Molecular dynamics: looking ahead to exascale

Steve Plimpton
Sandia National Laboratories

17th Annual Workshop on Charm++
and its Applications
May 2019 - University of Illinois Urbana-Champaign

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: SAND2018-5931C

Impact of advancing HPC on MD simulations

Impact of advancing HPC on MD simulations

- Most methods/models are $\sim$ **O(N)** cost in atom count
- Also scale as $\sim$ **O(N/P)** in parallel, for large enough N/P
- 1000x machine $\Rightarrow$ 1000x more atoms or time or combo

Impact of advancing HPC on MD simulations

- Most methods/models are $\sim$ **O(N)** cost in atom count
- Also scale as $\sim$ **O(N/P)** in parallel, for large enough N/P
- 1000x machine $\Rightarrow$ 1000x more atoms or time or combo

30 yrs ago:
my thesis
1000 atoms
50K steps

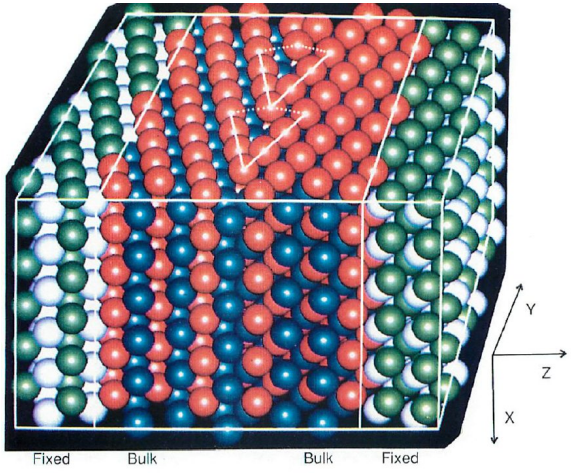

Impact of advancing HPC on MD simulations

- Most methods/models are $\sim$ **O(N)** cost in atom count
- Also scale as $\sim$ **O(N/P)** in parallel, for large enough N/P
- 1000x machine $\Rightarrow$ 1000x more atoms or time or combo

30 yrs ago:
my thesis
1000 atoms
50K steps

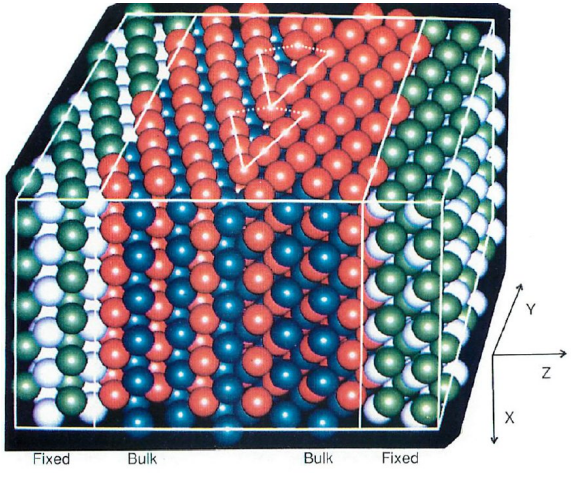

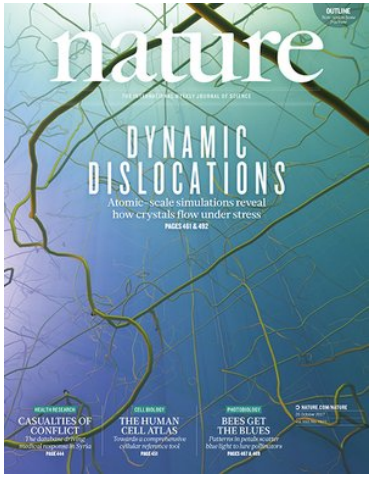

Today:
V Bulatov,
et al (LLNL)
2.1B atoms
460M steps

Impact of advancing HPC on MD simulations

- Most methods/models are $\sim$ **O(N)** cost in atom count
- Also scale as $\sim$ **O(N/P)** in parallel, for large enough N/P
- 1000x machine $\Rightarrow$ 1000x more atoms or time or combo

30 yrs ago:
my thesis
1000 atoms
50K steps

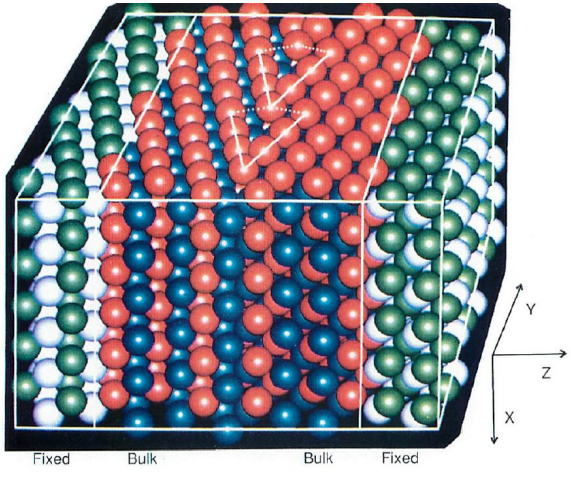

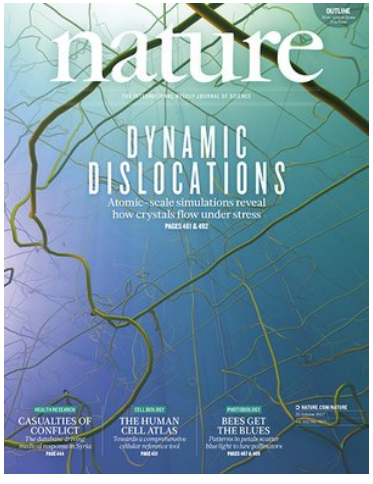

Today:
V Bulatov,
et al (LLNL)
2.1B atoms
460M steps

- Linpack: 1 BG/Q core / 1 Cray YMP proc = **41x** !!

Impact of advancing HPC on MD simulations

- Most methods/models are $\sim$ **O(N)** cost in atom count
- Also scale as $\sim$ **O(N/P)** in parallel, for large enough N/P
- 1000x machine $\Rightarrow$ 1000x more atoms or time or combo

30 yrs ago:
my thesis
1000 atoms
50K steps

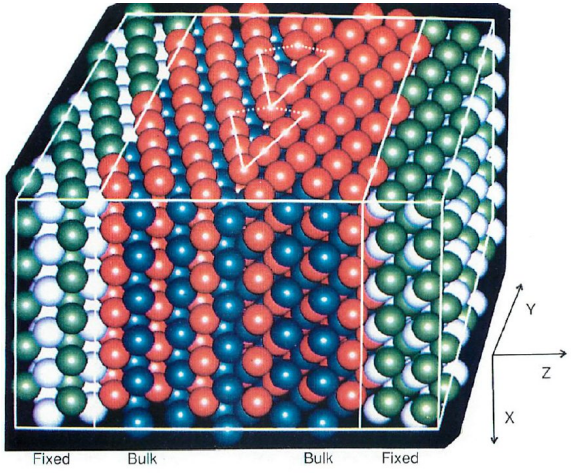

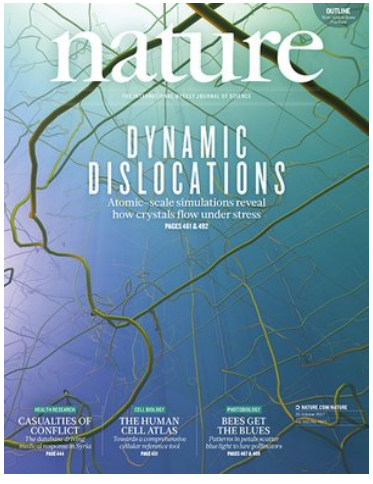

Today:
V Bulatov,
et al (LLNL)
2.1B atoms
460M steps

- Linpack: 1 BG/Q core / 1 Cray YMP proc = **41x** !!
- Cray YMP proc $\Rightarrow$ third of BG/Q Sequoia $\Rightarrow$ **21M** faster
- MD atom-steps/s $\Rightarrow$ **8.5M** faster

Impact of advancing HPC on MD simulations

- Most methods/models are $\sim$ **O(N)** cost in atom count
- Also scale as $\sim$ **O(N/P)** in parallel, for large enough N/P
- 1000x machine $\Rightarrow$ 1000x more atoms or time or combo

30 yrs ago:
my thesis
1000 atoms
50K steps

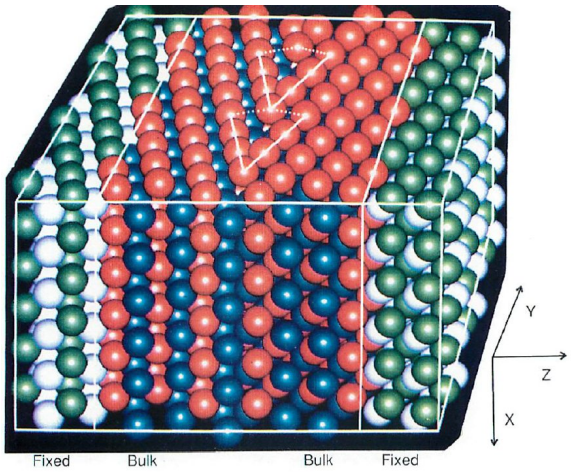

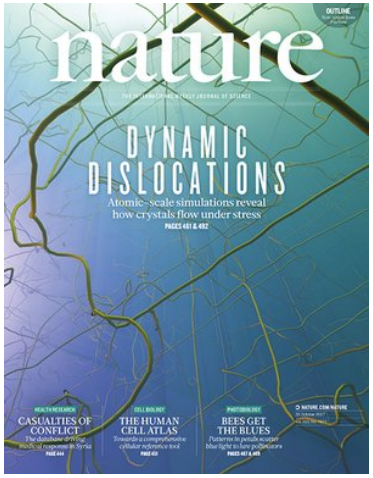

Today:
V Bulatov,
et al (LLNL)
2.1B atoms
460M steps

- Linpack: 1 BG/Q core / 1 Cray YMP proc = **41x** !!
- Cray YMP proc $\Rightarrow$ third of BG/Q Sequoia $\Rightarrow$ **21M** faster
- MD atom-steps/s $\Rightarrow$ **8.5M** faster
- Exascale is another **50x** beyond BG/Q $\Rightarrow$ 4 billion YMP procs

What will exascale computing mean for MD?

- 1000x machine $\Rightarrow$ 1000x more atoms or time ?

What will exascale computing mean for MD?

- 1000x machine $\Rightarrow$ 1000x more atoms or time ?

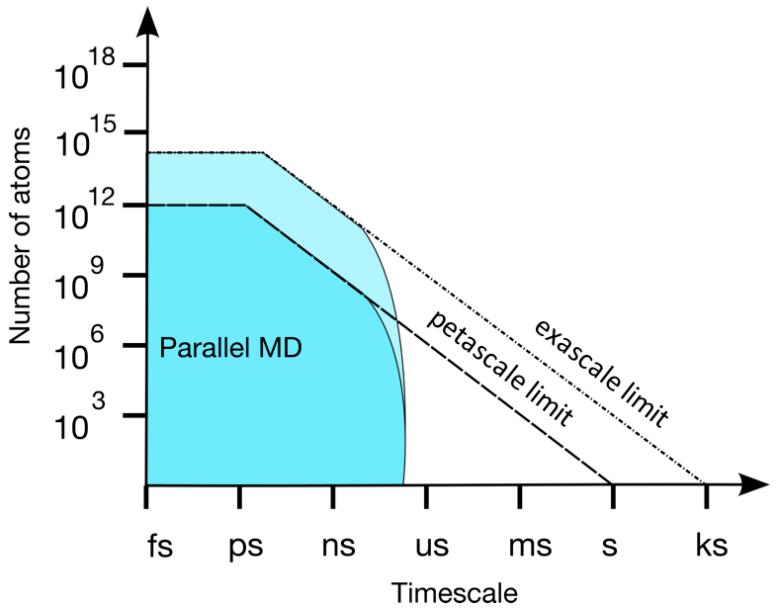

What will exascale computing mean for MD?

- 1000x machine $\Rightarrow$ 1000x more atoms or time ?

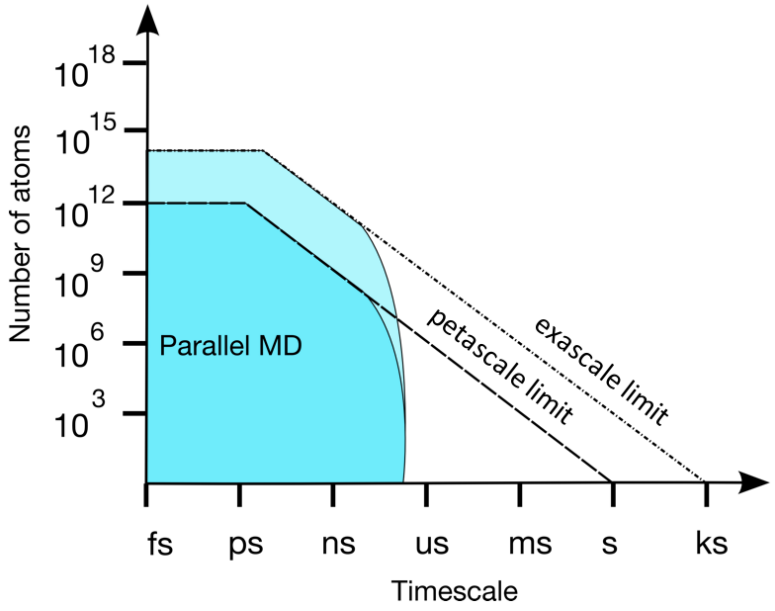

- Exascale can model systems 1000x bigger
- But can't run small systems 1000x longer
- **Why**: not enough parallel work, can't timestep any faster

A science motivation for long timescales

Modeling damage to materials in nuclear energy fusion reactors

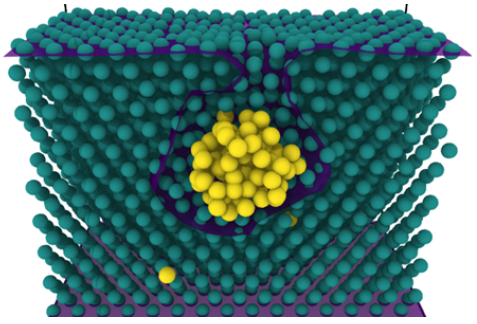

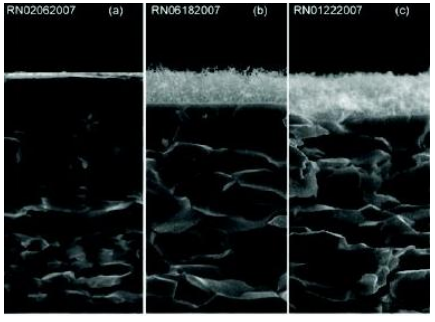

A science motivation for long timescales

Modeling damage to materials in nuclear energy fusion reactors

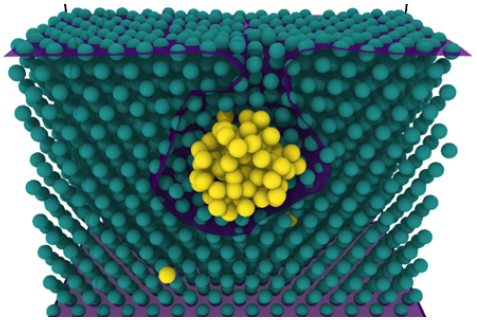

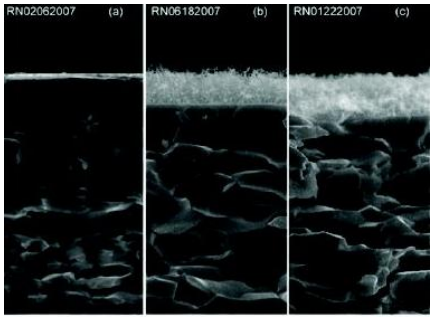

- EXA**ALT** = exascale atomistics for accuracy, length, time
- How EXAALT plans to model this problem at exascale
  - **not** a single large simulation with B or T atoms
  - millions of small MD replicas (few K to 1M atoms)
  - **ParSplice** code manages replicas:
    - chooses starting configurations
    - invokes LAMMPS as MD engine for each replica
    - creates distributed database of events
    - stitches together a long **statistically accurate** trajectory

Hyperdynamics (HD) can also extend MD timescales

- **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 a multi-replica or enhanced sampling approach

Hyperdynamics (HD) can also extend MD timescales

- **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 a multi-replica or enhanced sampling approach

- **Local** hyperdynamics
  - *Kim, Perez, Voter, J Chem Phys 139, 144110 (2013)*
  - **global**: bias one bond in entire system each timestep
  - **local**: bias multiple bonds separated by $R_{cut} = 10$ Å
  - tested correctness for simple, small systems
  - accelerated event rates match theory and experiment
  - biasing pairs of atoms $\Rightarrow$ **multi-atom events**

What kind of systems can benefit from HD

- Key **requirements**:
  - distinct, separated energy basins (solids, 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
  - time boost $\propto \exp(\Delta V/kT)$
- Complementary to **multi-replica** methods
  - each ParSplice replica could be running HD
  - time acceleration would be multiplicative

Pictorial view of hyperdynamics

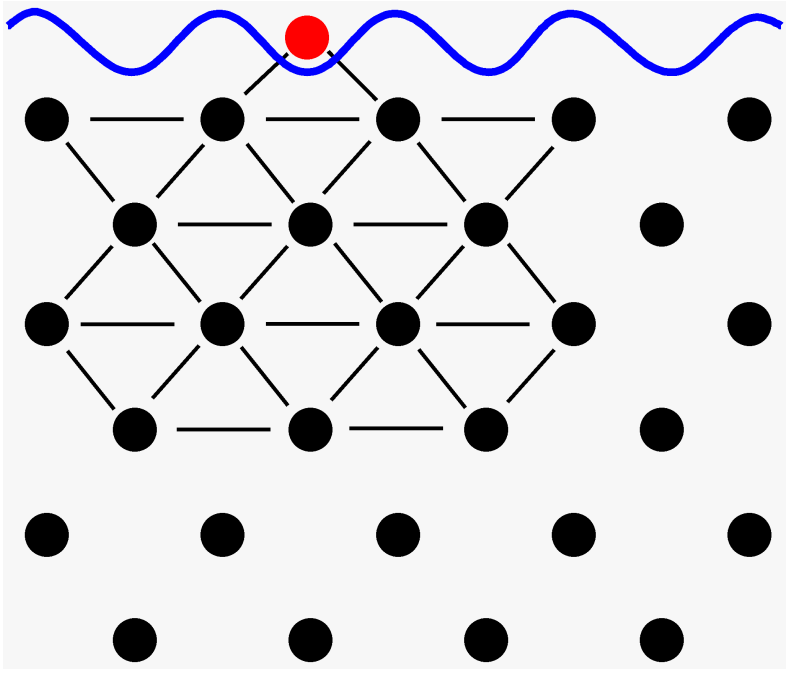

- Corrugated energy landscape for **adatom surface diffusion**
- Define (conceptual) **bonds** between all pairs of nearby atoms
  - e.g. $\sim$12 nearest neighbors per atom in fcc lattice

Zoom in to one adatom on surface

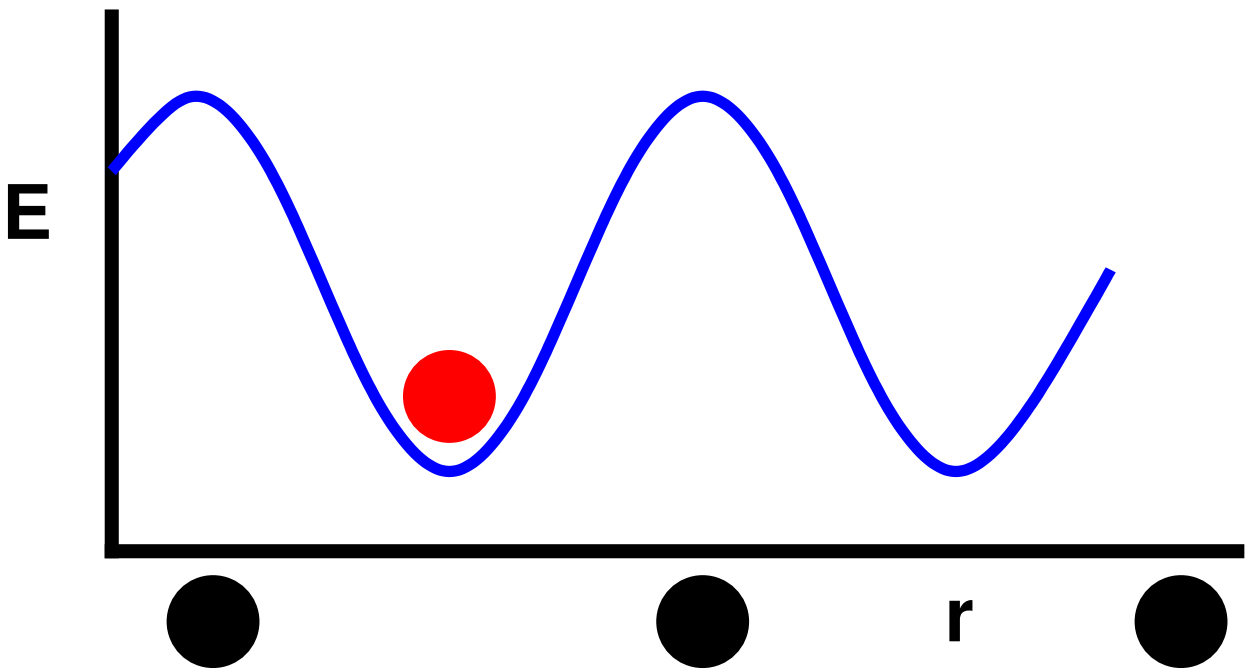

Added bias potential

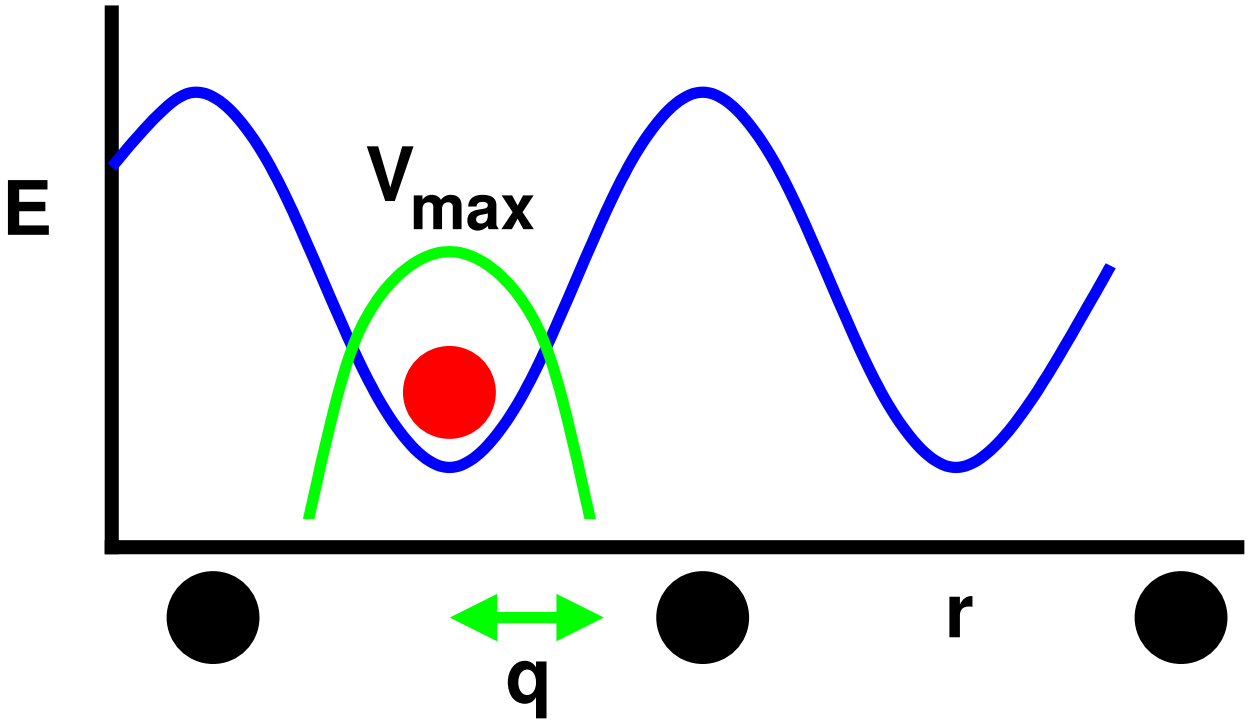

- Bond strain: $\epsilon_{ij} = (R_{ij} - Ro_{ij})/Ro_{ij}$
- Add **bias potential** to only the max-strain bond
- Bias: $V_{ij} = V_{max}[1 - (\epsilon_{ij}/q)^2], \quad |\epsilon_{ij}| < q, \quad \text{else} \quad$ **zero**
- **Different bond** may be biased at each timestep

Resulting potential energy surface

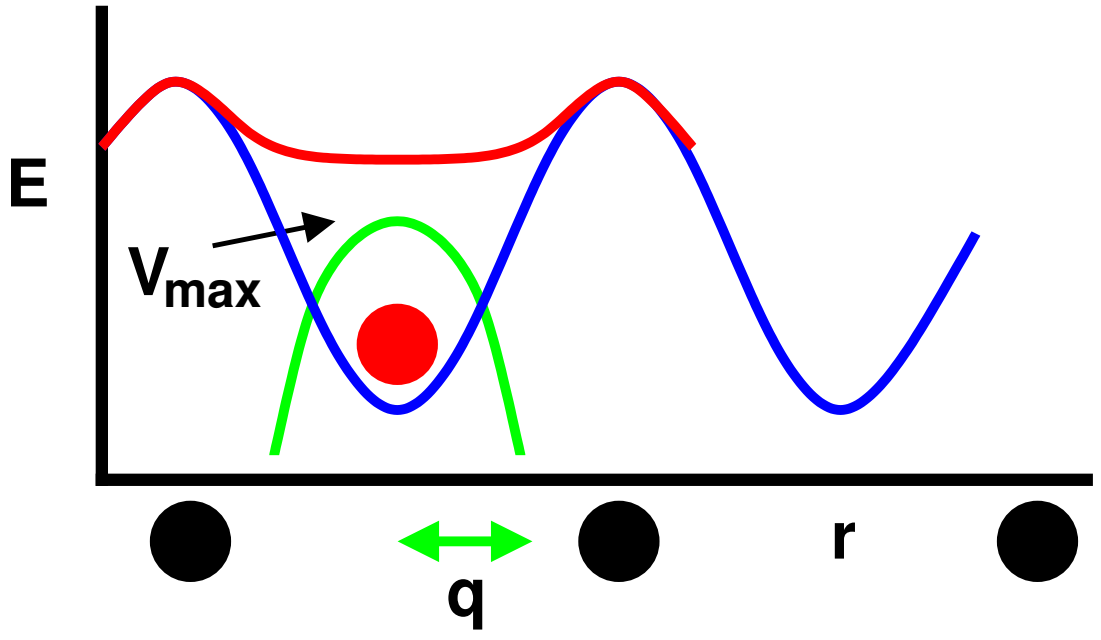

- Shallow well $\Rightarrow$ **faster transition** by I,J (and nearby) atoms

Resulting potential energy surface

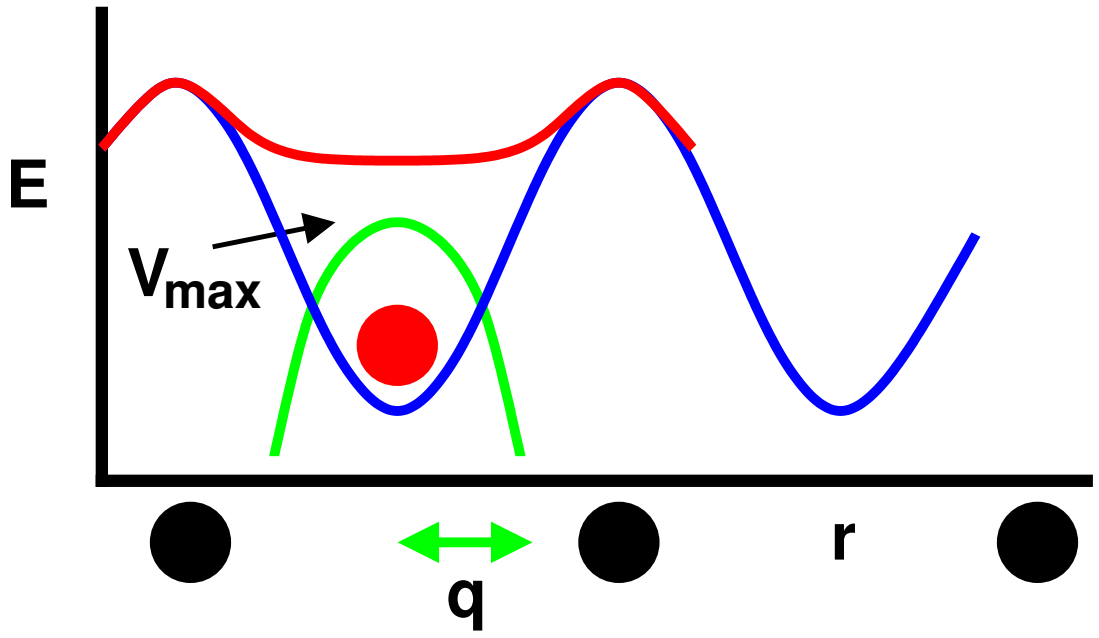

- Shallow well $\Rightarrow$ **faster transition** by I,J (and nearby) atoms
- Must choose V_{max} and q carefully:
  - if: zero bias at dividing surfaces (Q), no local minima (V_{max})
  - if: do not induce correlated events that violate TST
  - then: relative transition rates not altered for competing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

Surface diffusion modeling

- Pt (100) surface with 4% adatom coverage (random)
- HD: $V_{max} = 0.4$ eV, T = 400K $\Rightarrow$ **4000x boost**
- **1.2M** atoms, 50M timesteps $\Rightarrow$ **1 ms** of real time
- **48 hr run** on 128 Broadwell nodes (4K cores)

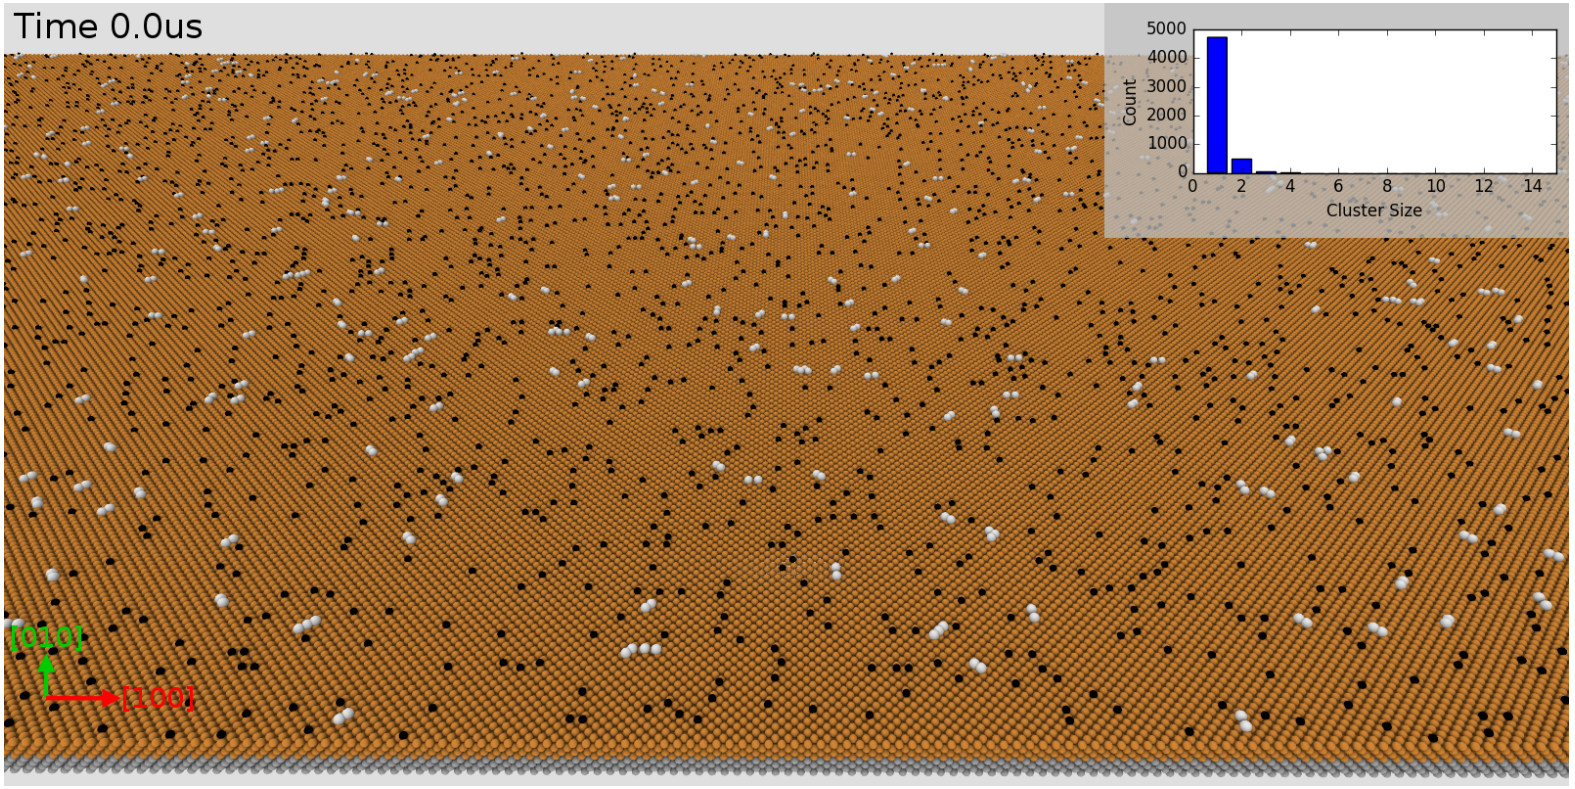

What movie will show

- Biasing ∼**3000** bonds each timestep, ∼**400K** diffusion events
- Versus 100 events with MD (one event per 60 adatoms)
- **Cluster formation**, monitored by size histogram
- Rich variety of events occur naturally, no *a priori* insight

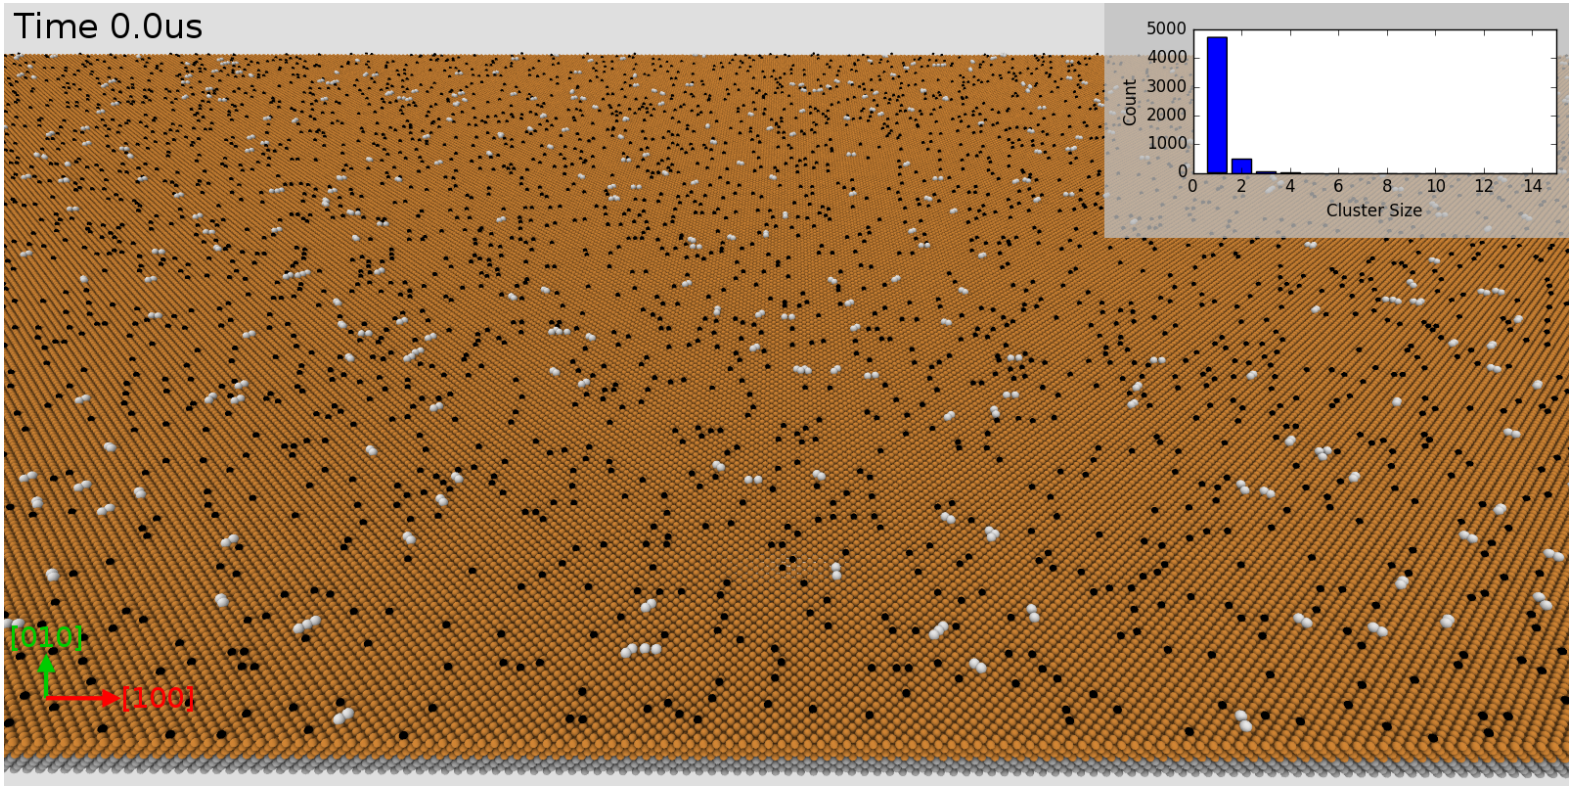

What movie will show

- Biasing ∼**3000** bonds each timestep, ∼**400K** diffusion events
- Versus 100 events with MD (one event per 60 adatoms)
- **Cluster formation**, monitored by size histogram
- Rich variety of events occur naturally, no *a priori* insight

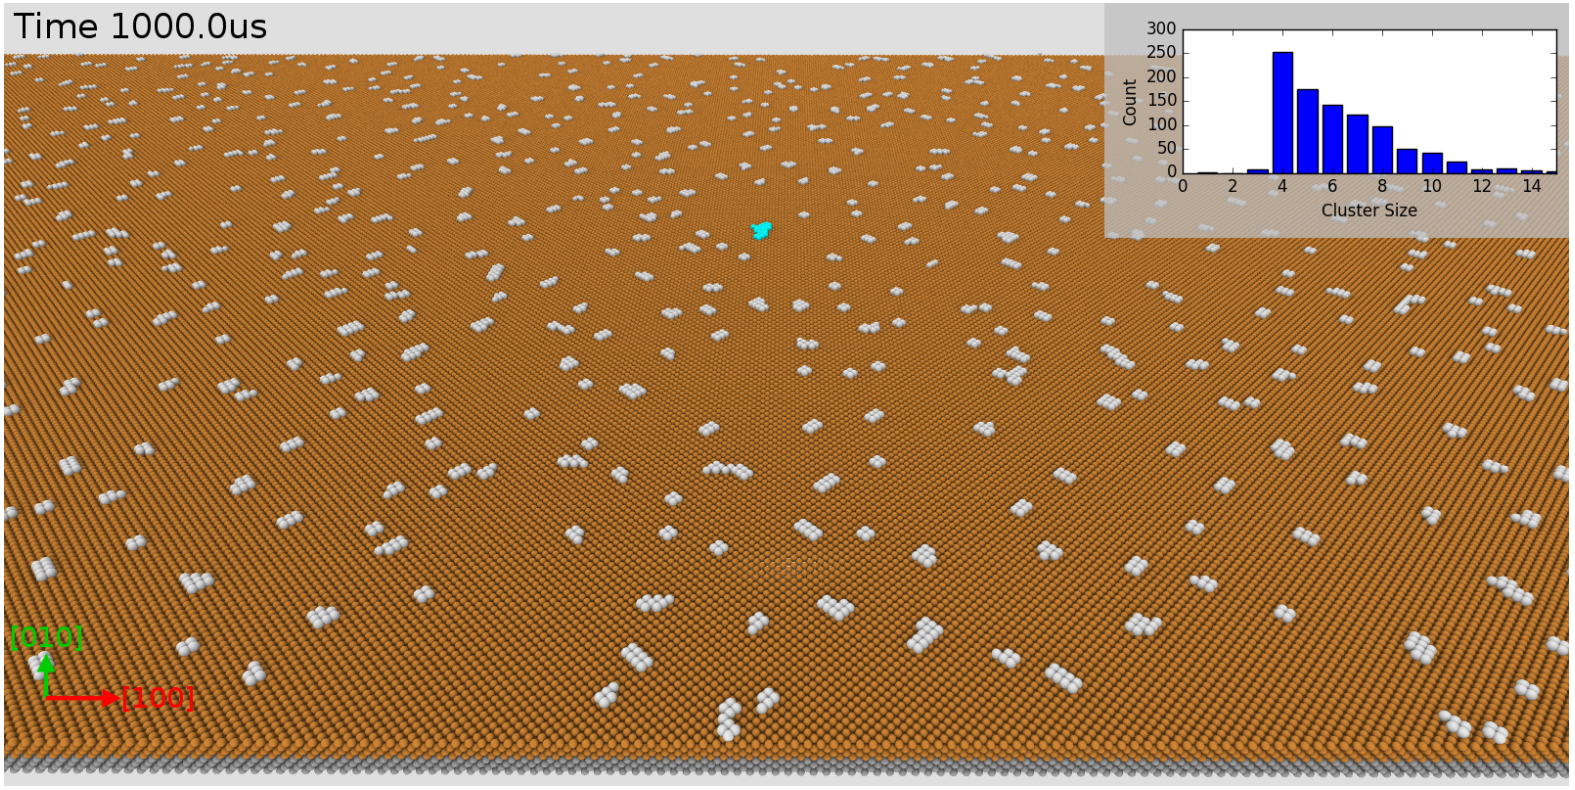

Movie

- Not just adatom motion, **substrate atoms** part of every event
- Mobile monomers, dimers, trimers
- Larger clusters are immobile, except around perimeter

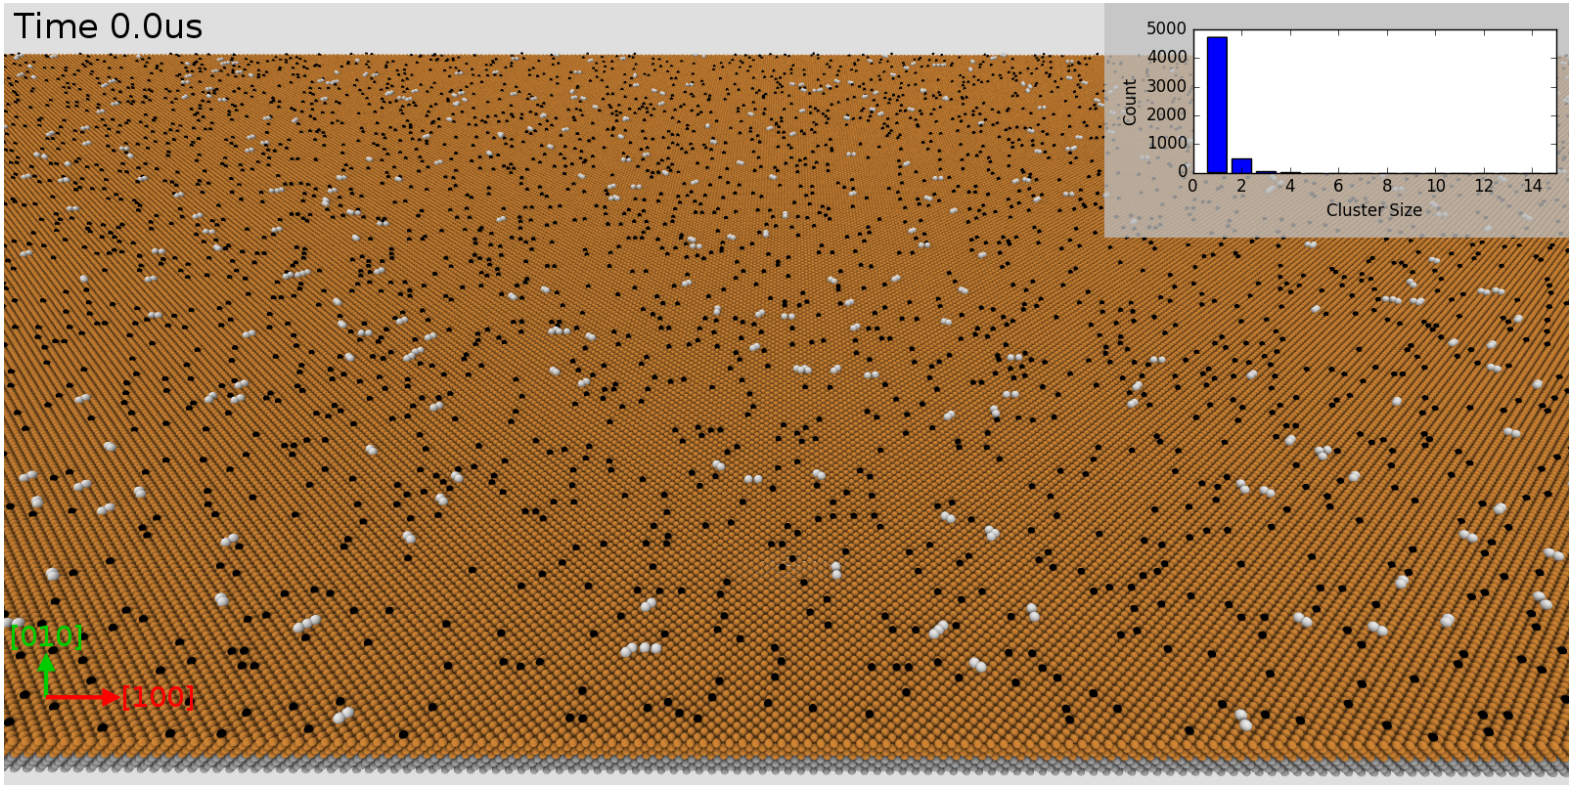

Movie

- Not just adatom motion, **substrate atoms** part of every event
- Mobile monomers, dimers, trimers
- Larger clusters are immobile, except around perimeter

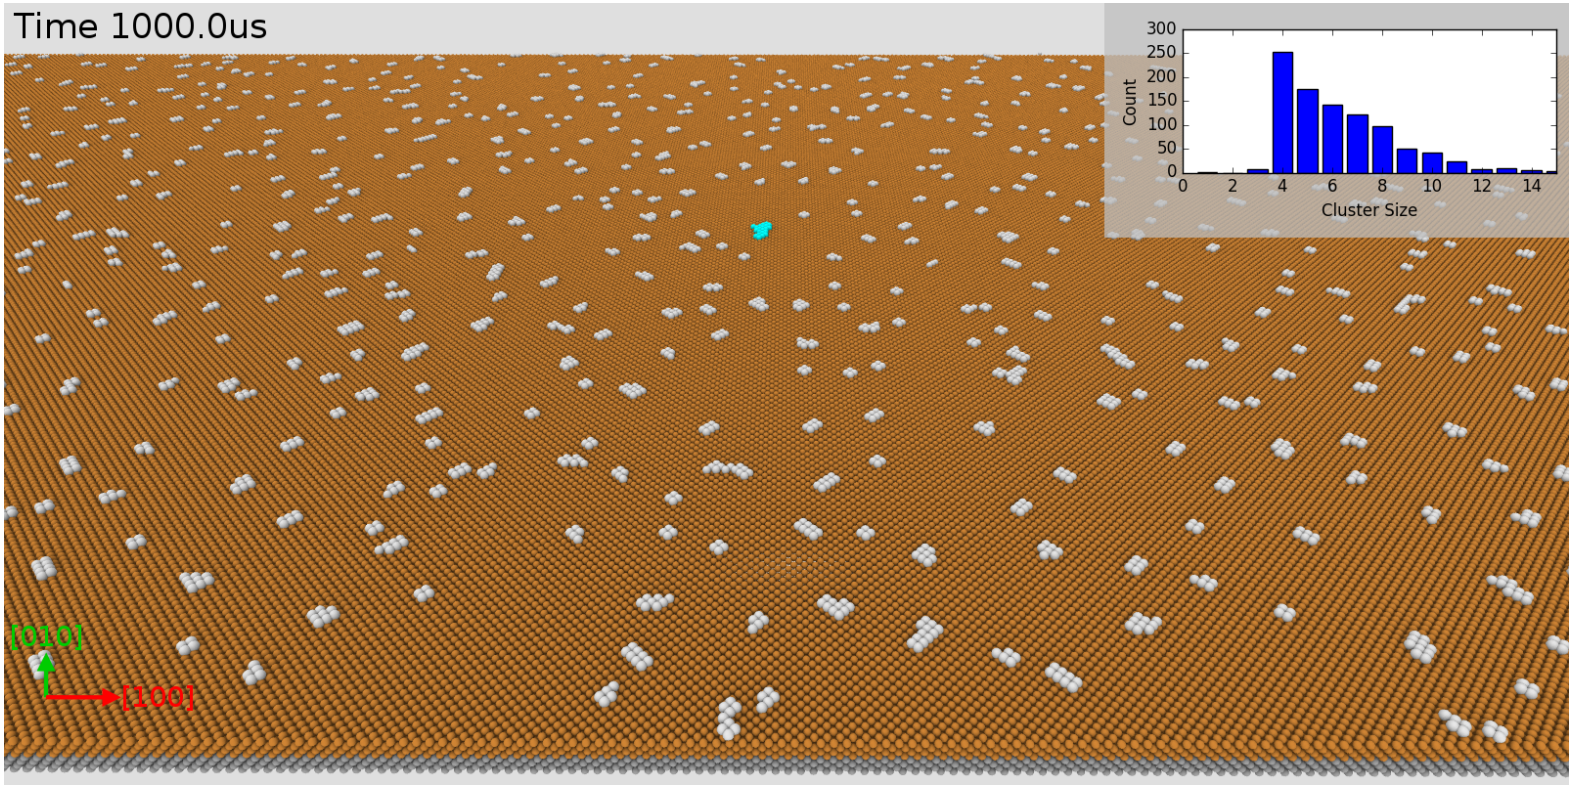

- **OVITO** help: thanks to Mitch Wood (Sandia)

Running a HD simulation in an MD code

Via new **hyper** command in LAMMPS

Running a HD simulation in an MD code

Via new **hyper** command in LAMMPS

- Choose V_{max}, q, and **T**
- Save initial quench state of system
- Loop:
  - run 100 steps of **MD** with Langevin thermostat
    add **HD bias** at every step to selected atom pair(s)
  - save dynamic state
  - perform **quench**
  - check if any **events** occurred (relative to previous quench)
  - if yes:
    archive event info
    save new quenched state
    recreate bond list = I,J pairs, equilibrium R_0
  - restore dynamic state

Running a HD simulation in an MD code

Via new **hyper** command in LAMMPS

- Choose V_{max}, q, and **T**
- Save initial quench state of system
- Loop:
  - run 100 steps of **MD** with Langevin thermostat
    add **HD bias** at every step to selected atom pair(s)
  - save dynamic state
  - perform **quench**
  - check if any **events** occurred (relative to previous quench)
  - if yes:
    archive event info
    save new quenched state
    recreate bond list = I,J pairs, equilibrium R_0
  - restore dynamic state

Usual **parallel** MD and quench (spatial partitioning of atoms)

Extra operations and data for computing HD bias

- Bias **every bond** that is local max-strain bond within R_{cut}
- R_{cut} = distance at which one event influences another
- $\sim$2x cutoff for EAM = 10 Å $\Rightarrow$ **700** neighbor bonds/bond

Extra operations and data for computing HD bias

- Bias **every bond** that is local max-strain bond within R_{cut}
- R_{cut} = distance at which one event influences another
- ~2x cutoff for EAM = 10 Å $\Rightarrow$ **700** neighbor bonds/bond

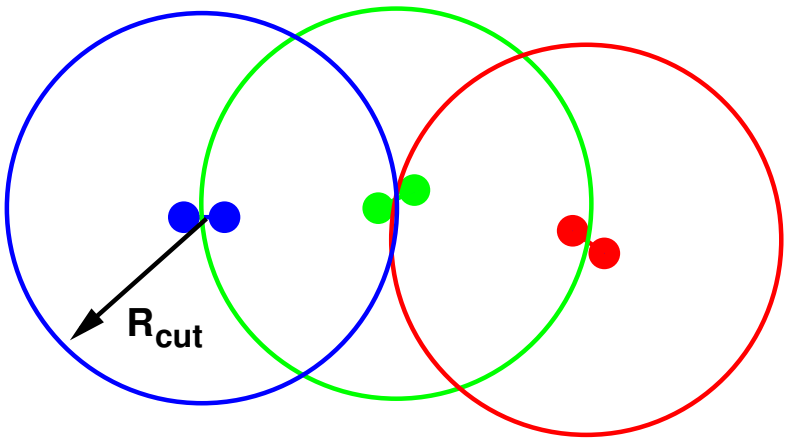

Extra operations and data for computing HD bias

- Bias **every bond** that is local max-strain bond within R_{cut}
- R_{cut} = distance at which one event influences another
- ~2x cutoff for EAM = 10 Å $\Rightarrow$ **700** neighbor bonds/bond

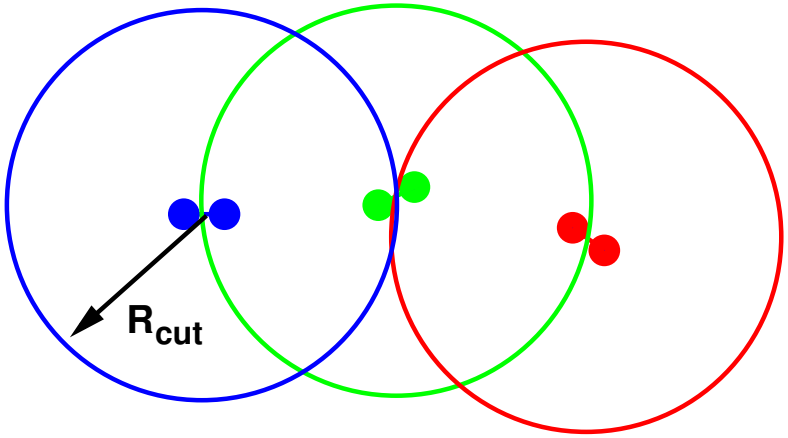

- Create and loop over **2nd neighbor list** out to R_{cut}
- **Communication** to acquire strain info for ghost atoms

Parallel scaling for local HD is similar to MD

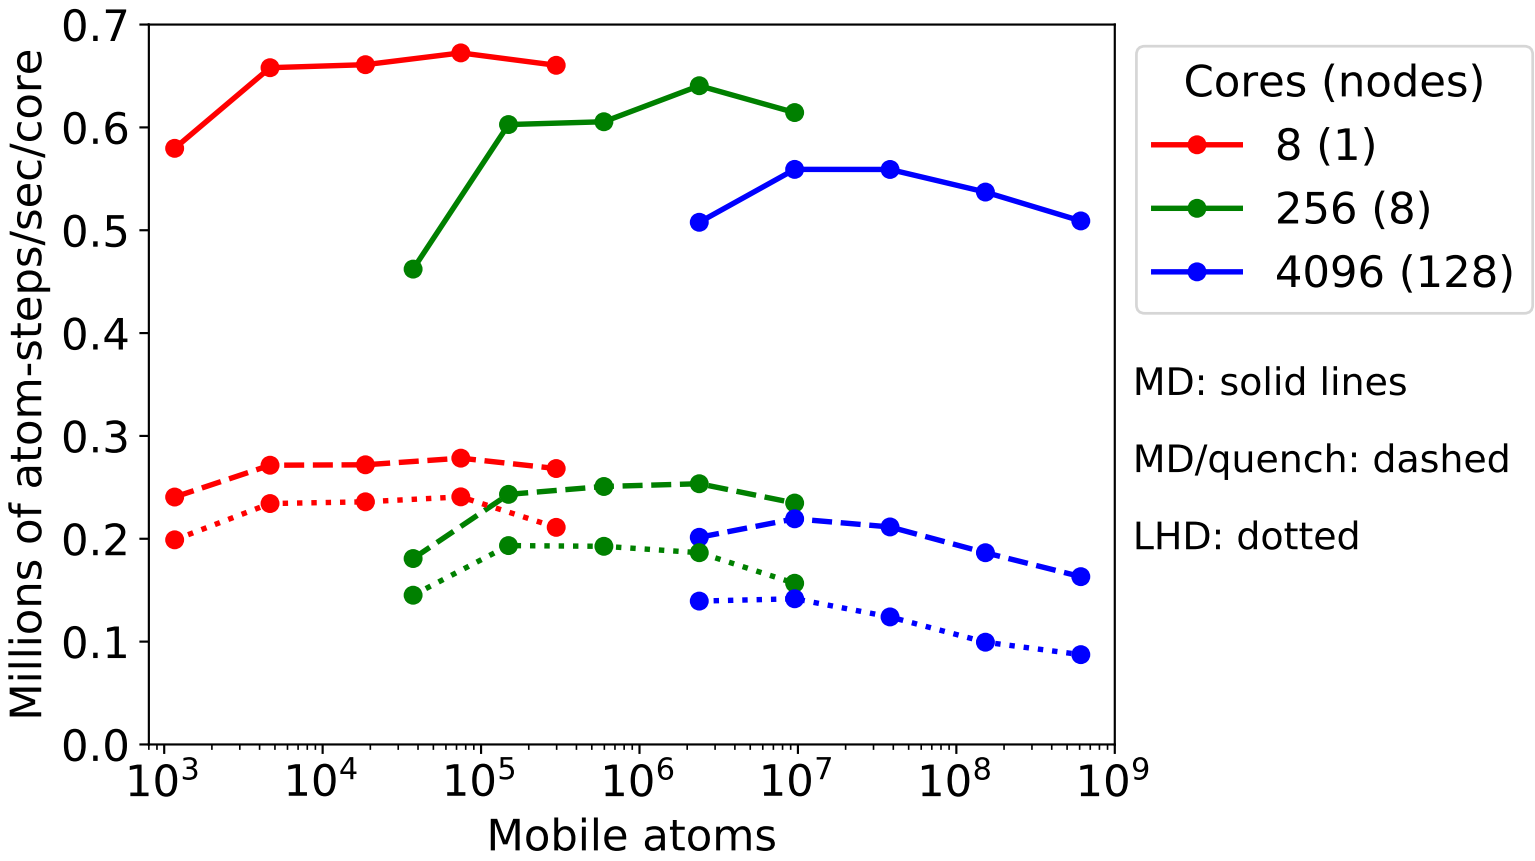

- For cheap EAM, HD is ∼**3x-5x** more expensive than MD
- Majority is careful **quench**, rest is comp/comm out to **Rcut**

Exchange event and dimer diffusion

Green: atom moves > 1.0 Å during event
Purple: > 0.2 Å, **Yellow**: > 0.1 Å, **Red**: < 0.1 Å

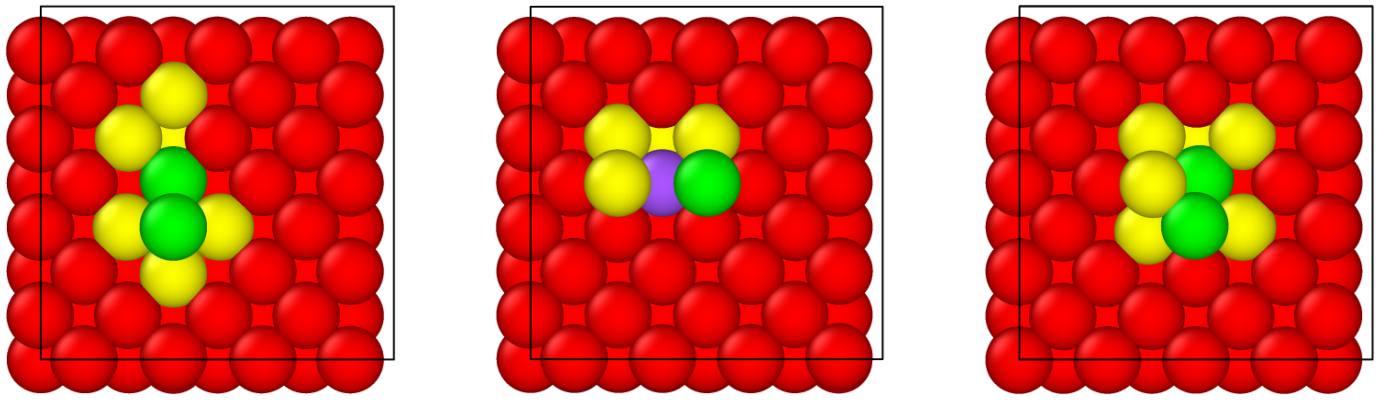

Exchange barrier = **0.656 eV**, hop barrier = 1.25 eV (too high)
Hop barrier when next to another adatom = **0.635 eV**
Successive exchanges enable **dimer diffusion**

Trimer duck-under and bend

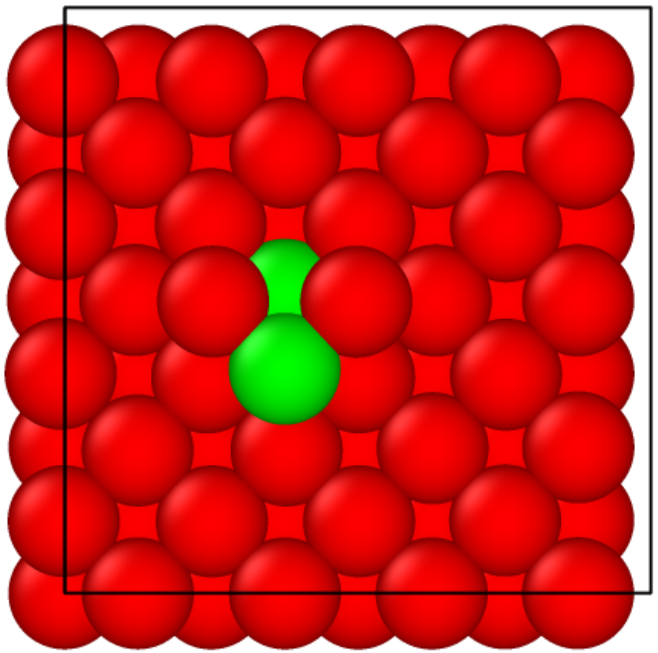

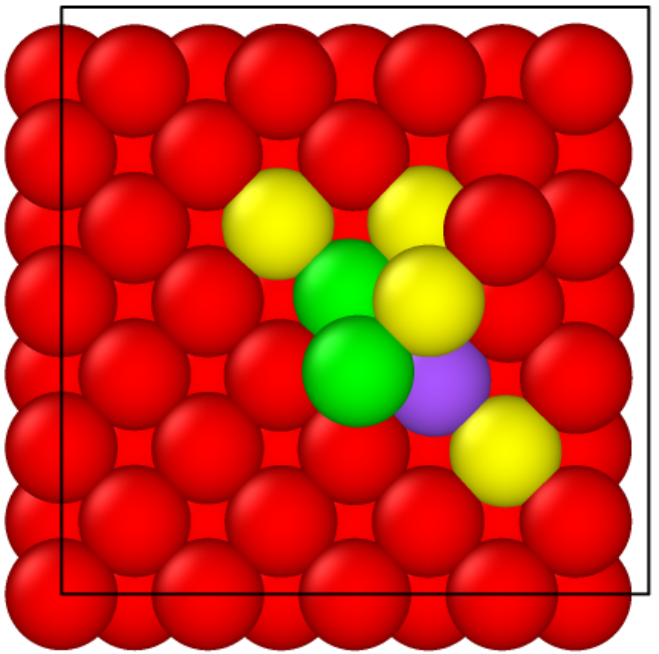

Duck-under barrier = 0.410 eV

Lowest barrier event, recall we chose $V_{max} = 0.4$ eV

Successive bends & duck-unders enable **trimer diffusion**

Flower formation event

Highly technical name!

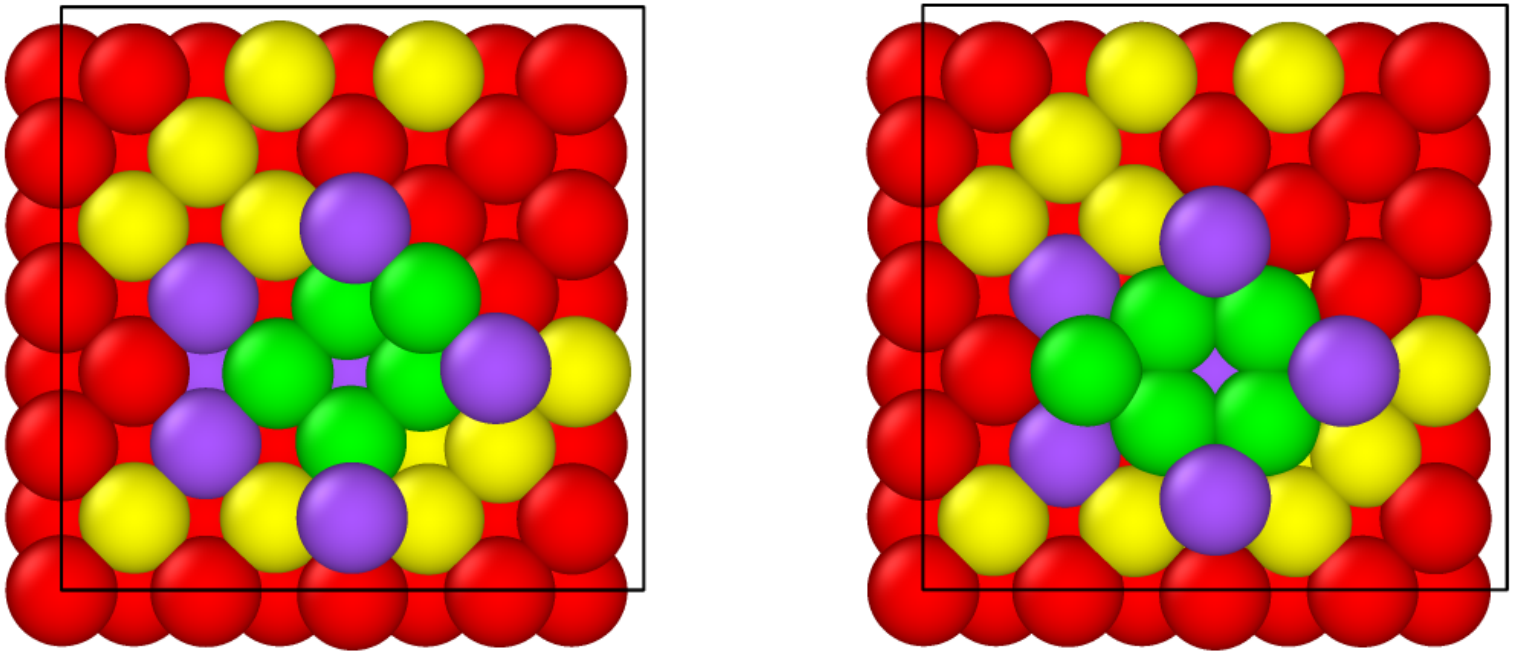

Barrier = 0.772 eV

Reverse event can result in **long-distance trimer move**

Crowdion event

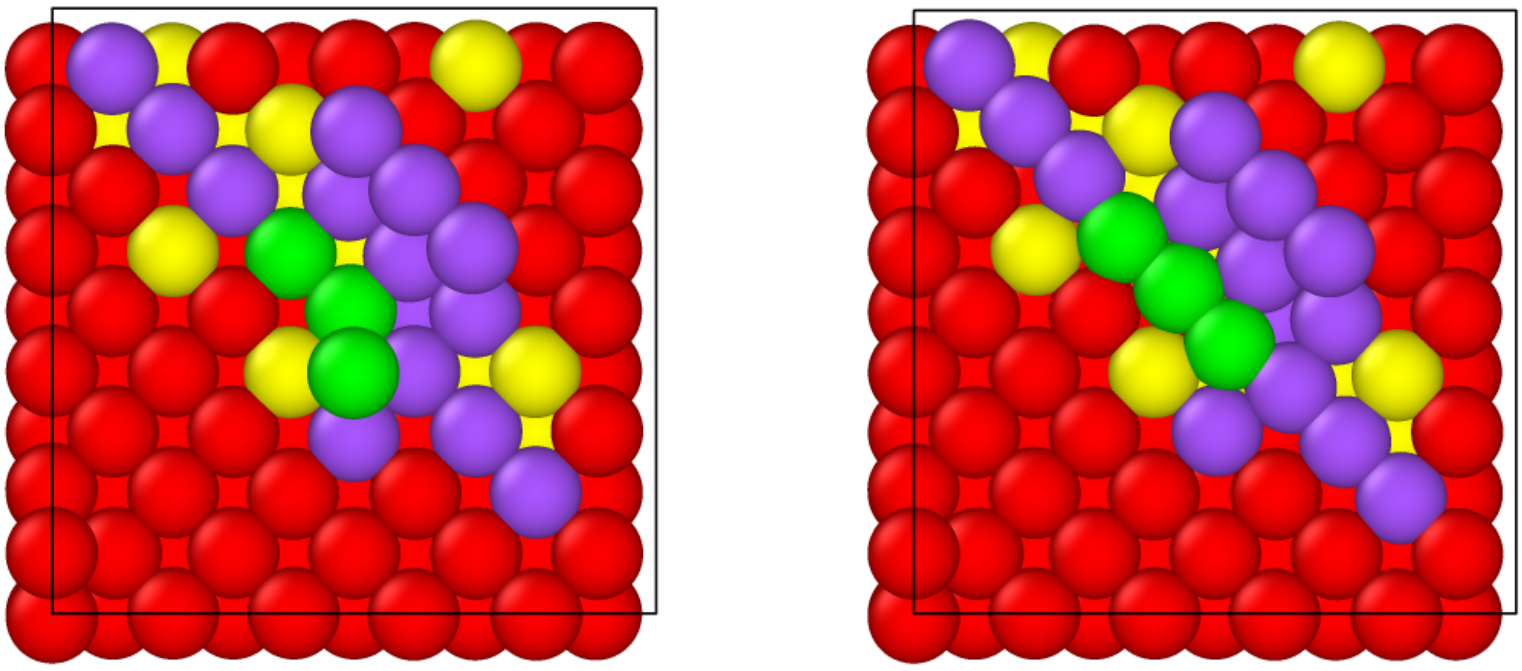

Barrier = 0.771 eV (induced by trimer)

Reverse event can displace adatom by 2 lattice sites in (110)

Hyperdynamics summary

Key points:

- Can use global/local HD with **any potential** in LAMMPS
- HD bias forces just added to interatomic forces
- Time boost is **free speed-up** for systems that allow for HD

Hyperdynamics summary

Key points:

- Can use global/local HD with **any potential** in LAMMPS
- HD bias forces just added to interatomic forces
- Time boost is **free speed-up** for systems that allow for HD

Lower temperatures:

- 400K $\Rightarrow$ 4000x boost $\Rightarrow$ 50M steps $\Rightarrow$ 1 ms
- 300K $\Rightarrow$ 120Kx boost $\Rightarrow$ 30 ms
- 200K $\Rightarrow$ 300Mx boost $\Rightarrow$ 75 s

Hyperdynamics summary

Key points:

- Can use global/local HD with **any potential** in LAMMPS
- HD bias forces just added to interatomic forces
- Time boost is **free speed-up** for systems that allow for HD

Lower temperatures:

- 400K $\Rightarrow$ 4000x boost $\Rightarrow$ 50M steps $\Rightarrow$ 1 ms
- 300K $\Rightarrow$ 120Kx boost $\Rightarrow$ 30 ms
- 200K $\Rightarrow$ 300Mx boost $\Rightarrow$ 75 s

Challenges:

- Can we perform smarter, cheaper quenches
- Often do not know all barrier heights *a priori*
  - allowed time boost is function of current lowest barrier height
  - **ideal**: on-the-fly adaptation of T_{boost}, V_{max}, q

Coding apps for the bleeding-edge of HPC

- **Vectorize** for YMP (medium vector length)
- **Vectorize** for SIMD (deja vu, long vectors)
- **Vectorize** for CPU/KNL (deja deja vu, short vectors)
- Learn **MPI** (distributed memory)
- Add **OpenMP** directives (modest threading)
- Learn **CUDA** for GPUs (massive threading)
- Overlap comp and comm (hide latencies)
- Manage memory for CPUs (4 level caches and growing)
- Hybrid nodes (CPU + multiple GPUs)
- Convert to **asynchronous multi tasking** (what?)
- Make codes **fault tolerant** (really?)
- **MPI may vanish** (#@!% really??)

Coding apps for the bleeding-edge of HPC

- **Vectorize** for YMP (medium vector length)
- **Vectorize** for SIMD (deja vu, long vectors)
- **Vectorize** for CPU/KNL (deja deja vu, short vectors)
- Learn **MPI** (distributed memory)
- Add **OpenMP** directives (modest threading)
- Learn **CUDA** for GPUs (massive threading)
- Overlap comp and comm (hide latencies)
- Manage memory for CPUs (4 level caches and growing)
- Hybrid nodes (CPU + multiple GPUs)
- Convert to **asynchronous multi tasking** (what?)
- Make codes **fault tolerant** (really?)
- **MPI may vanish** (#@!% really??)

- Hardware/Architects: this is the **price** apps have to pay
  to keep up with our amazing hardware

Coding apps for the bleeding-edge of HPC

- **Vectorize** for YMP (medium vector length)
- **Vectorize** for SIMD (deja vu, long vectors)
- **Vectorize** for CPU/KNL (deja deja vu, short vectors)
- Learn **MPI** (distributed memory)
- Add **OpenMP** directives (modest threading)
- Learn **CUDA** for GPUs (massive threading)
- Overlap comp and comm (hide latencies)
- Manage memory for CPUs (4 level caches and growing)
- Hybrid nodes (CPU + multiple GPUs)
- Convert to **asynchronous multi tasking** (what?)
- Make codes **fault tolerant** (really?)
- **MPI may vanish** (#@!% really??)

- Hardware/Architects: this is the **price** apps have to pay
  to keep up with our amazing hardware
- CS folks: these are really cool **research topics**

Coding apps for the bleeding-edge of HPC

- **Vectorize** for YMP (medium vector length)
- **Vectorize** for SIMD (deja vu, long vectors)
- **Vectorize** for CPU/KNL (deja deja vu, short vectors)
- Learn **MPI** (distributed memory)
- Add **OpenMP** directives (modest threading)
- Learn **CUDA** for GPUs (massive threading)
- Overlap comp and comm (hide latencies)
- Manage memory for CPUs (4 level caches and growing)
- Hybrid nodes (CPU + multiple GPUs)
- Convert to **asynchronous multi tasking** (what?)
- Make codes **fault tolerant** (really?)
- **MPI may vanish** (#@!% really??)

- Hardware/Architects: this is the **price** apps have to pay to keep up with our amazing hardware
- CS folks: these are really cool **research topics**
- App developers: this is a ton of not-so-useful **work**

Coding apps for the bleeding-edge of HPC

- **Vectorize** for YMP (medium vector length)
- **Vectorize** for SIMD (deja vu, long vectors)
- **Vectorize** for CPU/KNL (deja deja vu, short vectors)
- Learn **MPI** (distributed memory)
- Add **OpenMP** directives (modest threading)
- Learn **CUDA** for GPUs (massive threading)
- Overlap comp and comm (hide latencies)
- Manage memory for CPUs (4 level caches and growing)
- Hybrid nodes (CPU + multiple GPUs)
- Convert to **asynchronous multi tasking** (what?)
- Make codes **fault tolerant** (really?)
- **MPI may vanish** (#@!% really??)

- Hardware/Architects: this is the **price** apps have to pay to keep up with our amazing hardware
- CS folks: these are really cool **research topics**
- App developers: this is a ton of not-so-useful **work**
- Scientists: this is a **barrier to the science** I want to do

Qualitative history of apps on evolving HPC platforms

X-axis = paradigm shifts in HPC node hardware
Y-axis = percentage of scientific apps that adapt

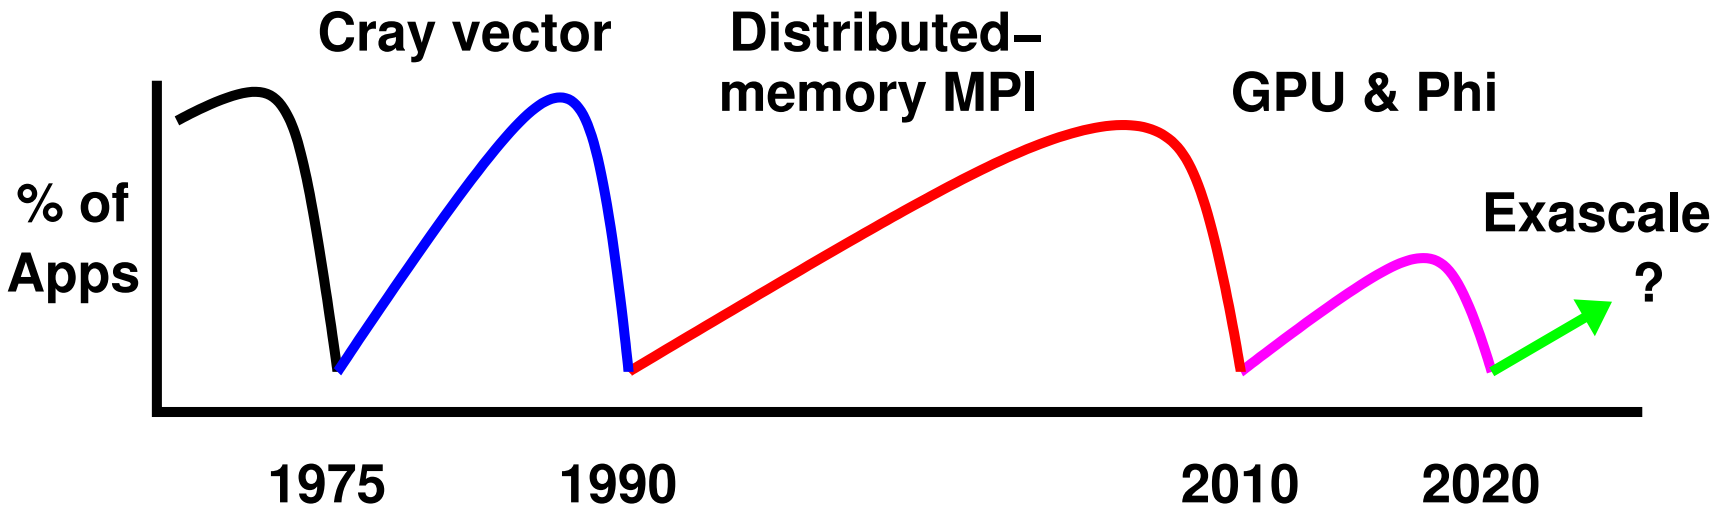

Qualitative history of apps on evolving HPC platforms

X-axis = paradigm shifts in HPC node hardware
Y-axis = percentage of scientific apps that adapt

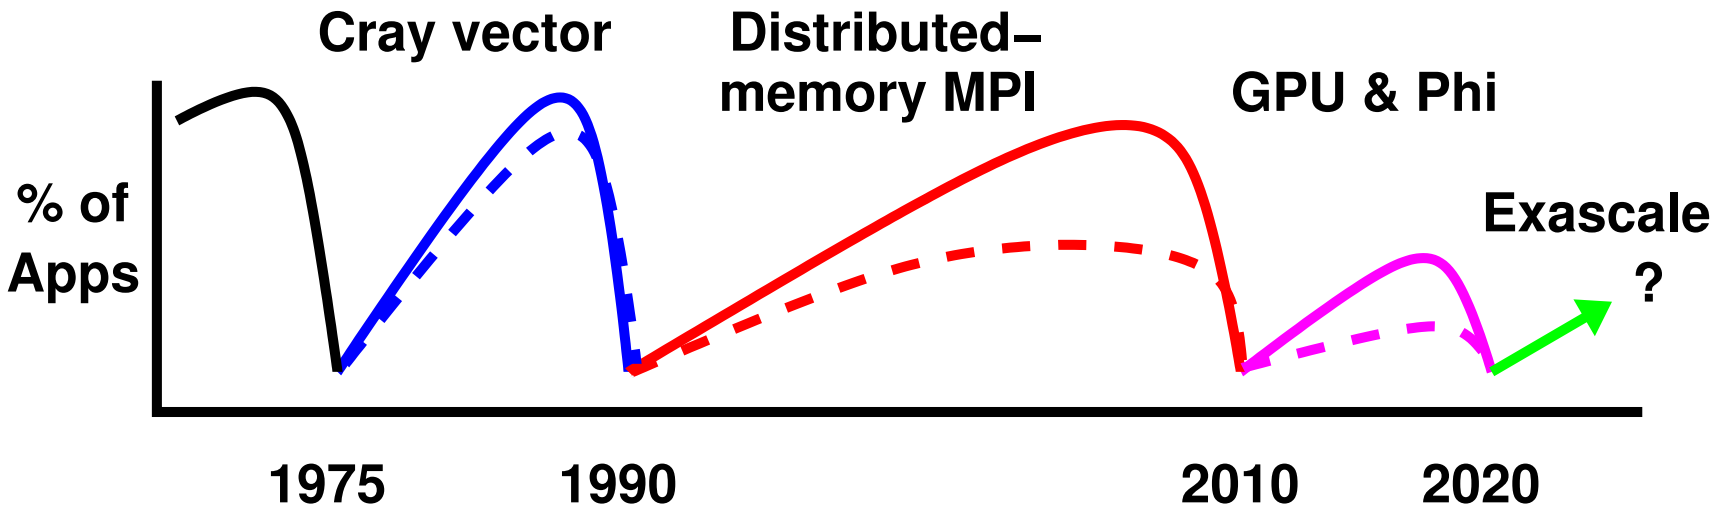

Y-axis = percentage of apps that adapt
and run efficiently on full machine

Why your app might be singing the HPC Blues

Balance ratios on past, present, future HPC platforms

Thanks to Si Hammond (Sandia) for this data!

Why your app might be singing the HPC Blues

Balance ratios on past, present, future HPC platforms

Thanks to Si Hammond (Sandia) for this data!

- **Local balance** = flops to pay for on-node word (8 bytes)
- **Remote balance** = flops to pay for off-node word

The Olde Timey Blues

- **Local balance** = flops to pay for on-node word (8 bytes)
- **Remote balance** = flops to pay for off-node word

Year	Machine	Linpack	Flops/ local	Flops/ remote
1988	Cray YMP	2.1 Gflops	0.52	0.52
1997	ASCI Red (SNL)	1.6 Tflops	8.3	20
2011	RoadRunner (LANL)	1.0 Pflops	6.7	170

Current blues

- **Local balance** = flops to pay for on-node word (8 bytes)
- **Remote balance** = flops to pay for off-node word

Year	Machine	Linpack	Flops/ local	Flops/ remote
1988	Cray YMP	2.1 Gflops	0.52	0.52
1997	ASCI Red (SNL)	1.6 Tflops	8.3	20
2011	RoadRunner (LANL)	1.0 Pflops	6.7	170
2012	Sequoia (LLNL)	17 Pflops	32	160
2013	Titan (ORNL)	18 Pflops	29	490
2018	Summit (ORNL)	122 Pflops	37	1060

Asian blues

- **Local balance** = flops to pay for on-node word (8 bytes)
- **Remote balance** = flops to pay for off-node word

Year	Machine	Linpack	Flops/ local	Flops/ remote
1988	Cray YMP	2.1 Gflops	0.52	0.52
1997	ASCI Red (SNL)	1.6 Tflops	8.3	20
2011	RoadRunner (LANL)	1.0 Pflops	6.7	170
2012	Sequoia (LLNL)	17 Pflops	32	160
2013	Titan (ORNL)	18 Pflops	29	490
2018	Summit (ORNL)	122 Pflops	37	1060
2011	K-Computer (Japan)	11 Pflops	15	95
2013	Tianhe-2 (China)	34 Pflops	22	2100
2016	Sunway TaihuLight (China)	93 Pflops	130	1500

Exascale blues

Year	Machine	Linpack	Flops/ local	Flops/ remote
1988	Cray YMP	2.1 Gflops	0.52	0.52
1997	ASCI Red (SNL)	1.6 Tflops	8.3	20
2011	RoadRunner (LANL)	1.0 Pflops	6.7	170
2012	Sequoia (LLNL)	17 Pflops	32	160
2013	Titan (ORNL)	18 Pflops	29	490
2018	Summit (ORNL)	122 Pflops	37	1060
2011	K-Computer (Japan)	11 Pflops	15	95
2013	Tianhe-2 (China)	34 Pflops	22	2100
2016	Sunway TaihuLight (China)	93 Pflops	130	1500
$\sim$2021	TBD Door #1	1.0 Eflops	80	3200
$\sim$2021	TBD Door #2	1.0 Eflops	300	10000

Exascale blues

Year	Machine	Linpack	Flops/ local	Flops/ remote
1988	Cray YMP	2.1 Gflops	0.52	0.52
1997	ASCI Red (SNL)	1.6 Tflops	8.3	20
2011	RoadRunner (LANL)	1.0 Pflops	6.7	170
2012	Sequoia (LLNL)	17 Pflops	32	160
2013	Titan (ORNL)	18 Pflops	29	490
2018	Summit (ORNL)	122 Pflops	37	1060
2011	K-Computer (Japan)	11 Pflops	15	95
2013	Tianhe-2 (China)	34 Pflops	22	2100
2016	Sunway TaihuLight (China)	93 Pflops	130	1500
$\sim$2021	TBD Door #1	1.0 Eflops	80	3200
$\sim$2021	TBD Door #2	1.0 Eflops	300	10000

- **Good news**: billion X speed-up in 30 years! (vs 4 YMP)

Interpretive blues

- Growing imbalance ratios mean:
  - fewer codes achieve high single-node performance
  - fewer codes achieve good scalability

Interpretive blues

- Growing imbalance ratios mean:
  - fewer codes achieve high single-node performance
  - fewer codes achieve good scalability

- Bottom line: HPC is **selecting** for certain kinds of apps that can withstand these high imbalance ratios

But hey ... growing imbalance is good news for MD

- MD and other **particle apps**:
  - lots of flops per memory access (expensive models)
  - particle/particle interactions are local (comm is local)
  - zillions of particles $\Rightarrow$ lots of threads

But hey ... growing imbalance is good news for MD

- MD and other **particle apps**:
  - lots of flops per memory access (expensive models)
  - particle/particle interactions are local (comm is local)
  - zillions of particles $\Rightarrow$ lots of threads

- So I shouldn't be complaining ...
  we're **thinning the herd** of apps, less competition for cycles

But hey ... growing imbalance is good news for MD

- MD and other **particle apps**:
  - lots of flops per memory access (expensive models)
  - particle/particle interactions are local (comm is local)
  - zillions of particles $\Rightarrow$ lots of threads
- So I shouldn't be complaining ...
  we're **thinning the herd** of apps, less competition for cycles
- But ...
  - particles don't represent **broad swath** of
    computational science, or majority of apps that need HPC
  - physics often isn't short-range
  - hard to reach long timescales with explicit timestepping

Cell biology

- PCR (1983) = polymerase chain reaction, DNA replication
- Microarray chips (1995) = parallel gene expression (millions)
- DNA sequencing (2001) = \$10K/Mb $\Rightarrow$, few \$0.01/Mb
- CRISPR (2012) = genome editing in living cells

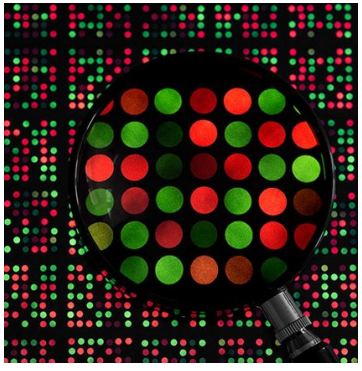

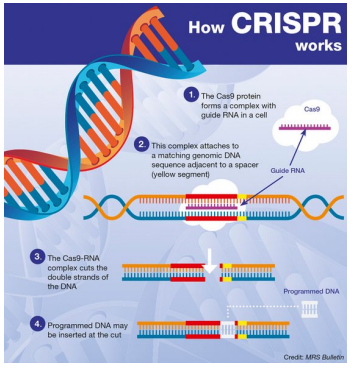

Cell biology

- PCR (1983) = polymerase chain reaction, DNA replication
- Microarray chips (1995) = parallel gene expression (millions)
- DNA sequencing (2001) = \$10K/Mb $\Rightarrow$, few \$0.01/Mb
- CRISPR (2012) = genome editing in living cells

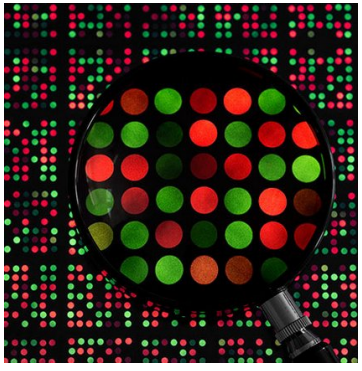

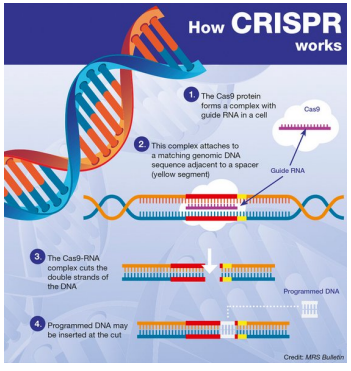

- All these technologies rapidly became **ubiquitous**
- Any lab, any grad student can use them
- Don't need add-on experts to write an NIH proposal

Cell biology

- PCR (1983) = polymerase chain reaction, DNA replication
- Microarray chips (1995) = parallel gene expression (millions)
- DNA sequencing (2001) = \$10K/Mb $\Rightarrow$, few \$0.01/Mb
- CRISPR (2012) = genome editing in living cells

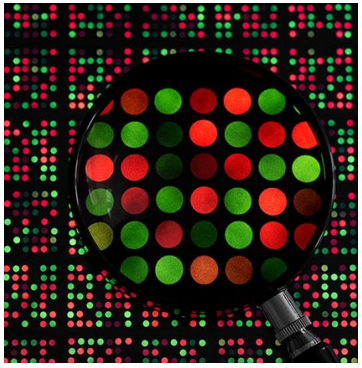

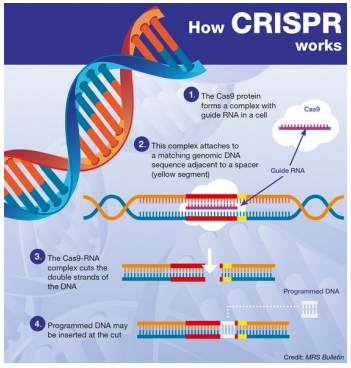

- All these technologies rapidly became **ubiquitous**
- Any lab, any grad student can use them
- Don't need add-on experts to write an NIH proposal
- Could we aspire to that **ease-of-use** for HPC machines?

User facilities with billion $ instruments

- Hubble telescope (NASA/ESA),
  SNS (ORNL), Z-machine (Sandia)

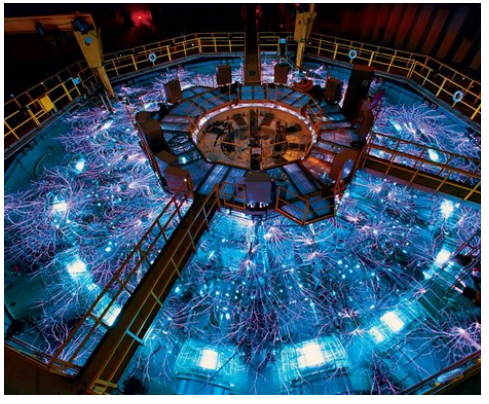

User facilities with billion $ instruments

- Hubble telescope (NASA/ESA),
  SNS (ORNL), Z-machine (Sandia)



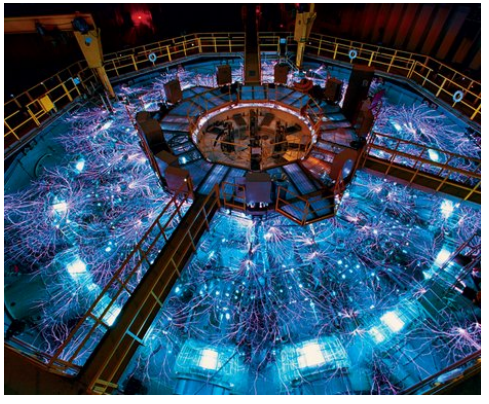

- Hubble: 1.3M observations, SNS: 20K users, Z: 3160 shots

User facilities with billion $ instruments

- Hubble telescope (NASA/ESA),
  SNS (ORNL), Z-machine (Sandia)



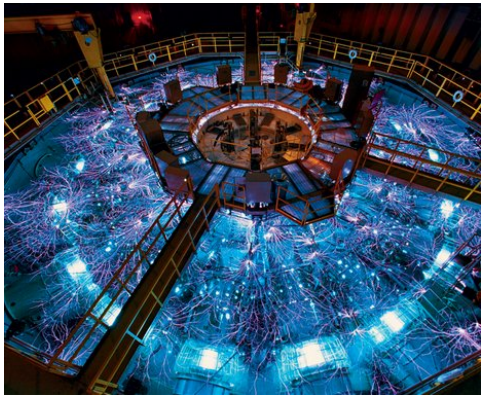

- Hubble: 1.3M observations, SNS: 20K users, Z: 3160 shots
- All solicit user proposals (Hubble from amateurs!)
- Facilities **shield users** from nearly all **complexity**

User facilities with billion $ instruments

- Hubble telescope (NASA/ESA),
  SNS (ORNL), Z-machine (Sandia)



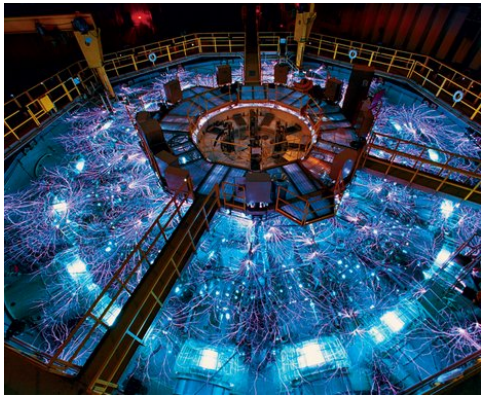

- Hubble: 1.3M observations, SNS: 20K users, Z: 3160 shots
- All solicit user proposals (Hubble from amateurs!)
- Facilities **shield users** from nearly all **complexity**
- What if 20x new HPC machine just gave all users 20x more?

High-energy particle physics

- CERN, FermiLab, etc

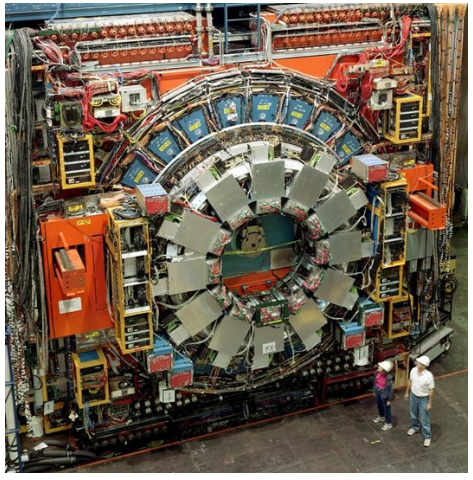

- Every new accelerator requires **one-of-a-kind new detectors** to be useful
- Detector = 100s of people, $100 million or more
- Performs handful of (high-impact, highly complex) science experiments in a narrow sub-field of physics

High-energy particle physics

- CERN, FermiLab, etc

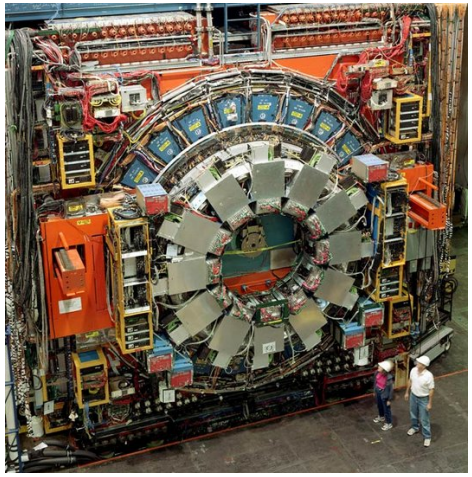

- Every new accelerator requires **one-of-a-kind new detectors** to be useful
- Detector = 100s of people, $100 million or more
- Performs handful of (high-impact, highly complex) science experiments in a narrow sub-field of physics
- **Is HPC more like** cell bio, user facilities, or HE physics?

Thanks

- Hope you view my remarks as **inducements** to:
  - insulate users from growing complexity of HPC machines
  - make life easier for the apps and the science

Thanks

- Hope you view my remarks as **inducements** to:
  - insulate users from growing complexity of HPC machines
  - make life easier for the apps and the science
- Funding from DOE **exascale computing program**

- **Hyperdynamics** collaborators:
  Art Voter, Danny Perez (LANL)
- **LAMMPS** collaborators:
  Aidan Thompson, Stan Moore, Mitch Wood (Sandia)
  Axel Kohlmeyer (Temple U)