Accelerated Molecular Dynamics Methods

Arthur F. Voter
Theoretical Division
Los Alamos National Laboratory

Monte Carlo Methods in the Physical and Biological Sciences
Institute for Computational and Experimental Research in Mathematics (ICERM)
Brown University
Providence, RI
October 29, 2012

LA-UR-12-26438

Los Alamos

Acknowledgments

Many people have contributed greatly

and in particular:
Blas Uberuaga (LANL)
Danny Perez (LANL)

DOE Office of Basic Energy Sciences
Los Alamos LDRD
ASCR (DOE)
SCIDAC (DOE)

Note

This set of slides contains some material beyond what I had time to cover during my presentation.

The time-scale problem

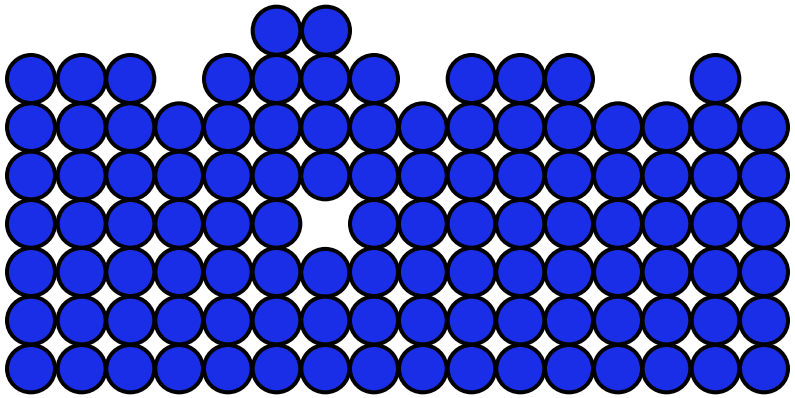

We have some system (e.g. atoms on a surface during growth).

Infrequent atomistic jumps move the system from state to state.

With molecular dynamics (MD), we can reach ~1 μs, but the interesting time scales are often much longer.

Individual transition events are sometimes complicated, involving many atoms, and the long-time evolution can be complex.

How do we accurately predict the long-time evolution?

Examples

Radiation damage annealing

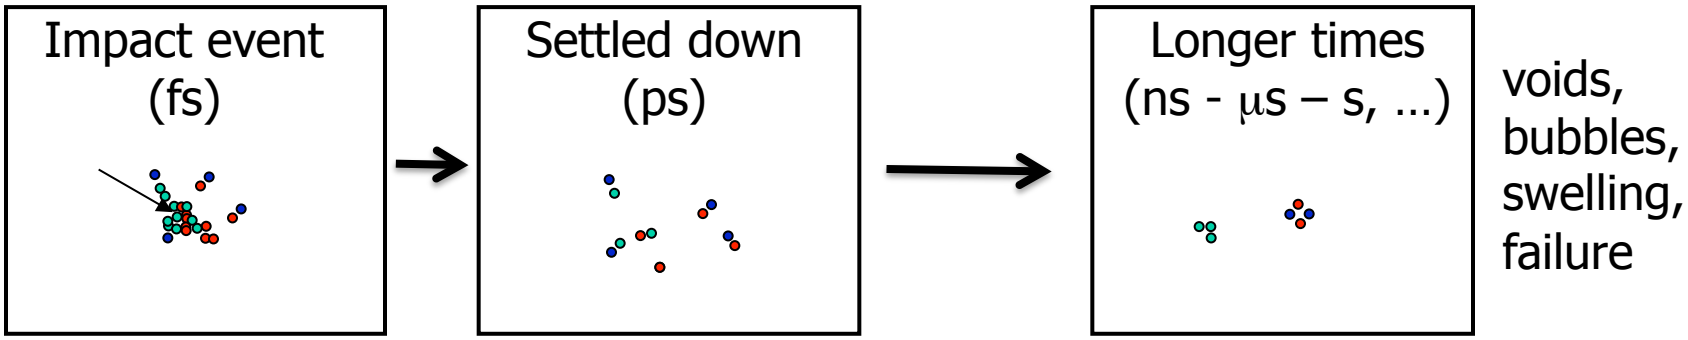

Vapor-deposited film growth (ms – s required)

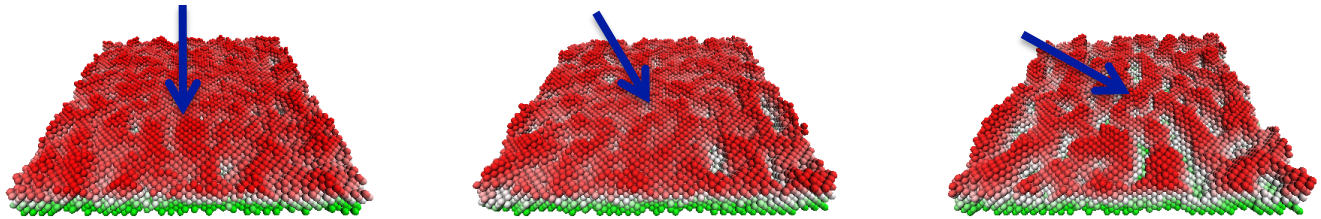

Evolution of a carbon nanotube fragment

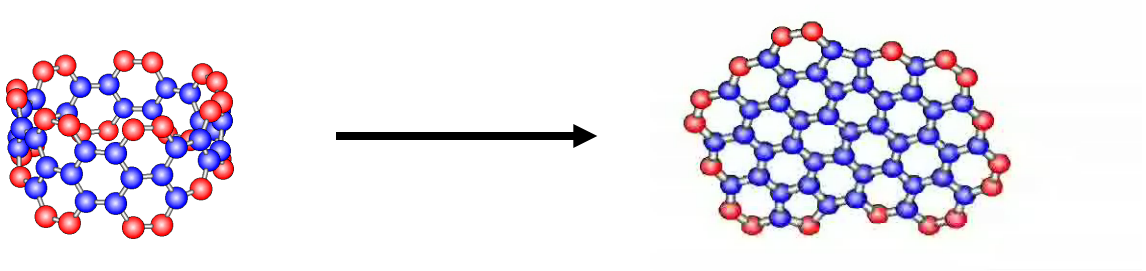

Infrequent Event System

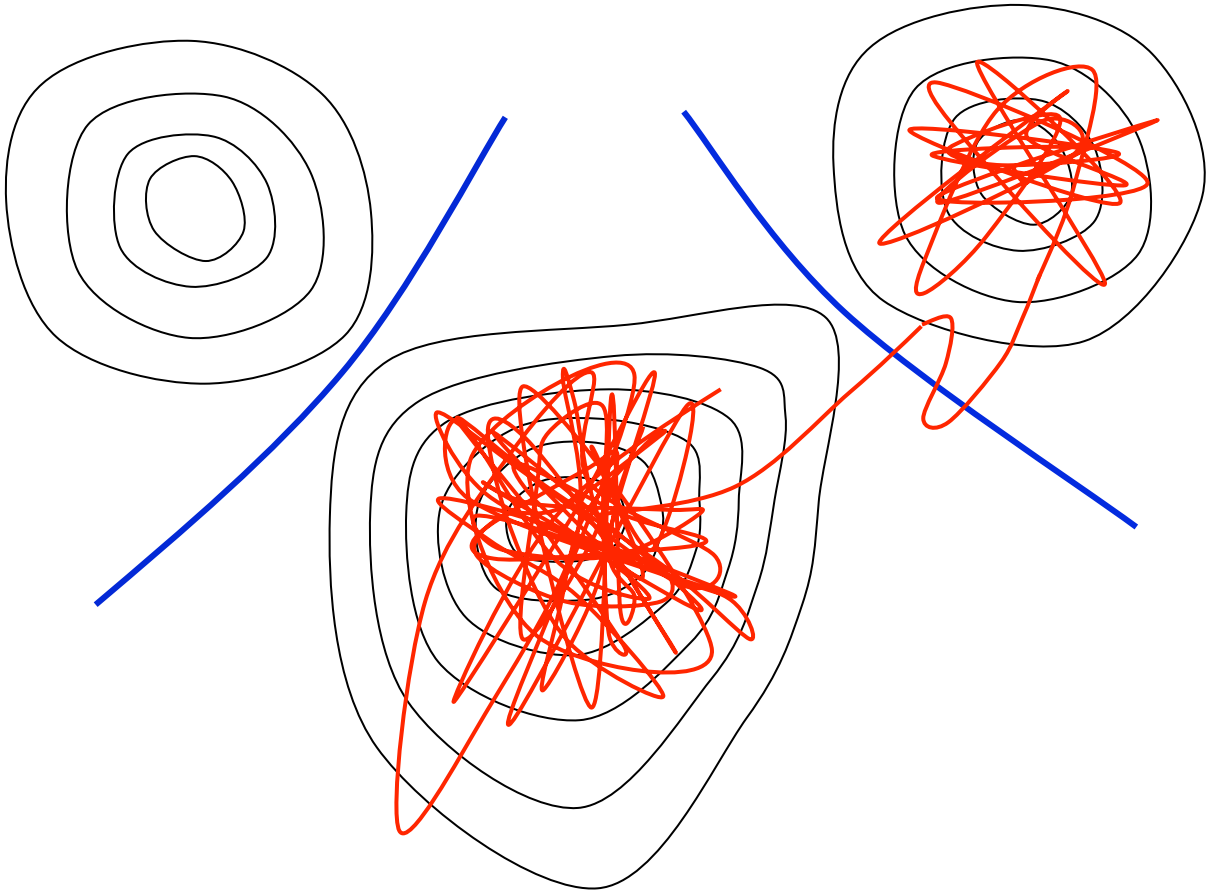

The system vibrates in 3N-dimensional basin many times before finding an escape path.

If we know the relevant pathway or pathways, we can use transition state theory to compute rates

Transition State Theory (TST)

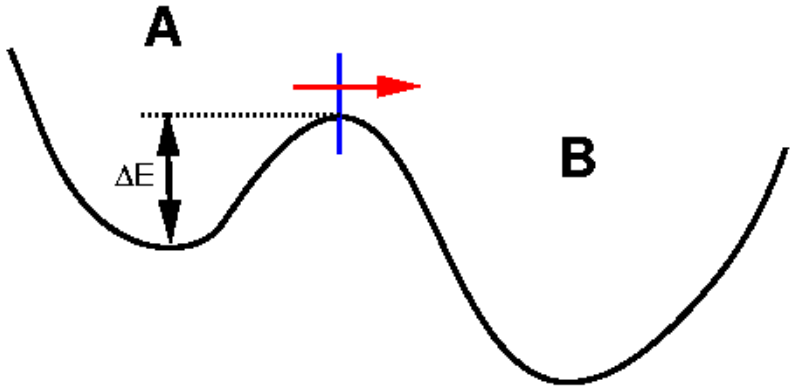

Marcelin (1915)
Eyring, Wigner,...

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

$$k_{A\to B}^{TST} = \left\langle \delta(x-q) \mid \dot{x} \mid \right\rangle \quad \text{(exact flux)}$$

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

- classically exact rate if no recrossings or correlated events
- no dynamics required to compute it
- very good approximation for activated events in materials

Cu/Cu(100) hop event, T=300K

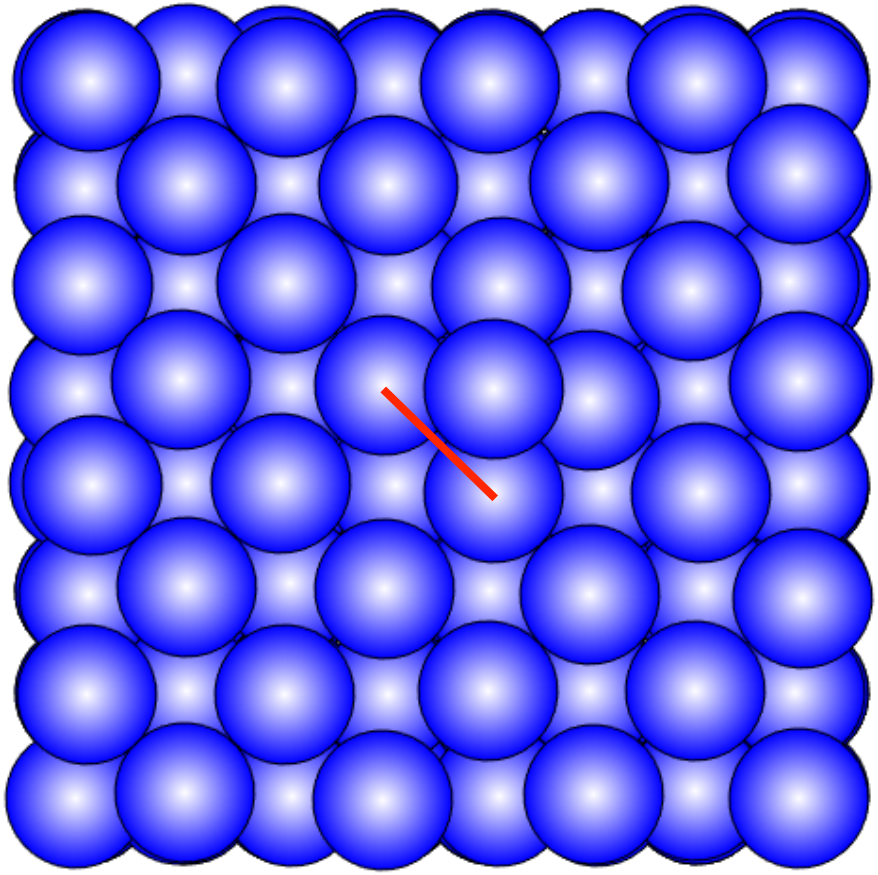

4 ps shown during transition event.
Rate at T=300K = once per 25 microseconds.

Cu/Cu(100) exchange event, T=300K

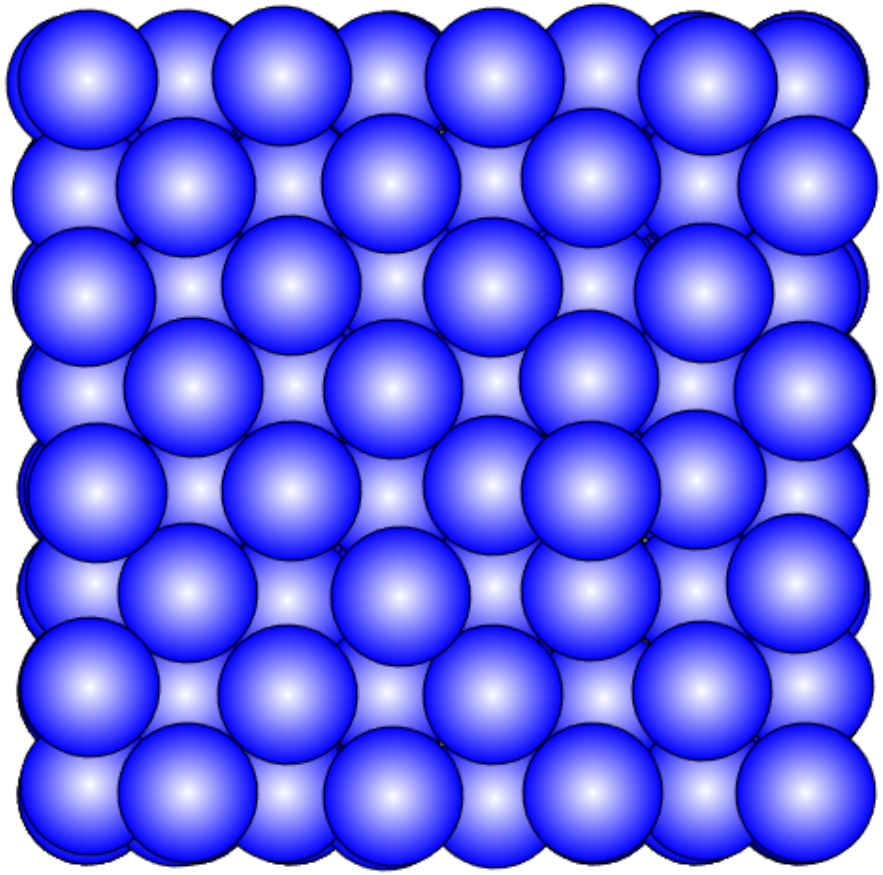

First seen by Feibelman, 1990.

4 ps shown during transition event.
Rate at T=300K = once per 14 seconds.
For Pt/Pt(100), exchange barrier is ~0.5 eV *lower* than hop barrier (10^8 x faster at 300K).

Accelerated Molecular Dynamics Approach

(not the only way, but our focus in this talk)

Accelerated molecular dynamics approach

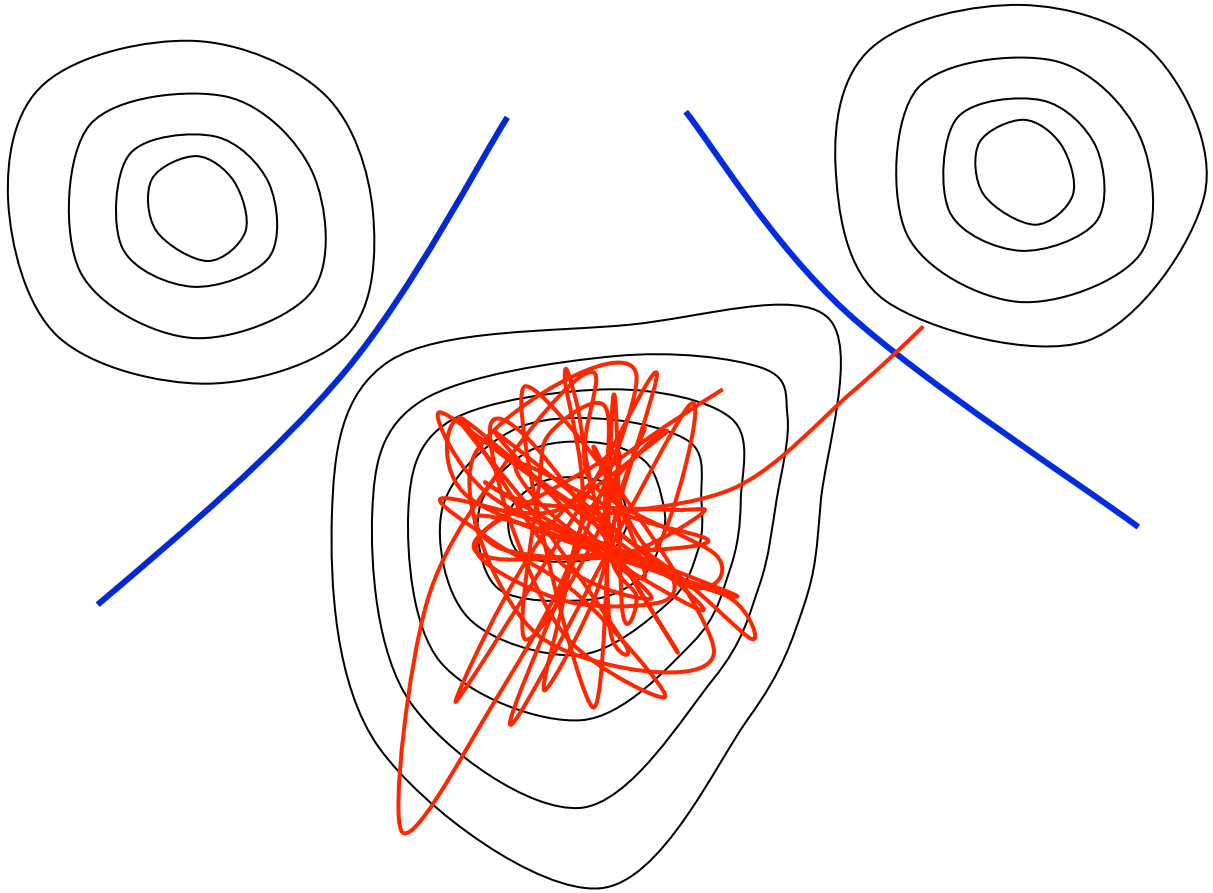

The system vibrates in 3N dimensional basin many times before finding an escape path. The trajectory finds an appropriate way out (i.e., proportional to the rate constant) without knowing about any of the escape paths except the one it first sees. Can we exploit this?

Accelerated molecular dynamics 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 accelerated molecular dynamics methods, we can follow a system from state to state, reaching time scales that we can't achieve with molecular dynamics.

As with regular MD, we can go back through the trajectory to determine rates and other properties in more detail, using conventional methods, and/or we can run more long trajectories to gather statistics.

Often, a single long trajectory can reveal some key behavior of the system, and often this behavior surprises us.

Accelerated Molecular Dynamics Methods

Hyperdynamics

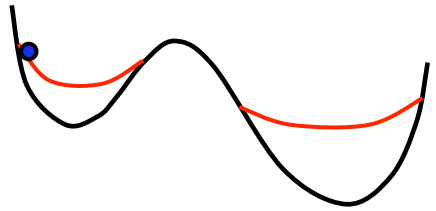

Parallel Replica Dynamics

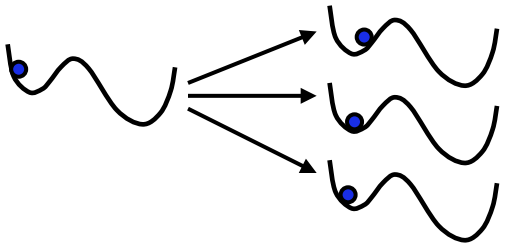

Temperature Accelerated Dynamics

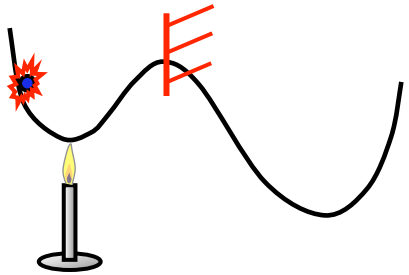

Accelerated Molecular Dynamics Methods

Hyperdynamics

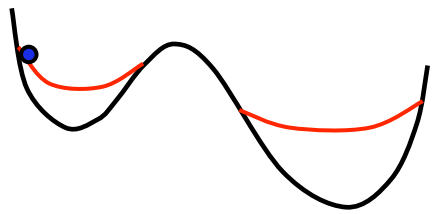

- Design bias potential that fills basins.
- MD on biased surface evolves correctly from state to state.
- Accelerated time is statistical quantity.

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

Parallel Replica Dynamics

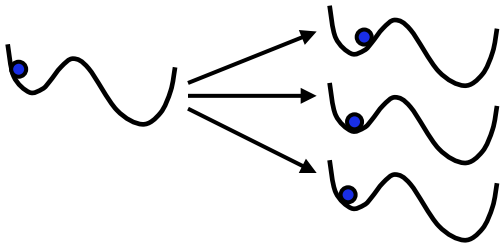

Temperature Accelerated Dynamics

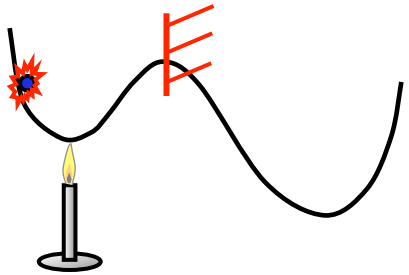

Accelerated Molecular Dynamics Methods

Hyperdynamics

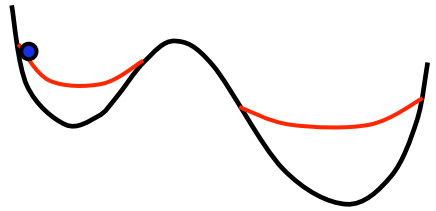

- Design bias potential that fills basins.
- MD on biased surface evolves correctly from state to state.
- Accelerated time is statistical quantity.

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

Parallel Replica Dynamics

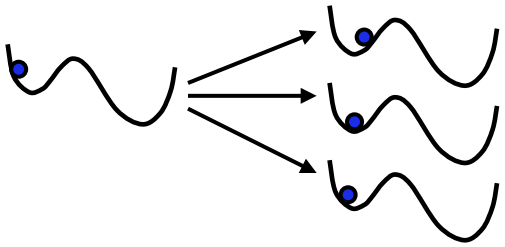

- Parallelizes time.
- Very general -- any exponential process.
- Gives exact dynamics if careful.
- Boost requires multiple processors

(AFV, Phys. Rev. B, 1998)

Temperature Accelerated Dynamics

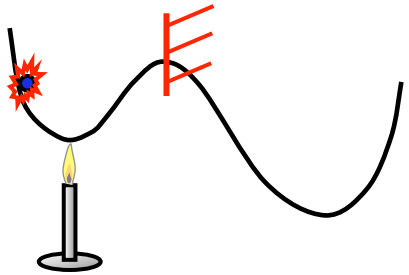

Accelerated Molecular Dynamics Methods

Hyperdynamics

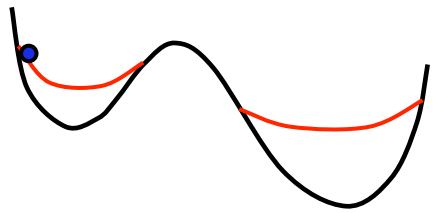

- Design bias potential that fills basins.
- MD on biased surface evolves correctly from state to state.
- Accelerated time is statistical quantity.

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

Parallel Replica Dynamics

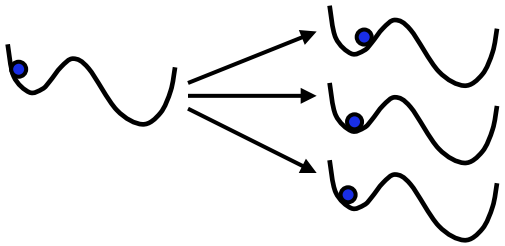

- Parallelizes time.
- Very general -- any exponential process.
- Gives exact dynamics if careful.
- Boost requires multiple processors

(AFV, Phys. Rev. B, 1998)

Temperature Accelerated Dynamics

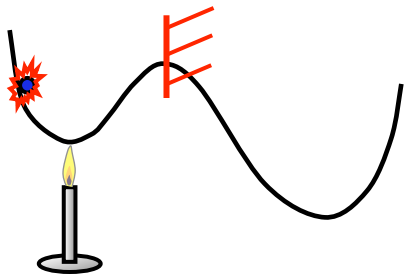

- Raise temperature of MD in this basin.
- Intercept and block every attempted escape.
- Accept event that would have occurred first at the low temperature.
- More approximate; good boost.

(M.R. Sorensen and AFV, J. Chem. Phys., 2000)

Hyperdynamics

LA-UR-12-26438

Hyperdynamics

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

Assumptions:

- infrequent events
- transition state theory (no recrossings)

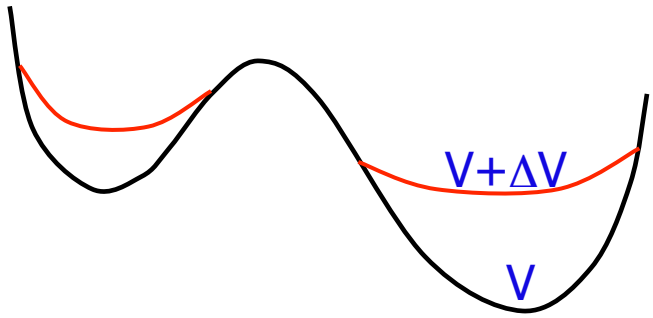

Procedure:

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

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

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

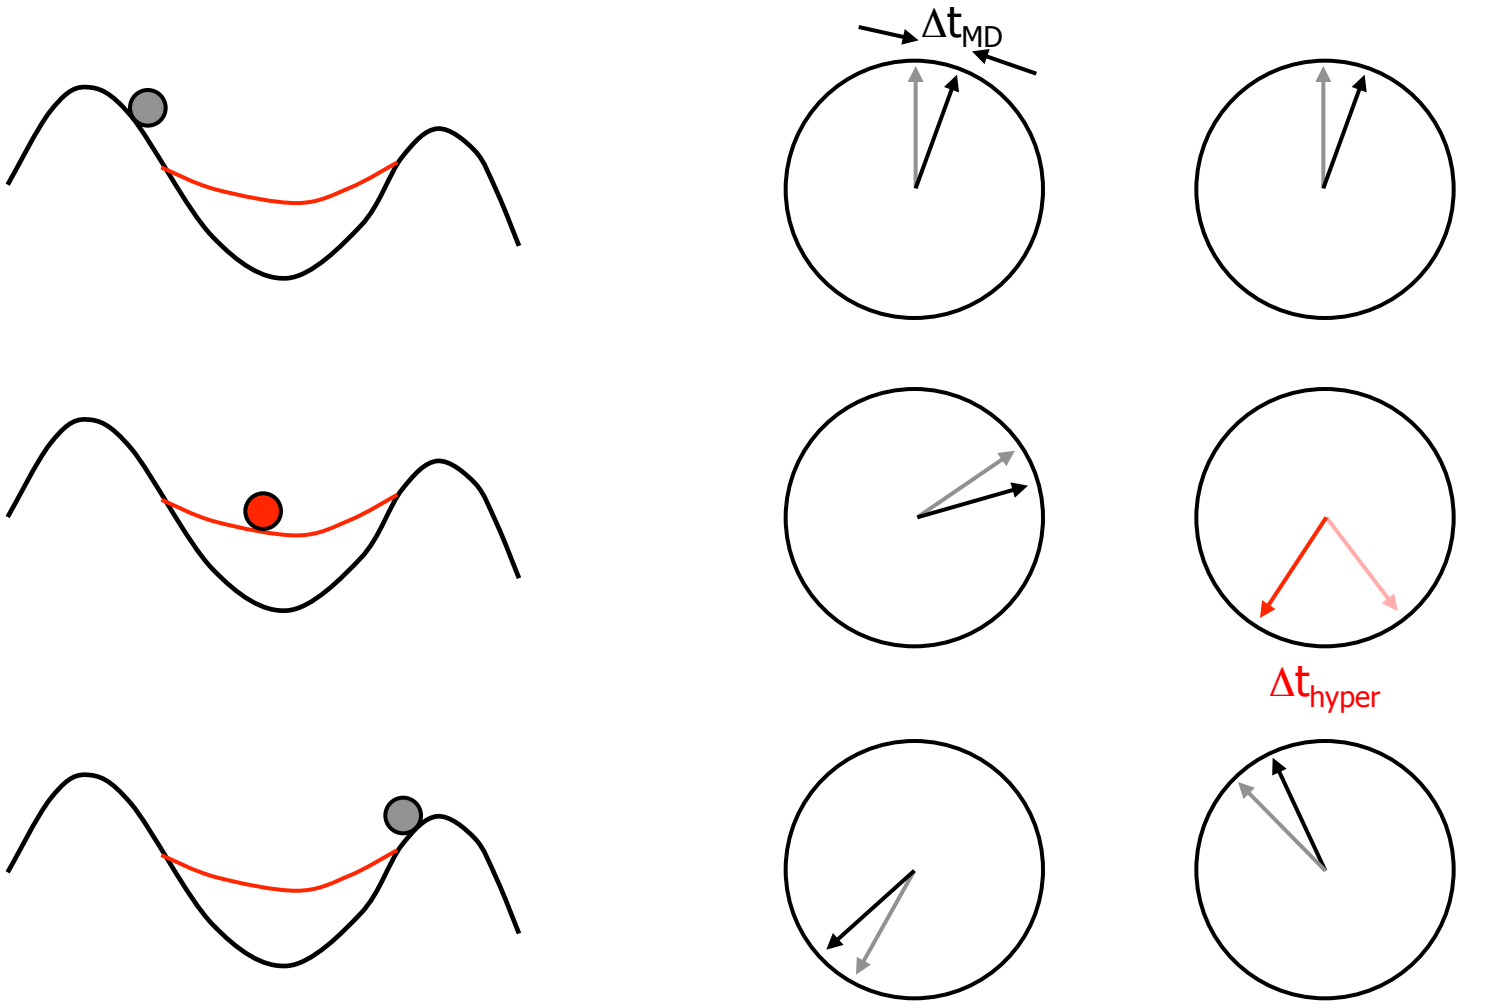

Boost = hypertime/(MD clock time)

The boost factor

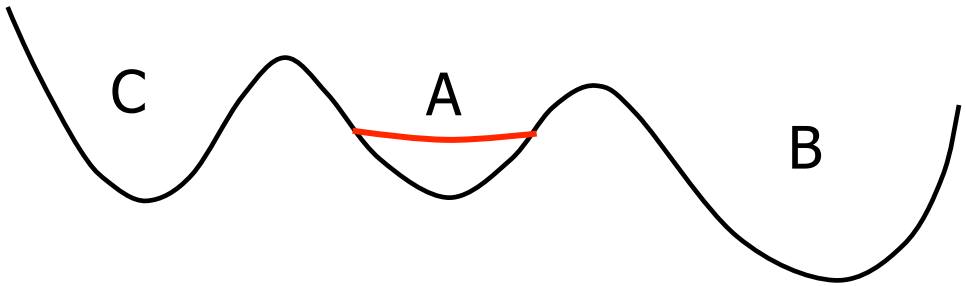

The boost factor (the hypertime over the MD time) is the average value of exp[+βΔV] on the biased potential:

$$\mathrm{Boost}_A = \left\langle e^{+\beta \Delta V(x)} \right\rangle_{Aboosted}$$

Hyperdynamics - characteristics

Designing valid and effective bias potential is the key challenge.

Bias potential can be a function of

- the shape of the energy surface (AFV, 1997)
- the energy (Steiner, Genilloud and Wilkins, 1998)
- the geometry
  - bond lengths, Miron and Fichthorn, 2003, 2005
  - local strain, Hara and Li, 2010

Must be careful that bias is zero on all dividing surfaces or dynamics will be wrong.

When barriers are high relative to T, boost can be many orders of magnitude.

Hyperdynamics bias potential

An extremely simple form: flat bias potential

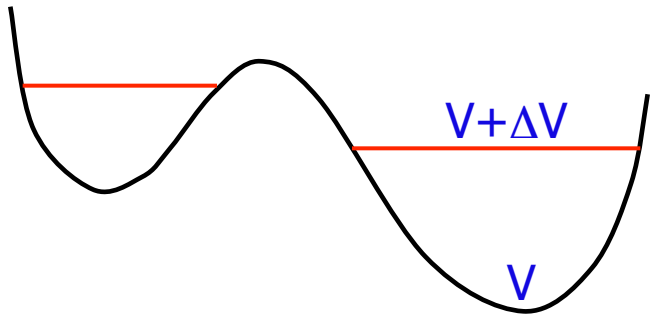

Steiner, Genilloud, and Wilkins, Phys. Rev. B **57**, 10236 (1998).

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

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

Bias potential is turned on near the minimum in the potential basin, but turns off when any bond is stretched beyond a threshold value

Very appealing approach:

- fairly general
- very low overhead
- purely geometric - behaves better than earlier bias potentials based on slope and curvature of potential

Miron and Fichthorn (JCP 2005) have used this effectively to study Co/Cu(001) film growth

Co/Cu(001) growth using bond-boost hyperdynamics

Miron and Fichthorn Phys. Rev. B 72, 035415 (2005)

Simulation of growth at 1 ML/s

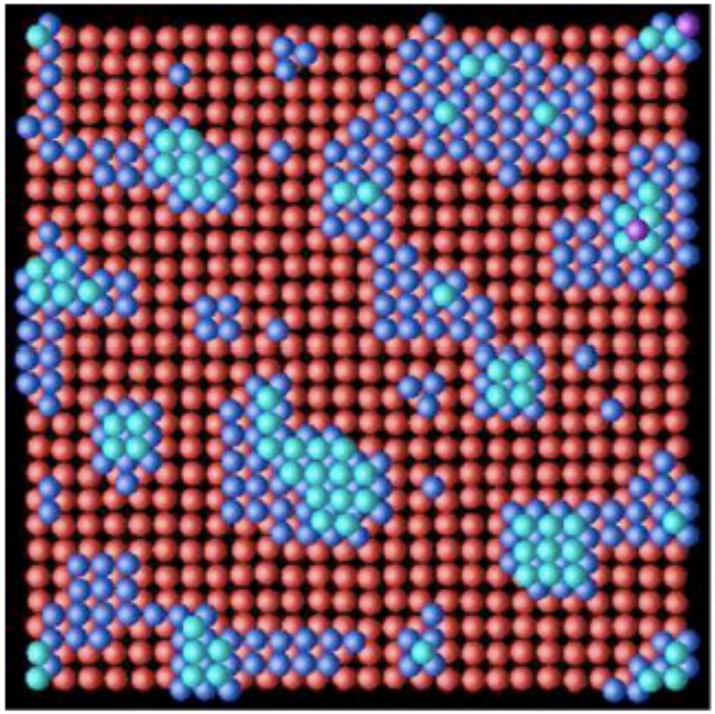

T=250K

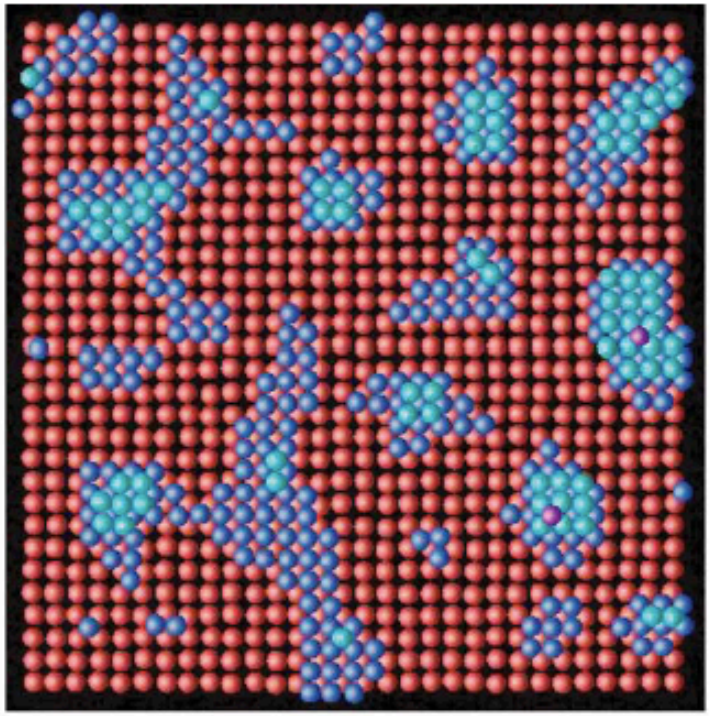

T=310K

Summary - Hyperdynamics

Powerful if an effective bias potential can be constructed

Need not detect transitions (though we sometimes do as part of the bias potential construction)

Boost factors climbs exponentially with inverse temperature

Especially effective if barriers high relative to T

Lots of possibilities for future development of advanced bias potential forms

Recently extended to large systems (S.Y. Kim, D. Perez, AFV, in preparation)

Limitations - Hyperdynamics

Must design bias potential

Assumes TST holds (though Langevin-noise recrossings may be OK)

Boost drops off when events are frequent (true of all the AMD methods)

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

Hyperdynamics on large systems

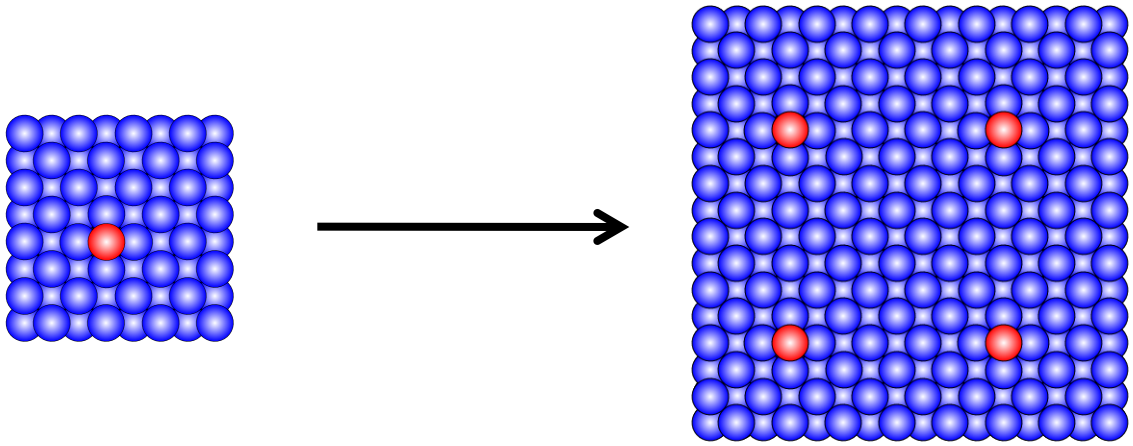

Whenever system is near a dividing surface, ΔV must be zero.

For a 4x larger system, the trajectory is near a dividing surface ~4x more often, causing a lower overall boost factor.

For very large systems, the boost decays to unity – i.e., there is no speedup, *no matter what form of bias potential is used*.

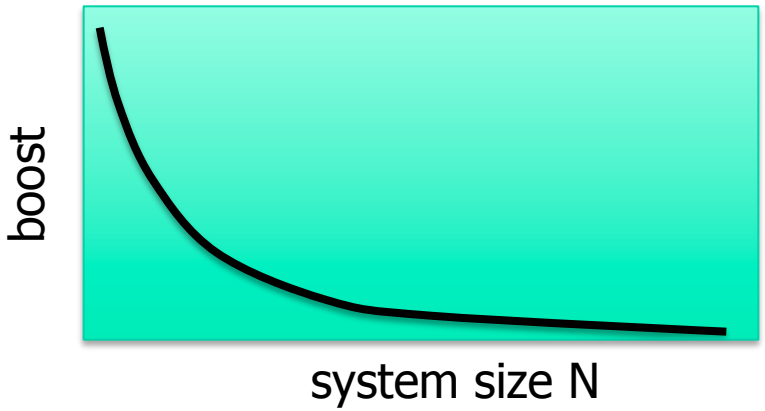

e.g.,
Miron and Fichthorn saw boost~$N^{-0.9}$
and Hara and Li saw boost~N^{-1}

Local Hyperdynamics

S.Y. Kim, D. Perez, and AFV (to be submitted).

Modified formulation of hyperdynamics that gives *constant boost for arbitrarily large systems.*

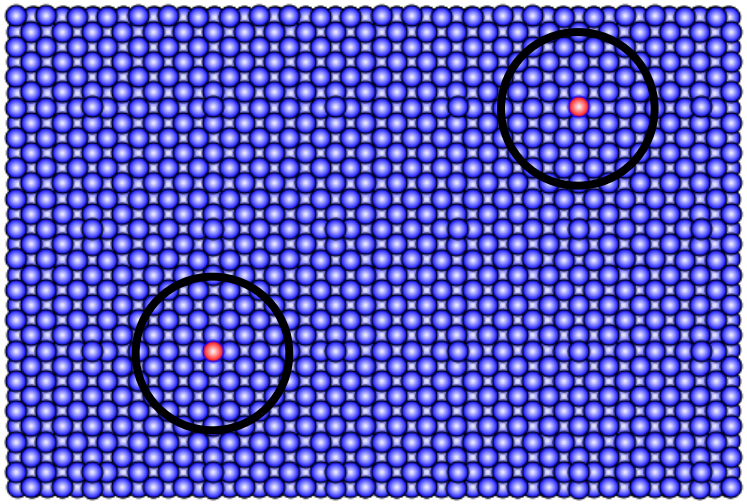

Key concept: Most systems we are interested in are intrinsically local in their behavior. A transition, or near-transition, in one region of system should not have any significant effect on atoms that are far away.

Local Hyperdynamics

- Define a *local* bias energy
- This local bias energy can be nonzero in one region even if there is a transition-causing distortion in another region, provided it is far away
- Differentiate the local bias to get local force (dynamics are no longer conservative)

The method is probably not exact, but we have some understanding of the error terms, and why they should largely cancel.

Tests on various systems show the method is very accurate.

We also see how to make it scale as N, with fixed boost factor, to arbitrarily large systems.

Local hyperdynamics tests for simple system

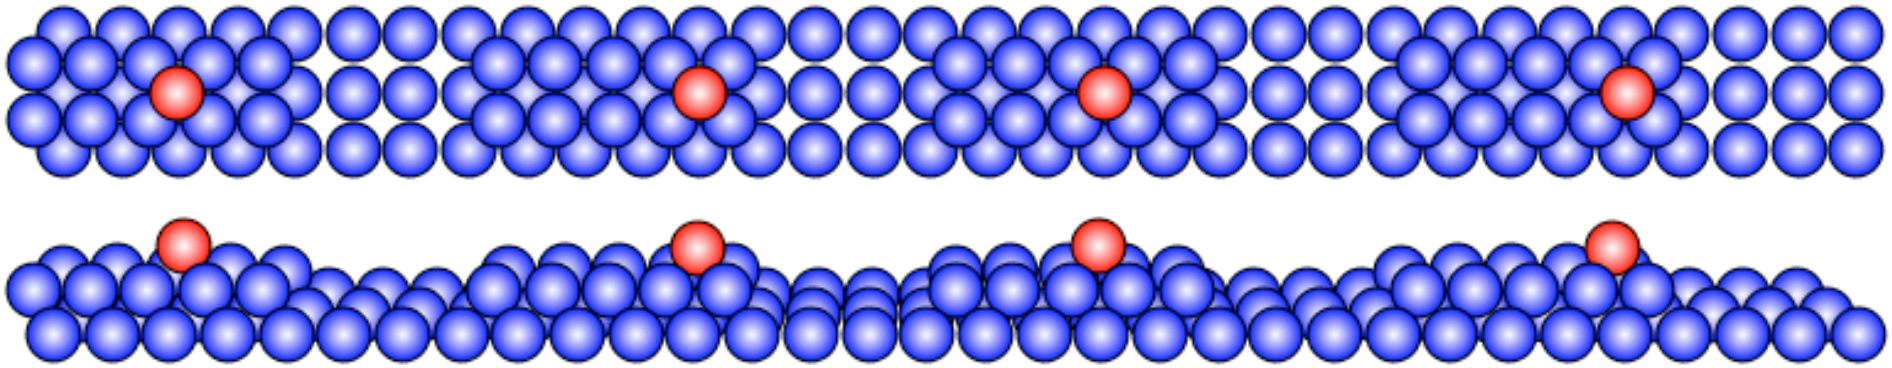

Tests inequivalent environments and multiple mechanisms

Main unique events (hops and exchanges):

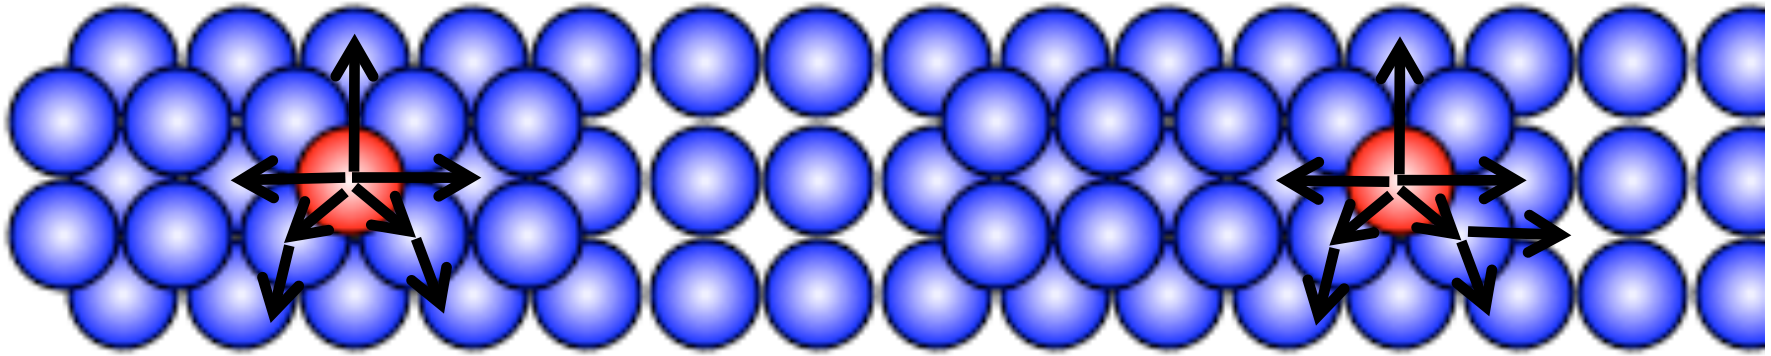

Local hyperdynamics tests for simple system

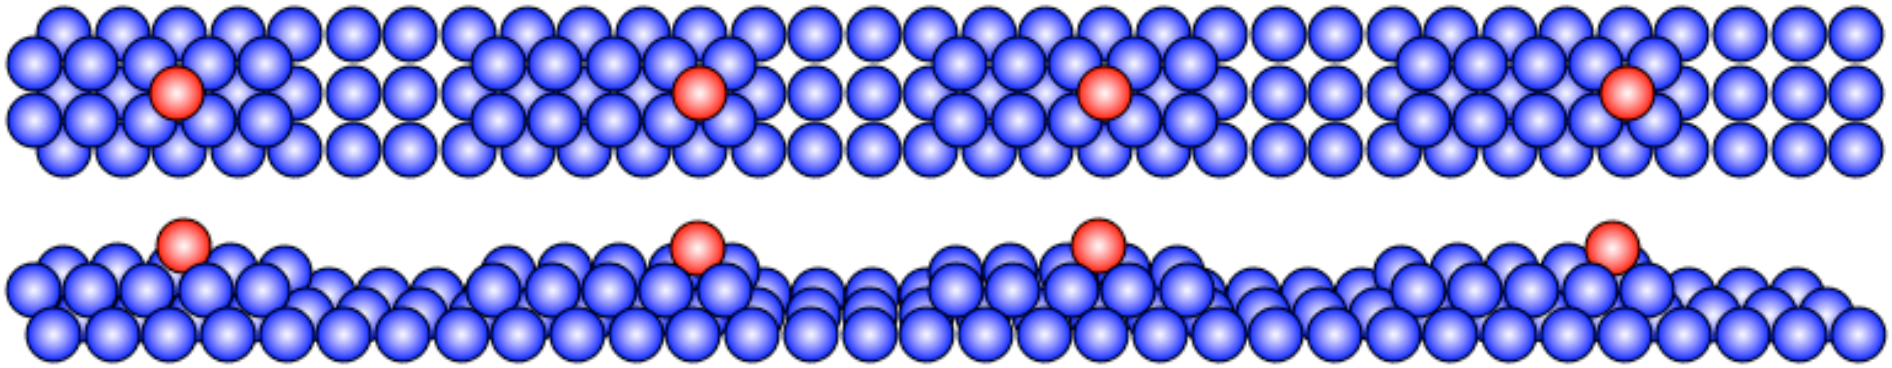

T=500K, B_{target}=100, α_B=5X10^9, range=10 Å, simple bond boost bias

Excellent agreement with direct MD

Boost=100

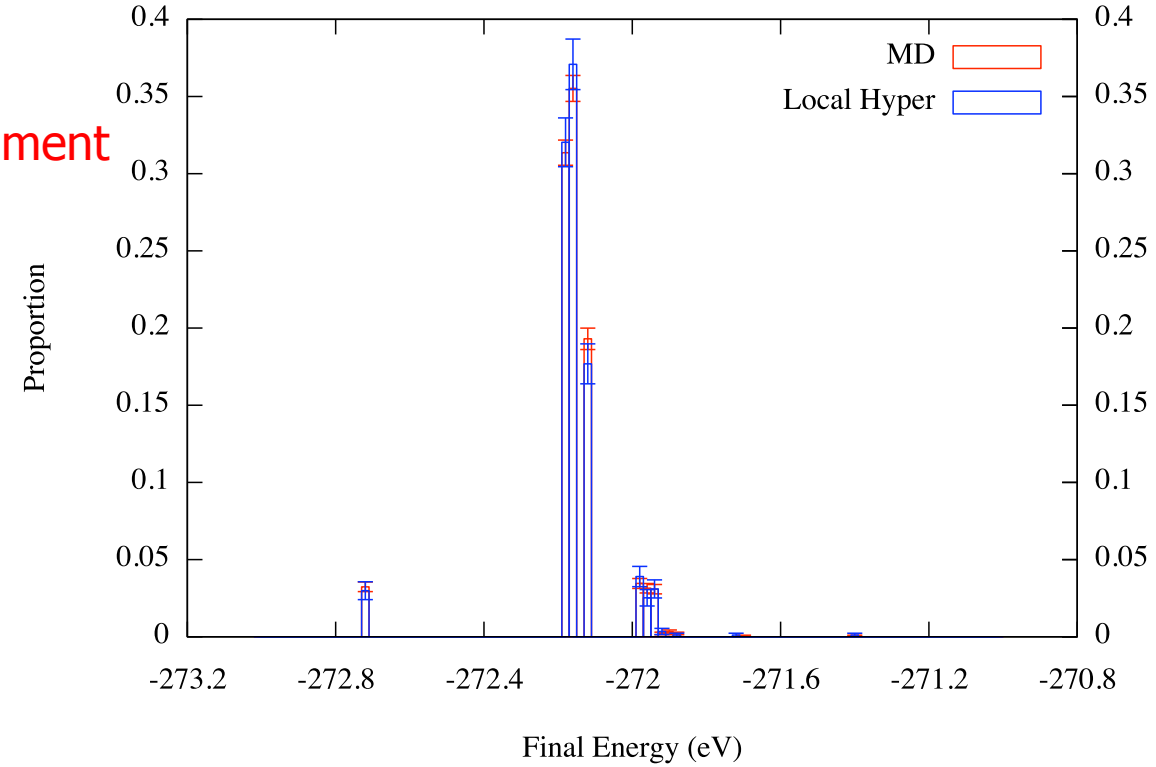

Parallel-Replica Dynamics

LA-UR-12-26438

Los Alamos

Parallel Replica Dynamics

Parallelizes time evolution

Assumptions:

- infrequent events
- exponential distribution of first-escape times

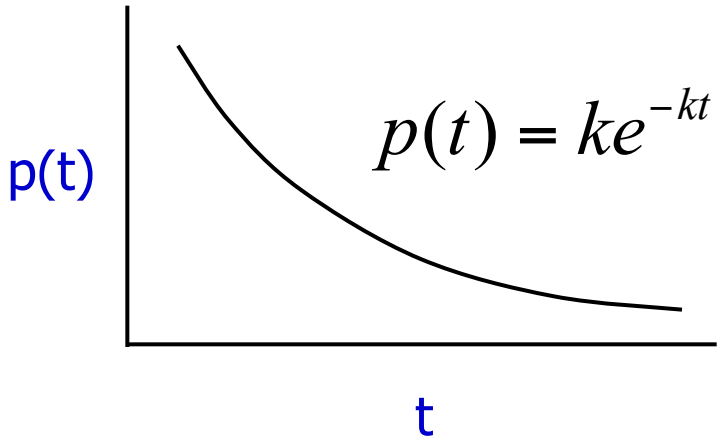

Must know:

- 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.

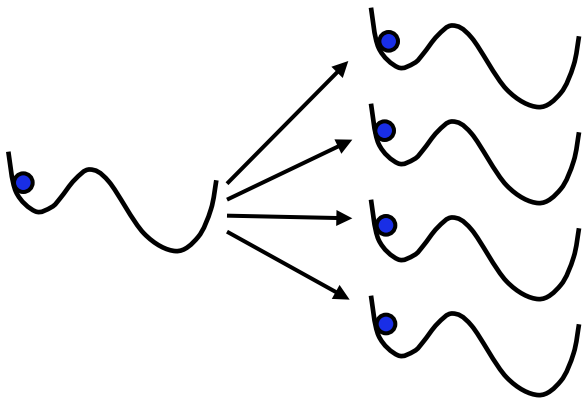

Parallel Replica Dynamics Procedure

Randomize momenta independently on each processor.

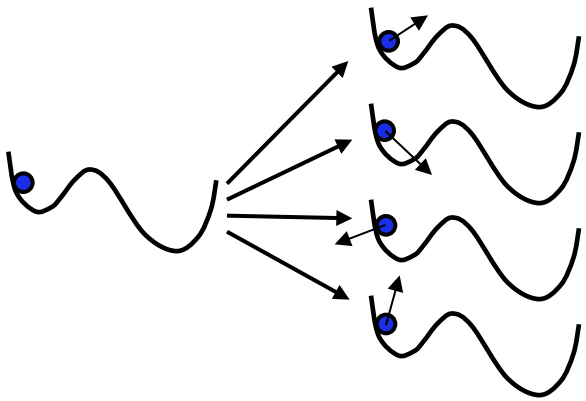

Parallel Replica Dynamics Procedure

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

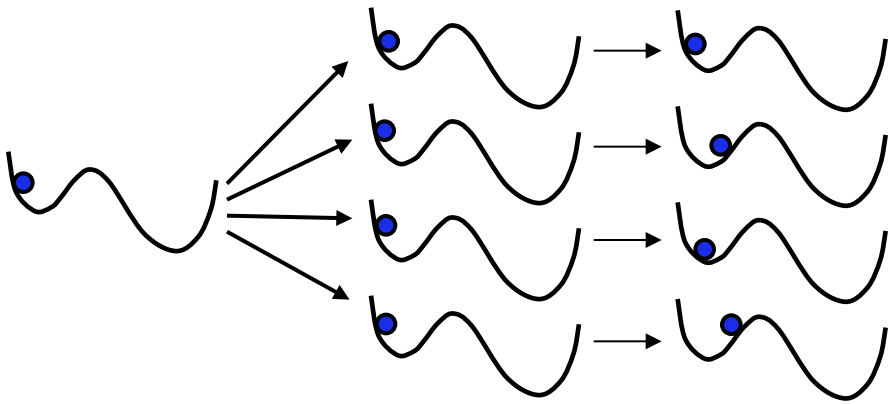

Parallel Replica Dynamics Procedure

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

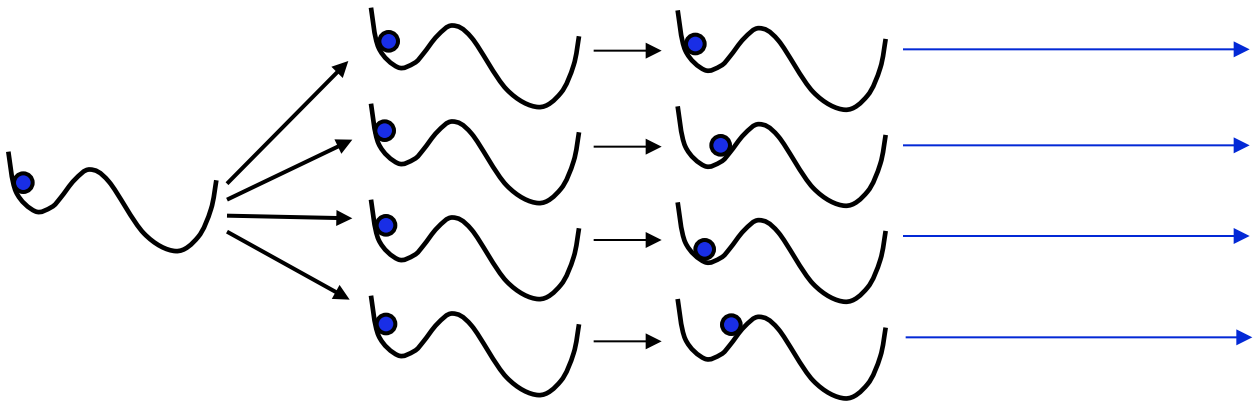

Parallel Replica Dynamics Procedure

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

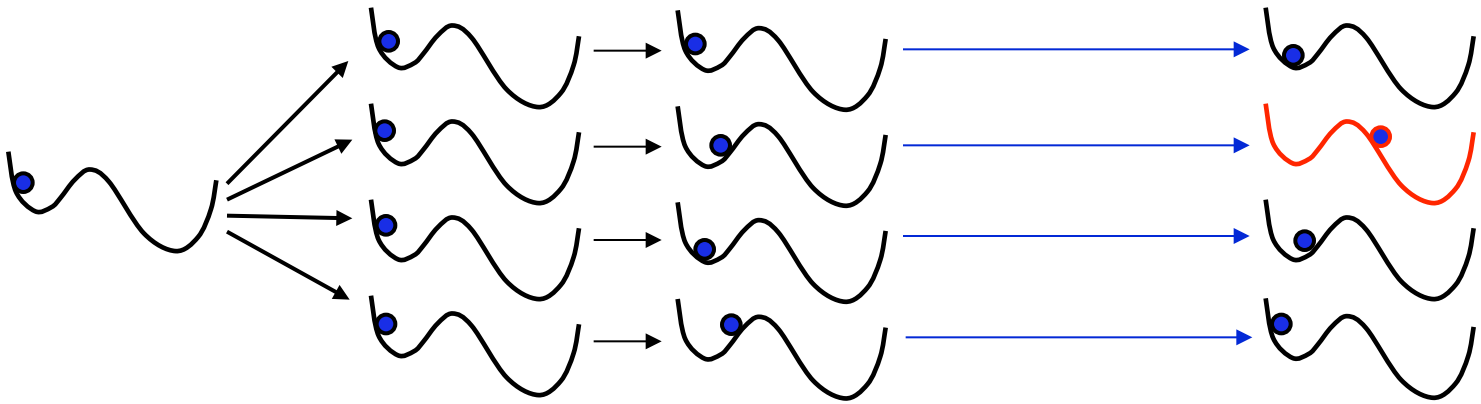

Parallel Replica Dynamics Procedure

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

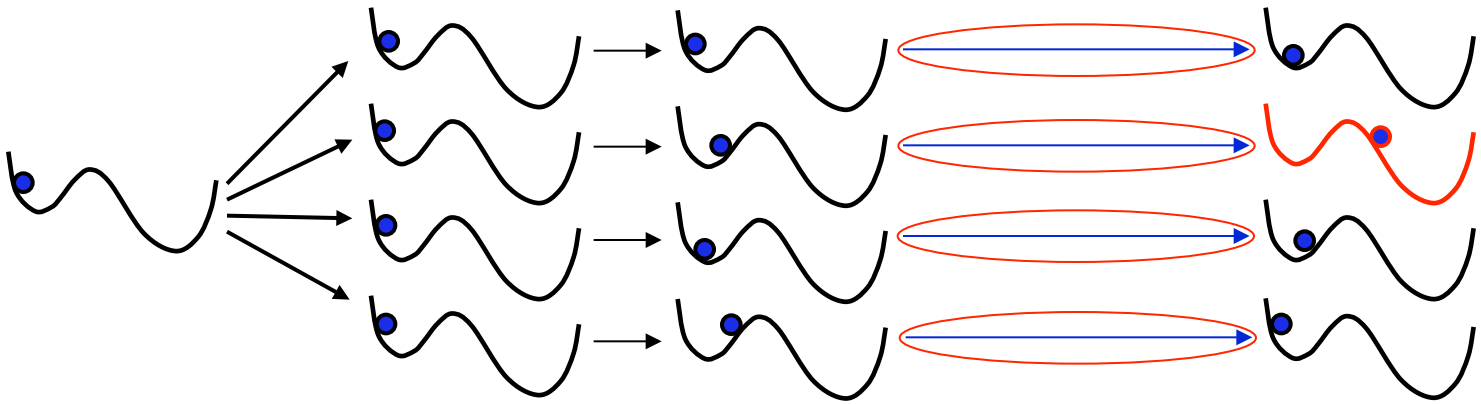

Parallel Replica Dynamics Procedure

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

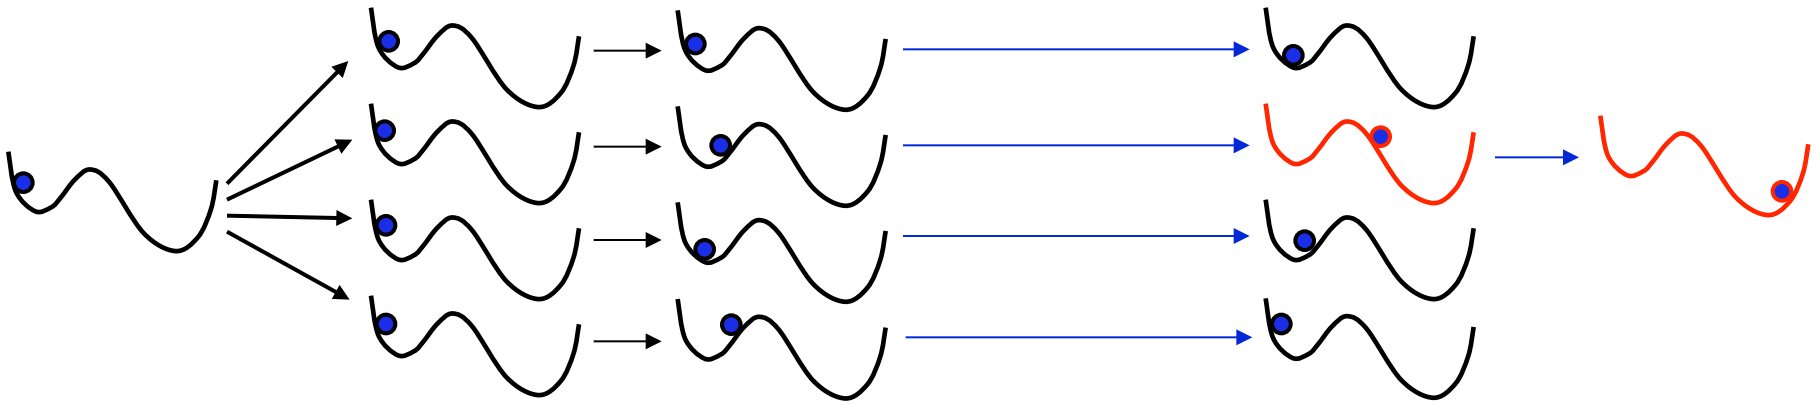

Parallel Replica Dynamics Procedure

Advance simulation clock by τ_{corr}.

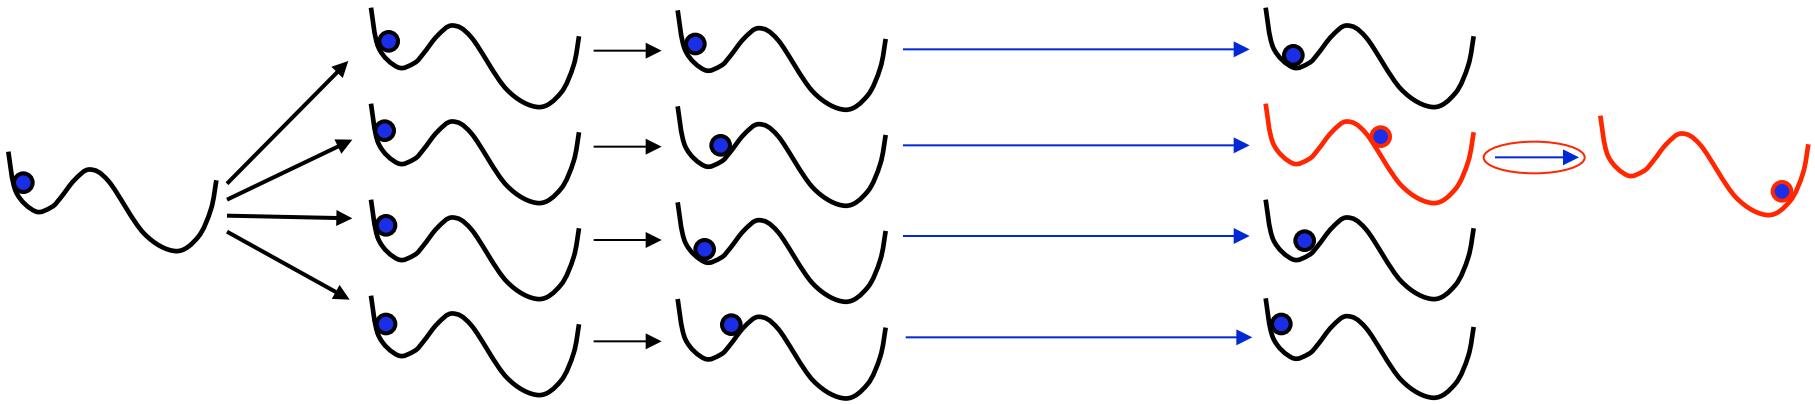

Parallel Replica Dynamics Procedure

Replicate the new state and begin procedure again.

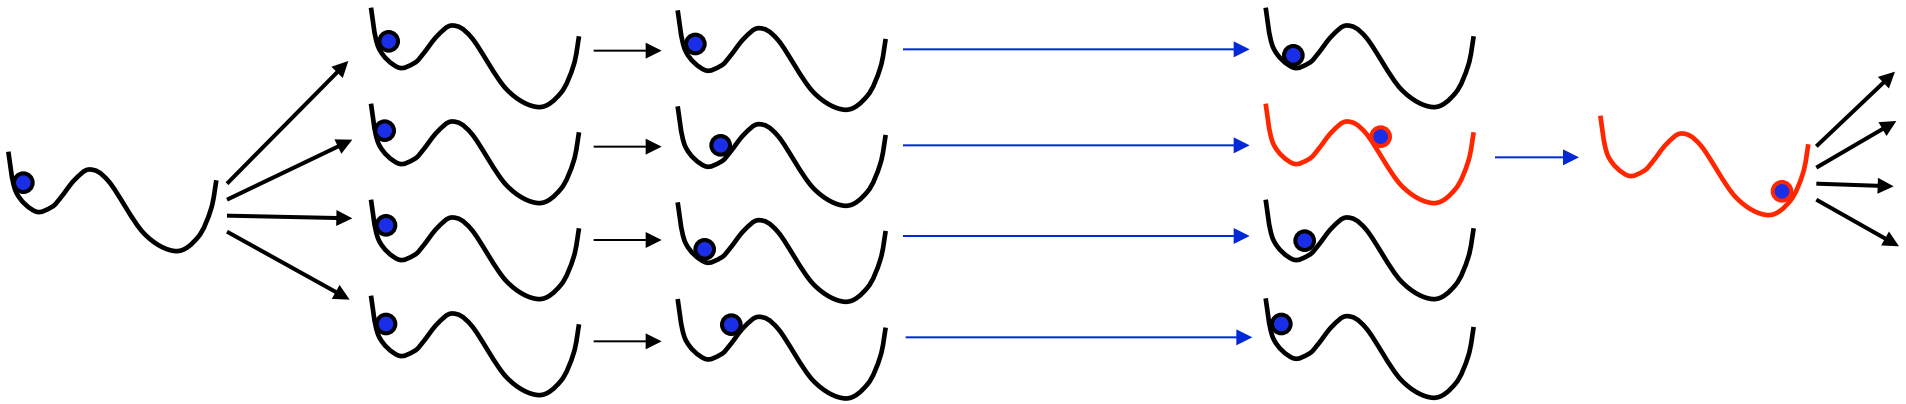

Parallel Replica Dynamics

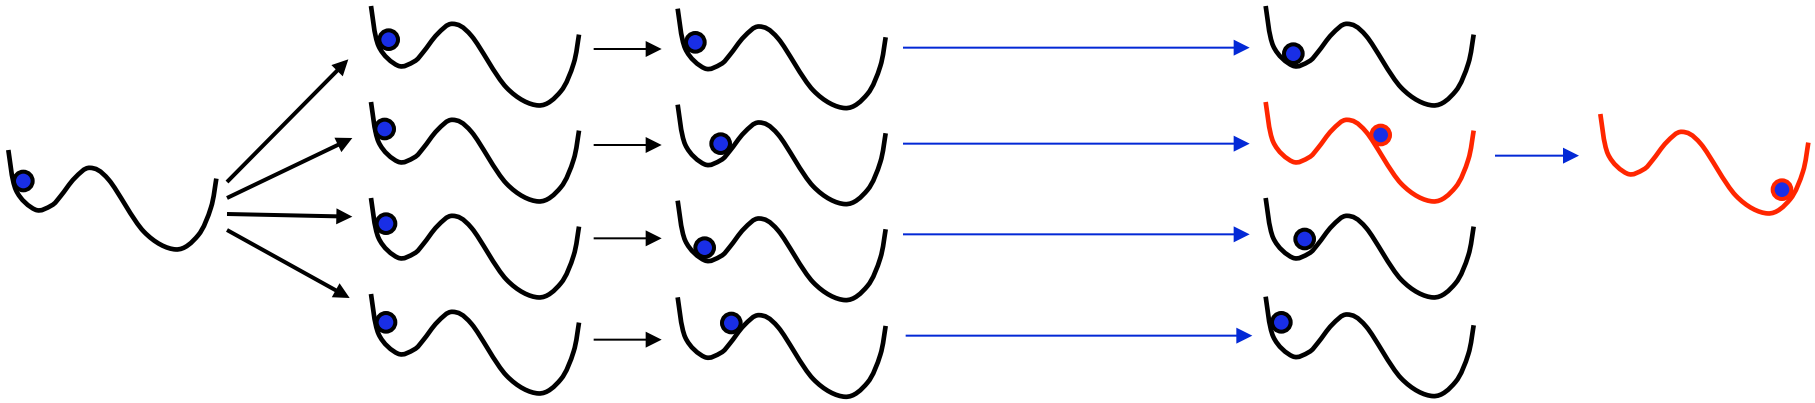

The summed time (t_{sum}) obeys the correct exponential distribution, and the system escapes to an appropriate state.

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

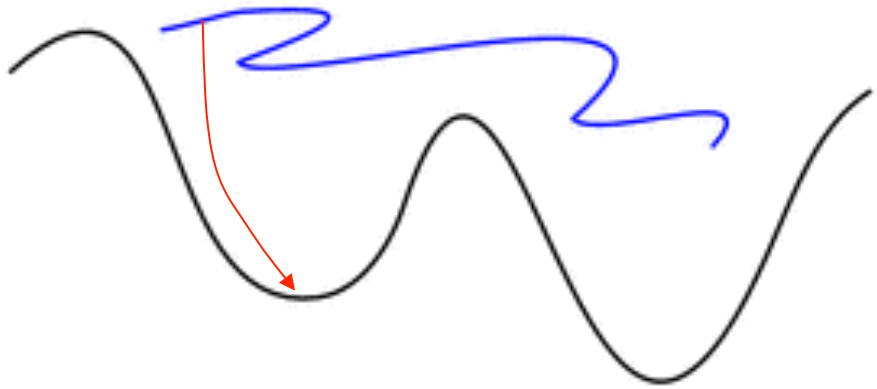

Detecting a transition

- 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

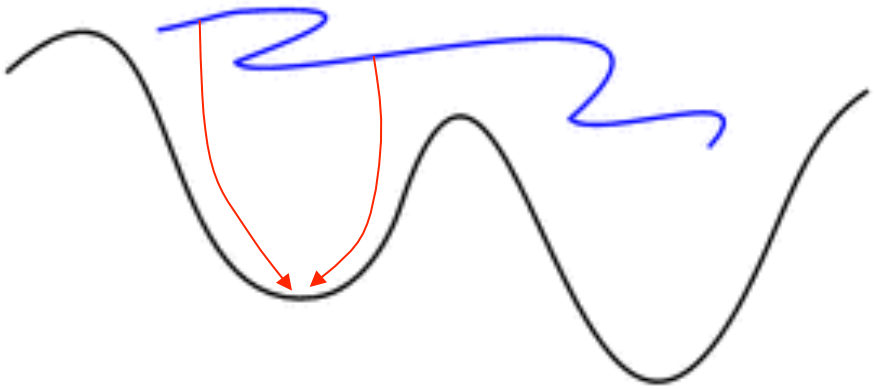

Detecting a transition

- 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

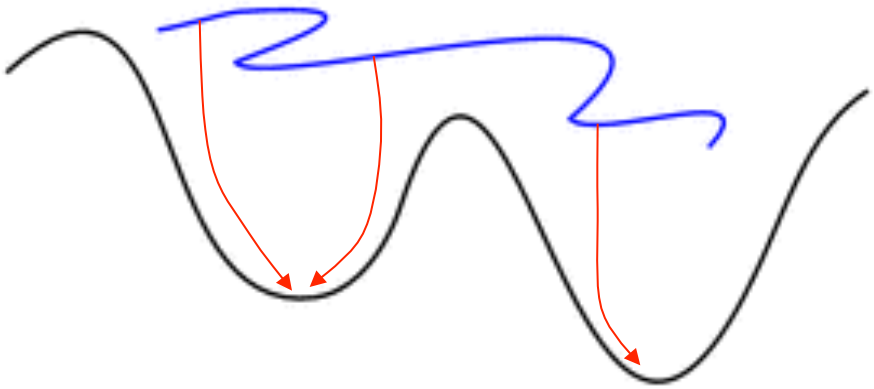

Examples of ParRep studies

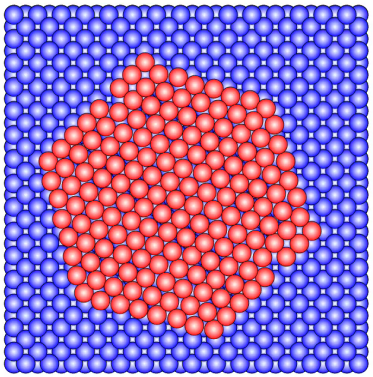

Ag_{169}/Cu(100), magic cluster, Uche et al, 2009.

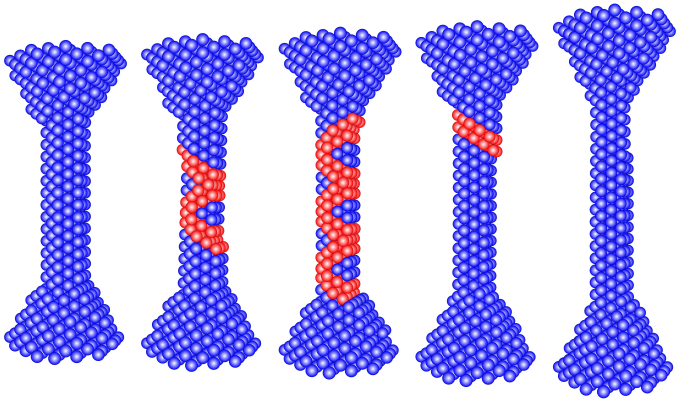

Ag nanowire stretch, μs - ms, Perez et al.

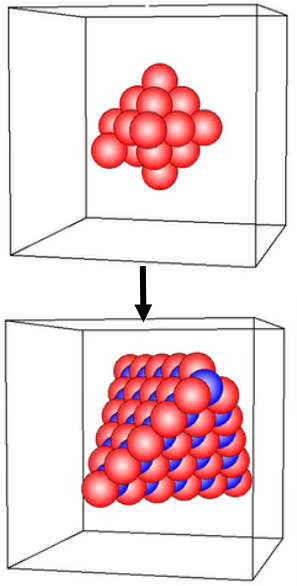

Cu void collapse to SFT, μs, Uberuaga et al, 2007.

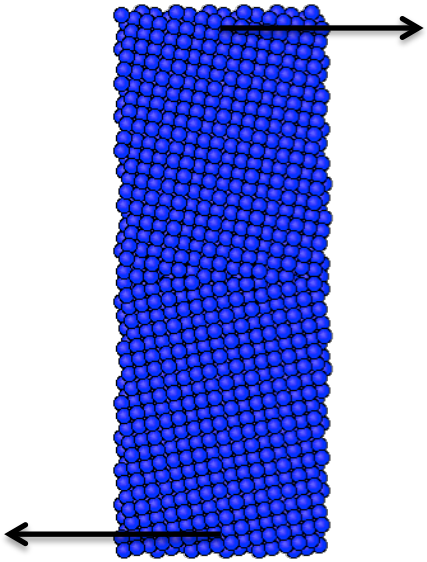

Driven Cu GB sliding, 500 μm/s Mishin et al, 2007.

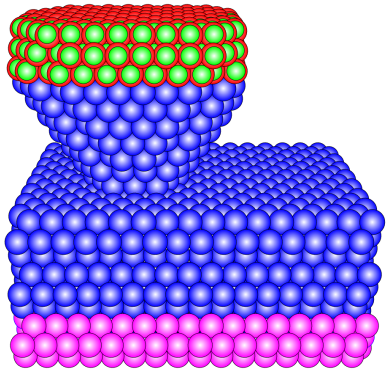

Friction force microscopy, Dong et al, 2009, 2010, 2011.

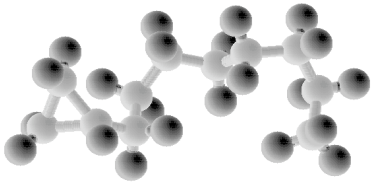

Hexadecane pyrolysis, μs, Kum et al, 2004.

Recent brief AMD review: Perez et al, Ann. Rep. Comp. Chem. **5**, 79 (2009)

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
- 79% efficiency on 39 processors
- At 1.69 μs, void transforms to SFT

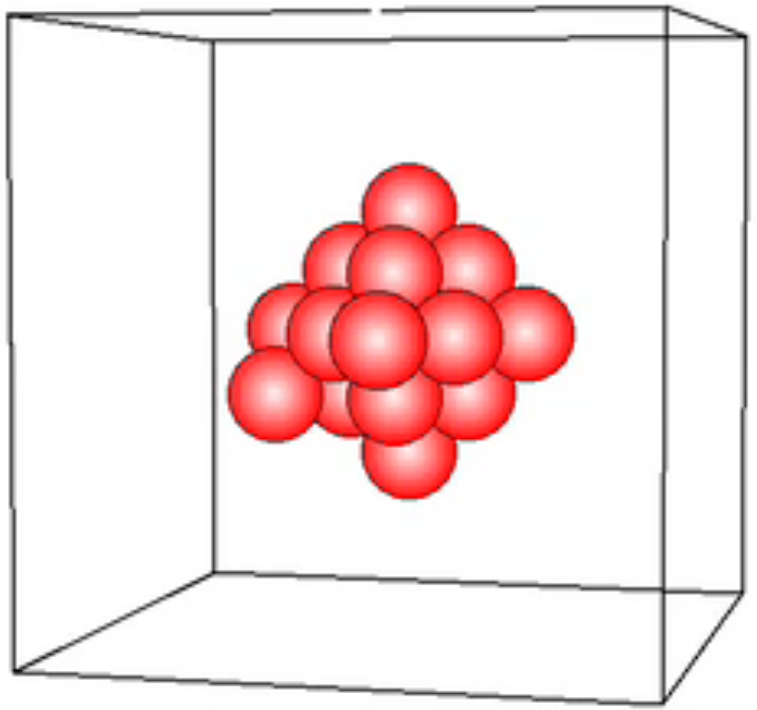

Red atoms=vacancies
Blue atoms=interstitials
Bulk atoms not shown

Uberuaga, Hoagland, Voter, Valone, PRL **99**, 135501 (2007)
LA-UR-12-26438

Ag nanowire using ParRep

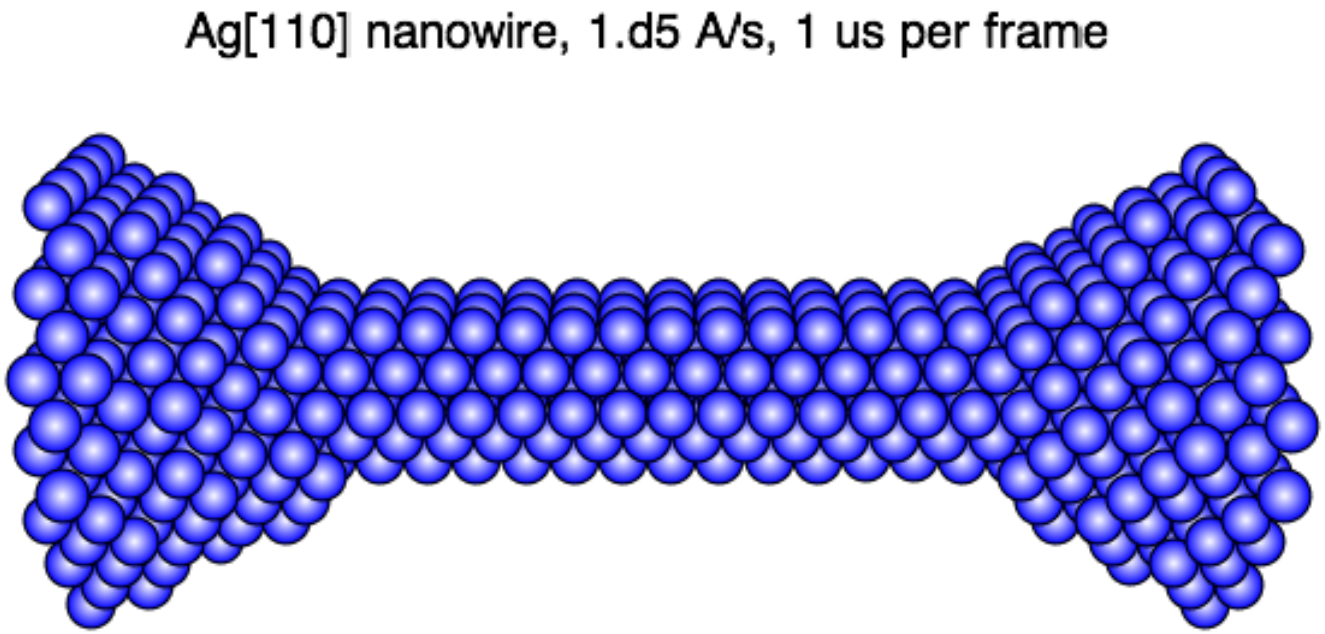

225 μs total; simulated on LANL Roadrunner computer

D. Perez, C.W. Pao, S. Swaminarayan, AFV (to be published)

Ag nanowire

It is possible that this icosahedral chain is what Rodriguez et. al. (PRL, 2002) saw in their Ag [110] nanowire experiments:

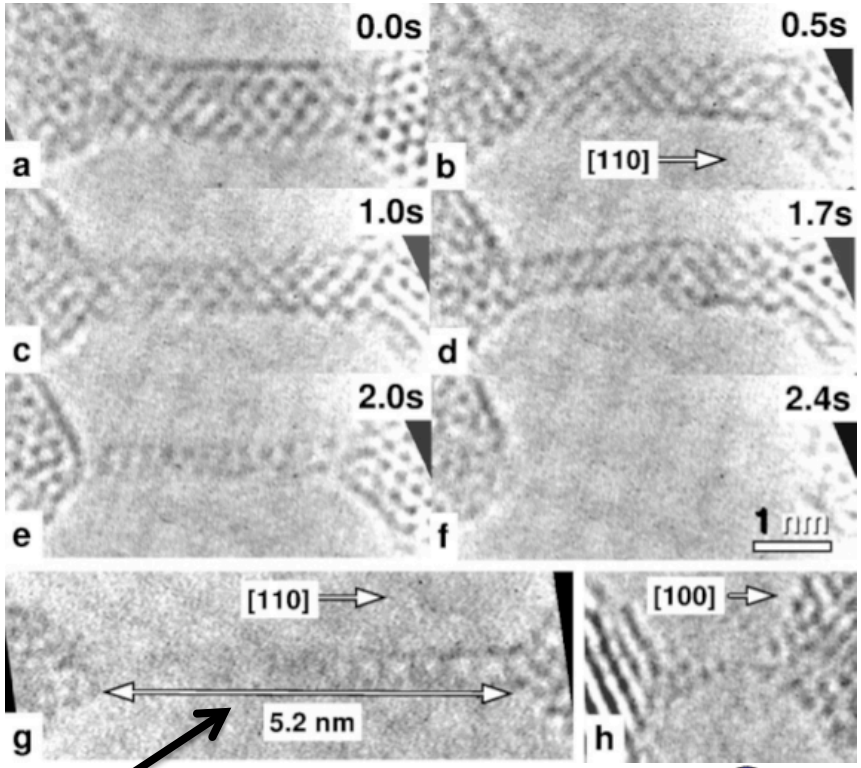

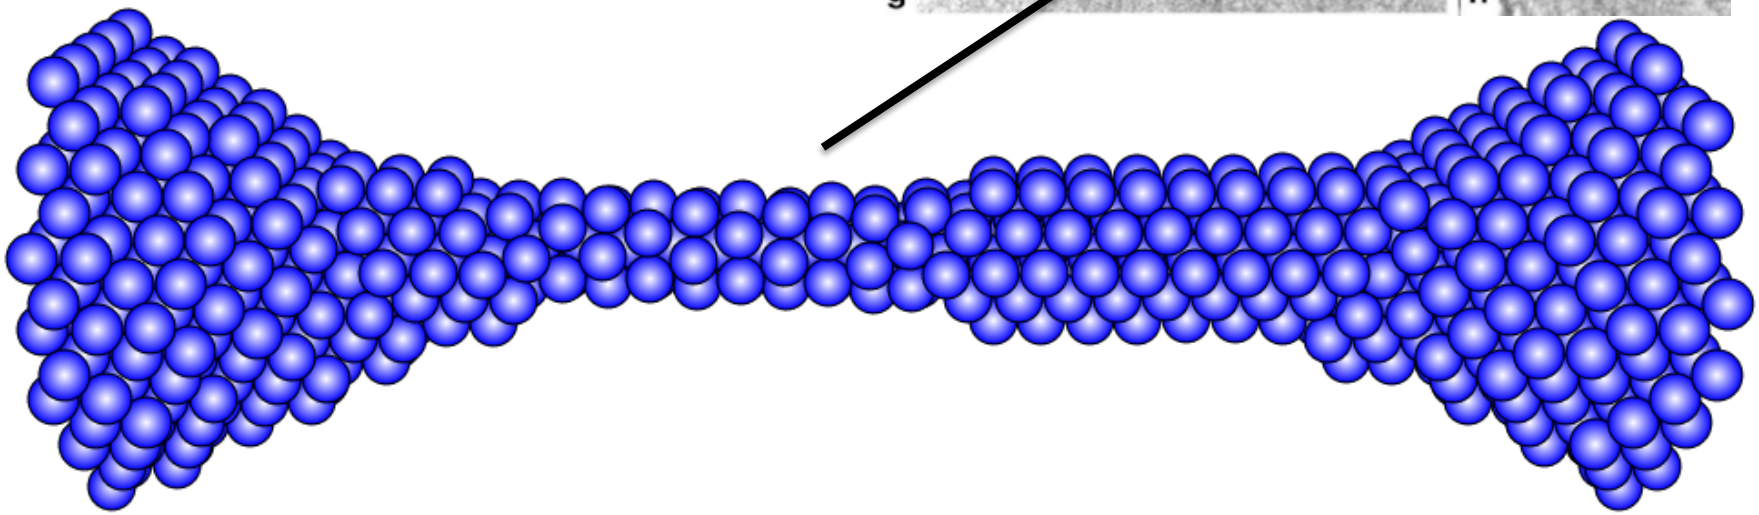

Ag nanowire

It is possible that this icosahedral chain is what Rodriguez et. al. (PRL, 2002) saw in their Ag [110] nanowire experiments:

"In fact, in many cases, we have observed that when the wire attains this peculiar structure, the apexes' retraction causes the nanowire to elongate by a factor ~1.5-3 without thinning. This lengthening reflects the enhanced strength of this atomic configuration."

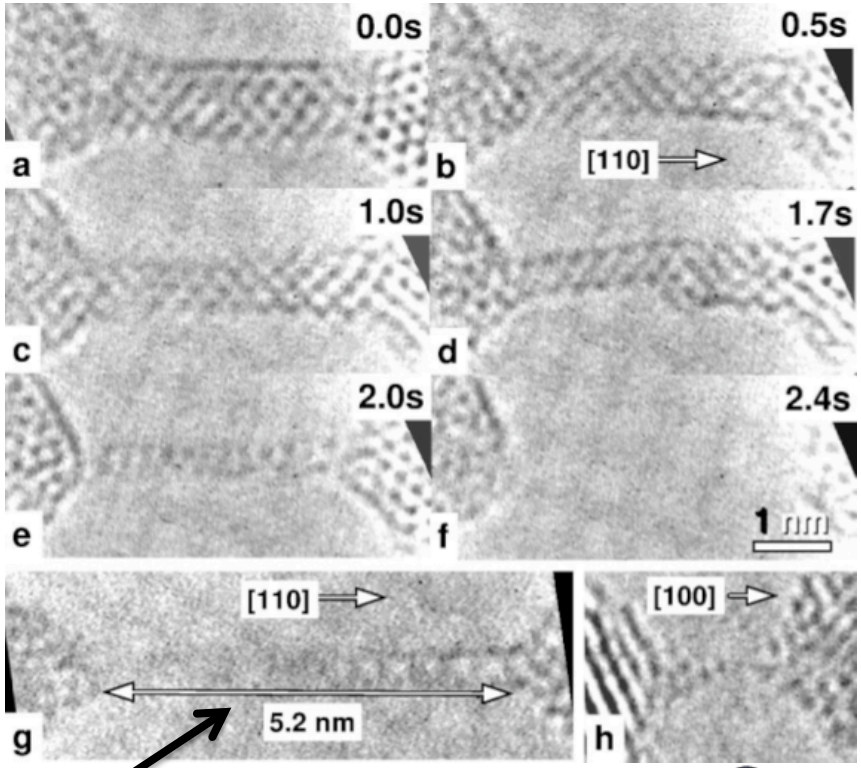

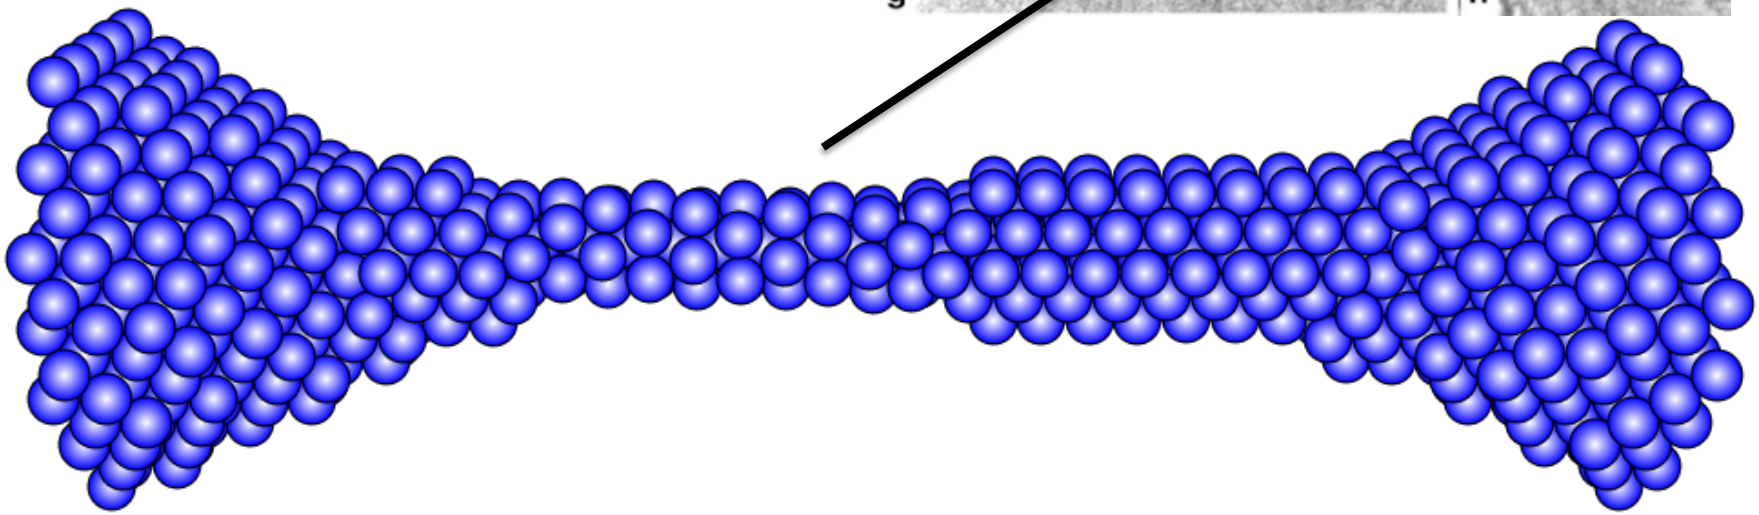

Basin complexity and ParRep

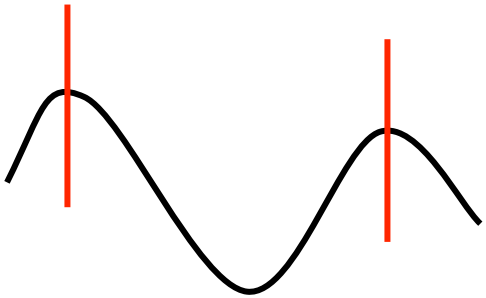

Straightforward: τ_{corr} = few vibrational periods.
Clean exponential escape distribution.

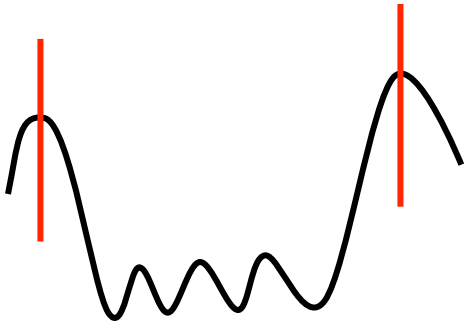

Longer τ_{corr} (multiple hops between sub-basin states).
Exponential escape distribution if allowed to equilibrate.

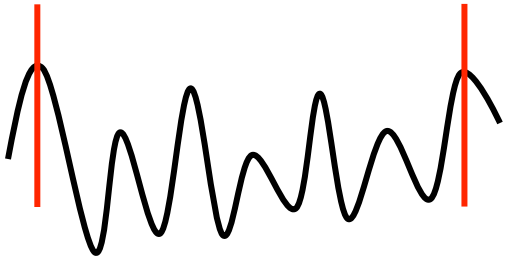

Here we seem to have a problem, since we don't expect it to be cleanly exponential.

The new understanding

The quasi-stationary distribution

During the ParRep dephasing step, we remove (and perhaps restart) any trajectories that escape from the state.

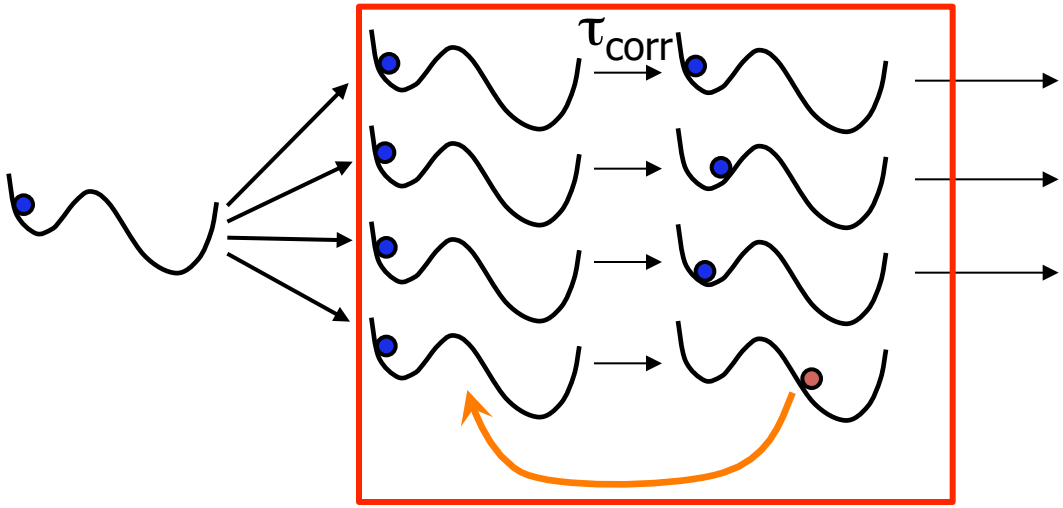

This ParRep dephasing procedure rigorously prepares a "quasi-stationary distribution" (QSD).

The quasi-stationary distribution (QSD)

The QSD is the distribution that results in the long time limit of dynamics in a potential with absorbing boundaries.

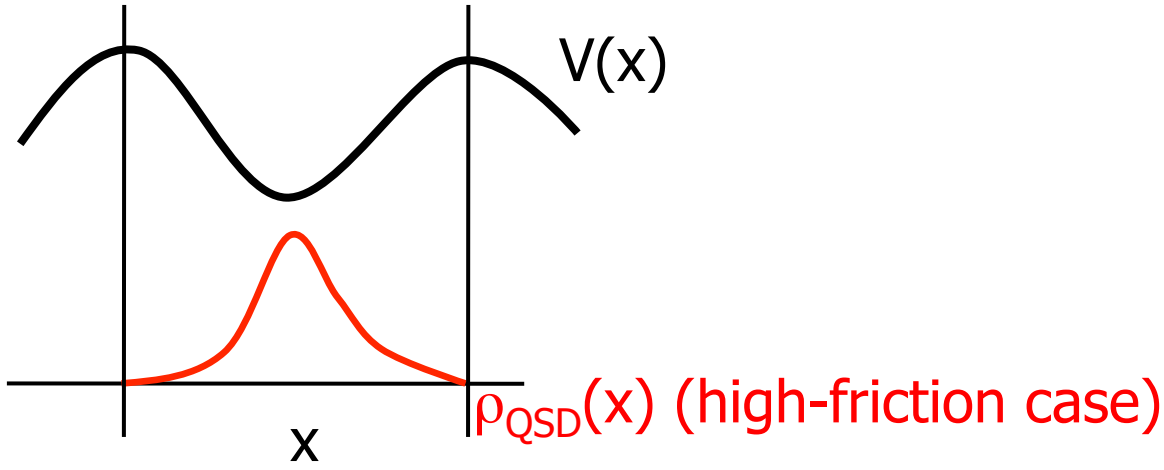

C. Le Bris, T. Lelièvre, M. Luskin, D. Perez, *A mathematical formalization of the parallel replica dynamics*, Monte Carlo Methods and Applications (in press); available as arXiv:1105.4636.

LA-UR-12-26438

The quasi-stationary distribution (QSD)

The QSD is the distribution that results in the long time limit of dynamics in a potential with absorbing boundaries.

Note that it is not the same as the canonical ensemble:

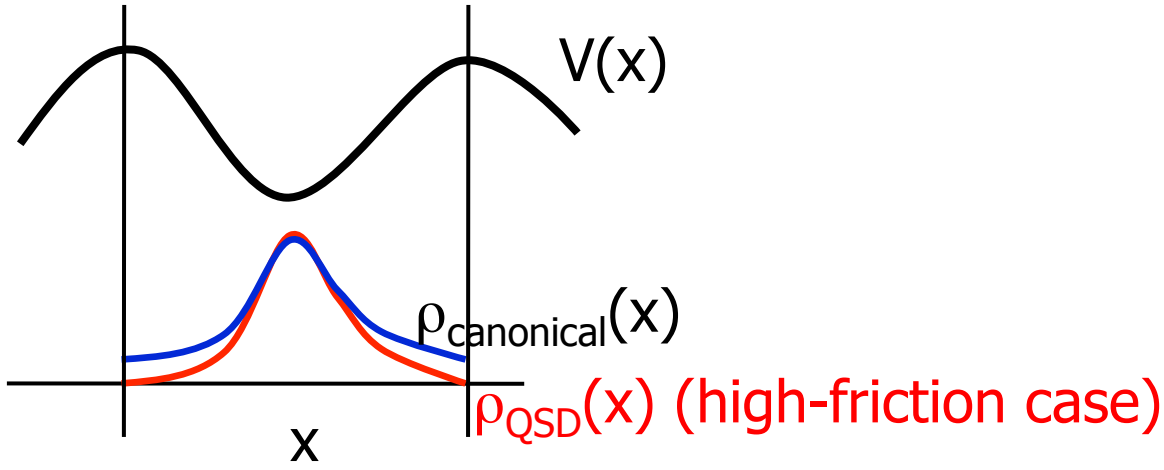

(although for a high barrier, it is virtually the same)

C. Le Bris, T. Lelièvre, M. Luskin, D. Perez, *A mathematical formalization of the parallel replica dynamics*, Monte Carlo Methods and Applications (in press); available as arXiv:1105.4636.

LA-UR-12-26438

Properties of the QSD

The probability distribution for the first escape from the QSD is an exponential, and the escape hitting points are independent of time.

Thus, it is appropriate for ParRep.

Moreover, it has these properties regardless of where the boundaries are positioned (!). E.g.

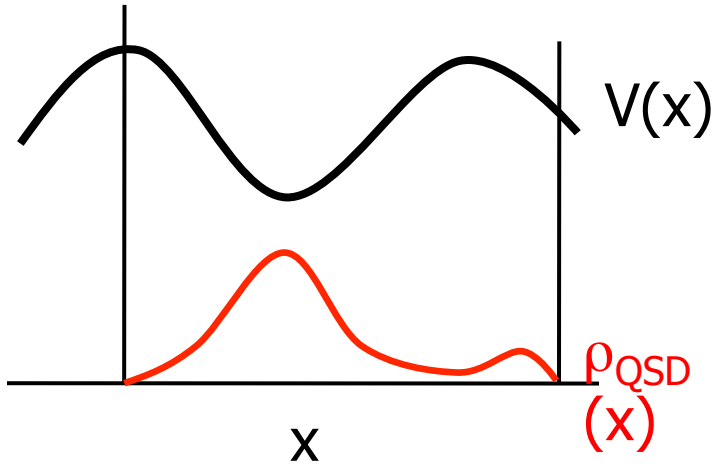

C. Le Bris, T. Lelièvre, M. Luskin, D. Perez, *A mathematical formalization of the parallel replica dynamics*, Monte Carlo Methods and Applications (in press); available as arXiv:1105.4636.

LA-UR-12-26438

Exploiting the QSD concept

This means we can even use ParRep on a "dangerous" case, provided we dephase long enough.

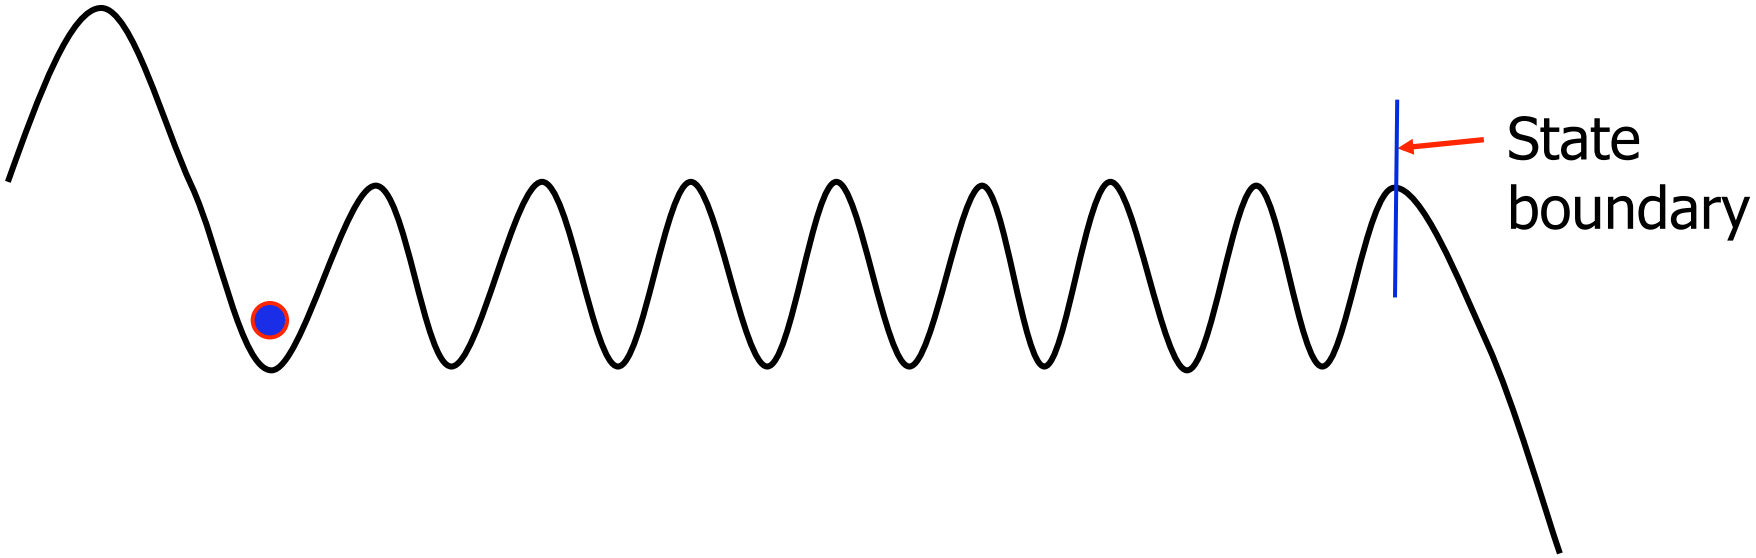

The required dephasing time may be so long that it is less efficient than simply continuing the ParRep in and out of each individual basin, or simply doing direct MD, *but the ParRep will be correct*.

Exploiting the QSD concept

This means we can even use ParRep on a "dangerous" case, provided we dephase long enough.

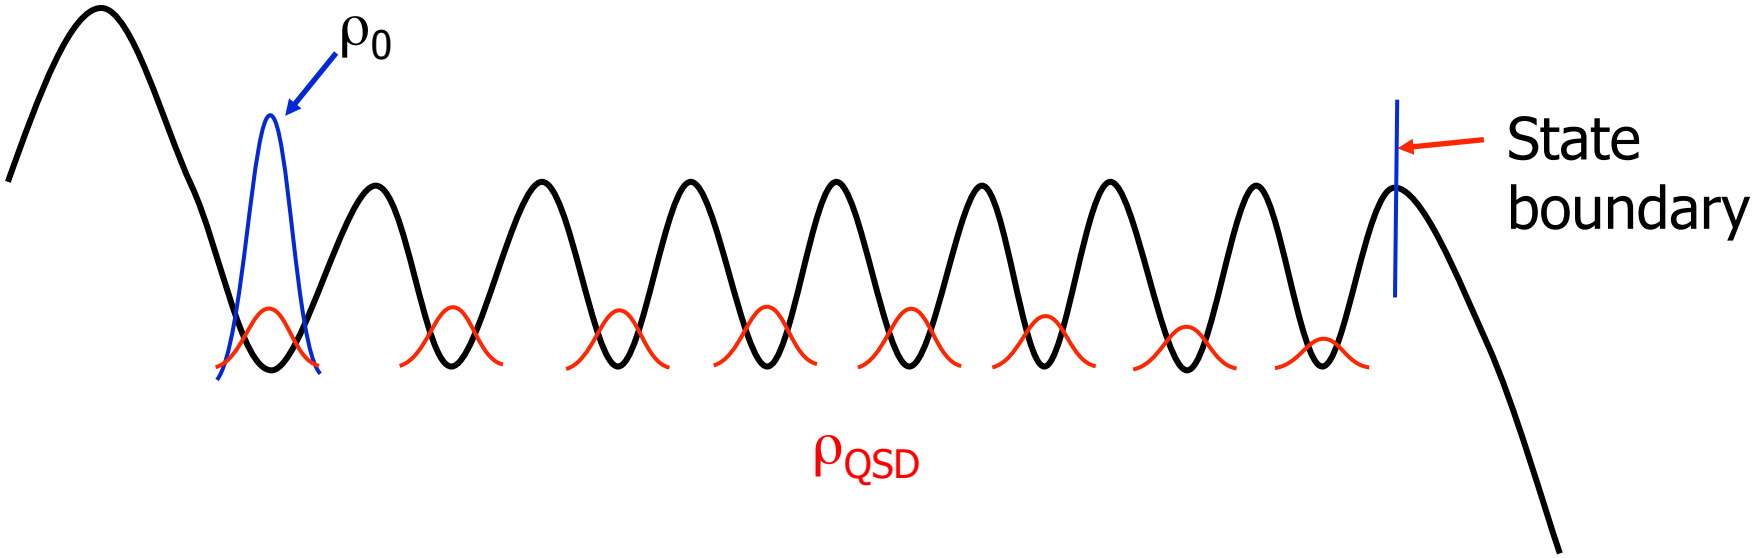

The required dephasing time may be so long that it is less efficient than simply continuing the ParRep in and out of each individual basin, or simply doing direct MD, *but the ParRep will be correct*.

Exploiting the QSD concept

This means we can even use ParRep on a "dangerous" case, provided we dephase long enough.

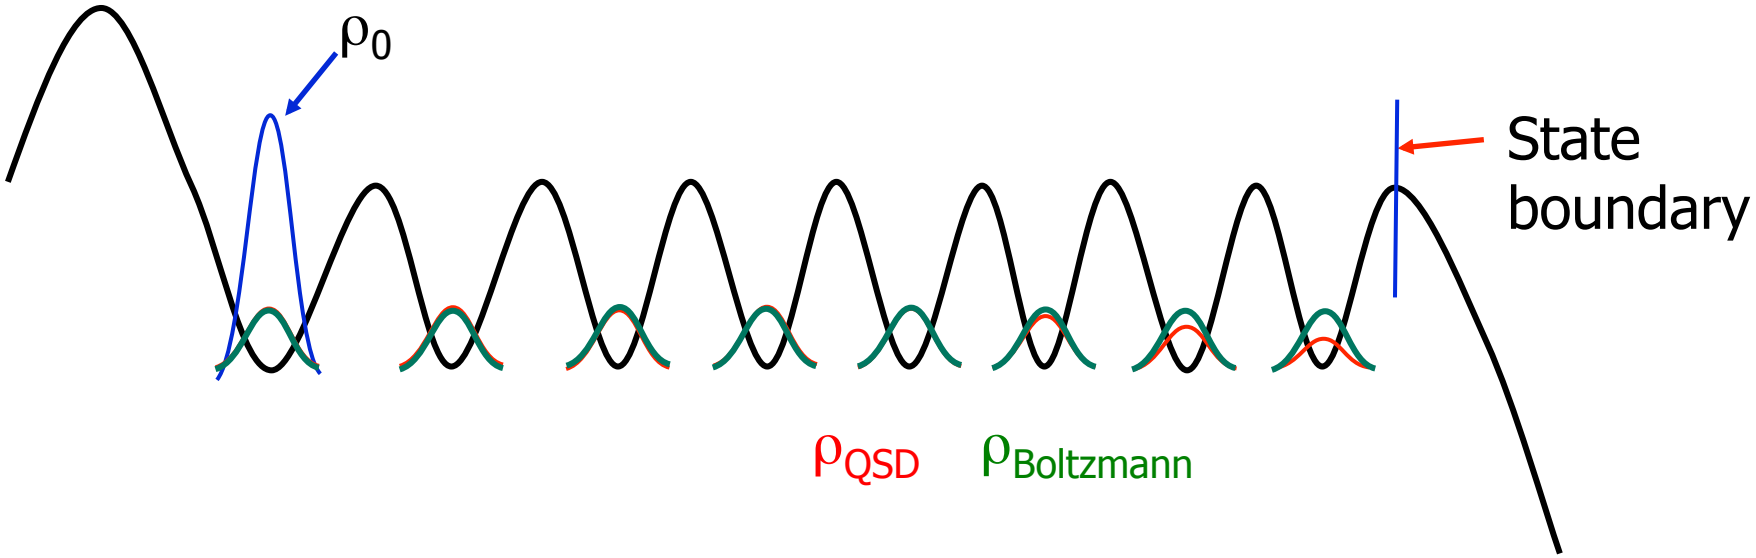

The required dephasing time may be so long that it is less efficient than simply continuing the ParRep in and out of each individual basin, or simply doing direct MD, *but the ParRep will be correct*.

Error in preparation of the QSD

For this overdamped Langevin case, LeBris et al* showed that the ParRep dephasing step prepares the QSD exactly as $\tau_{corr} \rightarrow$ infinity, with error at finite time τ_{corr} proportional to

$$\exp[-(\lambda_2-\lambda_1)\tau_{corr}]$$

where λ_1 and λ_2 are the 1st and 2nd eigenvalues of the Fokker-Planck equation with absorbing boundaries. Regular Langevin also gives a QSD, but it is harder to derive the error bound.

*C. Le Bris, T. Lelièvre, M. Luskin, D. Perez, *A mathematical formalization of the parallel replica dynamics*, Monte Carlo Methods and Applications (in press); available as arXiv:1105.4636.

The ParRep procedure, revisited

The procedure is exactly as stated originally.

No need to require it is a rare event to get exponential behavior. We force it to give us a good exponential, by dephasing long enough to obtain the quasi-stationary distribution (QSD) accurately before beginning the parallel stage.

The same τ_{corr} is used for both the dephasing stage and the correlation stage.

If the transitioning trajectory escapes from the superbasin before τ_{corr}, then there is no parallel step -- we just continue the MD into the next superbasin.

Parallel Replica Dynamics - revisited

Parallelizes time evolution

Assumptions:

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

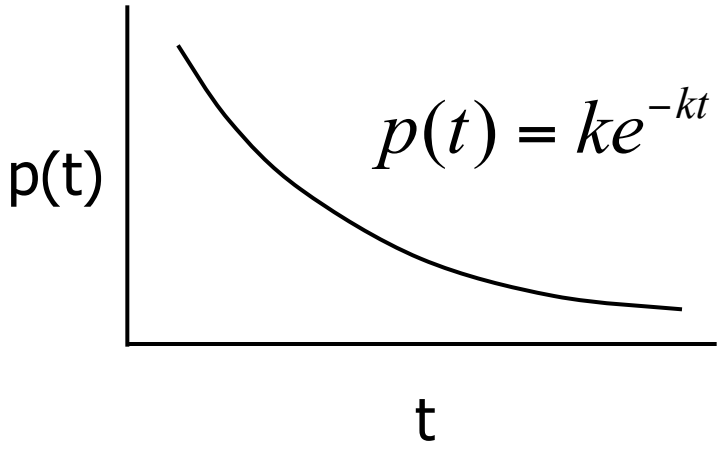

- correlation time known

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

Parallel Replica Dynamics - revisited

Parallelizes time evolution

Assumptions:

- ~~infrequent events~~ not required
- transitions can be detected

choose any state definition you want

- ~~exponential distribution of first-escape times~~ enforced automatically

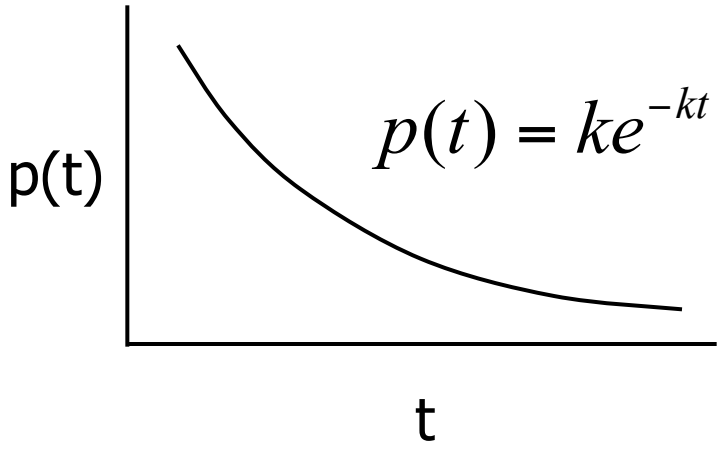

- correlation time known

This becomes the key requirement – τ_{corr} must be long enough.

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

Choosing the superbasin definition

We are free to choose any state definition we want.

If we can optimize the definition to maximize the separation of time scales, we will get more boost. (may or may not be difficult, depending on system)

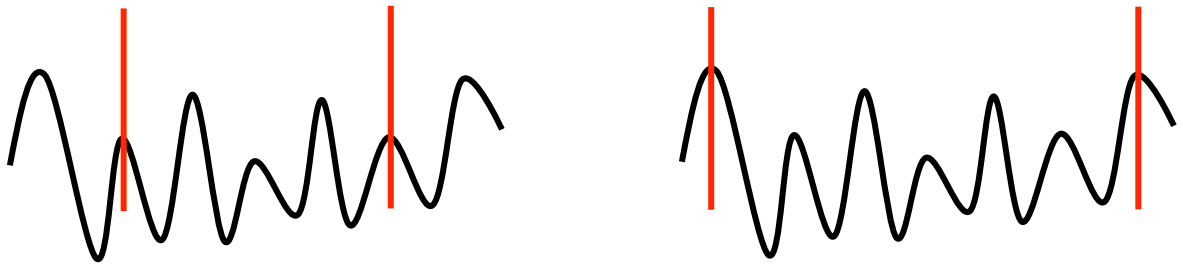

bad

better

Examples where QSD and general state definition will be powerful

Messy low-barrier situations (e.g., surface clusters, interstitials, grain boundaries, …)

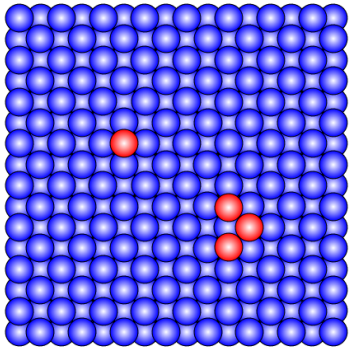

(just one simple example, but these show up *everywhere*)

Diffusion in a polymer

Soft matter dynamics

Glass dynamics

Protein folding 

Parallel-replica dynamics

Advantages:

Most general and accurate of the three methods.

Very easy to implement.

Can treat more complex systems; no saddles required, no TST assumption.

Good match to the massively parallel future of computing.

New QSD understanding (LeBris et al) allows arbitrary state definition, enforcing exponentiality automatically; burden falls to determining appropriate dephasing time.

Disadvantage:

Boost requires multiple processors.

TAD

LA-UR-12-26438

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.

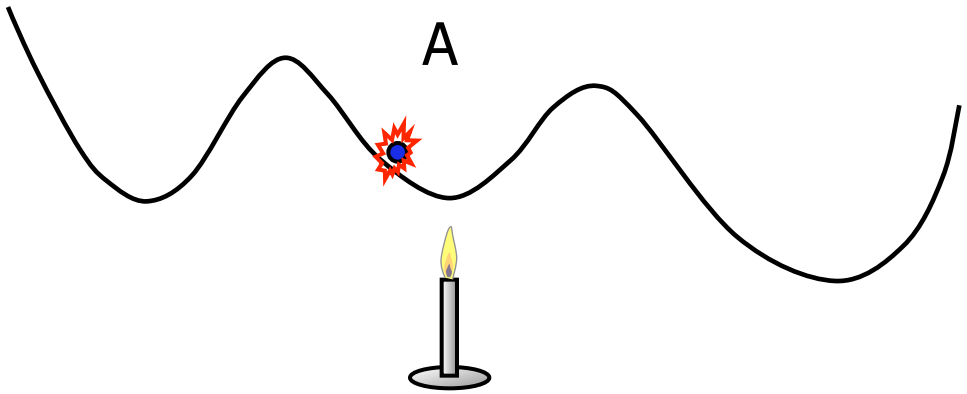

TAD Procedure

- 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.

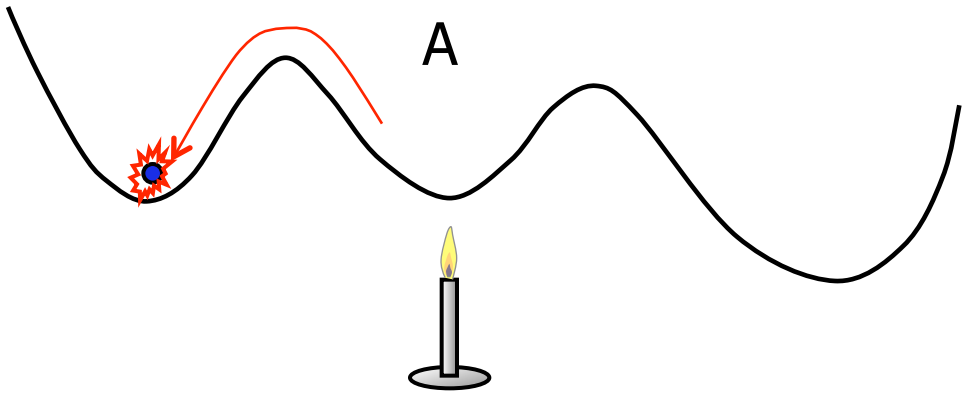

TAD Procedure

- 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.

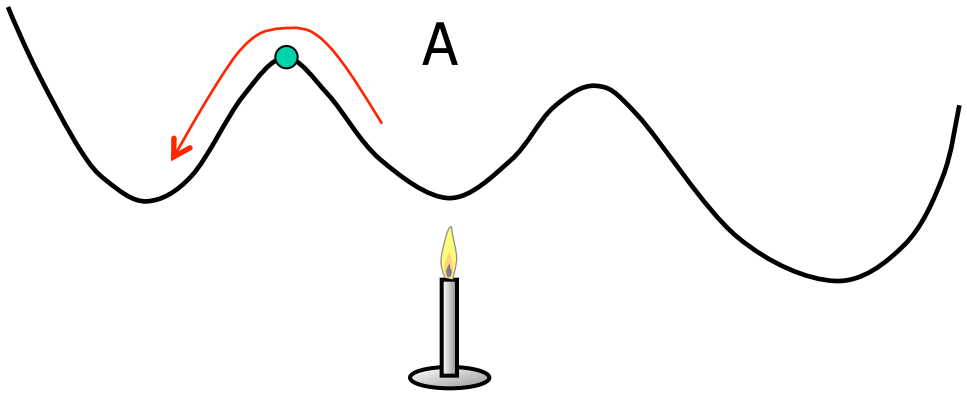

TAD Procedure

- 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.

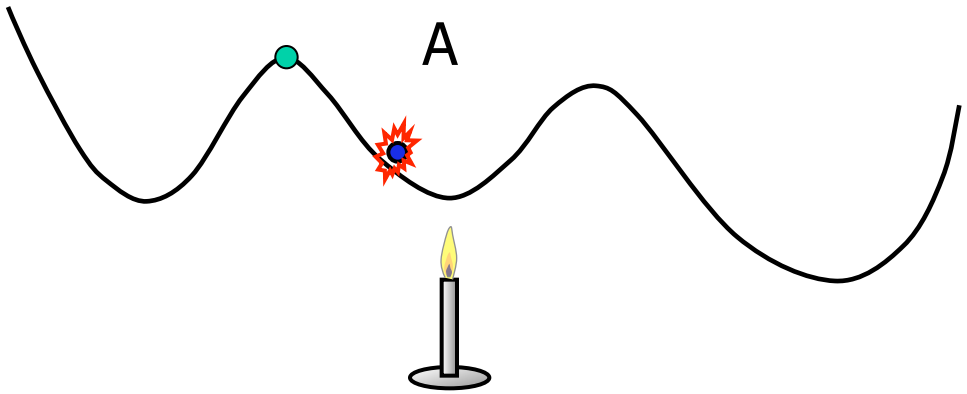

TAD Procedure

- 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.

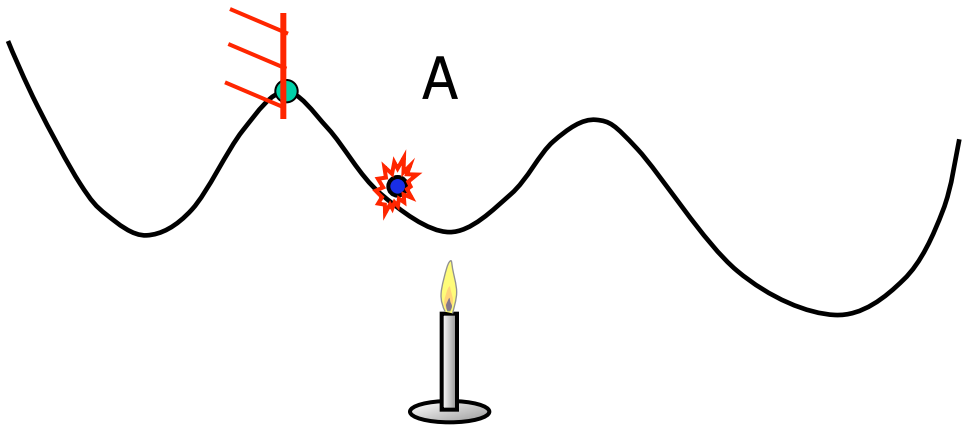

TAD Procedure

- 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.

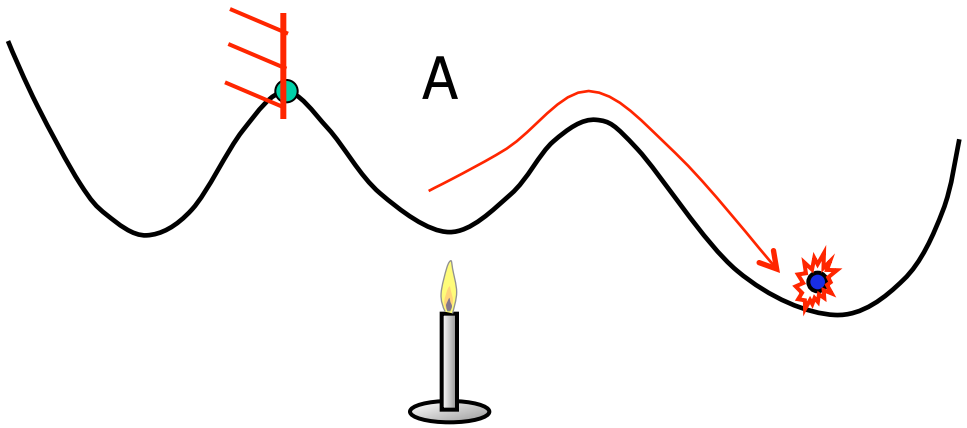

TAD Procedure

- 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.

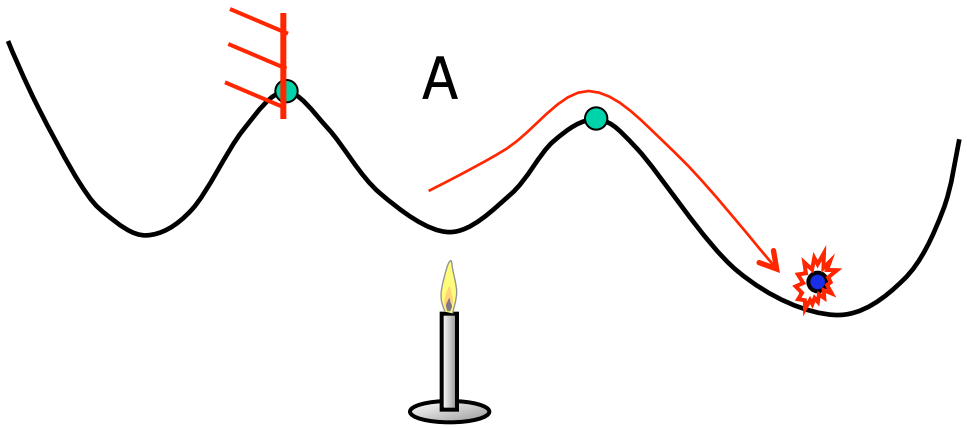

TAD Procedure

- 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.

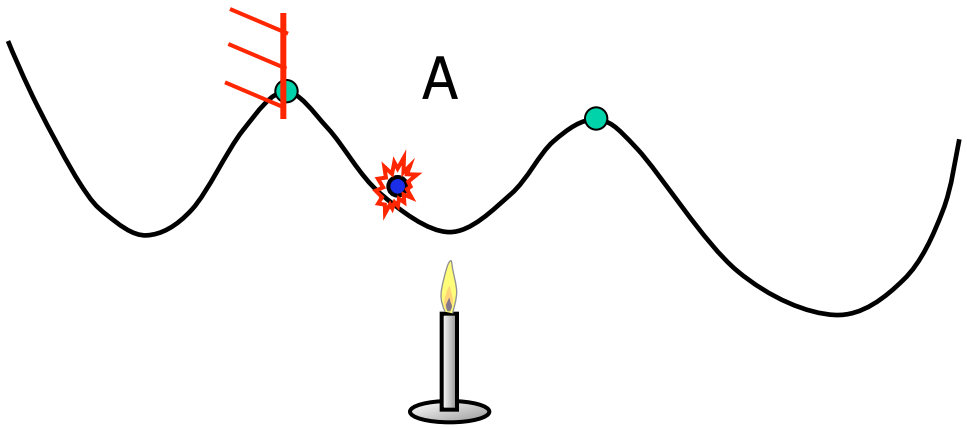

TAD Procedure

- 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.

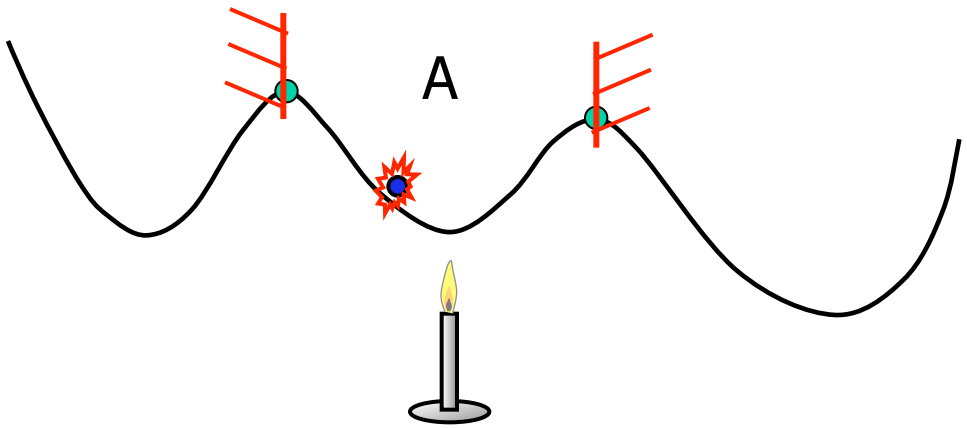

Nudged Elastic Band (NEB) to find saddle point

Construct chain of configurations connecting initial to final state.

Minimize energy of configurations plus springs between them.

Highest-energy point is the saddle point (basically).

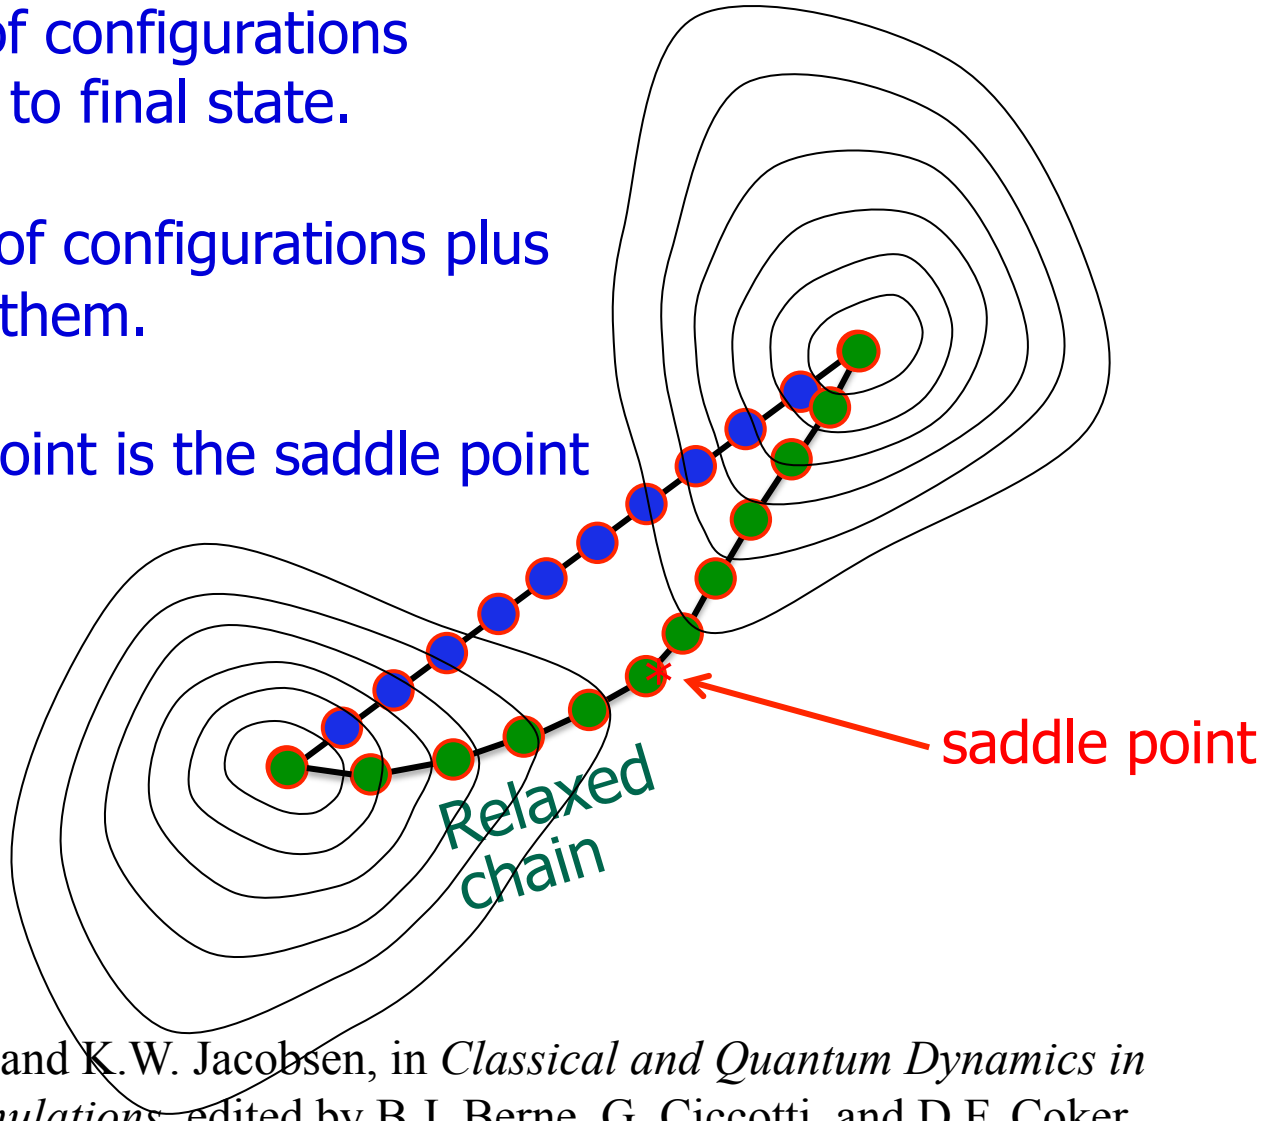

H. Jónsson, G. Mills, and K.W. Jacobsen, in *Classical and Quantum Dynamics in Condensed Phase Simulations*, edited by B.J. Berne, G. Ciccotti, and D.F. Coker (World Scientific, Singapore, 1998), p. 385.

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:

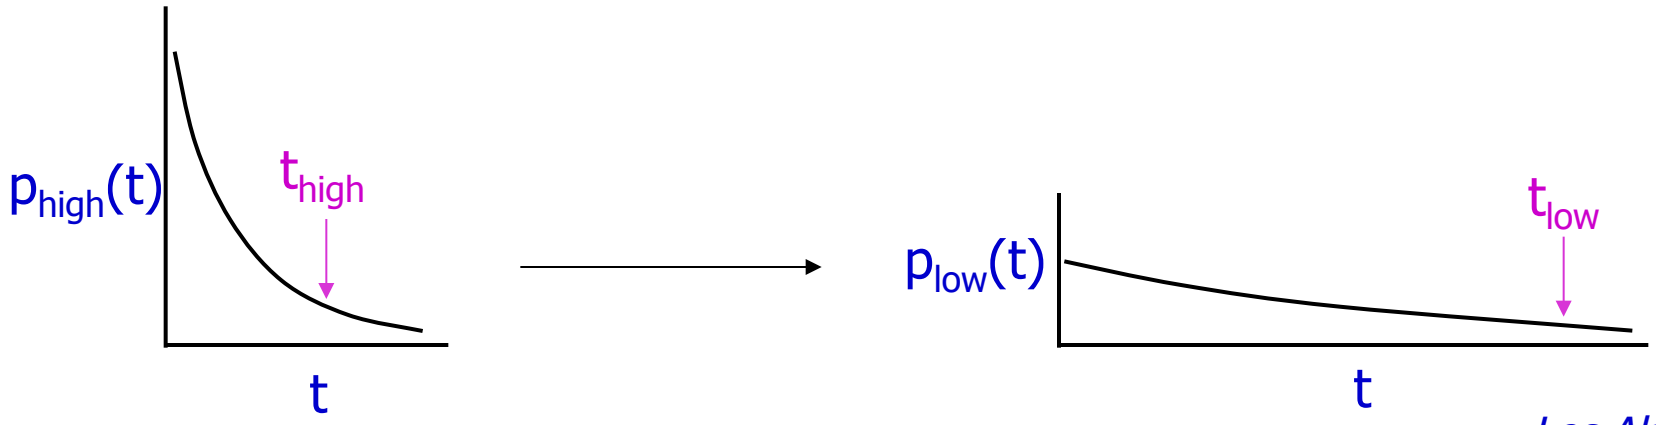

The Arrhenius view

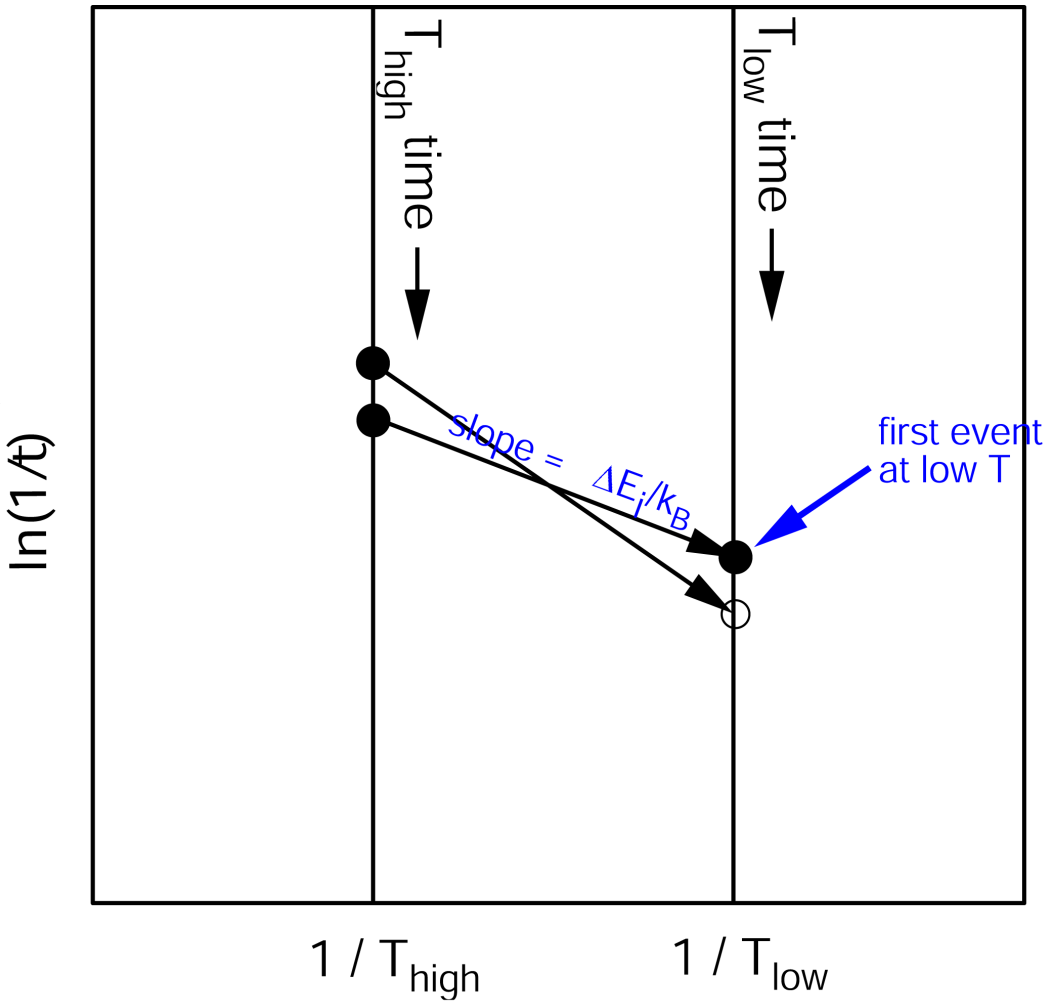

when can we stop?

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 $\ln[\nu_0/\ln(1/\delta)]$

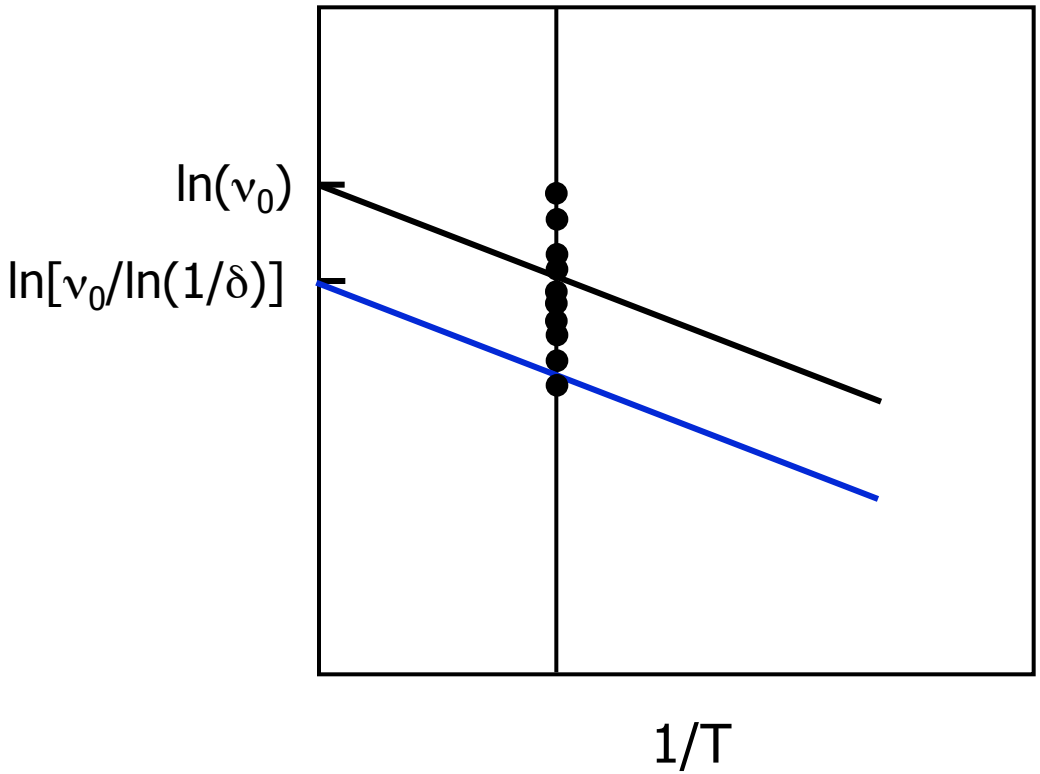

TAD - when can we stop the MD and accept an event?

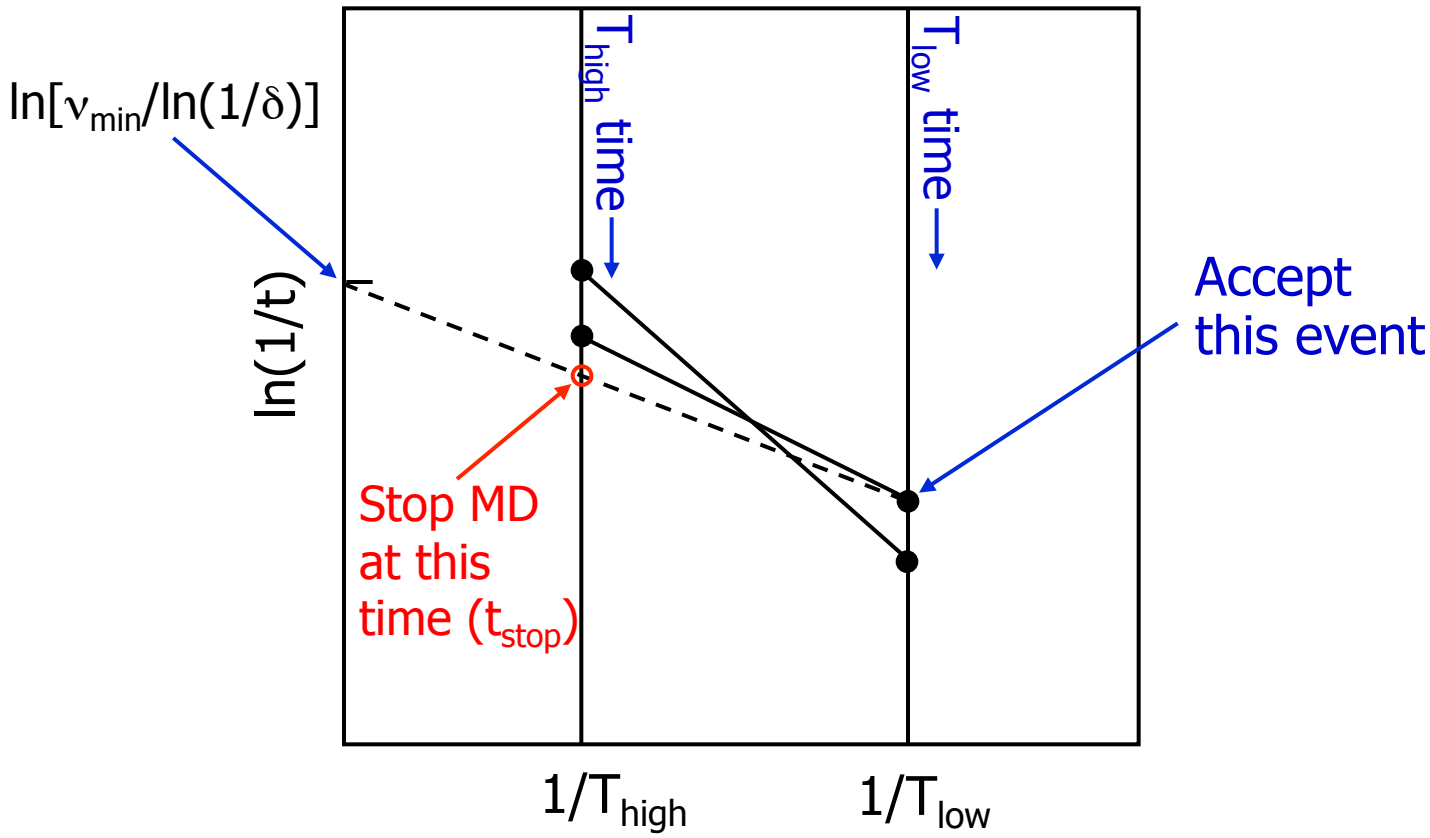

After time t_{stop}, with confidence 1-δ, no event can replace shortest-time event seen at low T.

Move system to this state and start again.

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

Examples of TAD Studies

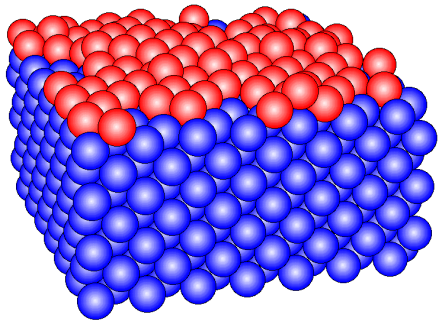

Cu/Ag(100), 1 ML/25 s
T=77K, Sprague et al, 2002.

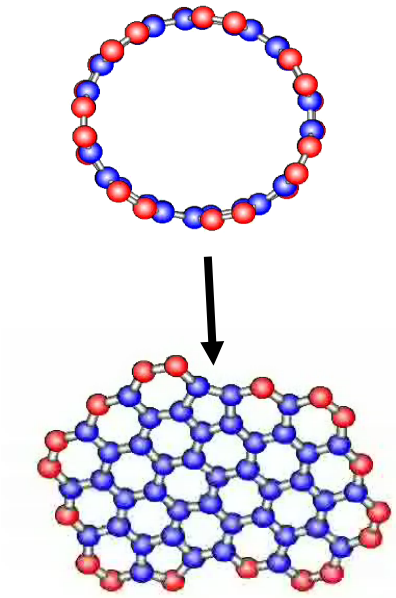

Annealing nanotube slices, μs, Uberuaga et al, 2011.

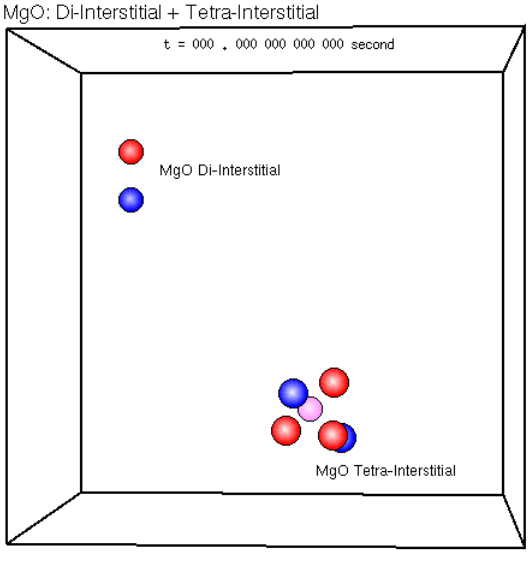

TAD simulation, Uberuaga et al, 2003

Interstitial defects in MgO, ps – s, Uberuaga et al, 2004.

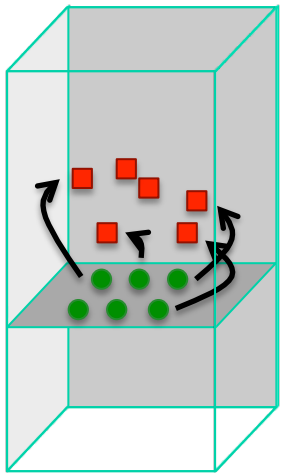

Interstitial emission from GB after cascade, μs, Bai et al, Science, 2010.

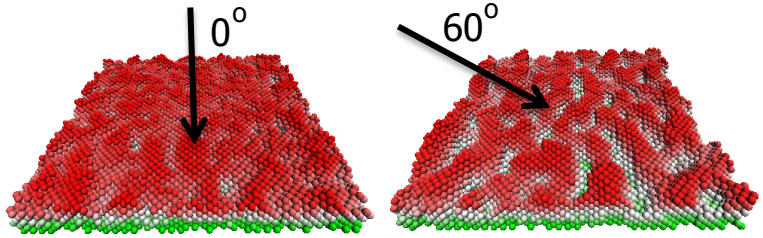

Growth of Cu(001), MD+ParTAD, 5 ML/ms, Shim, Amar et al, 2008.

Recent brief review: Perez et al, Ann. Rep. Comp. Chem. **5**, 79 (2009).

Radiation damage annealing near a GB

EAM copper, MD + TAD, "self healing" effect

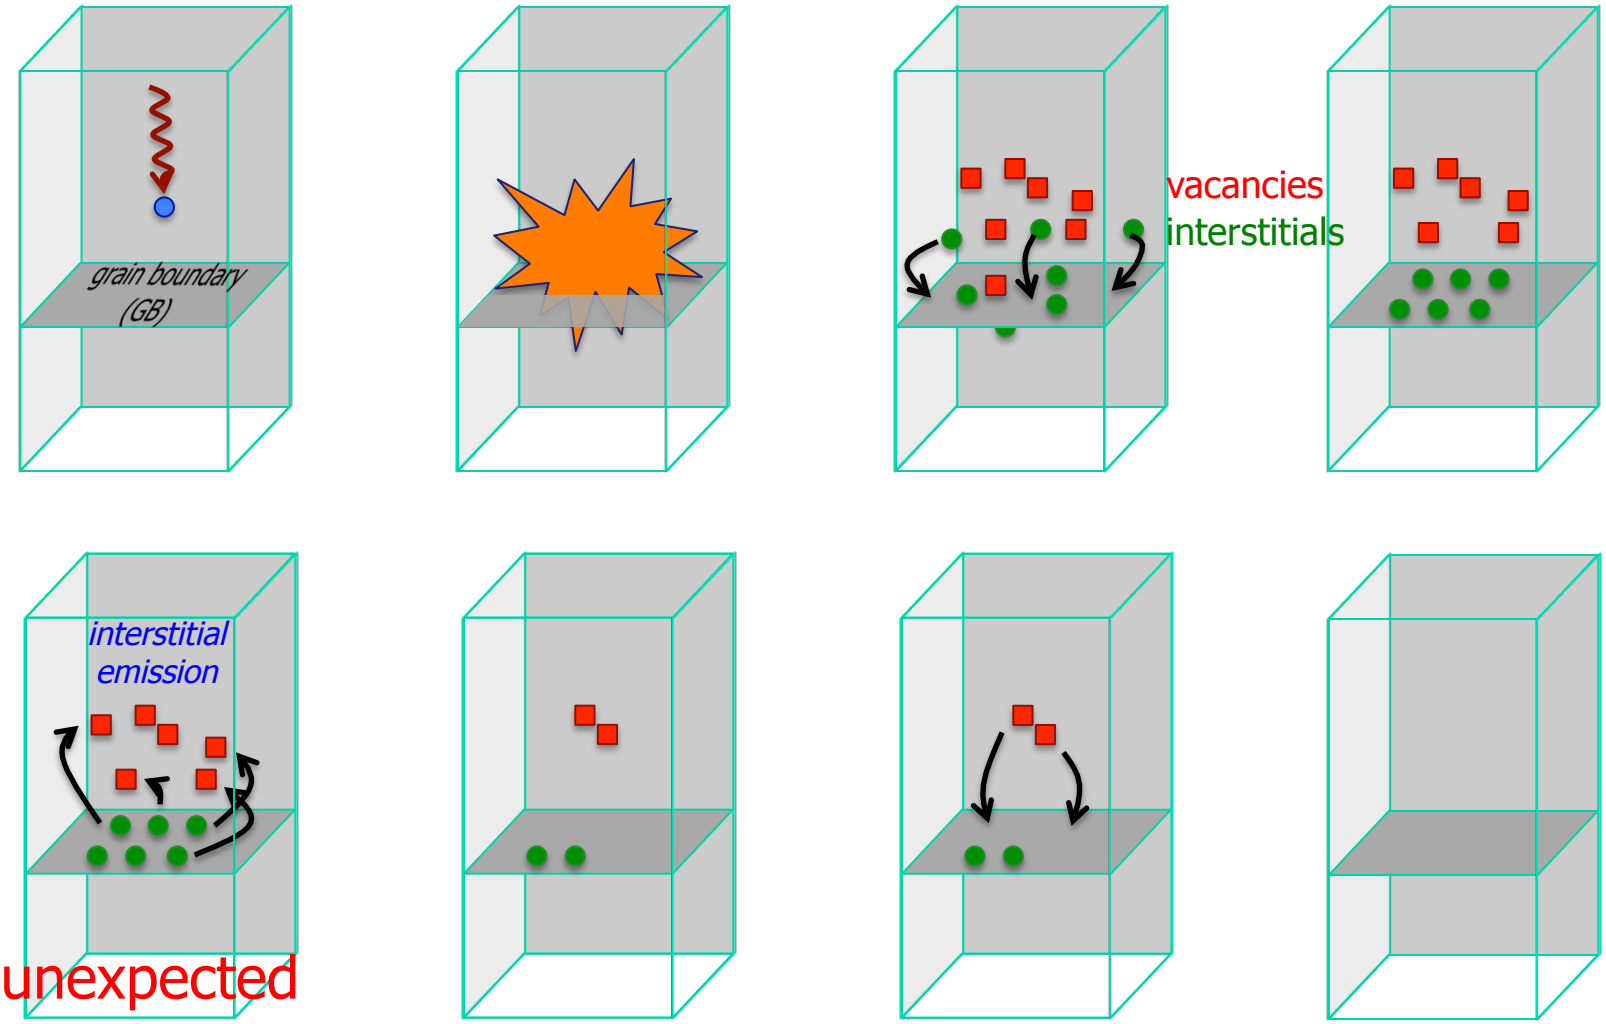

Xian-Ming Bai et al, Science 327, 1631 (2010).

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 (but making progress)
  - Low barriers
  - More compex systems
  - Scaling with system size

Reviews:

Perez et al, Ann. Rep. Comp. Chem. **5**, 79 (2009).
Voter, Montalenti, and Germann, Ann. Rev. Mater. Res. 32, 321 (2002).