Local Hyperdynamics

Arthur F. Voter
Theoretical Division
Los Alamos National Laboratory
Los Alamos, NM, USA

COSMOS Workshop
Paris
February 4, 2016

Acknowledgments:

Danny Perez (LANL, T-1), Soo Young Kim (former LANL postdoc)
Dipanjan Ray (LANL postdoc), Tim Germann (LANL, T-1)

DOE Office of Basic Energy Sciences, Los Alamos LDRD

Los Alamos
LA-UR-16-20856

The time-scale challenge

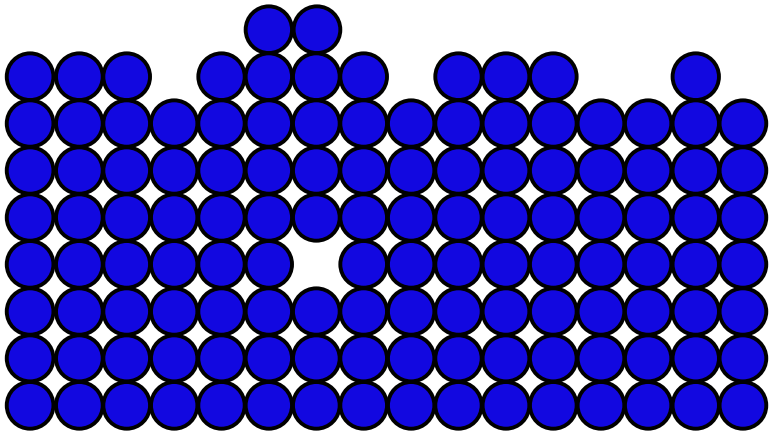

We have some system (e.g., adatoms on a surface).

We know that if we wait long enough, something will happen.
And then something else, and then something else,…

Using MD, we can run about 1 microsecond -- might not even see first event.

How do we accurately predict the long-time evolution?

Often the long-time evolution involves infrequent events…

Accelerated molecular dynamics approach

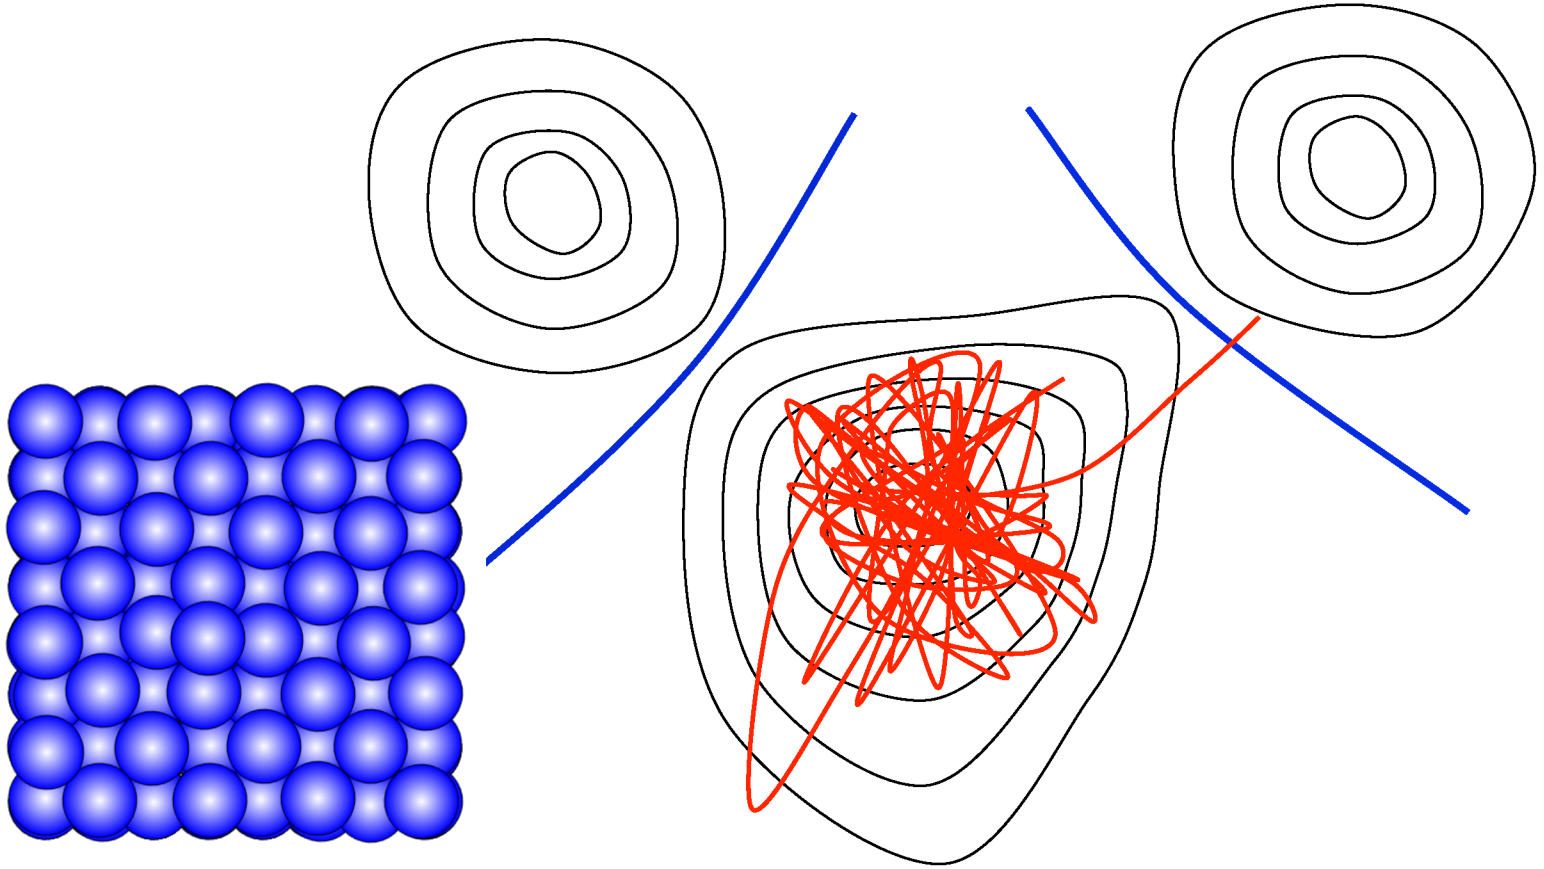

The trajectory finds an appropriate way out (i.e., proportional to the rate constant) without knowing about any of the escape paths except the one it first sees. In AMD, we trick it into doing this more quickly, trying very hard not to corrupt the relative probabilities.

Accelerated Molecular Dynamics Methods

Hyperdynamics

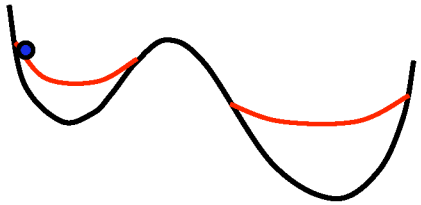

Parallel Replica Dynamics

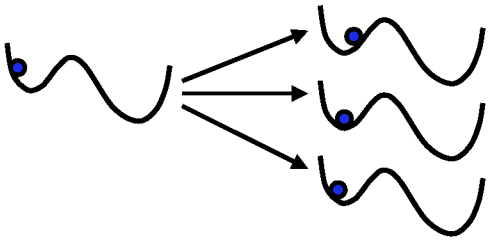

Temperature Accelerated Dynamics

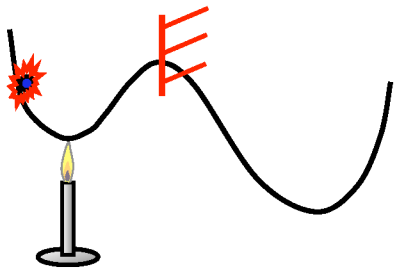

Recent review:
D. Perez, B.P. Uberuaga, and A.F. Voter, Computational Materials Science, **100**, 90 (2015)

Wide range of systems can be studied

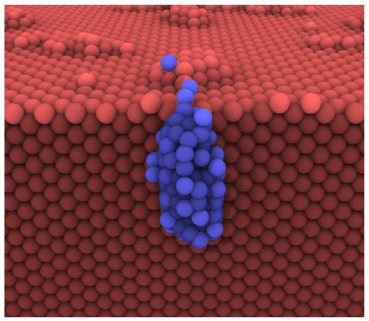

Growth/bursting of He bubble in W, μs, Sandoval et al, 2015.

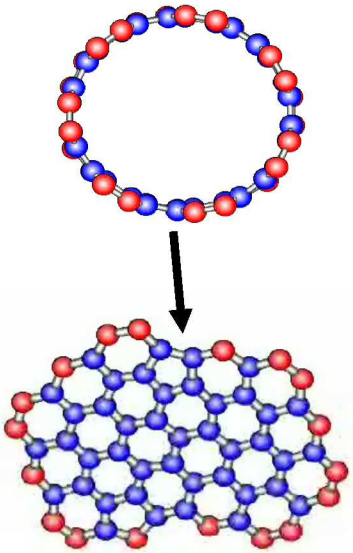

Annealing nanotube slices, μs, Uberuaga et al, 2011.

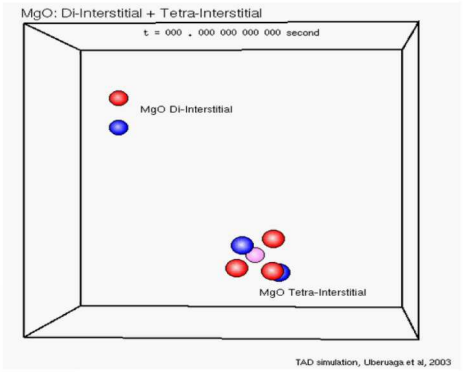

Interstitial defects in MgO, ps – s, Uberuaga et al, 2004.

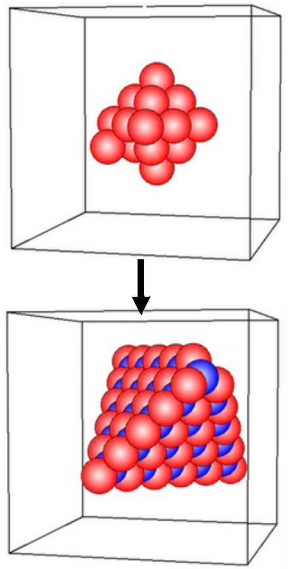

Cu void collapse to SFT, μs, Uberuaga et al, 2007.

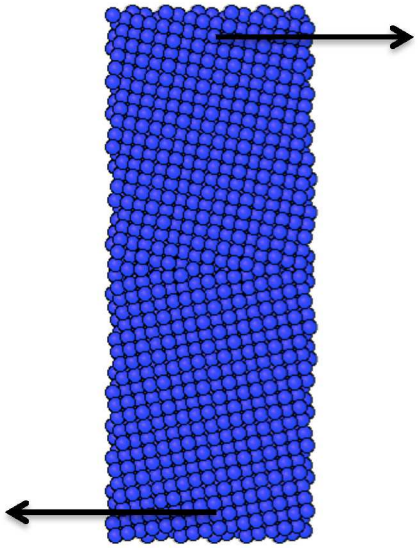

Driven Cu GB sliding, 500 μm/s Mishin et al, 2007.

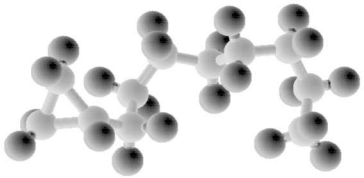

Hexadecane pyrolysis, μs, Kum et al, 2004.

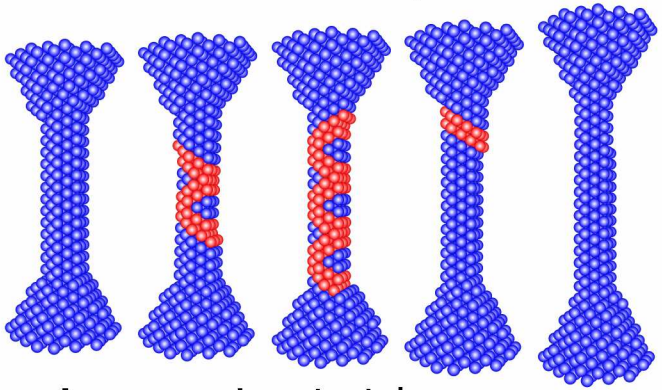

Ag nanowire stretch, μs - ms, Perez et al, 2015.

Characteristics of AMD methods

- In principle, the escape time can be collapsed to something approaching the correlation time – a few vibrational periods.
- When barriers are high, large boosts can be achieved
- When the barriers are low, boost is low.

Accelerated Molecular Dynamics Methods

Hyperdynamics

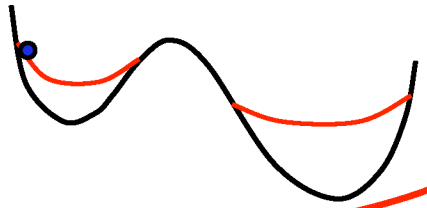

Parallel Replica Dynamics

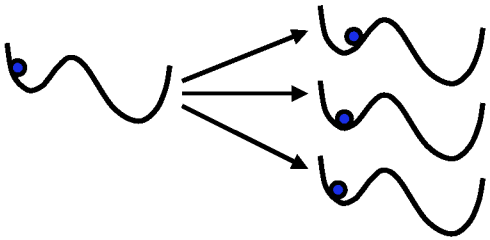

Temperature Accelerated Dynamics

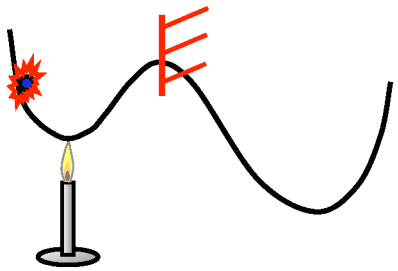

Hyperdynamics

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

Assumptions:

- infrequent events
- transition state theory (no recrossings)

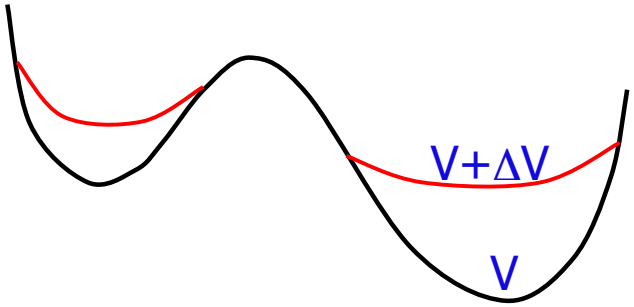

Procedure:

- 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 (relative escape rates are preserved)
- time converges on correct value in long-time limit (vanishing relative error)

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

The boost factor

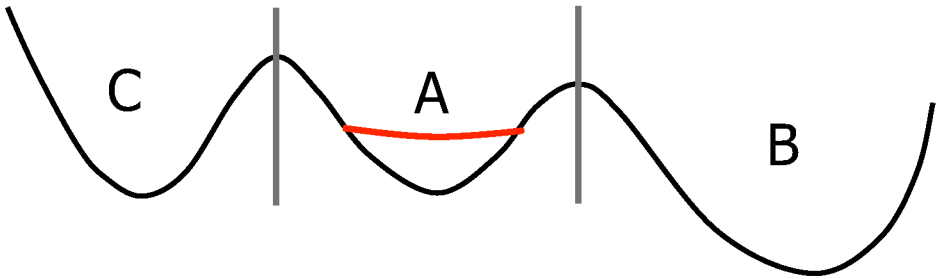

The boost factor (the hypertime divided by the MD time) is the average value of exp[+βΔV] on the biased potential:

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

Hyperdynamics – the bias potential

The hard part is designing the bias potential.

It needs to be zero on all dividing surfaces, even though we don't know in advance what the escape pathways are.

It also needs to maintain the TST-obeying nature of the dynamics – it cannot introduce any recrossings or other correlated events.

Bond boost method

Miron and Fichthorn, J. Chem. Phys. 119, 6210 (2003)

Bias is based on bond distortions; shuts off completely when the relative distortion ε of any "bond" exceeds a pre-chosen critical value q (e.g., q=0.3).

Simple and inexpensive to evaluate.

Probably the best existing bias potential.

Bond boost method

Miron and Fichthorn, J. Chem. Phys. 119, 6210 (2003)

Bias is based on bond distortions; shuts off completely when the relative distortion ε of any "bond" exceeds a pre-chosen critical value q (e.g., q=0.3).

Simple and inexpensive to evaluate.

Probably the best existing bias potential.

Simplified version ("simple bond boost"): ΔV depends purely on the coordinate (ε_{max}) of the most-distorted bond (Perez et al 2009).

Simplified bond-boost bias potential

ΔV depends purely on coordinate (ε_{max}) of most-distorted bond.
At most, one bond at a time has any bias force.

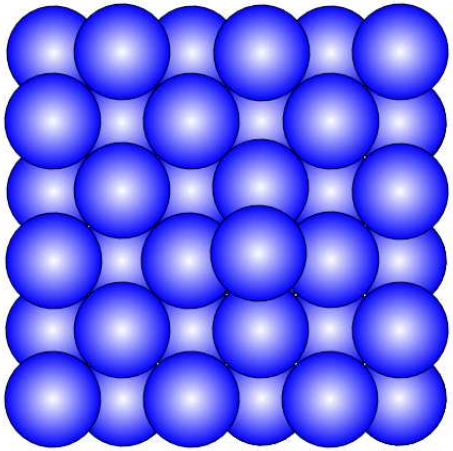

minimum

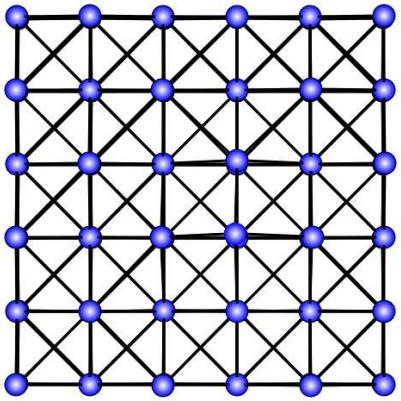

minimum

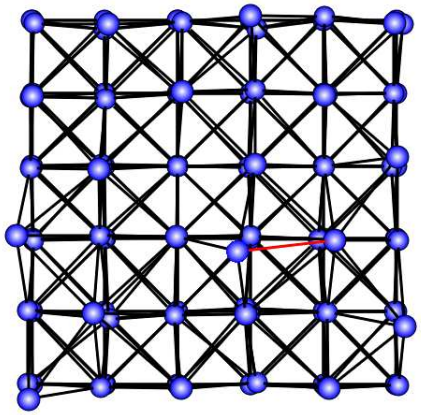

trajectory point

$\varepsilon_{ij} = (r_{ij}-r_{ij}^{min})/r_{ij}^{min}$ = relative bond distortion

$\Delta V(\varepsilon_{max}) = S[1-(\varepsilon_{max}/q)^2]$

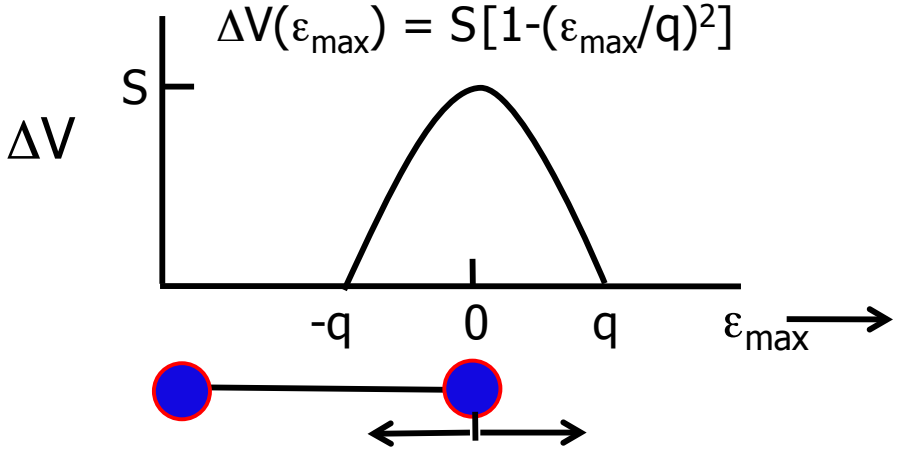

Simplified bond-boost bias potential

ΔV depends purely on coordinate (ε_{max}) of most-distorted bond.
At most, one bond at a time has any bias force.

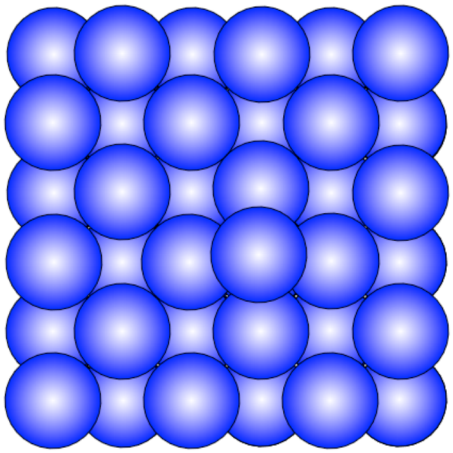

minimum

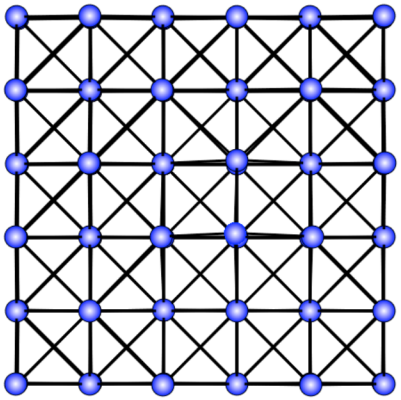

minimum

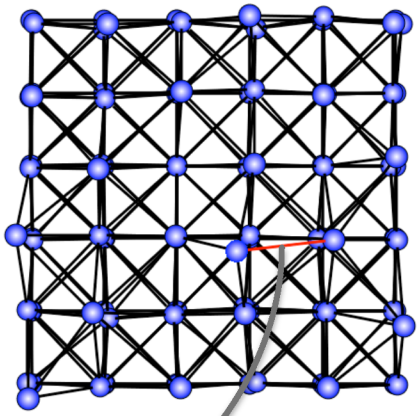

trajectory point

$\varepsilon_{ij} = (r_{ij}-r_{ij}^{min})/r_{ij}^{min}$ = relative bond distortion

$\Delta V(\varepsilon_{max}) = S[1-(\varepsilon_{max}/q)^2]$

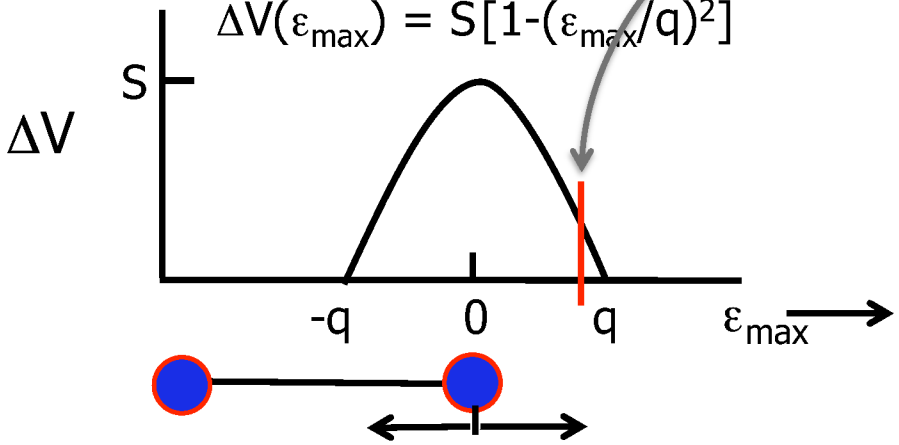

Simplified bond-boost bias potential

ΔV depends purely on coordinate (ε_{max}) of most-distorted bond.
At most, one bond at a time has any bias force.

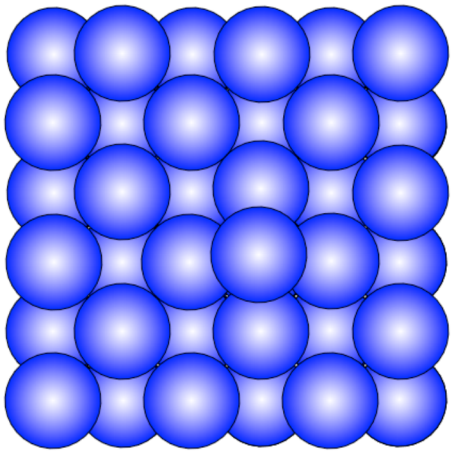

minimum

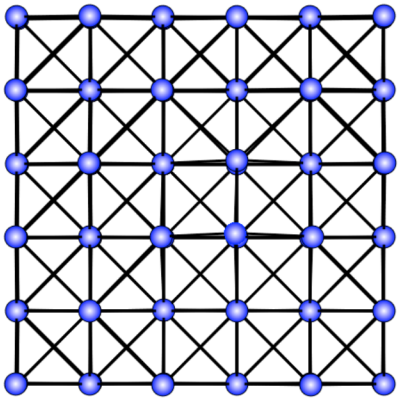

minimum

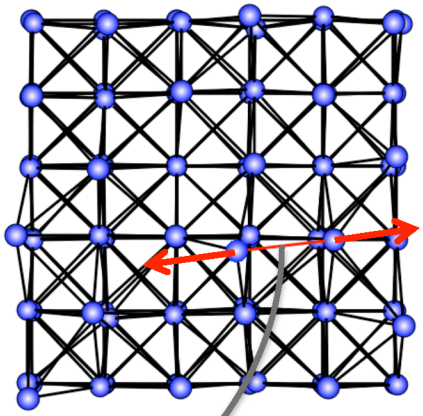

trajectory point

$\varepsilon_{ij} = (r_{ij}-r_{ij}^{min})/r_{ij}^{min}$ = relative bond distortion

$\Delta V(\varepsilon_{max}) = S[1-(\varepsilon_{max}/q)^2]$

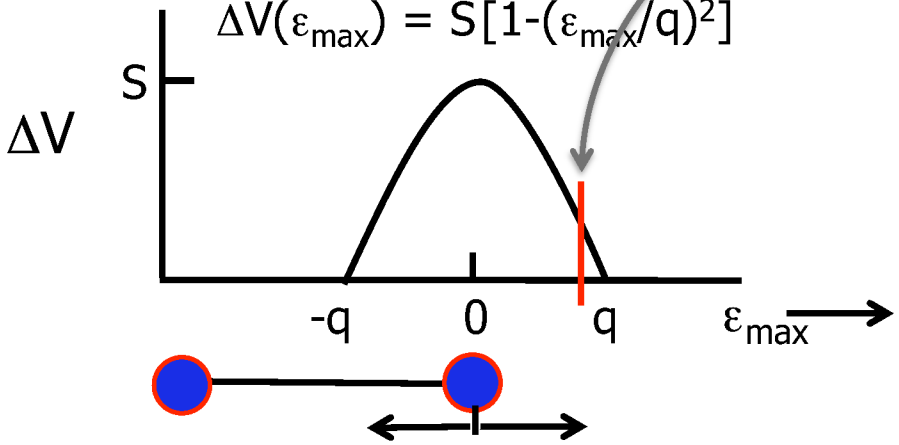

Simple bond-boost bias example

Cu adatom on Cu(100) surface

Hop barrier = 0.53 eV

T(K)	hop time	boost factor (S=0.4 eV)
300 K	27 μs	3.1×10^4
200 K	0.8 s	1.1×10^8

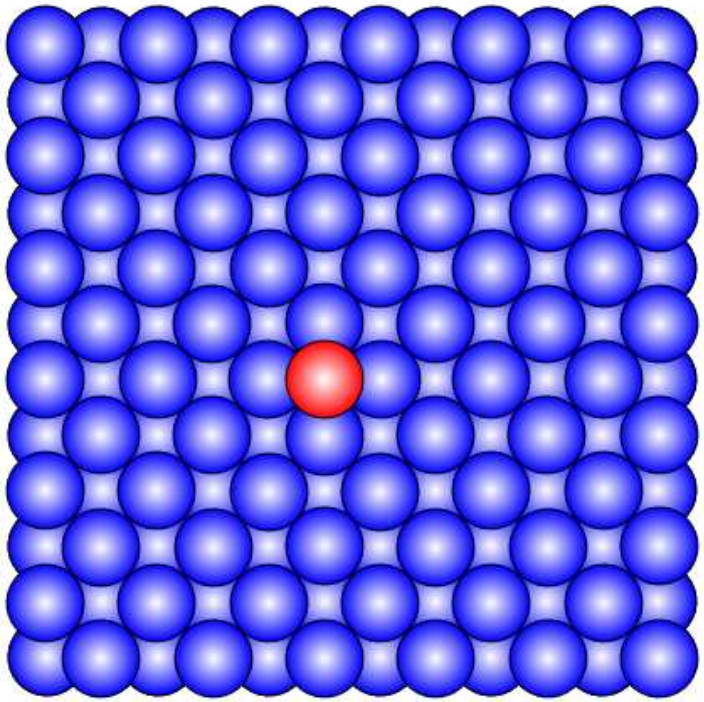

Local hyperdynamics for large systems

Los Alamos
LA-UR-16-20856

Hyperdynamics on large systems

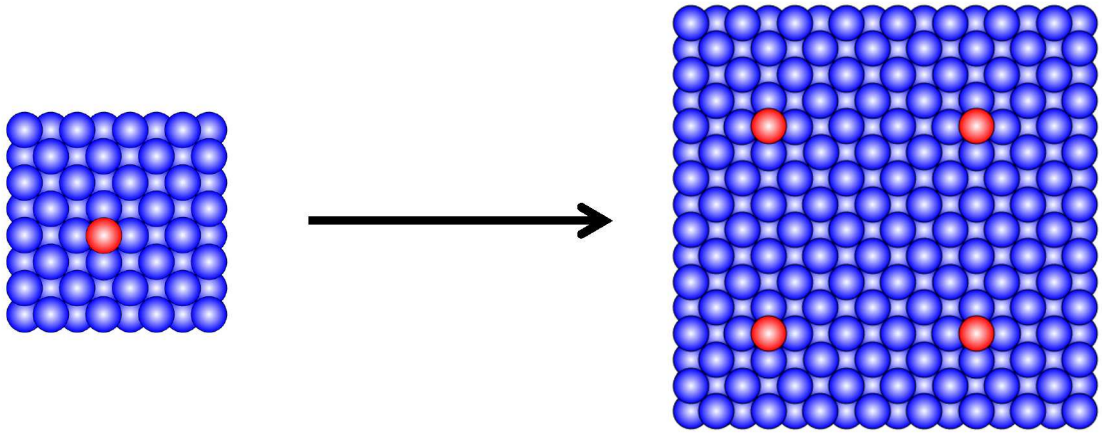

Whenever system is near a dividing surface, ΔV must be zero.

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*.

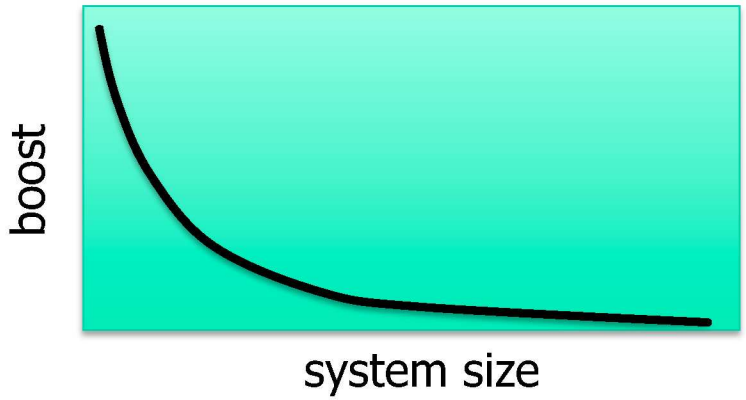

Hyperdynamics on large systems

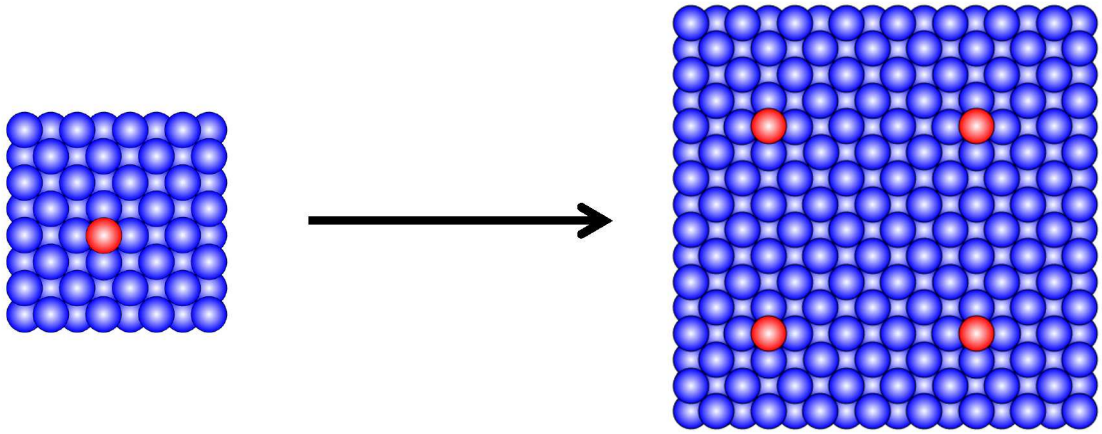

Whenever system is near a dividing surface, ΔV must be zero.

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*.

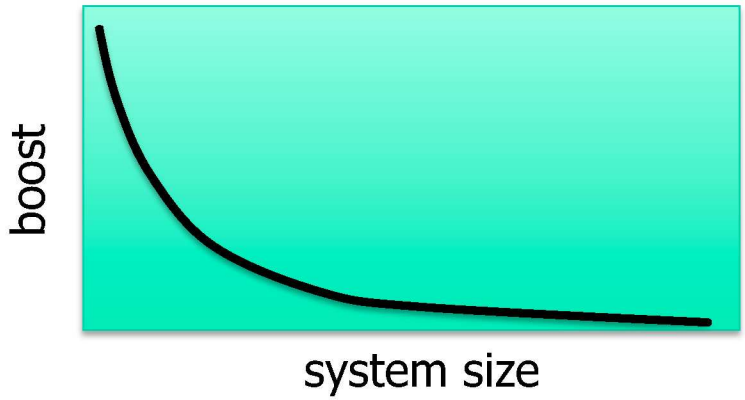

All the AMD methods in their primitive form show bad scaling with system size

Local Hyperdynamics

S.Y. Kim, D. Perez, and AFV, J. Chem. Phys. 139, 144110 (2013).

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

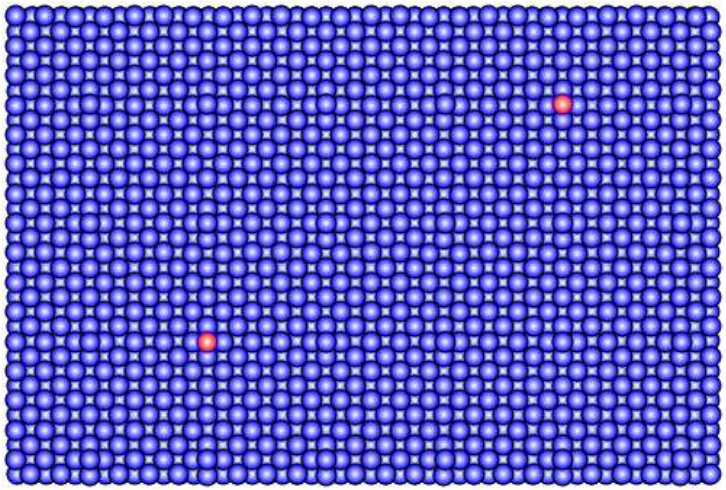

Key concept: Most systems we are interested in are intrinsically local in their behavior. A transition, or near-transition, in one region of system should not have any significant effect on atoms that are far away.

Local Hyperdynamics

S.Y. Kim, D. Perez, and AFV, J. Chem. Phys. 139, 144110 (2013).

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

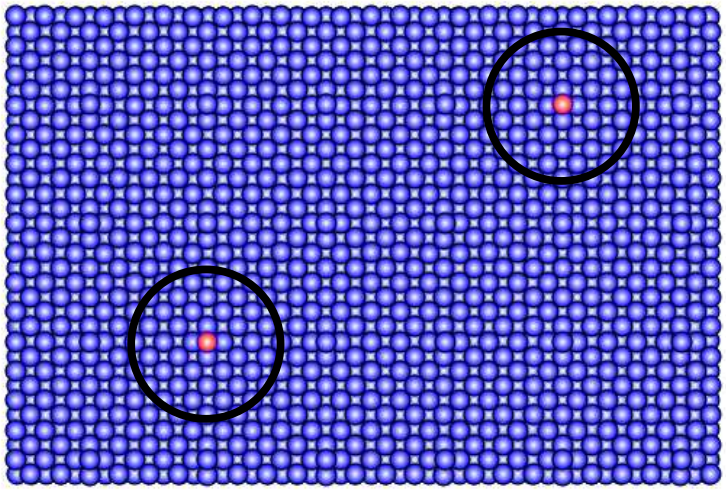

Key concept: Most systems we are interested in are intrinsically local in their behavior. A transition, or near-transition, in one region of system should not have any significant effect on atoms that are far away.

Local hyperdynamics – procedure

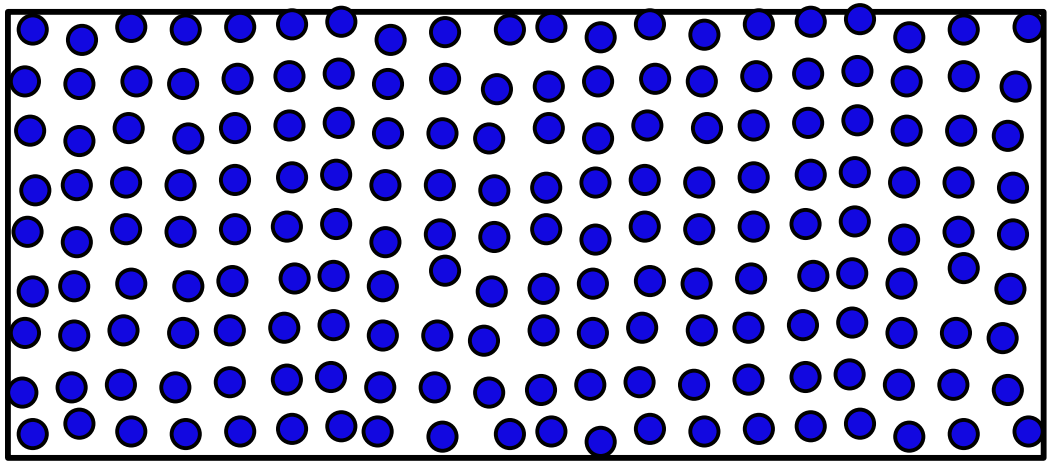

Each bond (i) has its own local domain (I).

Local hyperdynamics – procedure

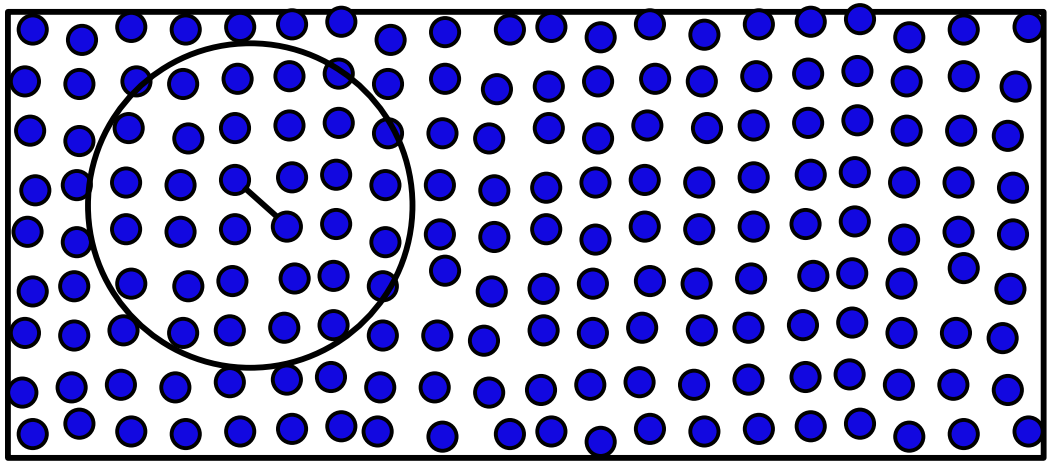

Each bond (i) has its own local domain (I).

Local hyperdynamics - procedure

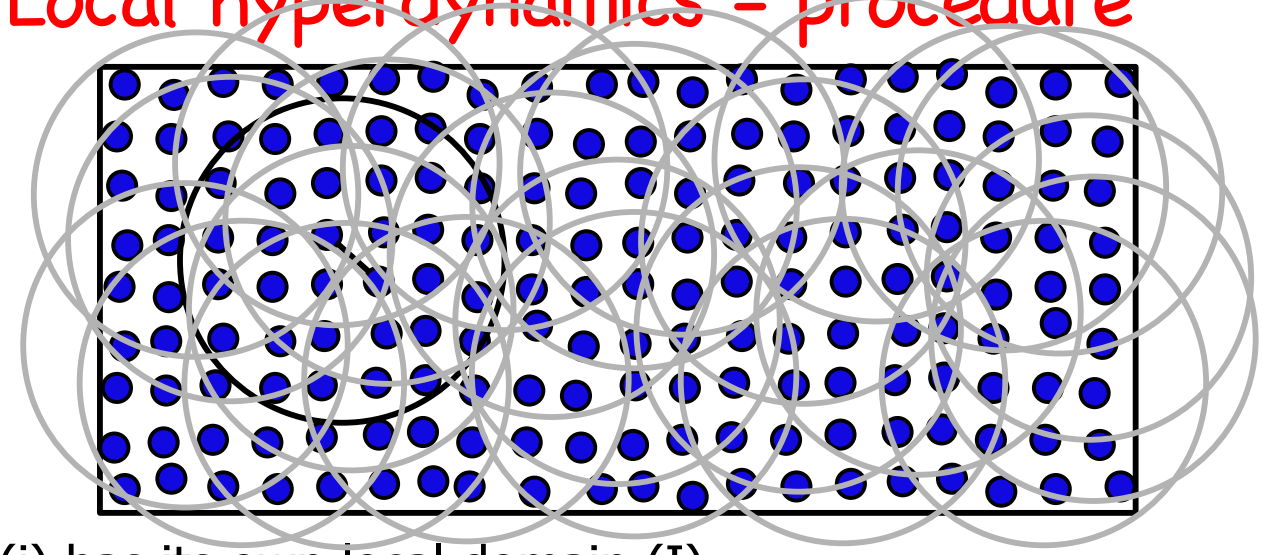

Each bond (i) has its own local domain (I).

Local hyperdynamics – procedure

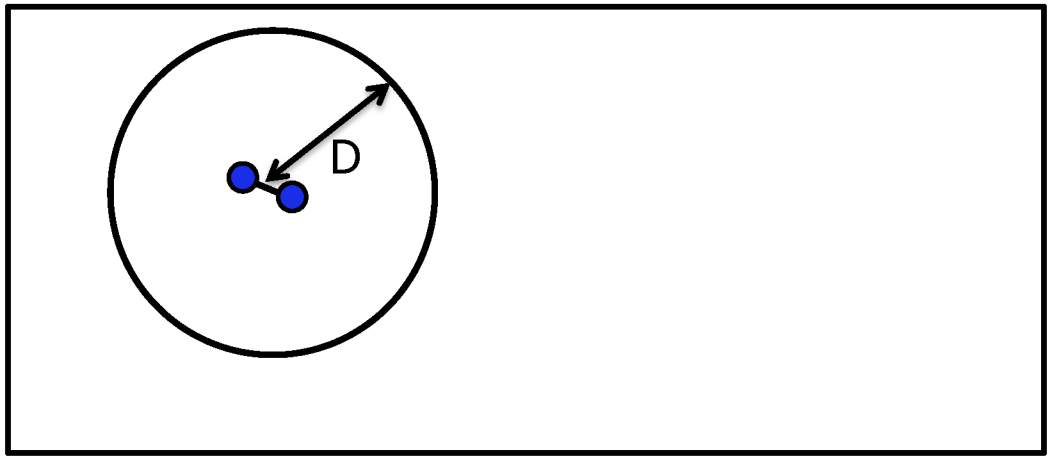

Each bond (i) has its own local domain (I).

Local hyperdynamics - procedure

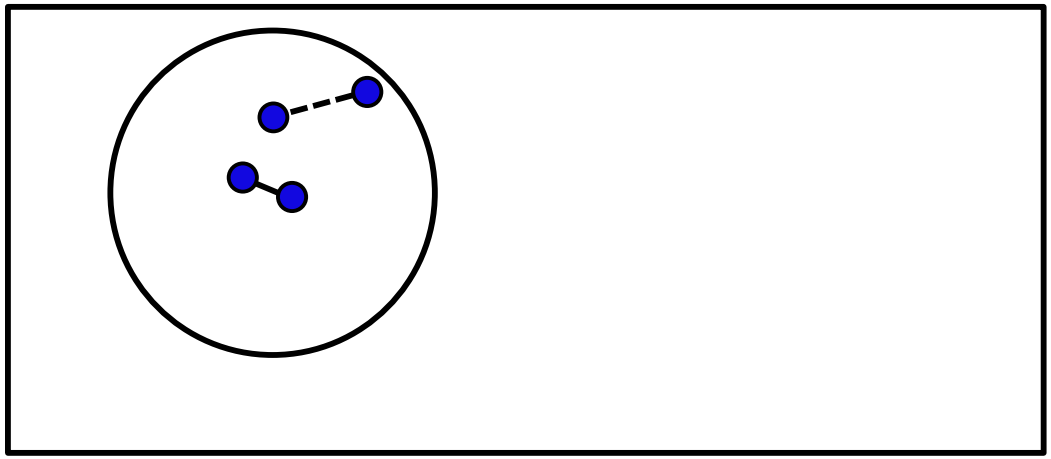

Each bond (i) has its own local domain (I).

Each domain has its own bias energy $\Delta V_I = C_I \Delta V_{bias}(R_I(t))$ and boost factor B_I based on the geometry R_I within the domain (exactly like a global hyperdynamics in that domain).

Local hyperdynamics - procedure

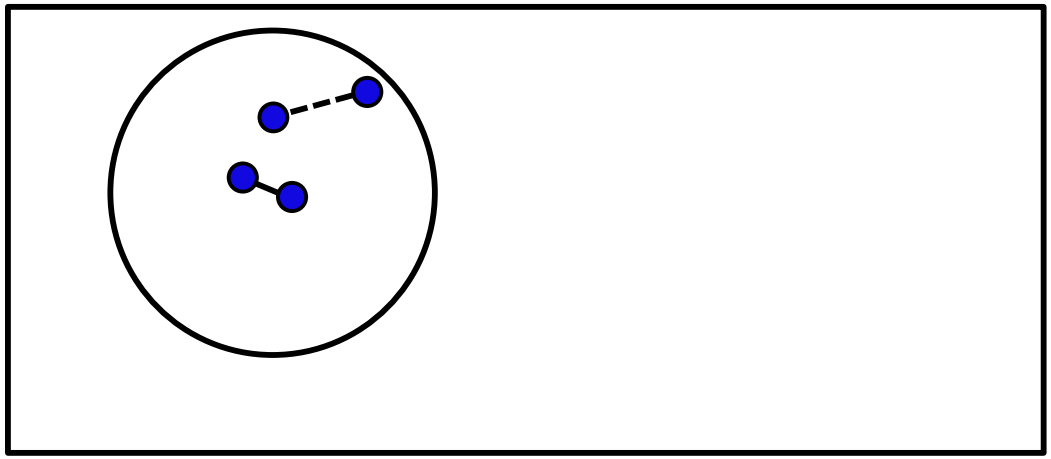

Each bond (i) has its own local domain (I).

Each domain has its own bias energy $\Delta V_I = C_I \Delta V_{bias}(R_I(t))$ and boost factor B_I based on the geometry R_I within the domain (exactly like a global hyperdynamics in that domain).

The bias energy in domain I is used to determine the force only on bond i. E.g., with a simple bond-boost bias there is a force on a bond if and only if it is the most distorted bond in its own domain.

Local hyperdynamics - procedure

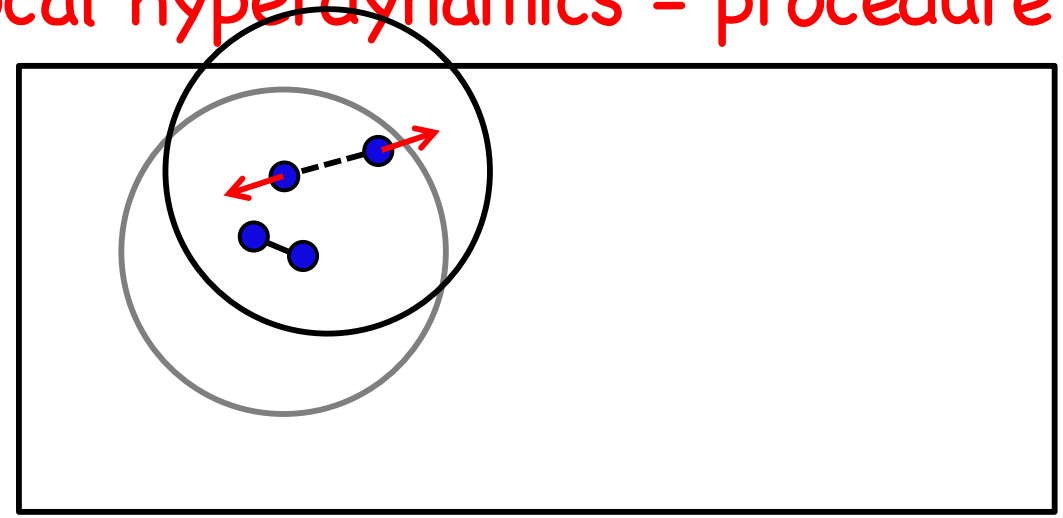

Each bond (i) has its own local domain (I).

Each domain has its own bias energy $\Delta V_I = C_I \Delta V_{bias}(R_I(t))$ and boost factor B_I based on the geometry R_I within the domain (exactly like a global hyperdynamics in that domain).

The bias energy in domain I is used to determine the force only on bond i. E.g., with a simple bond-boost bias there is a force on a bond if and only if it is the most distorted bond in its own domain.

Local hyperdynamics - procedure

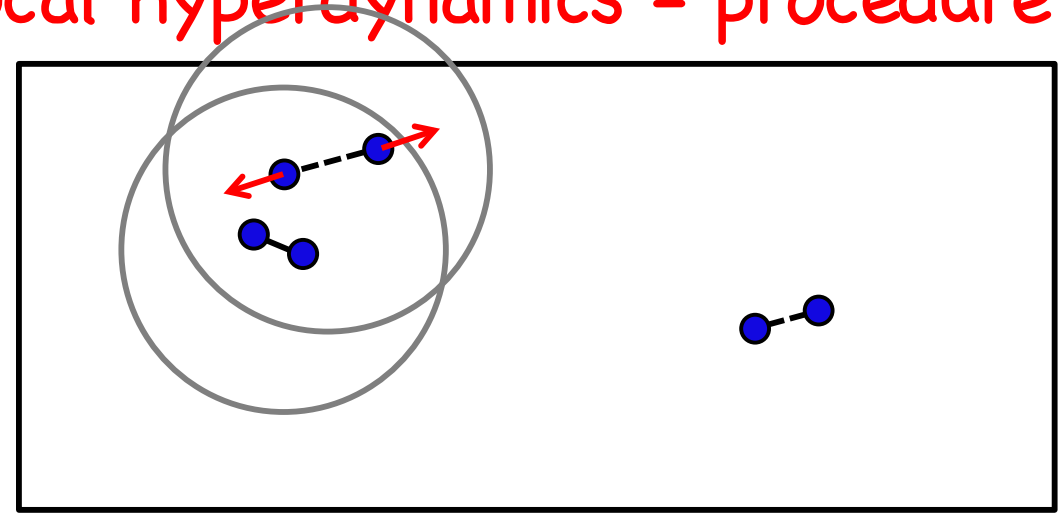

Each bond (i) has its own local domain (I).

Each domain has its own bias energy $\Delta V_I = C_I \Delta V_{bias}(R_I(t))$ and boost factor B_I based on the geometry R_I within the domain (exactly like a global hyperdynamics in that domain).

The bias energy in domain I is used to determine the force only on bond i. E.g., with a simple bond-boost bias there is a force on a bond if and only if it is the most distorted bond in its own domain.

Local hyperdynamics - procedure

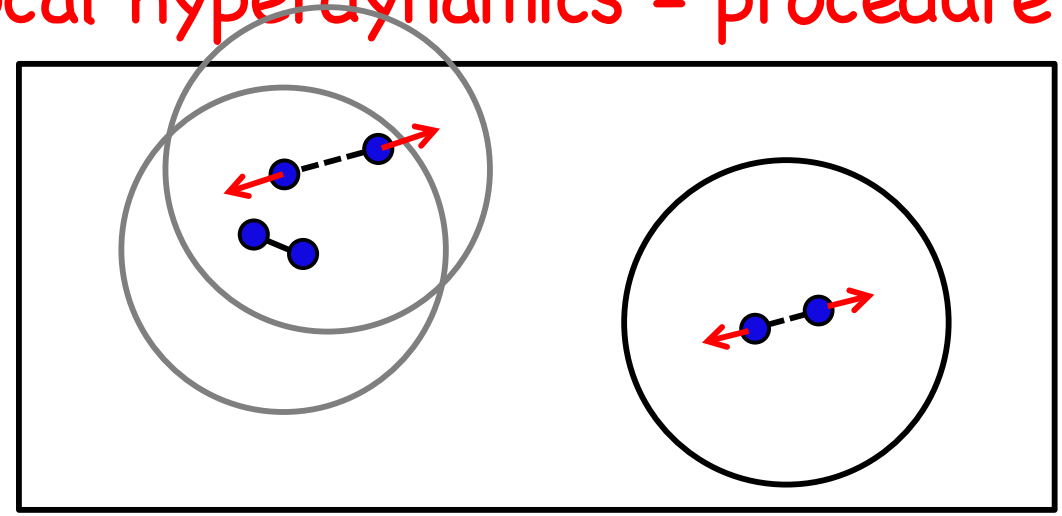

Each bond (i) has its own local domain (I).

Each domain has its own bias energy $\Delta V_I = C_I \Delta V_{bias}(R_I(t))$ and boost factor B_I based on the geometry R_I within the domain (exactly like a global hyperdynamics in that domain).

The bias energy in domain I is used to determine the force only on bond i. E.g., with a simple bond-boost bias there is a force on a bond if and only if it is the most distorted bond in its own domain.

Local hyperdynamics – procedure

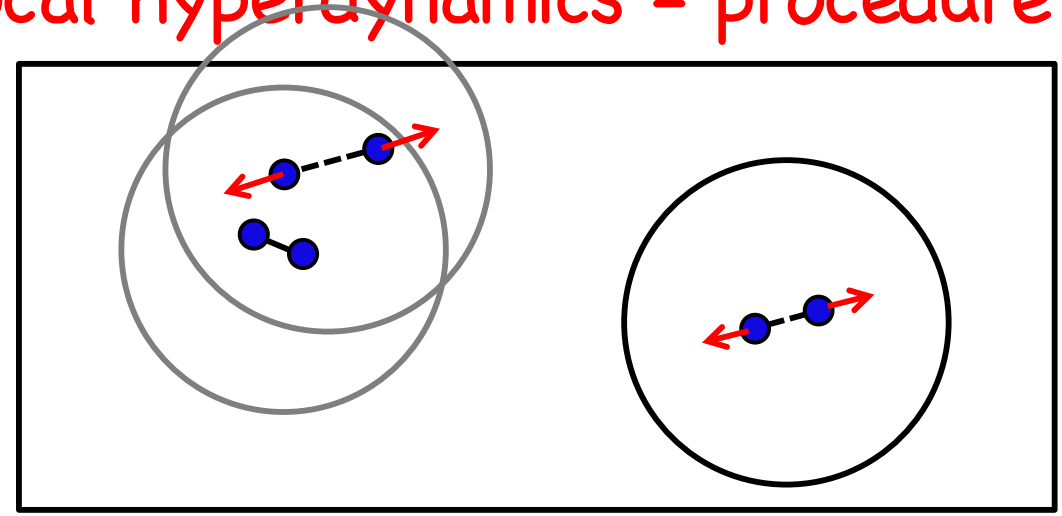

Each bond (i) has its own local domain (I).

Each domain has its own bias energy $\Delta V_I = C_I \Delta V_{bias}(R_I(t))$ and boost factor B_I based on the geometry R_I within the domain (exactly like a global hyperdynamics in that domain).

The bias energy in domain I is used to determine the force only on bond i. E.g., with a simple bond-boost bias there is a force on a bond if and only if it is the most distorted bond in its own domain.

A domain-bias multiplier (C_I) is adjusted for each domain to make its average boost $\langle B_I \rangle$ match the *target boost factor* B_{target}.

The boostostat

Instead of pre-adjusting the $\{C_I\}$ values for the current state, we have found that we can simply apply a "boostostat" to gently but constantly push on these $\{C_I\}$ values during the simulation to move the boost for each domain towards the correct target boost.

For each domain I at each MD step:

$$C_I(t+\Delta t) = C_I(t) - \alpha_B \Delta t_{MD}[B_I(t) - B_{target}]/B_{target}$$

where α_B = boostostat coupling strength ($\sim 10^9 - 10^{10}\ s^{-1}$)

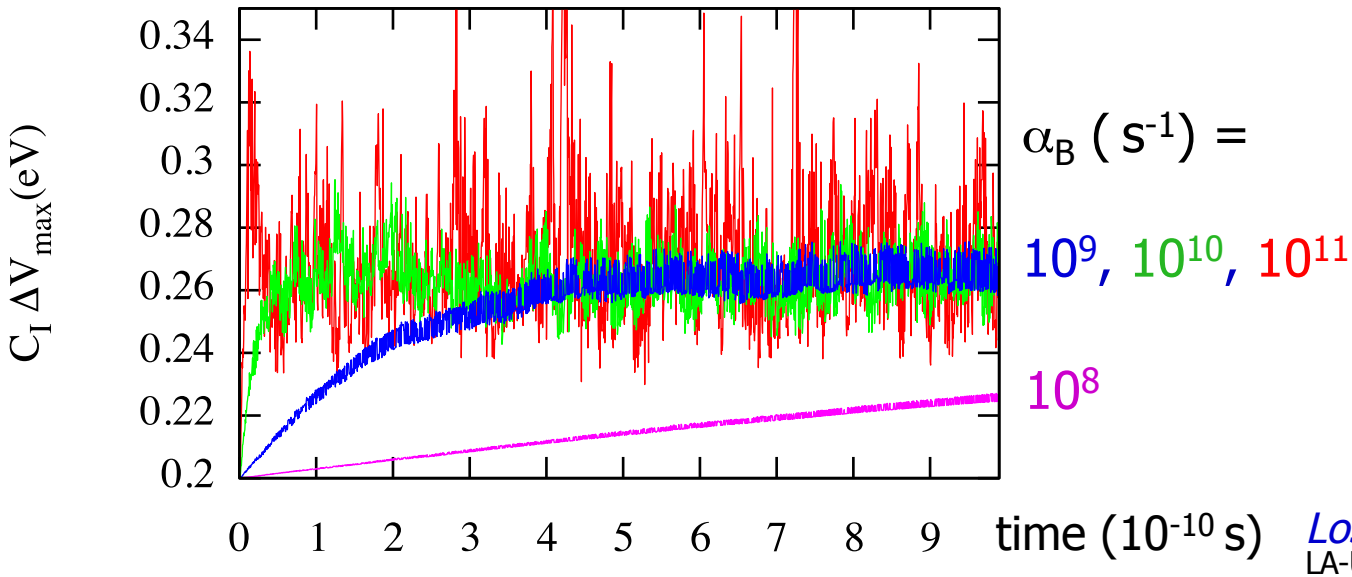

Note - the dynamics are not conservative

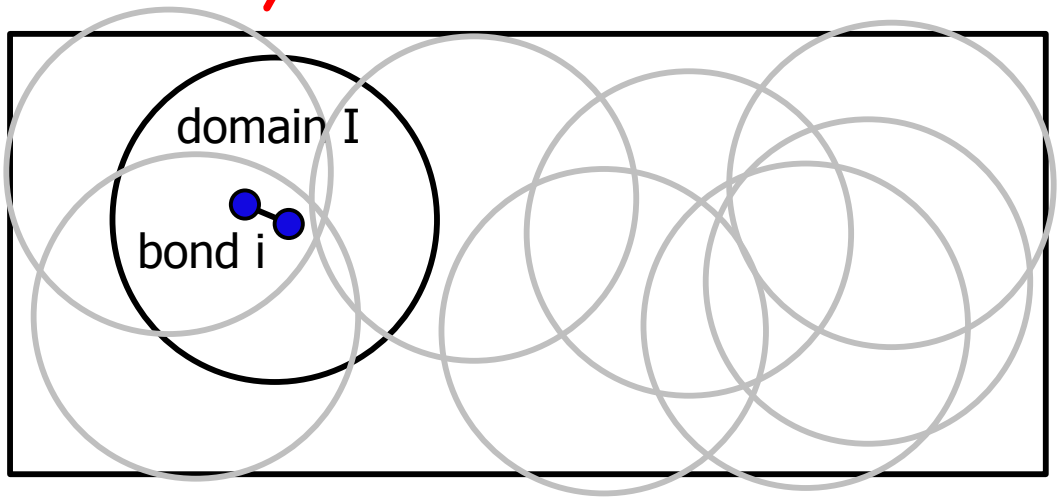

The force on bond i (the center of domain I) is taken to be

$$f_i = -\partial \Delta V_I / \partial x_i \ .$$

However, the force on nearby bond j is given by

$$f_j = -\partial \Delta V_J / \partial x_j \ ,$$

so this is not conservative dynamics.

Local hyperdynamics – scaling

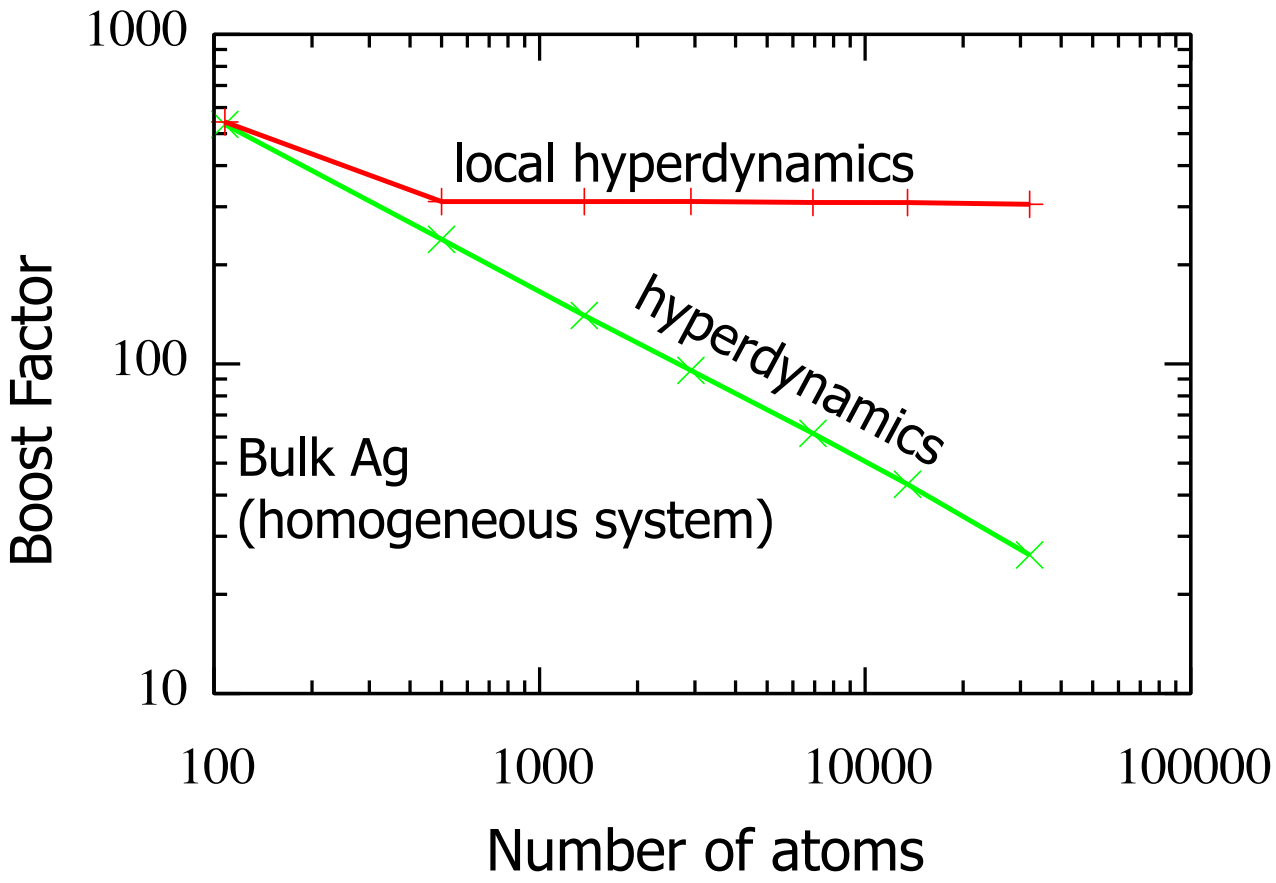

accuracy of the rates

Local hyperdynamics – rate tests

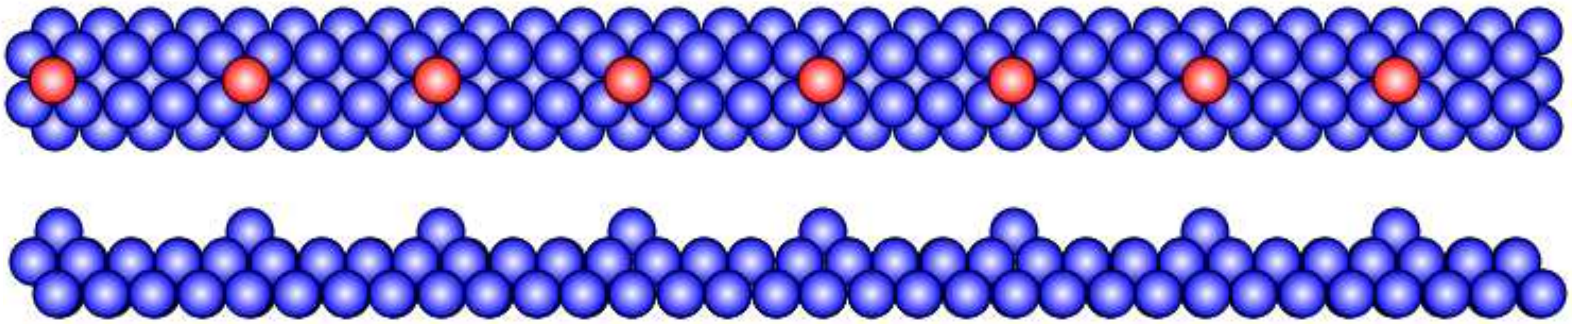

Narrow strip of Ag(100), 8 adatoms, top layer and adatoms free to move (72 moving atoms), periodic in *x*, EAM potential. Langevin friction = 10^{12} s^{-1}.

Small enough to run fast, but large enough and complex enough to test method.

Both hops and exchanges can occur.

Transitions observed and then rejected so the rates stay constant.

Local hyperdynamics – rate tests

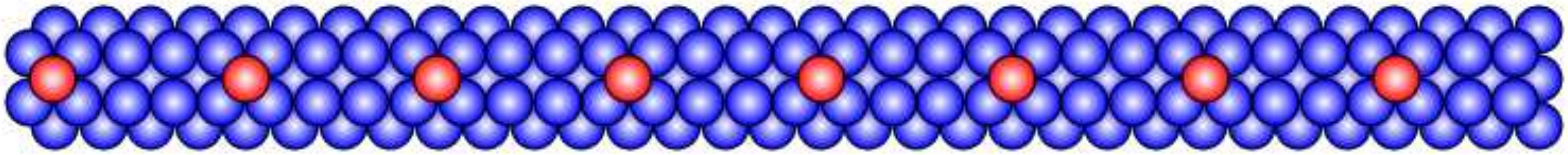

Hop pathway: Ea=0.504 eV, ν_0=3.53x10^{12} Hz

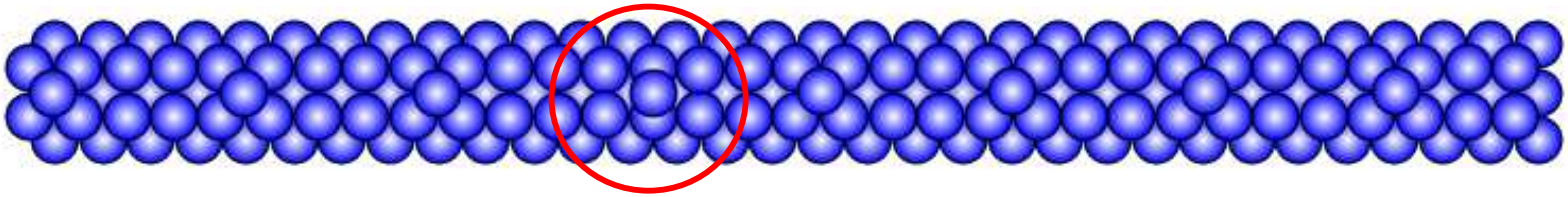

Exchange pathway: Ea=0.651 eV, ν_0=6.48x10^{12} Hz

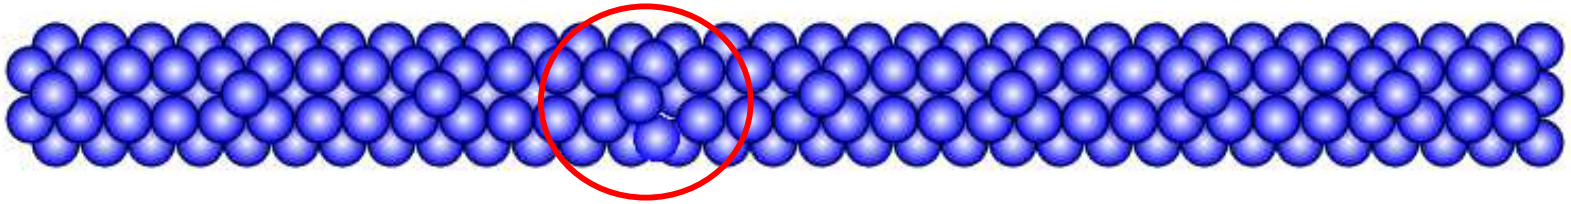

Side-hop pathway: Ea=0.743 eV, ν_0=3.02x10^{12} Hz

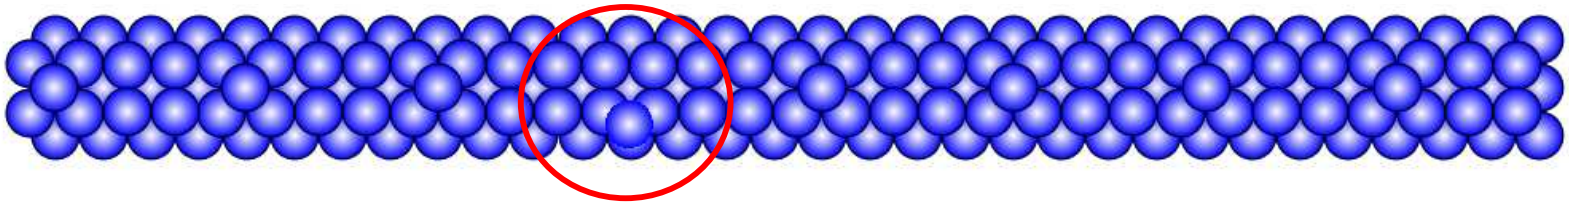

Local hyperdynamics test on strip system

T=500K, B_{target}=100, α_B=2X10^{10}, range D=10 Å

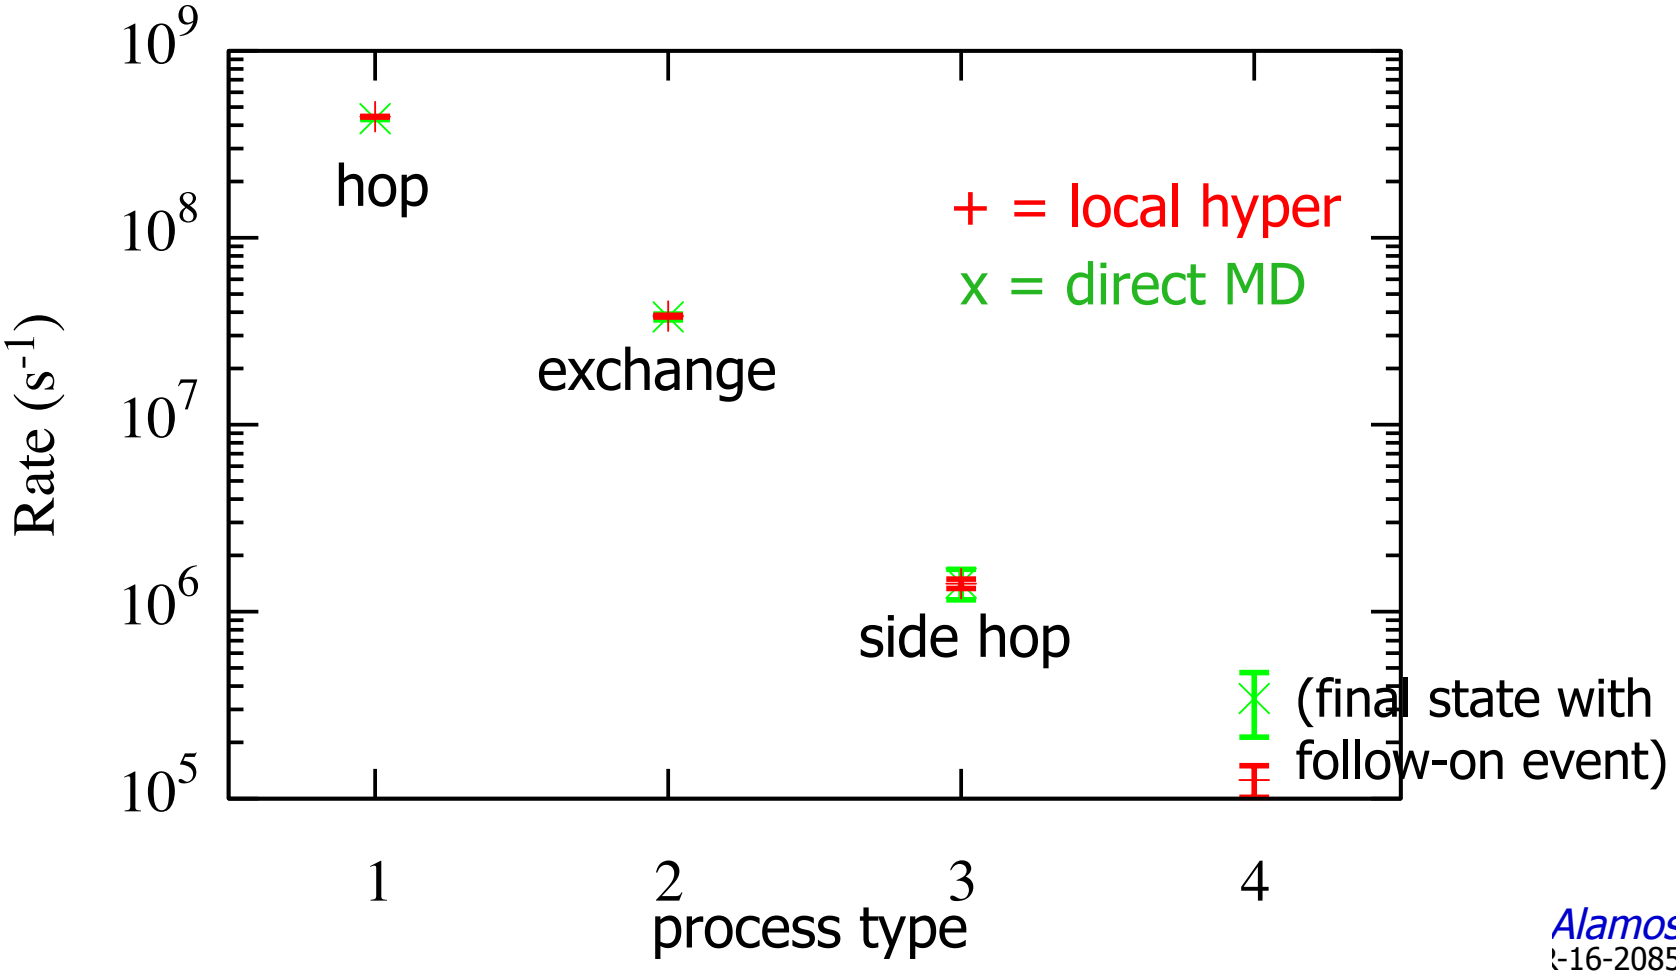

Rate accuracy w/ boostostat coupling strength

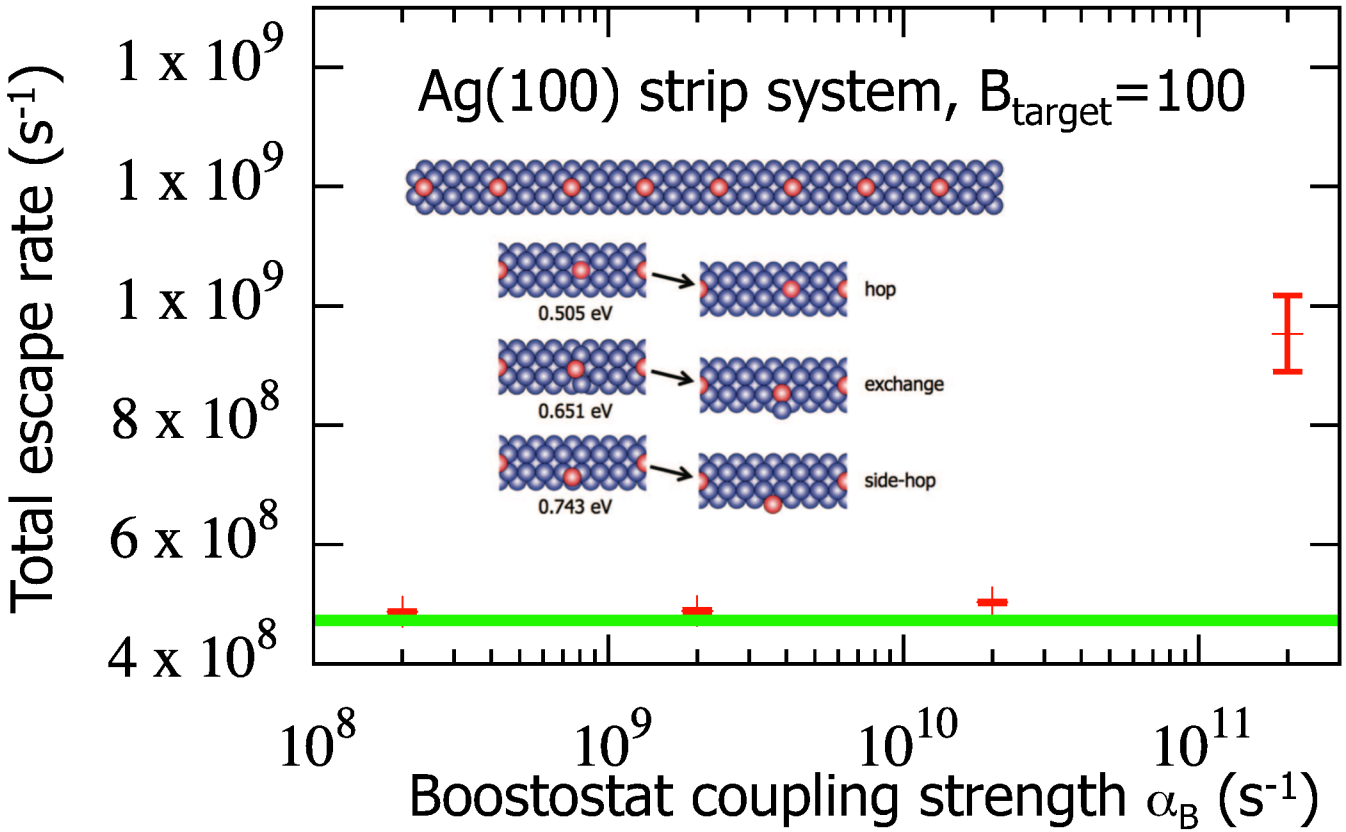

Boostostatting through a transition: Ag(100) test system

- Defective Ag surface: adatom, adatom dimer, vacancy
- 434 atoms, 218 moving
- T=325K, γ=5x10^{11} s^{-1}
- Locality radius D=10 Å
- On-the-fly boostostatting, α_B=2x10^{10} s^{-1} B_{target}=100
- Coefficients {C_I} reset locally to 0.2 eV after each transition
- About 25 thermally relevant transitions. Rates vary over three orders of magnitude.

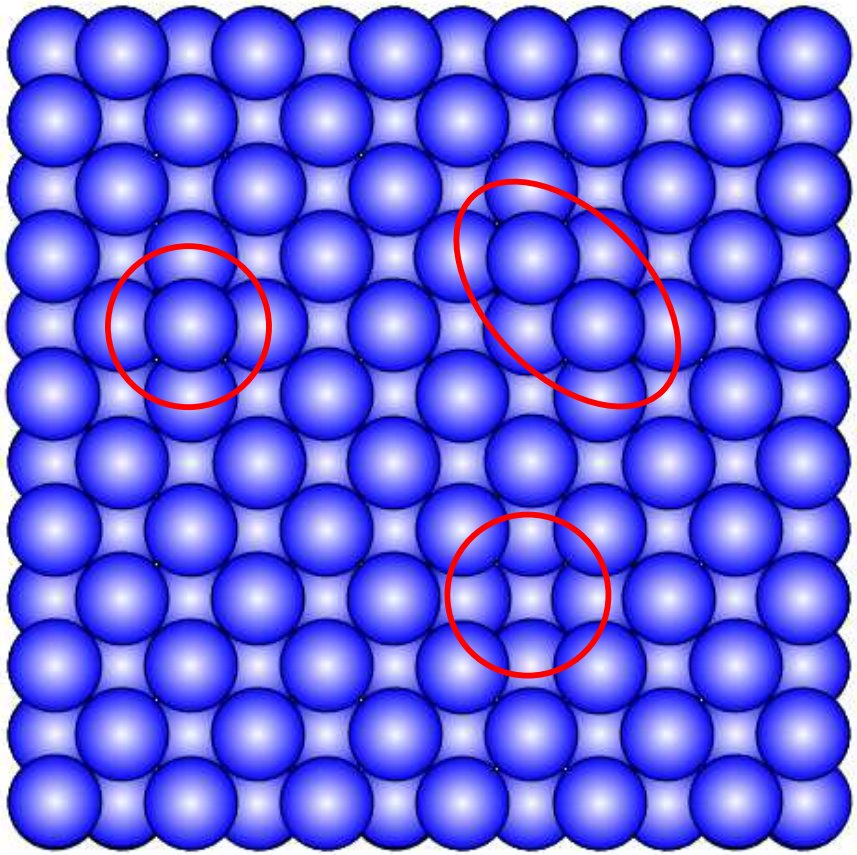

Local Hyperdynamics - Ag(100) tests

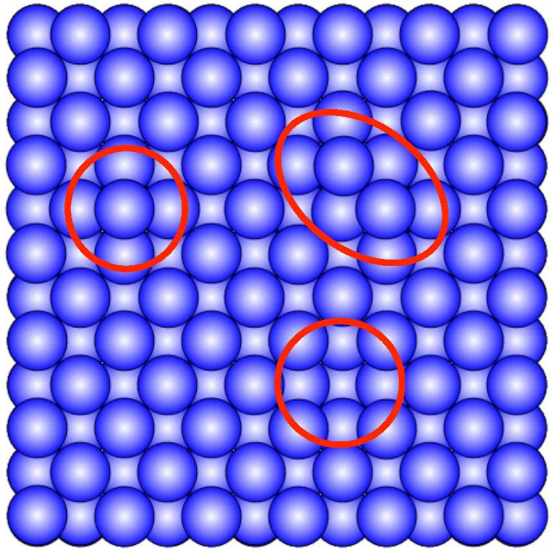

Near perfect agreement with exact results

Alamos
R-16-20856

Local Hyperdynamics - Ag(100) tests

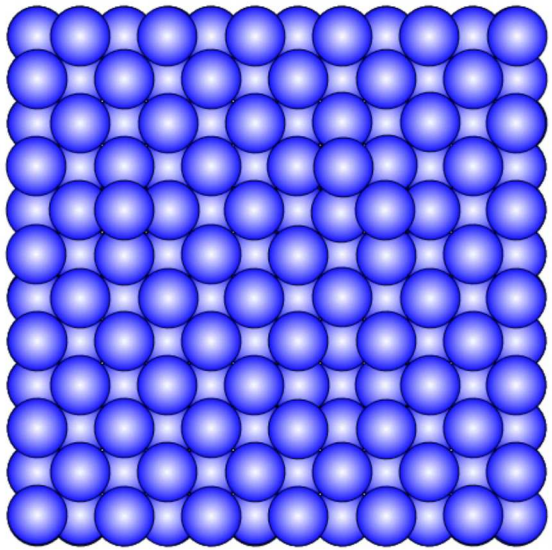

Target boost=100
T=325K

vacancy hops

adatom hops
dimer hops

exchanges and multi-atom events

(Results do not change with system size.)

+ = local hyperdynamics
x = global hyperdynamics

Rate (s^{-1})

10^7
10^6
10^5
10^4
10^3

0 5 10 15 20 25

Process Index

Near perfect agreement with exact results

Alamos
R-16-20856

Local hyperdynamics – could it be exact?

We can show that local hyperdynamics should give increasingly accurate results as the local bias range is increased, for any proper form of bias potential.

For the simple bond-boost bias potential:

We can show that for a homogeneous system (all bonds the same) the force "errors" arising from the non-conservative dynamics "cancel" (as discussed next).

For a non-homogeneous system, it does not seem that there should be exact cancellation.

Yet, every test we have done shows very high accuracy, making us suspect there may be something more general that can be derived about this kind of dynamics.

The force mismatch terms

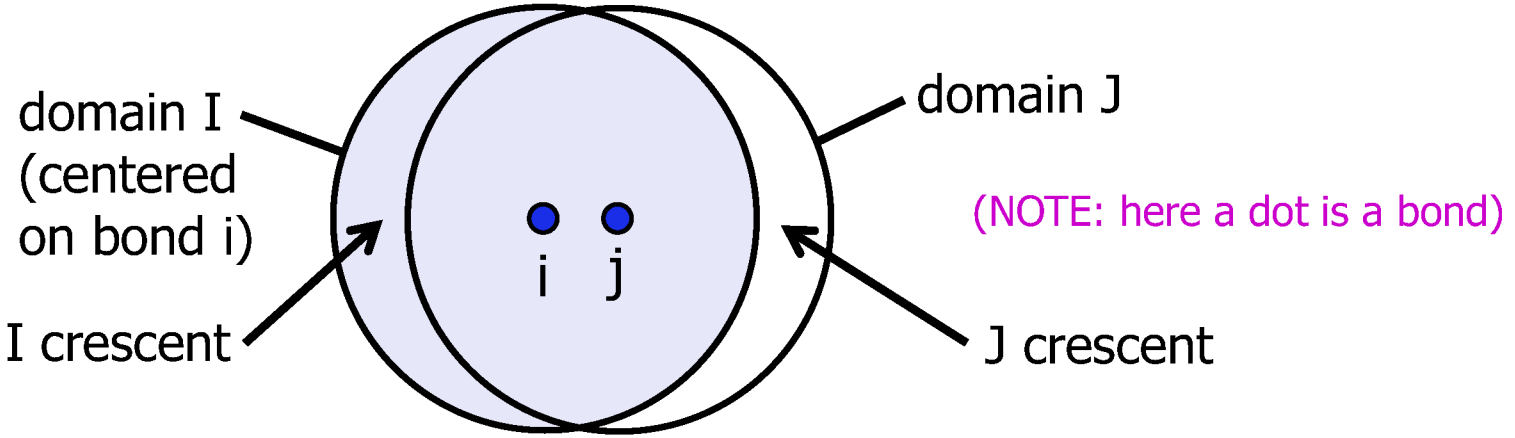

The force on bond j, which is based on domain J

$$f_j = -\partial \Delta V_J / \partial x_j$$

is different than the force that bond i (domain I) *thinks* that bond j will have

$$f_j^{(I)} = -\partial \Delta V_I / \partial x_j .$$

The size of this force mismatch ("error") depends on the "I crescent" and the "J crescent".

The force mismatch terms

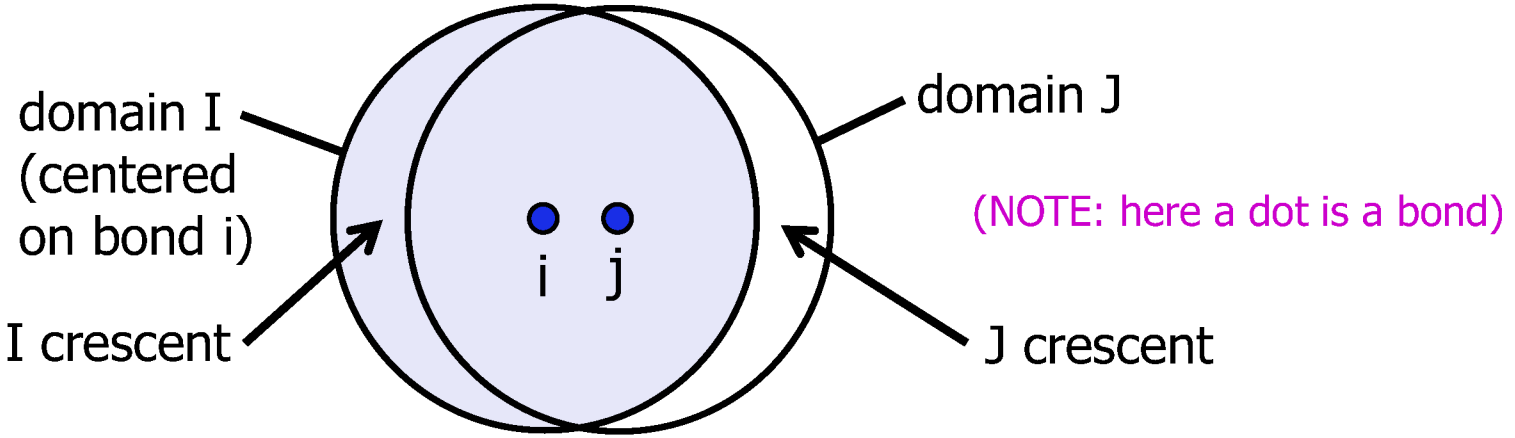

The force on bond j, which is based on domain J

$$f_j = -\partial \Delta V_J / \partial x_j$$

is different than the force that bond i (domain I) *thinks* that bond j will have

$$f_j^{(I)} = -\partial \Delta V_I / \partial x_j .$$

The size of this force mismatch ("error") depends on the "I crescent" and the "J crescent".

The bonds in the crescents are beyond the interaction range of bond i and bond j. If all bonds are equivalent, these force mismatches from the domain edges come and go in a way that is simply a reordering in time relative to what they would be if domain I was the same as domain J.

The force mismatch terms

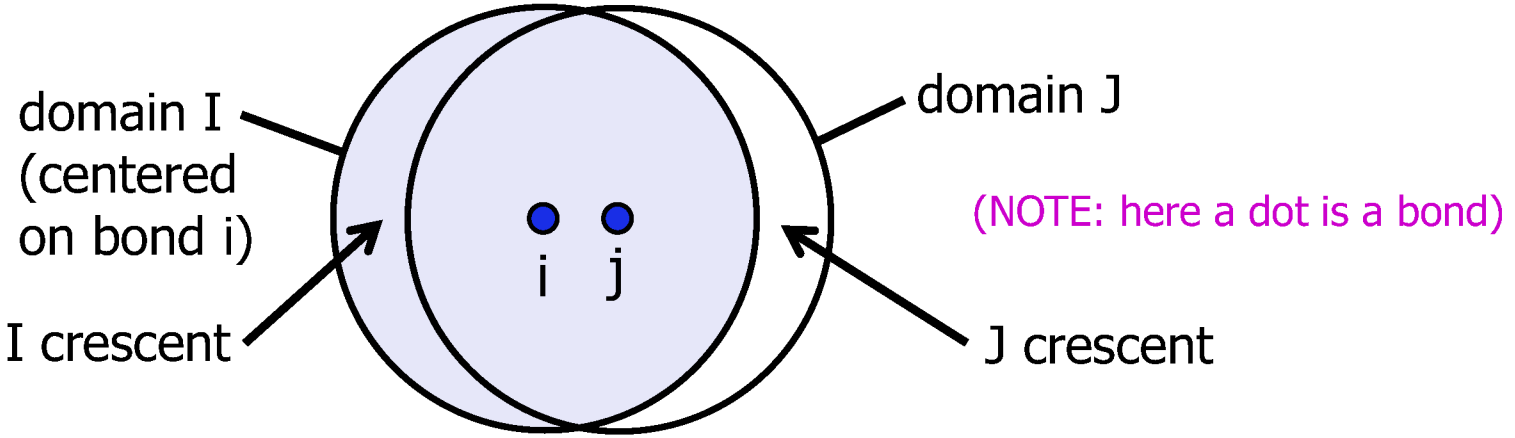

The force on bond j, which is based on domain J

$$f_j = -\partial \Delta V_J / \partial x_j$$

is different than the force that bond i (domain I) *thinks* that bond j will have

$$f_j^{(I)} = -\partial \Delta V_I / \partial x_j .$$

The size of this force mismatch ("error") depends on the "I crescent" and the "J crescent".

The bonds in the crescents are beyond the interaction range of bond i and bond j. If all bonds are equivalent, these force mismatches from the domain edges come and go in a way that is simply a reordering in time relative to what they would be if domain I was the same as domain J.

The domain crescent cases when $C_I=C_J$

(for simple bond boost bias)

case

1 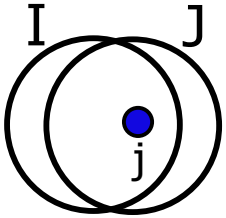

Bond j is the most distorted bond in domain I+J
→ exact ($f_j^{(I)}=f_j$)

2 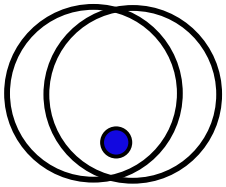

Some other bond k in I∧J is the most distorted in domain I+J
→ exact ($f_j^{(I)}=f_j=0$)

3 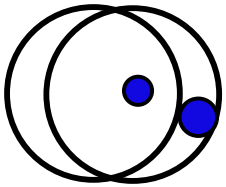

Bond j is the most distorted in domain I, but there is an even more distorted bond in the J crescent (that j sees but i does not see)
→ error ($f_j^{(I)}\neq 0$, $f_j=0$)

4 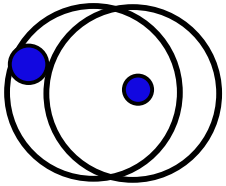

Bond j is the most distorted in domain J, but there is an even more distorted one in the I crescent (that i sees but j does not see)
→ error ($f_j^{(I)}=0$, $f_j\neq 0$)

5 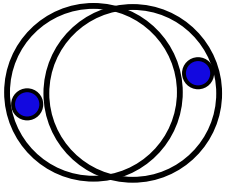

The most distorted bond in domain I is in the I crescent and most distorted bond in domain J is in the J crescent
→ exact ($f_j^{(I)}=f_j=0$)

The domain crescent cases when $C_I=C_J$

(for simple bond boost bias)

case

1 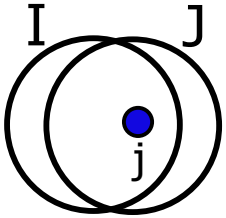

Bond j is the most distorted bond in domain I+J
→ exact ($f_j^{(I)}=f_j$)

2 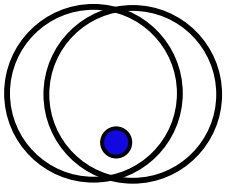

Some other bond k in I∧J is the most distorted in domain I+J
→ exact ($f_j^{(I)}=f_j=0$)

3 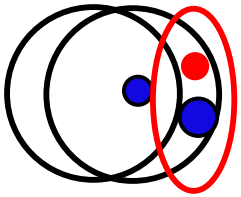

Bond j is the most distorted in domain I, but there is an even more distorted bond in the J crescent (that j sees but i does not see)
→ "error" ($f_j^{(I)}\neq 0$, $f_j=0$)

4 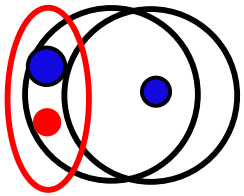

Bond j is the most distorted in domain J, but there is an even more distorted one in the I crescent (that i sees but j does not see)
→ error ($f_j^{(I)}=0$, $f_j\neq 0$)

5 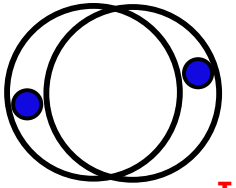

The most distorted bond in domain I is in the I crescent and most distorted bond in domain J is in the J crescent
→ exact ($f_j^{(I)}=f_j=0$)

If I-crescent and J-crescent are statistically equivalent,
then $C_I = C_J$, and error 3 and error 4 "cancel" in a time average.

The core argument – time reordering

If D > 2L (L=local range of interaction), then all bonds in both crescents are more than a distance D away from bond *i* or bond *j*.

Large bond distortions come and go in the I crescent, dictating, from a large distance away, what bond *i* thinks the force on bond *j* should be. Meanwhile, the *actual force* on bond *j* is dictated by the (distant) bonds in crescent J.

Because the bonds in crescent J are identical to the bonds in crescent I, the bond-length fluctuations in crescent J that can shut down bond *j* are also identical, other than a *random reordering in time,* to the bond-length fluctuations in crescent I.

A force-expectation discrepancy can arise only when there is no force on bond *i.* Thus, the abrupt turning-on and turning-off of the force on bond j, which also happens in regular hyperdynamics (though less often), will still give appropriate Langevin evolution of bond *j* from the point of view of bond *i,* whether or not the sequence is reordered in time. Bond *i* has no way of telling the difference. *Locally*, then, for all *j* bonds within a range L=D/2 of bond *i,* the evolution is equivalent to true Langevin evolution on $V+\Delta V$, and should give the same boost factor for that region as a global bias on domain I would give. And this same statement can be made for every bond *i* in the system.

Probing for accuracy problems in local hyper

We have found that it is difficult to even create a test case that can clearly distinguish right from wrong (e.g., to use in developing a more accurate variation on local hyperdynamics).

We have tested various inhomogeneous systems and find the results are still very accurate.

The errors we do see are so small that they might even be caused by setting the range D too small, or setting the target boost too high.

Conclusions – Local Hyperdynamics

A new, local formulation of hyperdynamics makes an advance on the size-scale problem in accelerated molecular dynamics methods.

Scalable, with constant boost, to arbitrarily large systems, provided the lowest barrier in system does not decrease with increasing system size.

We understand why it should be correct for homogeneous systems.

We are surprised how accurate it is for *every* system we have tried, homogeneous or not. Perhaps there is something deeper going on…

S.Y. Kim, D. Perez, and AFV, J. Chem. Phys. 139, 144110 (2013).