Crossing materials scales with LAMMPS

Vasily Bulatov

LLNL

LLNL-PRES-760376 This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Lawrence Livermore National Laboratory

Collaborators

Alexander Stukowski

Methods for data analysis and visualization

Invited talk to follow

Luis Zepeda

MD simulations
Data analysis

Nicolas Bertin

Data analysis

Tomas Oppelstrup

CS support
Runtime optimization

Support **NNSA ASC Program, Technische Universität Darmstadt**
Livermore Computing Grand Challenge Program
NNSA ATCC Campaign
DOE INCITE Program
Jülich Supercomputing Centre

Grand challenge: predict crystal plasticity from dislocation physics

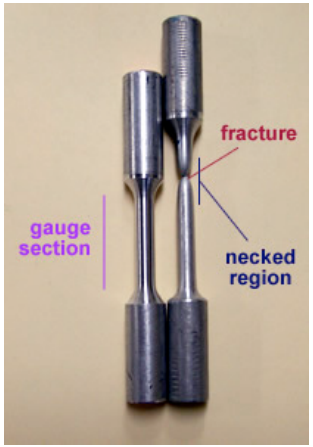

Discrete Dislocation Dynamics

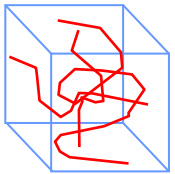

G. Canova and L. P. Kubin (1991)

Premise and promise of the DDD method

Dislocation theory, MD

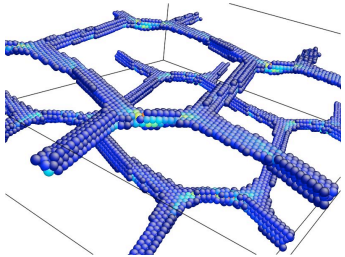

Dislocation mobility and interactions

Local rules

DDD

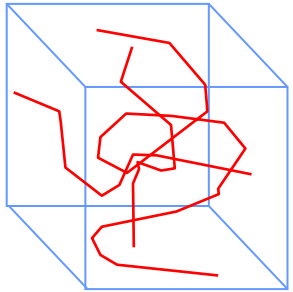

Collective response of statistically representative dislocation ensembles

Strain hardening

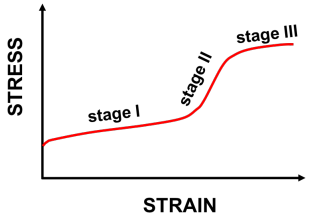

Patterns

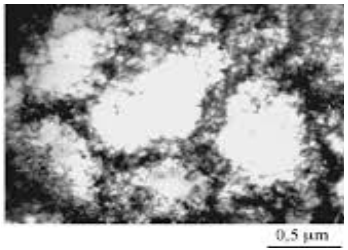

...

DDD: a darling of the multiscale materials modeling community

Huge reduction in DOFs
Microstructure → property
Physics → crystal plasticity

Two key challenges of DDD

Fidelity

Essential physics of dislocation behavior should not be "lost in translation"
(More on this later)

Computability

Computationally expensive
Still limited to small strains (<10%)
Many (solvable but unsolved) problems remain

Large-strain crystal plasticity, dislocation patterns, etc., are still 10-15 years away

Meanwhile

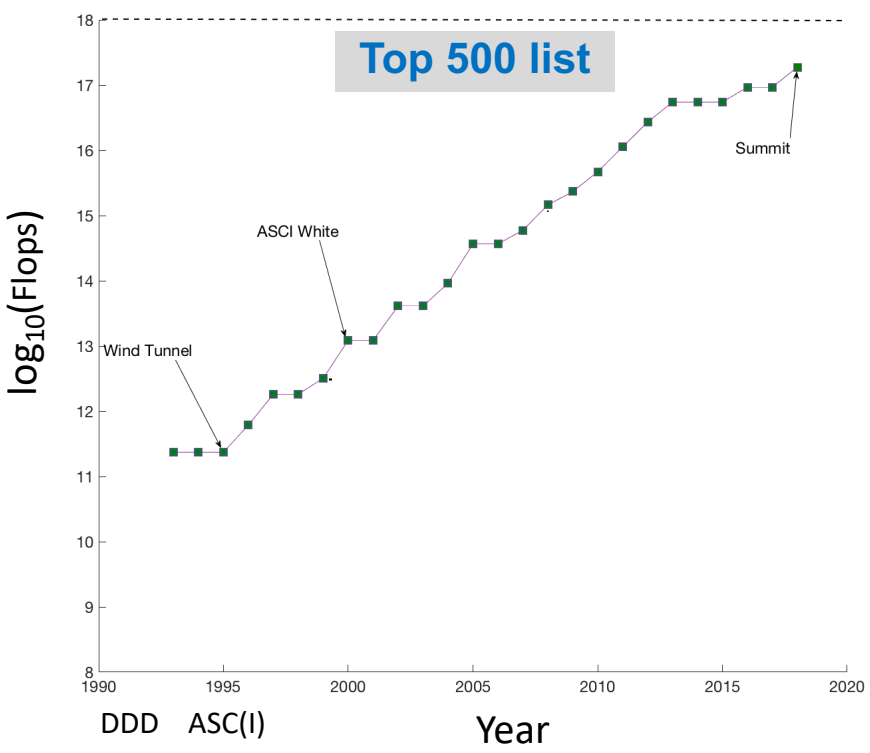

Since DDD launch in 1991

- Peak flops rate increased by 7 orders
- World's computing capacity increased by 8 orders (now ~ 10^{22} flops)

Crystal plasticity: multi-scale or cross-scale?

Multi-scaling is a necessity, not a must

Cross-scale == brute force: large enough to be statistically representative,
yet resolving every detail of atomic motion

Coarse-graining always introduces uncertainty
Unknown uncertainty is worse than a large uncertainty

By side-stepping coarse-graining inherent in DDD, CDD, etc., uncertainty of direct MD simulations of crystal plasticity reduces to the uncertainty in the interatomic potential

Development of multi-scale methods is (fun but) laborious, the MD method is mature and exceptionally robust

Back to good old Molecular Dynamics

Material dynamics in full glory: every atomic "jiggle and wiggle" (R. Feynman)

MD challenges

Severe limits on length and time scales (this is what DDD was developed to overcome)

Limited accuracy of atom-atom interaction models (shared with DDD)

Enormous amounts of data generated in fully dynamic MD simulations (will come back to this)

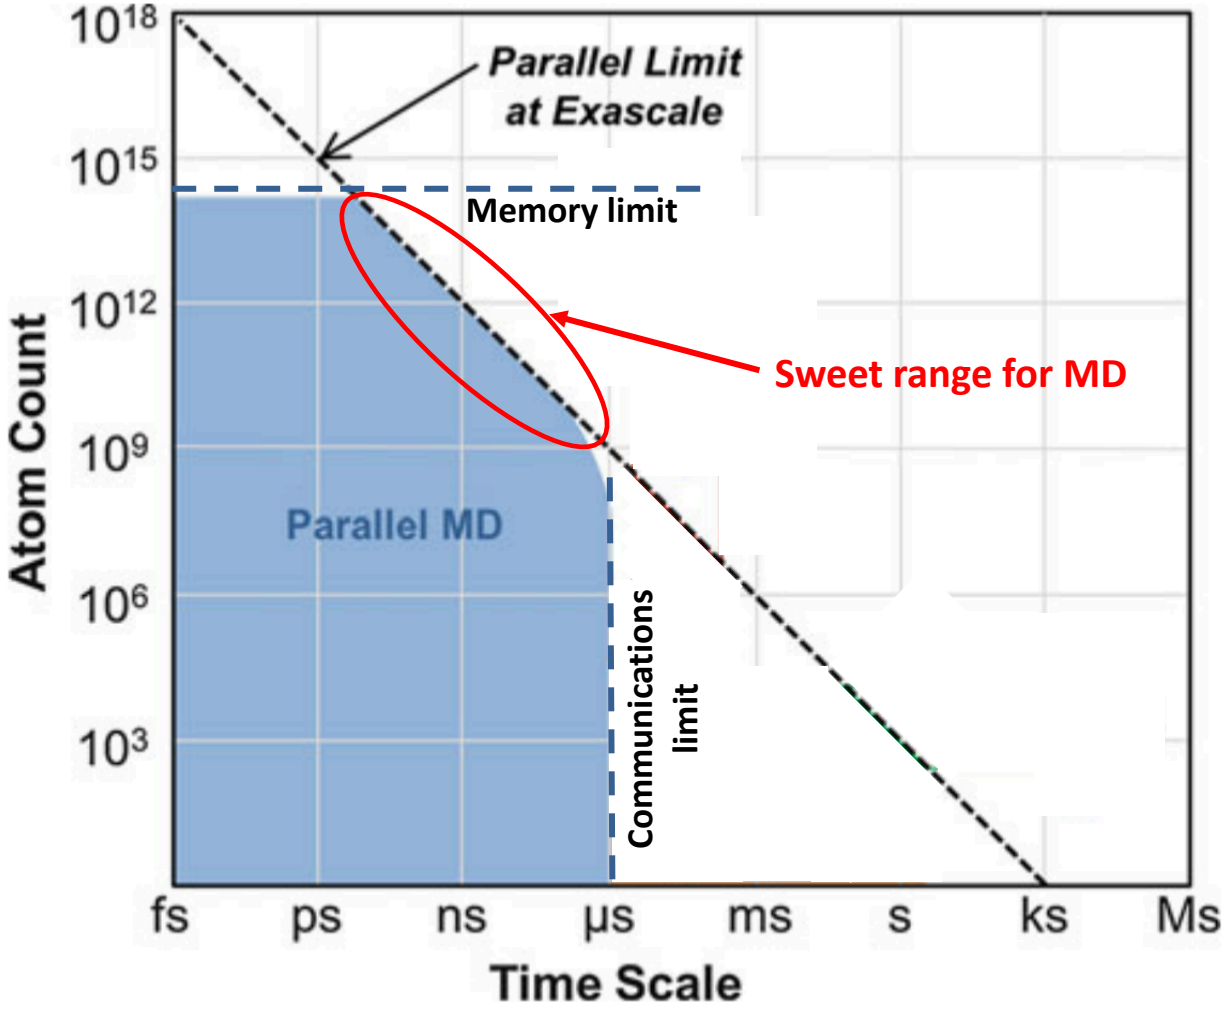
Parallel Limit
at Exascale
Memory limit
Sweet range for MD
Parallel MD
Communications
limit
Atom Count
10^18
10^15
10^12
10^9
10^6
10^3
fs
ps
ns
µs
ms
s
ks
Ms
Time Scale

Dislocation Extraction Algorithm (DXA)

Direct MD simulations of crystal plasticity

6M – 2.1B atoms, BCC lattice, 3D periodic

Li et al (2003) EAM potential model for Ta

Dislocations sourced through multiplication

Straining along one of the box axes

Constant true rate of straining

Constant temperature

Relaxation of lateral stress (Poisson effect)

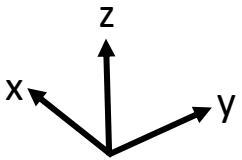

MD simulations of such magnitude were previously thought to be unthinkable

Steady/saturated flow is observed

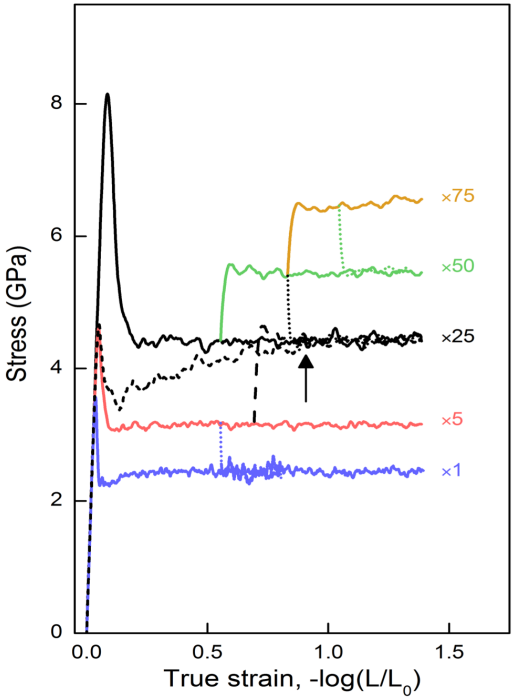

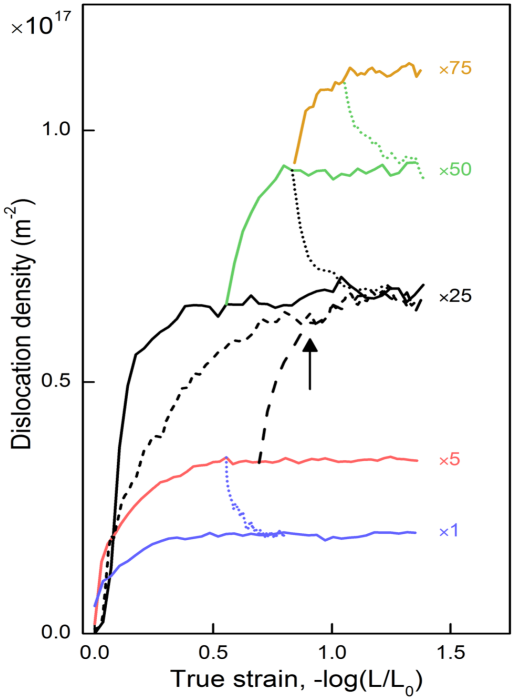

Path-independent plastic flow

All straining trajectories converging to the same ultimate straining conditions, lead to the same ultimate strength

dislocation density
line geometry
network topology
…

Kneading the metal

Brick geometry

1:2:4 ⇨ 2:4:1 ⇨ 4:1:2 ⇨ 1:2:4 ⇨ 2:4:1

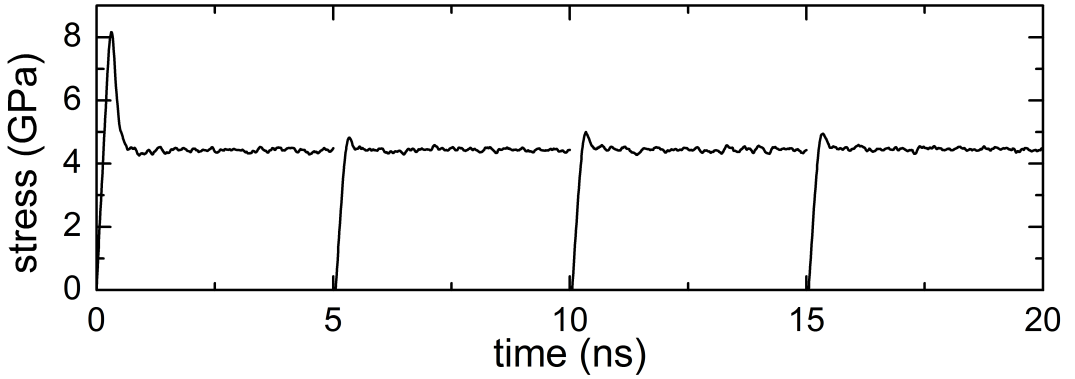

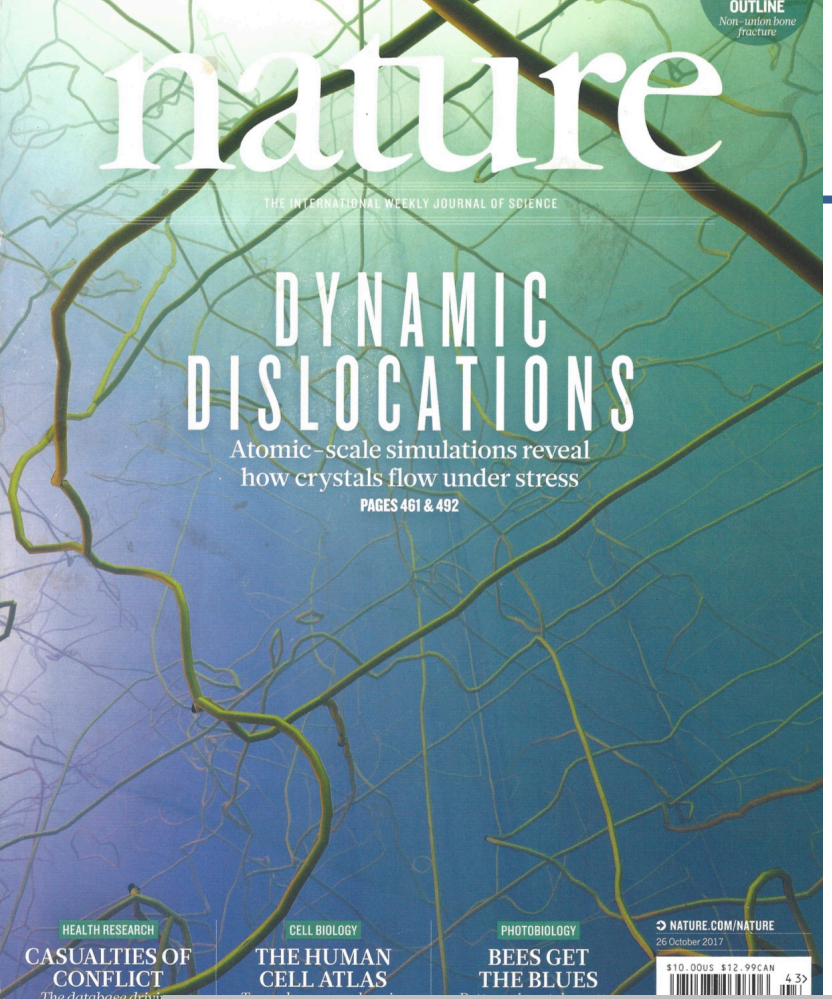

LETTER

doi:10.1038/nature23472

Probing the limits of metal plasticity with molecular dynamics simulations

Luis A. Zepeda-Ruiz[1], Alexander Stukowski[2], Tomas Oppelstrup[1] & Vasily V. Bulatov[1]

Ordinarily, the strength and plasticity properties of a metal are defined by dislocations—line defects in the crystal lattice whose motion results in material slippage along lattice planes[1]. Dislocation dynamics models are usually used as mesoscale proxies for true atomistic dynamics, which are computationally expensive to perform routinely[2]. However, atomistic simulations accurately capture every possible mechanism of material response, resolving strain-path-independent steady state of plastic flow in which the flow stress and the dislocation density remain constant as long as the conditions of straining thereafter remain unchanged. In this distinct state, tantalum flows like a viscous fluid while retaining its crystal lattice and remaining a strong and stiff metal.

The plasticity response of a metal depends critically on the presence or absence of dislocations before straining. Shown in Fig. 1

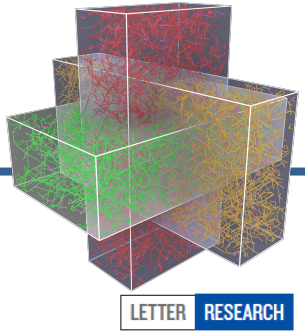

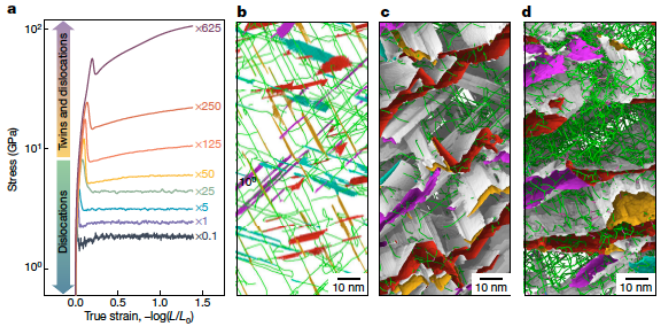

Figure 2 | Response to compression as a function of strain rate. a, Stress on a logarithmic scale as a function of true strain computed in molecular dynamics simulations of specimen compression at different strain rates (see colour-coded labels) and a temperature of 300 K. **b**, A snapshot of a simulation taken immediately after yield showing embryonic twins in the simulation performed at a rate of ×50. **c**, A later snapshot from the same simulation showing twin propagation. **d**, An even later snapshot in which the twins have grown to span the simulation volume. Colouring of dislocation lines and twin interfaces is the same as in Fig. 1.

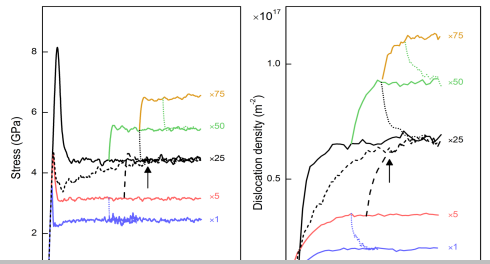

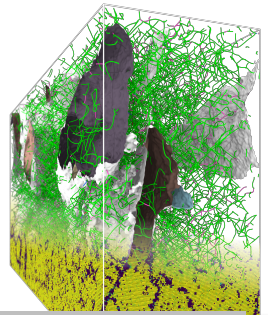

Dynamic equation of state (DEOS) defines material response to straining

Direct MD simulations of crystal plasticity

Solving two ages-old conundrums of physical metallurgy

3-stage hardening of metals

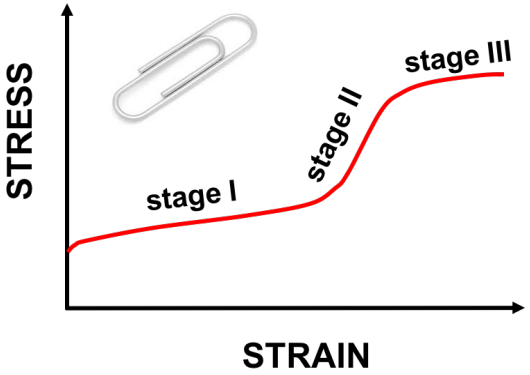

Work in progress

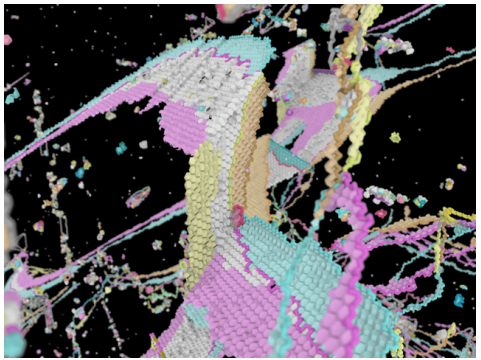

Dislocation patterns in crystals

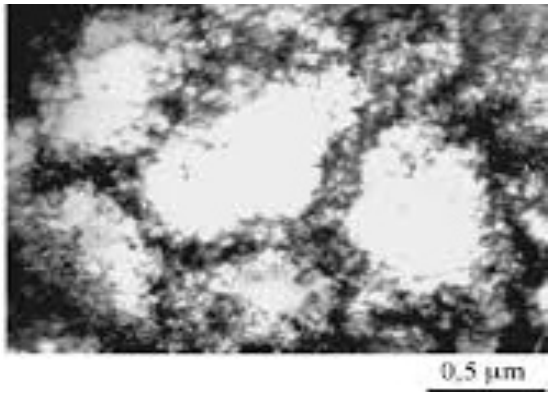

INCITE 2018-2019 Project: Crystal plasticity from first principles

INCITE 2018: Origin of 3-stage hardening

MD simulations

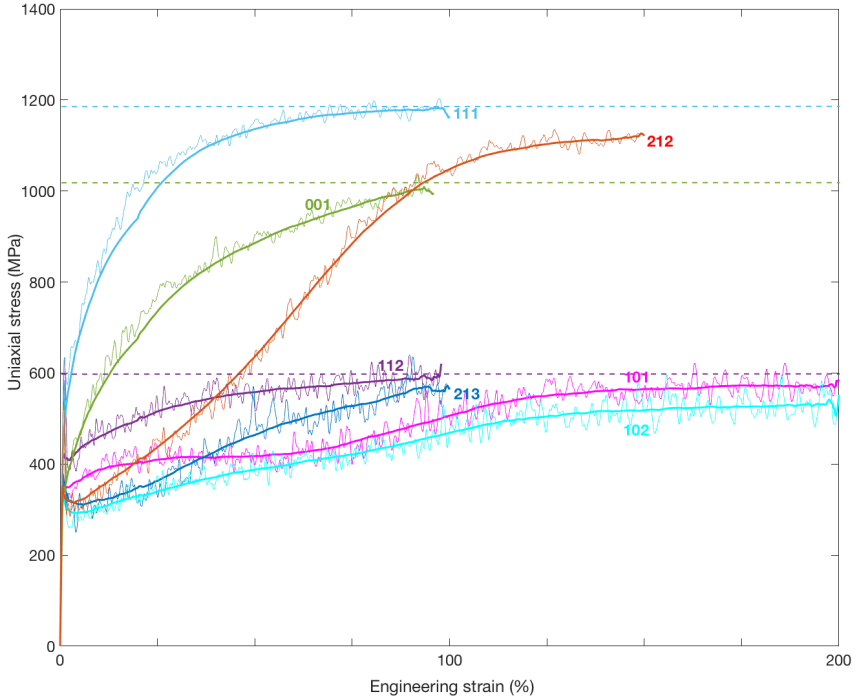

Experiments

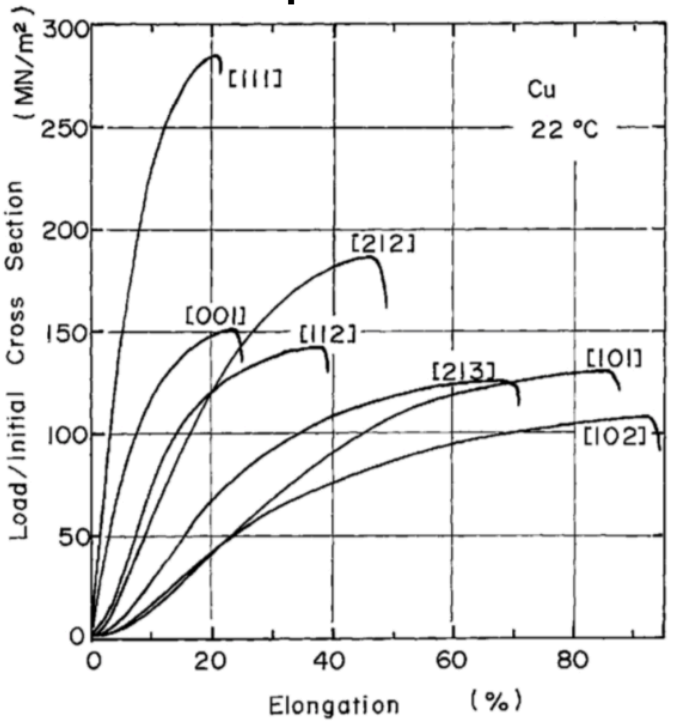

Summary: nature of 3-stage hardening

Initial orientation	Initial symmetry	Rotation observed	3-stage hardening
001	8-fold	no	no
111	6-fold	no	no
112	2-fold	no	no
102	2-fold, breaks	yes	yes
212	2-fold, holds	yes	yes
213	No symmetry	yes	yes
8,5,13	No symmetry	yes	yes
101	2-fold, breaks	yes	yes

- 3-stage hardening is caused by crystal rotation during straining
- Depending on the initial crystal orientation, straining axis asymptotically rotates to one of three stable ”attractor” orientations: 001, 111 or 112.

MD simulations reveal details not accessible in experiments

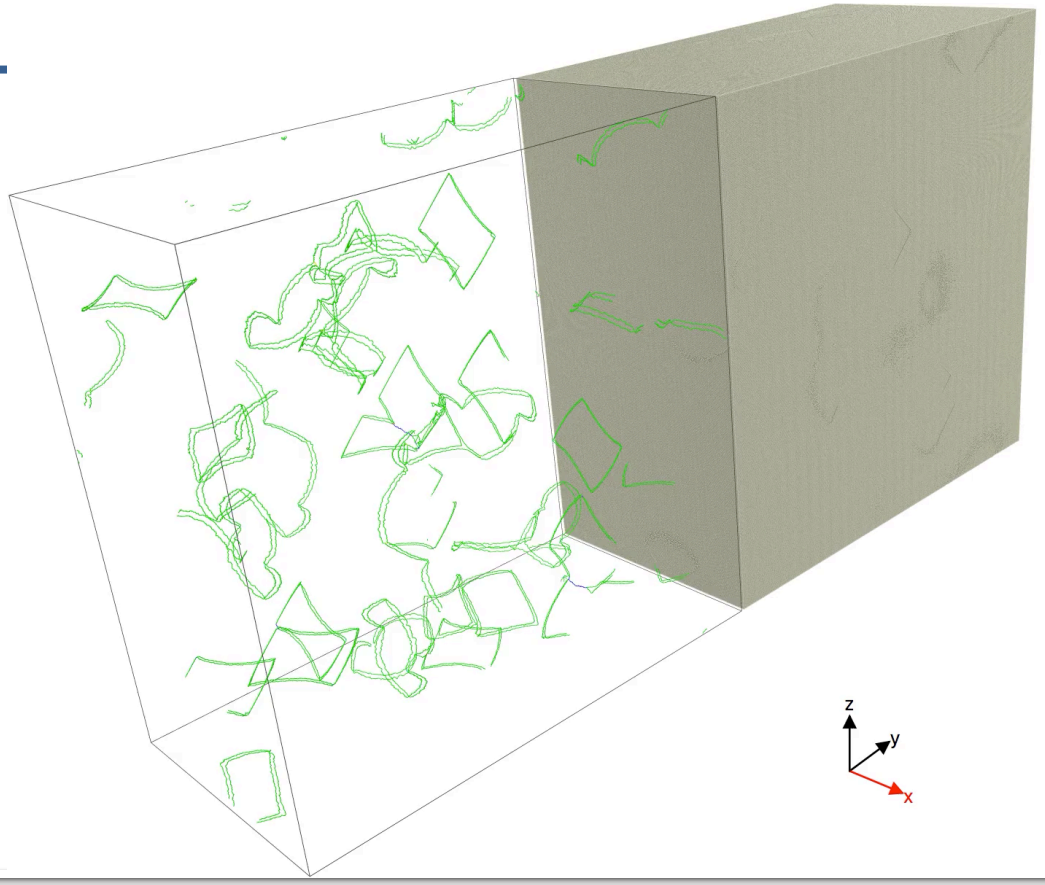

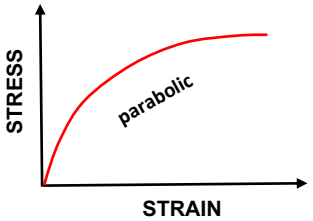

MD simulations reveal details not accessible in experiments

In situ in-bulk microscopy

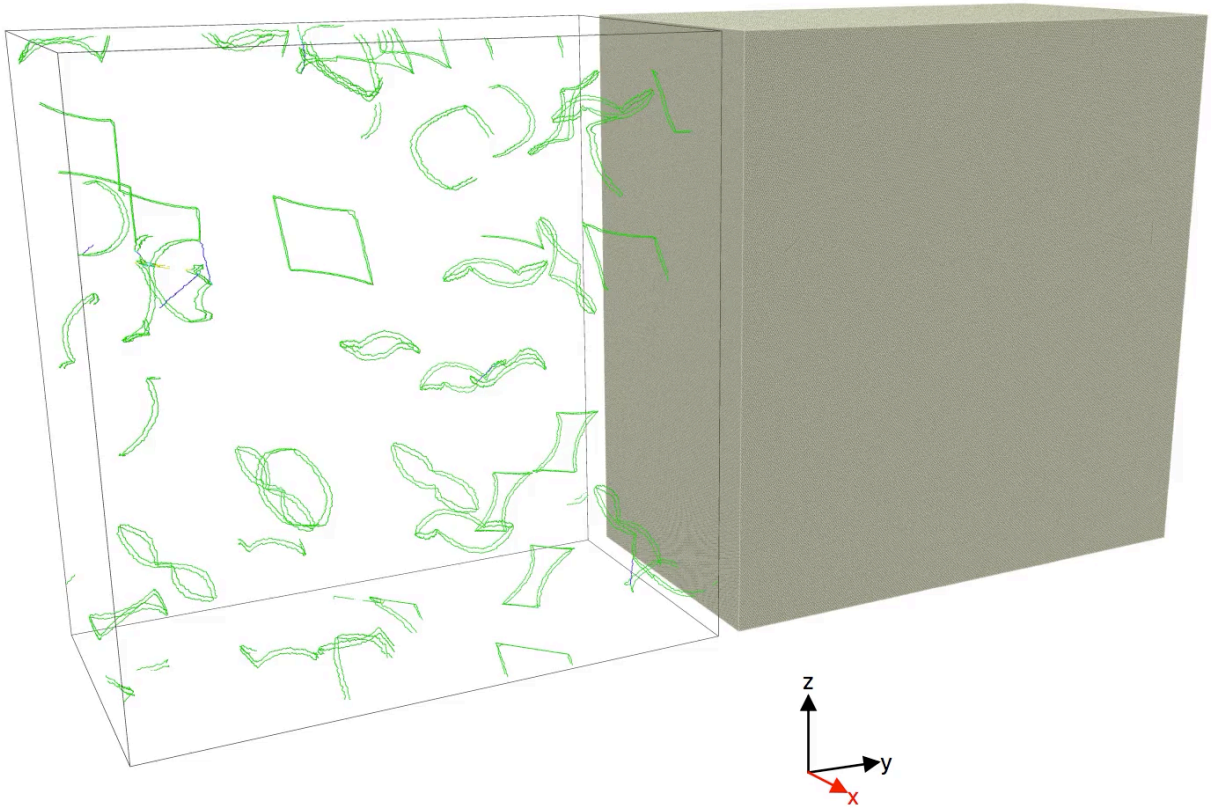

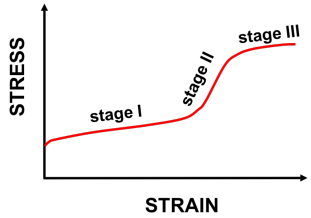

Origins of 3-stage hardening have been debated for over 60 years

Allan Cottrell: “Strain hardening is perhaps the most difficult remaining problem in classical physics. Harder than turbulence.”

Has been key aspiration of dislocation theory, thousands of papers published.

Direct MD simulations close the debate: it is all about crystal rotation.

Our most important observation so far

Physics of crystal plasticity scales: the mechanisms appear to remain the same over a vast range of straining rates from quasi-static experiments (10^{-5}/s) to MD (10^{4} – 10^{9}/s).

Livermore Big⊗Big simulation (LBB)

Developed at Sandia NL

LAMMPS on Sequoia

Massive BG/Q machine at LLNL

MD simulation with 2^{31} = 2,147,483,648 atoms

Simulated time = 5 μs ($5 \cdot 10^{-6}$ seconds)

Simulation size = 10,000 atoms·seconds (~100 times greater than any other MD simulation)

Produced $8 \cdot 10^{19}$ bytes = 80 exabytes of recordable trajectory data

~ 6 times (estimated) Google's worldwide storage capacity

LBB simulation as a benchmark for DDD

Atomistic input:

Elastic constants and dislocation mobility functions match the MD model at P=0 and T=300K

Simulation geometry:

Same size and shape of simulation volume as in MD

Initial configuration:

Same 32 hexagon-shaped prismatic loops placed in the same positions in the volume

Straining conditions:

Same 001 compression under the same straining rate 10^5 1/s

DDD model maximally matched to the MD model of tantalum

DDD vs MD comparison

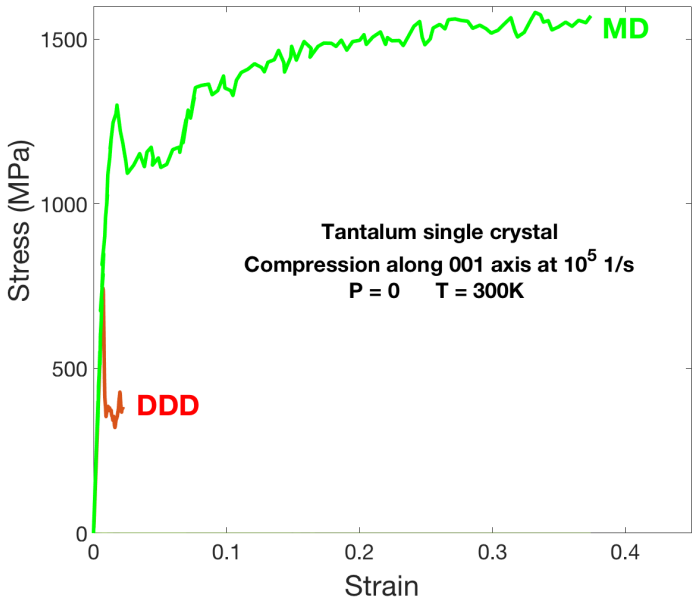

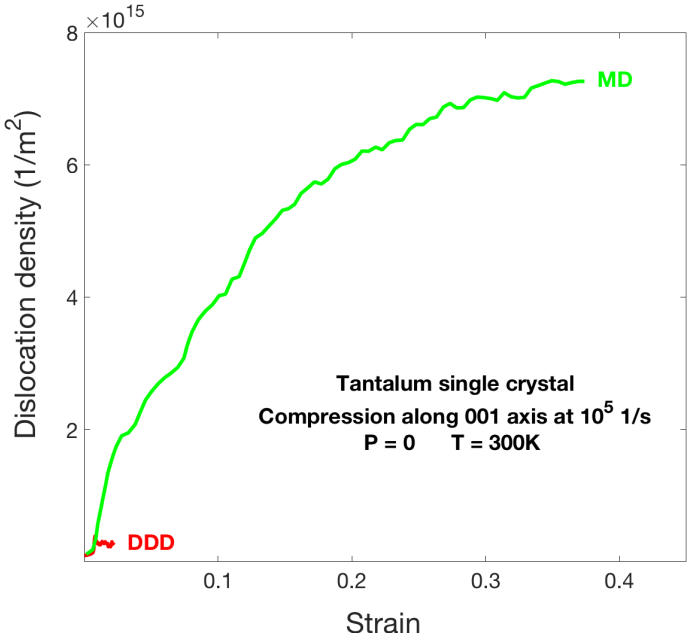

Strength is under-predicted by factor 3-4

Dislocation density is under-predicted by factor 15-20

Fidelity of MD → DDD workflow is uncertain

Too much ”lost in translation” ?

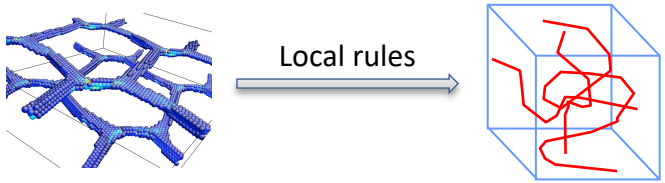

Challenge to DDD practitioners

Develop a DDD model to match LBB simulation predictions

Back to MD: tomorrow is now

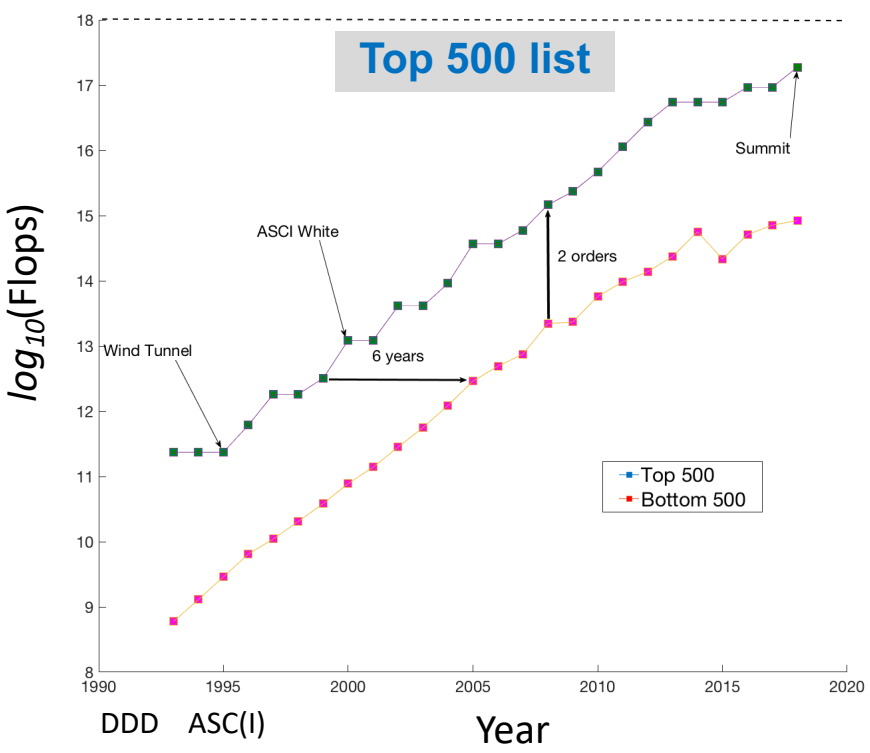

Future is bright for MD simulations

- In 6 years (on average) the Top500 top machine moves to the bottom
- Sequoia is already only 13th fastest
- 1000 machines faster than Sequoia expected in 4 years

Tomorrow is now

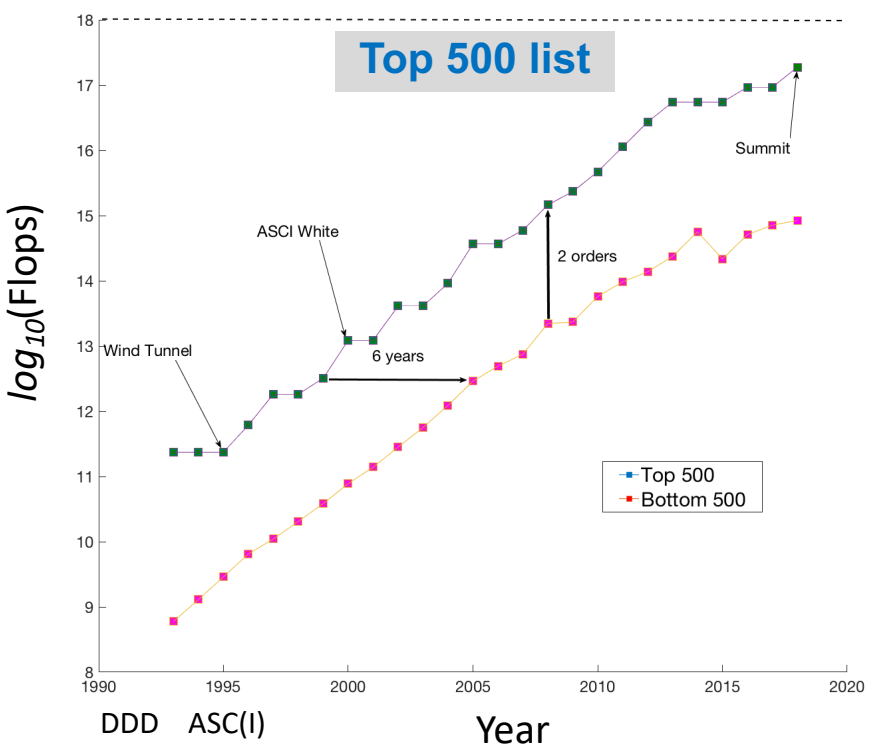

It takes ~ 6 years on average for the World's most powerful computer to move to the bottom of top500 list (Today all top 500 > 1 Pflops)

Cost of computing per Gflop

$165,000,000,000 in 1963
$400,000 in 1993
$0.01 in 2018

Power efficiency: 17 Gflops/Watt

Tomorrow is now

- The world's top machine Summit (ORNL) has 10 Pb of total memory:
  sufficient for LAMMPS simulations with ~10^{13} atoms or 10 x 10 x 10 μ^3.

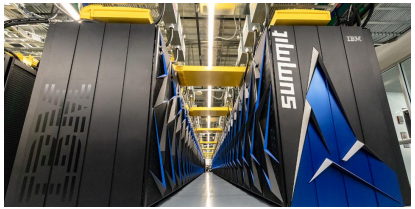

- Three major HPC hardware components:

  Integral processing speed (growing fastest)
  Memory (slower growth)
  Inter-node communication rate (slowest growth)

- Current limit is ~ 1-10 μs of simulated MD trajectory per compute day.

  Inter-node communication limit

- Methods of accelerated MD may speed up simulations by orders of magnitude.

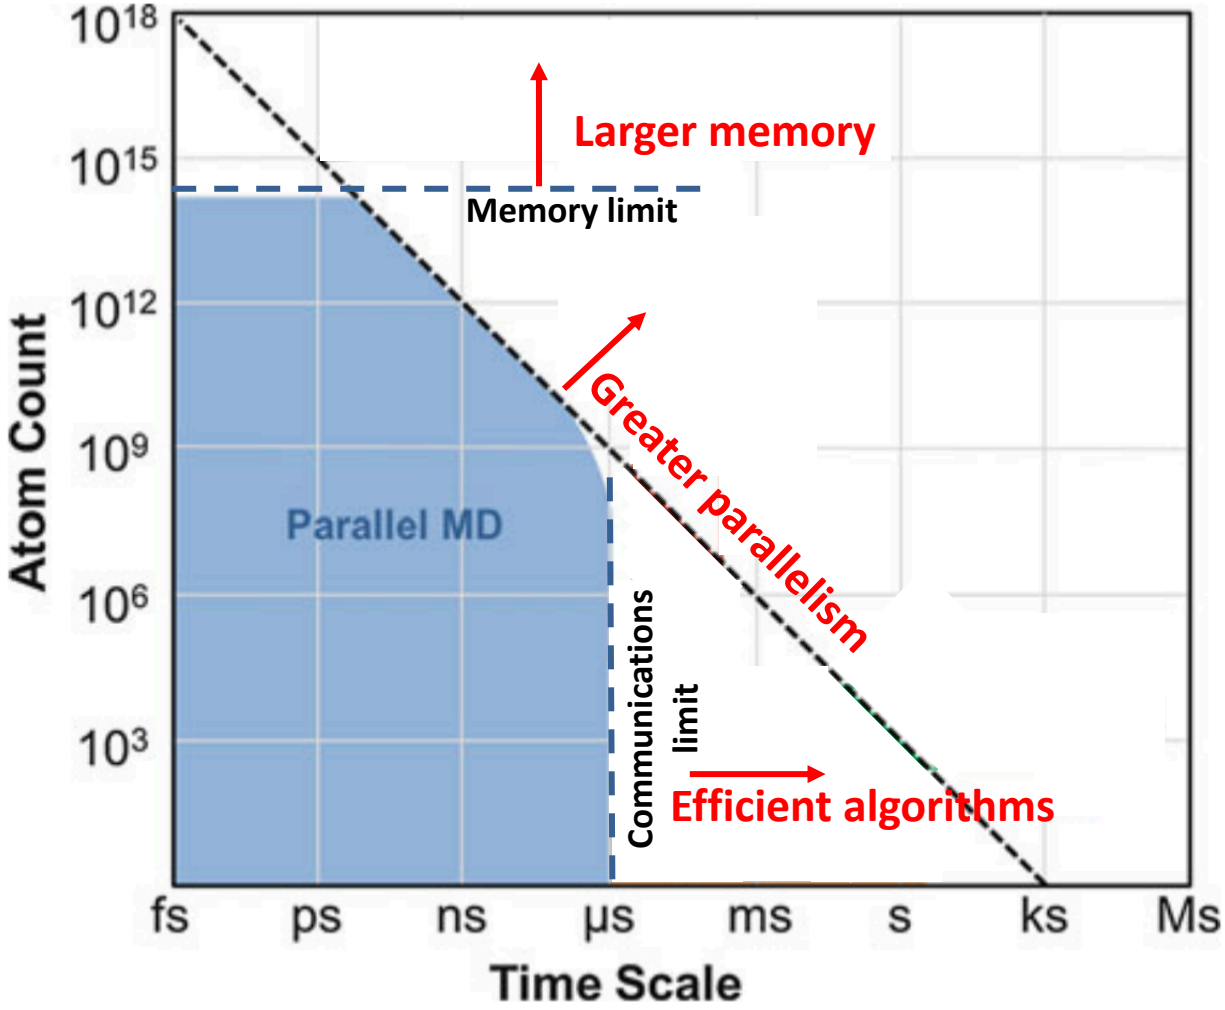
10^{18}
10^{15}
10^{12}
10^{9}
10^{6}
10^{3}
Atom Count
Larger memory
Memory limit
Greater parallelism
Parallel MD
Communications limit
Efficient algorithms
fs
ps
ns
μs
ms
s
ks
Ms
Time Scale

Cross-scale MD simulations

Sufficiently large to be statistically representative of the simulated model system and yet resolving every tiniest detail of atomic motion.

Key ingredients for success

1. Accurate, transferrable and computationally efficient interatomic potentials.

2. Increasingly large and long MD simulations.

3. Data management: on the fly analyses, reduction, compression, knowledge acquisition.

Cross-scale MD simulations

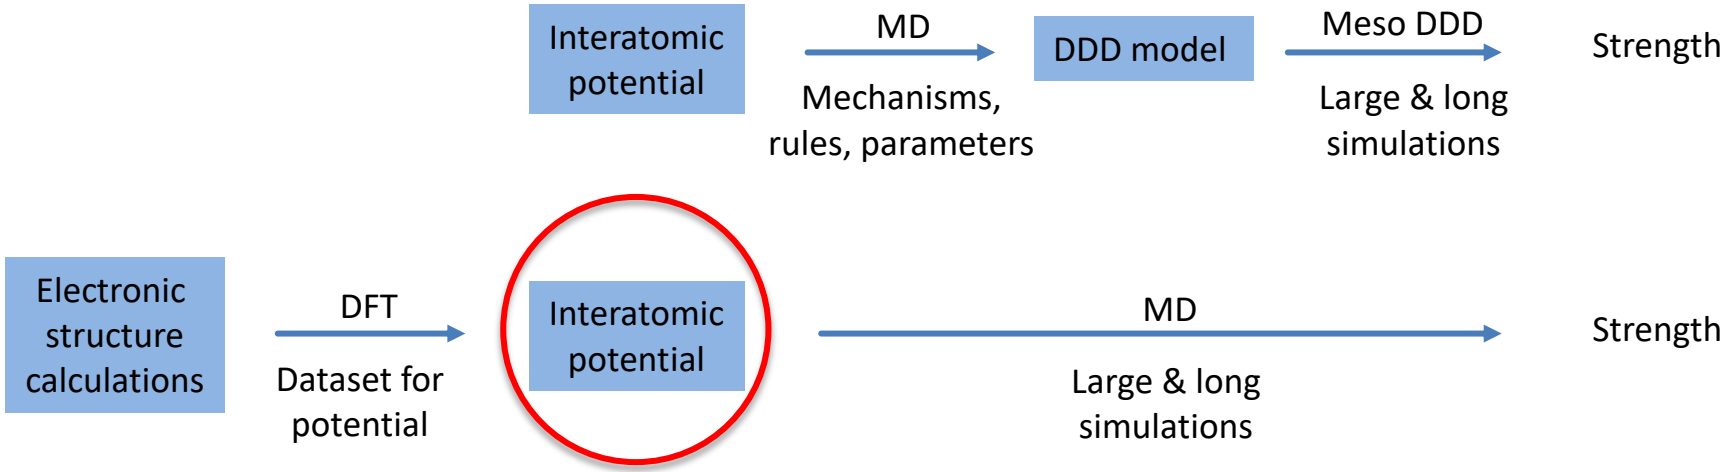

Interatomic potentials *à la carte*

Fidelity	Ability to accurately reproduce or predict a material property of interest.
Efficiency	Inverse of computational cost of force evaluations.
Transferability	Fidelity of predictions over a wide range of properties and conditions.

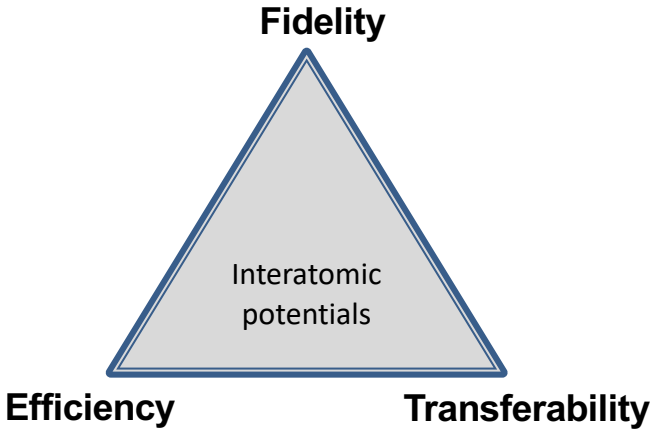

Automated on-demand development of interatomic potentials with optimal tradeoff between fidelity, efficiency and transferability

Increasingly large MD simulations

HPC capabilities continue to grow unabated

New efficient algorithms for massively parallel simulations

- Parallel-in-time integrators
- Asynchronous simulators
- Accelerated MD, parallel boost (already in LAMMPS)
- …

Simulations data management

A major challenge for cross-scale MD simulations

Exascale data challenge

One day of direct LAMMPS simulation of metal strength on Sequoia produces
$4 \cdot 10^{18}$ bytes = 4000 petabytes = 4 exabytes
of MD trajectory data

Comparable to Google's worldwide storage capacity

Data reduction is paramount to extract knowledge

In situ computational microscopy

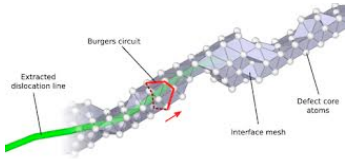

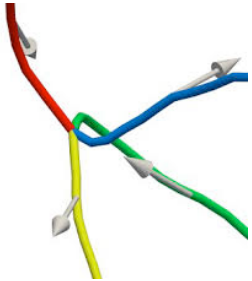

- Automated identification and precise indexing of crystal defects: dislocations, twins, grain boundaries, surfaces, point defects, etc.
- Simulation data reduction and analyses
- High-performance visualization tools

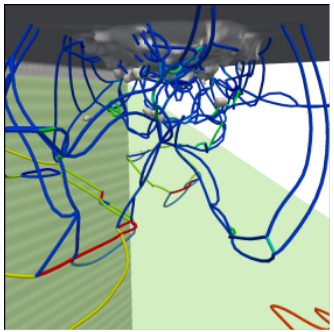

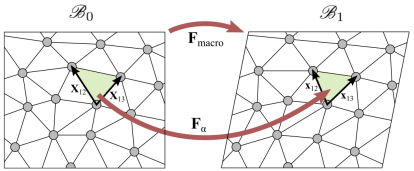

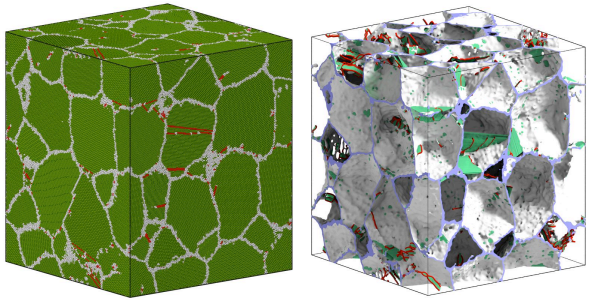

Alexander Stukowski
Technische Universität Darmstadt

Material dynamics in arbitrary detail

Alex Stukowski

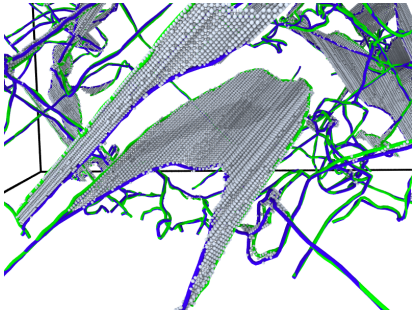

Edge dislocation motion

On-the-fly sweep-trace analysis (blur)

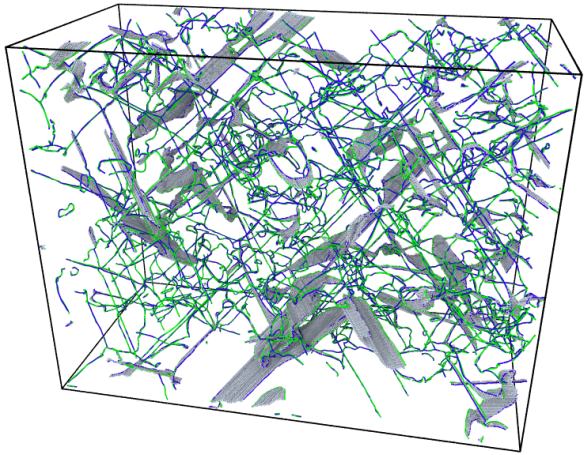

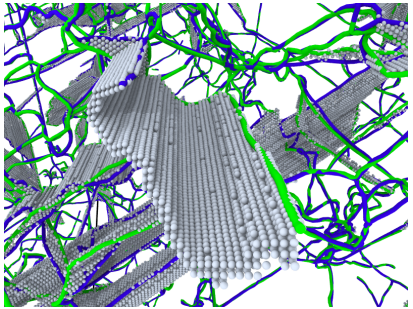

Screw dislocation motion

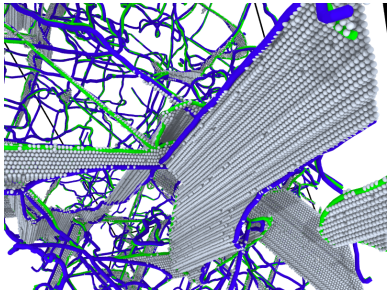

Dislocation annihilation

Data reduction ratio ~ 10^{-5}

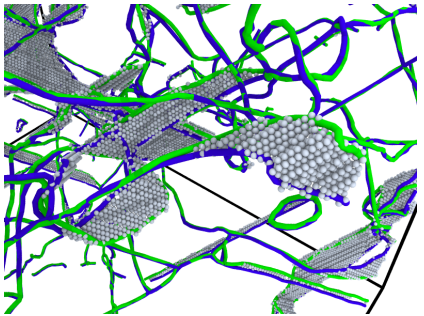

Dislocation intersection

Data reduction does not come cheap

Stukowski's classic DXA algorithm works in two stages

1. run fix *disloc/orig* in LAMMPS to build tracing mesh (runs in parallel)

2. run *trace_tool* to trace and stich the dislocation network (runs on single node only)

On a 24 billion atoms simulation disloc/orig takes minutes per snapshot, but trace_tool postprocessing takes hours and requires a lot of on-node memory.

56 billion atoms simulations in the pipeline, present a challenge to process.

Stukowski is hard at work developing scalable parallel on-the-fly DXA

Reduced data is still huge

Our ongoing 56 billion atoms simulations of Ta are producing dislocation networks containing 12M network nodes and about 20M dislocations in a single frame.

Even after 10^{-5} reduction such data is still too much for a human to comprehend.

To engage human vision and to feed intuition this data should be further reduced.

Multi-scale data reduction workflows are needed

0.99999 of MD trajectory data is lost forever

Can we machine learn on exabytes of trajectory data?

Learning itself is more expensive than MD simulation (per frame)

Much of MD trajectory data is redundant

Need near-zero-cost on-the-fly redundancy filtering

Data diversity pre-selection for subsequent ML

Information theory

Summary

Direct (cross-scale) MD simulations of crystal plasticity present an inviting alternative to multi-scaling.

Results of large-scale MD simulations can and should be used to improve fidelity of DDD and other multi-scale methods.

In our simulations so far, we observe that basic mechanisms of crystal plasticity appear to be the same over a vast range of straining rates from quasi-static (10^{-5}/s) to MD (10^{4} – 10^{9}/s).

Key challenge for cross-scale simulations: what to do with increasingly huge streams of MD data?