Challenges for Materials Modeling at the Atomic and Meso Scales

SAND2018-0867PE

Steve Plimpton
Sandia National Laboratories

Applied Math seminar
University of New Mexico – Jan 2018

CCR
Center for Computing Research

Sandia National Laboratories

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.
Presentation: SAND2015-9811PE

NNSA
National Nuclear Security Administration

See you at the movies ...

CGI modeling advances by Pixar

Bug's Life (1998)
vegetation

Monsters, Inc (2001)
hair

Finding Nemo (2003)
water

Cars (2006)
painted surfaces

Ratatouille (2007)
food

Wall-E (2008)
rust & decay

Quantum $\Rightarrow$ MD $\Rightarrow$ kMC $\Rightarrow$ Continuum

- Quantum density functional theory (DFT) = $\sim$100 atoms
  - **Schrodinger equation (PDE)** for electron density
  - $O(N^2)$ to $O(N^3)$ scaling with # of valence electrons

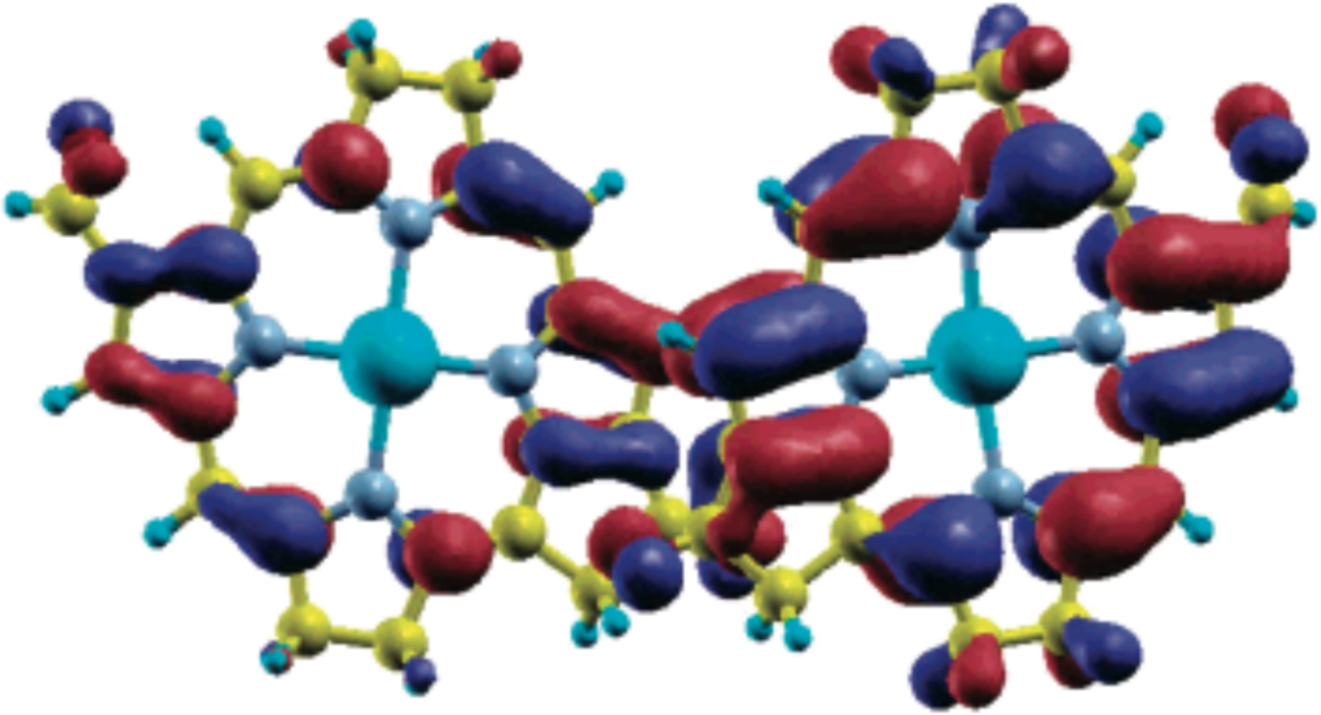

Quantum $\Rightarrow$ **MD** $\Rightarrow$ kMC $\Rightarrow$ Continuum

- Classical molecular dynamics (MD) $= \sim 10^3$ to 10^9 atoms
  - atoms or coarse-grained particles, empirical forces
  - time-integrate **ODEs of Newtonian mechanics**
  - $O(N)$ scaling with # of particles

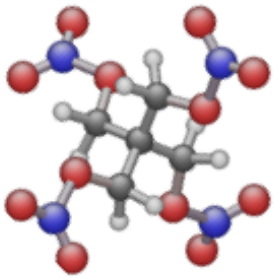

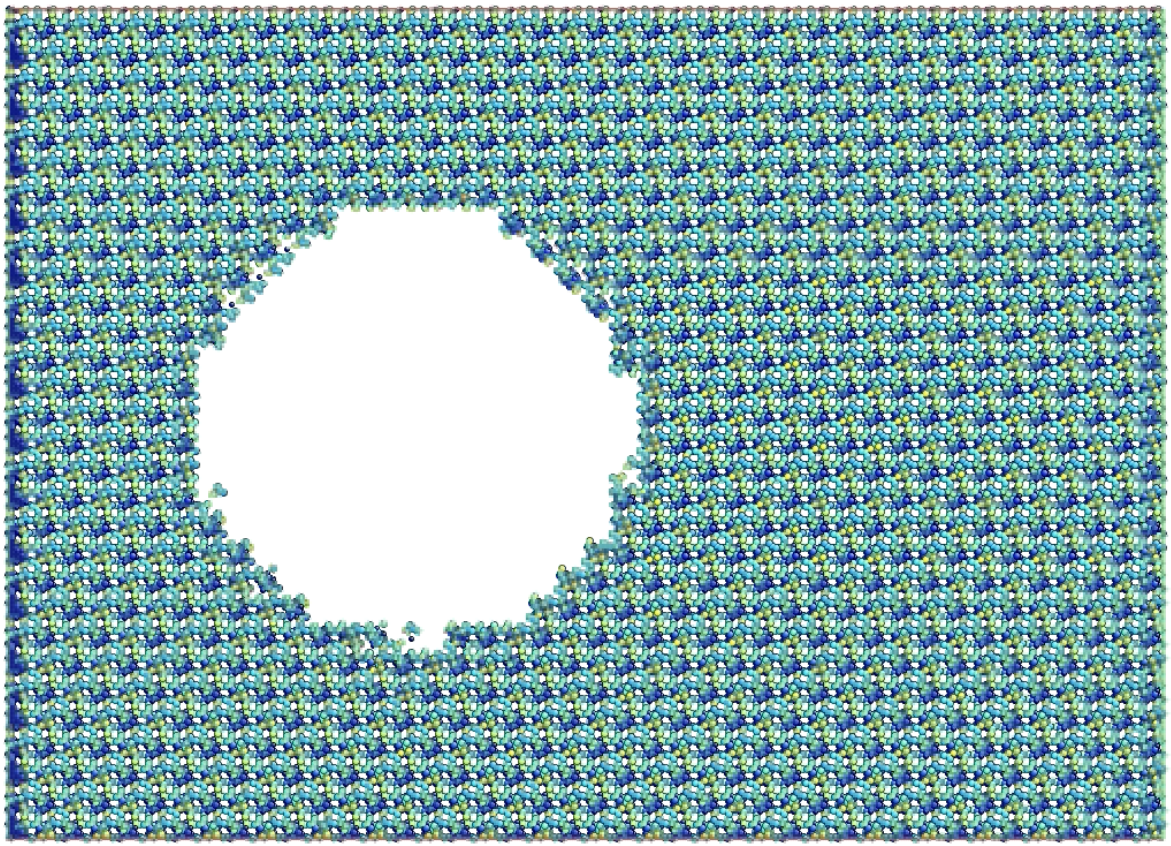

Quantum $\Rightarrow$ MD $\Rightarrow$ **kMC** $\Rightarrow$ Continuum

- Kinetic Monte Carlo (kMC) $= \sim 10^6$ to 10^9 lattice sites
  - site = atom or meso-scale chunk of material
  - **Hamiltonian** to describe site-site interaction energies
  - $O(N)$ scaling with # of sites, due to local interactions

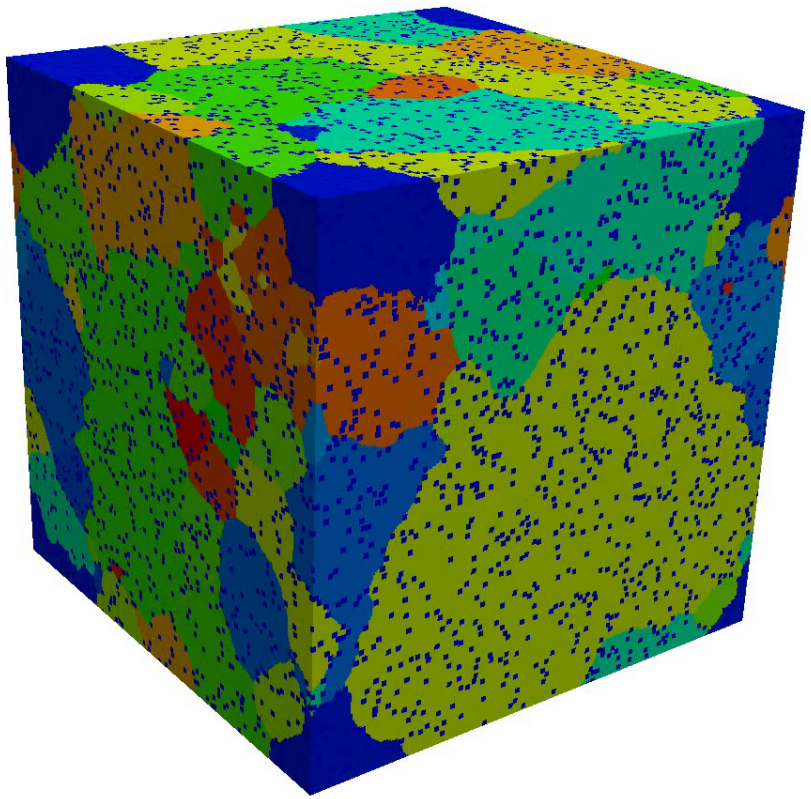

Quantum $\Rightarrow$ MD $\Rightarrow$ kMC $\Rightarrow$ **Continuum**

- Continuum PDEs $= \sim 10^6$ to 10^9 grid cells
  - deformation stress/strain, Navier-Stokes, etc
  - discretize by **finite difference (FD), finite elements (FE)**
  - $O(N)$ scaling with # of grid cells,
    explicit or implicit (optimal multigrid, Krylov DD)

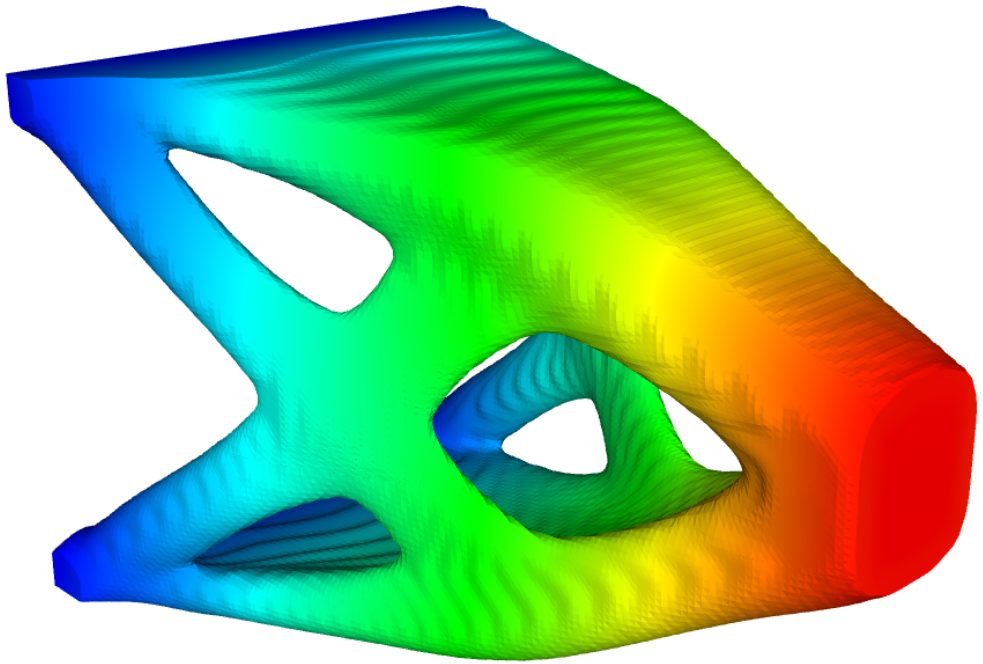

Recurrent challenges for 2 methods in the "middle"

- **Increase MD accuracy**
  - relative to quantum DFT = gold standard
- **Increase MD and kMC length & time scales**
  - via new models, new algorithms
  - aiming for continuum-scale validation & rigor

Recurrent challenges for 2 methods in the "middle"

- **Increase MD accuracy**
  - relative to quantum DFT = gold standard
- **Increase MD and kMC length & time scales**
  - via new models, new algorithms
  - aiming for continuum-scale validation & rigor

Remainder of talk: **3 examples**

1. Quantum-accurate empirical potentials for MD
2. Accelerated-time hyperdynamics algorithm for MD
3. Efficient and accurate algorithms for kMC

Moore's Law for interatomic potentials (force fields)

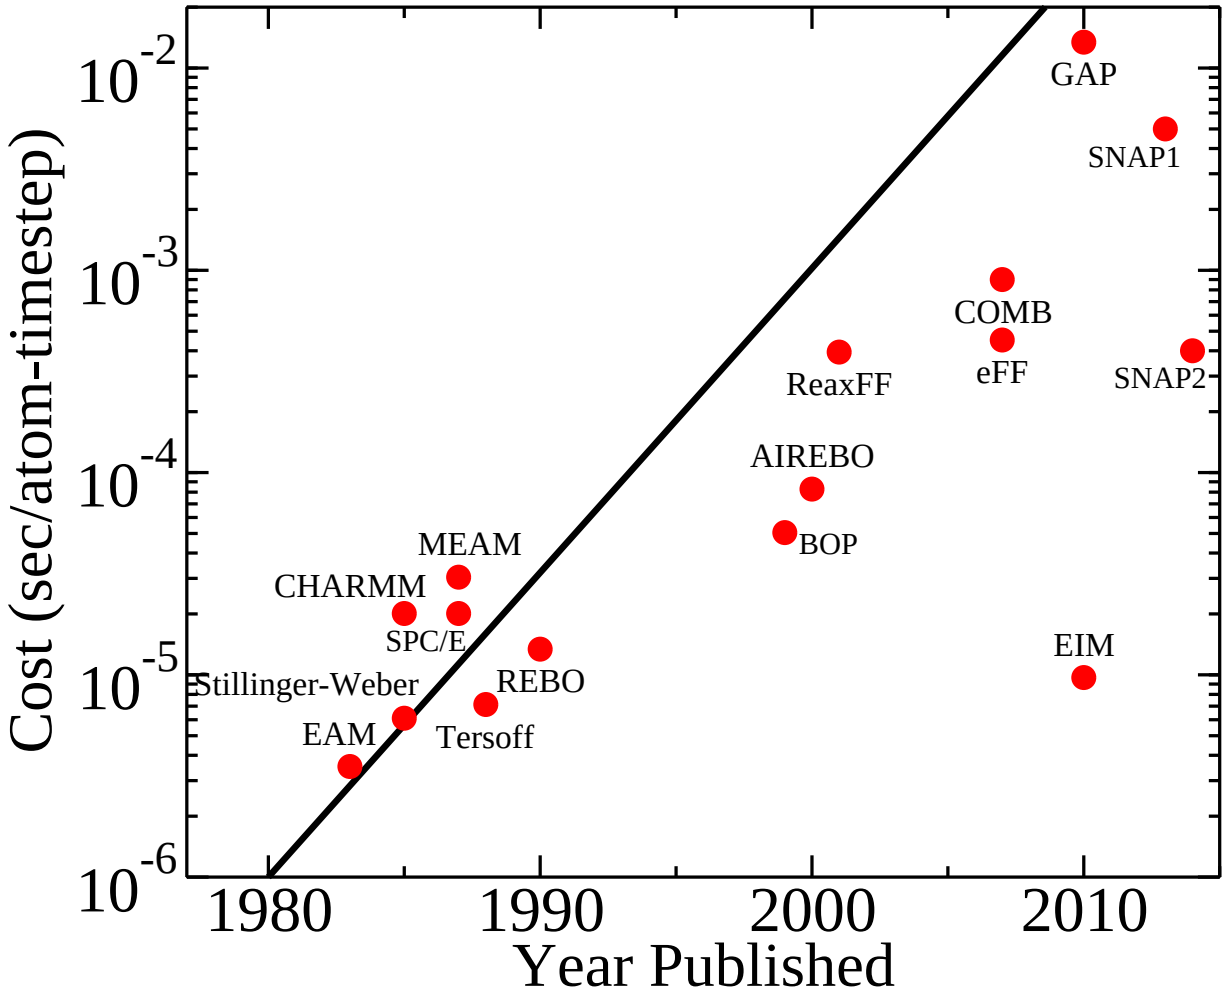

Moore's Law for interatomic potentials (force fields)

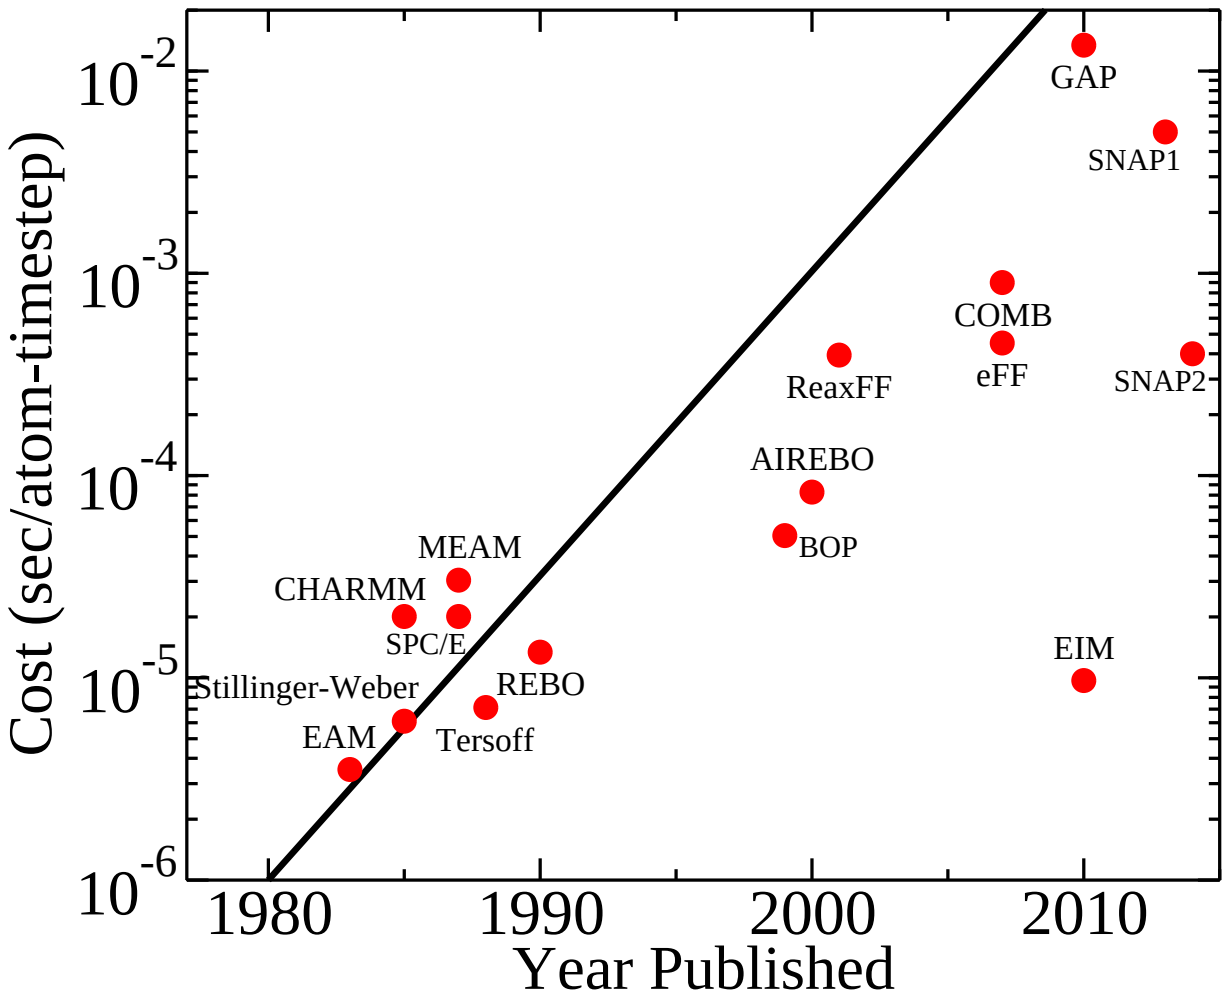

quantum DFT is 10^6 slower for 400 atoms, 10^9 for 40000, etc

Traditional interatomic potentials (force fields)

- For time-integration of N particles, just need ...
- Potential = energy of system, derivatives = forces

$$E_{total} = \Phi(\mathbf{R}_N) = \sum_{i,j} \phi_2(r_i, r_j) + \sum_{i,j,k} \phi_3(r_i, r_j, r_k) + \sum \phi_4() + \cdots$$

- Truncate many-body interactions, distance cutoff $\Rightarrow$ $O(N)$

- Lennard-Jones (2 params)
  to ReaxFF (100s of params)

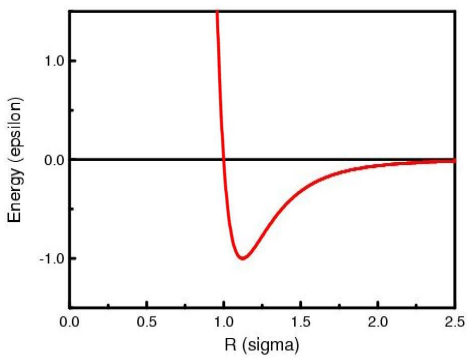

Traditional interatomic potentials (force fields)

- For time-integration of N particles, just need ...
- Potential = energy of system, derivatives = forces

$$E_{total} = \Phi(\mathbf{R}_N) = \sum_{i,j} \phi_2(r_i, r_j) + \sum_{i,j,k} \phi_3(r_i, r_j, r_k) + \sum \phi_4() + \cdots$$

- Truncate many-body interactions, distance cutoff $\Rightarrow$ $O(N)$

- Lennard-Jones (2 params)
  to ReaxFF (100s of params)

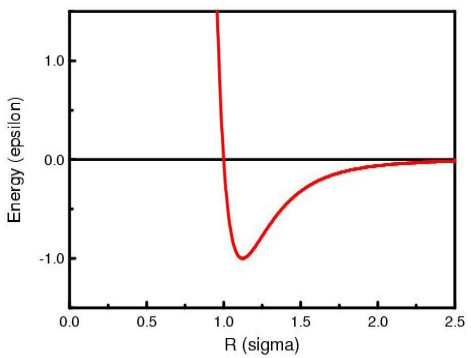

- Parameters fit to:
  - experimental data: density, elastic constants, melting temp, etc
  - quantum: bond stretch/stiffness, reaction barriers, etc
- **Big bottleneck** for materials modeling via classical MD

Quantum-accuracy with classical potentials?

- **Traditional** potential
  - functional form based on physical insight
  - handful of params fit to specific QM calcs
- **Machine-learned** potential
  - generic descriptors span conformational space
  - fit via ML to large database of QM calcs

Quantum-accuracy with classical potentials?

- **Traditional** potential
  - functional form based on physical insight
  - handful of params fit to specific QM calcs
- **Machine-learned** potential
  - generic descriptors span conformational space
  - fit via ML to large database of QM calcs
- **GAP** = Gaussian approximation potential
  - Csányi & Bartók-Partay (U Cambridge)
  - *Phys Rev Lett, 104, 136403 (2010)*
  - Gaussian process = high-dimensional interpolation method
- **SNAP** = spectral neighbor analysis potential
  - Aidan Thompson, Mitch Wood (Sandia)
  - *J Comp Phys, 285, 316-330 (2015)*
  - fit via simpler linear regression plus genetic algorithm wrapper

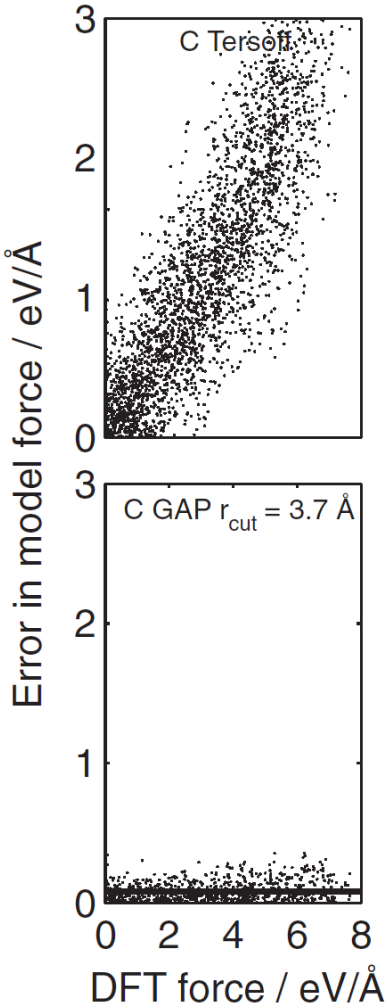

How SNAP potentials work

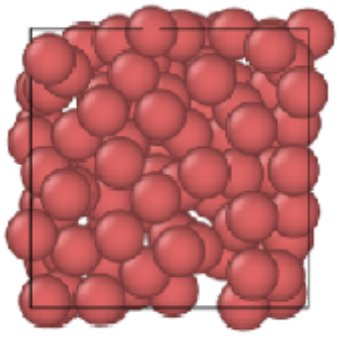

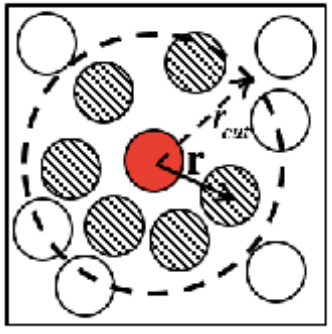

$$E^{SNAP} = \sum_{i=1}^{N} E_i^{SNAP} + \sum_{j<i}^{N} \phi_{ij}^{rep}(r_{ij})$$

$$E_i^{SNAP} = \beta_0 + \sum_{k \in \{J < J_{max}\}} \beta_k B_k^i$$

- Represent neighborhood via 4D hyperspherical harmonics
- B_k^i = **descriptors** = bispectrum components for atom I

How SNAP potentials work

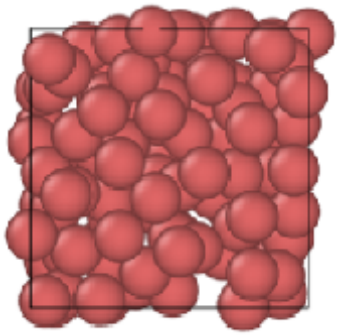

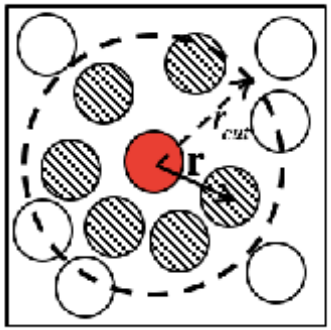

$$E^{SNAP} = \sum_{i=1}^{N} E_i^{SNAP} + \sum_{j<i}^{N} \phi_{ij}^{rep}(r_{ij})$$

$$E_i^{SNAP} = \beta_0 + \sum_{k \in \{J < J_{max}\}} \beta_k B_k^i$$

- Represent neighborhood via 4D hyperspherical harmonics
- B_k^i = **descriptors** = bispectrum components for atom I

- Pre-compute **database of DFT results** for small systems
  - 1000s of wisely-chosen configurations (up to 100s of atoms)
- Each config gives gold-standard F_i, E_{total}, $\mathbf{W}_{ij}$
- Compute optimal β_k (unknowns) via least-squares
- Embed fitting in **DAKOTA** GA to optimize hyper-parameters
  - weights on QM configs, cutoff, # of basis functions
- Expensive, but still linear $O(N)$ cost in # of atoms
- SNAP reduces cost by ~10x via symmetry exploitation

Success with SNAP potential for Tantalum

- Energy barrier for **screw dislocation migration** in bcc Ta
- *Thompson, et al, J Comp Phys, 285, 316-330 (2015)*

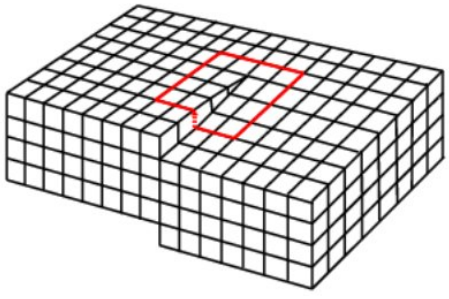

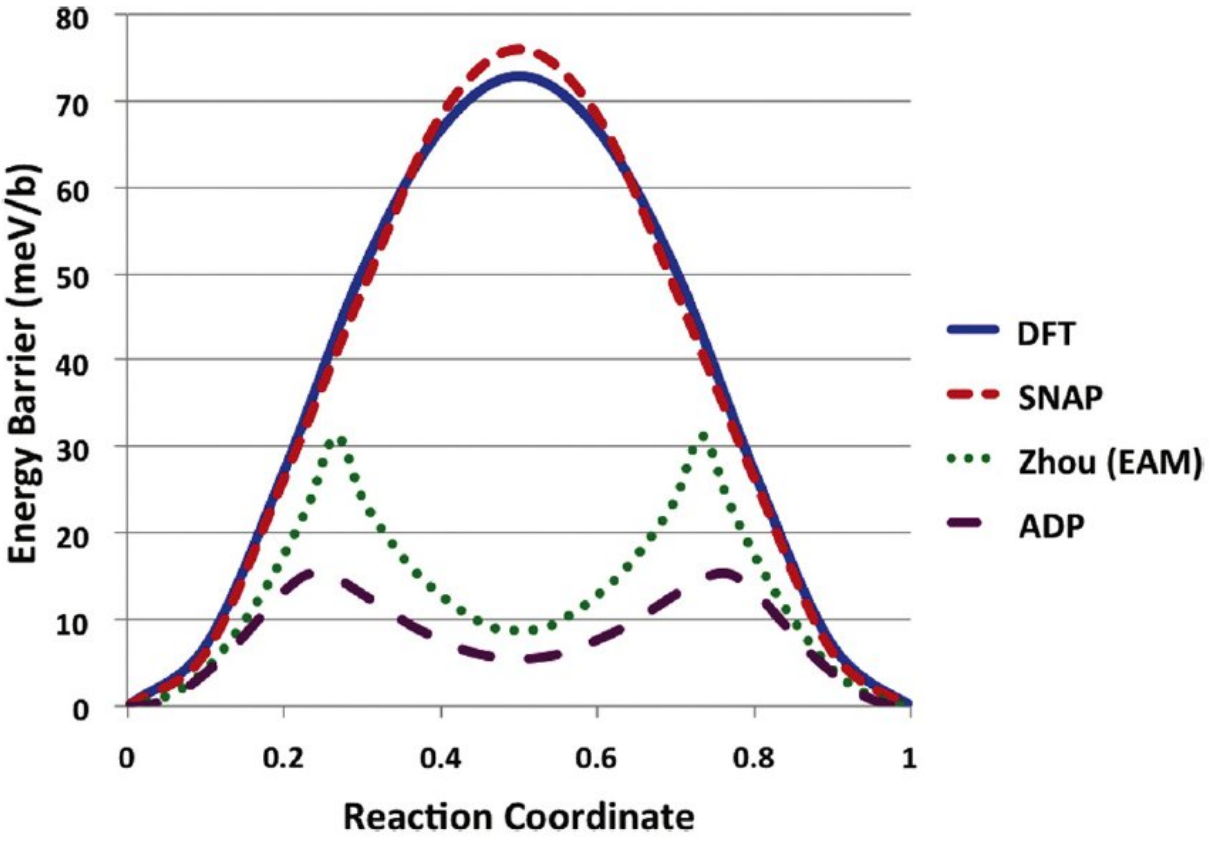

- Note: screw dislocations were **not** in QM database

Recent SNAP potential for W and W/He

	DFT[a]	EAM[b]	MEAM[c]	GAP[d]	SNAP[e]
a_0 (Å)	3.1803	3.165	3.164	3.1803	3.1806
C_{11} (GPa)	517	517	533	518	518
C_{12} (GPa)	198	200	205	198	196
C_{44} (GPa)	142	156	163	143	144

	DFT[a]	EAM[b]	MEAM[c]	GAP[d]	SNAP[e]	SNAP[f]
$E_{W-Tetra}$	11.1	10.4	11.2	12.6	9.8	9.8
E_{W-Octa}	11.7	10.4	9.0	12.9	11.3	11.3
$E_{W-[110]d}$	9.8	10.1	8.9	11.6	9.2	9.2
$E_{W-[111]d}$	9.6	9.4	9.0	11.0	9.5	9.5
$E_{W-vac.}$	3.3	3.6	3.9	3.3	3.2	3.2
$E_{W-divac.}$[g]	0.1	-0.4	-0.2	-0.3	0.1	0.1
$E_{He-Tetr}$	6.2	6.2	6.3	6.2	6.1	6.3
E_{He-Oct}	6.4	6.3	6.7	6.6	6.4	6.5
$E_{He-Subs}$	4.7	4.7	5.2	4.3	4.2	3.8
$E_{2He-Tetr}$[h]	1.0	0.9	0.7	0.7	0.8	1.1

- Only match to sign of di-vacancy binding

Wood & Thompson, "QM accurate MD potential for W and He", arXiv 1702.07042

Recent SNAP potential for W and W/He

	DFT[a]	EAM[b]	MEAM[c]	GAP[d]	SNAP[e]
a_0 (Å)	3.1803	3.165	3.164	3.1803	3.1806
C_{11} (GPa)	517	517	533	518	518
C_{12} (GPa)	198	200	205	198	196
C_{44} (GPa)	142	156	163	143	144

	DFT[a]	EAM[b]	MEAM[c]	GAP[d]	SNAP[e]	SNAP[f]
$E_{W-Tetra}$	11.1	10.4	11.2	12.6	9.8	9.8
E_{W-Octa}	11.7	10.4	9.0	12.9	11.3	11.3
$E_{W-[110]d}$	9.8	10.1	8.9	11.6	9.2	9.2
$E_{W-[111]d}$	9.6	9.4	9.0	11.0	9.5	9.5
$E_{W-vac.}$	3.3	3.6	3.9	3.3	3.2	3.2
$E_{W-divac.}$[g]	0.1	-0.4	-0.2	-0.3	0.1	0.1
$E_{He-Tetr}$	6.2	6.2	6.3	6.2	6.1	6.3
E_{He-Oct}	6.4	6.3	6.7	6.6	6.4	6.5
$E_{He-Subs}$	4.7	4.7	5.2	4.3	4.2	3.8
$E_{2He-Tetr}$[h]	1.0	0.9	0.7	0.7	0.8	1.1

Wood & Thompson, "QM accurate MD potential for W and He", arXiv 1702.07042

- Only match to sign of di-vacancy binding
- **Q**: why SNAP better than GAP?
- **A**: additional QM conformations, including multiple vacancies
- 3.8 $\Rightarrow$ 4.9 for combined W/He potential

Complications

- How to insure a **sufficiently complete** database of quantum calculations?
- **Knobs** to turn when creating SNAP potential:
  - weighting of different quantum conformations
  - how many bispectrum components (20,30,more)
  - how big a cutoff $\Rightarrow$ # of influencing neighbors

- Use DAKOTA optimization toolkit to automate knob tuning
  - iterate over fitting procedure many times
  - **https://dakota.sandia.gov**

Research challenges for quantum-accurate potentials

1. Can we answer **UQ**-style questions?
   - quantify errors induced by fitting & database lookup
   - how do errors depend on quantity/coverage of quantum data?
   - how do errors propagate to MD statistics:
     barrier heights, diffusion coeffs, etc

Research challenges for quantum-accurate potentials

1. Can we answer **UQ**-style questions?
   - quantify errors induced by fitting & database lookup
   - how do errors depend on quantity/coverage of quantum data?
   - how do errors propagate to MD statistics:
     barrier heights, diffusion coeffs, etc
2. **On-demand quantum** calculations?
   - detect when interpolation of MD config is inaccurate
   - trigger new "just-in-time" QM calculation to reduce error
   - re-fit to more quantum data incrementally

Research challenges for quantum-accurate potentials

1. Can we answer **UQ**-style questions?
   - quantify errors induced by fitting & database lookup
   - how do errors depend on quantity/coverage of quantum data?
   - how do errors propagate to MD statistics:
     barrier heights, diffusion coeffs, etc
2. **On-demand quantum** calculations?
   - detect when interpolation of MD config is inaccurate
   - trigger new "just-in-time" QM calculation to reduce error
   - re-fit to more quantum data incrementally
3. **Fully automated** creation of new transferable potentials?
   - start with small database of QM results, inaccurate potential
   - run MD simulation, trigger new QM as needed
   - initially: lots of new QM calculations, refits of potential
   - over time: database fills, do only classical MD

Research challenges for quantum-accurate potentials

1. Can we answer **UQ**-style questions?
   - quantify errors induced by fitting & database lookup
   - how do errors depend on quantity/coverage of quantum data?
   - how do errors propagate to MD statistics:
     barrier heights, diffusion coeffs, etc
2. **On-demand quantum** calculations?
   - detect when interpolation of MD config is inaccurate
   - trigger new "just-in-time" QM calculation to reduce error
   - re-fit to more quantum data incrementally
3. **Fully automated** creation of new transferable potentials?
   - start with small database of QM results, inaccurate potential
   - run MD simulation, trigger new QM as needed
   - initially: lots of new QM calculations, refits of potential
   - over time: database fills, do only classical MD
4. **Holy grail**: for any new material, automatic quantum-accurate potentials with cost/scaling of classical MD

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
- **Bond-boost** formulation
  - *Miron & Fichthorn, J Chem Phys, 119, 6210 (2003)*
  - variant: apply bias to only one or few 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

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
- **Bond-boost** formulation
  - *Miron & Fichthorn, J Chem Phys, 119, 6210 (2003)*
  - variant: apply bias to only one or few 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
- Key **caveat**:
  - applicable to solids with distinct energy basins, rare transitions from one basin to another
- Effective speed-up can be **orders of magnitude**
  - especially for high barriers and low temperatures
  - accelerating a single trajectory

Pictorial view of bond-boost hyperdynamics

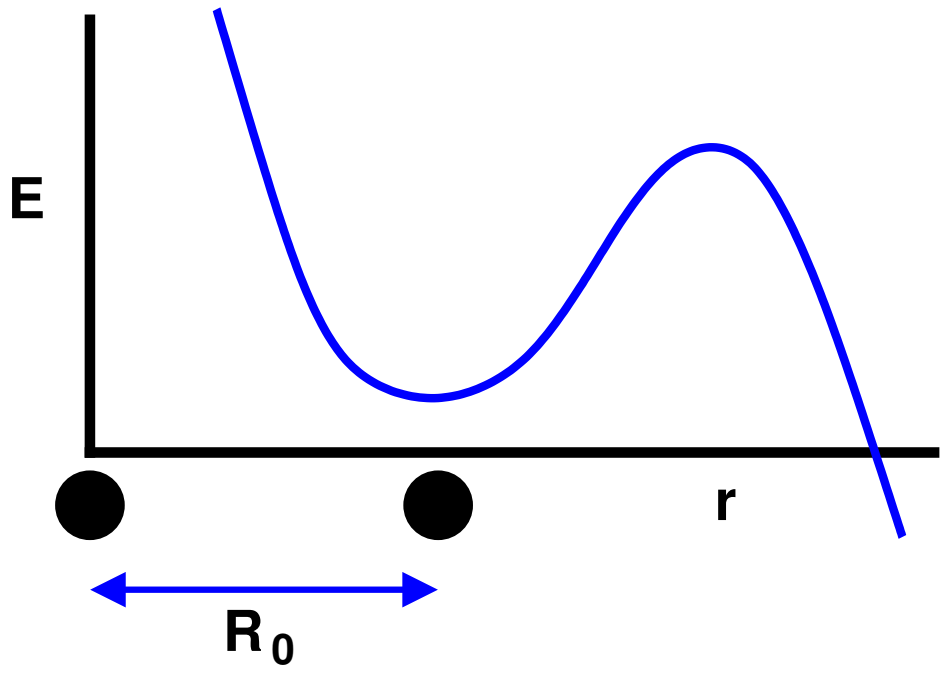

- Define (conceptual) **bonds** between all pairs of nearby atoms
- For example, $\sim$12 nearest neighbors per atom in fcc lattice

Pictorial view of bond-boost hyperdynamics

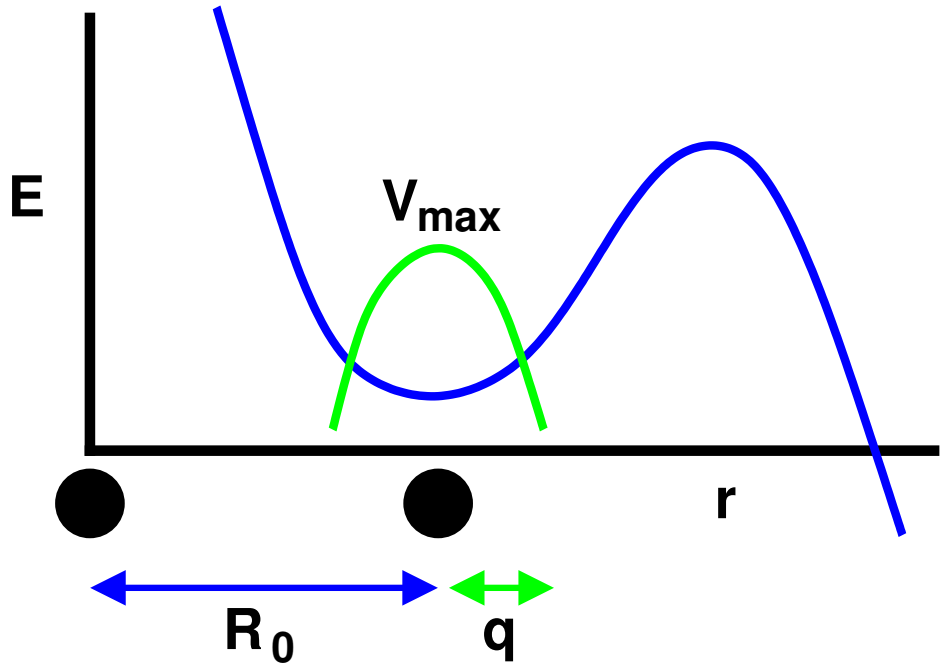

- Bond strain: $\epsilon_{ij} = (R_{ij} - Ro_{ij})/Ro_{ij}$
- Add **bias potential** to two atoms in max-strain bond:
  $$V_{ij} = V_{max}[1 - (\epsilon_{ij}/q)^2], \quad |\epsilon_{ij}| < q, \quad \text{else } 0$$
- Bias may be added to different bond each timestep

Pictorial view of bond-boost hyperdynamics

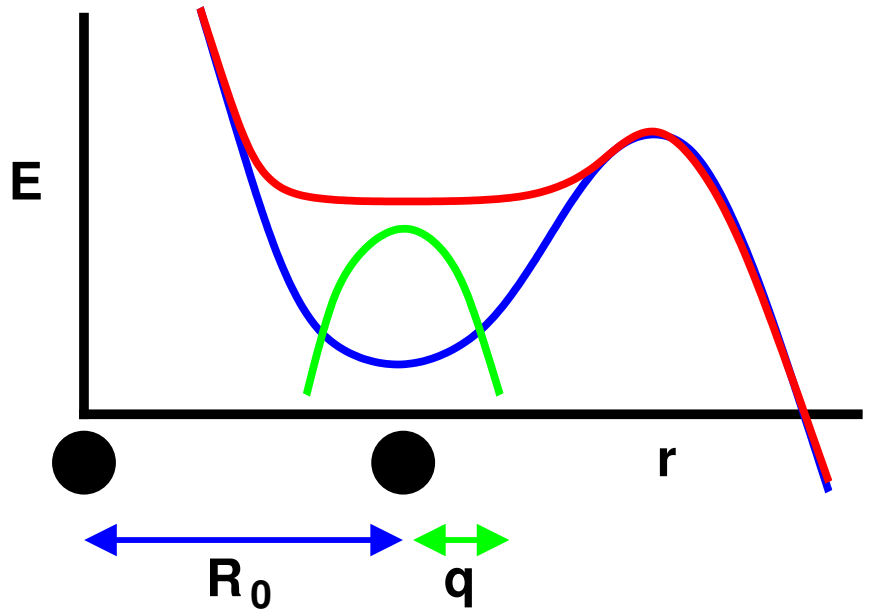

- Shallow V_{ij} allows faster transition by I,J (and nearby) atoms
- Must choose V_{max} and q carefully:
  - if: zero bias at dividing surfaces
  - if: do not induce TST-violating correlated events
  - then: relative transition rates not altered
  - then: quantifiable **time boost factor** each timestep
  - then: resulting traj is time-accurate (unlike enhanced sampling)

Local hyperdynamics in parallel

Local HD **algorithm** loop:

- Run 100 steps of MD with local HD bias
  - bias all bonds which are max strain within D_{cut} region
  - adjust per-bond V_{max} to boost time uniformly everywhere
- Extra operations for HD:
  - **2nd neighbor list** out to D_{cut}
  - **double loops** over atoms and big neighbor list
  - **comm** to acquire bond-strain info for ghost atoms
- Quench to check if **event** has occurred
  - if yes: re-form bond list, preserving bond prefactors

Local hyperdynamics in parallel

Local HD **algorithm** loop:

- Run 100 steps of MD with local HD bias
  - bias all bonds which are max strain within D_{cut} region
  - adjust per-bond V_{max} to boost time uniformly everywhere
- Extra operations for HD:
  - **2nd neighbor list** out to D_{cut}
  - **double loops** over atoms and big neighbor list
  - **comm** to acquire bond-strain info for ghost atoms
- Quench to check if **event** has occurred
  - if yes: re-form bond list, preserving bond prefactors

Cost of HD versus regular MD:

- For cheap EAM, about 2x more expensive
  - half is quench & event search every 100 steps
  - half is comp/comm to find max-strain bonds
  - 2x shrinks for more expensive potentials
- Strong and weak scaling essentially same as MD

Pt surface diffusion model with EAM

- 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 $\Rightarrow$ **4000x boost**
- 1.2M atoms, 50M timesteps $\Rightarrow$ **1 ms** of real time

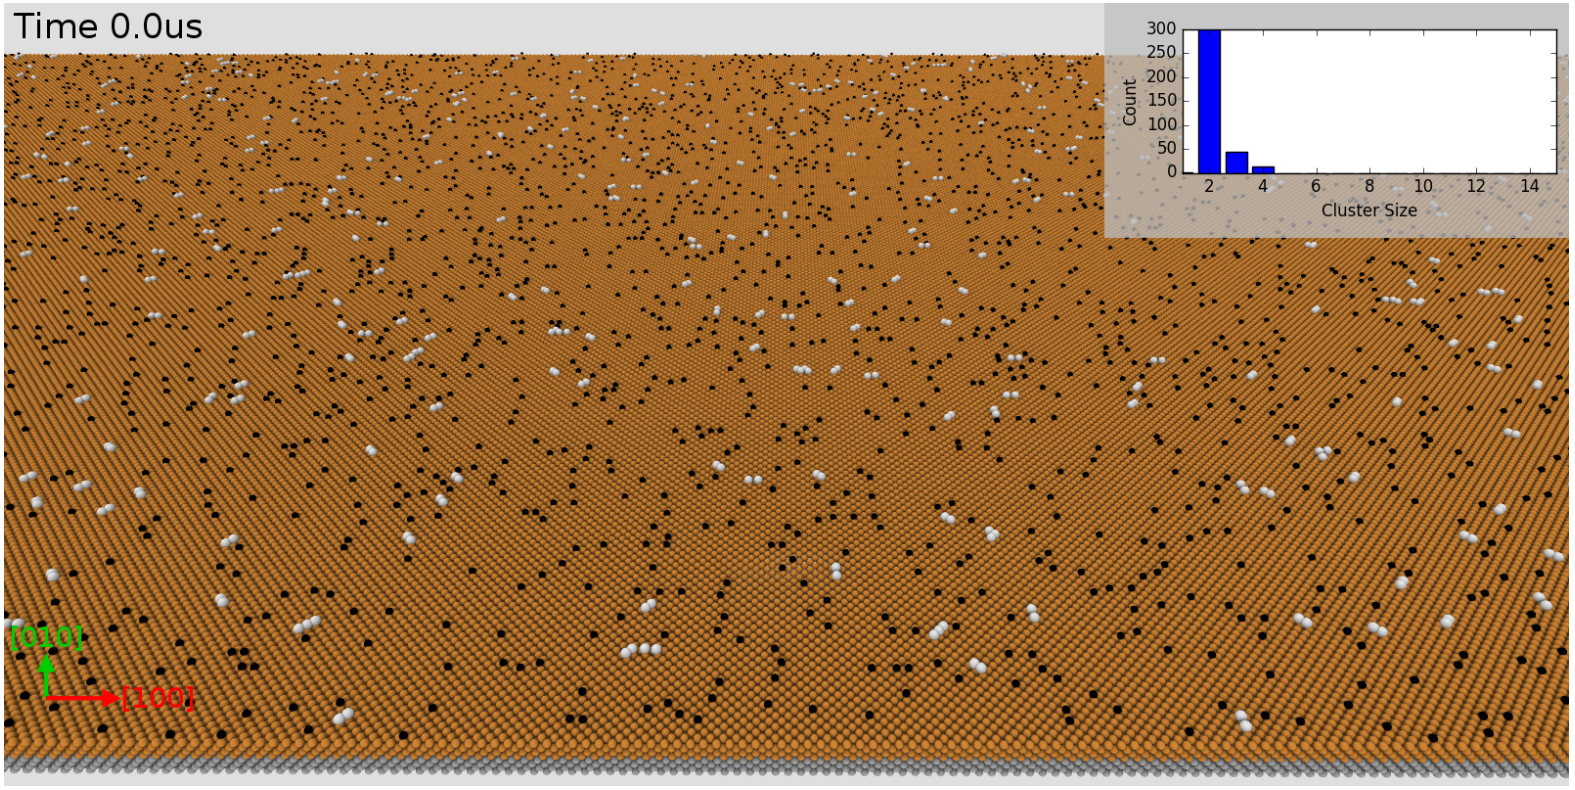

Pt surface diffusion model with EAM

- 24 hrs on 128 Broadwell nodes (4096 cores, 300 atoms/core)
- Observed ~400K diffusion events with 5900 adatoms

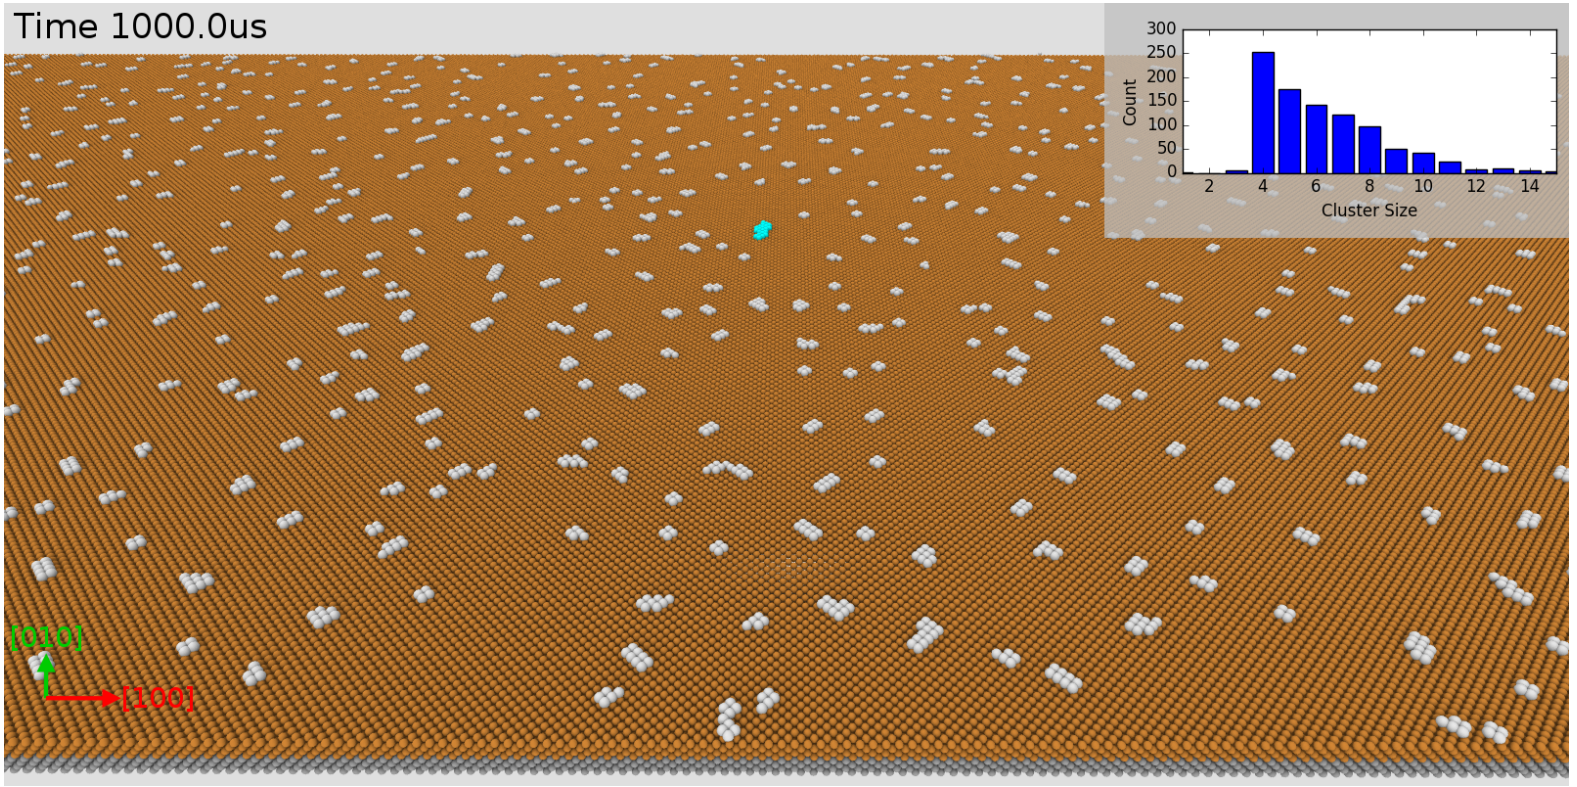

- Can use global/local HD with **any potential** in LAMMPS

Materials modeling via on-lattice kMC

kMC = kinetic Monte Carlo

- Variety of models for time evolution of processed materials
- Length scale = atomic to mesoscale
- Time scale = microseconds to days
- Hamiltonian encodes physics of local interactions
- Time accurate (via event rates) versus Metropolis MC

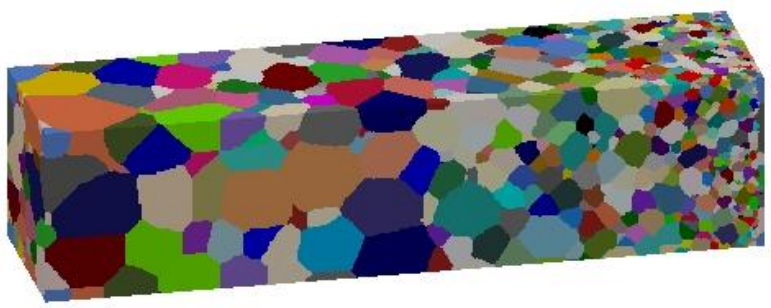

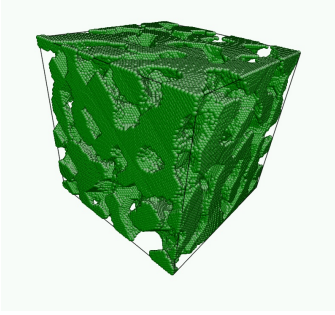

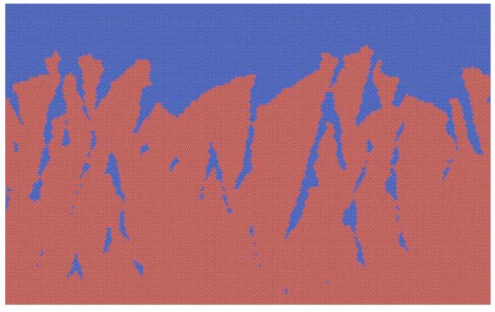

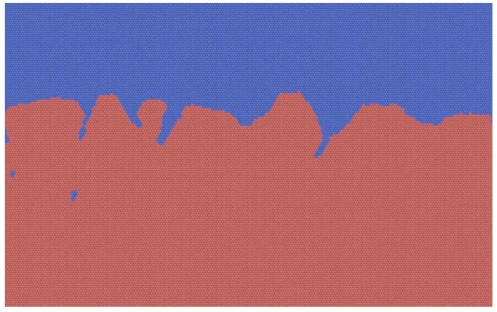

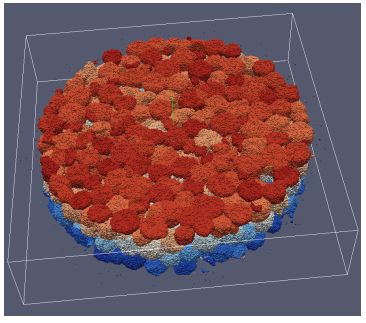

kMC solvers

- **KMC algorithm**:
  - select one of N events with correct relative probability
  - $\Delta t = -\ln(\mathrm{RN})/P_{total}$
  - update affected events

$P_{total} = \sum P_i$

kMC solvers

- **KMC algorithm**:
  - select one of N events with correct relative probability
  - $\Delta t = -\ln(\mathrm{RN})/P_{total}$
  - update affected events

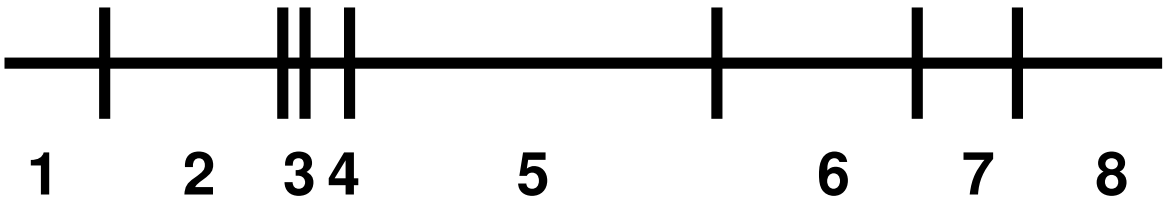

$P_{total} = \sum P_i$

- Serial algorithms are **exact**, but not parallel
- Parallel algorithms are **approximate**
  - deviate from exact event selection with "acceptable" error
  - acceptable to me, may not be acceptable to you
  - exploit **local** nature of many materials models

kMC algorithm arises in many contexts

- BKL algorithm for **Ising model spin updates**
  - *Bortz, Kalos, Lebowitz, J Comp Phys, 17, 10-18 (1975)*
  - event = flip one or a set of spins on lattice
- SSA for stochatic integration of coupled ODEs
  - *D Gillespie, J Chem Phys, 22, 403-434 (1976)*
  - event = which reaction occurs next
  - **biochemical reaction pathways**
- TC (time counter) algorithm for DSMC
  - *G Bird, Phys Fluids, 6, 1518 (1963)*
  - event = collision between 2 particles
  - **low density fluid flow**

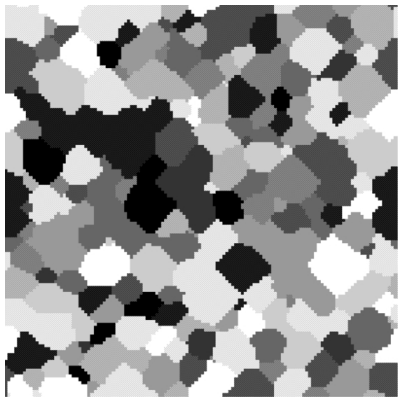

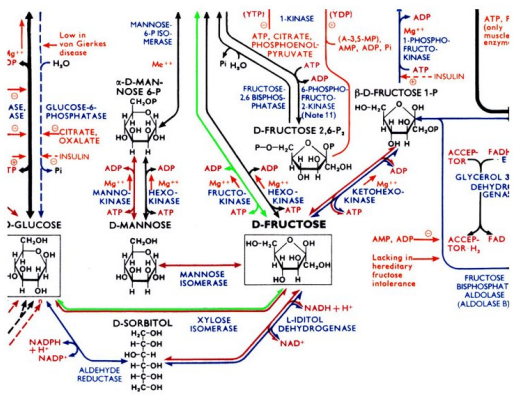

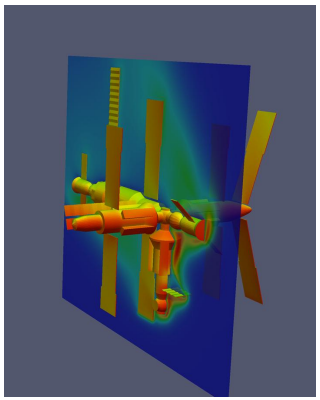

kMC algorithm arises in many contexts

- BKL algorithm for **Ising model spin updates**
  - *Bortz, Kalos, Lebowitz, J Comp Phys, 17, 10-18 (1975)*
  - event = flip one or a set of spins on lattice
- SSA for stochatic integration of coupled ODEs
  - *D Gillespie, J Chem Phys, 22, 403-434 (1976)*
  - event = which reaction occurs next
  - **biochemical reaction pathways**
- TC (time counter) algorithm for DSMC
  - *G Bird, Phys Fluids, 6, 1518 (1963)*
  - event = collision between 2 particles
  - **low density fluid flow**

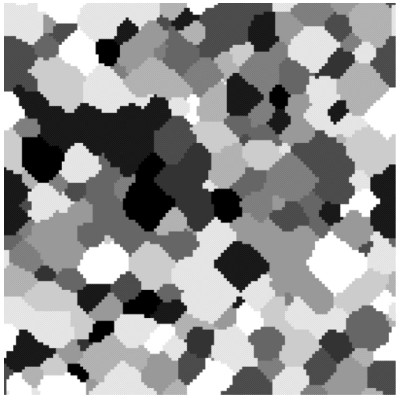
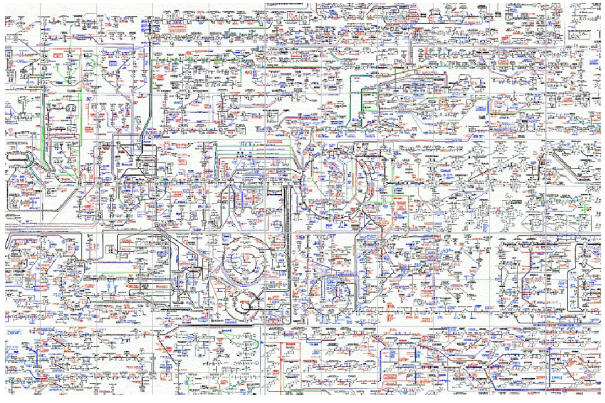
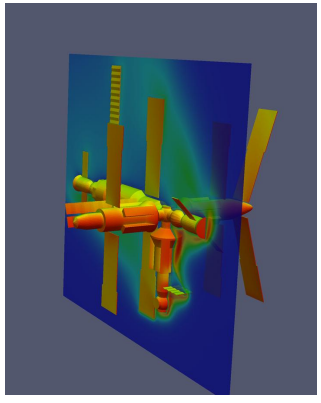

Exact linear and logarithmic kMC algorithms

- $O(N)$ **linear** algorithm is simple (Gillespie)
  - worked fine when N was small (early biochemical networks)
  - O(N) to select from list, O(1) to update
- $O(log(N))$ **logarithmic** algorithm
  - *Gibson & Bruck, J Phys Chem A, 104, 1876-89 (2000)*
  - method-of-choice as network sizes grew
  - binary tree: leaves = events, branch nodes = summed probs

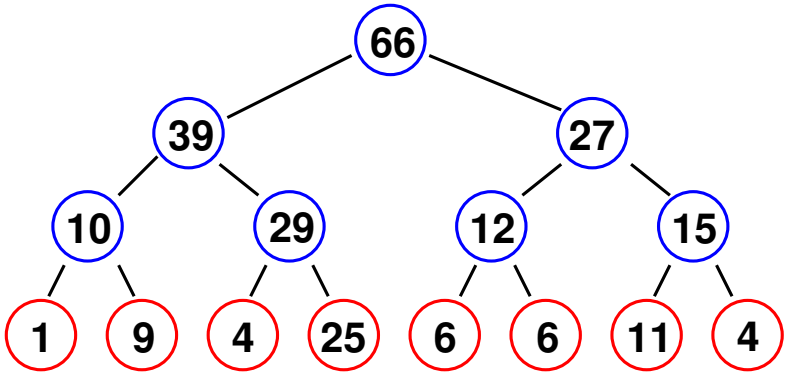

Exact linear and logarithmic kMC algorithms

- $O(N)$ **linear** algorithm is simple (Gillespie)
  - worked fine when N was small (early biochemical networks)
  - O(N) to select from list, O(1) to update
- $O(log(N))$ **logarithmic** algorithm
  - *Gibson & Bruck, J Phys Chem A, 104, 1876-89 (2000)*
  - method-of-choice as network sizes grew
  - binary tree: leaves = events, branch nodes = summed probs

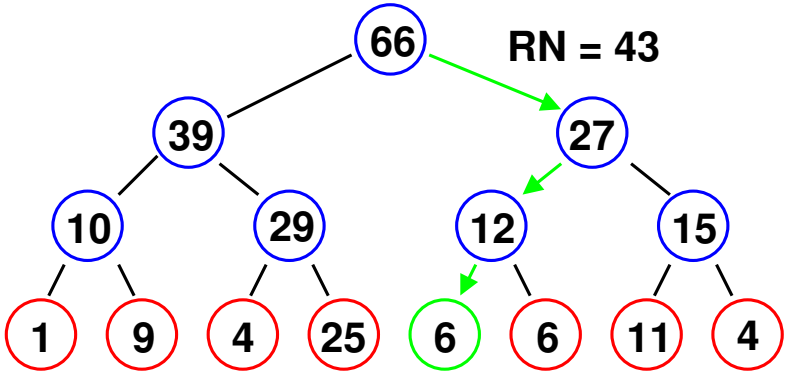

Exact linear and logarithmic kMC algorithms

- $O(N)$ **linear** algorithm is simple (Gillespie)
  - worked fine when N was small (early biochemical networks)
  - O(N) to select from list, O(1) to update
- $O(log(N))$ **logarithmic** algorithm
  - *Gibson & Bruck, J Phys Chem A, 104, 1876-89 (2000)*
  - method-of-choice as network sizes grew
  - binary tree: leaves = events, branch nodes = summed probs

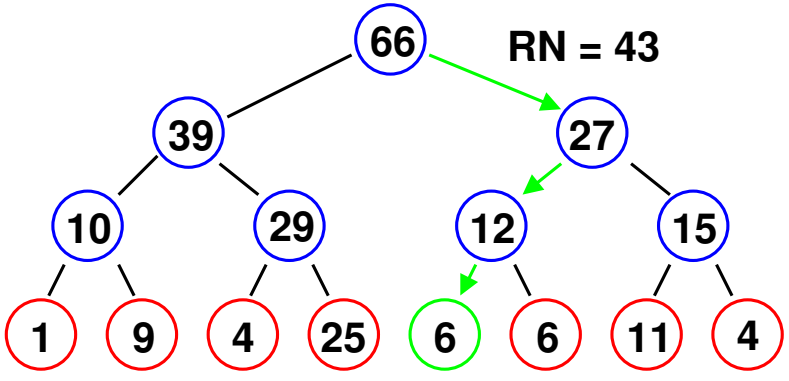

- Is it possible to have O(1) **constant-time** algorithm?

Exact constant-time kMC algorithm

- *Devroye, Non-uniform random variate generation (1986)*
- Based on **compostion/rejection** idea
- Applied to bio: *Slepoy et al, J Chem Phys 128, 205101 (2008)*

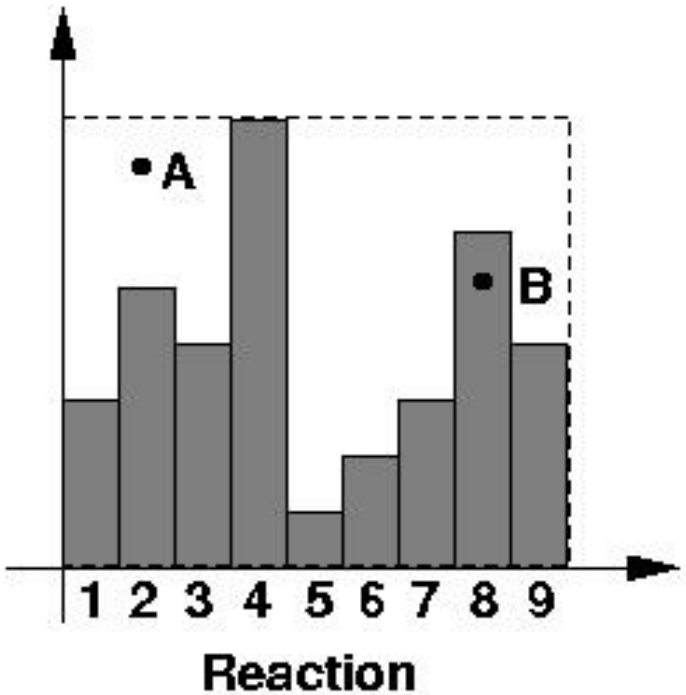

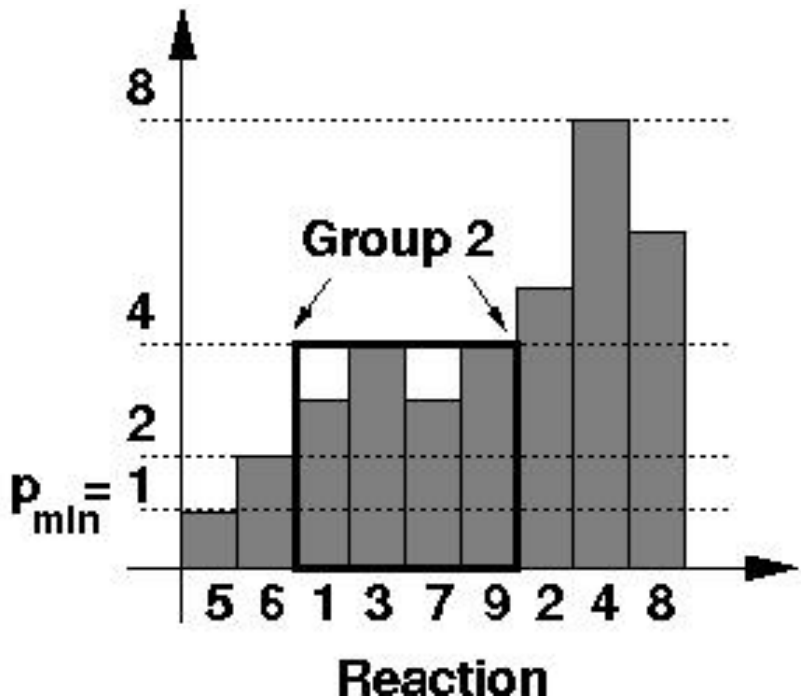

- Group bounds increase in **powers of 2**
- Rejection coverage $> 50\% \Rightarrow < 2$ RN on average

A few more details

- Constant-time requires $O(1)$ **operations** to:
  - find specific event within group (two-way pointers)
  - add/delete event to group (append to list, swap with end)
- **Condition #1**:
  - # of groups must be independent of N
  - bound ratio of max/min probabilities
- **Condition #2**:
  - # of affected-events per event does not grow with N

- Both typically satisfied for materials & bio network models

Performance of constant-time algorithm versus $\log(N)$ tree

- Network of cellular bio-chemistry reactions $A + B \xrightarrow{k} C + D$
- Biological cell = small well-mixed reaction volume

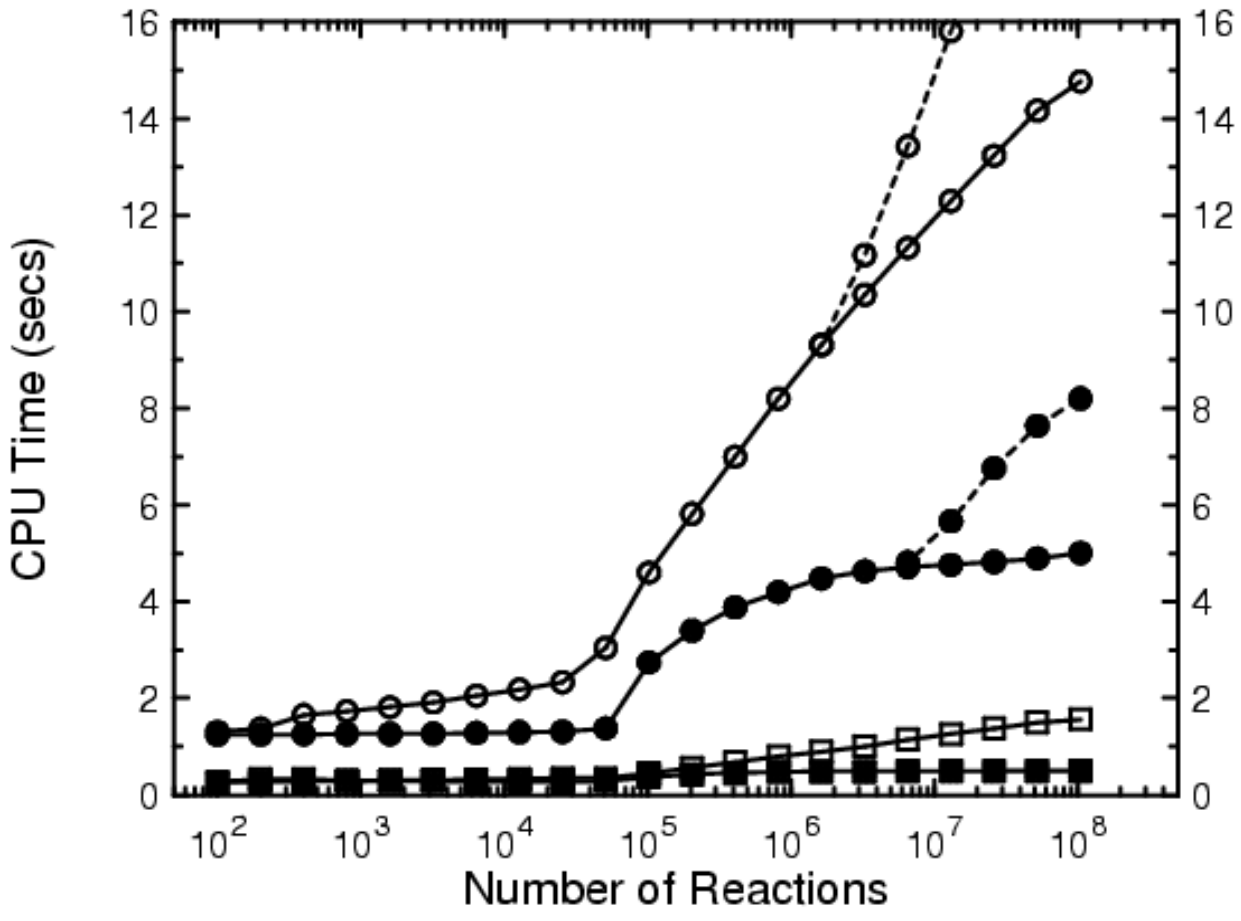

Parallel approximate kMC algorithm in SPPARKS

- Adapted from semi-rigorous synchronous sub-lattice idea
  Shim & Amar, *Phys Rev B 71, 125432 (2005)*

Parallel approximate kMC algorithm in SPPARKS

- Adapted from semi-rigorous synchronous sub-lattice idea
  Shim & Amar, *Phys Rev B 71, 125432 (2005)*
- Split domain into **sub-domains** and **sectors**: 4 in 2d, 8 in 3d
- Each processor works on same sector at same time
- kMC chooses site, sites can have multiple events
- Alternate kMC within sector, **comm of boundary sites**

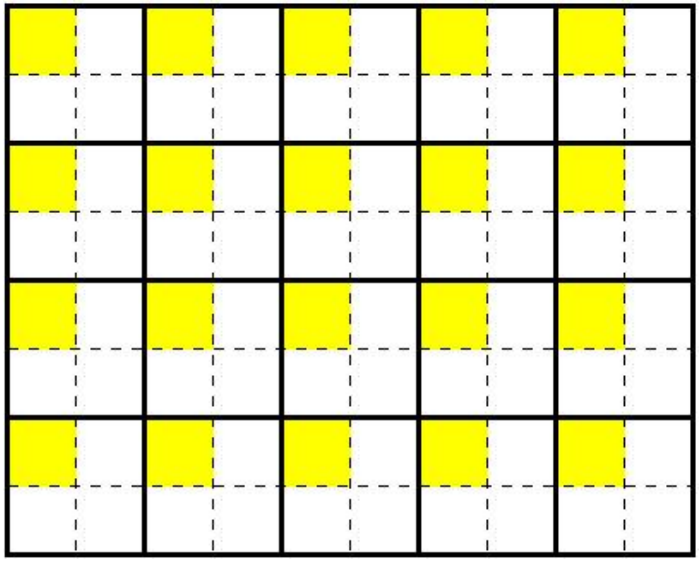

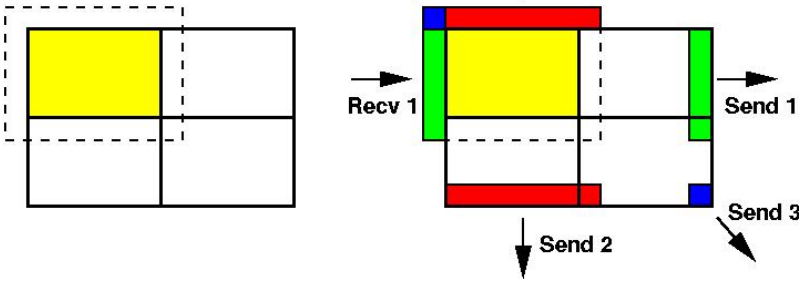

Sources of error

Approximate method for several reasons:

- Events occur simultaneously on different processors
- Order dependence induced by sector ordering
- Consecutive events in same sector are over-sampled
- Consecutive cross-sector events are under-sampled
- Boundary sites are older/younger than sector sites

Biggest issue: boundary sites are **frozen** during per-sector kMC

Can you find any bad grains?

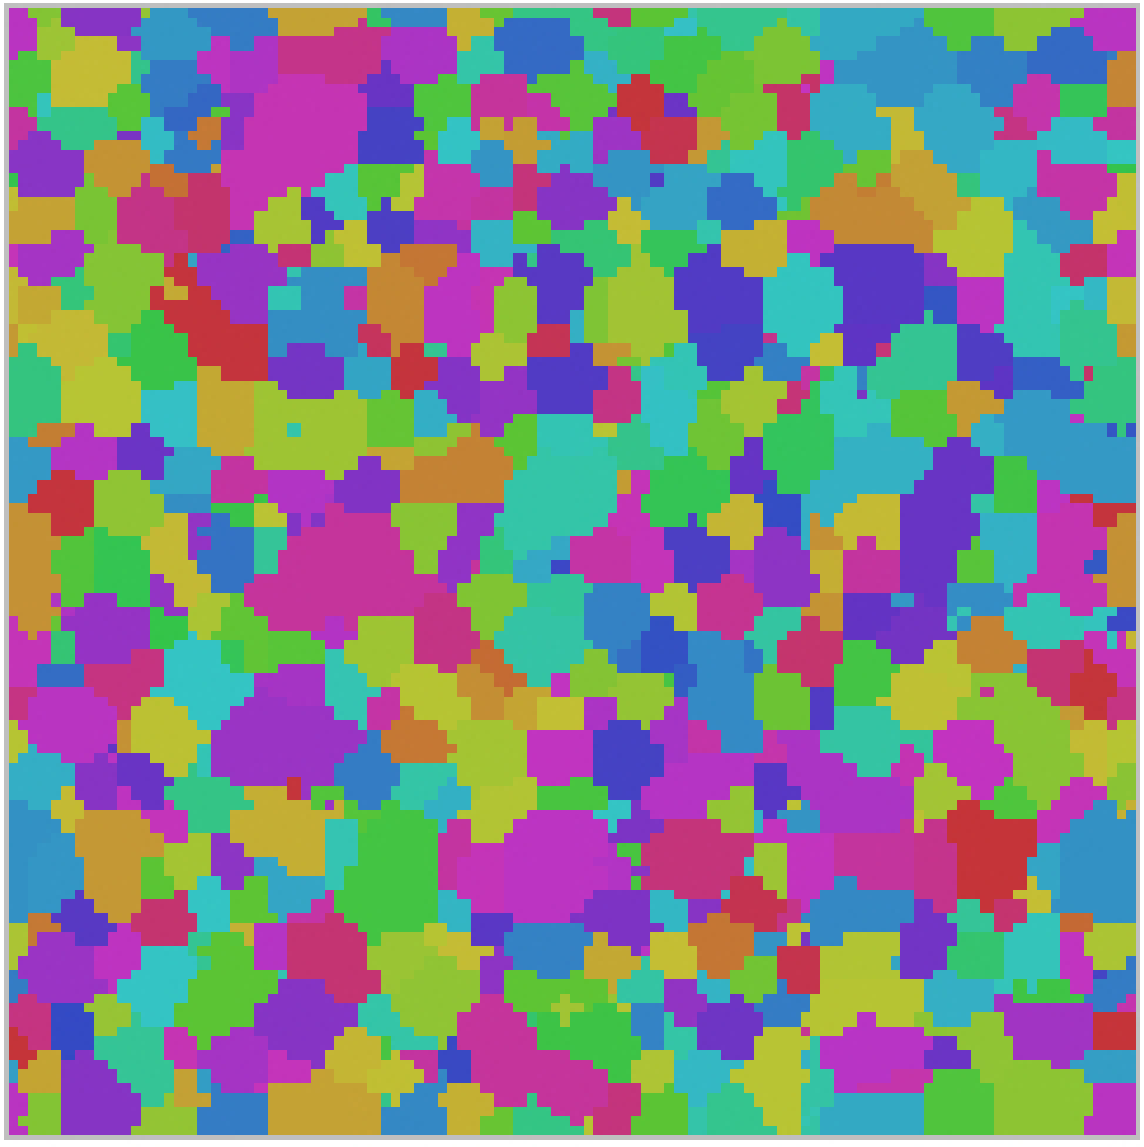

Shish-kebab effect

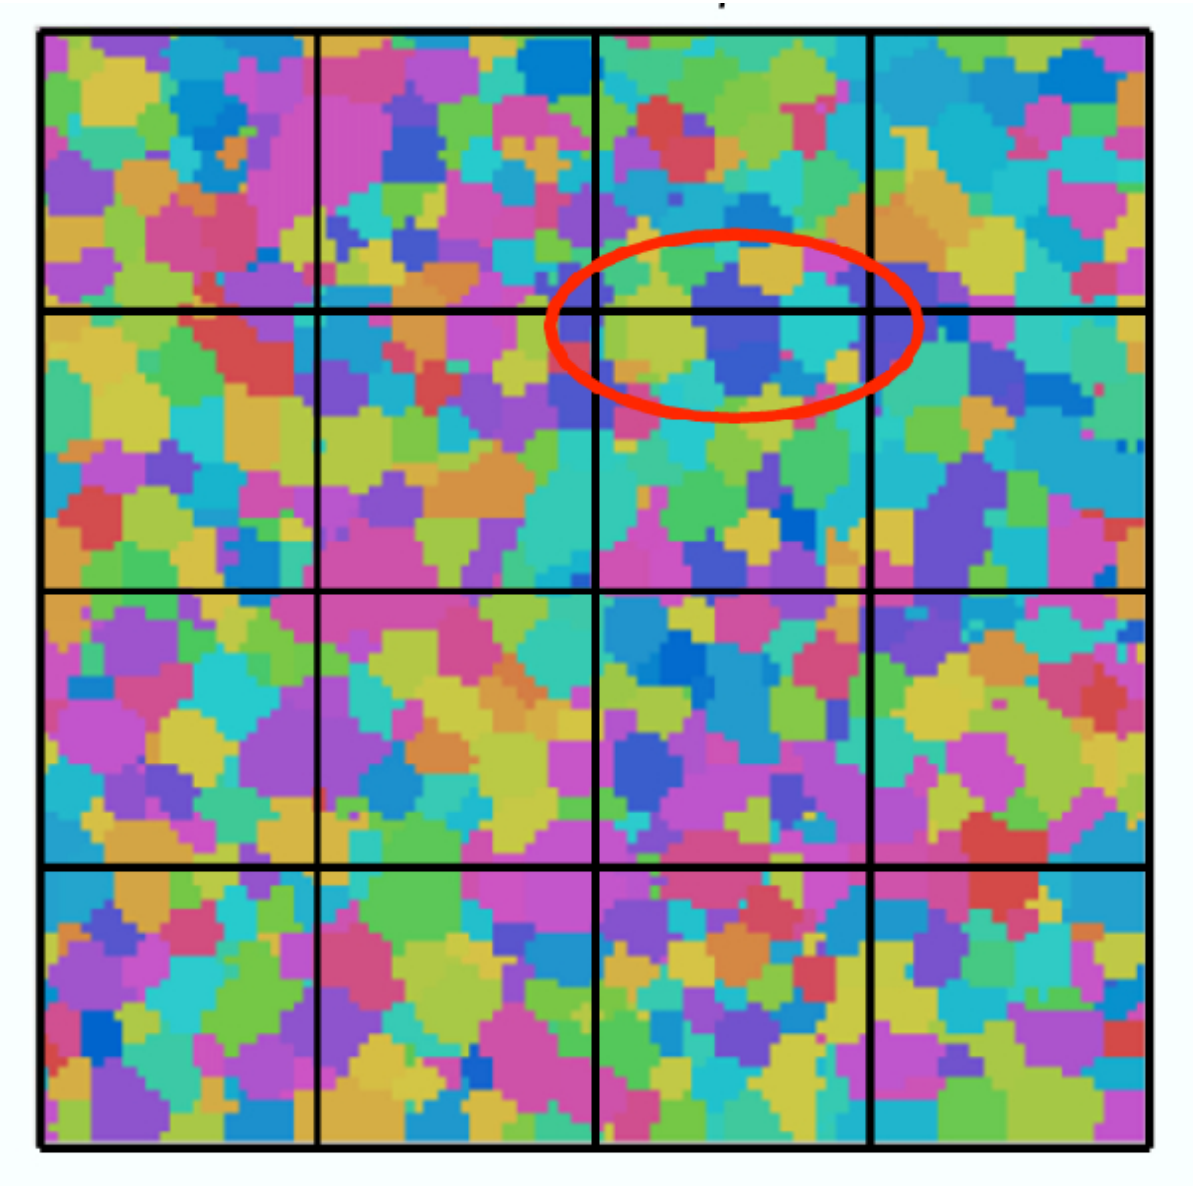

Solution: adaptive heuristic using events/site metric

Basic trade-off:

- Perform more events within a sector:
  - less communication $\Rightarrow$ good parallel scalability
  - more sector boundary effects
- Perform less events within a sector:
  - more communication $\Rightarrow$ poor parallel scalability
  - fewer sector boundary effects

Good compromise:

- User picks **nstop** = events per active site within sector
- Nstop = 1 is often a good rule-of-thumb
- **Adaptive** because active sites may vary greatly over time

Materials processing for space applications

- John Mitchell & Veena Tikare for NASA, DOE/ONE
- Radioisotope thermal generators (**RTG**) for Mars rovers

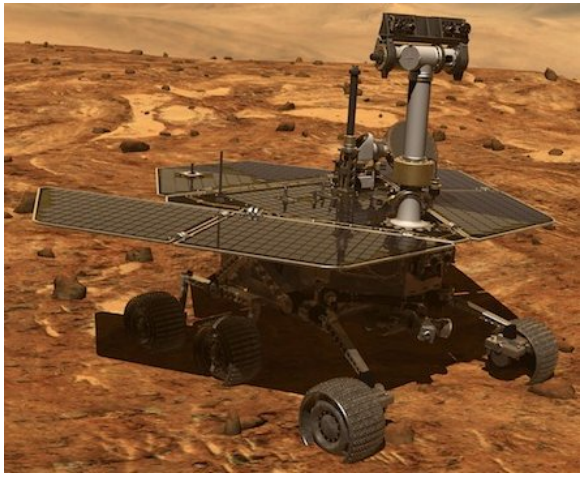

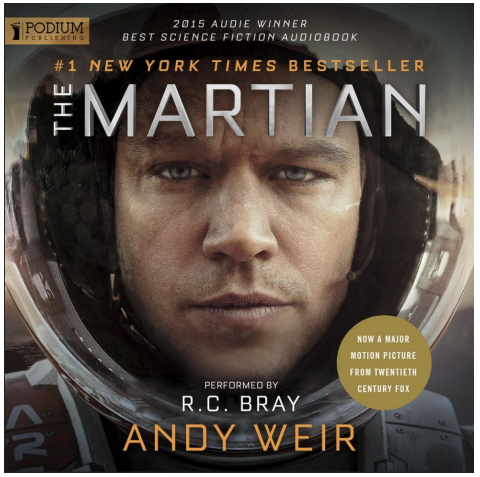

Materials processing for space applications

- John Mitchell & Veena Tikare for NASA, DOE/ONE
- Radioisotope thermal generators (**RTG**) for Mars rovers

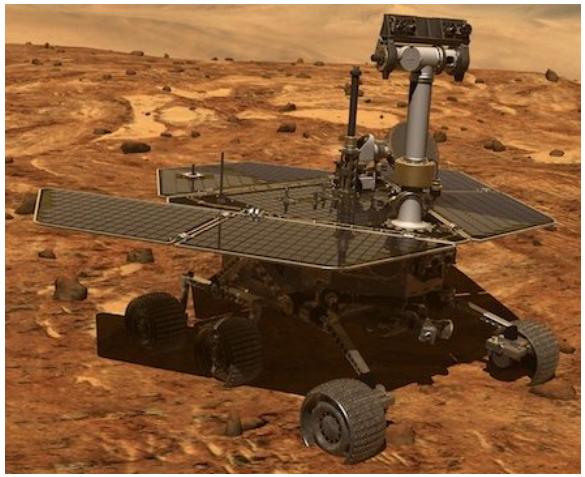

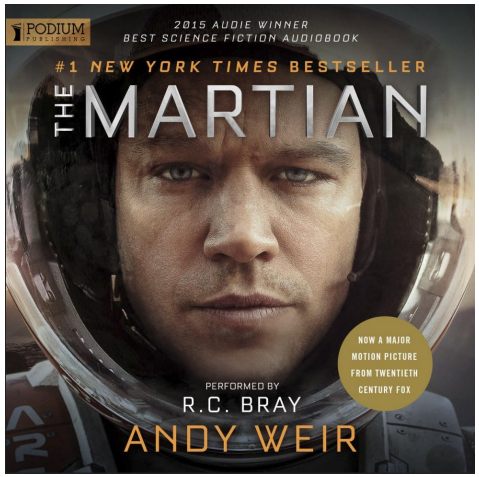

- Decay heat $\Rightarrow$ electrical power via thermoelectric effect
- Powers experiments, analysis (not spinning wheels)
- RTG **fuel pellet** = plutonia (PuO_2) in **Iridium container**
  - 3cm x 3cm cylinder, wrapped with 600 μm thick Ir
  - Ir = refractory metal, melting pt near 2500° C
- Tungsten gas arc used to **weld container shut**

kMC model of welding process

- How a weld joins two pieces of metal?
  - metals are **polycrystalline**
  - arc spot melts the two pieces
  - recrystallizes to grain structure that spans the gap
  - good = small, equi-axial grains
- Arc melts a spot with 3d profile
- Nearby **heat-affected zone** where T drops off
- **Potts model** for grain growth
  - spin = crystal lattice orientation
  - grain = cluster of same spins
  - curvature-driven grain growth
  - add T(x,t)-dependent mobility term

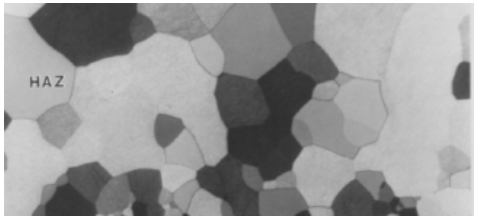

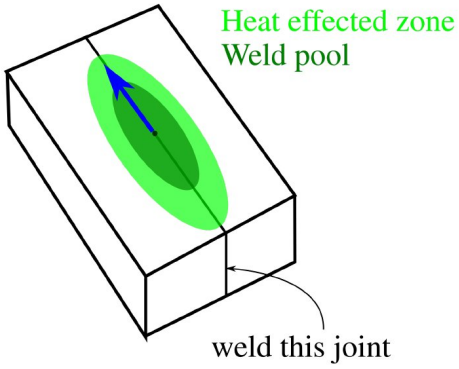

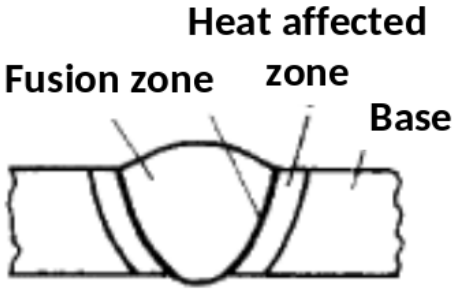

Model details

- kMC event = **spin flip** to neighbor site value
- Hamiltonian and event probability with **mobility prefactor**

$$H_i = \sum_{j=1}^{N} J(1 - \delta_{ij}), \quad \delta_{ij} = 1 \text{ if } s_i = s_j, \text{ else } 0$$

$$P = \exp\left(\frac{-Q}{T(x,t)}\right) \exp\left(\frac{-\Delta E}{k_B T}\right)$$

Model details

- kMC event = **spin flip** to neighbor site value
- Hamiltonian and event probability with **mobility prefactor**

$$H_i = \sum_{j=1}^{N} J(1 - \delta_{ij}), \quad \delta_{ij} = 1 \text{ if } s_i = s_j, \text{ else } 0$$

$$P = \exp\left(\frac{-Q}{T(x,t)}\right) \exp\left(\frac{-\Delta E}{k_B T}\right)$$

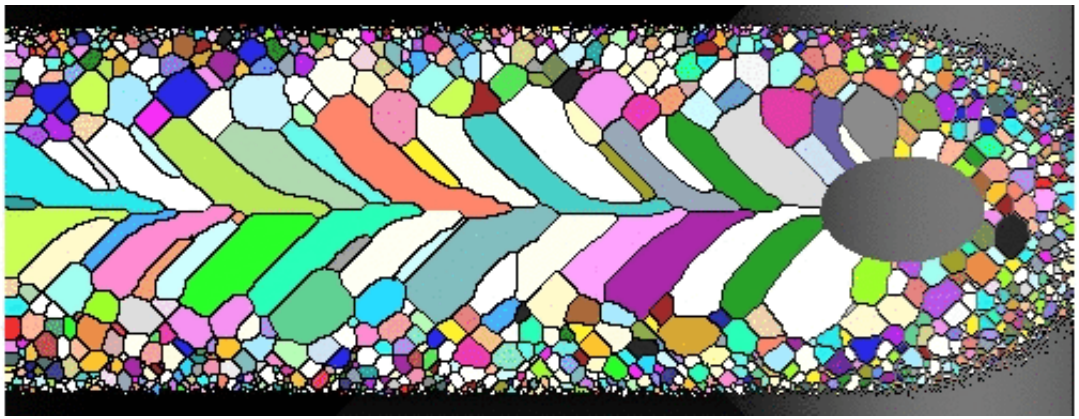

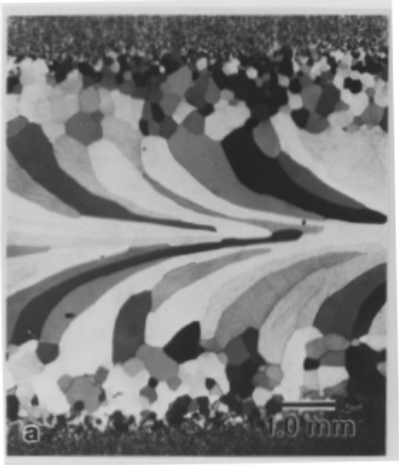

- **Banana shaped** grains!

Movie of arc spot motion

2d view of 3d lattice, tens of millions of lattice sites

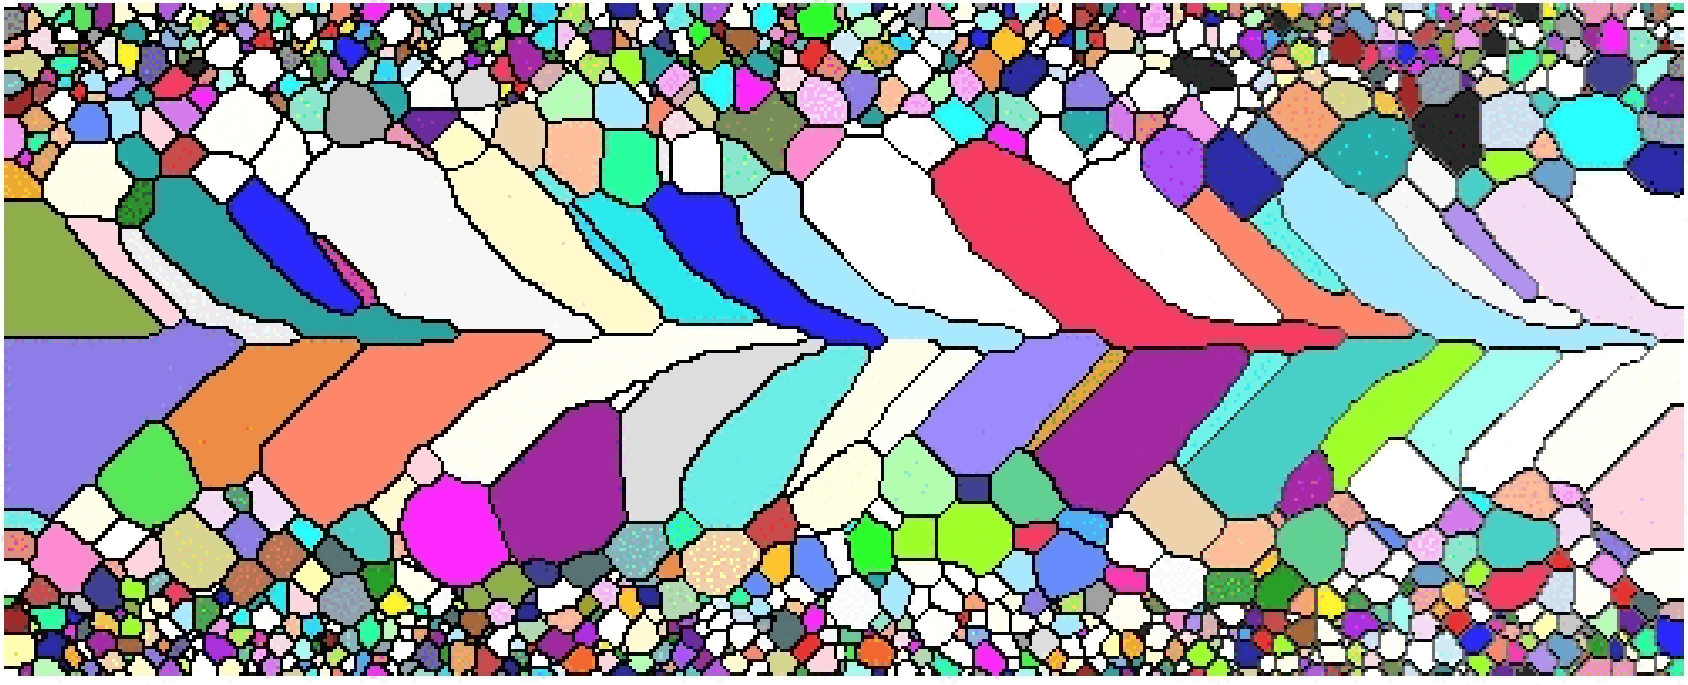

Time-varying arc pulse $\Rightarrow$ more equi-axial grains

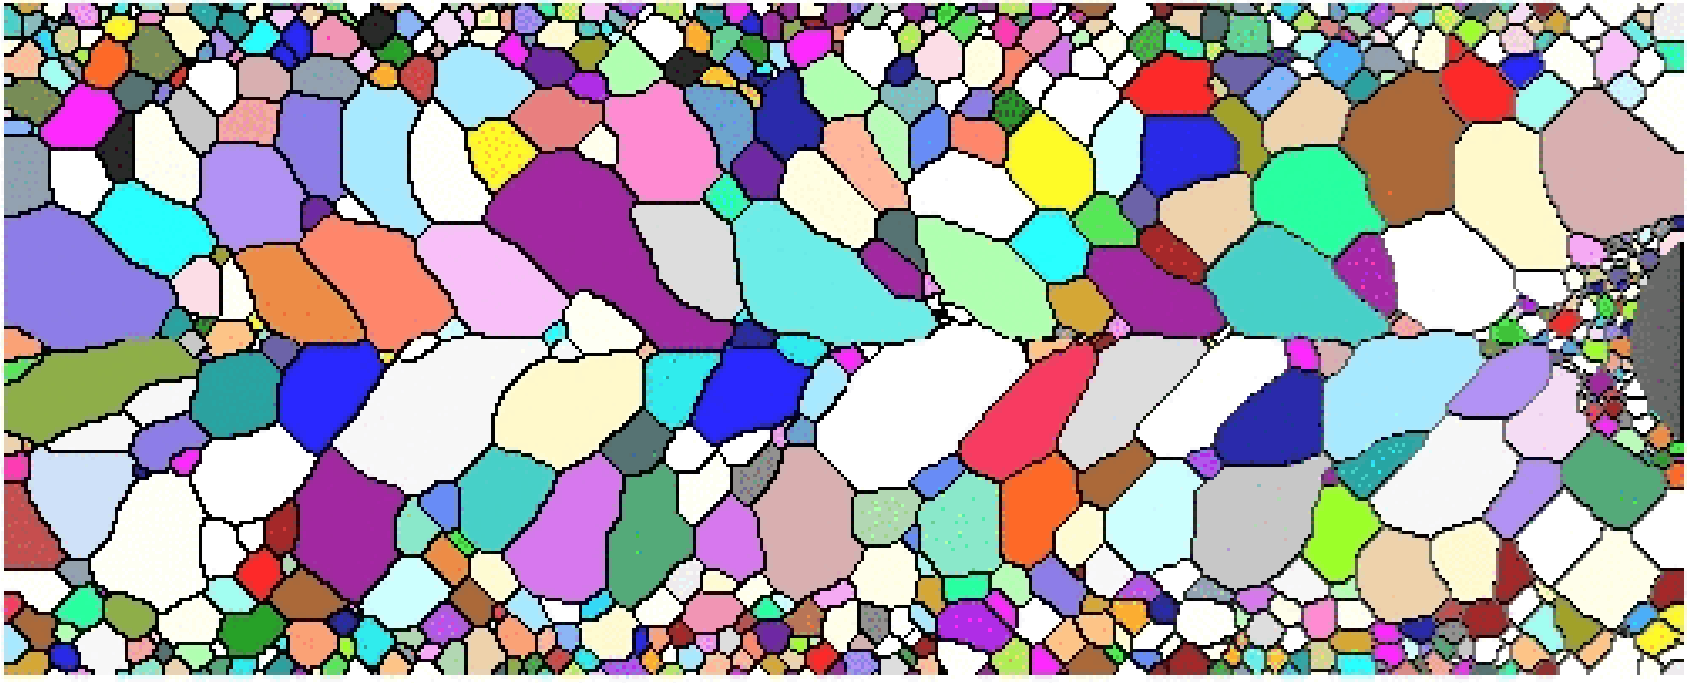

Time-varying arc pulse $\Rightarrow$ more equi-axial grains

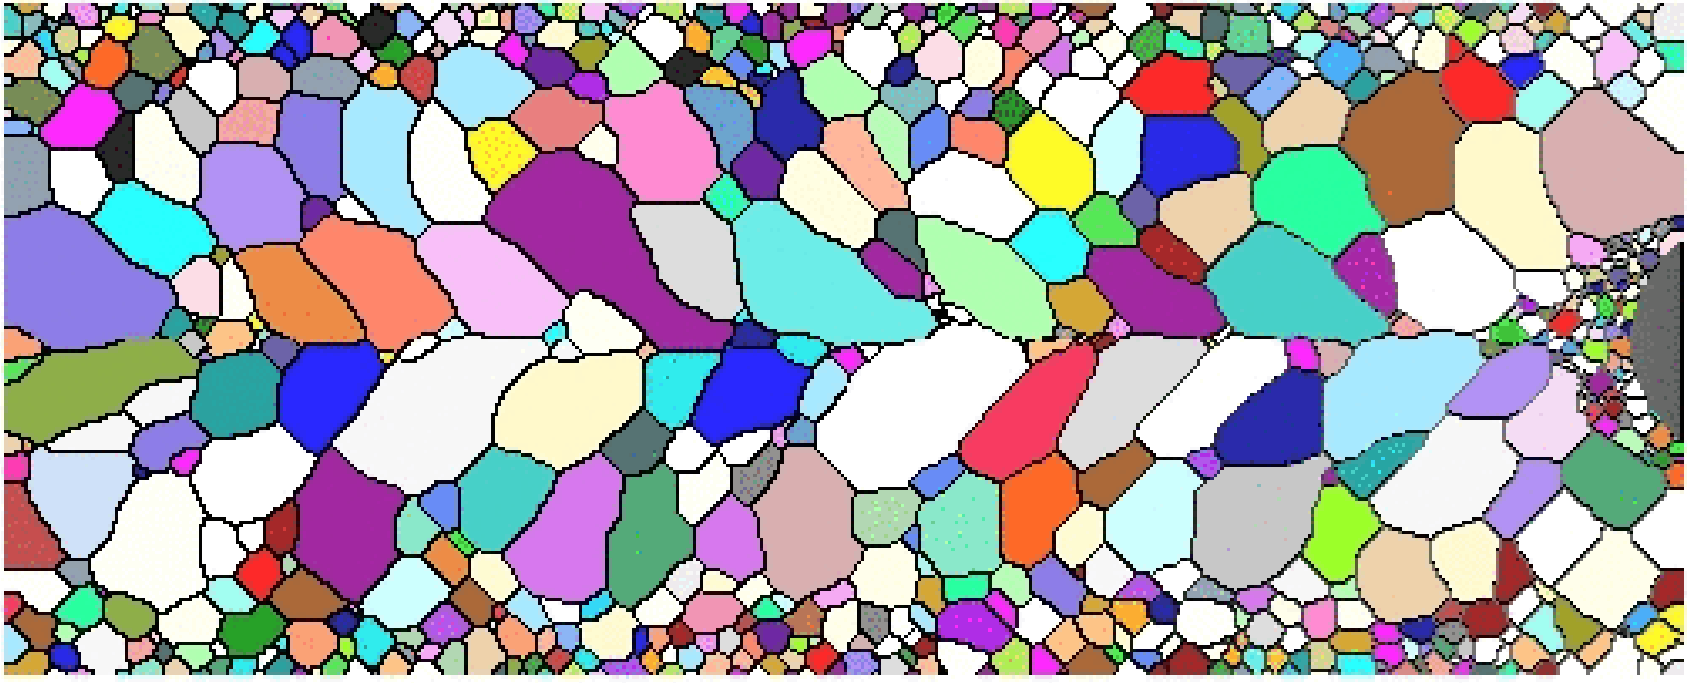

- With $\sim 10^{10}$ sites $\Rightarrow$ experimental length/time scales
  - 10 μm grain = 10x10x10 lattice sites
  - arc spot moving at $\sim$1 cm/sec
  - dependence on speed, spot size/shape/time-variability

Research challenges for kMC modeling

1. **Rigorous error bounds** on parallel approximate algorithms
   - for any Hamiltonian and model
   - Markos and collaborators work on this

Research challenges for kMC modeling

1. **Rigorous error bounds** on parallel approximate algorithms
   - for any Hamiltonian and model
   - Markos and collaborators work on this
2. New algorithms for **off-lattice kMC**:
   - most applicable to atomic-scale
     - huge timescale advantage vs MD
     - μsec events vs fmsec timestepping
   - events = diffusive hops, rearrangements of surfaces/interfaces
   - identifying events is more challenging
     - irregular neighbors, no preferred direction, multi-atom moves
     - barrier heights via nudged-elastic band (NEB)
     - accelerated MD methods can help

Research challenges for kMC modeling

1. **Rigorous error bounds** on parallel approximate algorithms
   - for any Hamiltonian and model
   - Markos and collaborators work on this
2. New algorithms for **off-lattice kMC**:
   - most applicable to atomic-scale
     - huge timescale advantage vs MD
     - μsec events vs fmsec timestepping
   - events = diffusive hops, rearrangements of surfaces/interfaces
   - identifying events is more challenging
     - irregular neighbors, no preferred direction, multi-atom moves
     - barrier heights via nudged-elastic band (NEB)
     - accelerated MD methods can help
3. **Coupled kMC/MD** algorithms (akin to MD/QM):
   - use MD to find kMC events
   - initially run mostly at tiny MD timescale (expensive)
   - until catalog all relevant events
   - then run at large kMC timescale (cheap)

Thanks and links

- Faster computers enabling advances in materials modeling
- Leads to **new challenges** for algorithms, models, mathematics

Thanks and links

- Faster computers enabling advances in materials modeling
- Leads to **new challenges** for algorithms, models, mathematics

- **Joint work** with many folks at Sandia:
  - LAMMPS: Aidan Thompson, Stan Moore,
    Mitch Wood, Axel Kohlmeyer (Temple U)
  - SPPARKS: Aidan Thompson, John Mitchell, Veena Tikare

- Open-source codes:
  - LAMMPS for MD: **http://lammps.sandia.gov**
  - SPPARKS for kMC: **http://spparks.sandia.gov**