Molecular Dynamics and Accelerated Molecular Dynamics

Arthur F. Voter
Theoretical Division
Los Alamos National Laboratory

Lecture 4

Tutorial Lecture Series
Institute for Pure and Applied Mathematics (IPAM)
UCLA
September 13-16, 2005

Acknowledgment: DOE/BES

Overall Outline

Day 1: Molecular dynamics methods

Day 2: Molecular dynamics methods

Day 3: Infrequent events, transition state theory, saddle finding, etc.

Day 4: Accelerated molecular dynamics methods

Overall Outline

Day 1: Molecular dynamics methods

Day 2: Molecular dynamics methods

Day 3: Infrequent events, transition state theory, saddle finding, etc.

Day 4: Accelerated molecular dynamics methods

- Quick review of key concepts
- Parallel-replica dynamics
- Hyperdynamics
- Temperature accelerated dynamics (TAD)

Infrequent Event System

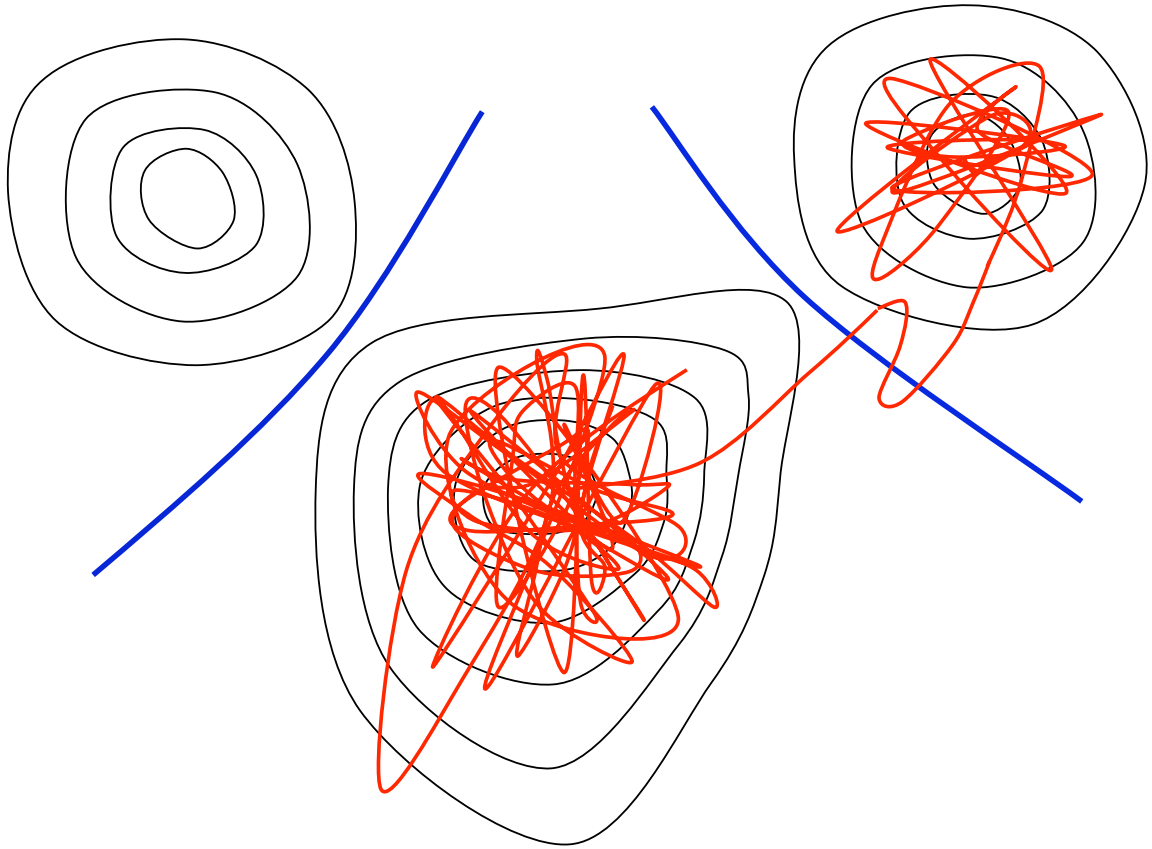

Review of some key points

For an infrequent-event system, the probability of escaping to any particular final state is proportional to the rate constant for that escape path.

The bottleneck for escape may be energetic (a barrier) or entropic (a free-energy barrier), or some of both.

In transition state theory (TST), we approximate this rate constant as the flux through the dividing surface that separates the two states.

The TST rate is an equilibrium property of the system. We could run a very long MD trajectory and count the crossings per time, or (for the canonical ensemble) we can simply determine the Boltzmann probability of being in the dividing surface subspace and multiply it by the average velocity normal to that dividing surface $[(2k_BT/\pi m)^{1/2}]$.

Successive crossings may be correlated; so k^{exact} is less than or equal to k^{TST}.

Review of some key points (cont.)

We can correct for these recrossings, to compute k^{exact}, by initiating "half trajectories" at the dividing surface (the same one we used to compute the TST rate), and tallying up where they come to rest. This assumes there is a separation of time scales -- that the correlation time (τ_{corr}) is much shorter than the average reaction time ($\tau_{rxn}=1/k$).

We can actually pick almost any dividing surface we want, and the TST + dynamical corrections formalism will still give the same rate constant.

If there is not a separation of time scales, even the rate itself is ill defined, since there is no longer a first-order (exponential) population decay.

For a true infrequent-event system, after a correlation time has passed since entering the state, the probability distribution function for the time to the next escape event is exponentially distributed (i.e., a poisson process).

Exponential escape-time distribution

Infrequent events give an exponential distribution of first-escape times (after correlation time has passed since entering state).

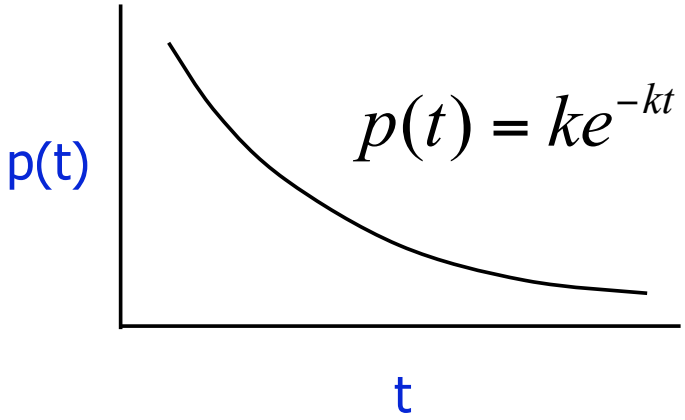

$\tau_{rxn} = 1/k$

We will exploit this in some of the accelerated MD methods.

Accelerated molecular dynamics concept

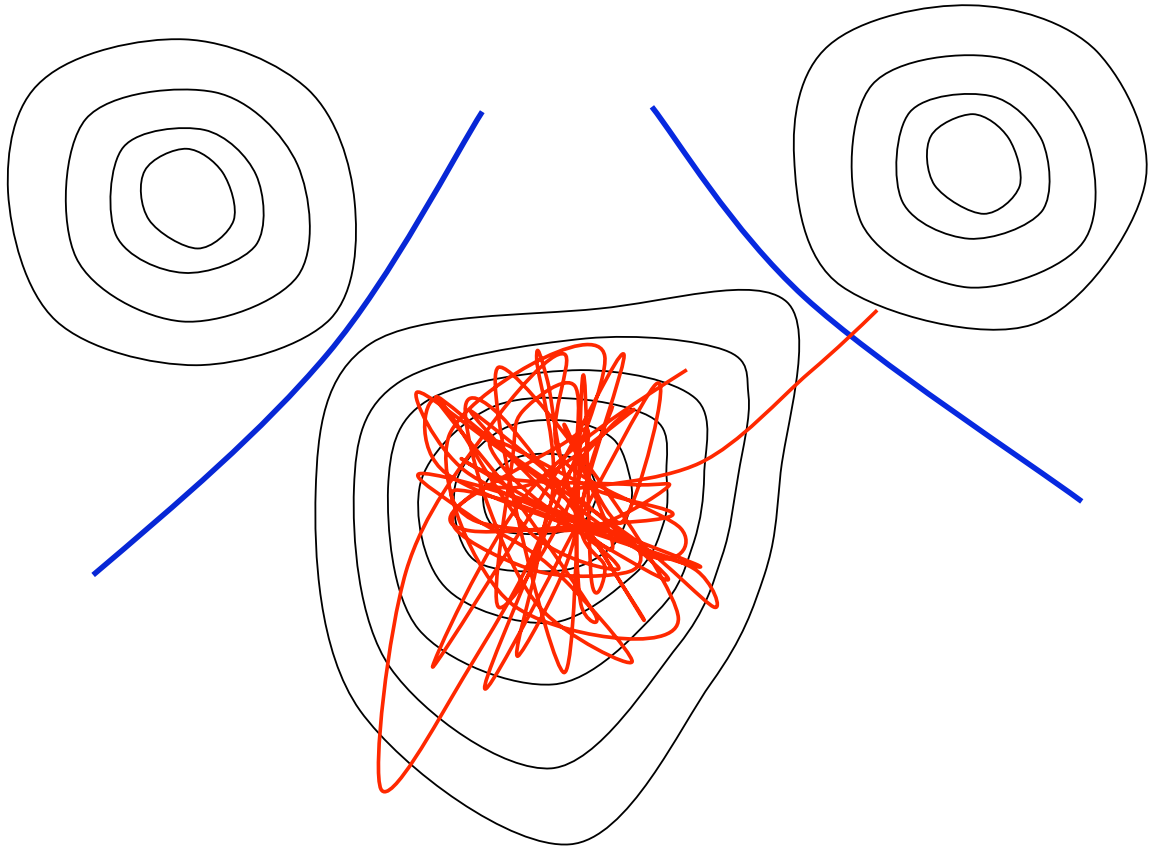

The trajectory finds an appropriate way out (i.e., proportional to the rate constant) without knowing about any of the escape paths except the first one it sees. In accelerated dynamics we try to maintain this key characteristic, while making the first escape happen sooner.

Parallel Replica Dynamics

Parallel Replica Dynamics

Parallelizes time evolution

Assumptions:

- infrequent events
- transitions can be detected
- exponential distribution of first-escape times

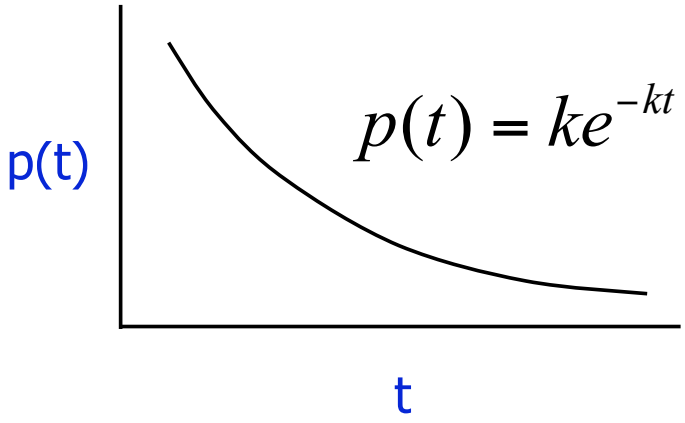

- correlation time known

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

Parallel Replica Dynamics - derivation

Consider an infrequent-event system with a set of possible escape events

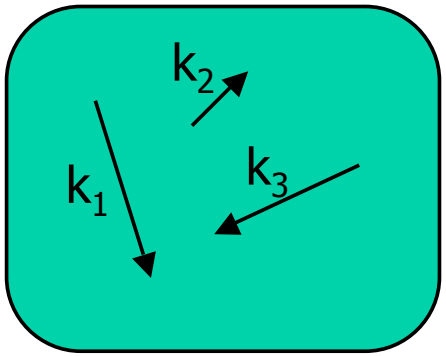

Total escape rate:

$$k_{tot} = \sum_i k_i$$

Probability distribution function for time to first event:

$$p(t) = k_{tot} \exp(-k_{tot} t)$$

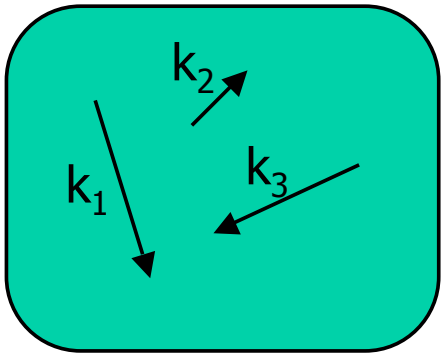

k_{tot}

$p(t) = k_{tot} \exp(-k_{tot} t)$

Now consider a set of M equivalent systems, each in the same state, evolving independently on M identical processors.

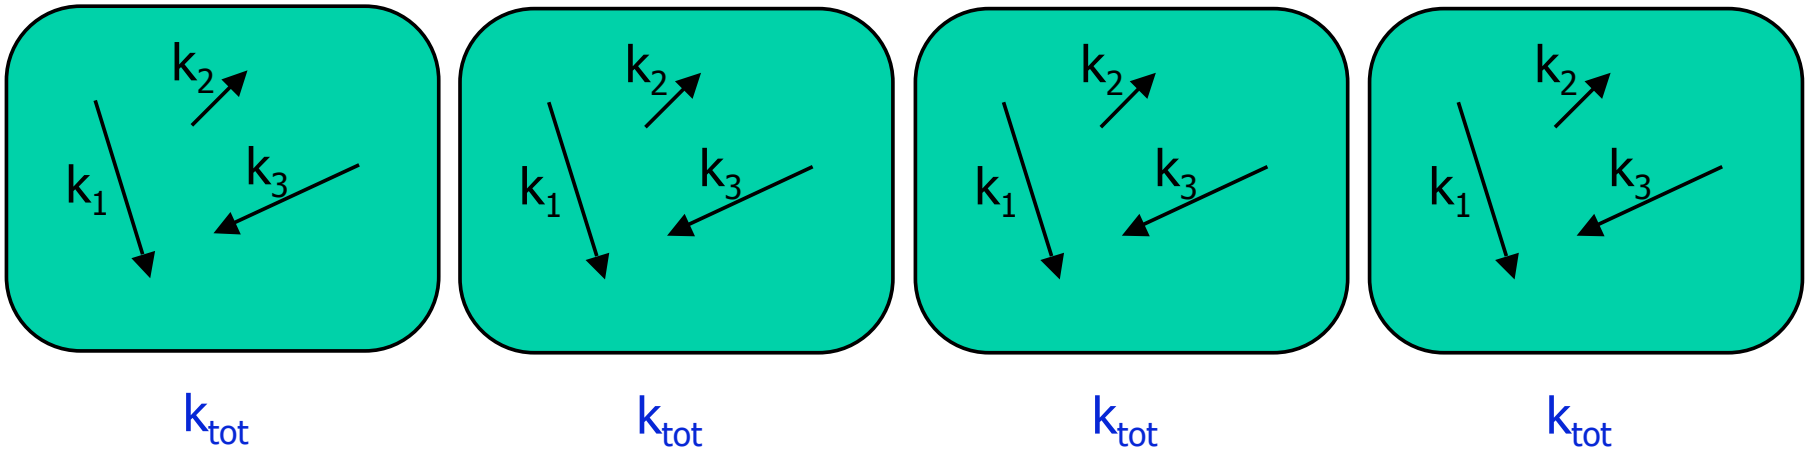

Now consider a set of M equivalent systems, each in the same state, evolving independently on M identical processors.

This behaves like a physical supersystem that is M times larger, with the simulation time on processor 1 representing the time (t) for the whole system.

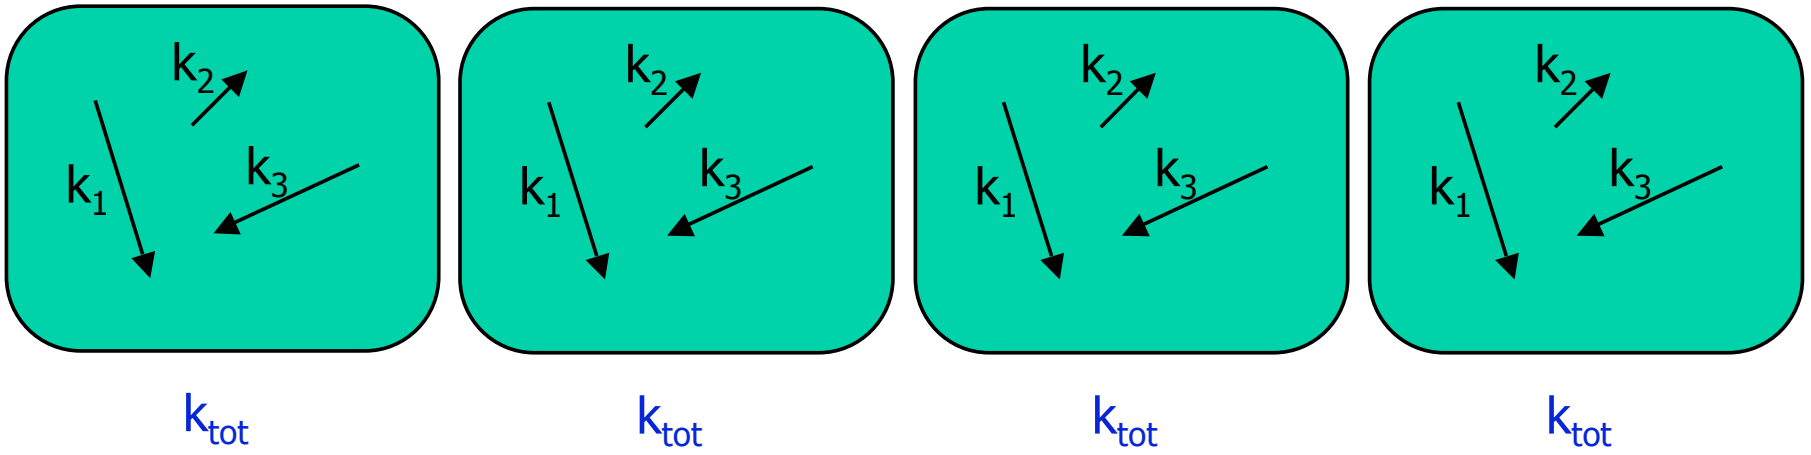

Now consider a set of M equivalent systems, each in the same state, evolving independently on M identical processors.

This behaves like a physical supersystem that is M times larger, with the simulation time on processor 1 representing the time (t) for the whole system.

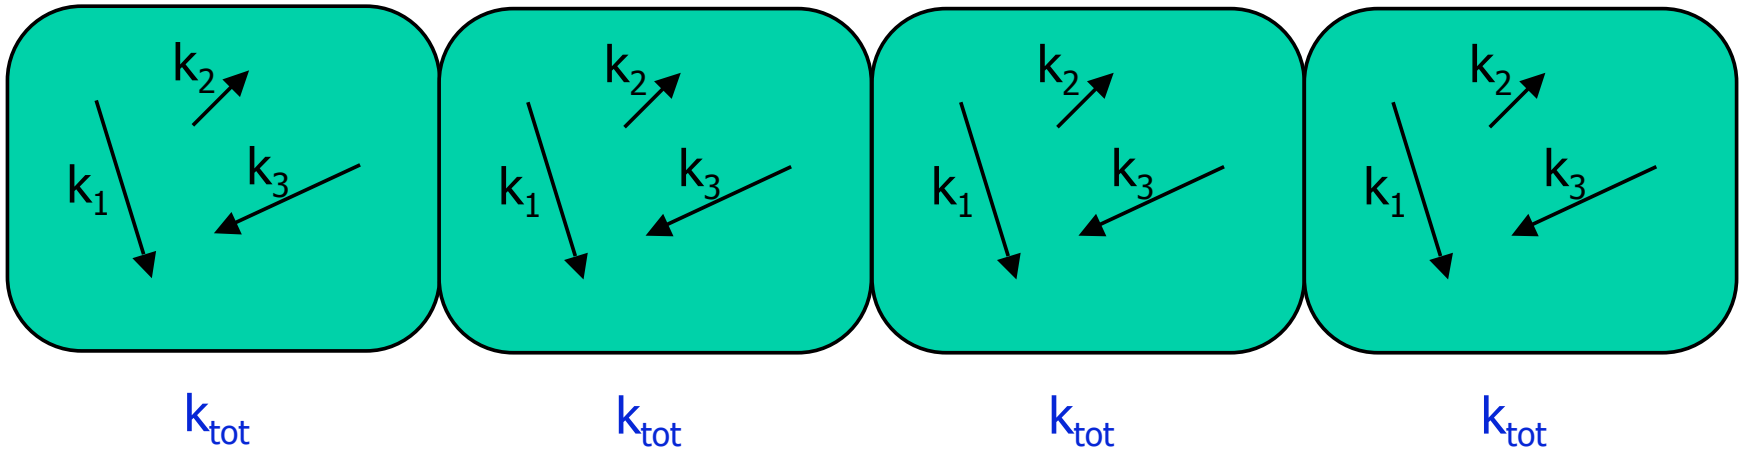

Now consider a set of M equivalent systems, each in the same state, evolving independently on M identical processors.

This behaves like a physical supersystem that is M times larger, with the simulation time on processor 1 representing the time (t) for the whole system.

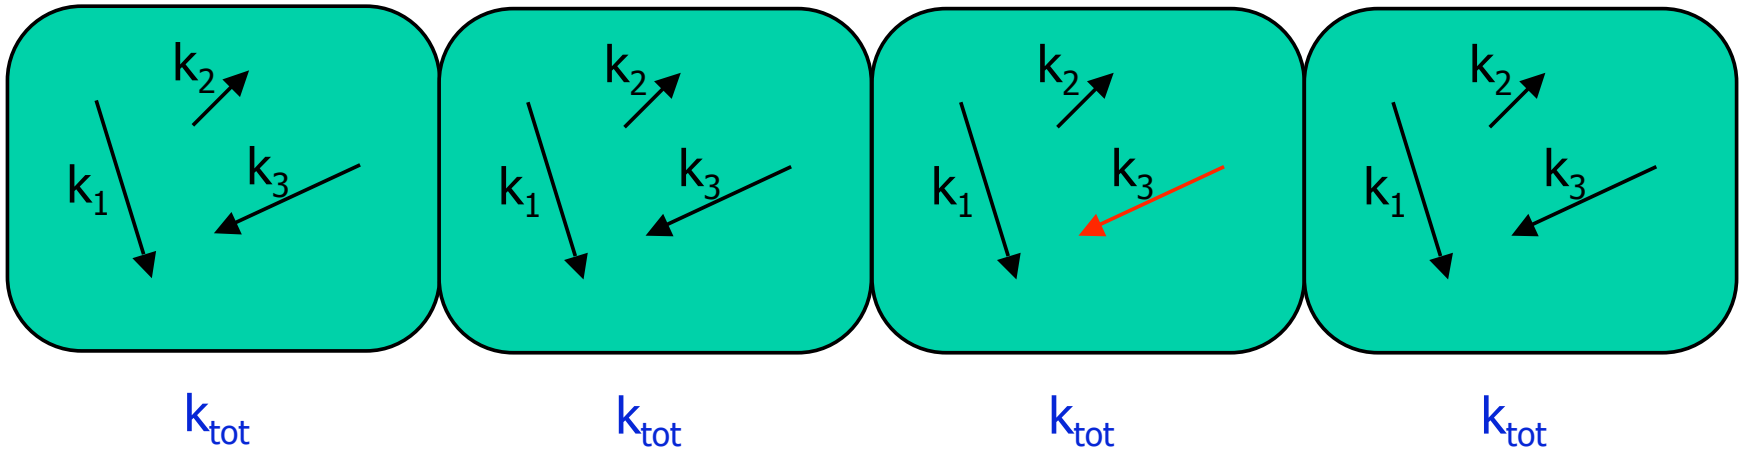

At some point, a transition will occur on one of the processors. What can we say about it?

Now consider a set of M equivalent systems, each in the same state, evolving independently on M identical processors.

This behaves like a physical supersystem that is M times larger, with the simulation time on processor 1 representing the time (t) for the whole system.

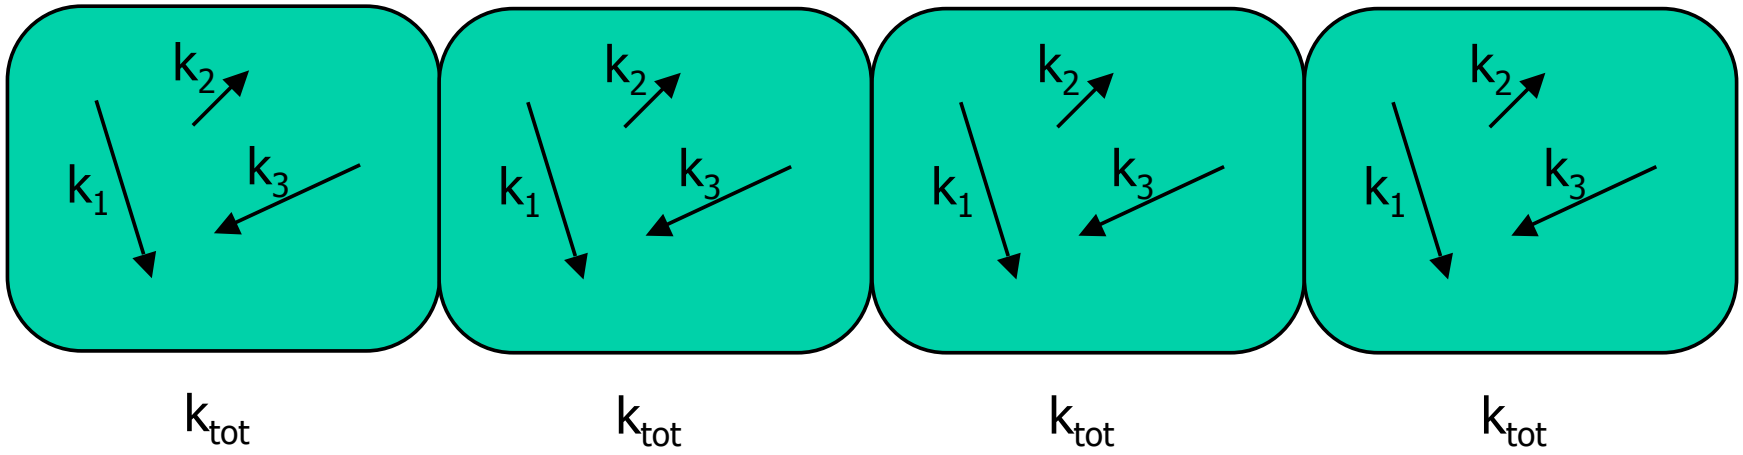

$$k_{super} = Mk_{tot}$$

Now consider a set of M equivalent systems, each in the same state, evolving independently on M identical processors.

This behaves like a physical supersystem that is M times larger, with the simulation time on processor 1 representing the time (t) for the whole system.

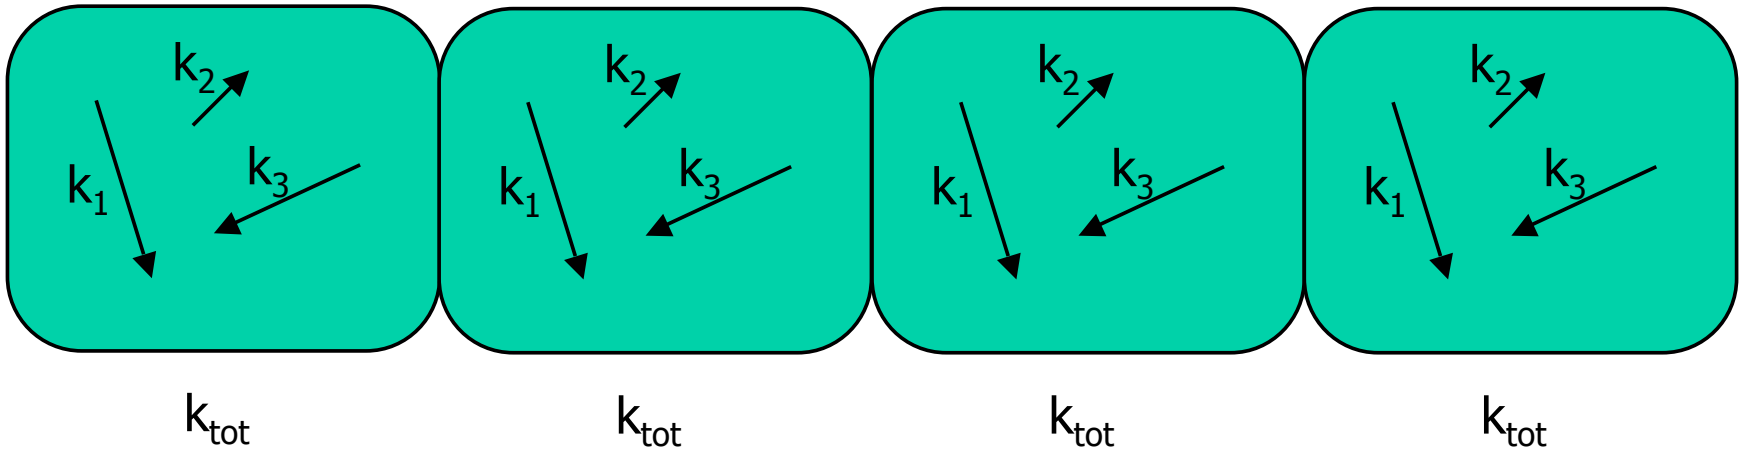

$$k_{super} = Mk_{tot}$$

$$p(t) = k_{super} \exp(-k_{super} t)$$

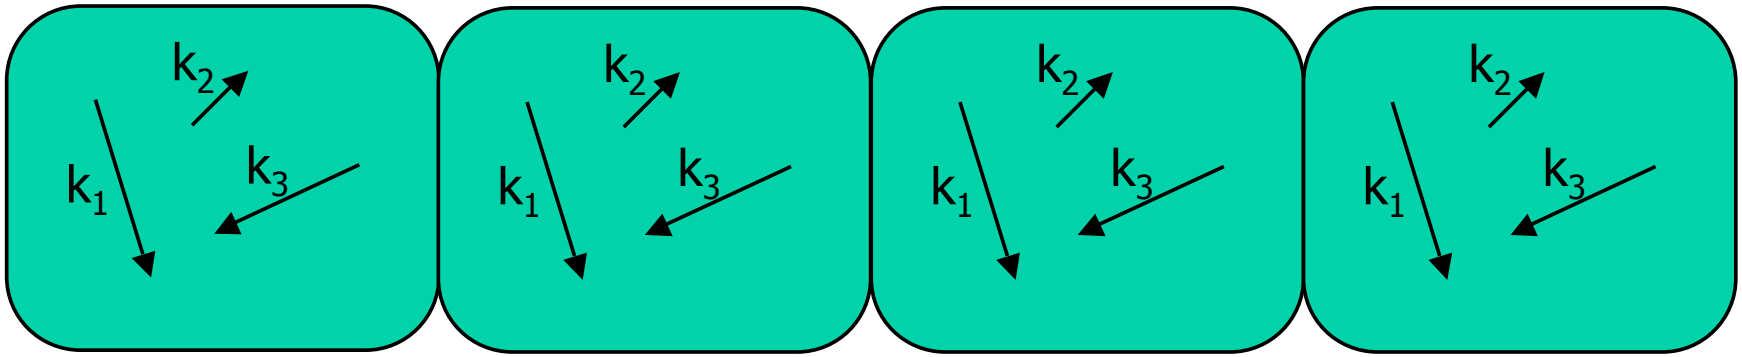

Define t_{sum} as the sum of times on all processors

$t_{sum} = Mt$

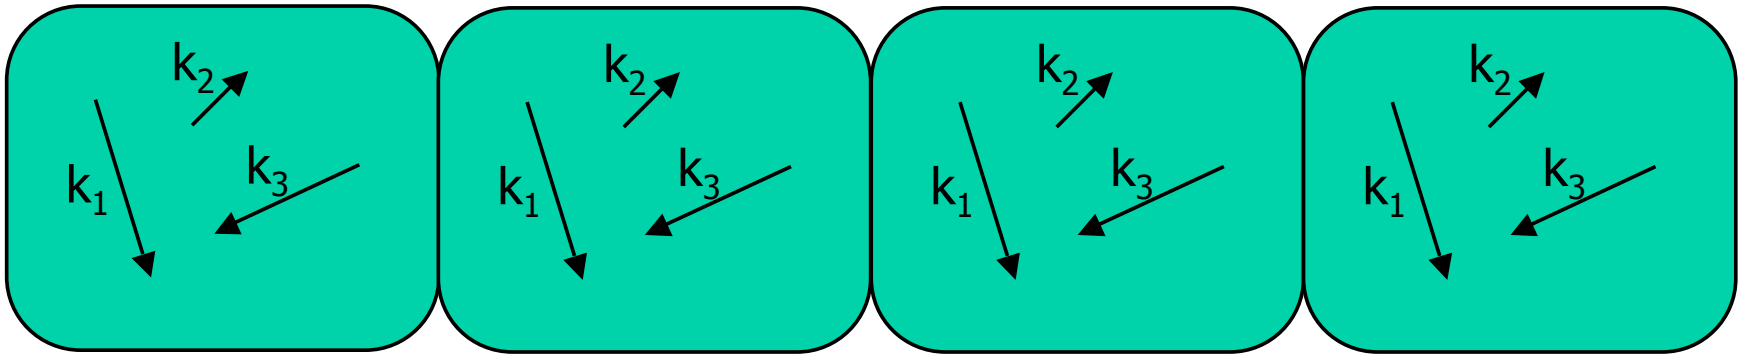

Define t_{sum} as the sum of times on all processors

$$t_{sum} = Mt$$

Then

$$p(t) = k_{super} \exp(-k_{super} t)$$

$$p(t_{sum}/M) = Mk_{tot} \exp(-Mk_{tot} t)$$

$$(1/M)p(t_{sum}/M) = k_{tot} \exp(-k_{tot} Mt)$$

$$p(t_{sum}) = k_{tot} \exp(-k_{tot} t_{sum})$$

Recall

$$k_{super} = Mk_{tot}$$

And note

$$(1/a)p(t'/a)dt' = p(t)dt$$

Notice

$$p(t_{sum}) = k_{tot} \exp(-k_{tot} t_{sum}) \quad <==> \quad p(t) = k_{tot} \exp(-k_{tot} t)$$

Expressed in this summed time, the probability distribution function for the first escape time from this state is exactly that for the original system.

We have done nothing to corrupt the relative escape rates for the different pathways.

The time (t_{sum}) for the supersystem accumulates M times faster.

We have parallelized the time leading up to the first escape event.

Comment

Processors need not run at the same speed,
nor at constant speed

Parallel Replica Dynamics Procedure

Replicate entire system on each of M processors.

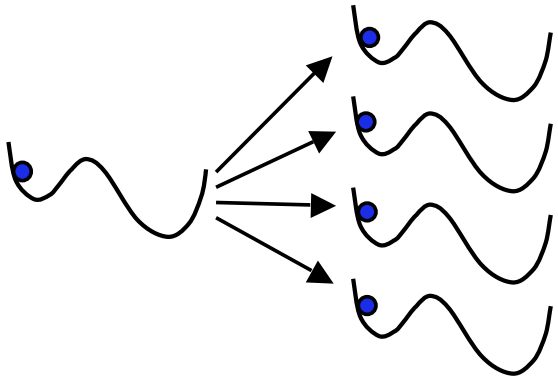

Parallel Replica Dynamics Procedure

Randomize momenta independently on each processor.

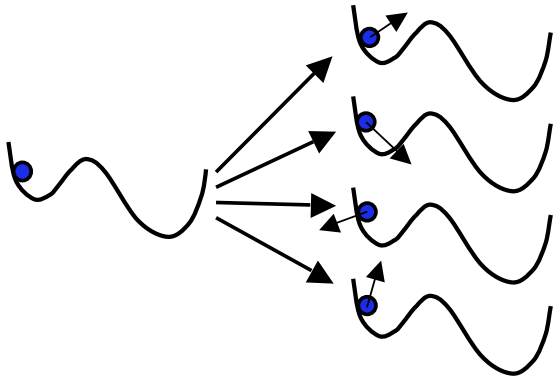

Parallel Replica Dynamics Procedure

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

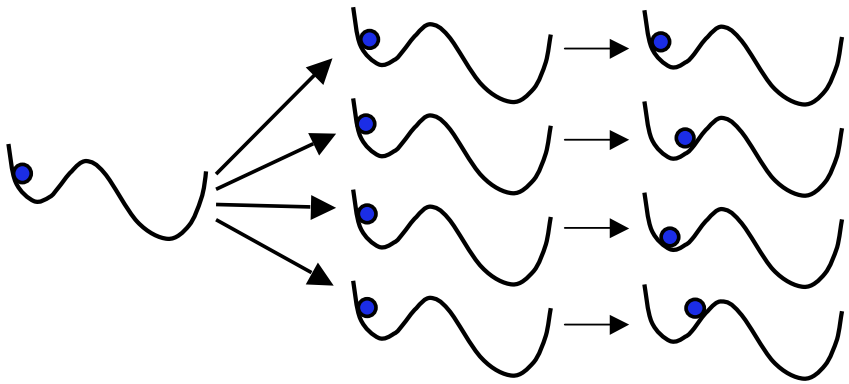

Parallel Replica Dynamics Procedure

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

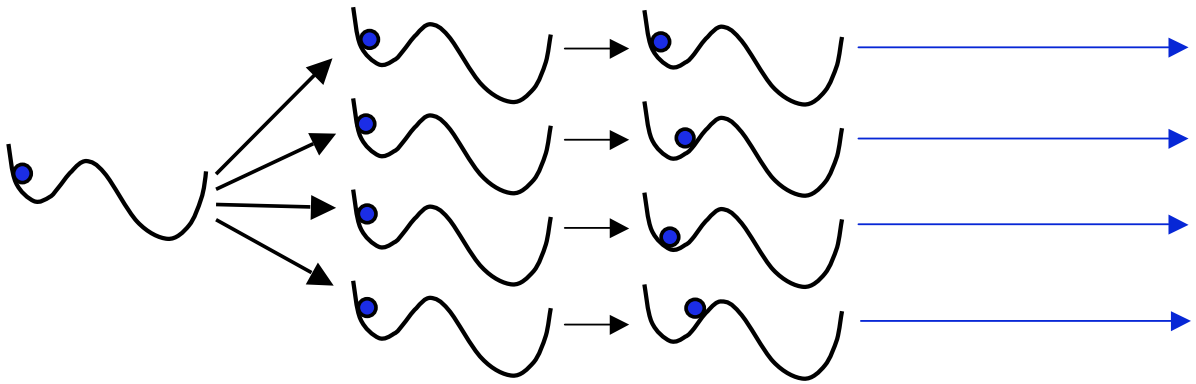

Parallel Replica Dynamics Procedure

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

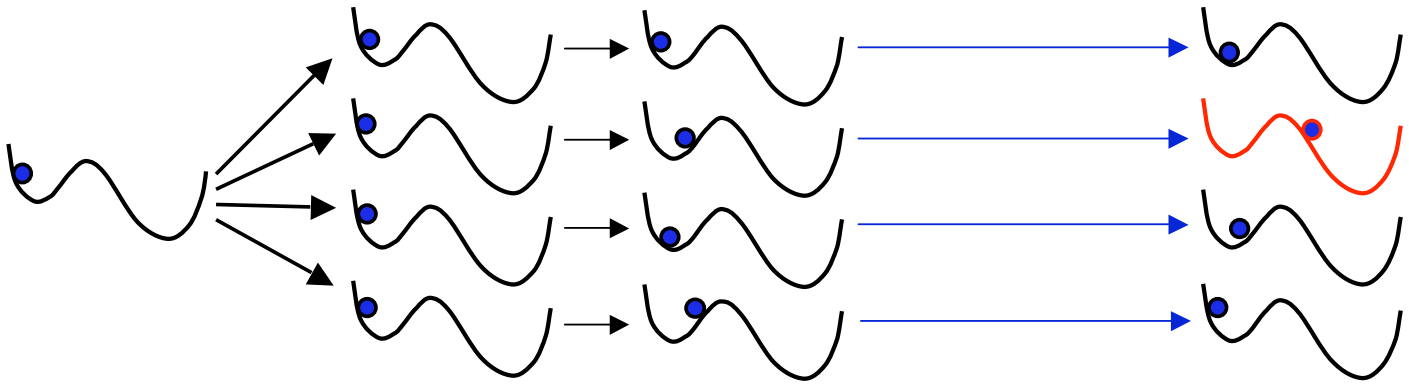

Parallel Replica Dynamics Procedure

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

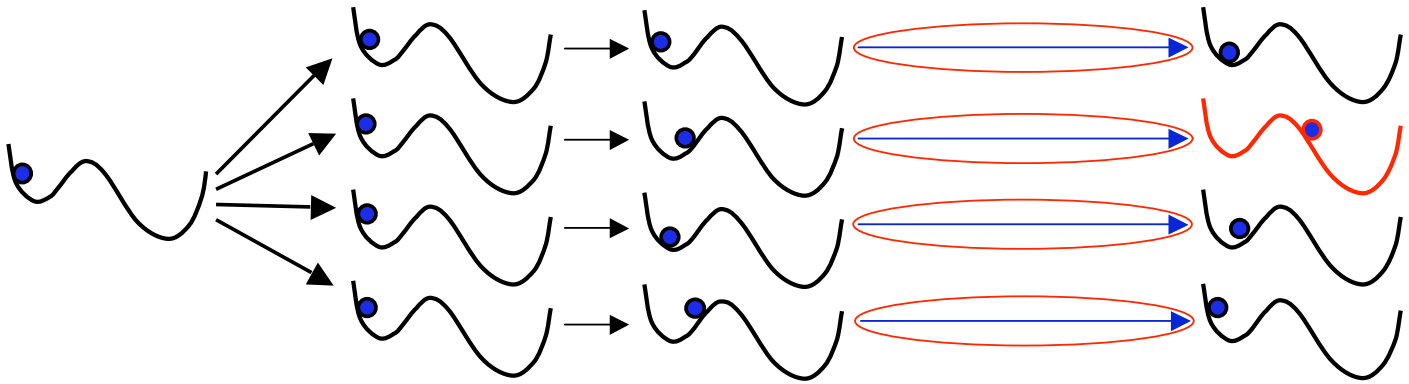

Parallel Replica Dynamics Procedure

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

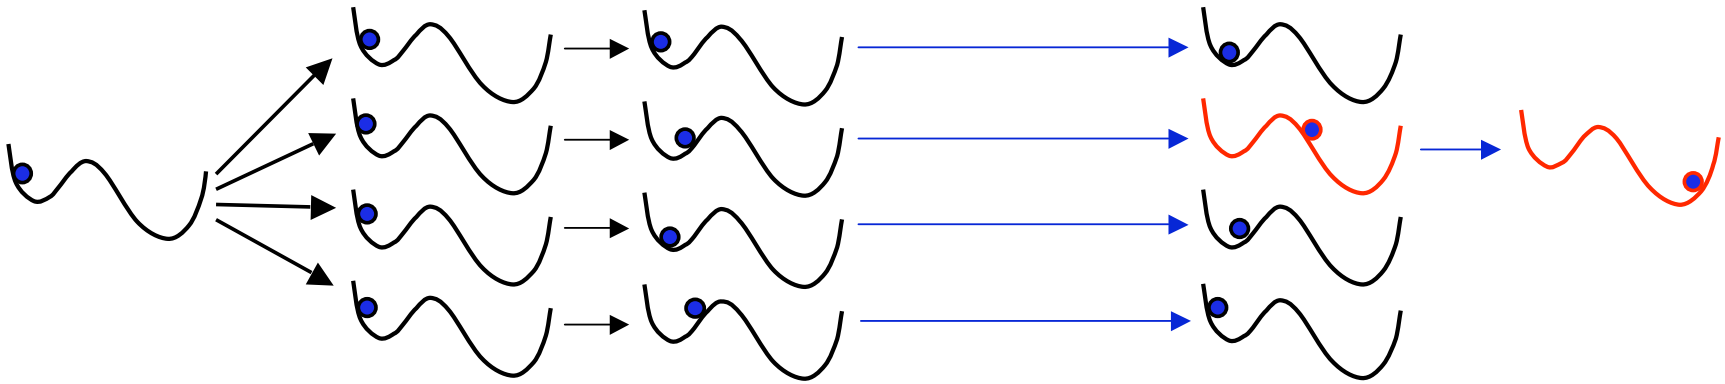

Parallel Replica Dynamics Procedure

Advance simulation clock by τ_{corr}.

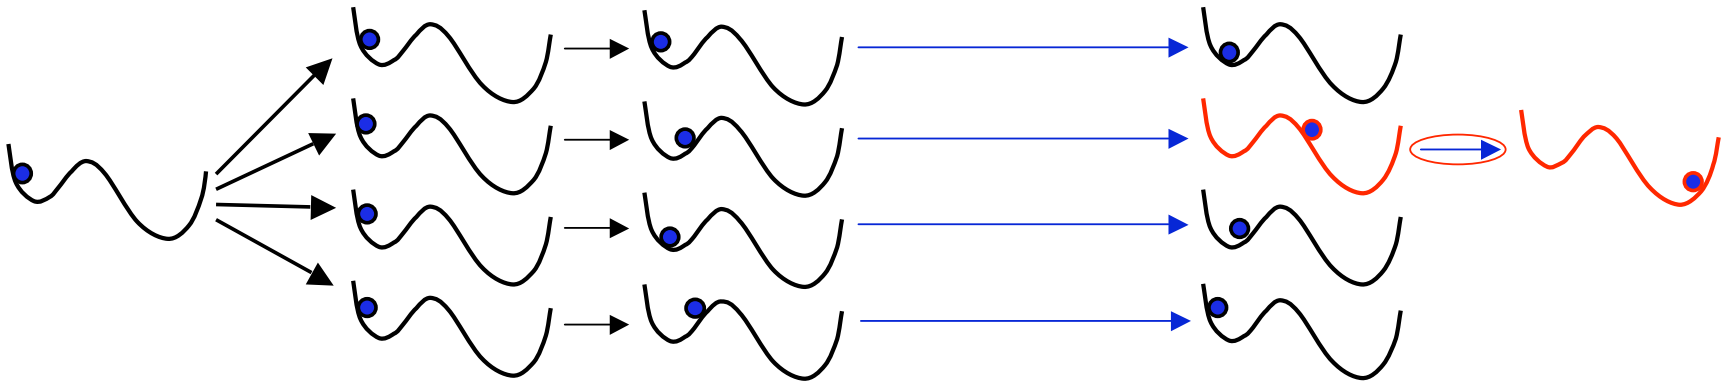

Parallel Replica Dynamics Procedure

Replicate the new state and begin procedure again.

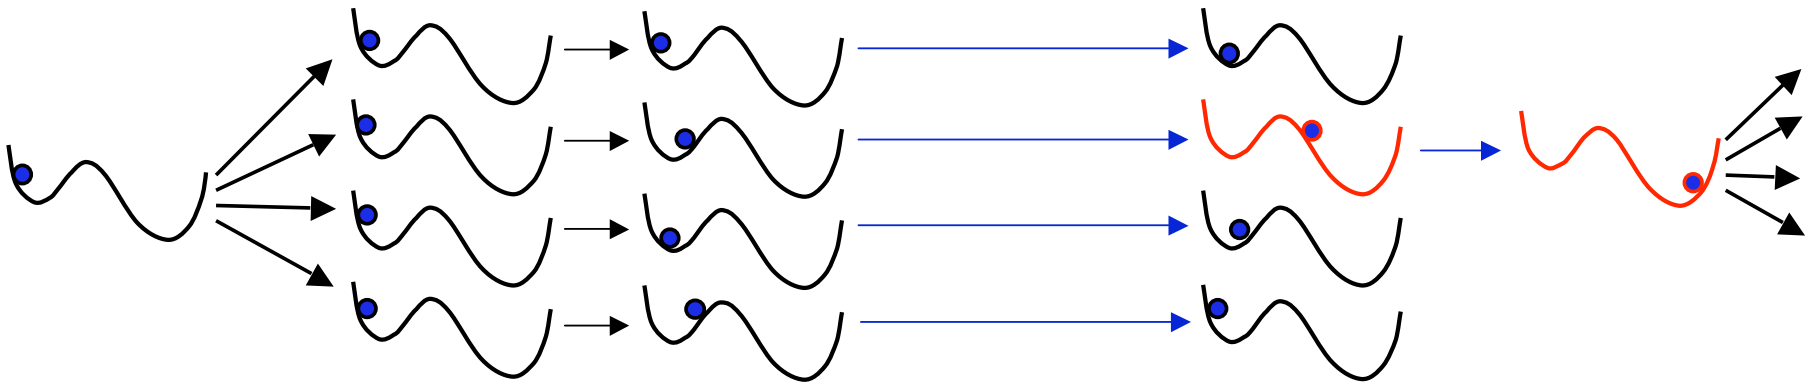

Parallel Replica Dynamics

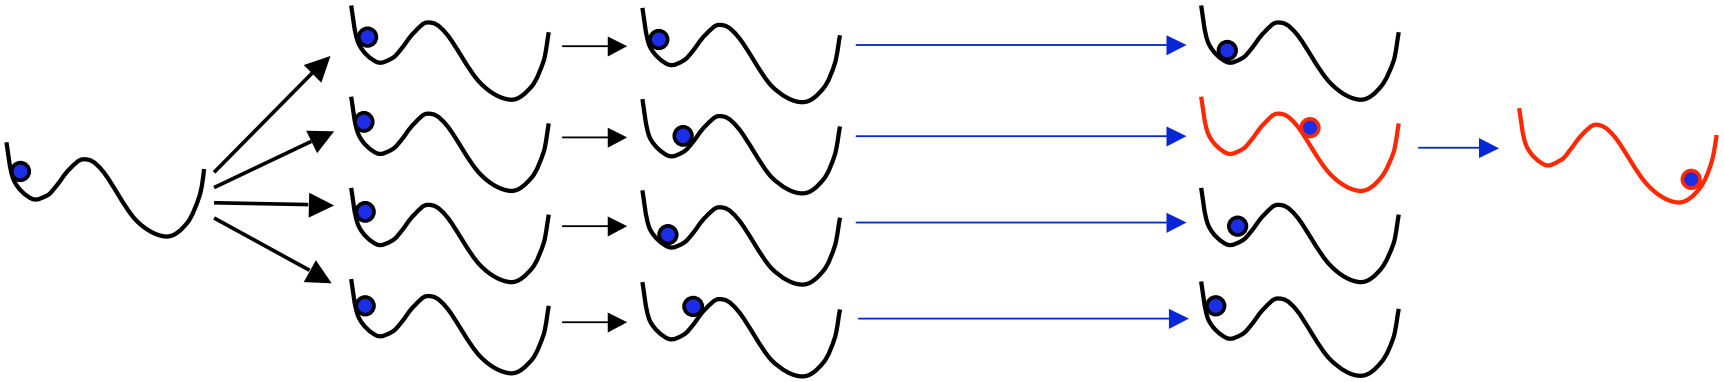

The summed time (t_{sum}) obeys the correct exponential distribution, and the system escapes to an appropriate state.

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

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

Applicable to any system with exponential first-event statistics

Detecting a transition

- best method depends on the system

- simple method for EAM metal systems:
  periodically perform steepest-descent quench;
  see if geometry at basin minimum has changed

Detecting a transition

- best method depends on the system

- simple method for EAM metal systems:
  periodically perform steepest-descent quench;
  see if geometry at basin minimum has changed

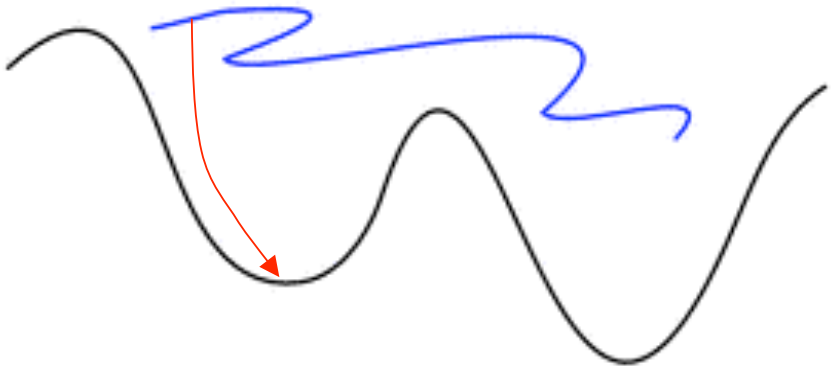

Detecting a transition

- best method depends on the system

- simple method for EAM metal systems:
  periodically perform steepest-descent quench;
  see if geometry at basin minimum has changed

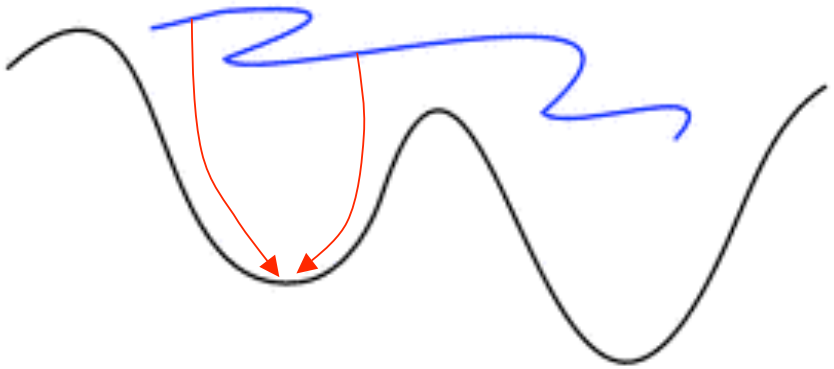

Detecting a transition

- best method depends on the system

- simple method for EAM metal systems:
  periodically perform steepest-descent quench;
  see if geometry at basin minimum has changed

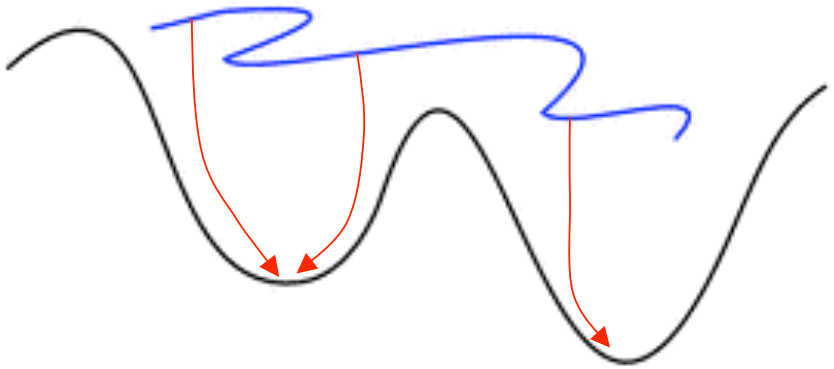

Parallel-replica dynamics example

Ag(111) island-on-island decay

Embedded atom (EAM) potential

Temperature = 400K

5 days on 32 processors
(1 GHz Pentium-IIIs)

Upper island decays into lower island via step-edge exchange events.

For this case, parallel replica is even faster than temperature-accelerated dynamics, since barriers are low relative to T.

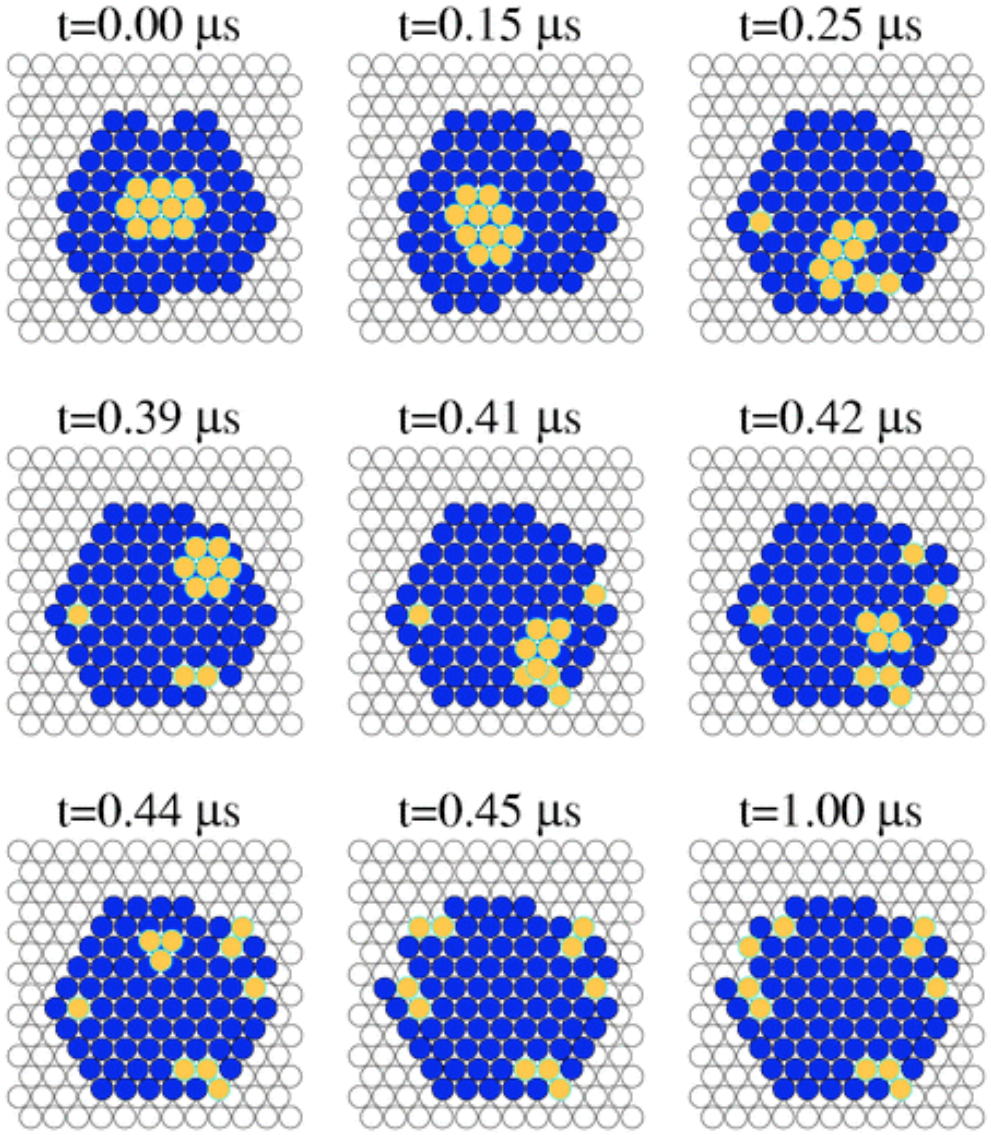

AFV, F. Montalenti and T.C. Germann, Ann. Rev. Mater. Res. 32, 321 (2002).

Must p(t) be exponential?

What if p(t) looks like this?

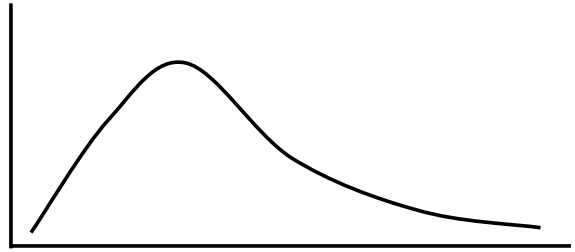

Must p(t) be exponential?

What if p(t) looks like this?

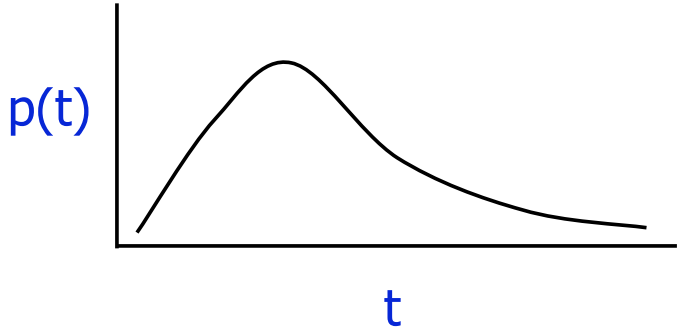

This is a problem.

Imagine the limiting case:

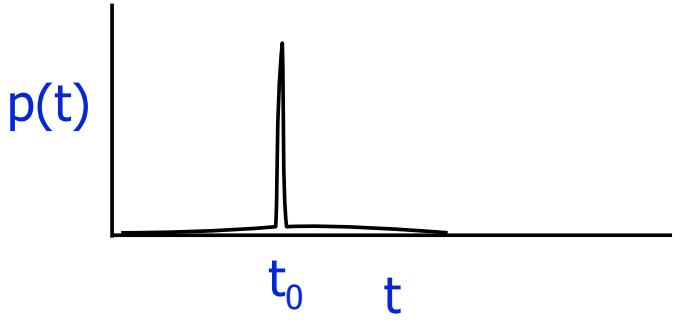

If $p(t)=\delta(t-t_0)$, parallel replica simulation will detect an event on all M processors at the same time (t_0), giving no boost. Worse, the summed time will predict a transition time too large by exactly a factor of M.

General comments about non-exponential p(t)

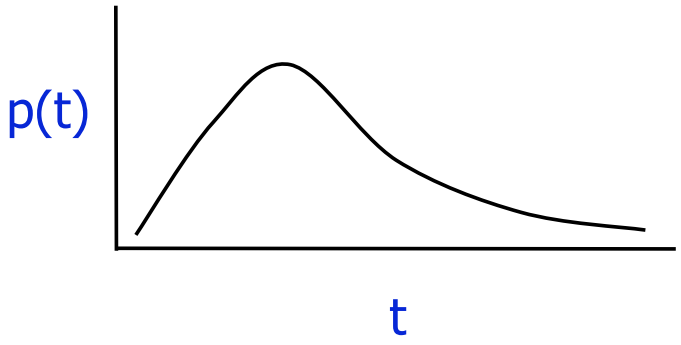

A non-exponential p(t) is almost certainly caused by "hidden" transitions, or, similarly, a diffusive character to some coordinate(s).

Thus, applying parallel-replica dynamics when p(t) is non-exponential will give incorrect times, and incorrect dynamics.

If every transition is detected, each p(t) will be exponential for a (non-diffusive) infrequent-event system, and parallel-replica dynamics will be valid.

Parallel Replica Dynamics
Can we overlook fast transitions?

The maximum boost is limited by the fastest processes in the system.

If the rapid transitions seem to be unimportant, can we ignore them?

Yes - but be careful…

Safe case

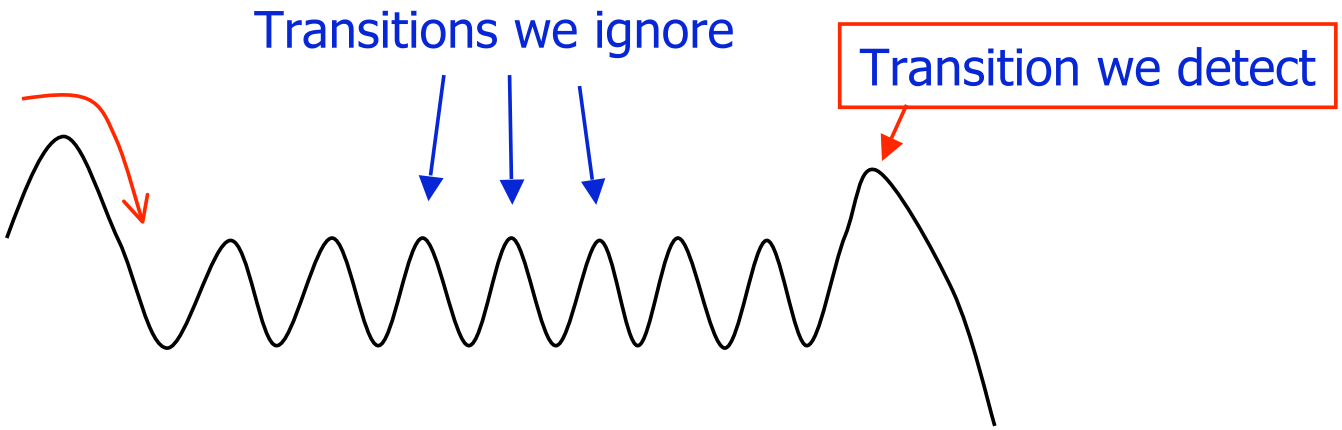

Trajectory will go back and forth many times before escaping

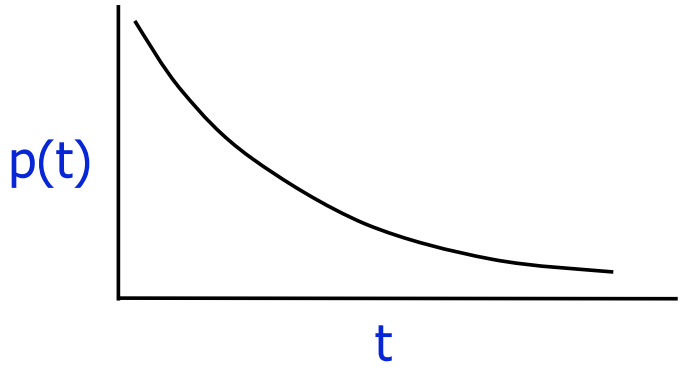

Dangerous case

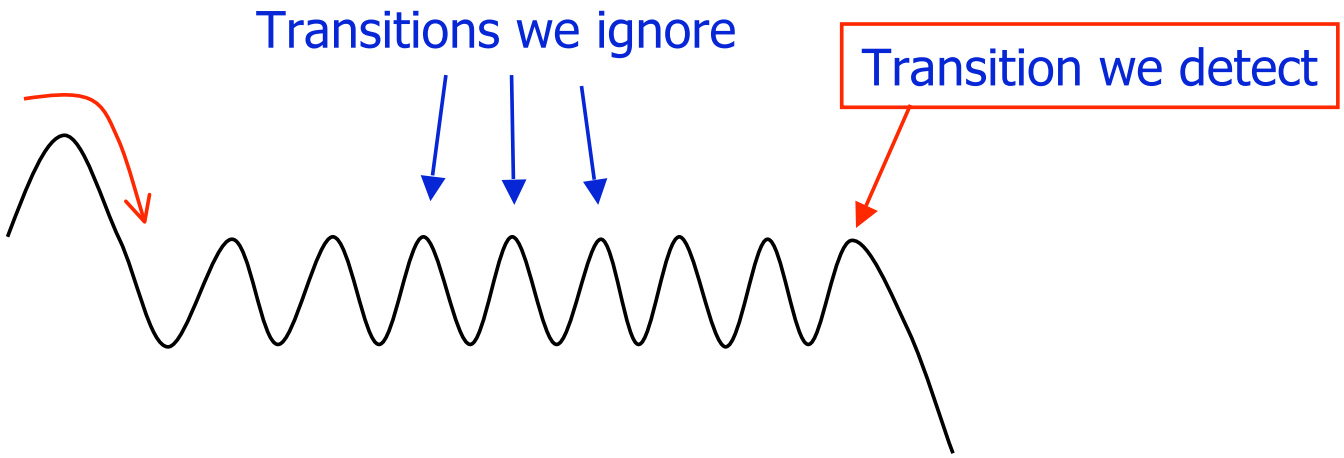

Trajectory may fall off right-hand end on first pass - acts like an exponential preceded by a lag time while system diffuses from left to right

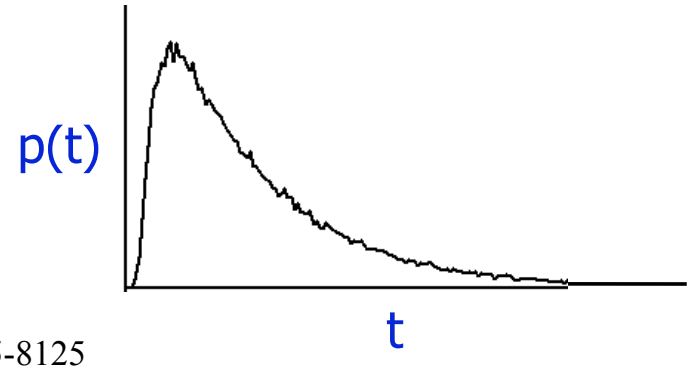

Parallel replica on floppy systems
(ignoring fast conformational transitions)

Small "fast folding" proteins (Pande et al, 2000, 2001,...)

- folding transitions detected based on energy fluctuations
- screen-saver based (SETI-at-home-like) parallelization
- microsecond time scales, folding seen

Hydrocarbon pyrolysis (Kum, Stuart et al, Clemson, 2002)

- bond-breaking transitions detected
- microseconds at T=2000K (usually T=3000K with MD)
- unexpected ring-formation observed

Pyrolysis of n-hexadecane ($C_{16}H_{34}$)

3-carbon-ring formation (not in standard models)

Subsequent ring opening leads to branched structure

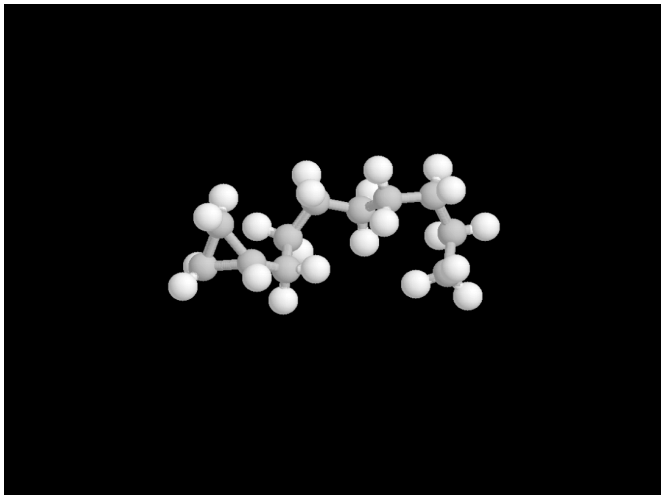

Cyclic intermediate
(11 carbons left)

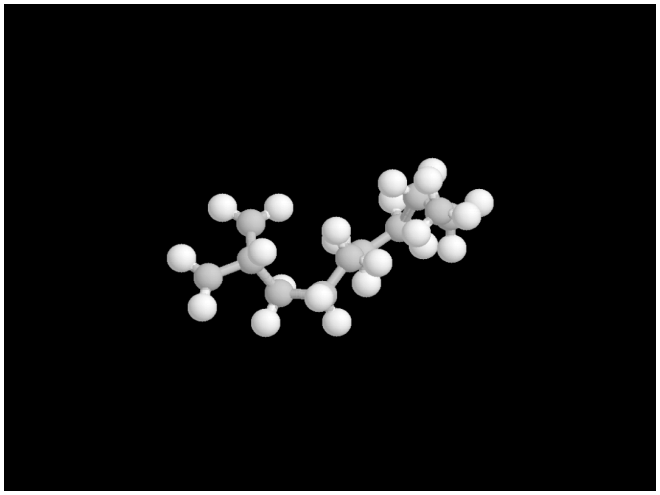

Branched radical after
ring reopens.

REBO potential

Reactive Bond-Order potential

Brenner et al, J. Phys.: Condens. Matter 14, 783 (2002)

- fit to lots of C and C-H data
- improved property predictions relative to "Tersoff Brenner" form

  Brenner, Phys. Rev. B 42, 9458 (1990).
- dihedral angle term for C=C bonds
- good description of graphite

Summary - Parallel replica dynamics

Most exact of the accelerated dynamics methods
- no harmonic approximation
- goes beyond TST to include correlated dynamical events
- no assumption that barrier is energetic - can be entropic

Easy to implement; Requirements:
- transition detection
- good estimate of correlation time

Very general applicability
- any system with exponentially distributed events

Good match to increasing availability of parallel processing power, distributed computing, etc.

Limitations - Parallel replica dynamics

Only get parallel boost

Boost drops off when events are frequent (or become frequent because so many processors)

Boost especially limited if correlation time is long

For complex systems (e.g., proteins) even transition detection can be tricky.

Hyperdynamics

Hyperdynamics

Builds on umbrella-sampling techniques (e.g., Valleau 1970's)

Assumptions:

- infrequent events
- transition state theory (no recrossings)

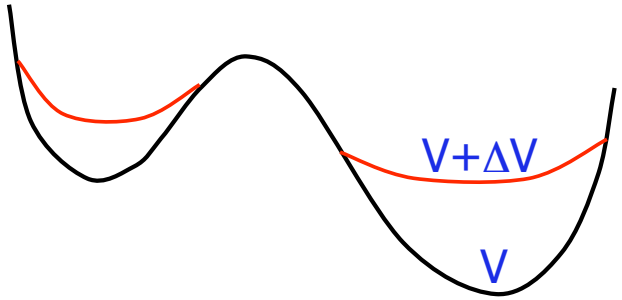

Procedure:

- design bias potential ΔV (zero at dividing surfaces; causes no recrossings)
- run thermostatted trajectory on the biased surface ($V+\Delta V$)
- accumulate hypertime as

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

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

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

The hypertime clock

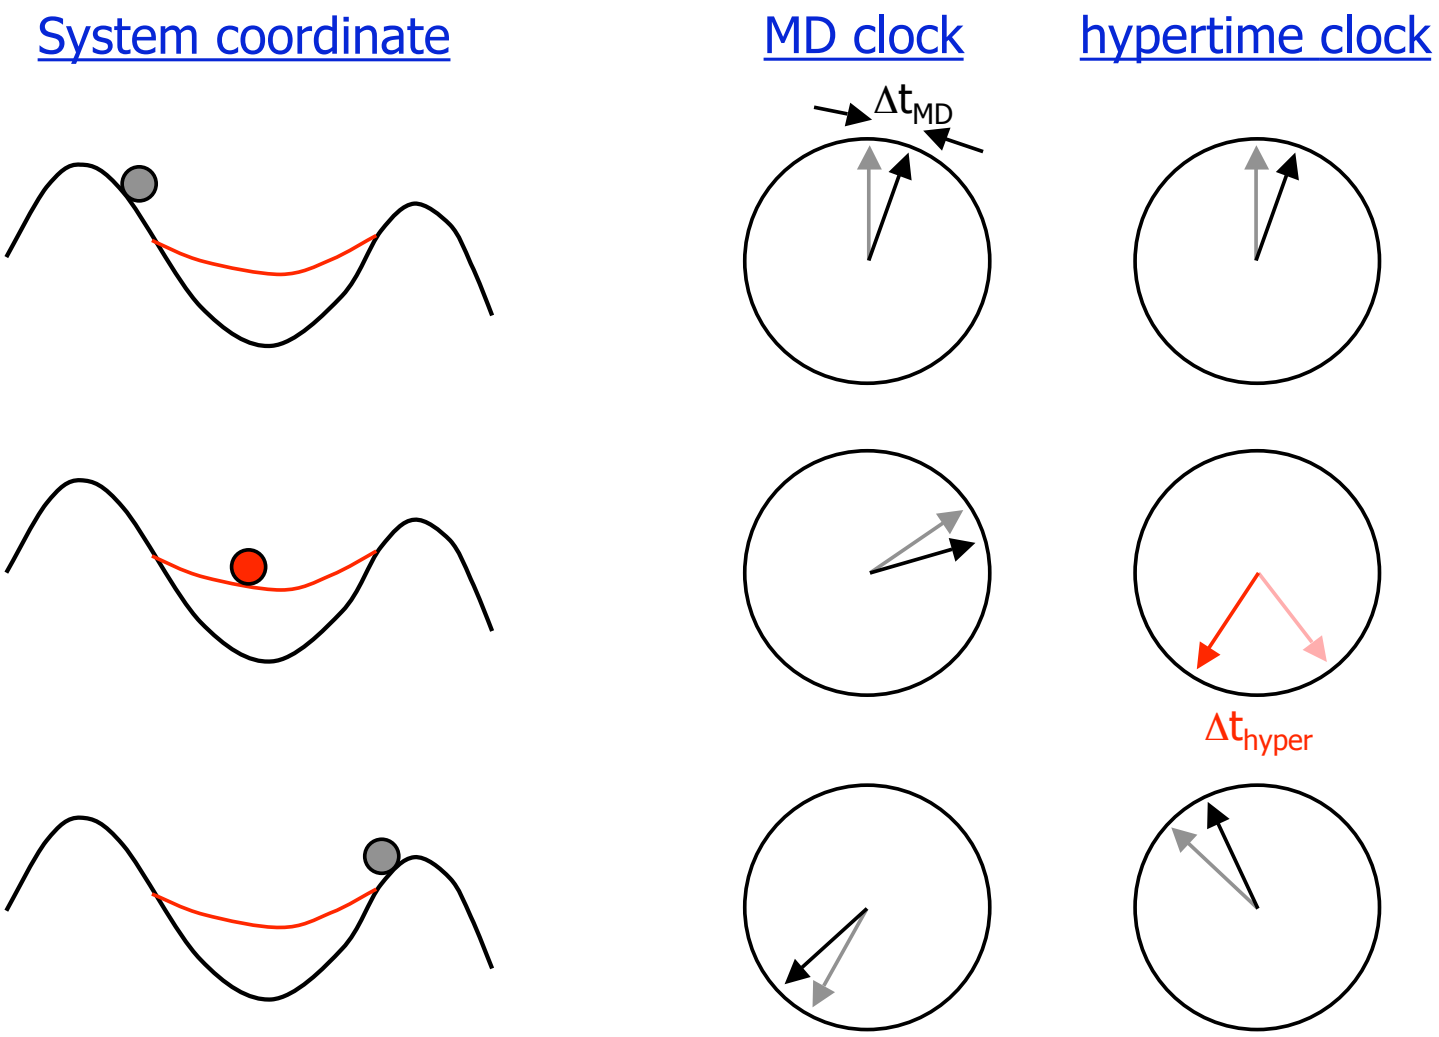

Boost = hypertime/(MD clock time)

The hypertime clock

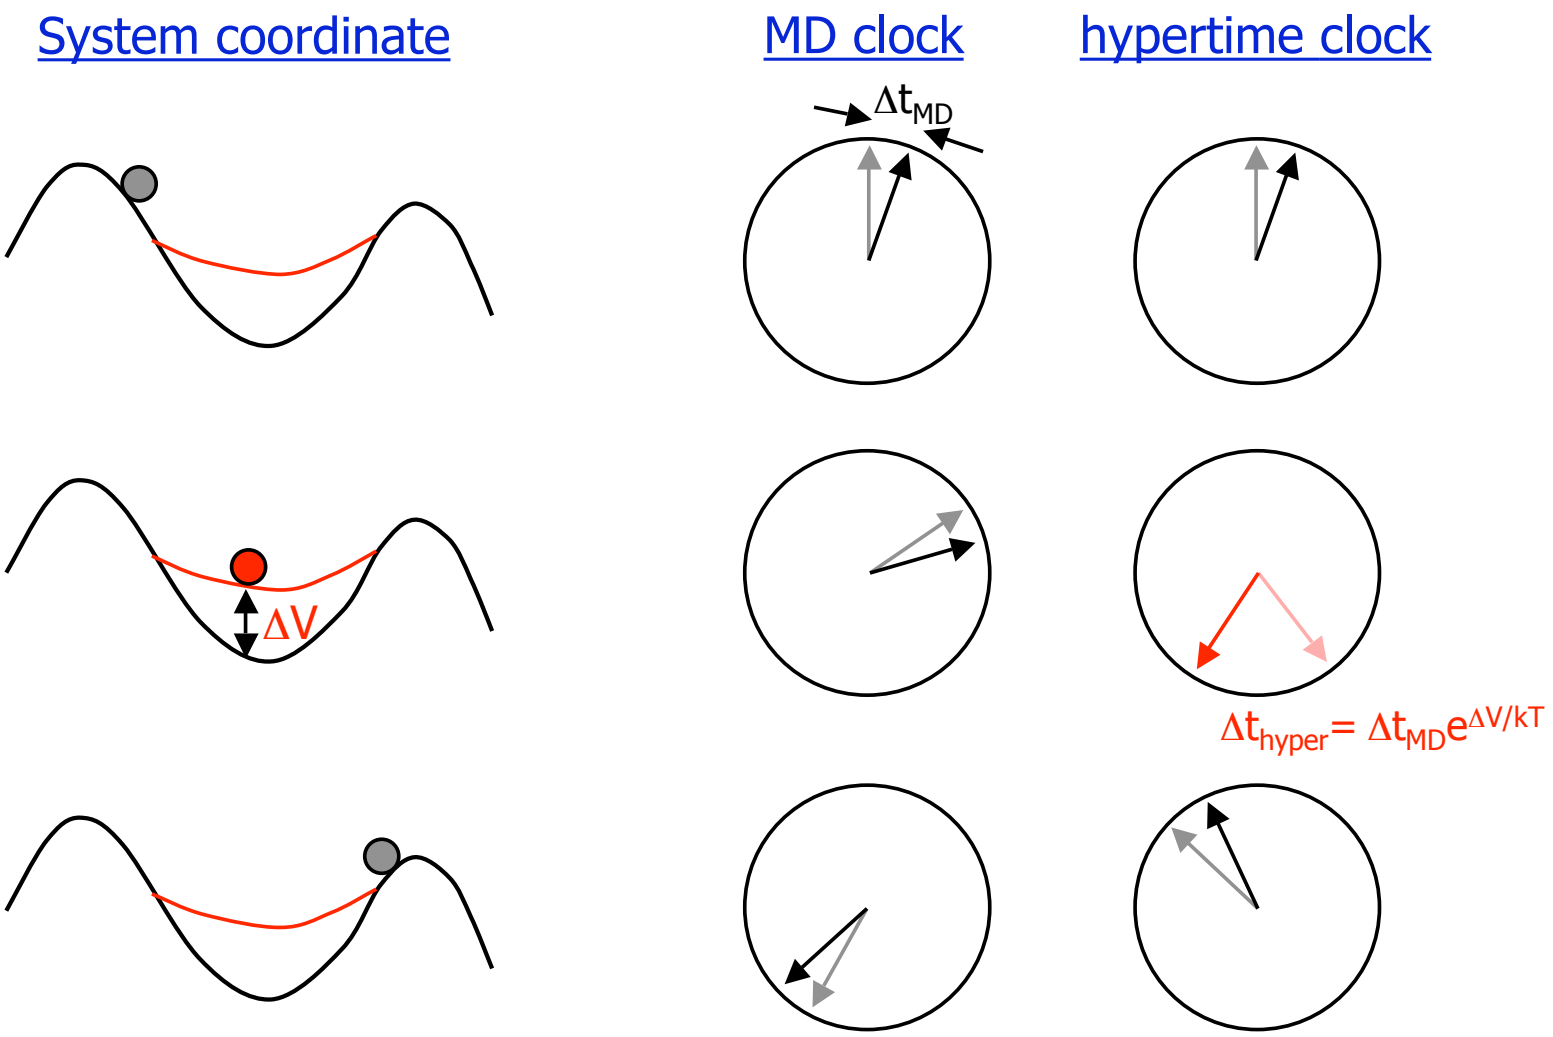

Boost = hypertime/(MD clock time)

Hyperdynamics - derivation

Viewgraphs here

Hyperdynamics

Key challenge is designing a bias potential that meets the requirements of the derivation and is computationally efficient.

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

Flat bias potential

An extremely simple form: flat bias potential

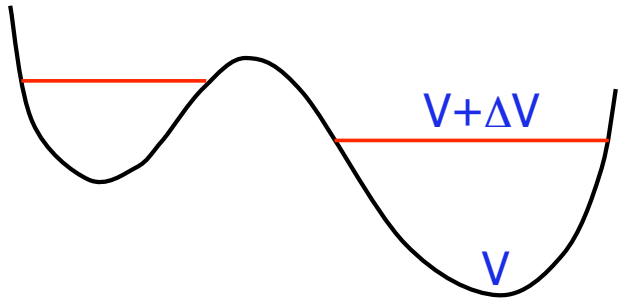

M. M. Steiner, P.-A. Genilloud, and J. W. Wilkins, Phys. Rev. B **57**, 10236 (1998).

- 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

Total ΔV --> 0 when any bond is stretched beyond some threshold value

Looks promising:

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

Summary - Hyperdynamics

Powerful if an effective bias potential can be constructed

Need not detect transitions

Exponential boost factors possible

Especially effective if barriers high relative to T

Possibilities for powerful system-specific bias potentials

Lots of possibilities for future development of advanced bias potential forms

Limitations - Hyperdynamics

Must design bias potential

Assumes TST holds (though Langevin jiggle recrossings may be OK)

Boost drops off when events are frequent

Harder to implement properly if bottlenecks are entropic (but possible)

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.

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

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.

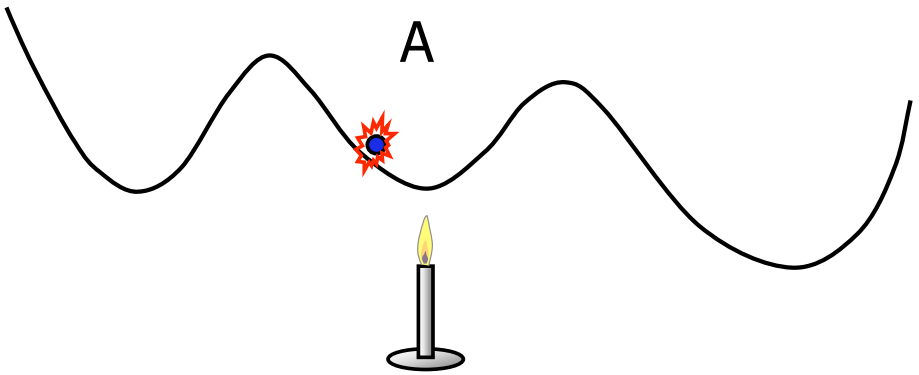

TAD Procedure

- 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.

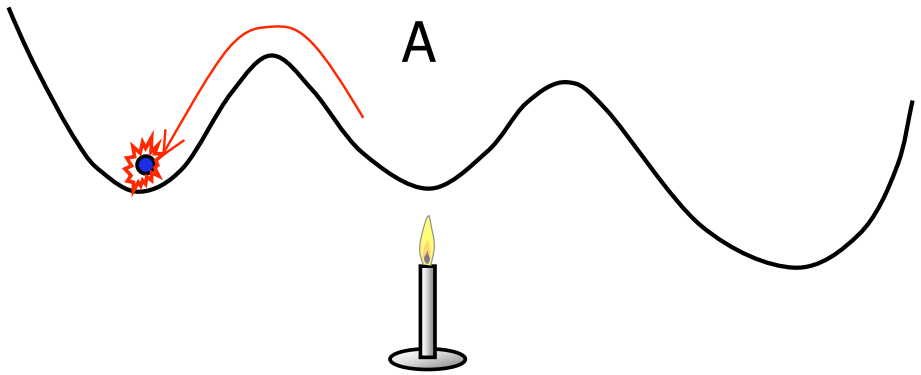

TAD Procedure

- 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.

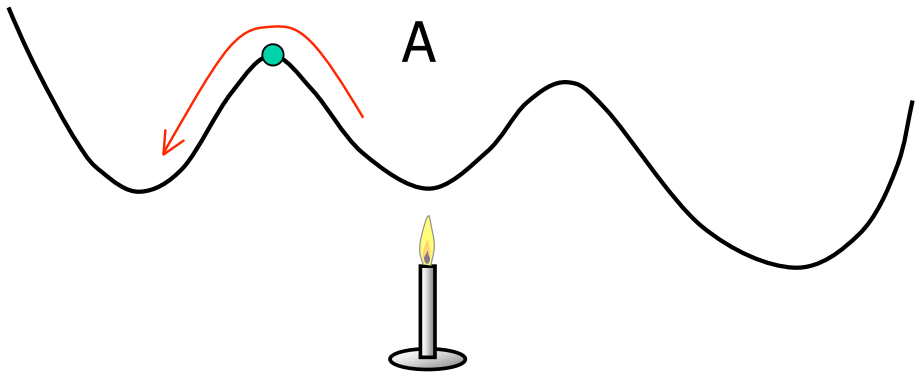

TAD Procedure

- 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.

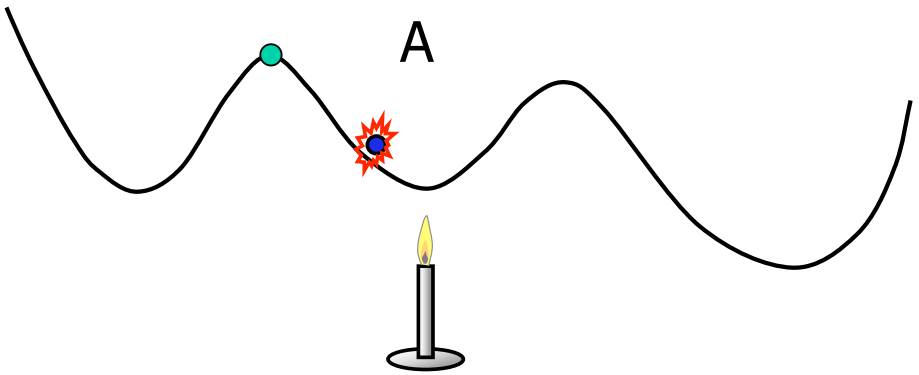

TAD Procedure

- 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.

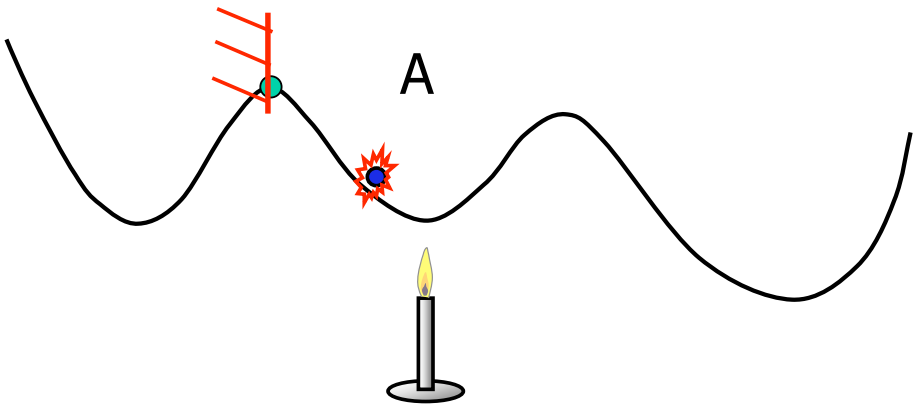

TAD Procedure

- 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.

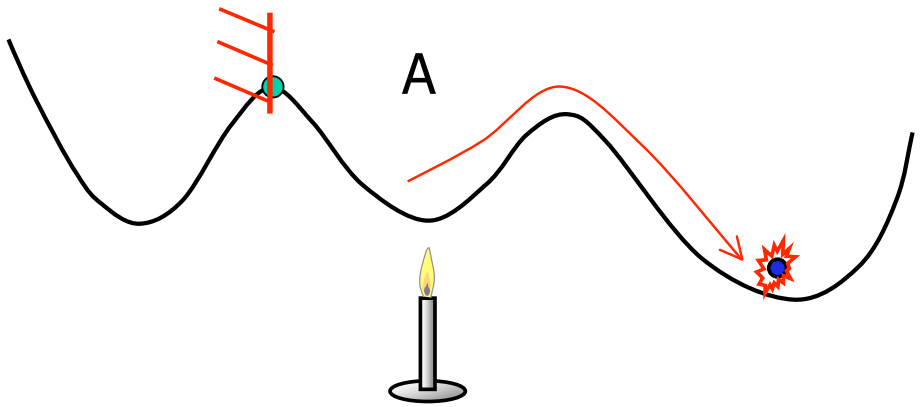

TAD Procedure

- 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.

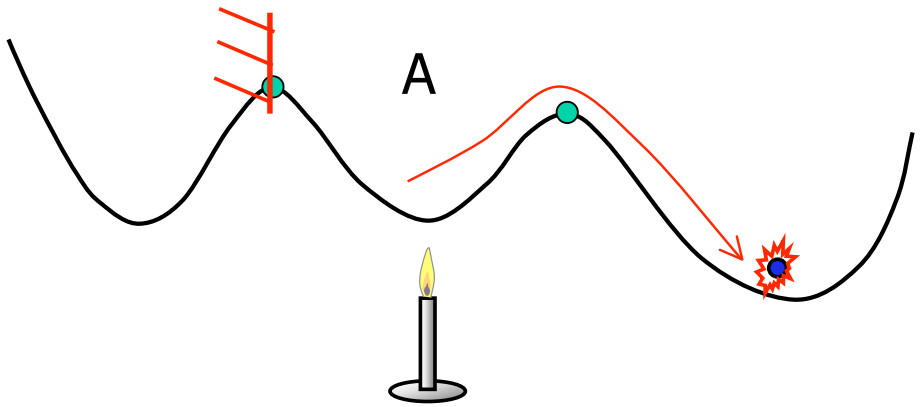

TAD Procedure

- 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.

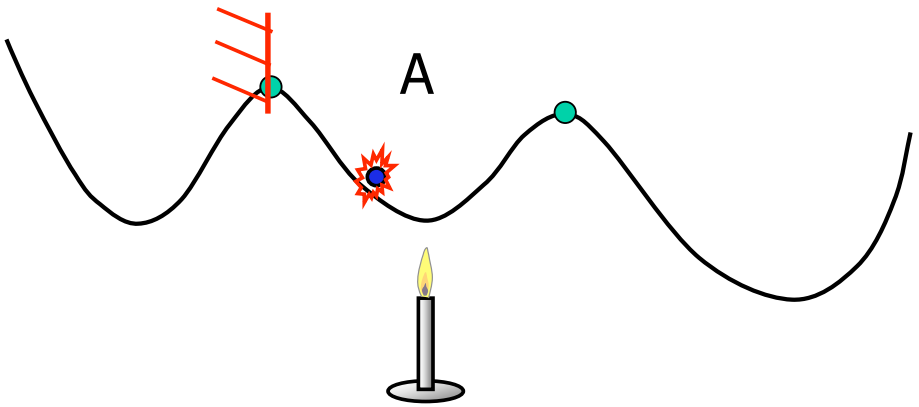

TAD Procedure

- 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.

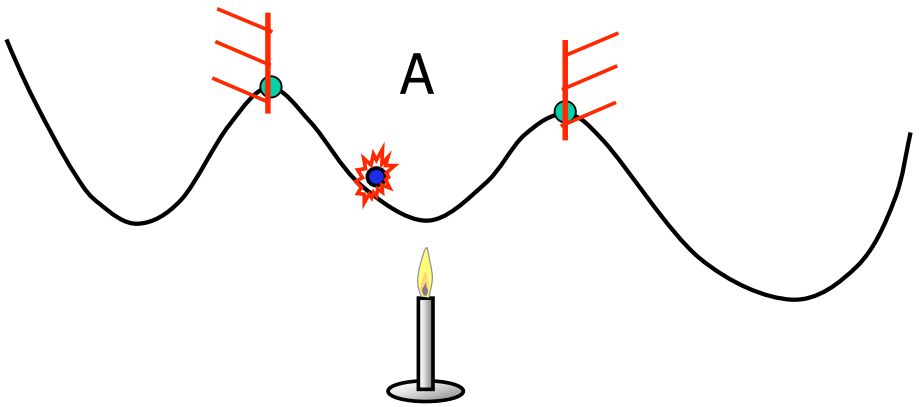

TAD temperature-extrapolated time

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:

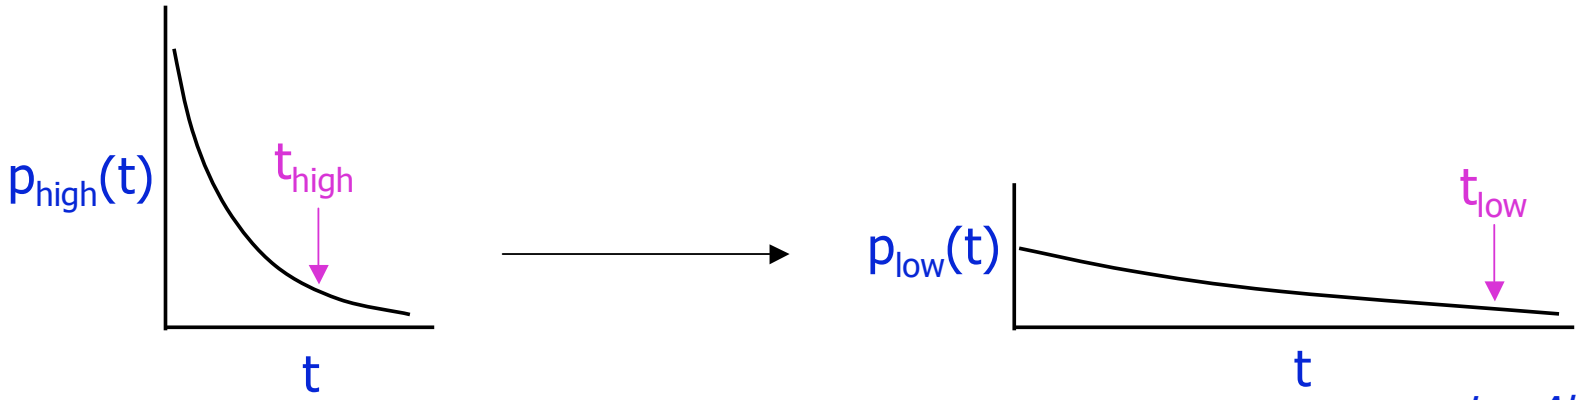

The Arrhenius view

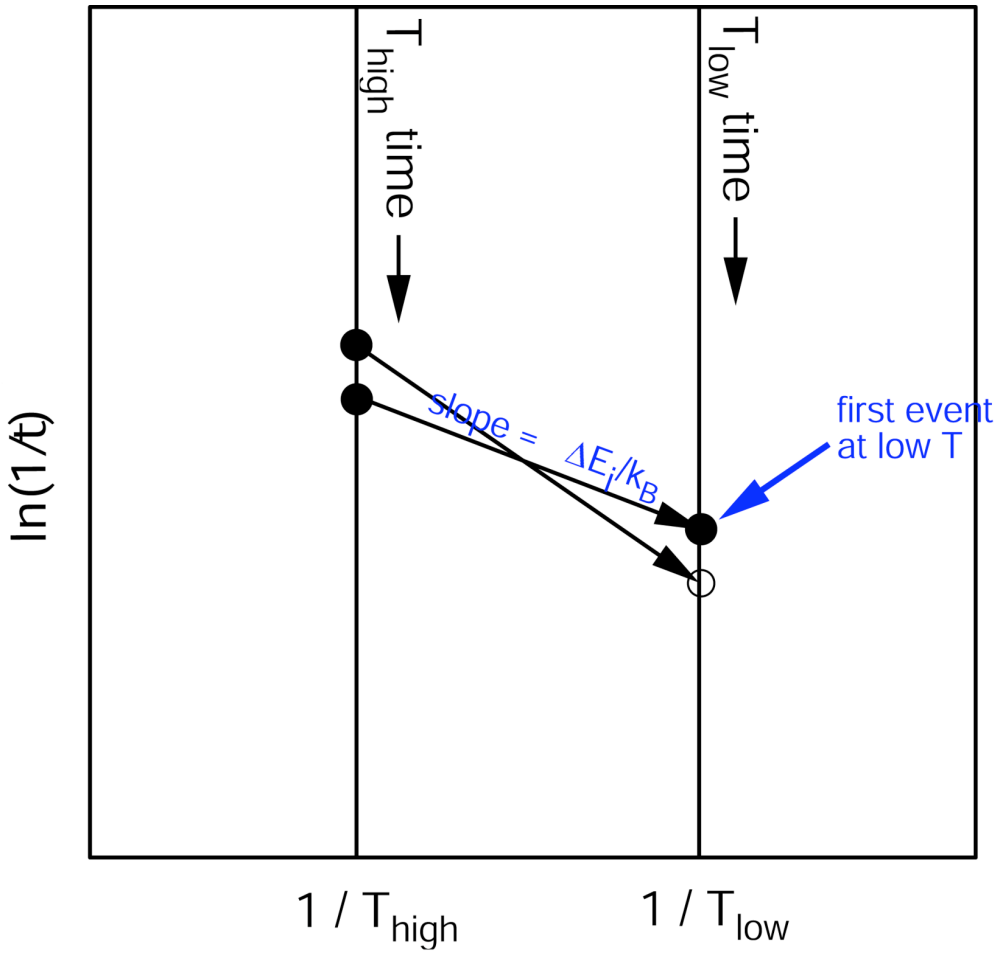

when can we stop?

The confidence line

For a given rate k_i, the time required to be f-confident to see the first event (e.g., f=0.99 for 99% confidence) is

$$\tau_i^f = (1/k_i)\ \ln[1/(1-f)] \qquad \text{(or } \delta = 1-f\text{)}$$

If rate is Arrhenius ($k = \nu_0 \exp[-\Delta E/k_B T]$), then fraction f of first attempts will occur *above* the line with slope $-\Delta E/k_B$ and intercept $\nu_0^f = \nu_0/\ln[1/(1-f)]$

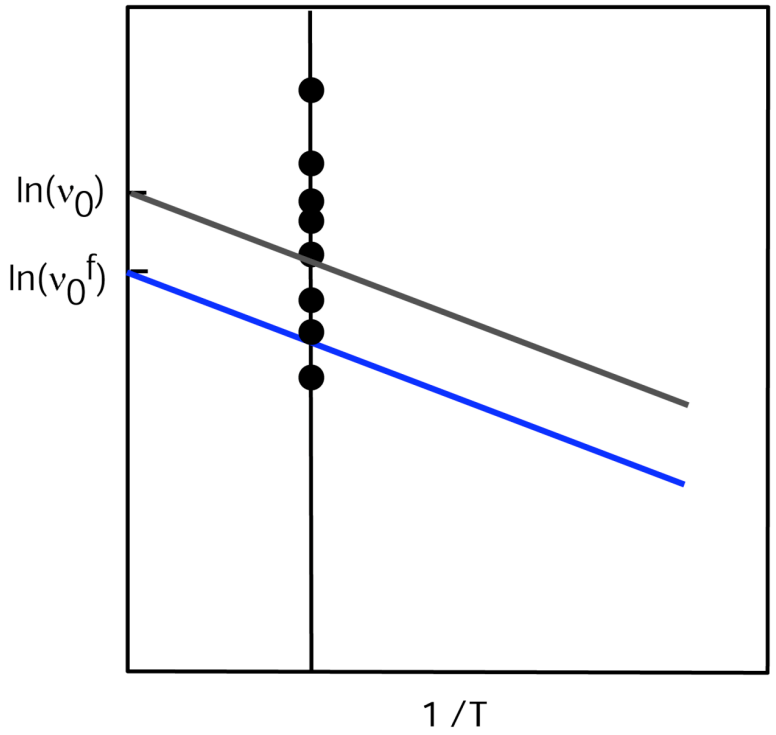

TAD - when can we stop the MD and accept an event?

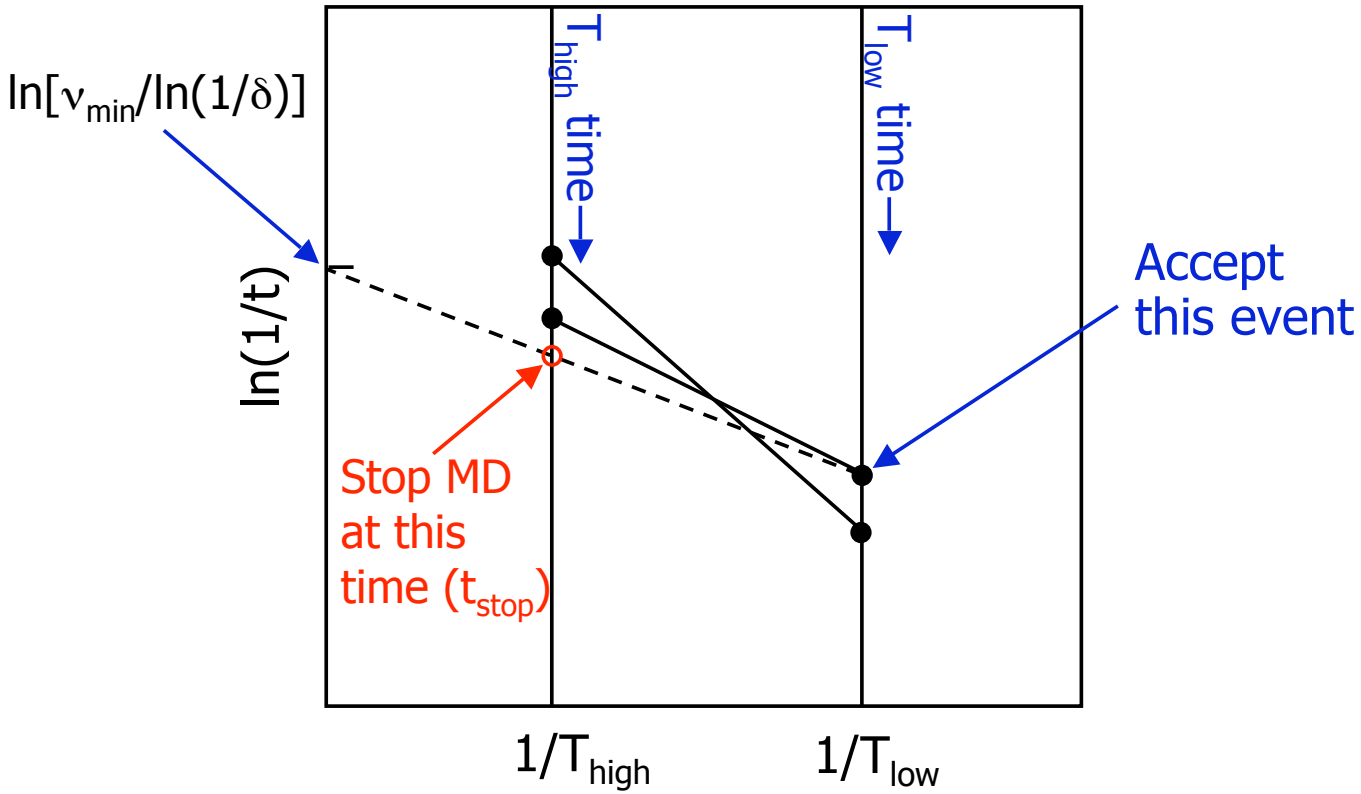

After time t_{stop}, with confidence 1-δ, no event can replace shortest-time event seen at low T.

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)

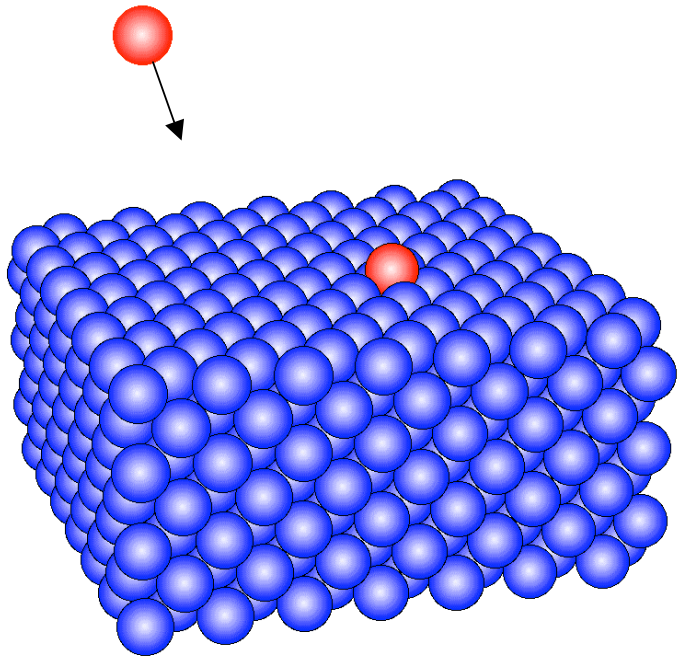

MD+TAD deposition of Cu/Cu(100)

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

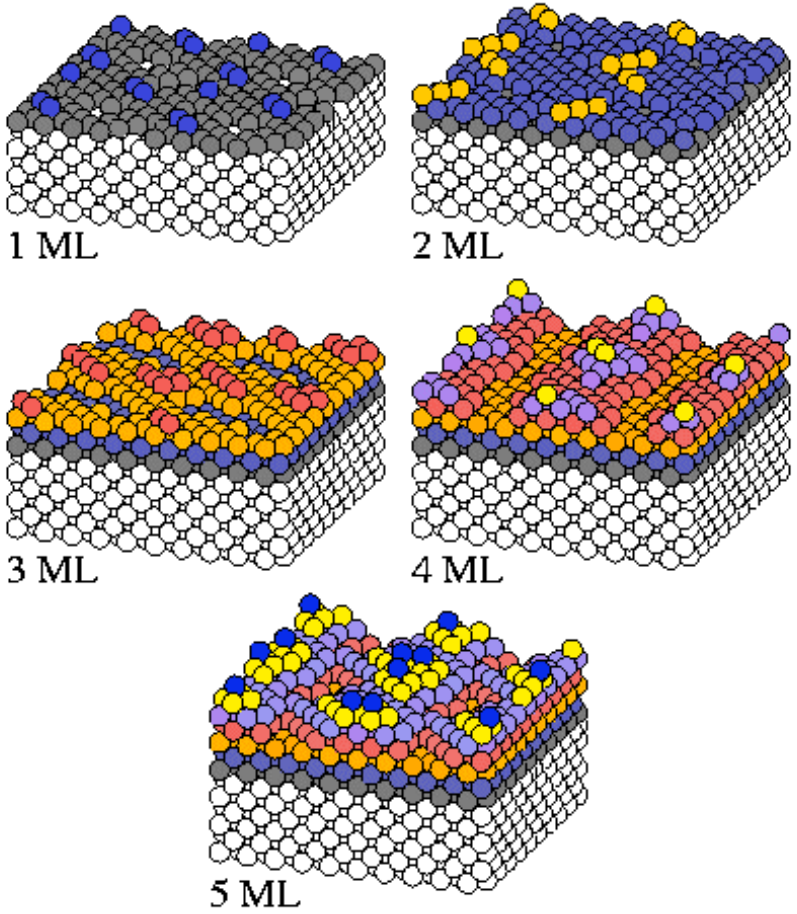

boost factor ~10^7

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

MD+TAD deposition of Cu/Cu(100)

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

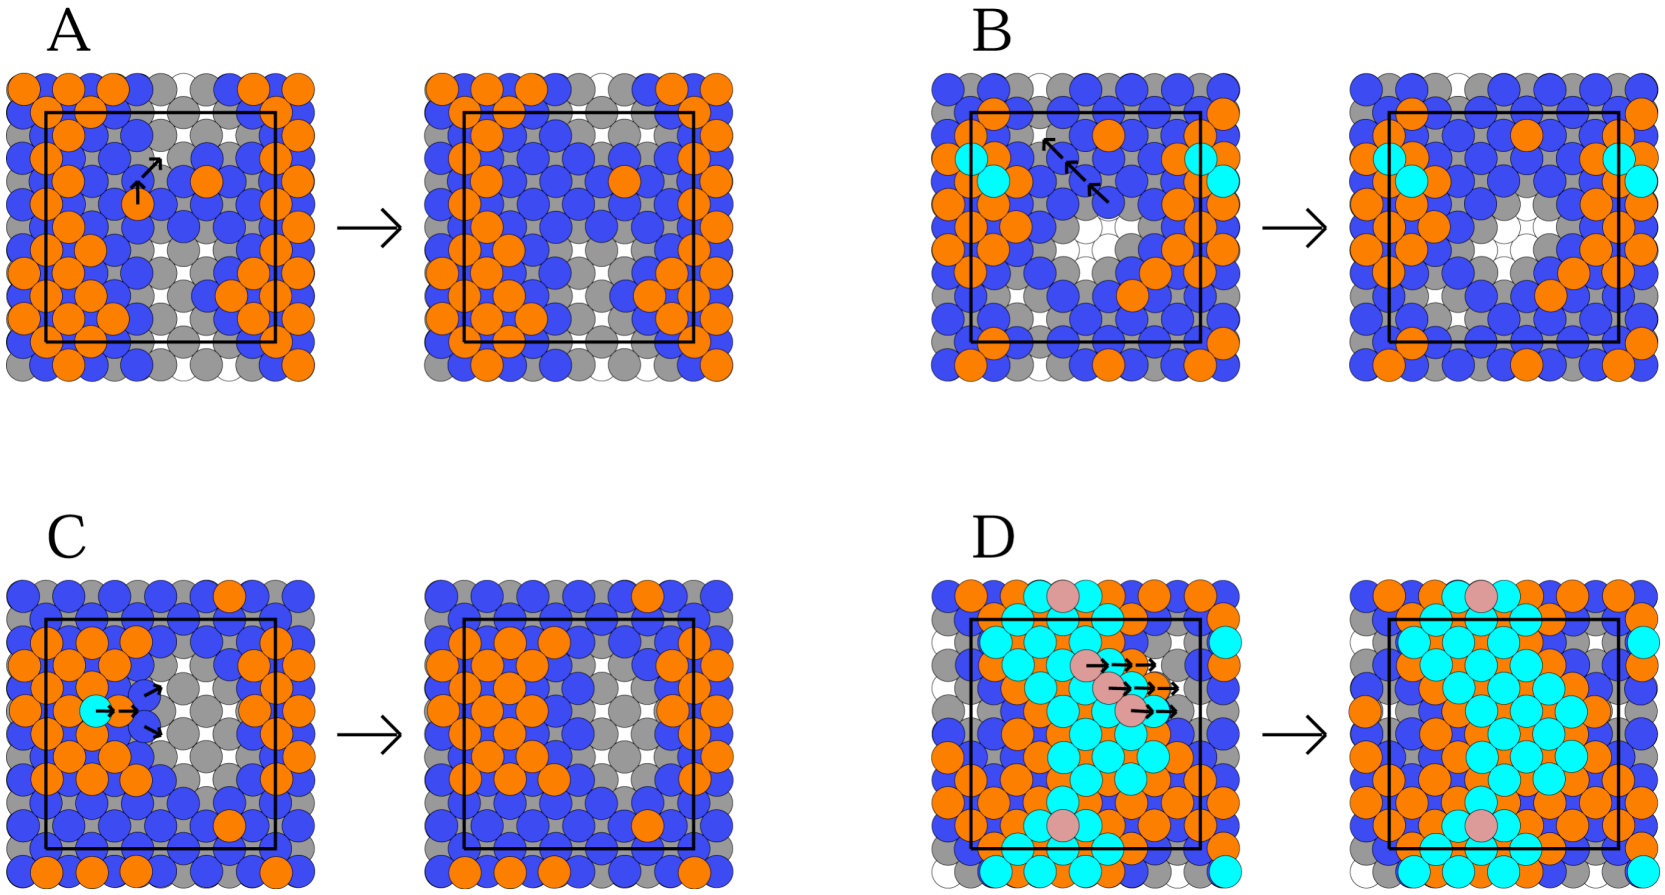

Slice of a nanotube evolved with TAD

High Temperature: 3000K

Low Temperature: 1500K

States Visited: 100 shown

Boost: 129

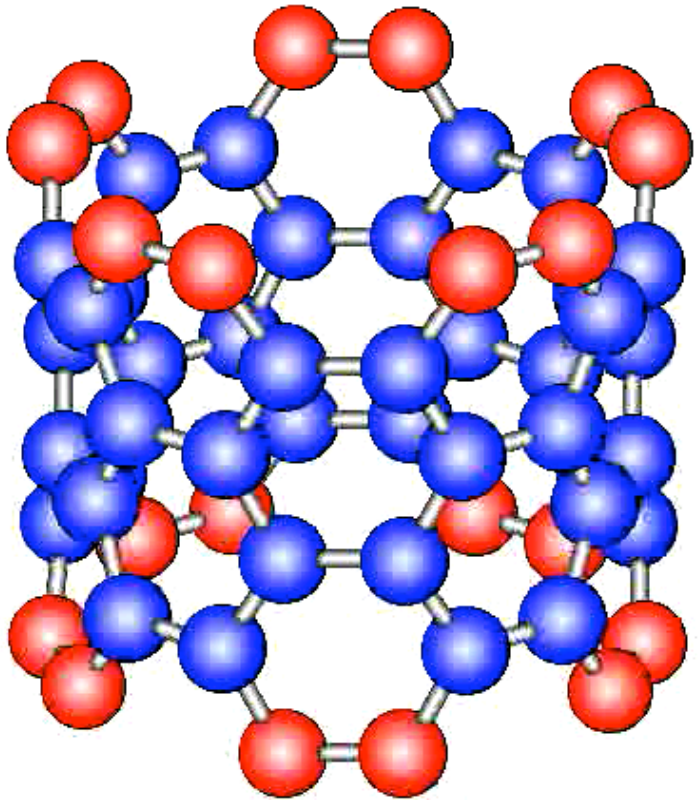

0 - 2 μs

Blas Uberuaga
Steve Stuart

Slice of a nanotube evolved with TAD

High Temperature: 3000K

Low Temperature: 1500K

States Visited: 100 shown

Boost: 129

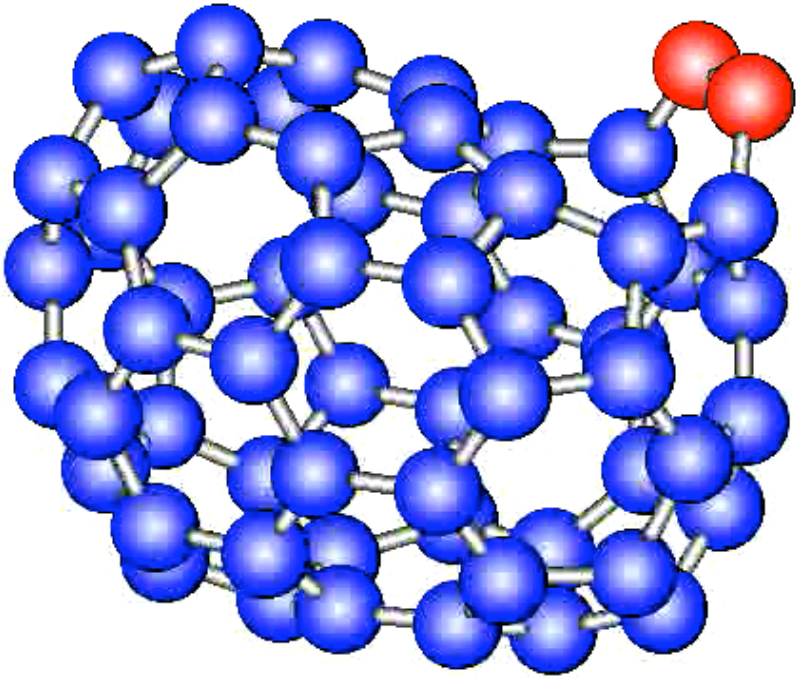

after 14 μs

Blas Uberuaga
Steve Stuart

TAD: More 60-atom Nanotube Fragments

T=1500K (High T = 3000K), after first 1000 transitions

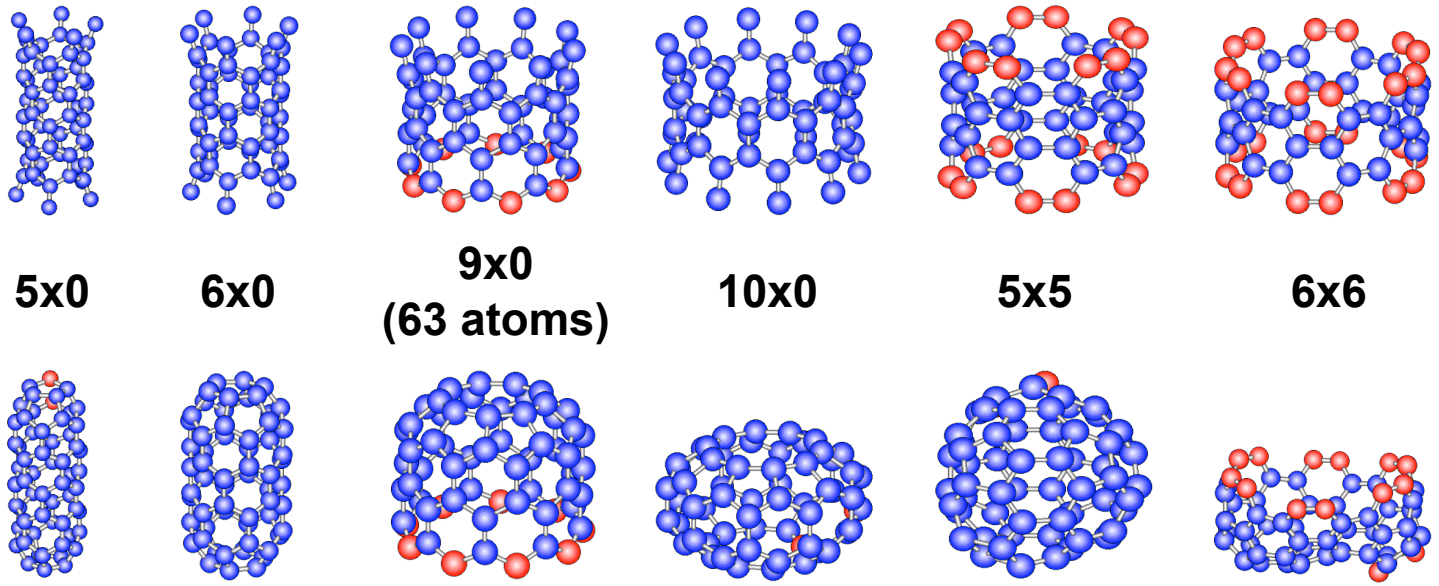

5x0	6x0	9x0 (63 atoms)	10x0	5x5	6x6
Total Simulation Time (μs)					
59	20	2.7	0.2 (133 states)	4	3
Boost					
1445	796	80	7	108	100

Blas Uberuaga *Los Alamos*

TAD: More 60-atom Nanotube Fragments

T=1500K (High T = 3000K), after first 1000 transitions

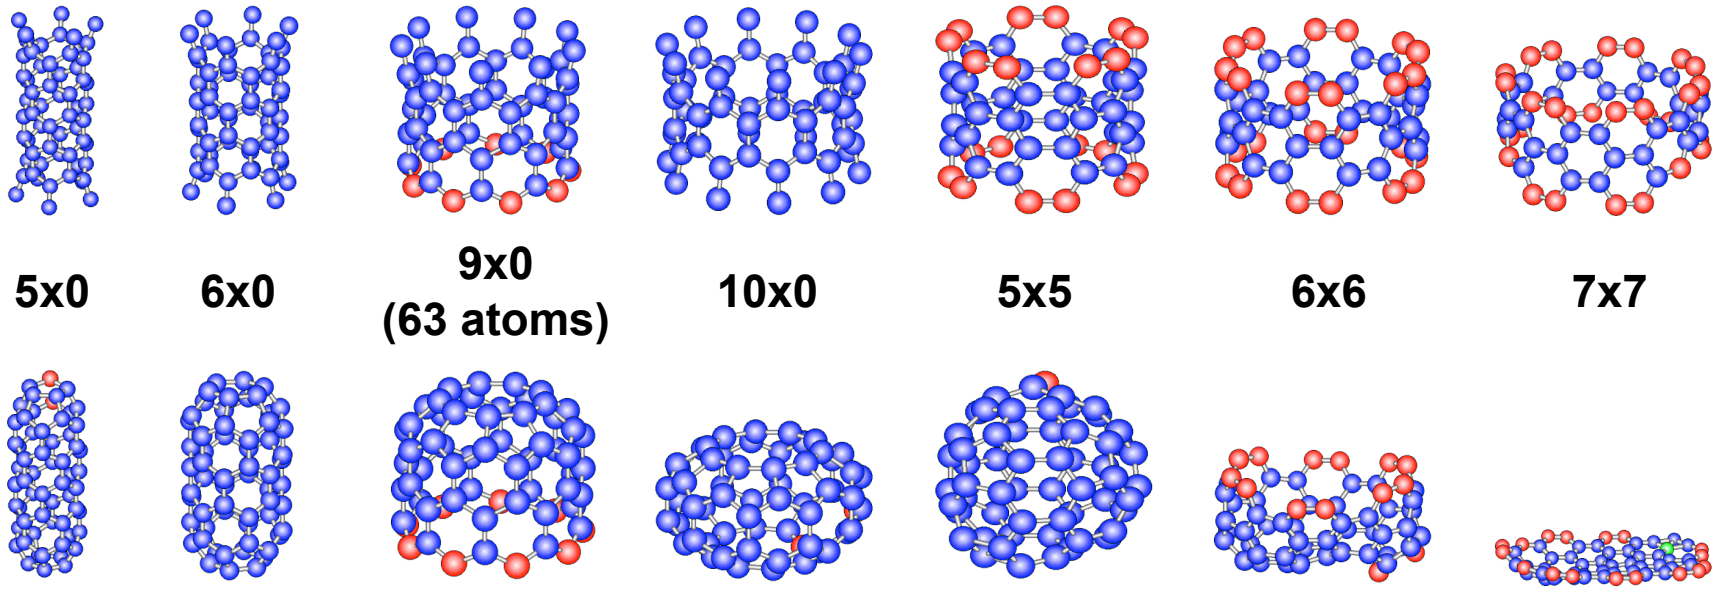

5x0	6x0	9x0 (63 atoms)	10x0	5x5	6x6	7x7
Total Simulation Time (μs)						
59	20	2.7	0.2 (133 states)	4	3	1
Boost						
1445	796	80	7	108	100	29

Blas Uberuaga *Los Alamos*

Slice of a 7x7 nanotube evolved with TAD

High Temperature: 3000K

Low Temperature: 1500K

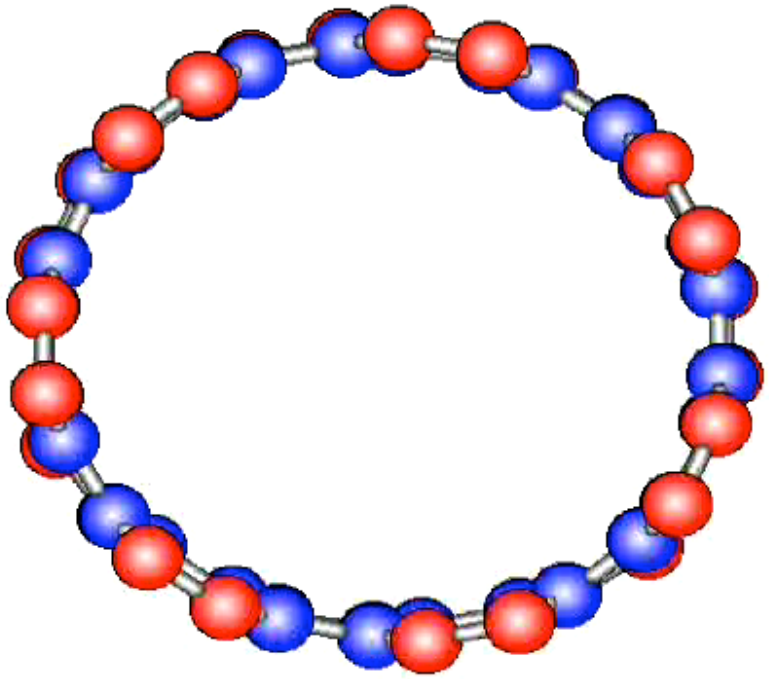

first few
transitions

Slice of a 7x7 nanotube evolved with TAD

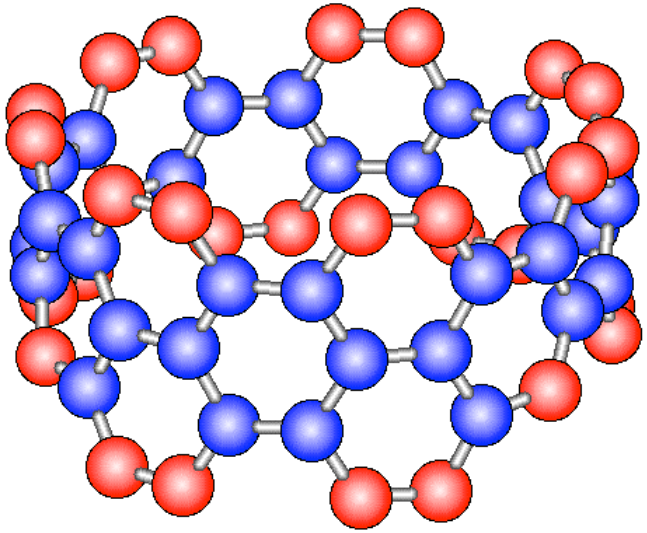

t=0

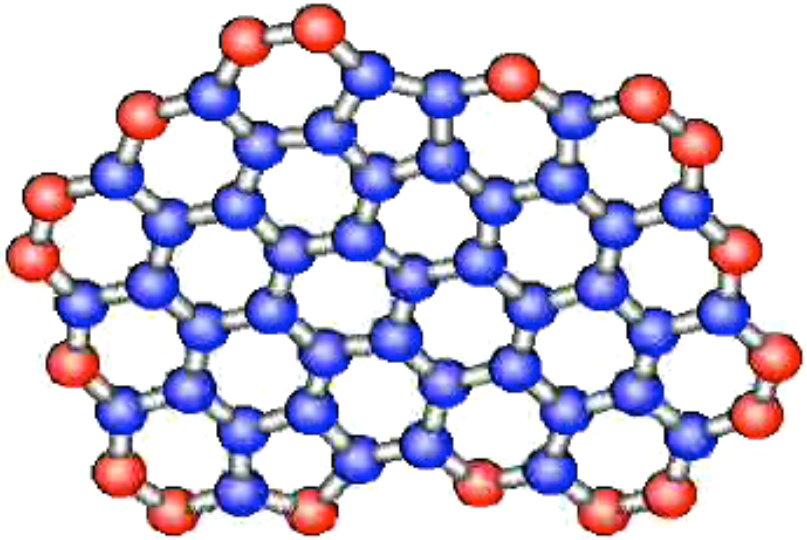

8.5 μs

(~25 eV above buckyball)

Blas Uberuaga
A.F. Voter, September, 2005 LAUR-05-8125

Using minimum barrier to stop TAD sooner

Assume we know the minimum barrier (ΔE_{min}) for escape from this state

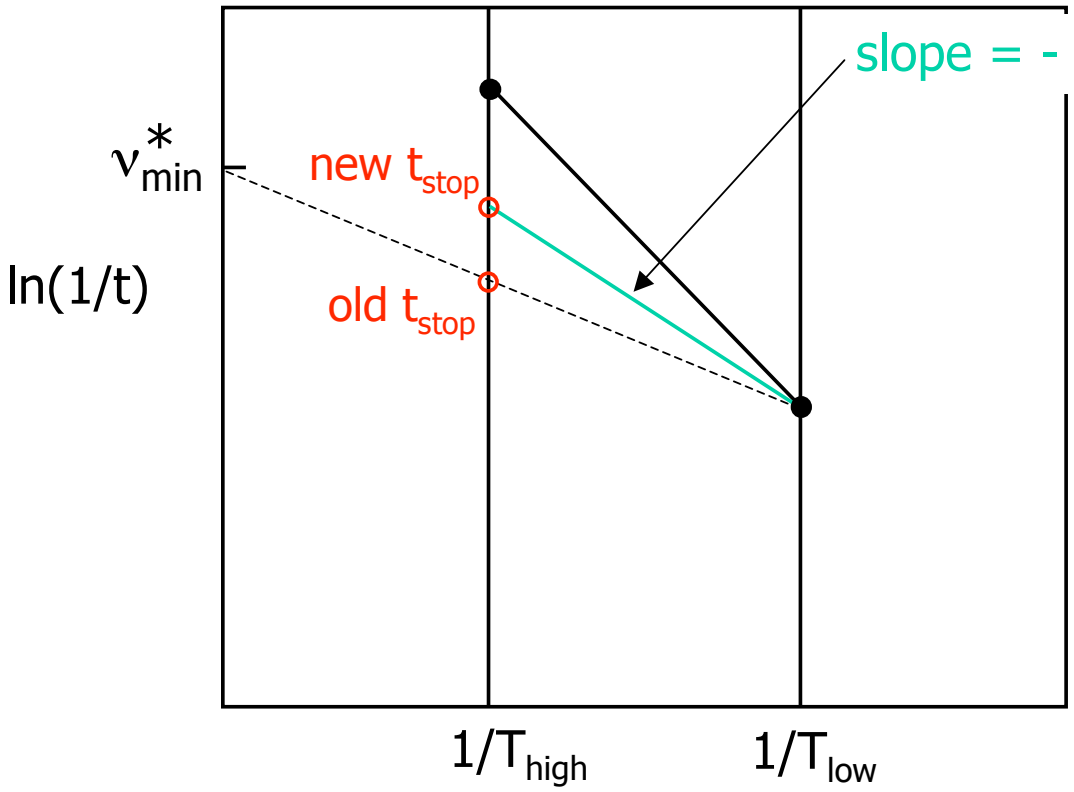

"ETAD"

F. Montalenti and AFV, J. Chem. Phys. 116, 4819 (2002).

Combining dimer method with TAD

Dimer-TAD

- Use dimer method (Henkelman and Jonsson, 1999) to find a number of saddles and assume the lowest barrier (ΔE_{min}) is among them
- Supply this ΔE_{min} to ETAD for this state

----> accuracy of TAD (unless lowest barrier missed), with roughly the speed of the dimer method

MgO Radiation Damage Annealing T=300K

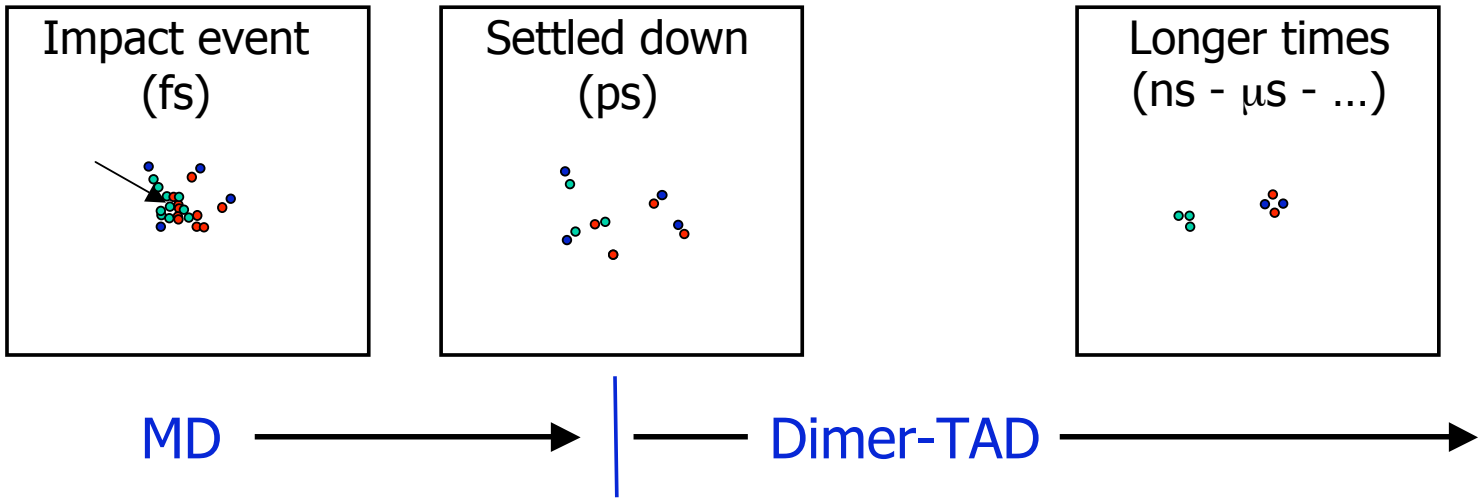

Coulombic Buckingham potential

Interesting picture emerges for annealing after 400 eV cascade.

Uberuaga, Smith, Cleave, Montalenti, Henkelman, Grimes, Voter, and Sickafus, Phys. Rev. Lett., **92**, 115505 (2004); Phys. Rev. **71**, 104102 (2005).

Growth of interstitial clusters in MgO, T=300K

Typical 400 eV collision event forms a few vacancies and interstitials

Diffusing interstitials coalesce into clusters (vacancies are immobile)

Mono-interstitial - diffuses on ns-μs time scale

Di-interstitial - diffuses on s time scale

Tetra-interstitial - immobile (years)

Growth of interstitial clusters in MgO, T=300K

Typical 400 eV collision event forms a few vacancies and interstitials

Diffusing interstitials coalesce into clusters (vacancies are immobile)

Mono-interstitial - diffuses on ns-μs time scale

Di-interstitial - diffuses on s time scale

Tetra-interstitial - immobile (years)

Is the tetramer a sink
for all larger clusters?

Growth of interstitial clusters in MgO, T=300K

Typical 400 eV collision event forms a few vacancies and interstitials

Diffusing interstitials coalesce into clusters (vacancies are immobile)

Mono-interstitial - diffuses on ns-μs time scale

Di-interstitial - diffuses on s time scale

Tetra-interstitial - immobile (years)

Is the tetramer a sink
for all larger clusters? No!

TAD Simulation: dimer + tetramer interstitial clusters

- In this case, dimer + tetramer forms hexamer in metastable state
- Metastable hexamer exhibits fast one-dimensional diffusion!
  - ns timescale
  - diffusion is 1D along <110>
  - decay to ground state takes years

(perfect bulk atoms not shown, red=O--, blue=Mg++)

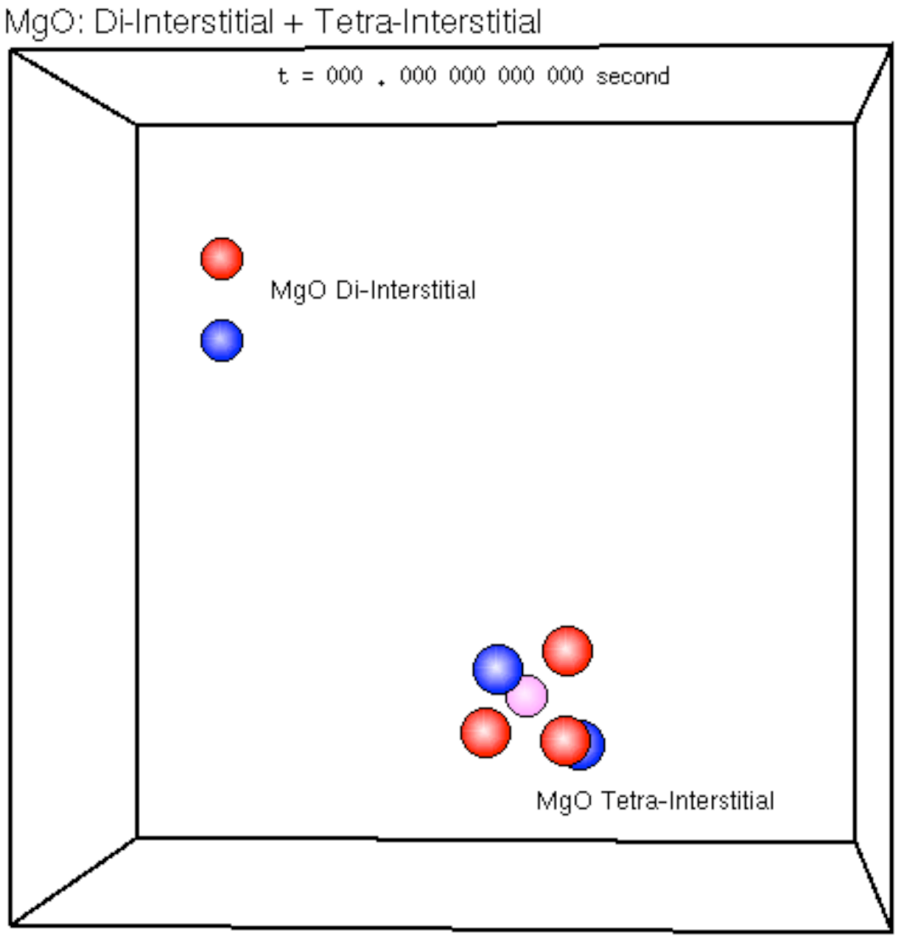

TAD simulation, Uberuaga et al, 2003

Interstitial Cluster Mobility

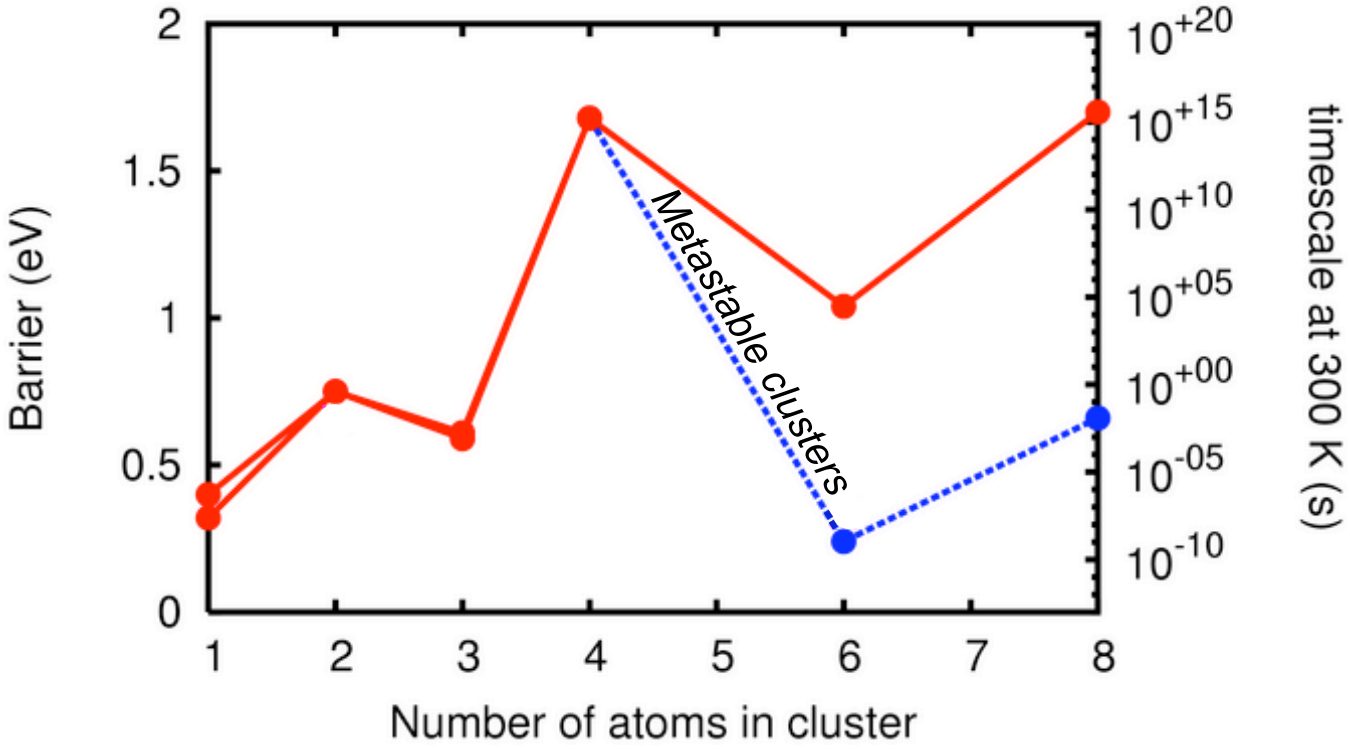

- Diffusion barrier follows no clear pattern versus cluster size
- Metastable clusters can be surprisingly mobile
- Metastable clusters can be very stable:
  - metastable I_6 lifetime is years at room temperature

Interstitial Cluster Mobility

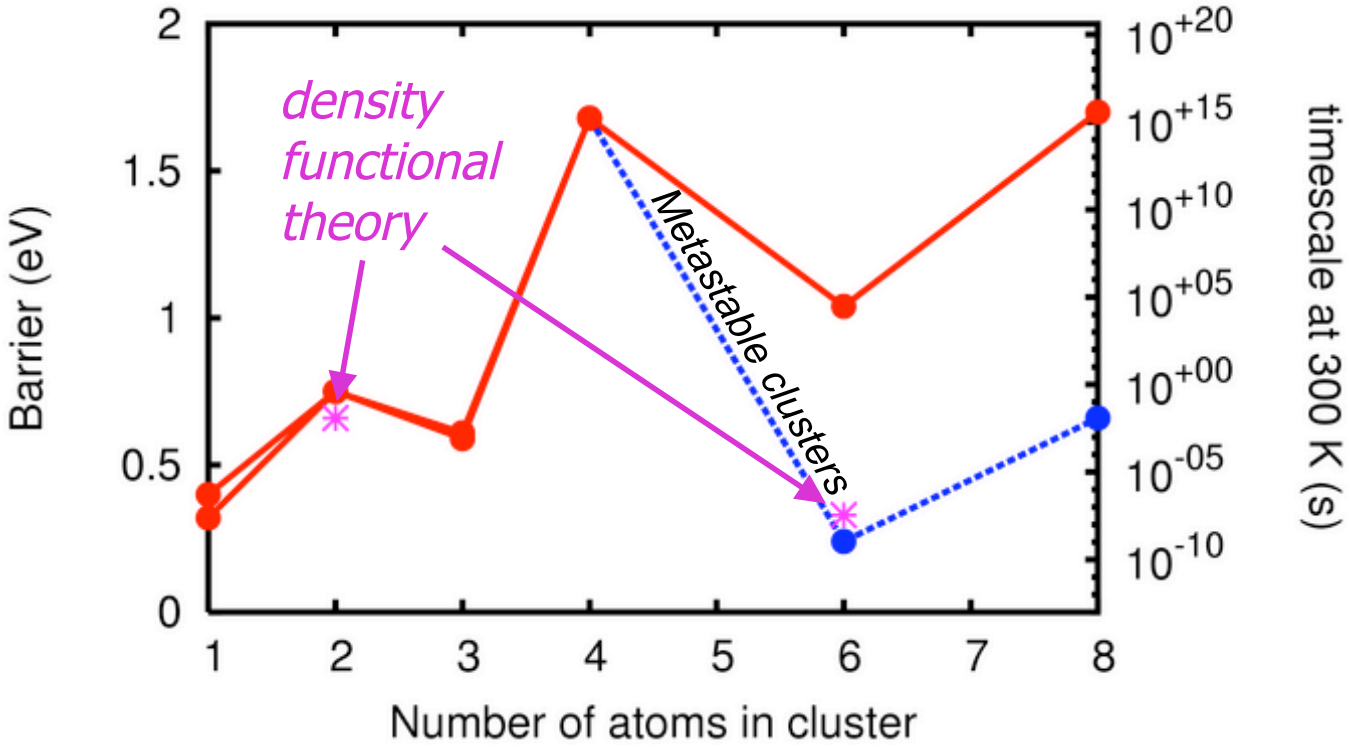

- Diffusion barrier follows no clear pattern versus cluster size
- Metastable clusters can be surprisingly mobile
- Metastable clusters can be very stable:
  - metastable I_6 lifetime is years at room temperature

Issues for future

- Boost limited by lowest barrier - general problem for many realistic systems (e.g., interstitials in metals)
- Roughness in potentials lowers effectiveness/efficiency of saddle searches (e.g., needed in TAD) and transition detection
  - shallow local minima
  - cusps at cutoffs and bond-switching points
- Currently limited to small systems
  - Trying to achieve N scaling
  - Spatial parallelization (with Jacques Amar at U. Toledo)
- Feeding information about atomistic behavior to 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
- Significant speedup over standard MD when barriers are high relative to temperature
- Often encounter unexpected behavior
- Ongoing challenges
  - low barriers and pesky local minima
  - cuspy potentials
  - scaling with system size

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