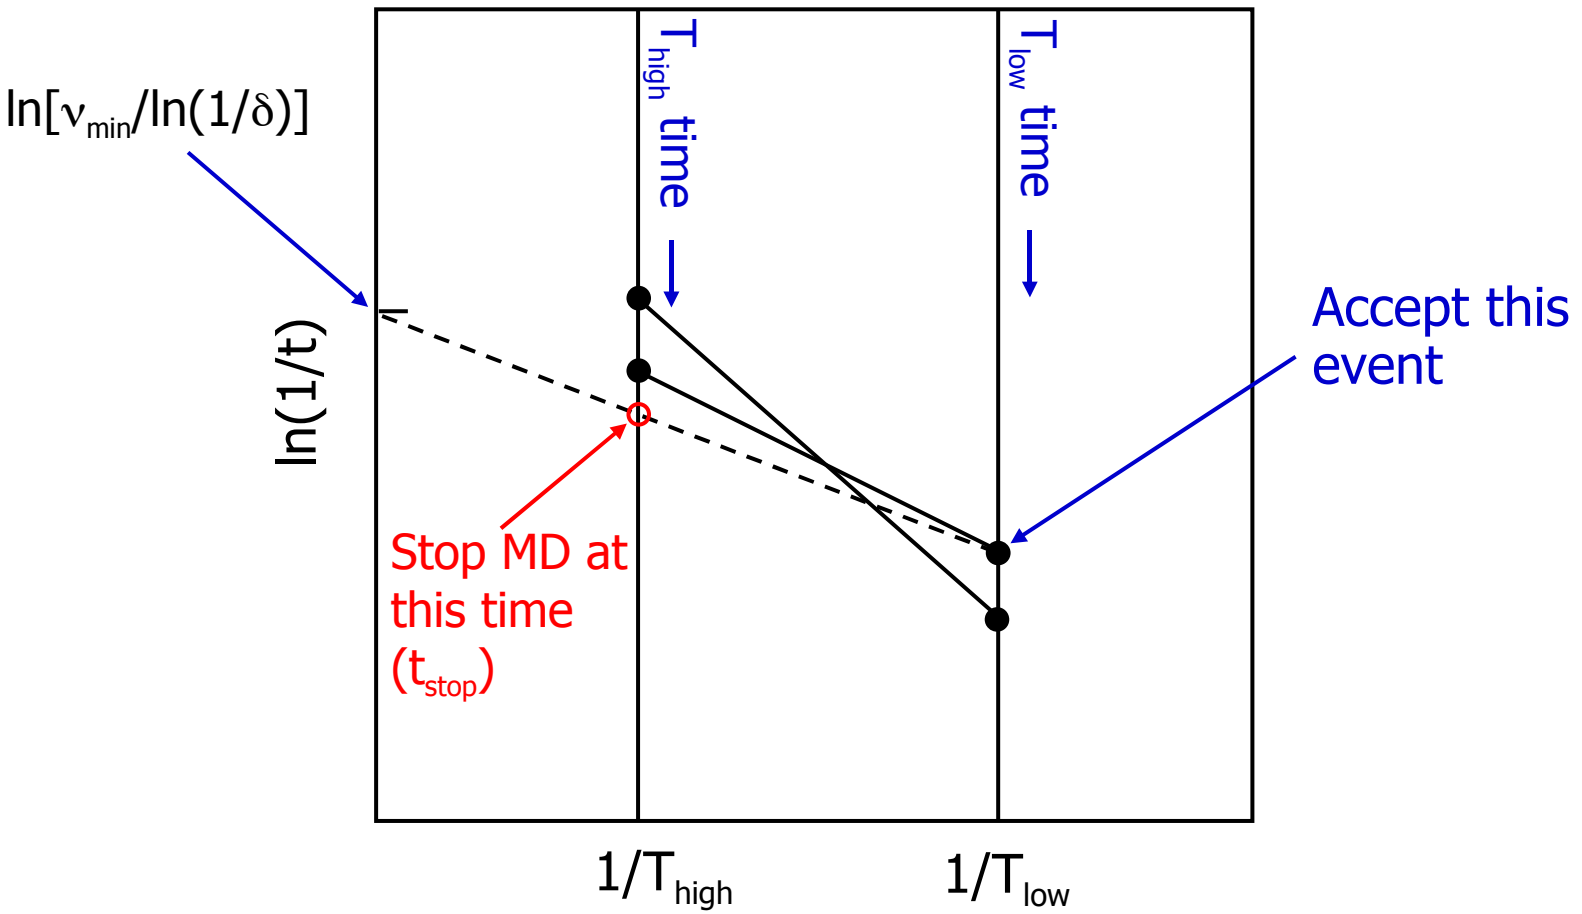

After time t_{stop}, with confidence 1-δ, no event can replace shortest-time event seen at low T.

Move system to this state and start again.

Exact dynamics, assuming harmonic TST, ν_{min}, uncertainty δ.

MD+TAD metal deposition simulation

- MD for each deposition event (2 ps)
- TAD for intervening time (~1 s)
- Embedded atom method (EAM) for fcc metals (e.g., Cu, Ag, …; LANL fit)

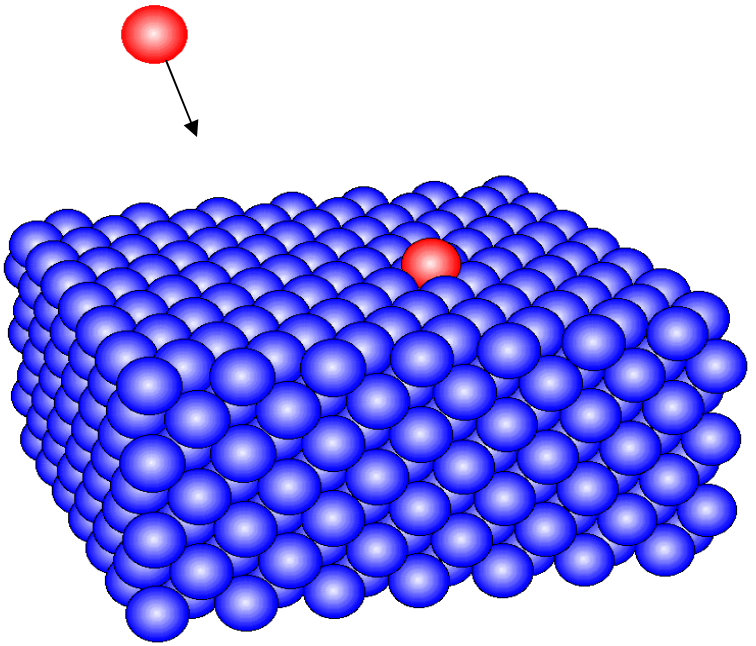

MD+TAD deposition of Cu/Ag(100)

T=77K, flux= 0.04 ML/s, matching deposition conditions Of Egelhoff and Jacob (1989).

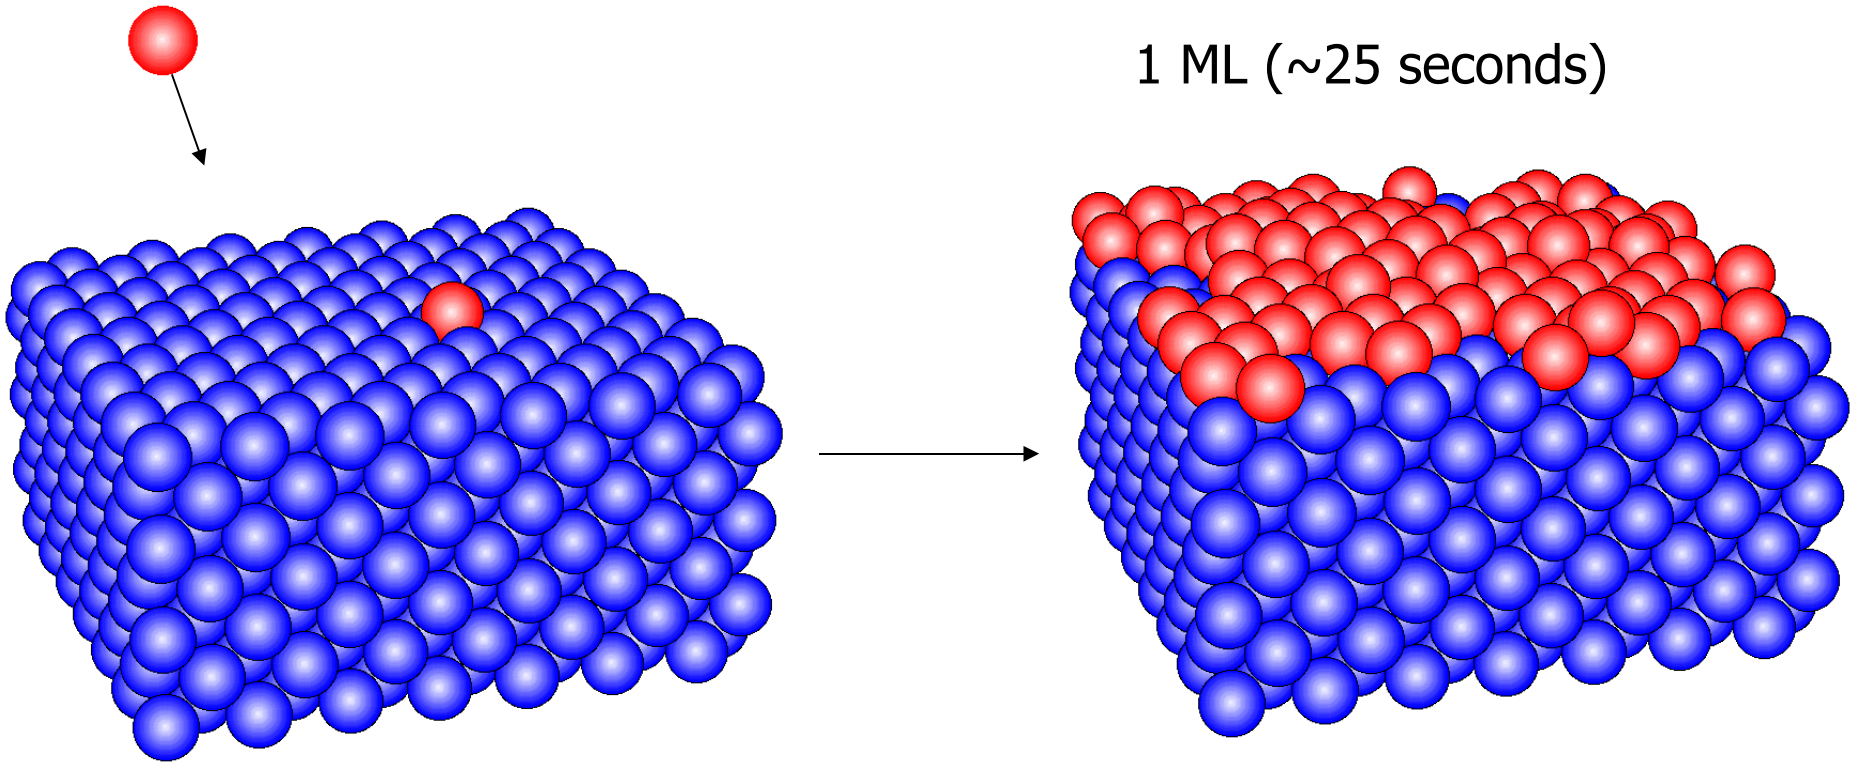

Second-layer Cu atoms exhibit mobility at T=77K, due to epitaxial strain of Cu on Ag(100).

Sprague, Montalenti, Uberuaga, Kress and Voter, Phys. Rev. B 66, 205415 (2002)

MD+TAD deposition of Cu/Ag(100)

T=77K, flux= 0.04 ML/s, matching deposition conditions of Egelhoff and Jacob (1989).

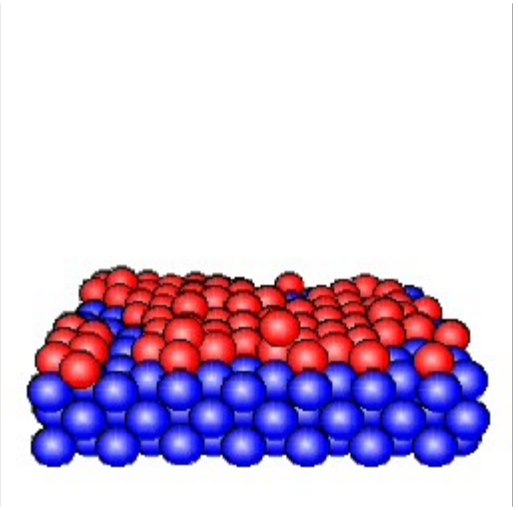

Second-layer Cu atoms exhibit mobility at T=77K, due to epitaxial strain of Cu on Ag(100).

Sprague, Montalenti, Uberuaga, Kress and Voter, Phys. Rev. B 66, 205415 (2002)

MD+TAD deposition of Cu/Cu(100)

Concerted events observed at T=77K and T=100K:

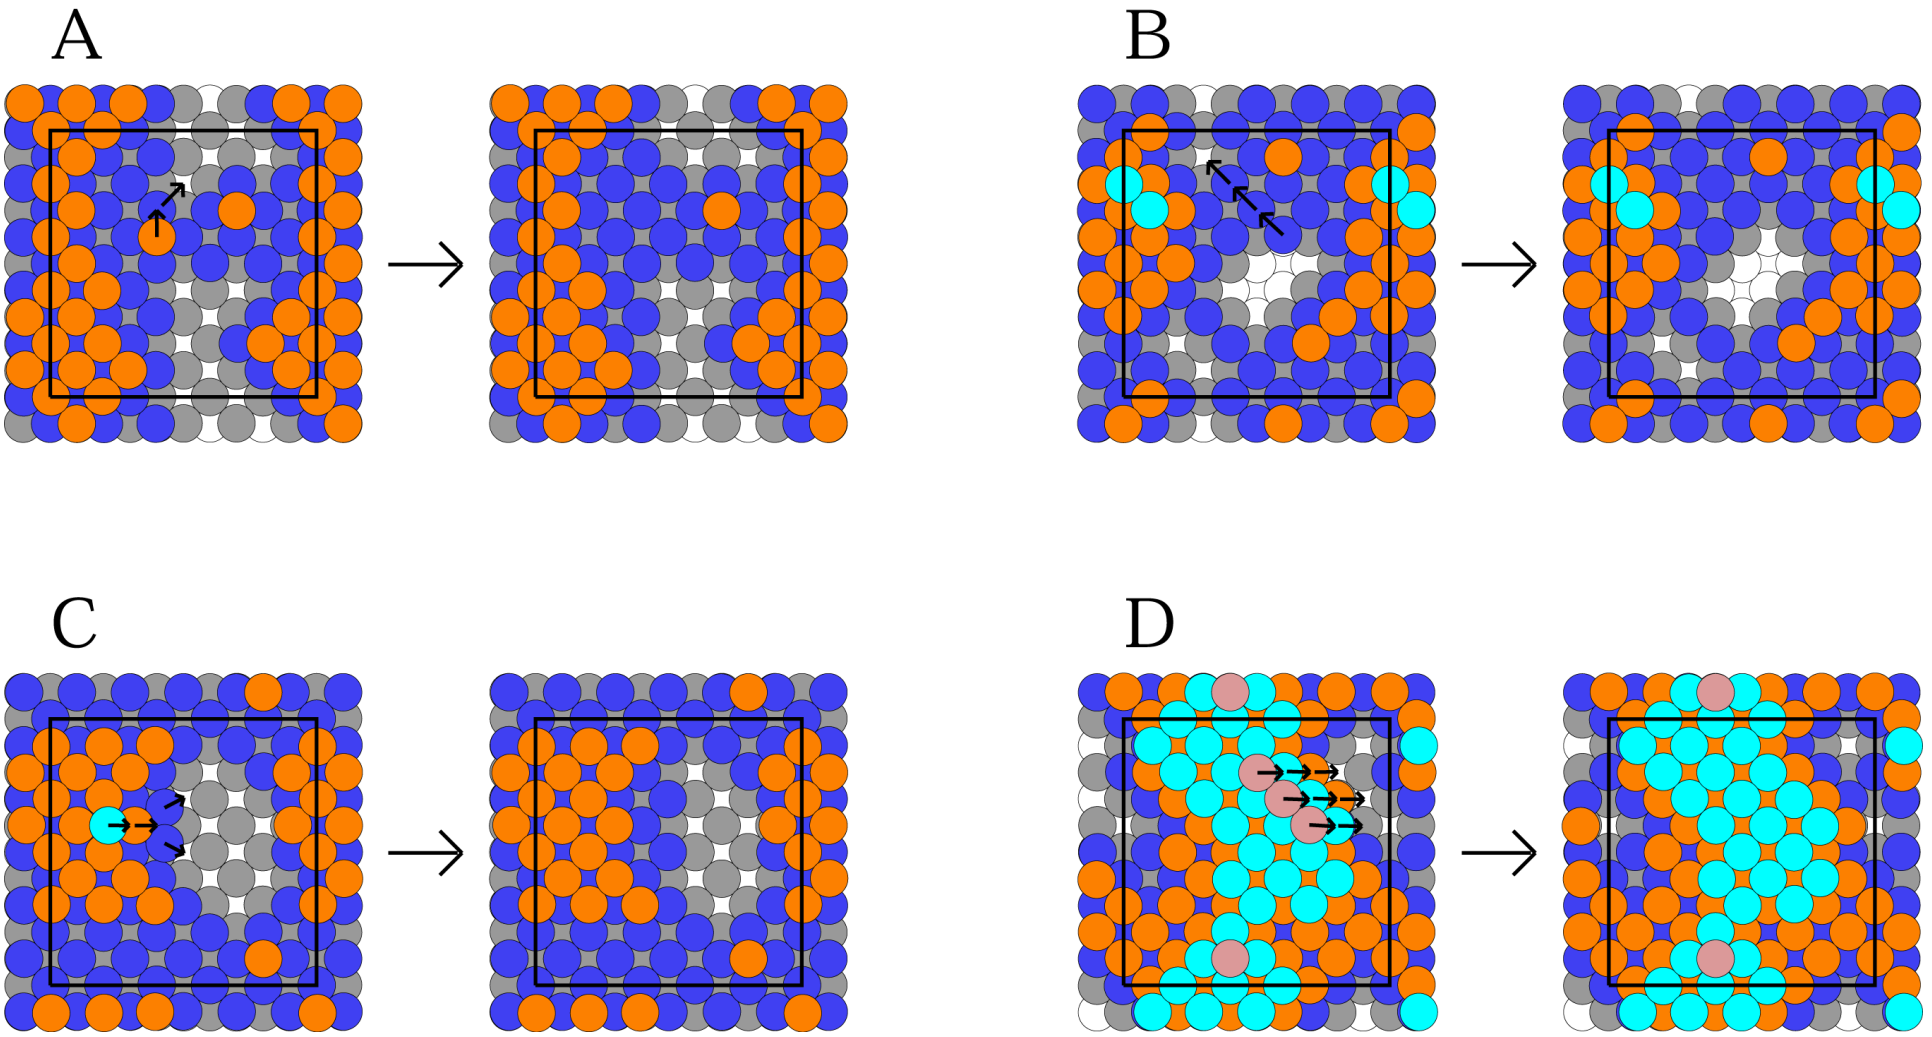

Summary - TAD

Very powerful is all barriers are relatively high relative to T.

Can reach boost factors in the thousands or millions.

Complex to implement if we want to play every trick.

Can be generalized to work in other ensembles.

Current challenges

- The low-barrier problem - boost is limited by lowest barrier - problem for many realistic systems. Detecting equilibration within meta-basins could really help us.
- Improving scaling with system size - methods as described are currently limited to small systems ($\sim 10^3$ atoms)
- Treating more complex systems (e.g., solid-liquid interface) where we don't even know what are the slow variables.
- Using ab initio or DFT force calls for higher accuracy, eliminating potentials
- Feeding information about atomistic behavior to higher-level models and combining with higher-level models

Summary

- Accelerated molecular dynamics concept:
  - Let the trajectory find an appropriate way out or state, but coax it into doing so more quickly
  - This way, we include all possible transitions, irrespective of their complexity.
- Significant speedup over standard MD when barriers are high relative to temperature (from 10x to 1,000,000x)
- Often encounter unexpected behavior

Recent review: B.P. Uberuaga, F. Montalenti, T.C. Germann, and A.F. Voter, Handbook of Materials Modeling, Part A - Methods (Springer, 2005)