Physics-based modeling of materials and uncertainty quantification as drivers for exascale computing

Jaime Marian

Department of Materials Science and Engineering
Department of Mechanical and Aerospace Engineering
IDRE Executive Committee Member
University of California Los Angeles

UCLA

Department of Materials Science and Engineering

IPAM, February 3, 2017

UCLA Communications & Public Outreach • 1147 Murphy Hall, Box 951436 • Los Angeles, CA 90095-1436

‘Discovery class’ DOE facilities approaching exascale capabilities


LLNL (500 teraFLOPs)

ORNL (10 petaFLOPs)

LLNL (120 petaFLOPs)

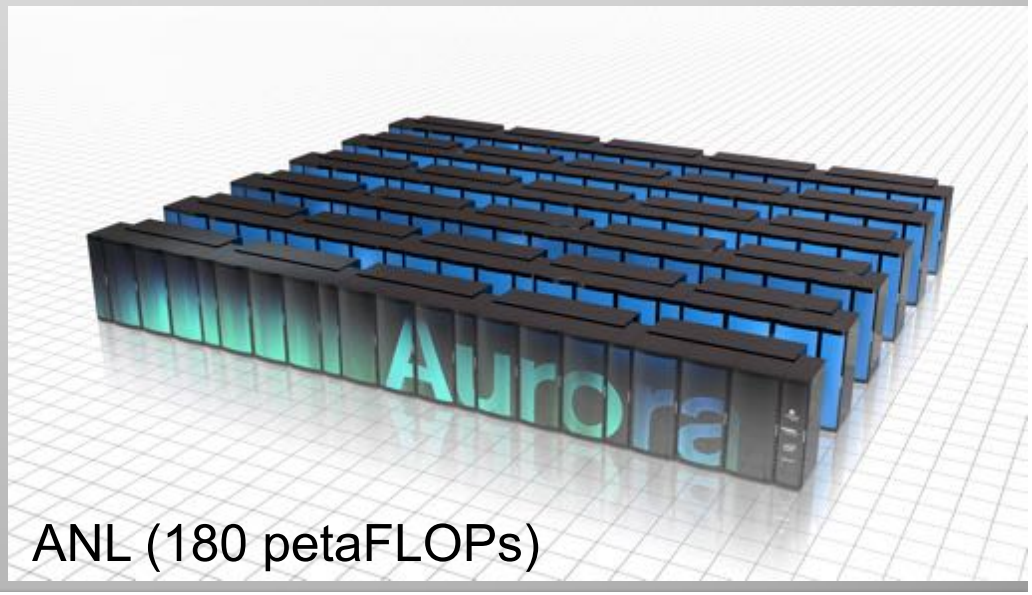

ANL (180 petaFLOPs)

The predictive science paradigm

- **Aim:** *Predict the behavior of complex physical/engineered systems* ***with quantified uncertainties***
- **Paradigm shift** in experimental science, modeling and simulation, scientific computing (***predictive science***):
  - Deterministic → Non-deterministic systems
  - Mean performance → Mean performance + Uncertainty

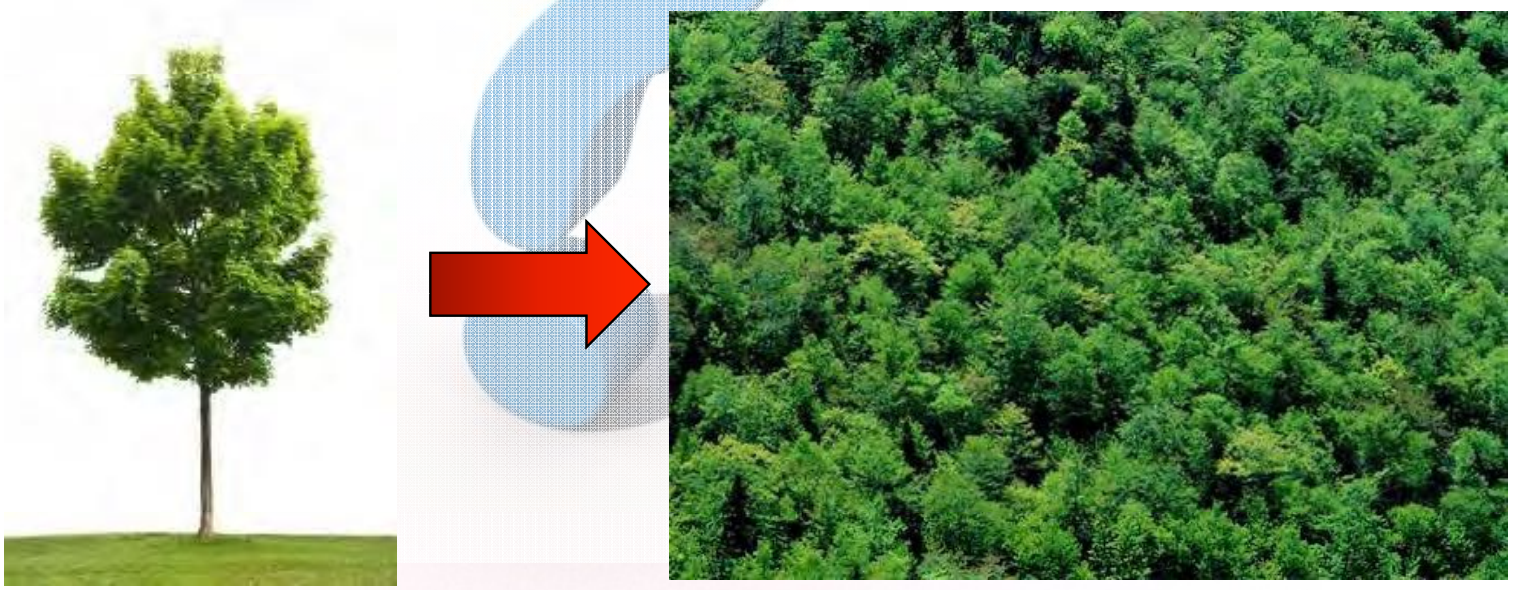

Old (single-calculation) paradigm

New(ensemble of calculations) paradigm

The predictive science paradigm

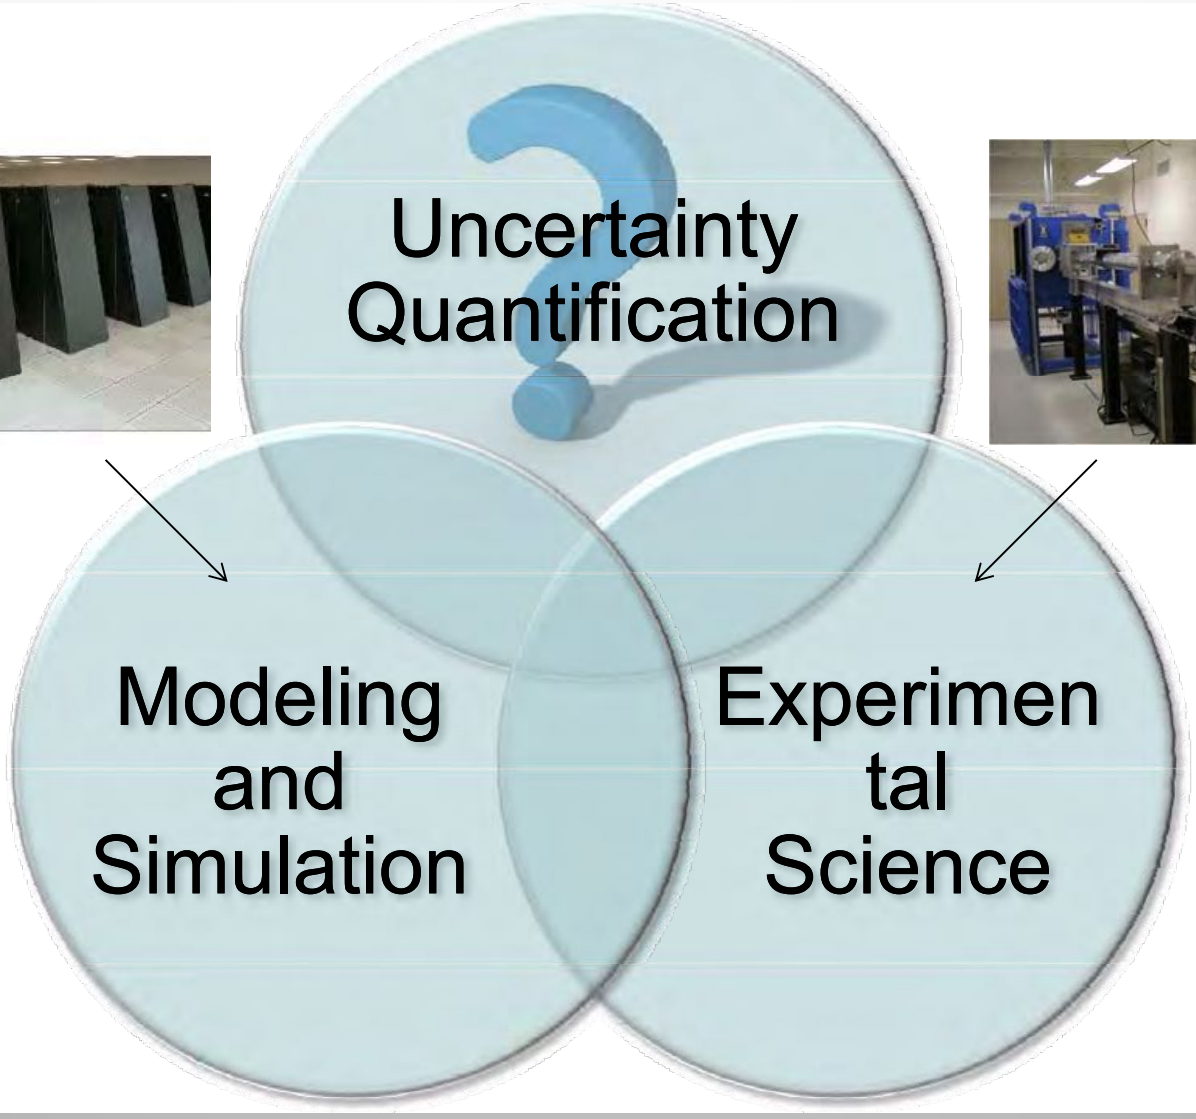

Materials science is a great example of massively-parallel computation thrusts

- Siesta, VASP, Quantum Espresso, QBox: **first-principles atomistic simulations**
- Moldy, MDCASK, LAMMPS: **domain decomposition molecular dynamics**
- ParaDis, Numodis, MODEL: **continuum dislocation dynamics**
- Diablo, Abaqus: **finite elements, continuum mechanics**
- Ale3D: **hydrodynamics**
- etc…

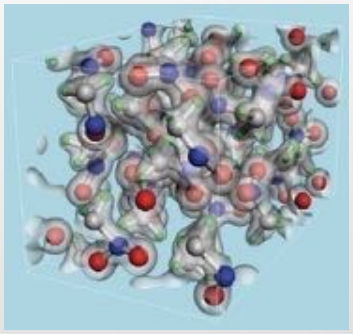

First principles

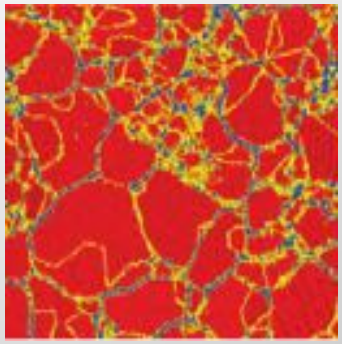

Classical MD

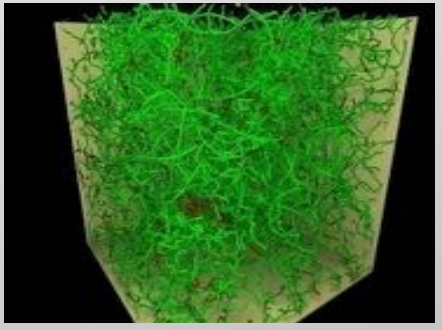

Dislocation dynamics

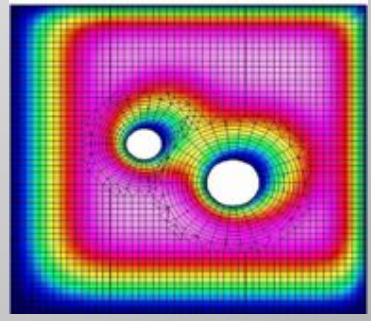

Finite elements

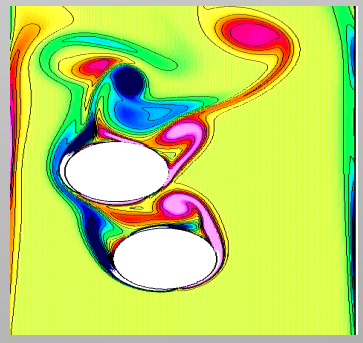

Hydrodynamics

Some materials science drivers for petascale computing

Application area	Science driver	Science objective	Impact
Nanoscale science	Material-specific understanding of high-temperature superconductivity theory	Understand the quantitative differences in the transition temperatures of high-temperature superconductors	Macroscopic quantum effect at elevated temperatures (>150K) New materials for power transmission and oxide electronics
	Thermodynamics of nanostructures	Understand and improve colossally magneto-resistive oxides and magnetic semiconductors Develop new switching mechanism in magnetic nanoparticles for ultrahigh-density storage Simulate and design molecular-scale electronics devices	Magnetic data storage Economically viable ethanol production Energy storage via structural transitions in nanoparticles
	Evolution of an understanding of biological system behavior	Elucidate the physical-chemical factors and mechanisms that control damage to DNA	Medicine, biomimetics, sequence dependencies, and inhibiting agents of hazardous bioprocesses
Material response	Elucidation of the causes leading to eventual brittle or ductile fragmentation and failure of a solid	Understand macro-cracking due to coalescence of subscale cracks, local deformation due to void coalescence, and dynamic propagation of cracks or shear bands	Reduction of engineering margins to within required safe operating envelop

Metal plasticity – Multiscale analysis

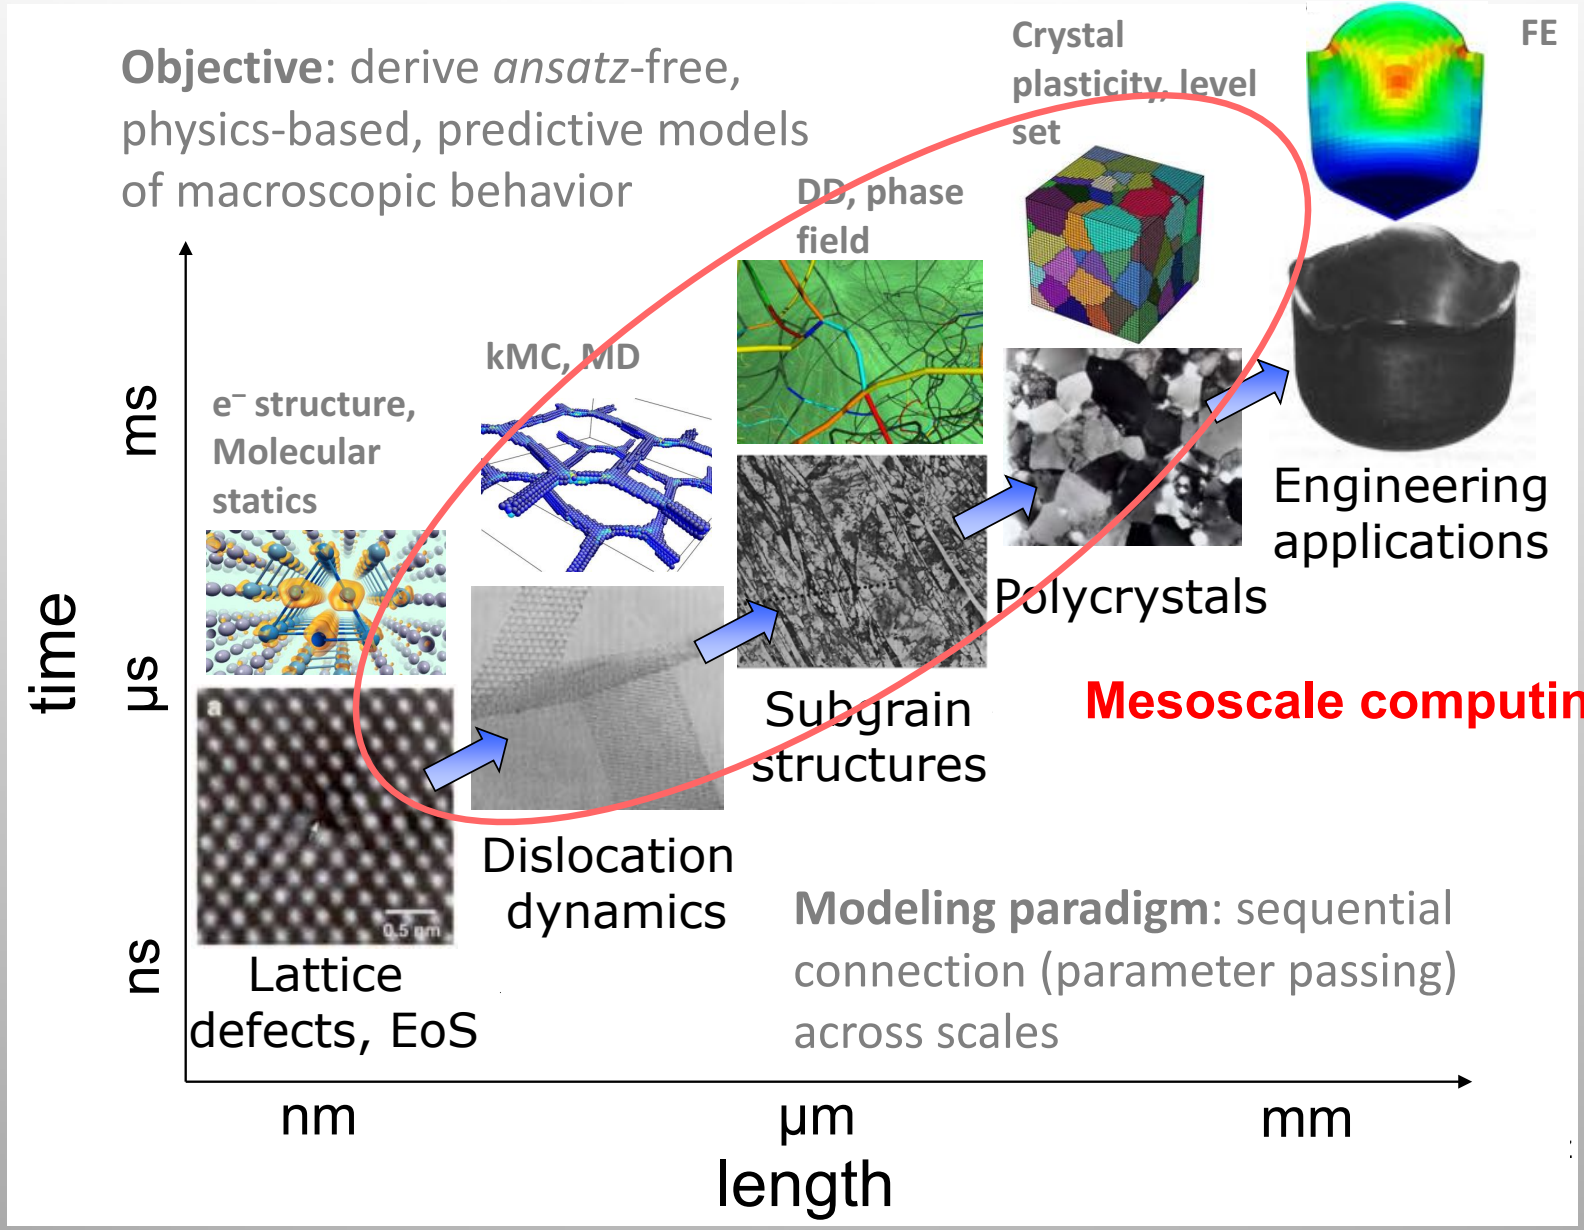

Atomistic simulations are pushing the envelope of available computational resources

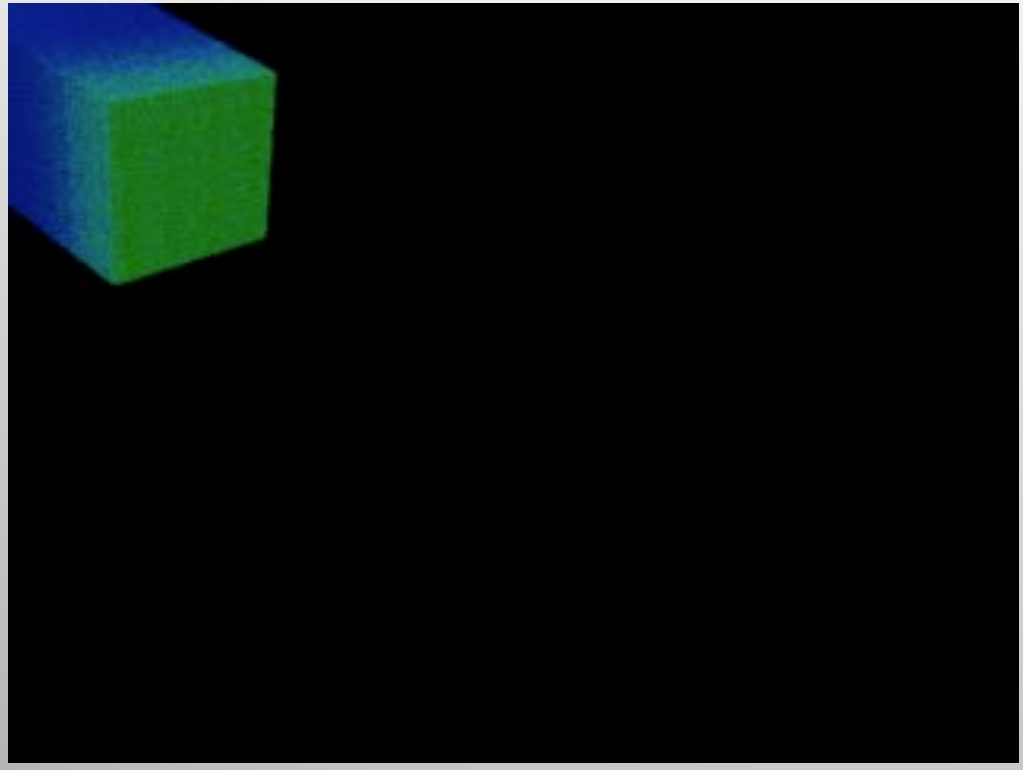

Laser ablation of Au column
(Gilmer et al)

We routinely simulate systems with several billion atoms using $10^{5\sim6}$ processors

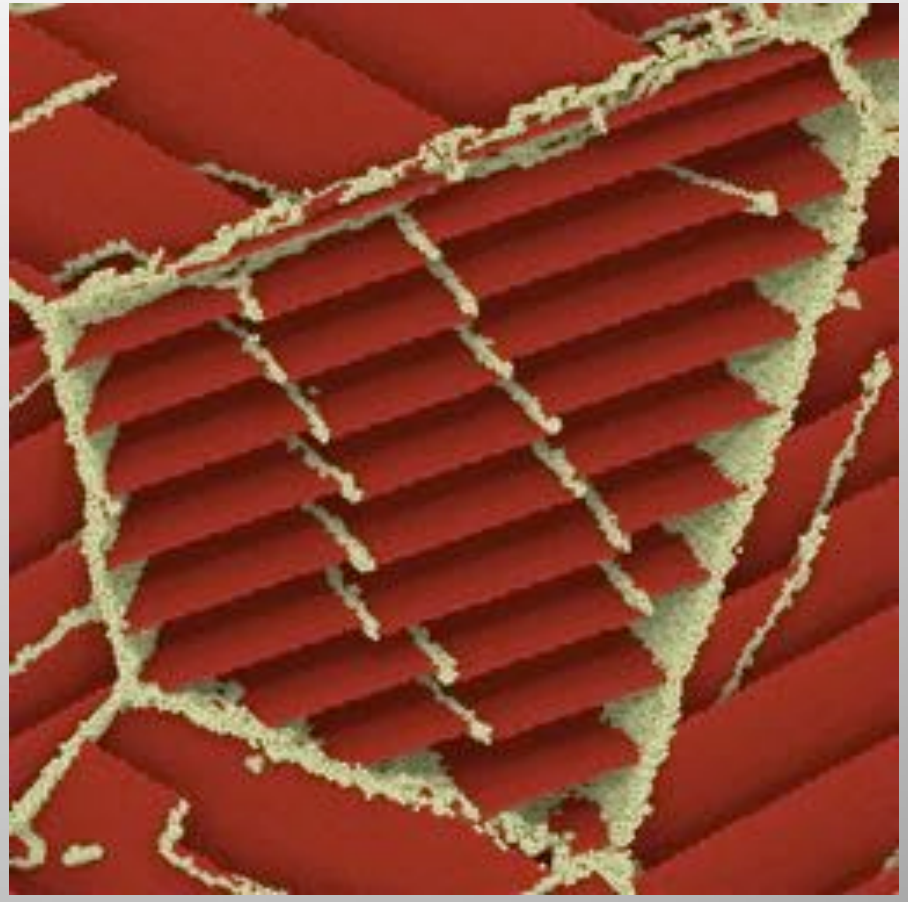

Deformation of nano-twinned metals (Sansoz et al)

Dislocation ensembles control high-rate deformation in metallics alloys

Ian Robertson (UIUC)

Atomistic simulation of dislocation interactions and density evolution

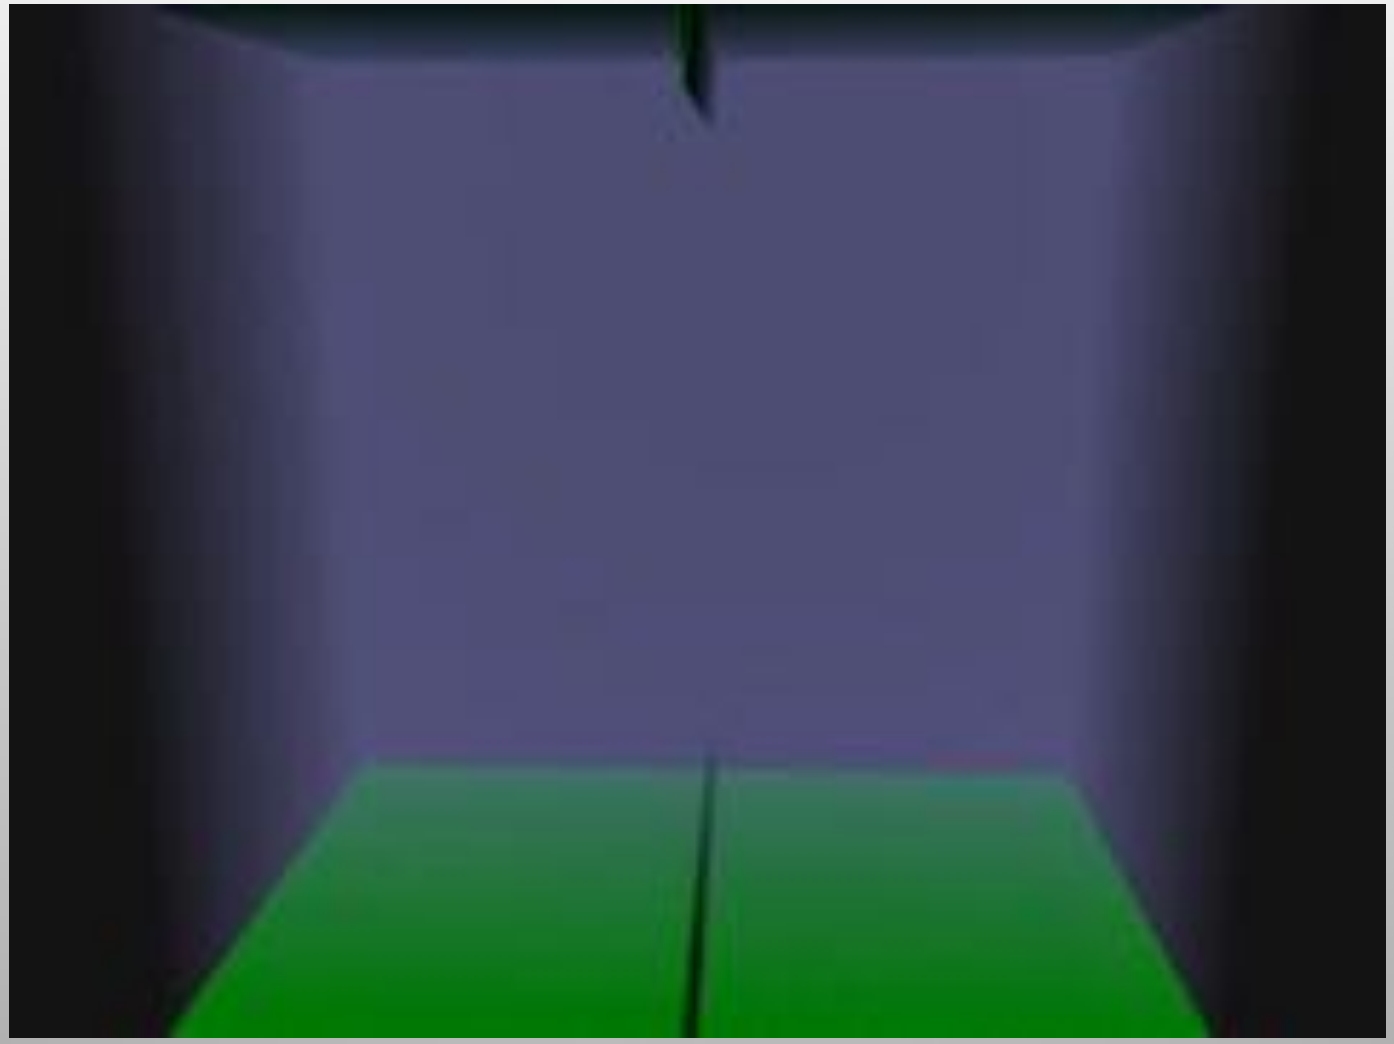

10^9 degrees of freedom from atomistic calculations → too much information to store

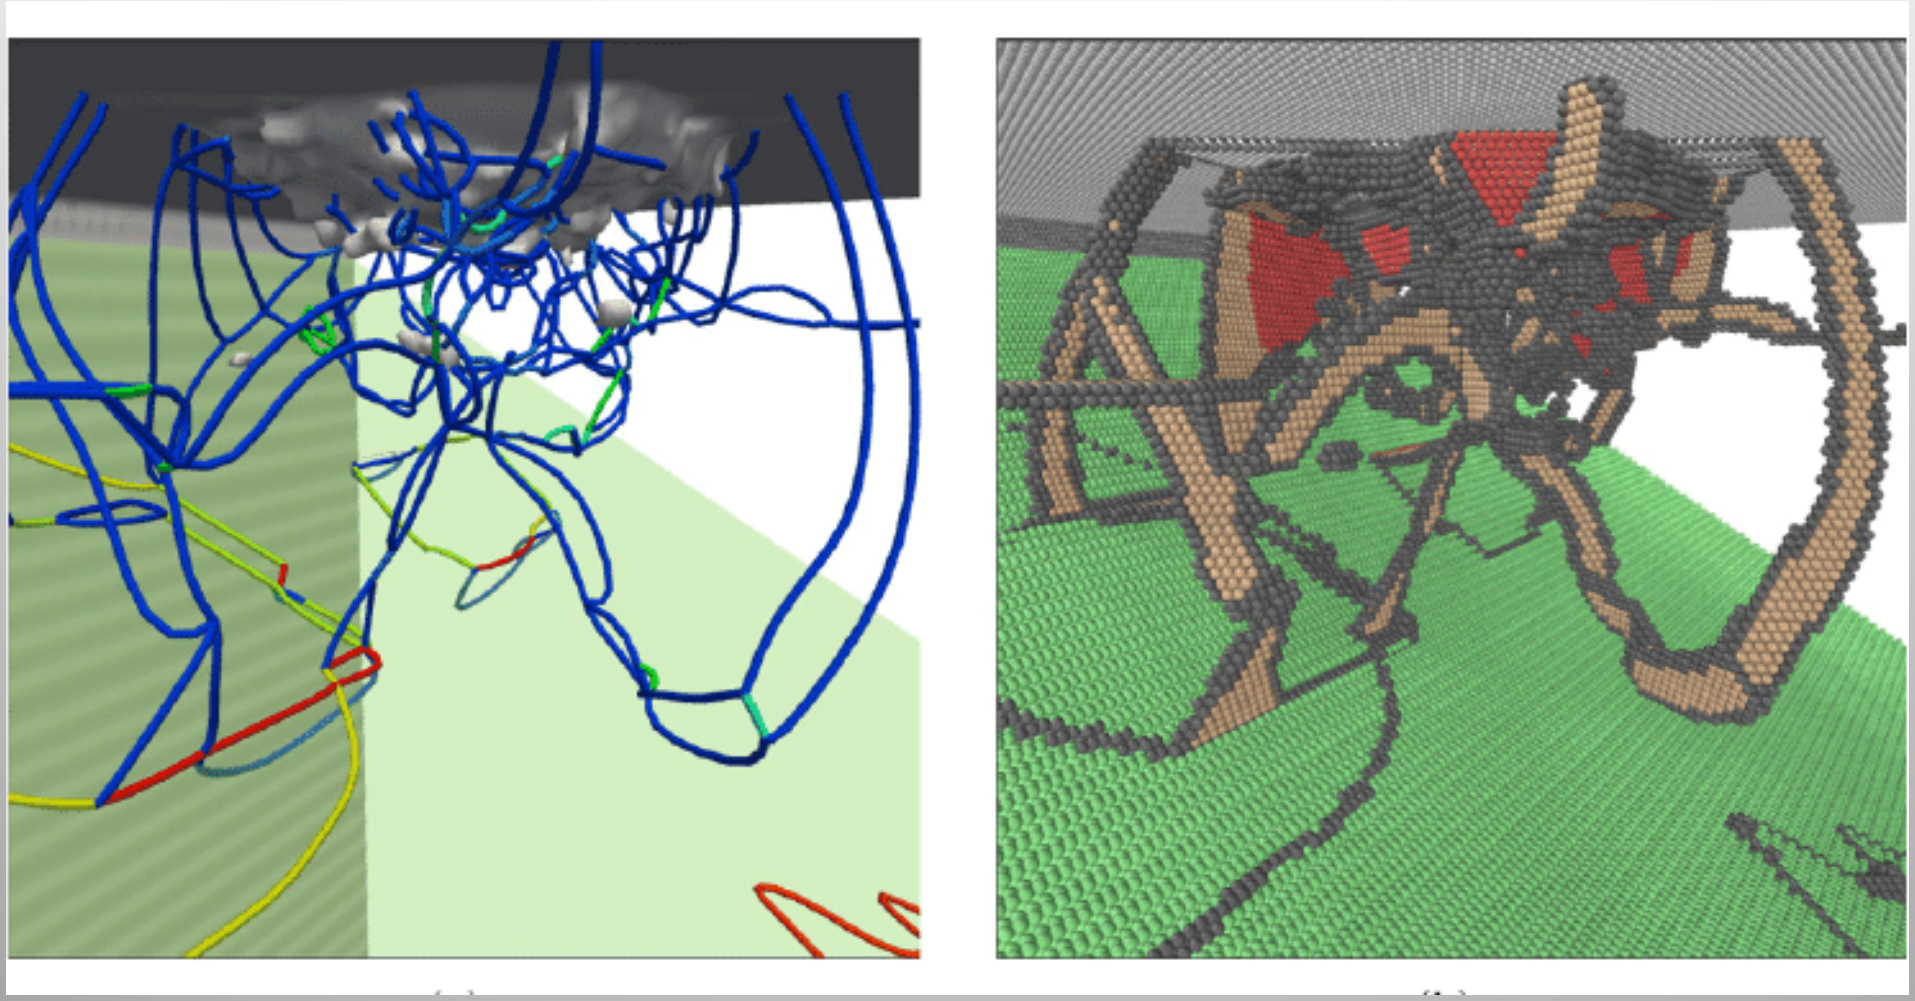

Use appropriate filters that help with visualization while preserving key physical features (DXA algorithm, Stukowski et al)

The Quasicontinuum Method (QC)

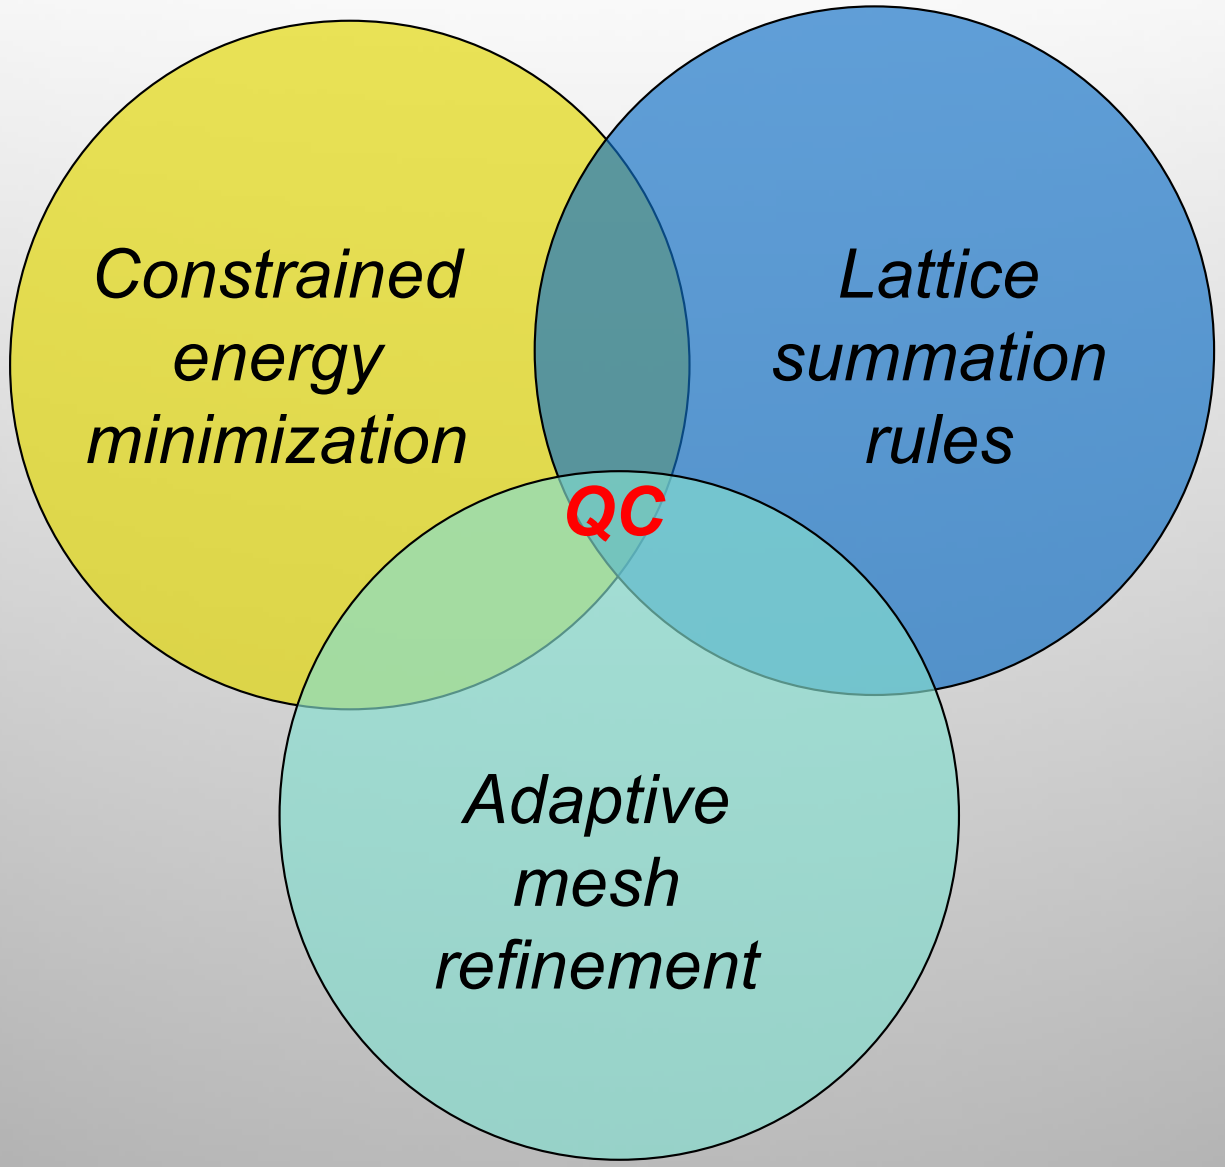

The Quasicontinuum Method (QC)

Lattice statics – problem definition

- The reference configuration: atoms arranged on a subset of a ***simple Bravais lattice.***
- $\mathcal{L} \subset Z^3 \equiv$ enumeration of atoms in the ensemble.
- Atomic positions in the reference configuration:

$$X(l) = l^i \boldsymbol{a}_i, \quad l \in \mathcal{L}$$

- Positions of atoms ***after deformation***

$$\boldsymbol{q} \equiv \{\boldsymbol{q}(l), l \in \mathcal{L}\} \in R^{3N} \equiv X$$

- $X \equiv$ Configuration space of the ensemble.

- Wish to determine the ***equilibrium configuration*** of the ensemble.
- Total energy: $E(\boldsymbol{q})$, $\boldsymbol{q} \in X$
- Problem:

$$\inf_{\boldsymbol{q} \in X} E(\boldsymbol{q})$$

- Difficulties:
  i. N ***very large*** $\sim 10^{23}$
  ii. $E(\boldsymbol{q})$ highly **non-convex** $\longrightarrow$ lattice defects, defect structures

The Quasicontinuum Method (QC)

QC-reduction

- Representative atoms: $\mathcal{L}_h \subset \mathcal{L}$

$$N_h = \text{card}(\mathcal{L}) \ll N$$

- Introduce triangulation $\mathcal{T}_h$ of $\mathcal{L}_h$
- Basis functions: for $l_h \in \mathcal{L}_h$

i. $\varphi_h(X|l_h)$ continuous

ii. Linear over simplices $K \in \mathcal{T}_h$

iii. $\varphi_h(X(l_h)|l'_h) = \begin{cases} 1, & \text{if } l_h = l'_h \\ 0, & \text{otherwise} \end{cases}$

$$q_h(l) = \sum_{l_h \in \mathcal{L}_h} \varphi_h(l|l_h)\, q_h(l_h)$$

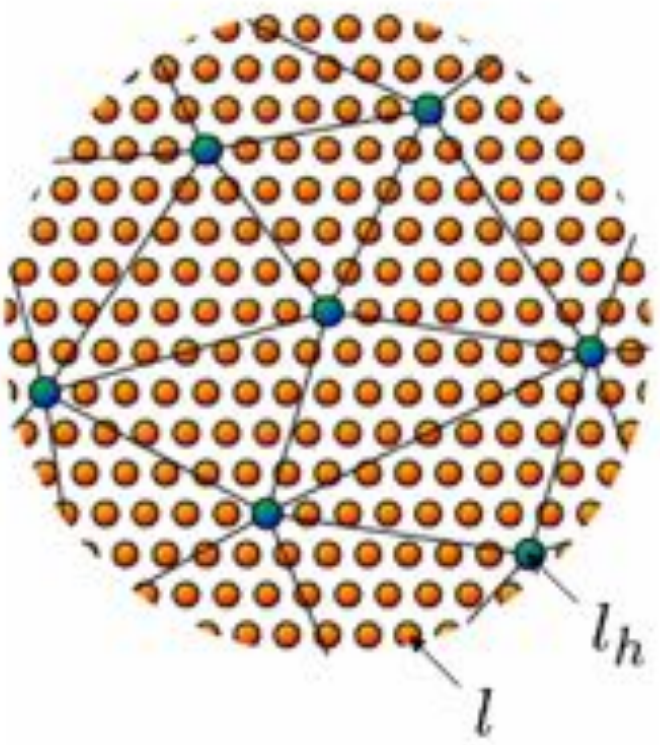

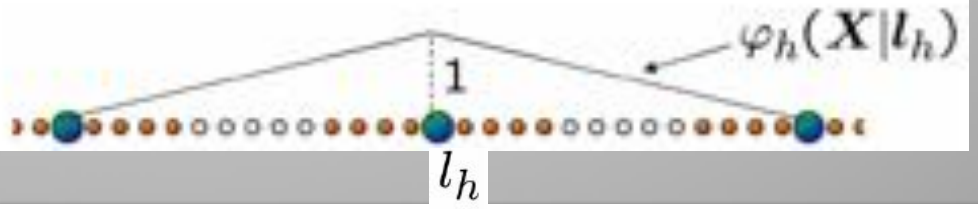

The Quasicontinuum Method (QC)

QC-reduction

- Reduced problem: $\inf_{q \in X_h} E(q)$
- Reduced equilibrium equations:

$$f_h(l|q_h) = \sum_{l \in \mathcal{L}} f(l|q_h)\varphi_h(l|l_h) = 0$$

where

$$f(l|q_h) = \frac{\partial E}{\partial q(l)}(q_h).$$

- number of equilibrium equations $3N_h = \dim(X_h)$
- ***But***: Calculation of $f(l|q_h)$ entails sum over ***entire*** lattice $\mathcal{L}$!

The Quasicontinuum Method (QC)

Cluster–based lattice summation rules

- $C_h(\boldsymbol{l}_h) = \{\boldsymbol{l} \in \mathcal{L} \mid \|\boldsymbol{X}(\boldsymbol{l}) - \boldsymbol{X}(\boldsymbol{l}_h)\| \leq r\} \equiv$ cluster of sites centered at $\boldsymbol{l}_h$.
- Cluster summation rule:

$$S_h = \sum_{\boldsymbol{l}_h \in \mathcal{L}_h} n_h(\boldsymbol{l}_h)[\sum_{\boldsymbol{l} \in C(\boldsymbol{l}_h)} f(\boldsymbol{l})]$$

$n_h(\boldsymbol{l}_h)$ s.t. $\varphi_h(\boldsymbol{l}|\boldsymbol{l}_h)$ summed exactly.

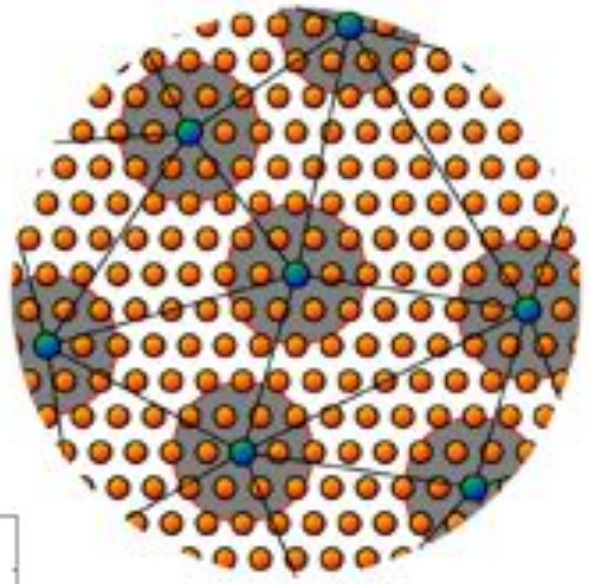

- Effect of cluster size:

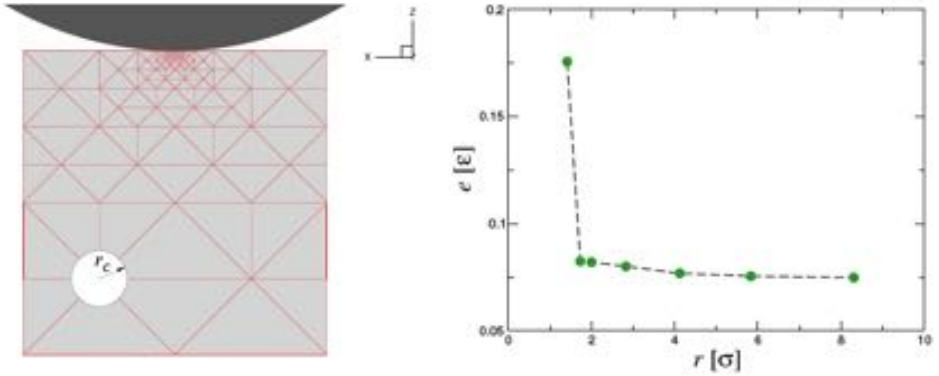

Effect of cluster size on energy error for 0.1σ indentation of $64 \times 64 \times 64$ fcc cell sample of Lennard–Jones crystal.

Quasicontnuum simulations of nanoindentation: 10^9 degrees of freedom

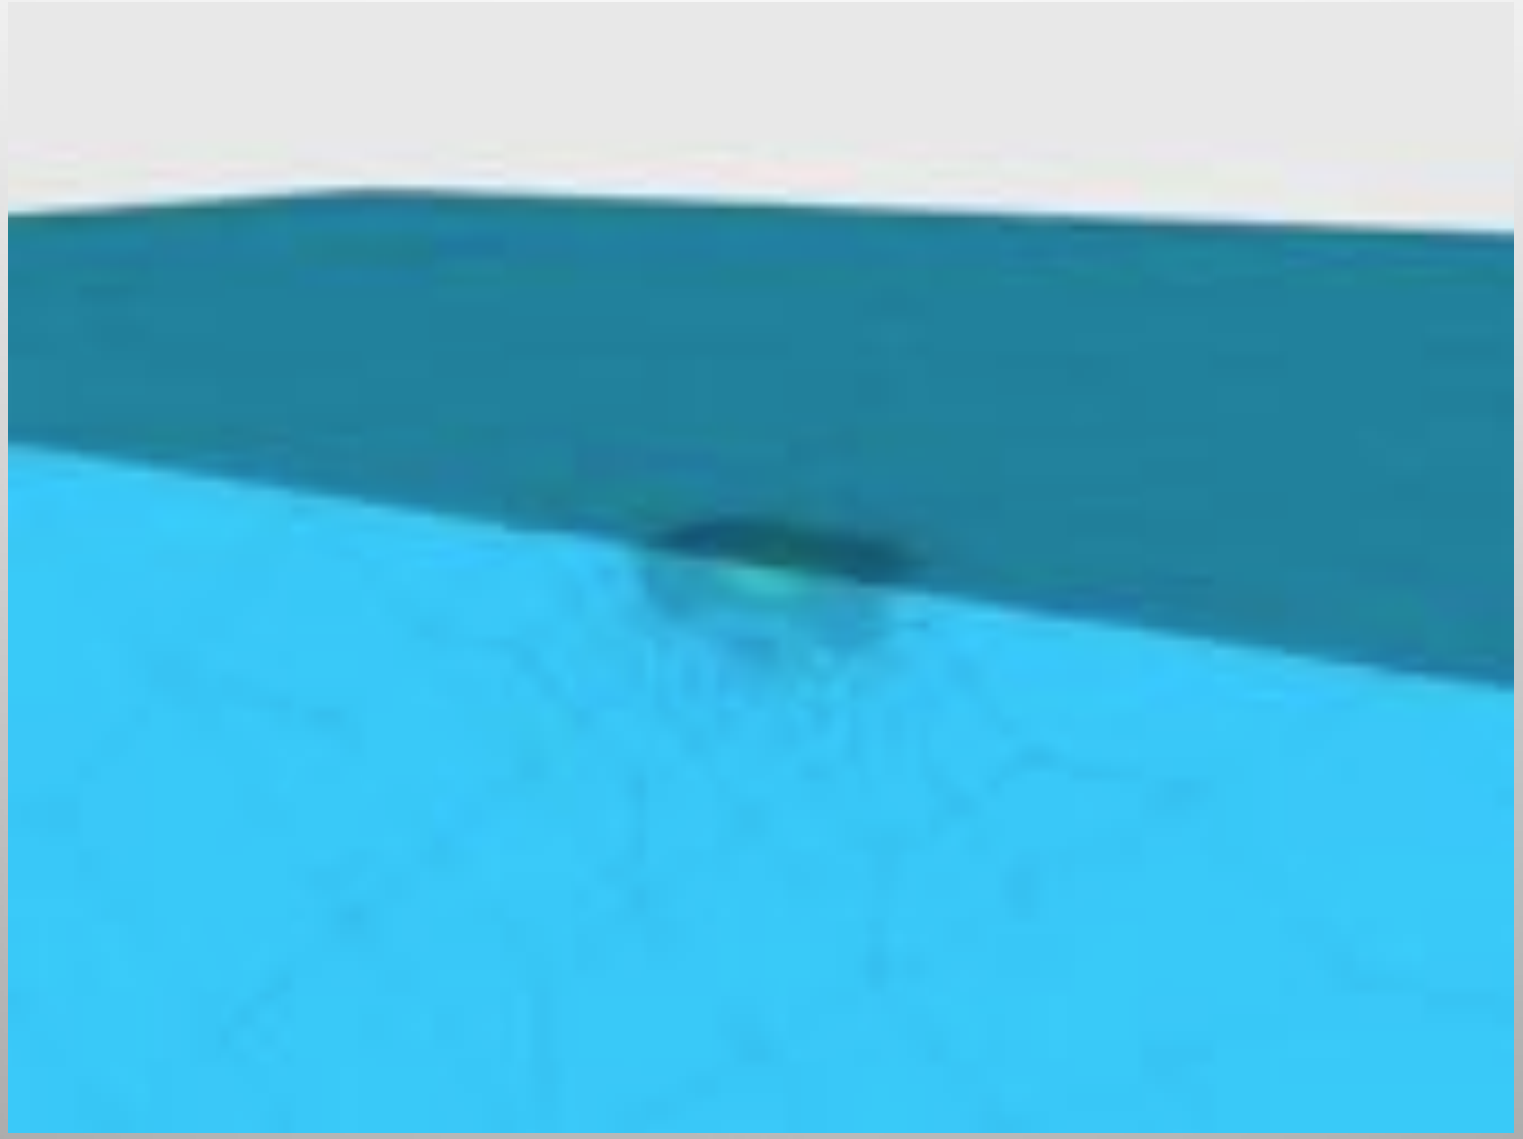

Marian, Knap, Ortiz, 2004

We use massively parallel codes for dislocations dynamics simulations

Dislocation network represented by interconnected line segments:

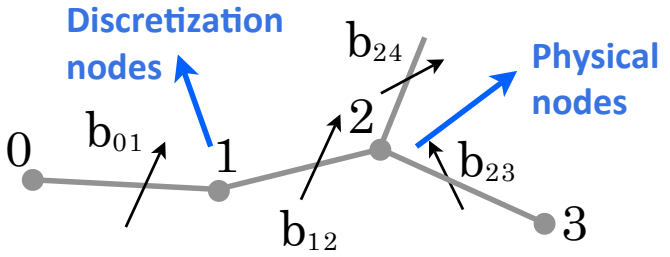

$$\sum_i \mathbf{b}_i = 0$$

Node force
Isotropic elasticity

$$\mathbf{f}_i = -\frac{\partial E(\mathbf{r}_i)}{\partial \mathbf{r}_i}$$

Node velocity
Obtained from MD

$$\mathbf{v}_i = \mathbf{M}\mathbf{f}_i$$

Move nodes
Topology changes

$$\mathbf{r}_i = \mathbf{r}_i + \mathbf{v}_i \delta t$$

ParaDiS integrates the multiplication and interactions of dislocations for simulating evolution of strength

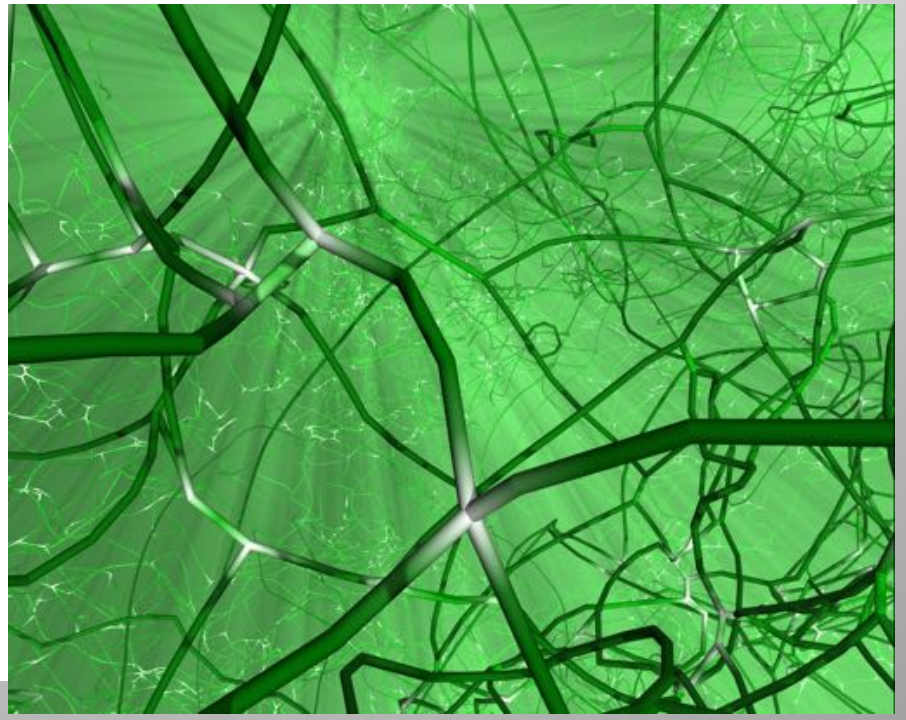

Bulatov *et al* , Nature (2006)

Large-scale simulations of irradiation hardening in Fe

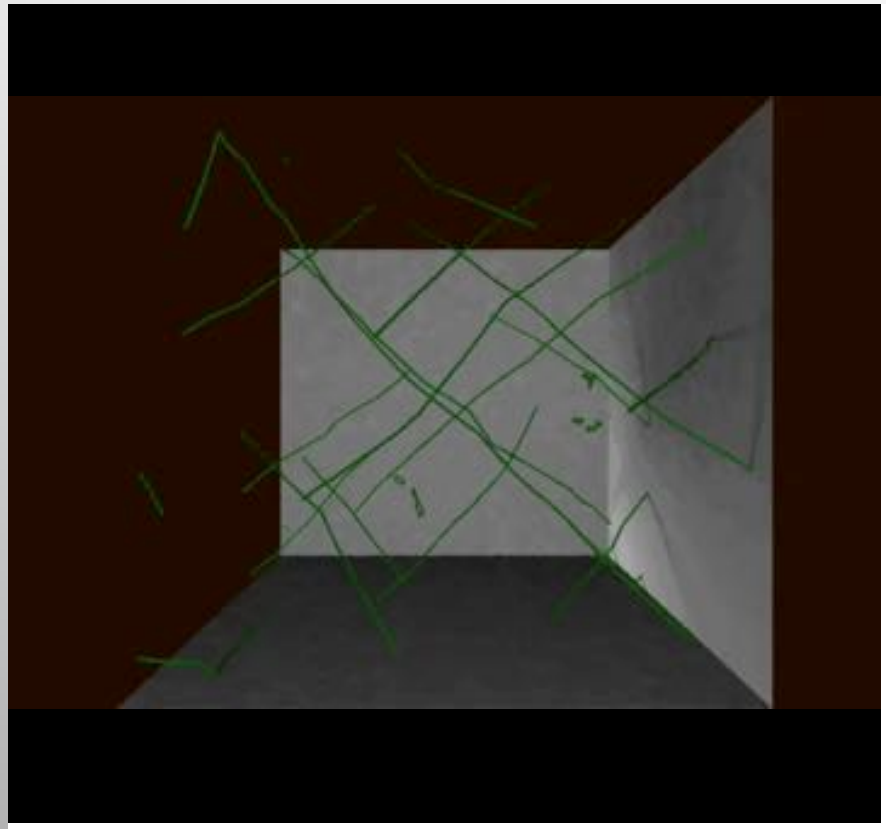

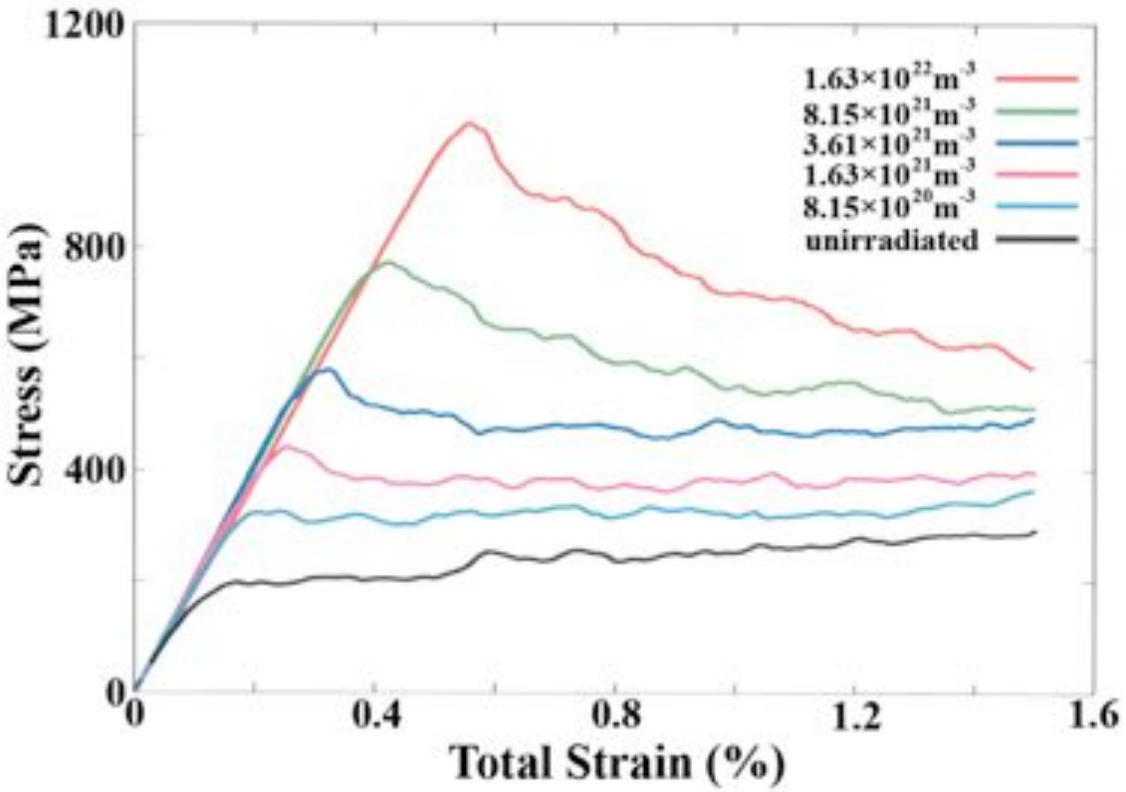

- Progressive hardening occurs as defect density (irradiation dose) increases.
- Strain hardening behavior is eventually lost.
- Unstable (softening) behavior seen at high defect densities (doses).

Simulations of plastic localization in irradiated steels

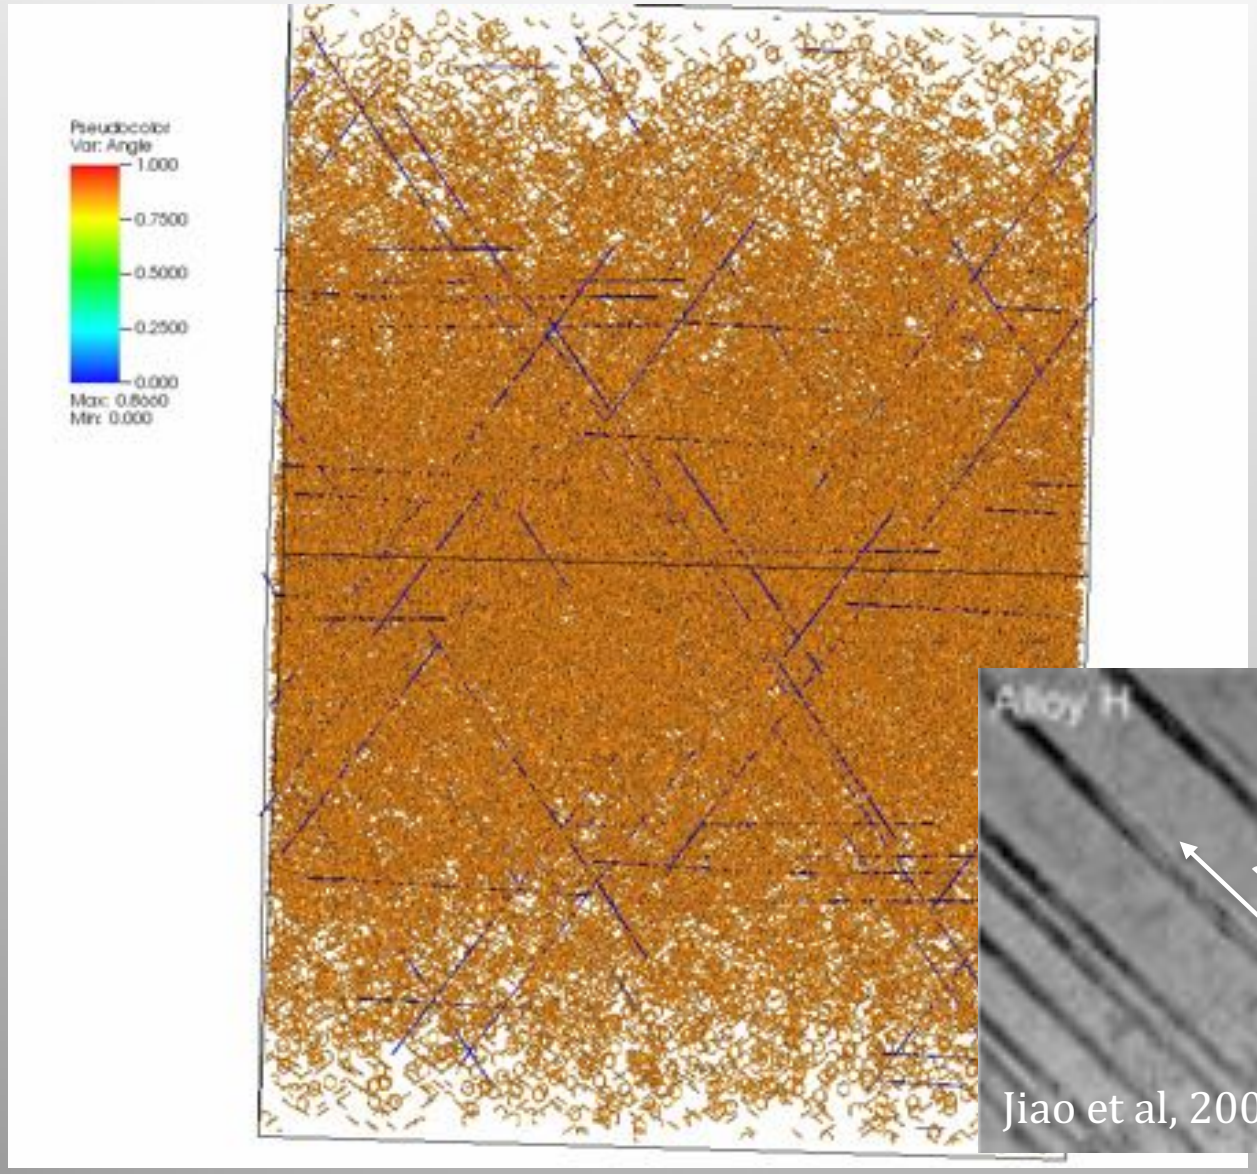

- Channels form above a critical obstacle density.
- Channels are oriented along ⟨112⟩ directions.
- Channel width around 180 nm.

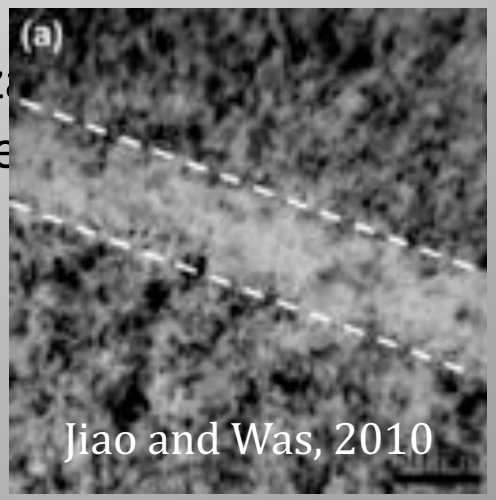

Single Crystal Plasticity Model contains Independent network and loop strength contributions

$$\mathbf{F} = \mathbf{F}^p \mathbf{F}^e$$ Multiplicative decomposition

$$\dot{\mathbf{F}}^p = \left(\sum_\alpha \dot{\gamma}^\alpha \mathbf{s}^\alpha \otimes \mathbf{n}^\alpha \right) \mathbf{F}^p$$ Flow rule

$$\dot{\gamma} = \sum_\alpha |\dot{\gamma}^\alpha| \longrightarrow \dot{\gamma}^\alpha = \dot{\gamma}_\mathrm{o} \left(\frac{\tau^\alpha}{g^\alpha} \right)^{\frac{1}{m}}$$ Power law kinetics. τ^α is RSS on slip system α

$$g^\alpha = g_\mathrm{o} + \mu b \left(\sqrt{h_\mathrm{n} \rho_\mathrm{n}} + \sqrt{h_\mathrm{d} \mathbf{N}^\alpha : \mathbf{H}} \right)$$ Slip system strength. $\mathbf{N}^\alpha \equiv \mathbf{n}^\alpha \otimes \mathbf{n}^\alpha$

Peierls resistance — strain hardening — radiation hardening

$$\mathbf{H} = \frac{3}{2V} \sum_i d_l^i \left(\mathbf{I} - \mathbf{n}_l^i \otimes \mathbf{n}_l^i \right)$$ Irradiation loop density as a tensorial quantity.

$$\dot{\mathbf{H}} = -\eta \sum_\alpha \left(\mathbf{N}^\alpha : \mathbf{H} \right) \mathbf{N}^\alpha |\dot{\gamma}^\alpha|$$ Loop density evolution.

Network density evolution:

$$h \equiv \rho_\mathrm{n} / \rho_\mathrm{o}^\perp$$

$$\dot{h} = \dot{\gamma} \left(k_1 \sqrt{h} - k_2 h \right),$$ where: $$k_2 = k_{2\mathrm{o}} \left(\frac{\dot{\gamma}_\mathrm{ko}}{\dot{\gamma}} \right)$$

Kocks-Mecking hardening law

Calibration and fitting of crystal plasticity model to dislocation dynamics simulations (lower scale)

Crystal plasticity model against dislocation dynamics results

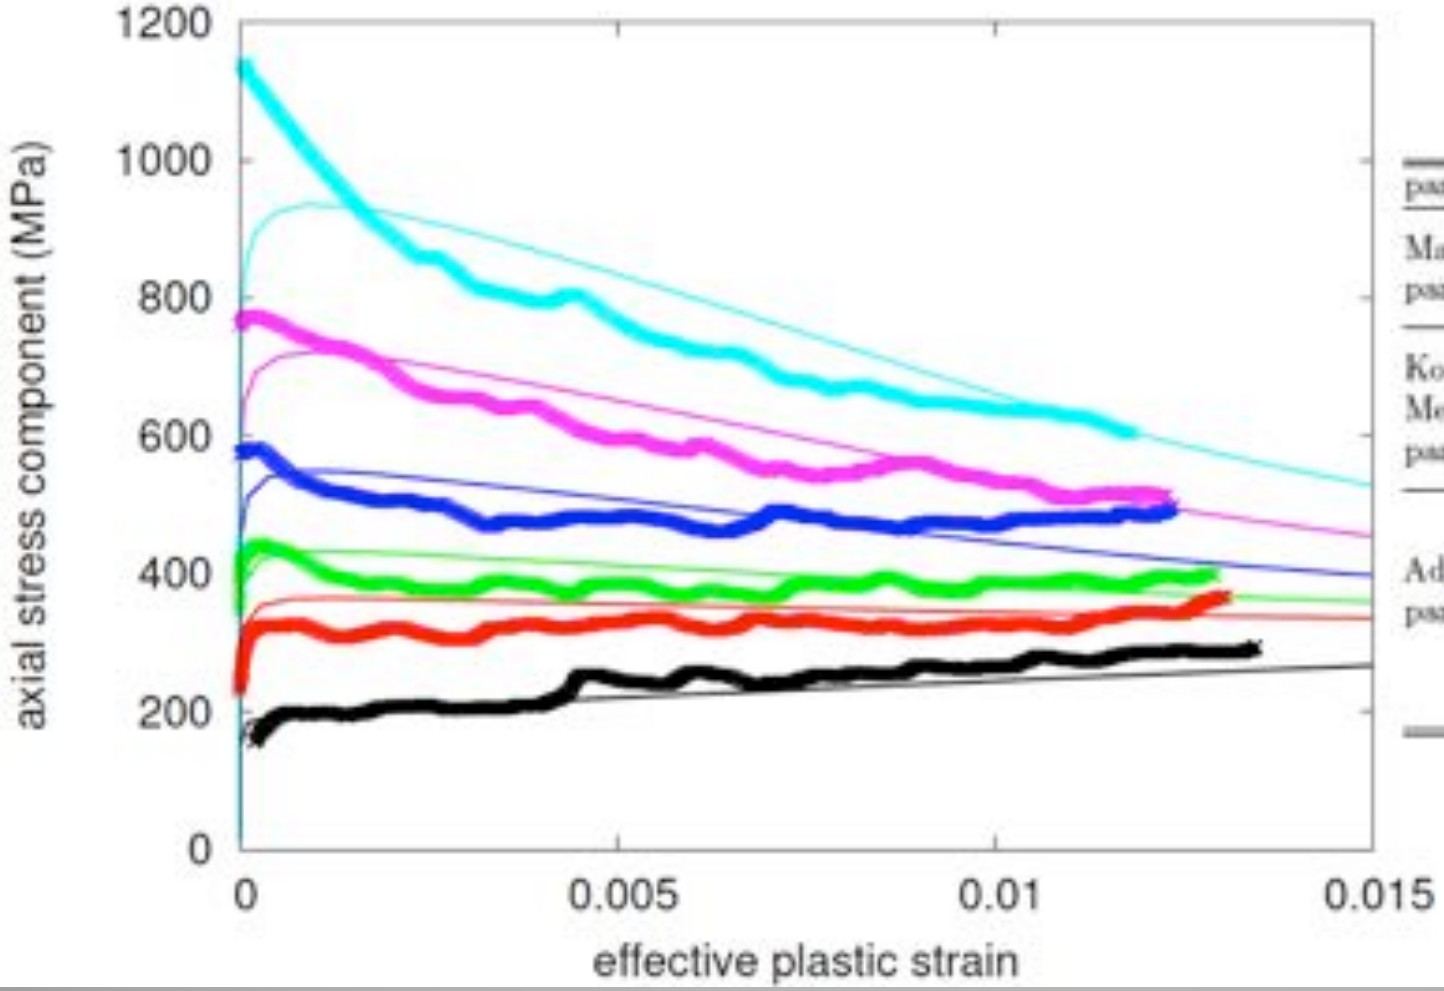

parameter	description	value	units
Material parameters	b	2.48	Å
	ρ_0	10^{12}	m^{-2}
	μ_0	86	GPa
Kocks-Mecking parameters	k_1	450	
	k_{20}	14	
	n	100	
	$\dot{\gamma}_{s0}$	10^{10}	s^{-1}
Adjustable parameters	η	100	
	h_a	0.125	
	h_d	0.675	
	g_0	90	MPa
	$\dot{\gamma}_0$	3×10^4	s^{-1}
	m	0.05	

FEM Polycrsytal Simulation of the Single Crystal Constitutive Model

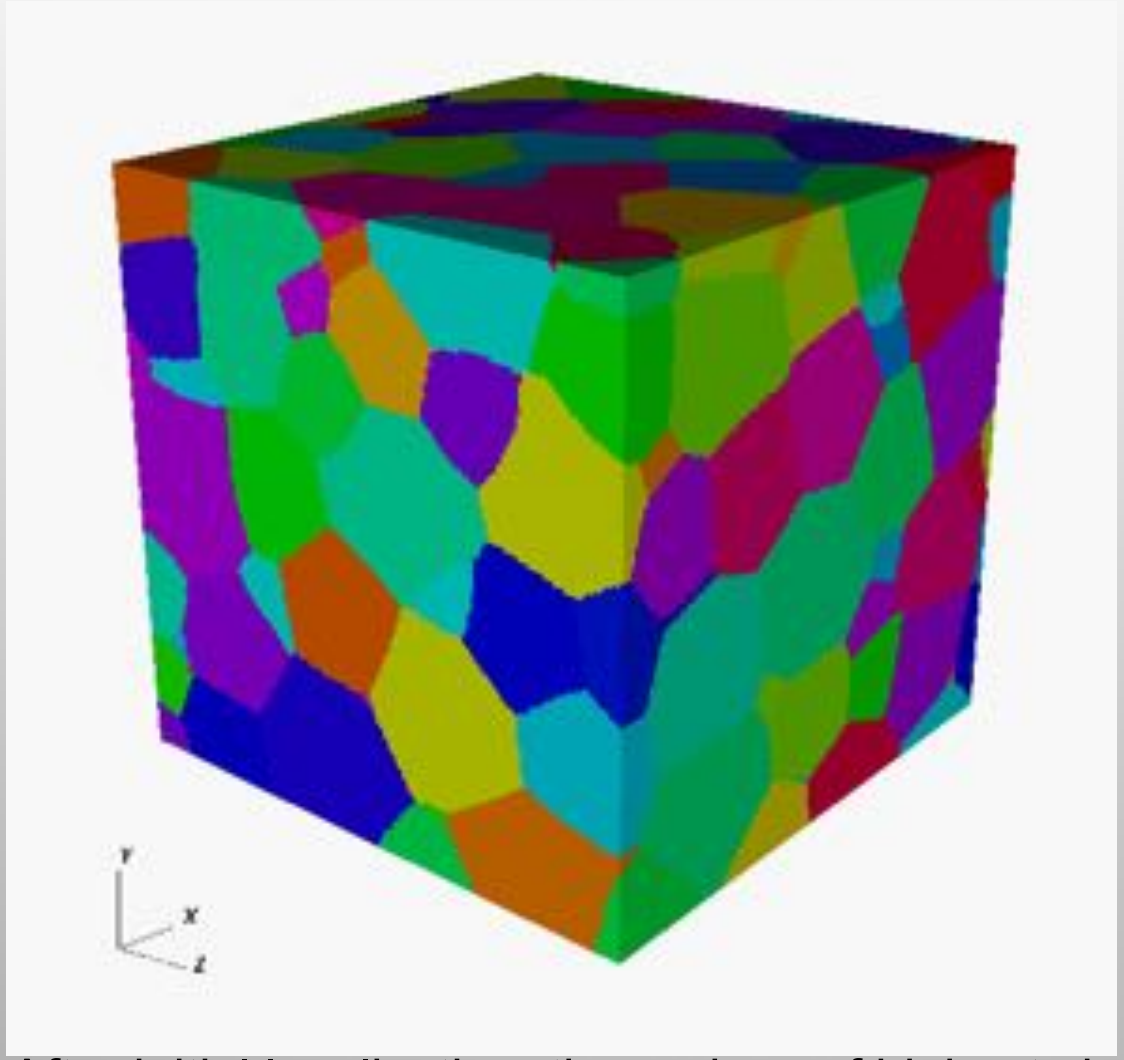

- Simulations done with implicit quasi-statics (inertial terms neglected)
- Compatibility enforced by construction
- Stress equilibrium enforced in the weak sense
- Simulations conducted under isothermal conditions
- 8-node hexahedral elements used (8 quadrature points)

After initial localization, the regions of highest shearing rate (darker contrast) split and spread across the grains

Stress-Strain Curve with Spatial Deformation Observation (Unirradiated)

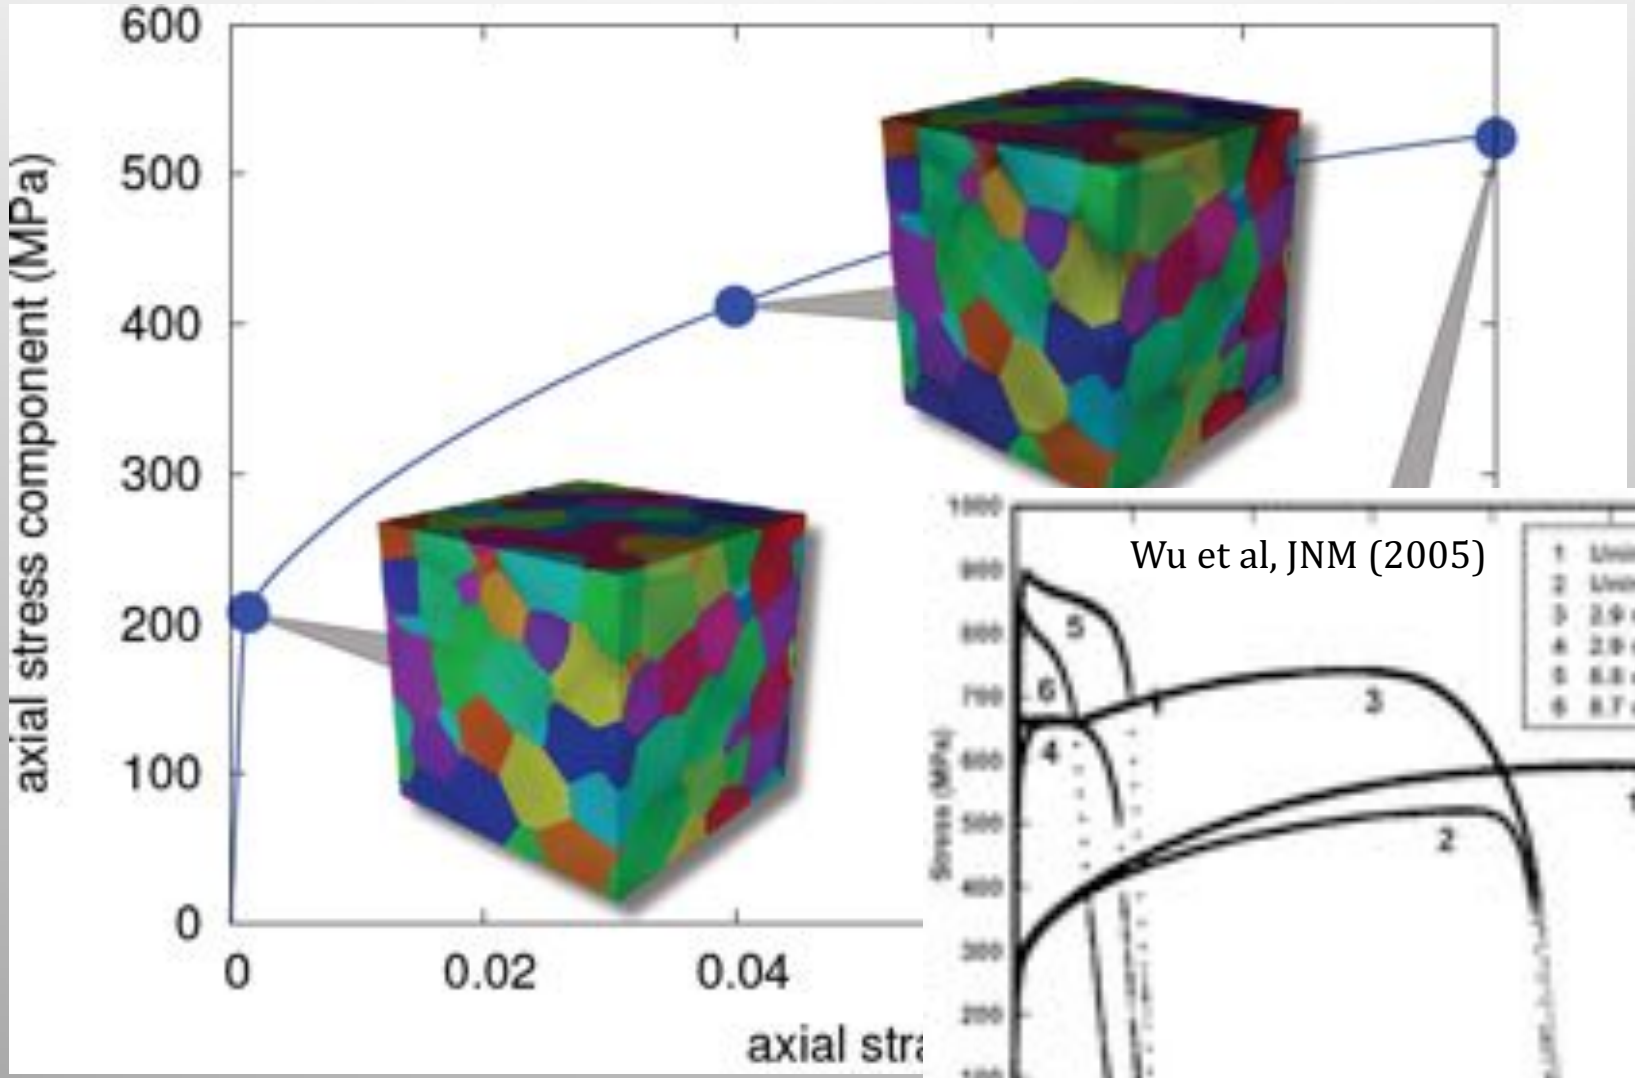

Correlation of Features of Stress-Strain Curve with Spatial Deformation Observations

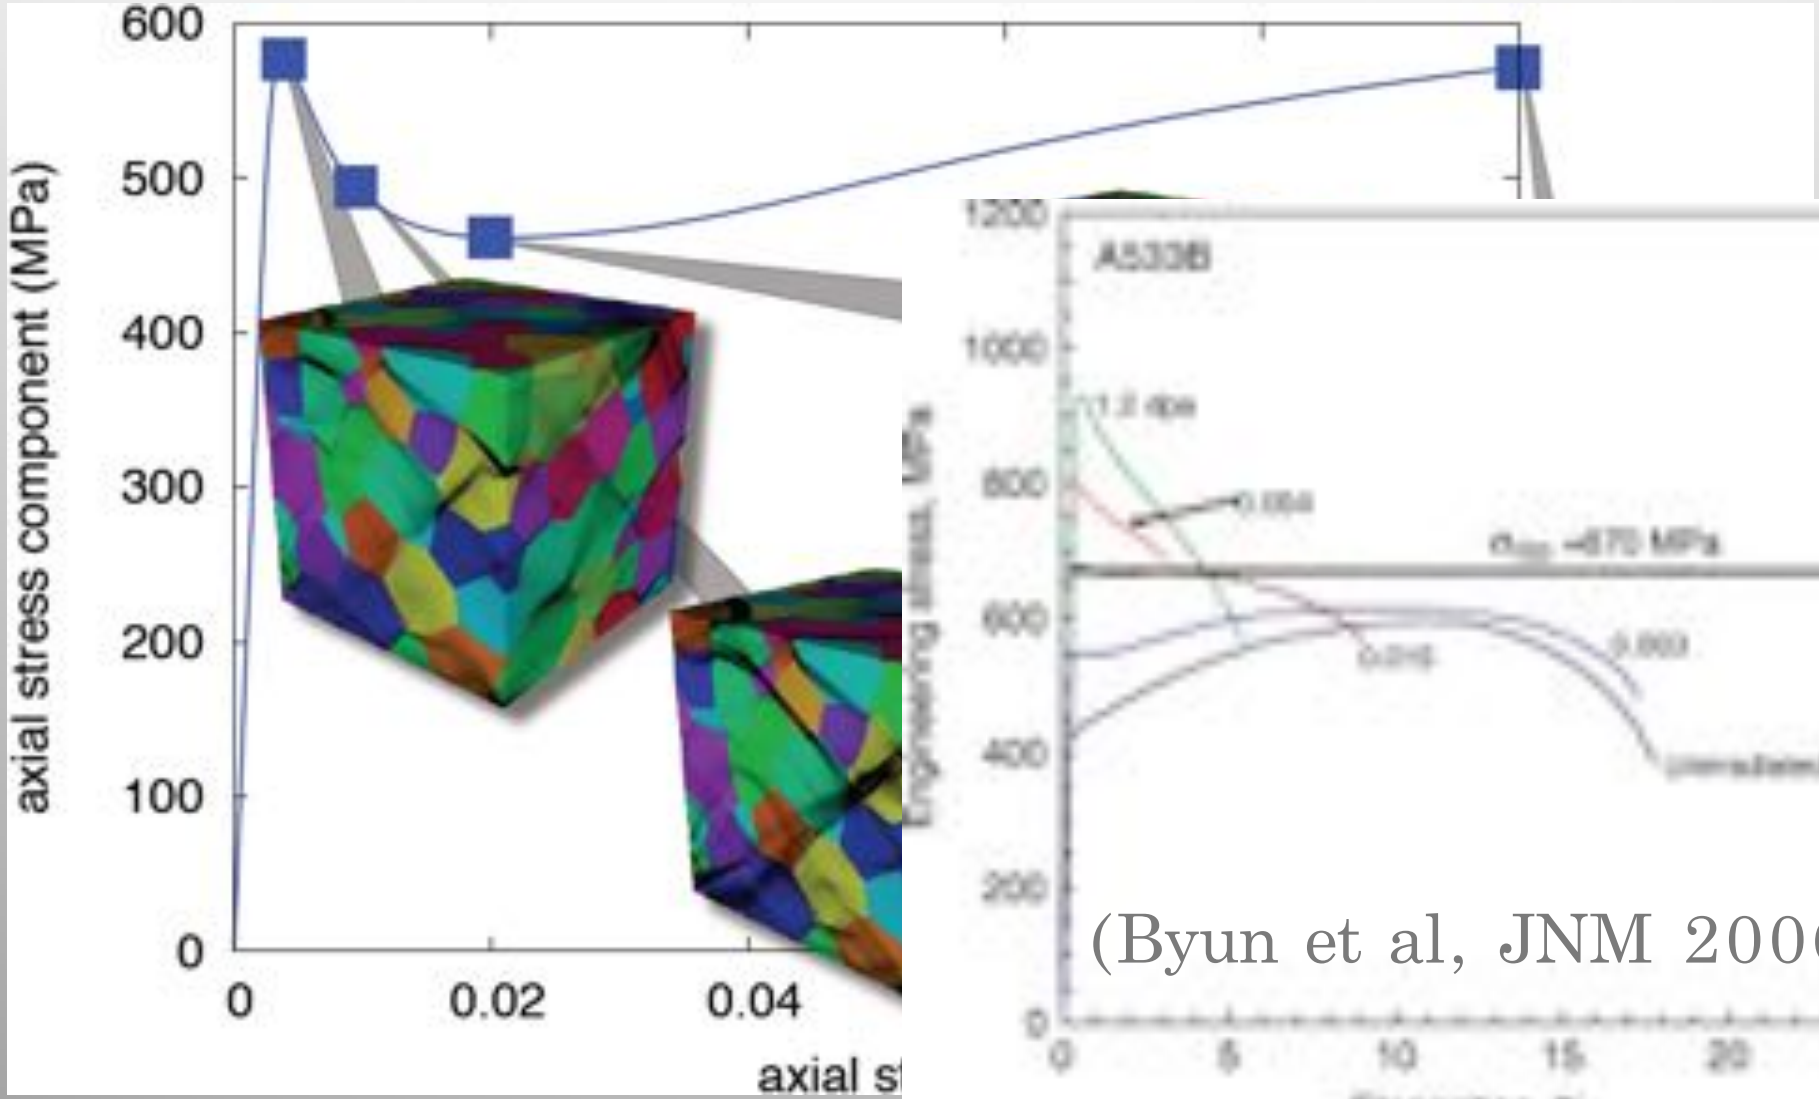

Deterministic methods will have to become non-deterministic for UQ purposes

- Stochasticity associated with UQ is not an intrinsic feature in most models.
- Using ‘naturally-stochastic’ methods such as kinetic Monte Carlo (KMC) is gaining more traction but serial versions are slow.
- Can KMC take advantage of tera/peta/exa scale computing?

What is the place of kMC in materials science simulations?

If we look at the number of research papers published over the last half a century in a number of sub-disciplines of materials science simulations, we observe several features:

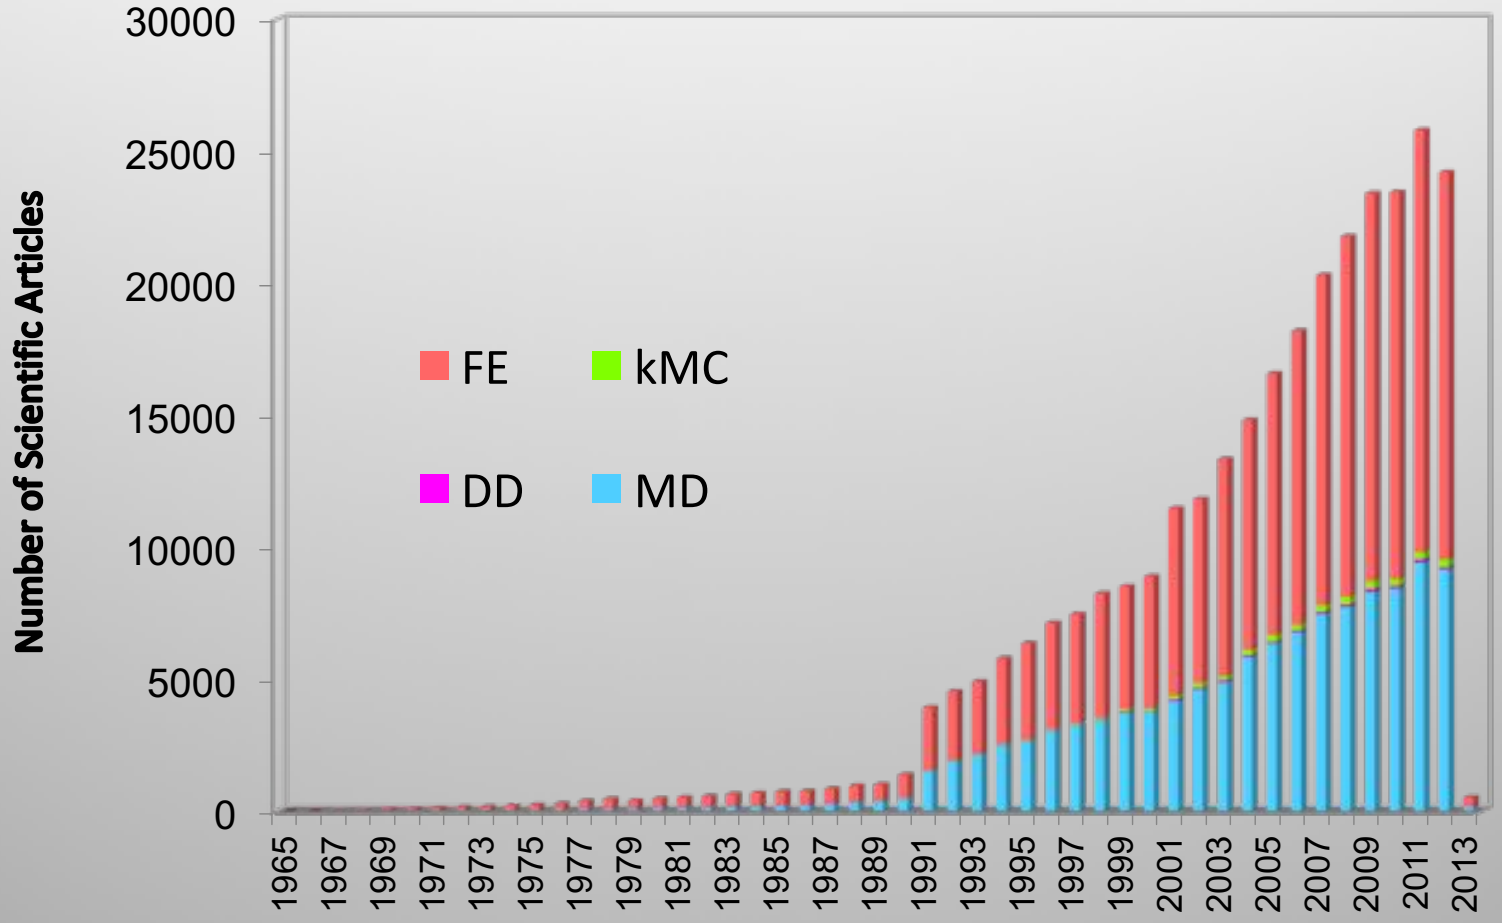

Of these, MD, FE, and DD are amenable to large-scale parallelization.

What about parallelization of kMC?

The parallelization effort is just marginal relative to the volume of kMC work.

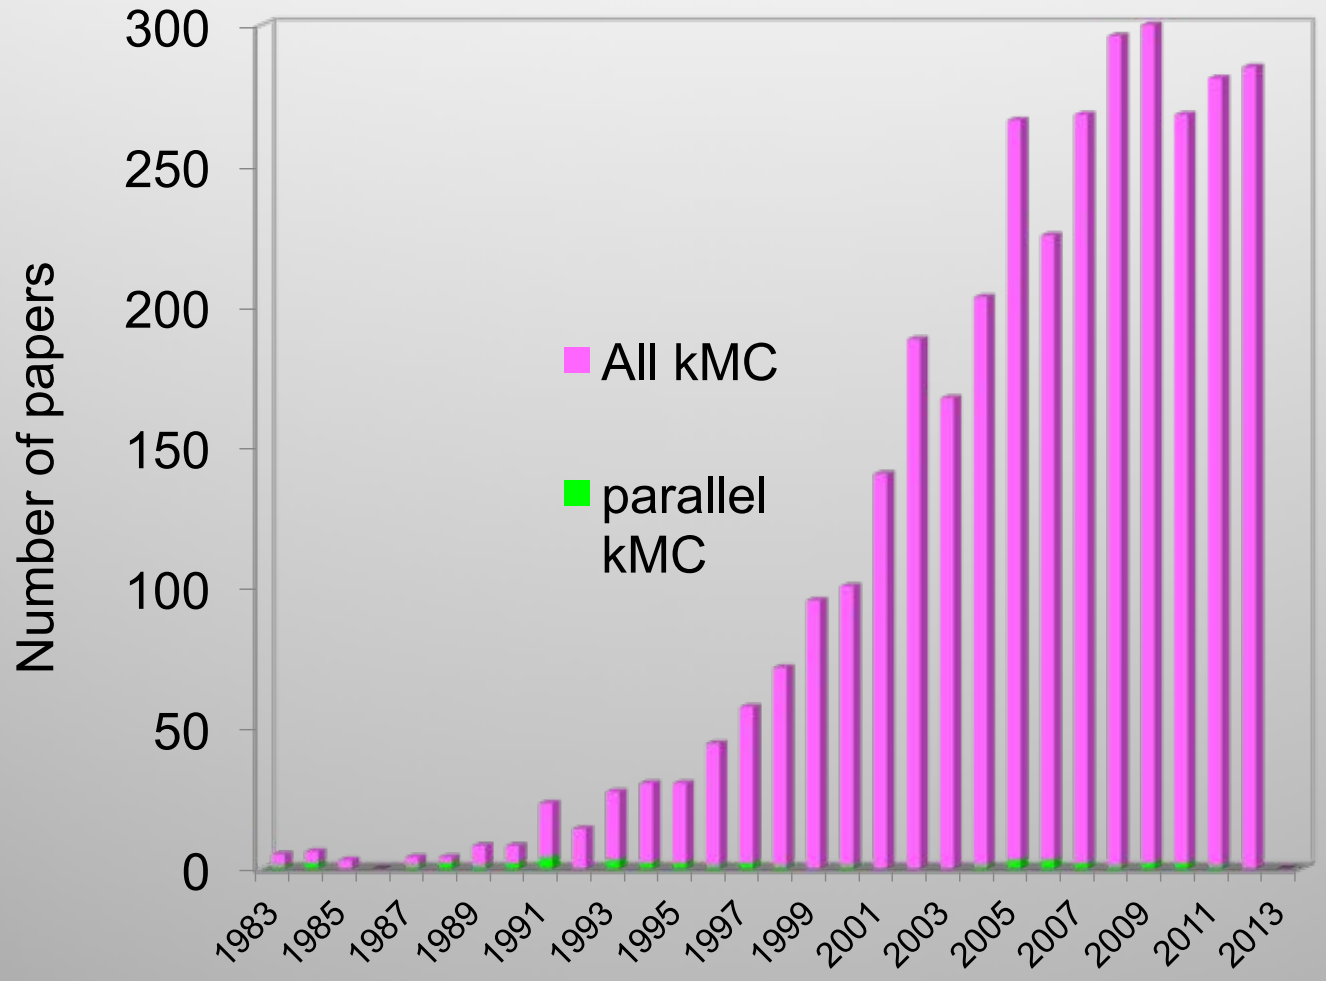

Parallel kMC algorithms to study large kinetic systems

- Discrete event kinetics are inherently difficult to parallelize.
- Traditional parallelization approaches based on asynchronous kinetics (Lubachevsky 1988, Jefferson 1985, Korniss 2003).
- Causality errors arise with these approaches: mutually affecting events occurring in different domains.
- This requires 'roll-back' techniques to reconcile the time evolution of different processors.
- This leads to implementation complexity and regions of low efficiency.
- Conservative and optimistic (semi-rigorous) algorithms have been proposed (Amar and Shim 2003, Shim and Amar 2005, Fichthorn et al.)

Asynchronous methods give rise to spatial fluctuations of time advancement

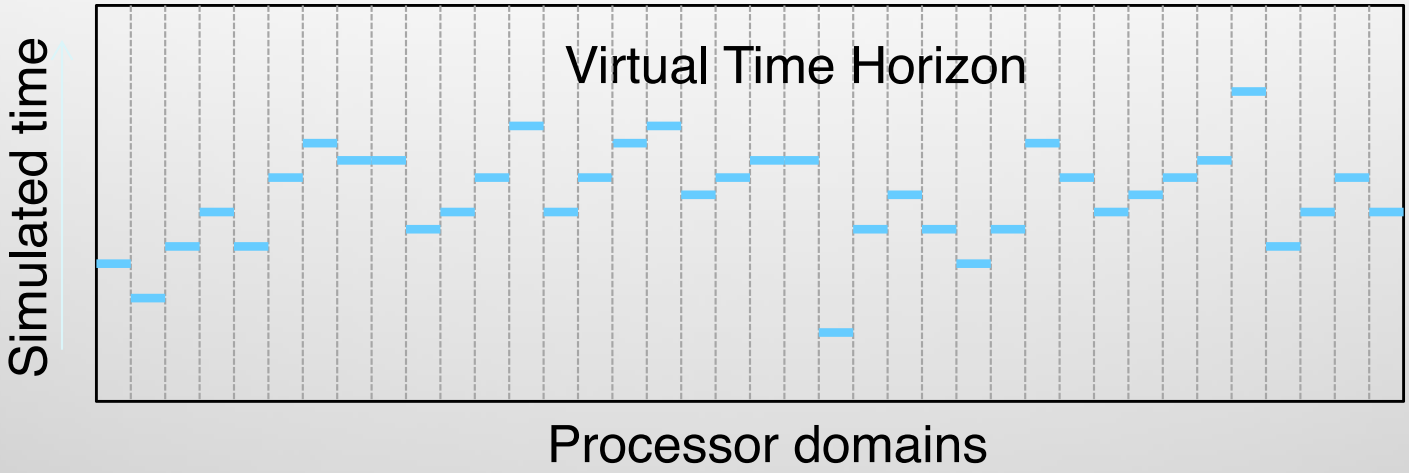

Suppressing Roughness of Virtual Times in Parallel Discrete-Event Simulations

G. Korniss,[1]* M. A. Novotny,[2] H. Guclu,[1] Z. Toroczkai,[3] P. A. Rikvold[4]

www.sciencemag.org SCIENCE VOL 299 31 JANUARY 2003

Fig. 1. Snapshots of the virtual time horizons in the steady state for the one-site-per-PE conservative PDES scheme with $N_{PE} = 10{,}000$ after $t = 10^6$ parallel algorithmic steps for (A) the original algorithm ($p = 0.00$) and (B) the modified one ($p = 0.10$) with quenched random connections. The vertical scale is the same in (A) and (B).

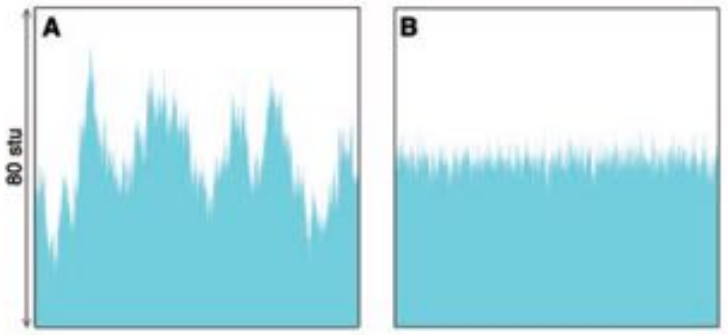

Asynchronous methods are highly complex

- The VTH problem is a bottleneck of parallel discrete event simulations.
- What about *synchronous* methods?

We have developed a synchronous parallel kMC algorithm to study general systems

Assume a spatial domain containing *N* walkers:

Each walker defined by a rate q_i, $R_{\text{tot}} = \sum_i q_i$

Perform spatial domain decomposition:

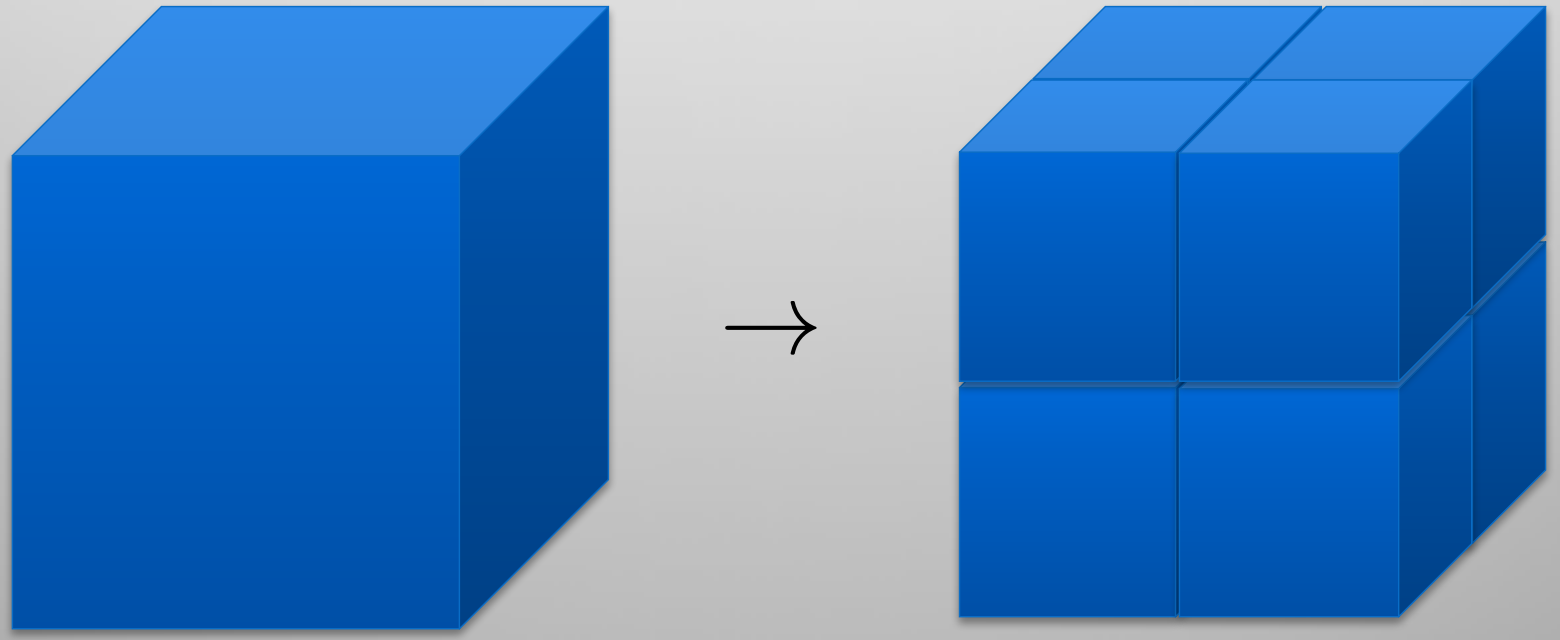

We have developed a synchronous parallel kMC algorithm to study general systems

Now, for parallel kMC, perform K (4) domain partitions and construct frequency lines:

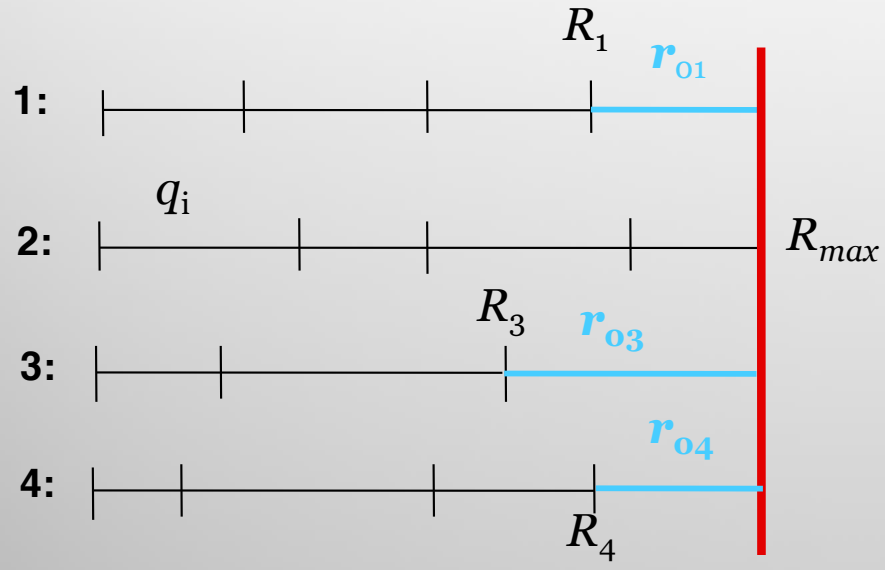

The r_{0k} are the 'dummy' rates (no event) that ensure synchronicity:

$$R_{max} = r_k + r_{0k}$$

$$R_{tot} = \sum_k r_k \leq K R_{max}$$

$$\left.\begin{aligned} \delta t_p &\approx \frac{1}{R_{max}} \\ \delta t_s &\approx \frac{1}{R_{tot}} \end{aligned}\right\} \delta t_p \leq K \delta t_s$$

For optimum scalability, perform domain decomposition subject to the following constraint:

$$\min \left[\sum_k r_{0k} \right]$$

Parallel kMC algorithm

1. Perform spatial decomposition into *K* Ω_k domains.

$$R_{\text{tot}} = \sum_{k}^{K} r_k$$

2. Define partial aggregate rates in each Ω_k:

$$r_k = \sum_{i}^{n_k} q_{ik}$$

3. Choose the maximum partial rate as:

$$R_{\max} = \max_k \{r_k\}$$

4. Assign 'null' rates to each Ω_k such that:

$$r_{k0} = R_{\max} - r_k$$

5. Sample event from each subdomain with probability

$$q_{ik}/R_{\max}$$

6. Execute event and advance time by $\delta t_p = -\frac{\ln \xi}{R_{\max}}$

Temperature behavior of Ising systems

- At the critical temperature T_c, domains of aligned spins are created.
- These domains are defined by a correlation length ξ:

$$\xi \propto |T - T_c|^{\nu}$$

- ν is the 'scale' critical exponent.
- Critical exponents not converged for 3D.

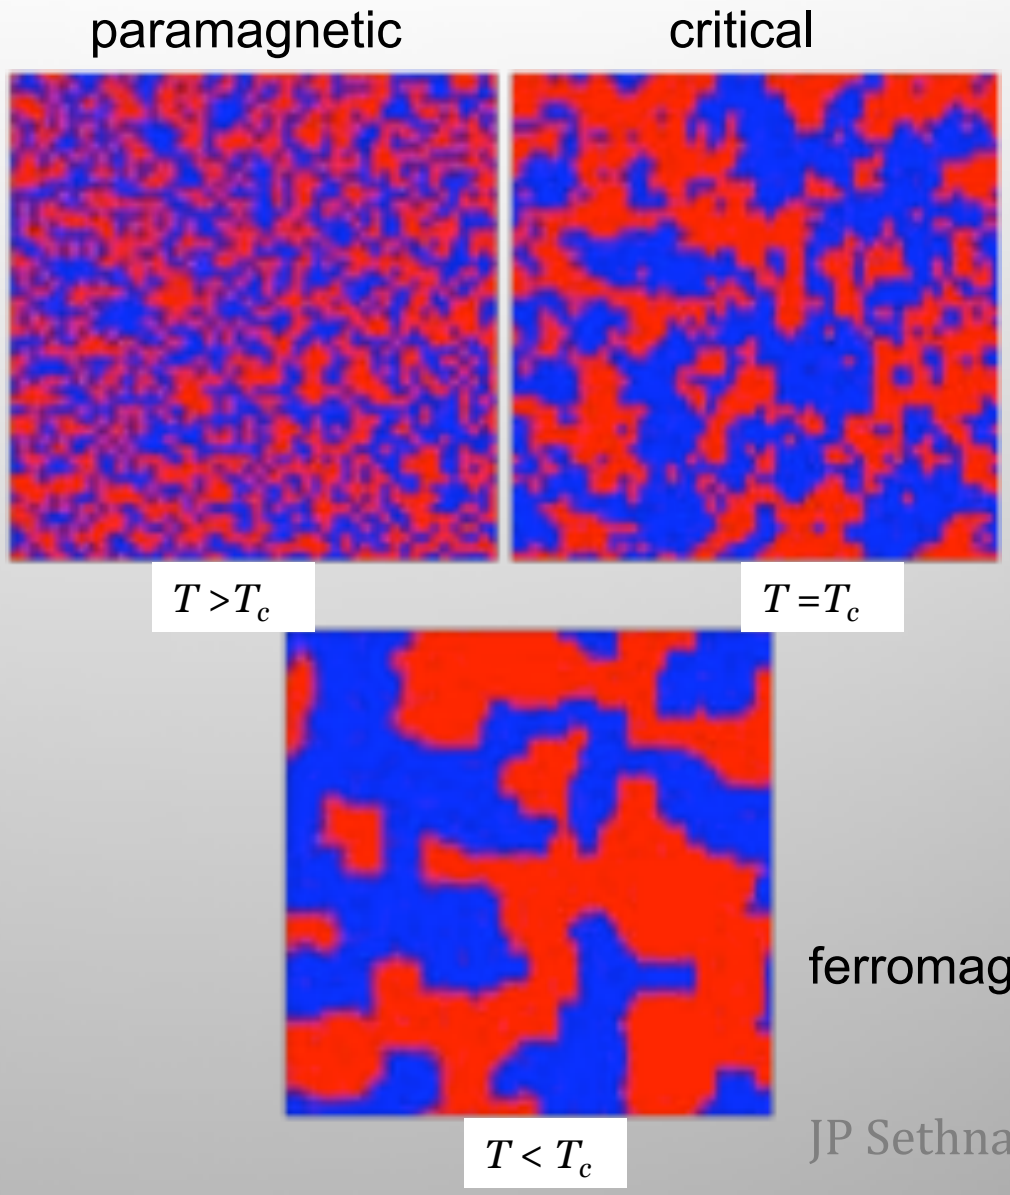

The net magnetization is the order parameter of the Ising system

$$m(\sigma) = \frac{1}{N}\sum_i \sigma_i$$

$$m(t) \propto t^{-\kappa/z\nu}$$

The value of the critical exponents in 1D and 2D is analytically known and can be converged for 4 and higher dimensions.

In 3D: no converged numerical solution.

$$z = -\frac{\kappa}{\nu}\left[\frac{d(\log m)}{d(\log t)}\right]^{-1}$$

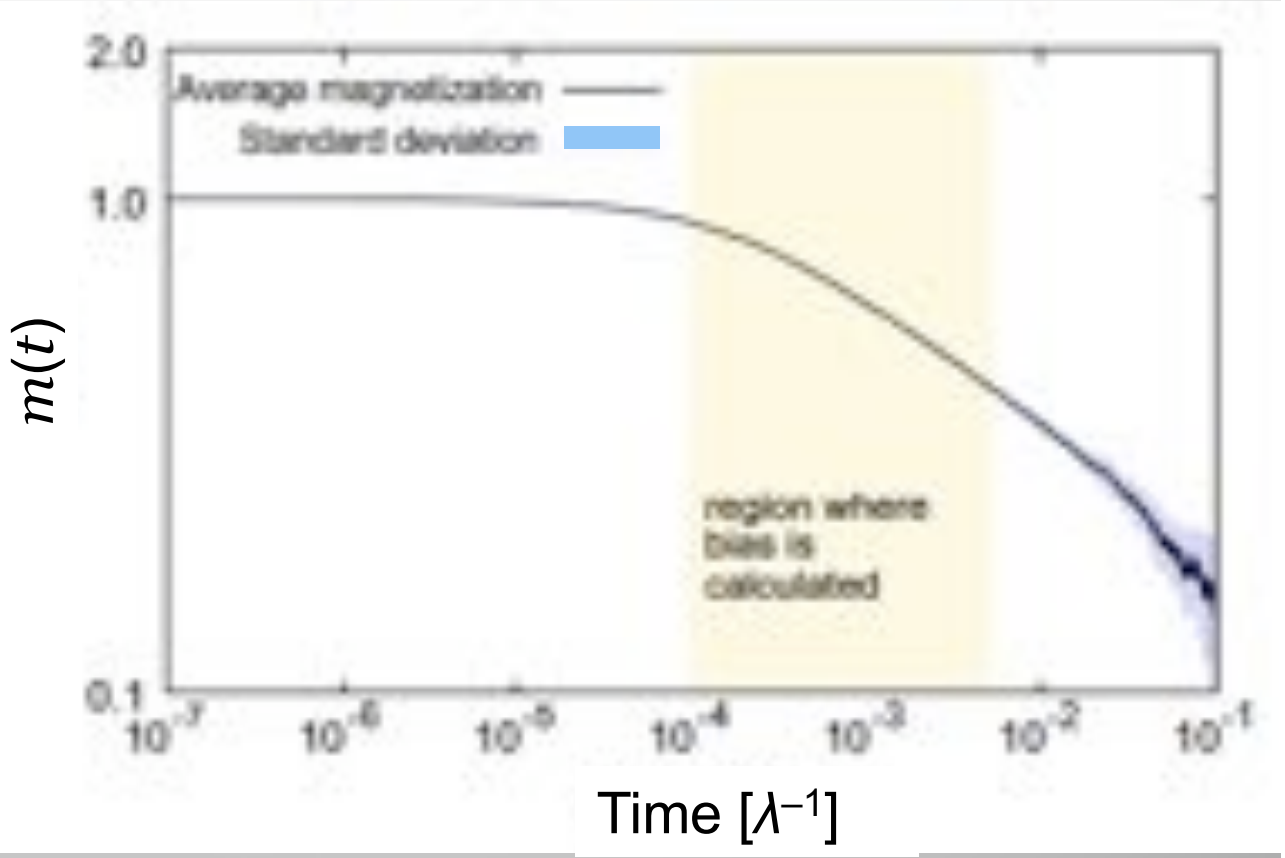

Serial kMC not sufficient

Boundary conflicts appear when using interacting particles

Boundary conflicts appear when mutually-influencing events occur simultaneously on different domains

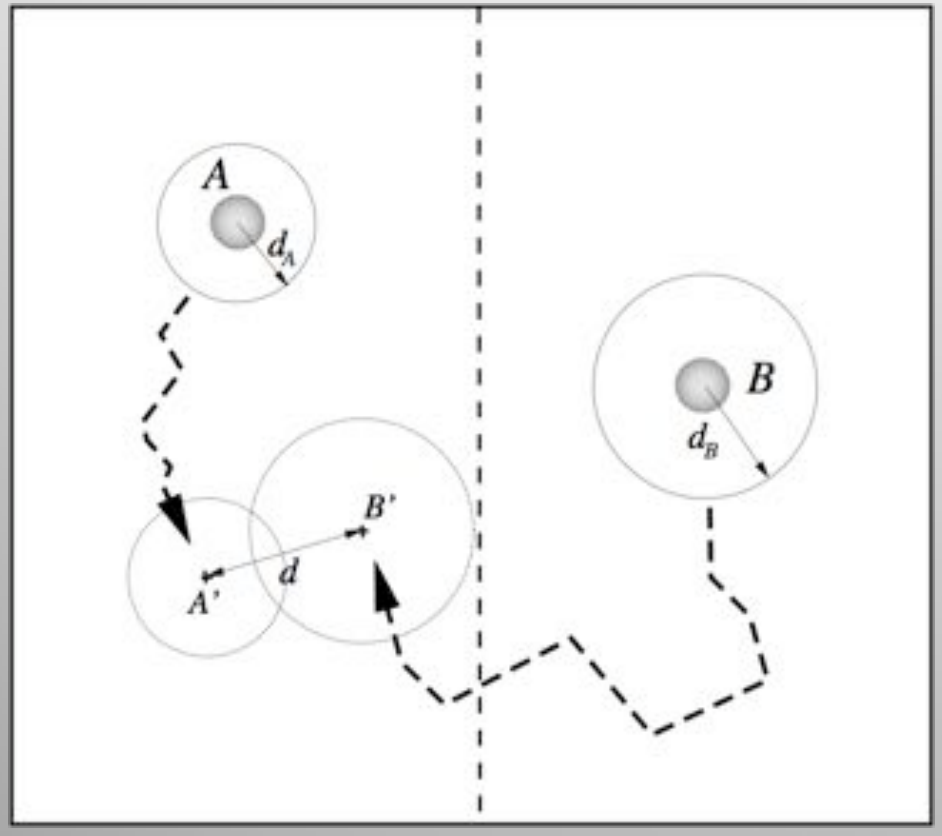

A simple solution is to use a sublattice decomposition (chess method in 2D)

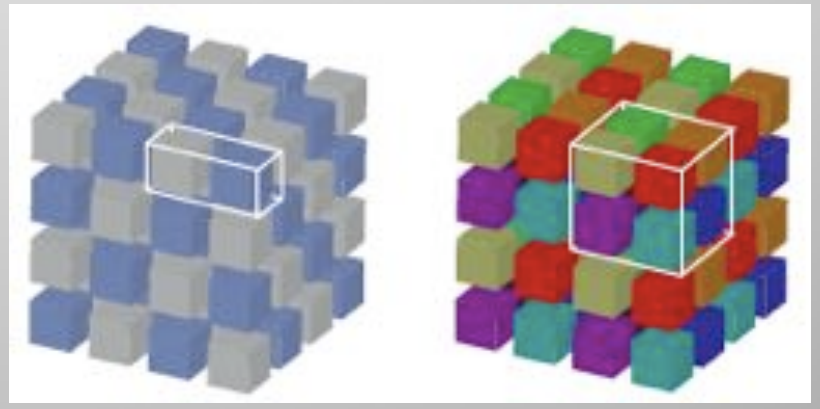

Amar *et al.* (2004, 2005)

Co-occurring events only on identically-colored subcells

Calculation of critical exponent z

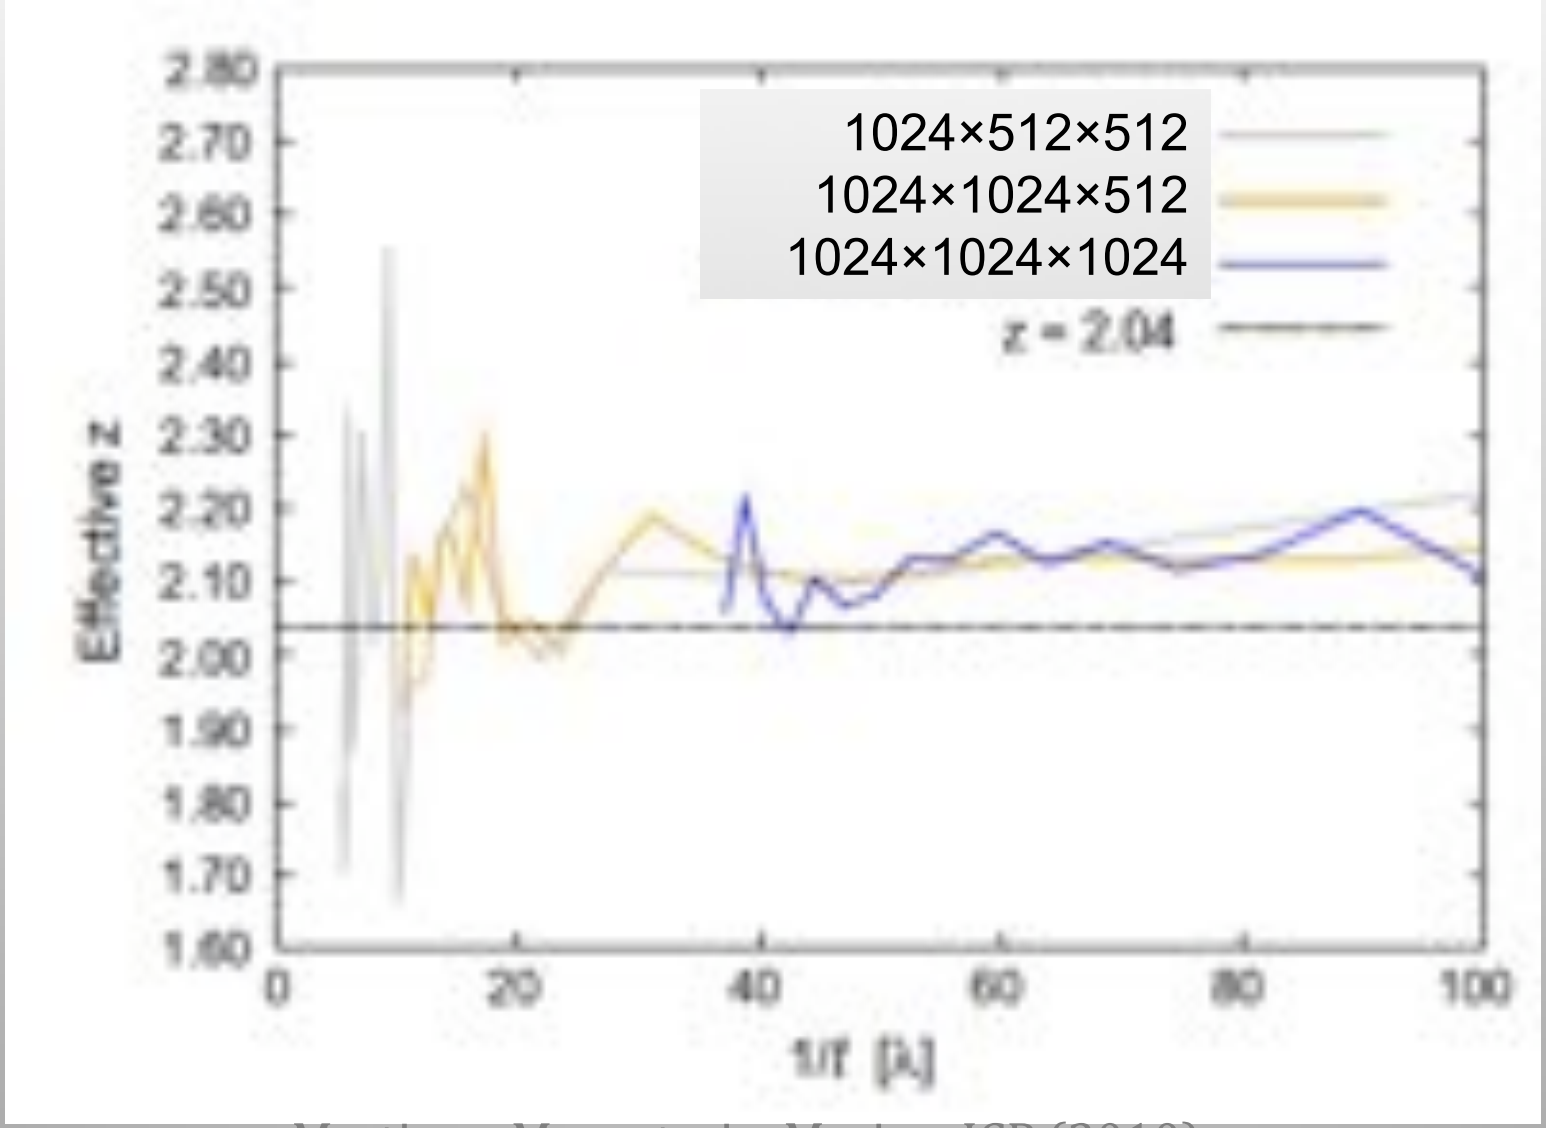

Martinez, Monasterio, Marian, JCP (2010)

Weak scaling of parallel kMC algorithm is good.

Parallel efficiency governed by local MPI calls:

$$\eta = \frac{1}{1 + a \log K + b}$$

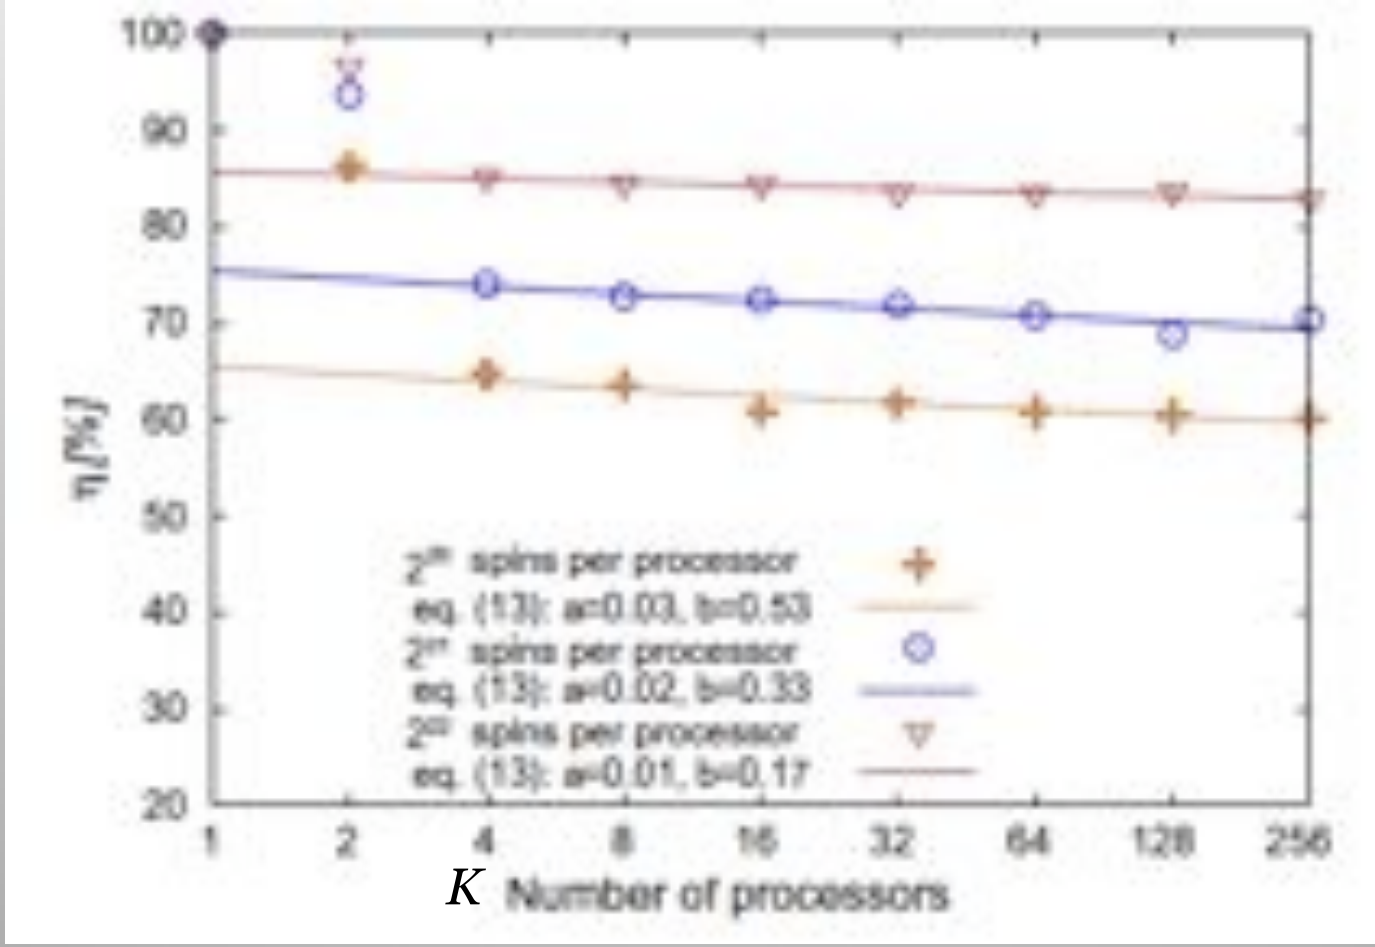

Outlook for exascale computing

- Computer architectures are becoming increasingly heterogeneous and hierarchical, with greatly increased flop/byte ratios, architectural design uncertain.
- The algorithms, programming models, and tools that will thrive in this environment must mirror these characteristics, codes will need to be rewritten.
- Standard bulk synchronous parallelism (message passing, MPI) will no longer be viable.
- Power, energy, and heat dissipation are increasingly important, presently unsolved technological bottleneck.
- Traditional global checkpoint/restart is becoming impractical (fault tolerance and resilience!).
- Analysis and visualization.

Evolution of predictive science