Parallel Algorithms for Hyperdynamics in LAMMPS

SAND2018-2542C

Steve Plimpton (Sandia)
Danny Perez and Art Voter (LANL)

TMS 2018
Symposium on Algorithm Development in Materials Science and Engineering
March 2018 - Phoenix, AZ

CCR
Center for Computing Research

Sandia National Laboratories

Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. Presentation: SAND2017-5724C

NNSA
National Nuclear Security Administration

What is hyperdynamics (HD)

- **Accelerated time method** for MD
  - *Voter, J Chem Phys, 106, 4665 (1997)*
  - bias the PE surface to enable more rapid transitions
  - time-accurate speed-up of a single trajectory
  - not an enhanced sampling or multi-replica approach
- **Bond-boost** formulation
  - *Miron & Fichthorn, J Chem Phys, 119, 6210 (2003)*
  - variant: apply bias to only one or handful of bonds
- **Local** hyperdynamics
  - *Kim, Perez, Voter, J Chem Phys 139, 144110 (2013)*
  - global: bias one bond in entire system at a time
  - local: bias multiple bonds separated by $D_{cut} = 10$ Å
  - tested correctness for simple systems
  - accelerated event rates match theory and experiment

What kind of systems can benefit from HD

- Key **requirements**:
  - solids with distinct energy basins (not soft matter)
  - equilibrium MD with rare transitions from one basin to another
- Effective speed-up can be **orders of magnitude**
  - especially for high barriers and low temperatures

What kind of systems can benefit from HD

- Key **requirements**:
  - solids with distinct energy basins (not soft matter)
  - equilibrium MD with rare transitions from one basin to another
- Effective speed-up can be **orders of magnitude**
  - especially for high barriers and low temperatures

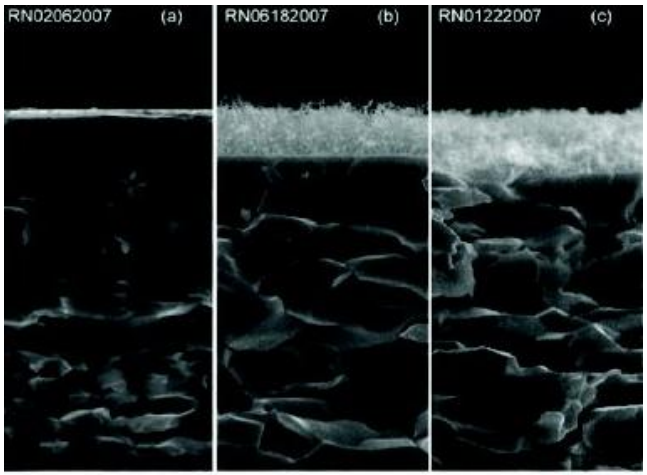

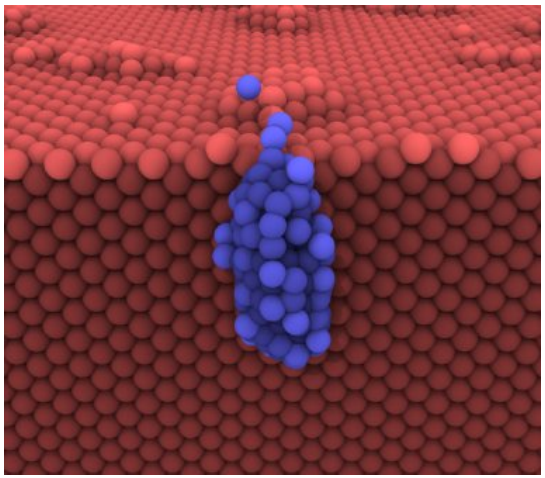

- **ITER**: degradation of W surface exposed to He plasma
- Due to sub-surface He bubbles producing nanoscale **fuzz**
- Modeling as part of DOE exascale computing project (ECP)

Pictorial view of hyperdynamics

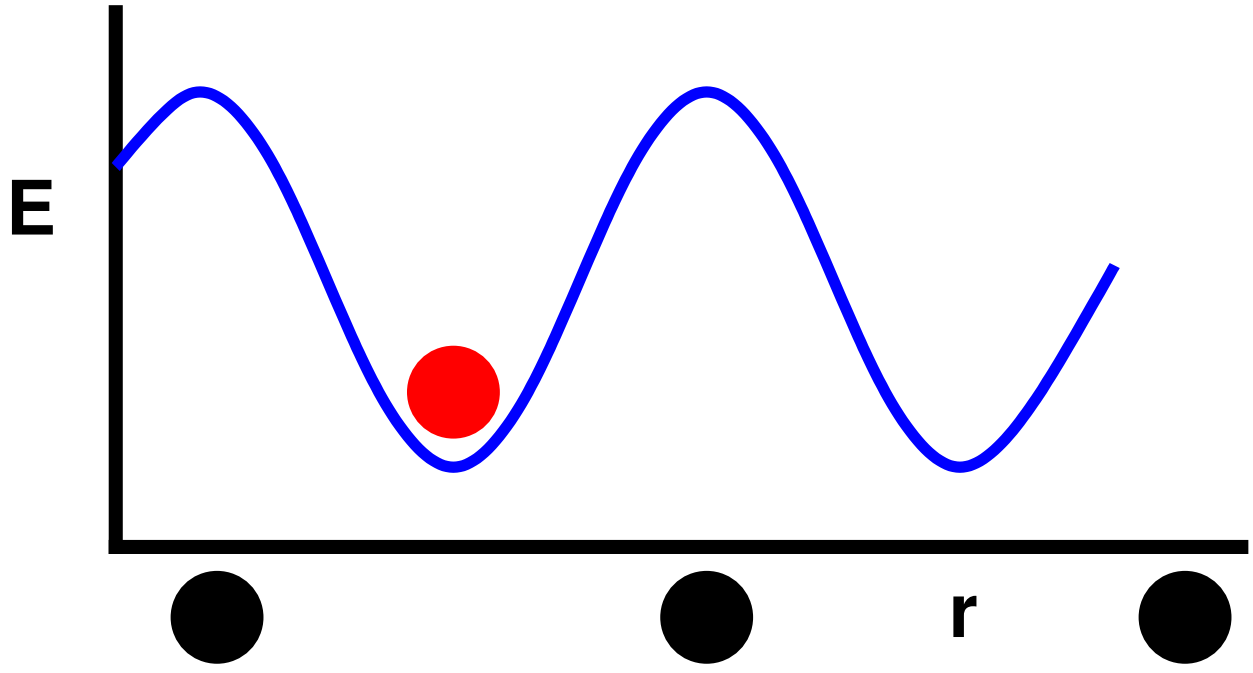

- Corrugated energy landscape for **adatom surface diffusion**
- Define (conceptual) **bonds** between all pairs of nearby atoms
  - e.g. $\sim$12 nearest neighbors per atom in fcc lattice

Pictorial view of hyperdynamics

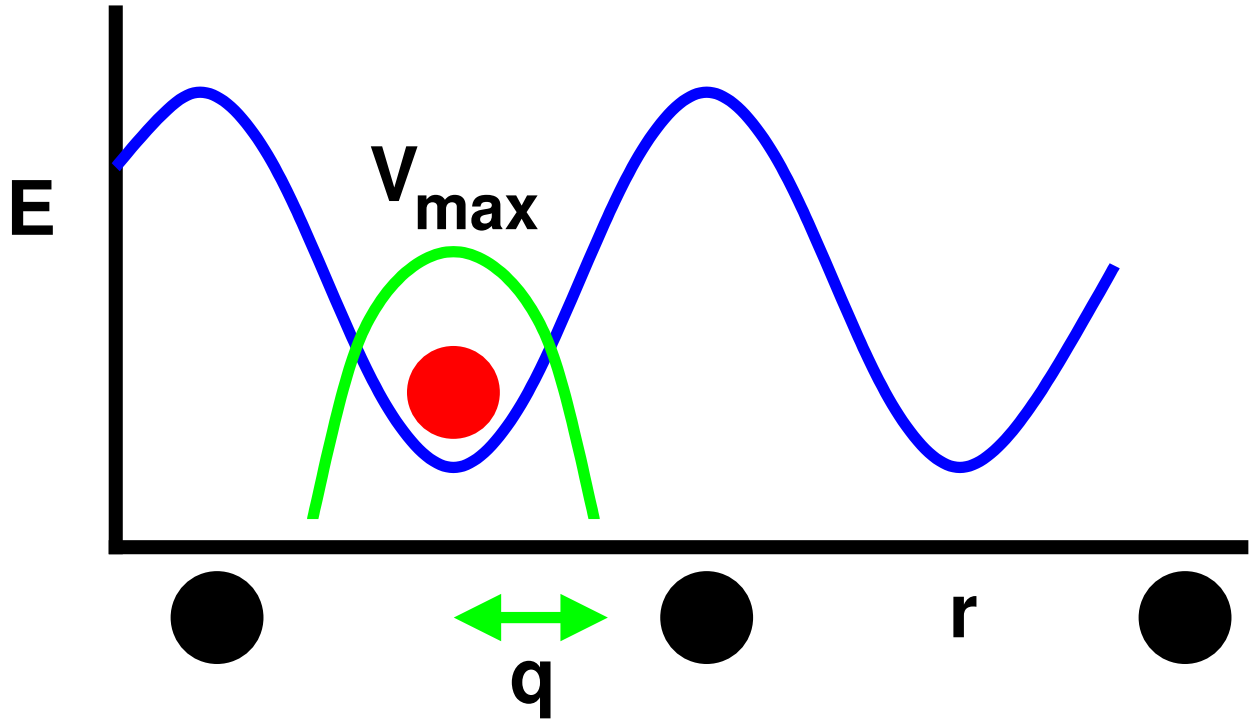

- Add **bias potential** to (only) selected bonds
- Bond strain: $\epsilon_{ij} = (R_{ij} - Ro_{ij})/Ro_{ij}$
- Bias: $V_{ij} = V_{max}[1 - (\epsilon_{ij}/q)^2], \quad |\epsilon_{ij}| < q, \quad \text{else} \mathbf{zero}$
- **Different bond** may be biased at each timestep

Pictorial view of hyperdynamics

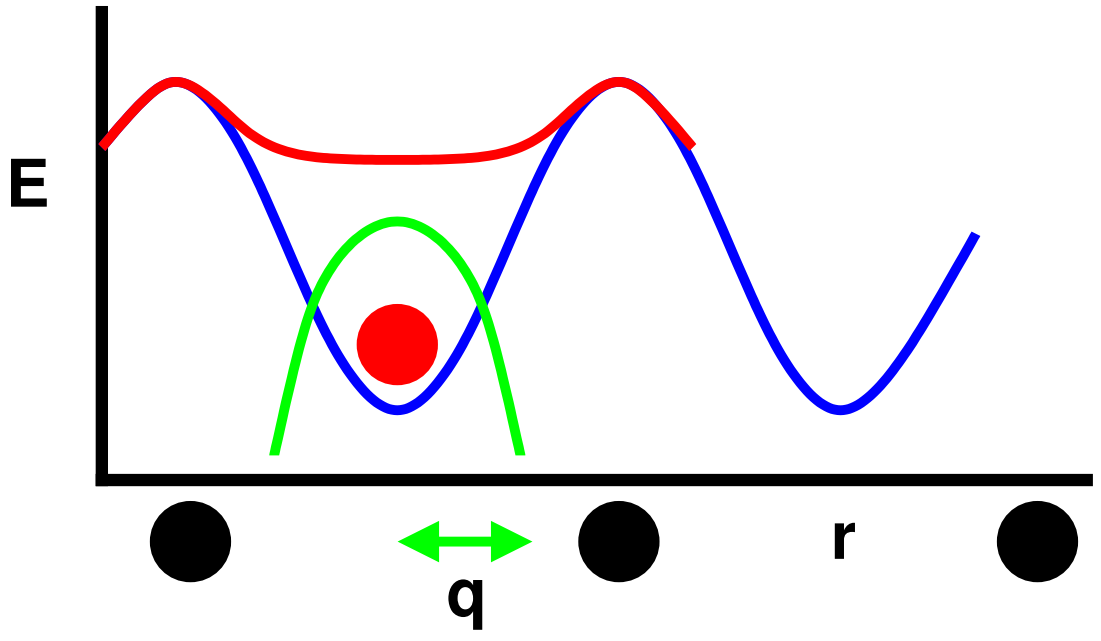

- Shallow well $\Rightarrow$ **faster transition** by I,J (and nearby) atoms

Pictorial view of hyperdynamics

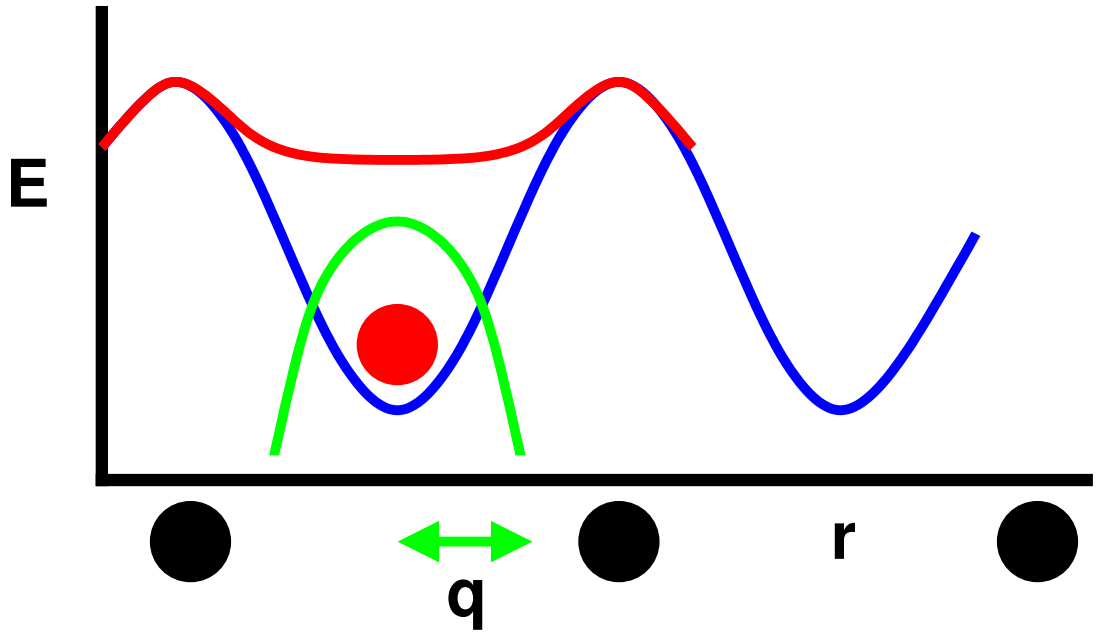

- Shallow well $\Rightarrow$ **faster transition** by I,J (and nearby) atoms
- Must choose V_{max} and q carefully:
  - if: zero bias at dividing surfaces (Q), no local minima (V_{max})
  - if: do not induce TST-violating correlated events
  - then: relative transition rates not altered
  - then: quantifiable **time boost factor** each timestep
  - then: trajectory is **time-accurate** (unlike enhanced sampling)

Surface diffusion modeling

- Pt (100) surface with 4% adatom coverage (random)
- Exchange barrier = 0.64 eV, hop barrier = 1.25 eV
- HD Vmax = 0.4 eV, T = 400K ⇒ **4000x boost**
- 1.2M atoms, 50M timesteps ⇒ **1 ms** of real time

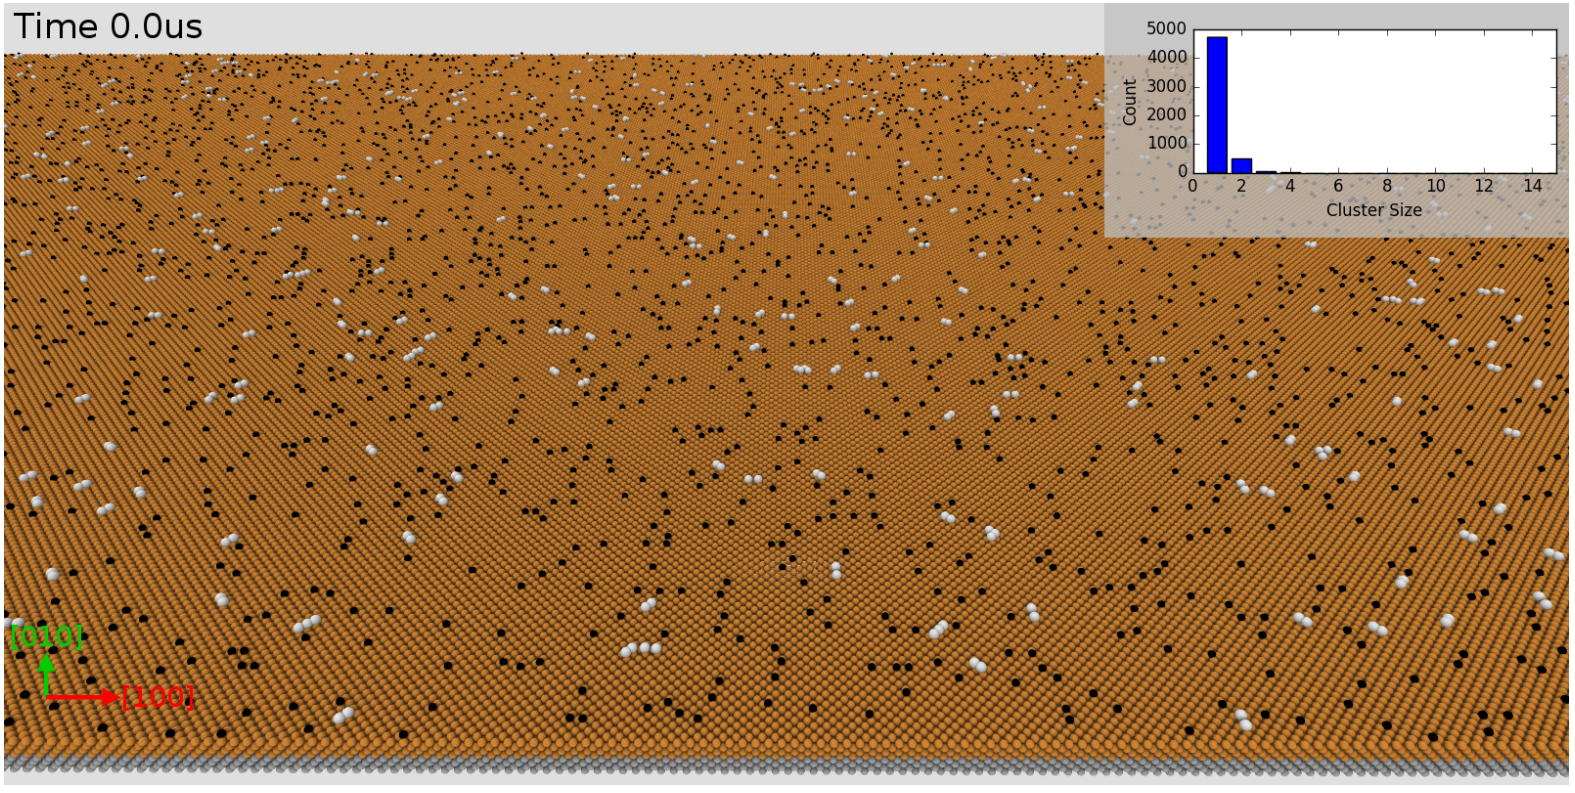

Surface diffusion modeling

- Pt (100) surface with 4% adatom coverage (random)
- Exchange barrier = 0.64 eV, hop barrier = 1.25 eV
- HD Vmax = 0.4 eV, T = 400K ⇒ **4000x boost**
- 1.2M atoms, 50M timesteps ⇒ **1 ms** of real time

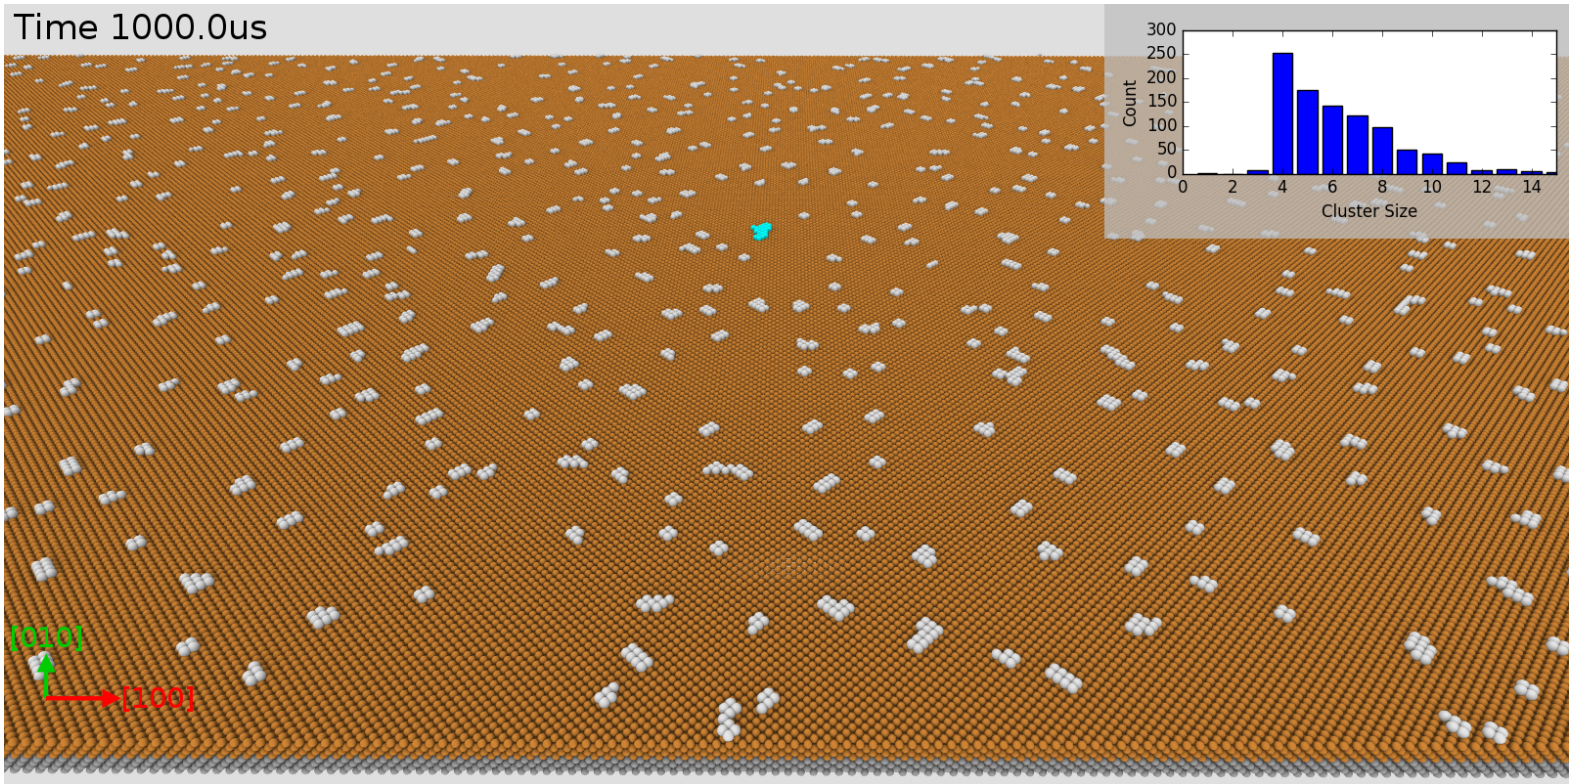

- **OVITO**: thanks to Mitch Wood and Alex Stukowski!

Surface diffusion modeling

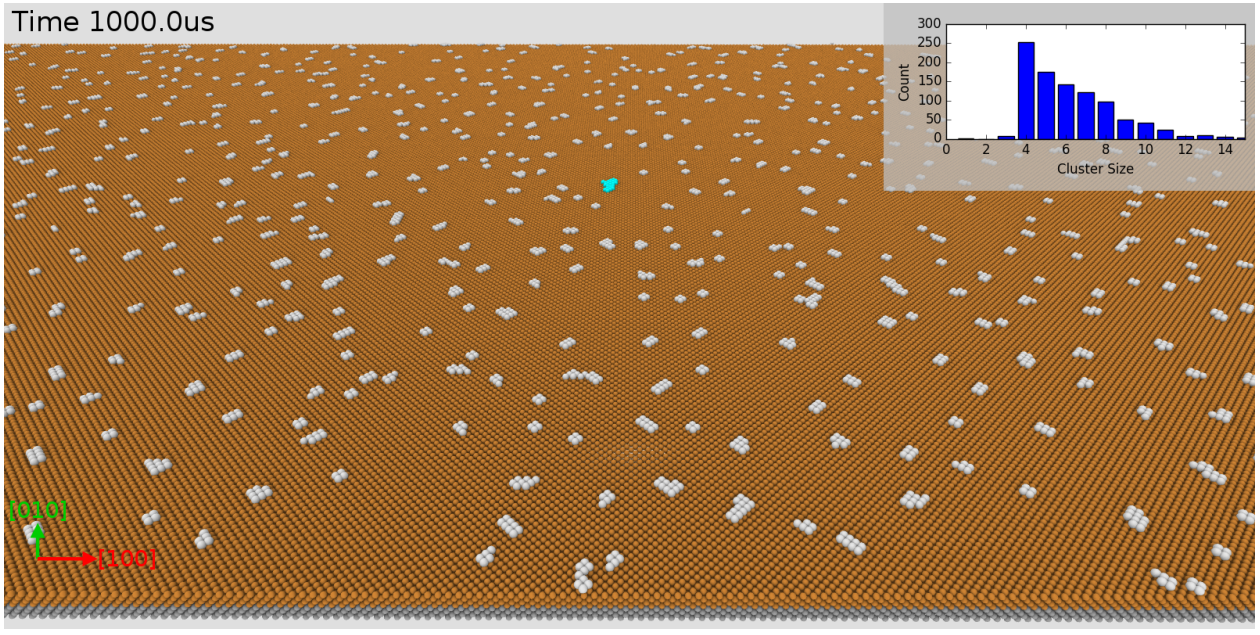

- Observe ∼**400K diffusion events**
- Versus 100 events (one event for every 60 adatoms)

Surface diffusion modeling

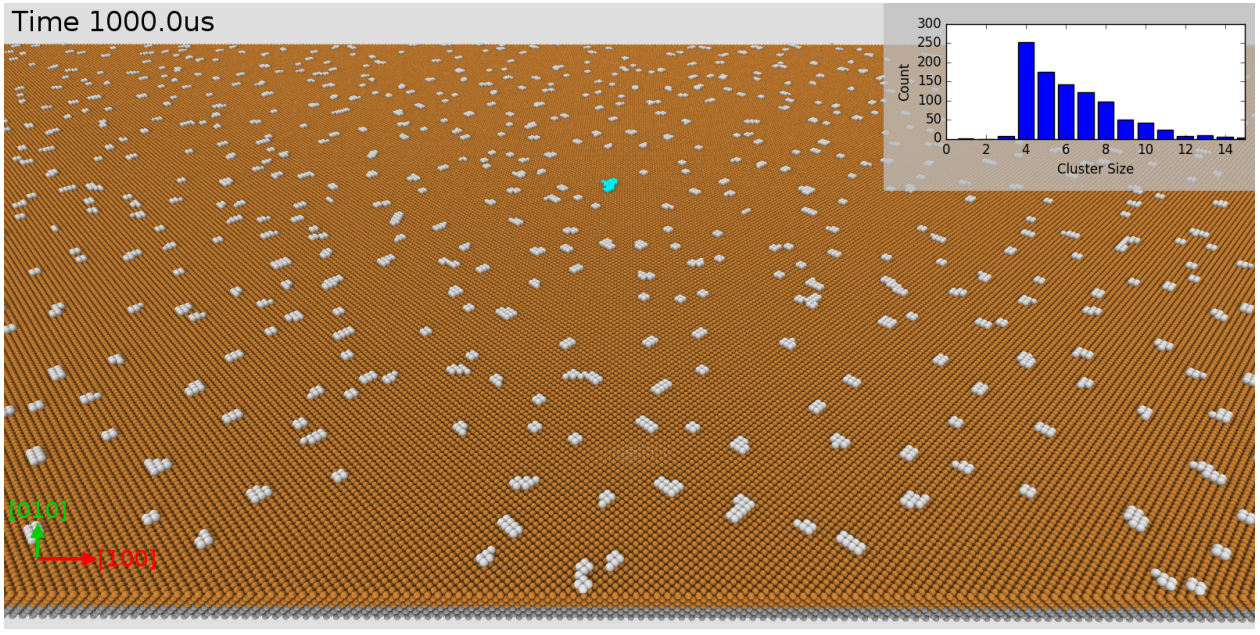

- Observe ∼**400K diffusion events**
- Versus 100 events (one event for every 60 adatoms)
- **24 hr run** on 4K Broadwell cores (128 nodes)
- Just 300 atoms/core ⇒ maximize time scale

Surface diffusion modeling

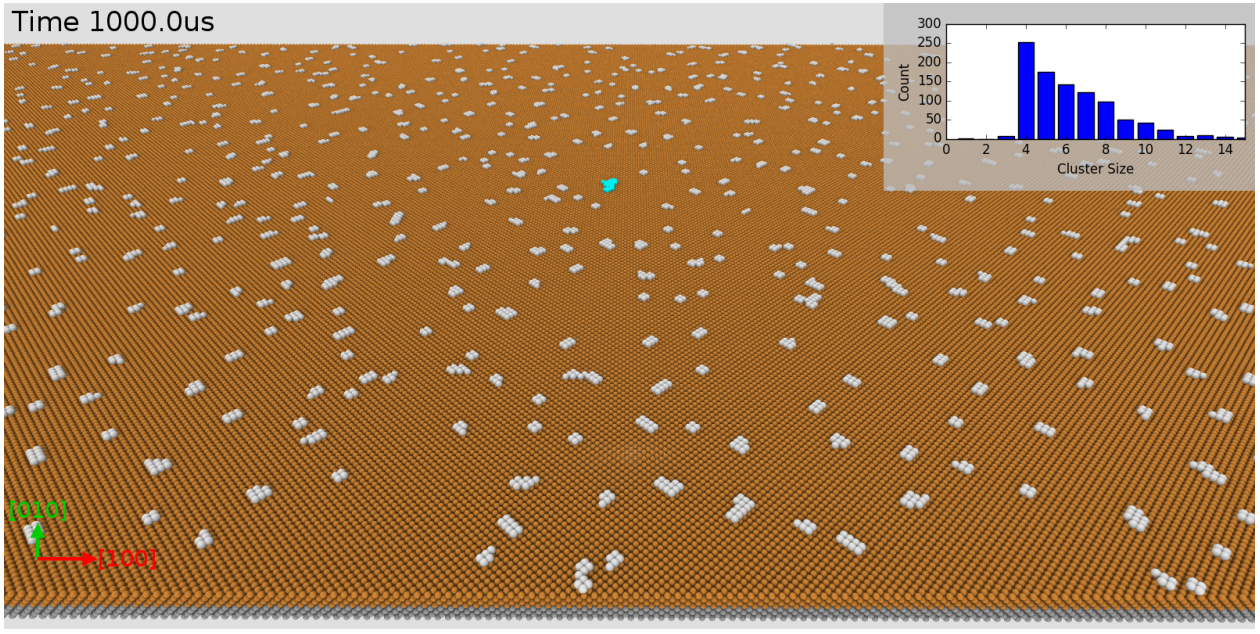

- Observe ∼**400K diffusion events**
- Versus 100 events (one event for every 60 adatoms)
- **24 hr run** on 4K Broadwell cores (128 nodes)
- Just 300 atoms/core ⇒ maximize time scale
- Pt cluster shapes affected by exchange mechanism
- By contrast to Cu where diffusion is hop dominated

Local hyperdynamics in parallel

Rule: Bias **all** bonds that are max-strain within red $D_{cut} = 10$ Å

- Run 100 steps of **MD with HD bias**, every step:
  - compute strain of all bonds
  - identify locally max bonds, bias them
  - adjust per-bond V_{max} to boost time uniformly everywhere
    - for every bond, based on max-strain bond within its D_{cut}
- Quench to check if **event** has occurred
  - if yes: re-form bond list, preserving bond prefactors
  - repeat MD/quench sequence

Local hyperdynamics in parallel

Rule: Bias **all** bonds that are max-strain within red $D_{cut} = 10$ Å

- Run 100 steps of **MD with HD bias**, every step:
  - compute strain of all bonds
  - identify locally max bonds, bias them
  - adjust per-bond V_{max} to boost time uniformly everywhere
    - for every bond, based on max-strain bond within its D_{cut}
- Quench to check if **event** has occurred
  - if yes: re-form bond list, preserving bond prefactors
  - repeat MD/quench sequence
- Extra operations for HD:
  - **2nd neighbor list** out to D_{cut}
  - **double loops** over atoms and big neighbor list
  - **comm** to acquire strain/max-strain for ghost atoms

Strong & weak scaling for local HD is similar to MD

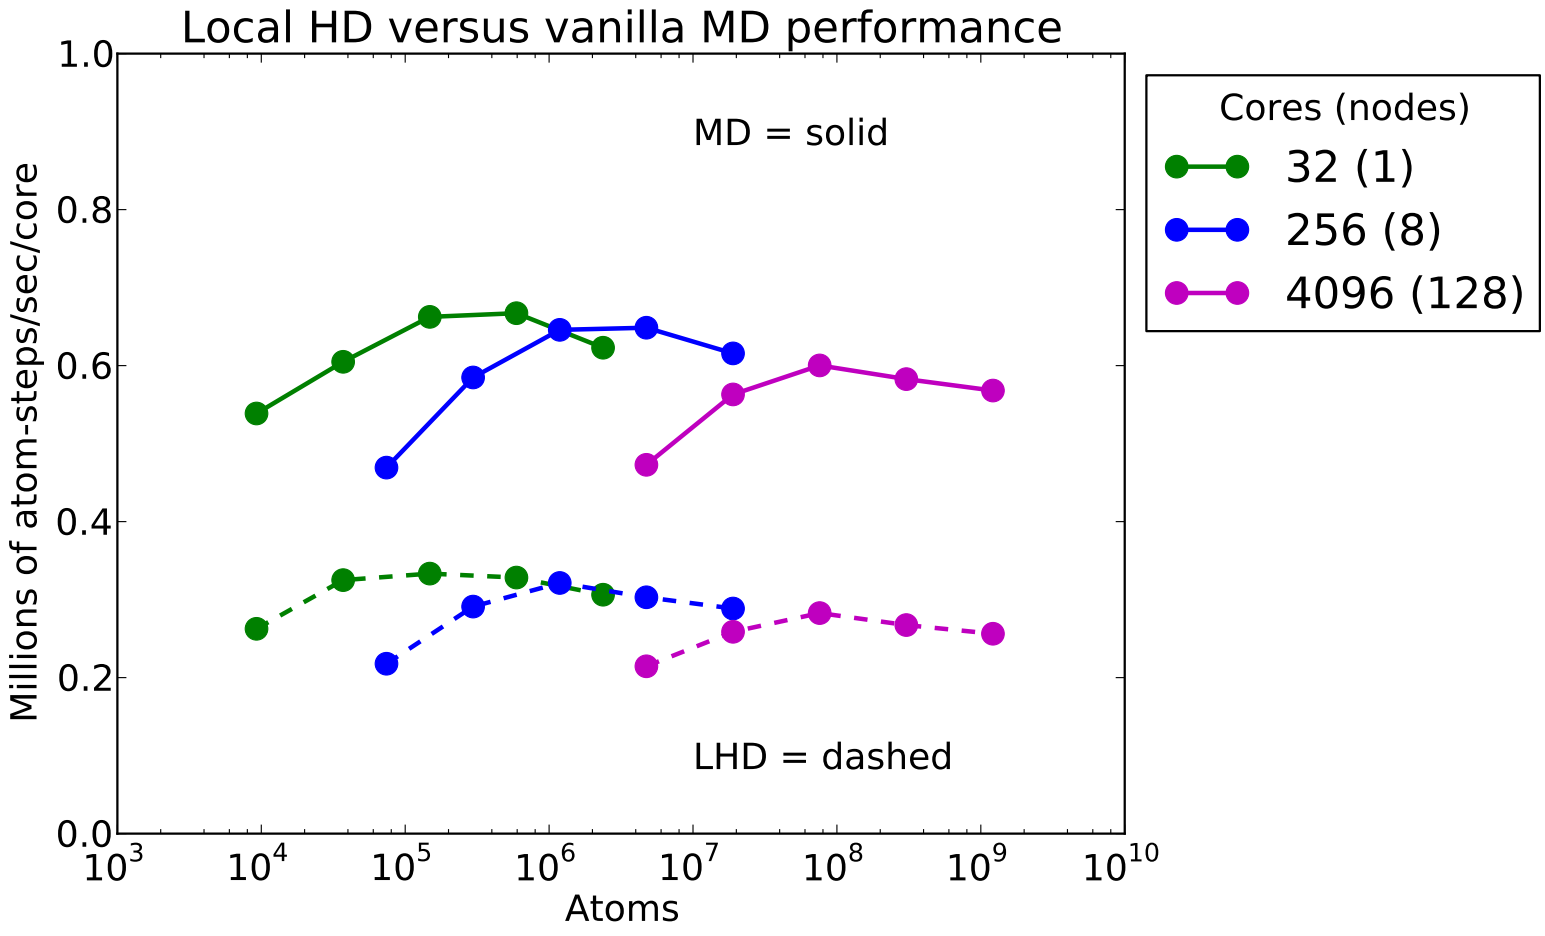

- For cheap EAM, HD is ~2x more expensive than MD

Cost of HD versus regular MD

Key point:

Can use global/local HD with **any potential** in LAMMPS

- Why **2x** more expensive for EAM?
  - half is quench & event search every 100 steps
  - half is comp/comm to find max-strain bonds
  - 2nd half shrinks for more expensive potentials $\Rightarrow$ **1.5x**

Cost of HD versus regular MD

Key point:

Can use global/local HD with **any potential** in LAMMPS

- Why **2x** more expensive for EAM?
  - half is quench & event search every 100 steps
  - half is comp/comm to find max-strain bonds
  - 2nd half shrinks for more expensive potentials $\Rightarrow$ **1.5x**

Bottom line:

- For systems that allow for HD, boost is free speed-up

Open issues - **suggestions** welcome!

- **Automated selection** of time boost
  - allowed boost is function of current lowest barrier height
  - but often do not know barriers **a priori**
  - **ideal**: on-the-fly adaptation
- **Smarter quench**:
  - converge based on bond lengths
  - quick convergence even w/out frozen substrate ?

Open issues - **suggestions** welcome!

- **Automated selection** of time boost
  - allowed boost is function of current lowest barrier height
  - but often do not know barriers **a priori**
  - **ideal**: on-the-fly adaptation
- **Smarter quench**:
  - converge based on bond lengths
  - quick convergence even w/out frozen substrate ?

- Run multi-replica methods, each with HD ?
- Application to **non-equilibrium** models ?
- Use this for **ab initio MD** with quantum force fields ?
- **CSlib**: code coupling via files/sockets/MPI,
  enable any MD code to talk to any QM code

Thanks and links

- Funding from DOE **exascale computing program** (ECP)
- EXA**ALT** = exascale atomistics for accuracy, length, time

- **LAMMPS** collaborators:
  Aidan Thompson, Stan Moore, Mitch Wood (Sandia)
  Axel Kohlmeyer (Temple U)

- LAMMPS MD package: **http://lammps.sandia.gov**
  - hyperdynamics to be released soon