A growing challenge: what are we going to do with data generated in exascale simulations?

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

A. Stukowski	A. Metere	T. Oppelstrup	N. Bertin	T. Hsu	R. Freitas
					
Ovito GmbH	Samsung Semiconductor, Inc.	LLNL	LLNL	LLNL	MIT

Support NNSA ASC Program, Technische Universität Darmstadt
Livermore Computing Grand Challenge Program (2015 - pres)
DOE INCITE Program (2018-2019)
NNSA ATCC Program (2017 – pres)
Jülich Supercomputing Centre (2016)

My MD storyline

In the 80’s	MD simulations of materials
90 - 2015	Multiscale simulations of materials
2015 - present	Back to MD simulations

NNSA Advanced Simulation and Computing

Predicting, with confidence, the behavior of nuclear weapons through comprehensive, science-based simulations.

Computational Systems & Software
Integrated Codes
Physics and Engineering Models
…

Materials aging
Equation of State
Material Strength and Damage
…

Predictive multiscale simulations of metal strength and plasticity

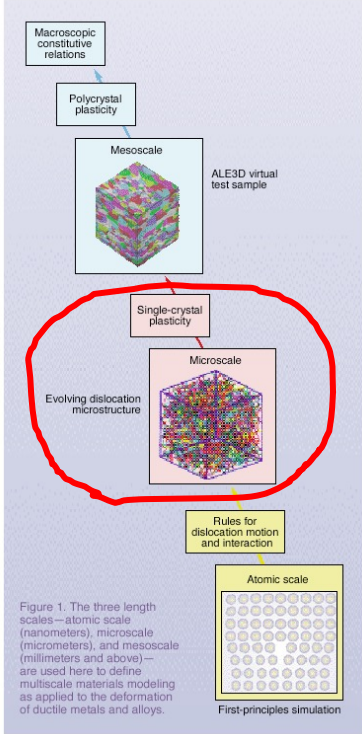

Figure 1. The three length scales—atomic scale (nanometers), microscale (micrometers), and mesoscale (millimeters and above)—are used here to define multiscale materials modeling as applied to the deformation of ductile metals and alloys.

Bottom-up modeling hierarchy
Compute and pass parameters from the lower to the higher scales to bridge length- and time-scale gaps

Collective dynamics of dislocations at micron scales
The weakest, most uncertain and highest potential pay-off link

In the past, connection of continuum engineering models to underlying physics of dislocation motion has been only inspirational/motivational.

Grand challenge: predict crystal strength from dislocation physics

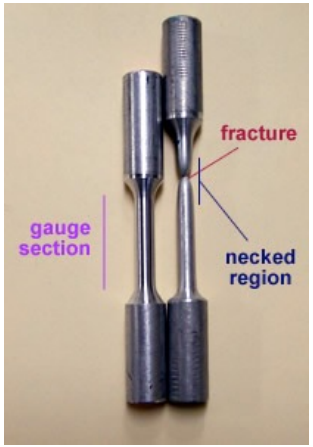

Discrete Dislocation Dynamics

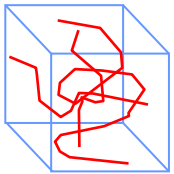

G. Canova and L. P. Kubin (1991)

Premise and promise of the DDD method

Dislocation theory, MD

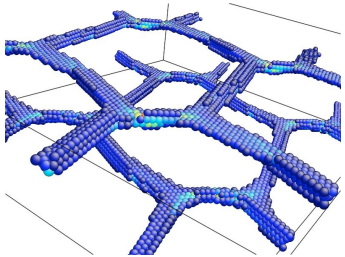

Dislocation mobility and interactions

Local rules

DDD

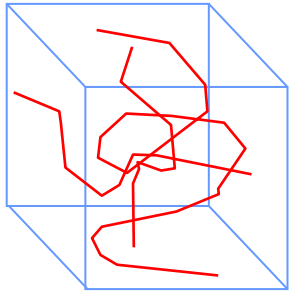

Collective response of statistically representative dislocation ensembles

Strain hardening

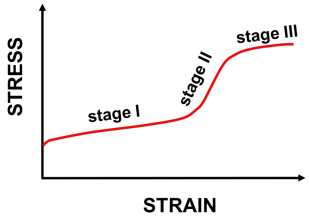

Patterns

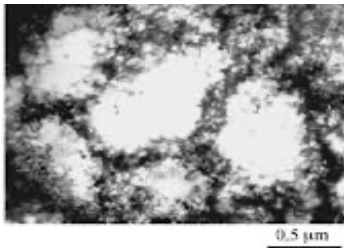

...

Lots of fun and hard work in DDD method development, but coveted connection has remined 10-15 years away!

Meanwhile

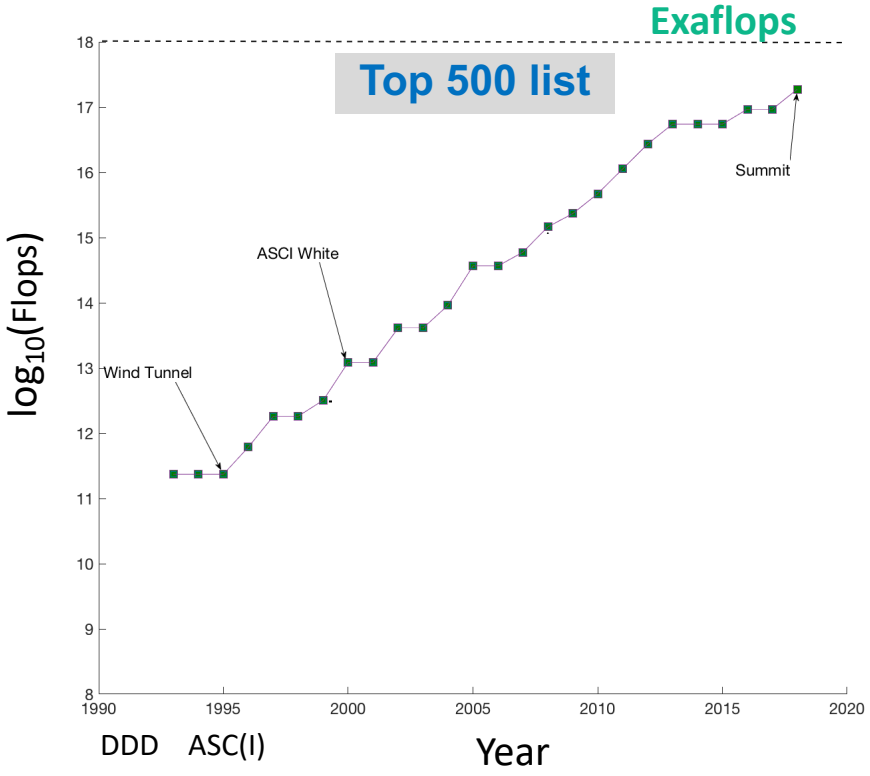

Since DDD launch in 1991 to 2019

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

Cost of computing per Gflop

1963: $165,000,000,000
1993: $400,000
2023: $0.000000003

Back to good old Molecular Dynamics?

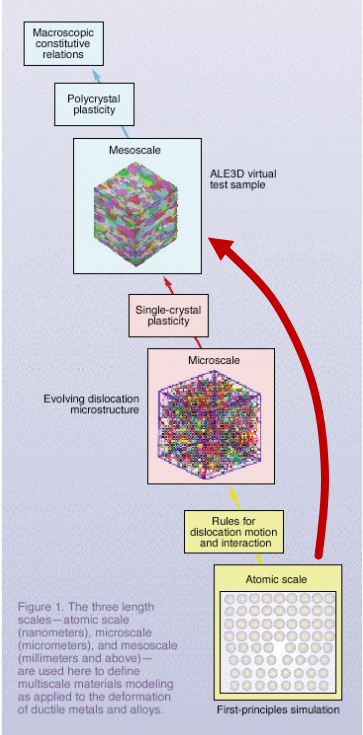

Figure 1. The three length scales—atomic scale (nanometers), microscale (micrometers), and mesoscale (millimeters and above)—are used here to define multiscale materials modeling as applied to the deformation of ductile metals and alloys.

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

Dislocation Extraction Algorithm (DXA)

Video by Alex Stukowski

Direct MD simulations of crystal plasticity

6M – 56B atoms, BCC lattice, 3D periodic

Interatomic potential models for Ta, W, Al, Cu, …

Dislocations sourced for multiplication

Straining along one of the periodic box axes

Constant true rate of straining

Constant temperature

Relaxation of lateral stress (Poisson effect)

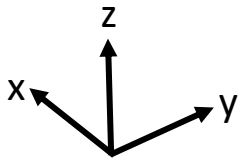

MD simulations of such magnitude were previously thought to be unthinkable

A strong metal flows like anisotropic viscous fluid

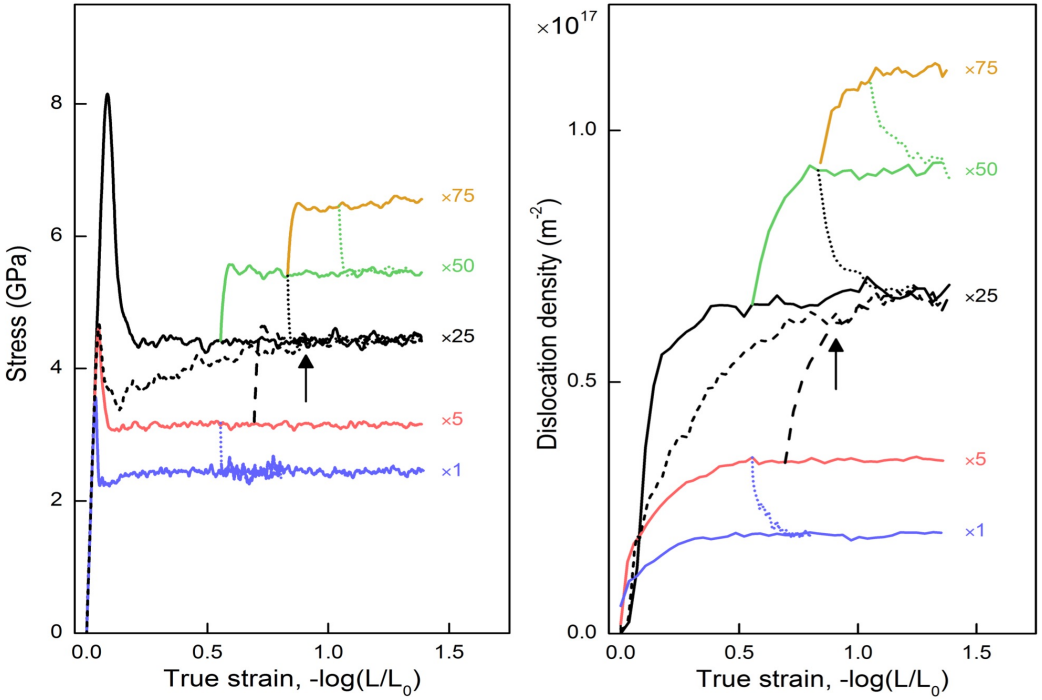

Path-independent plastic flow

All straining trajectories converge to the same ultimate flow stress (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

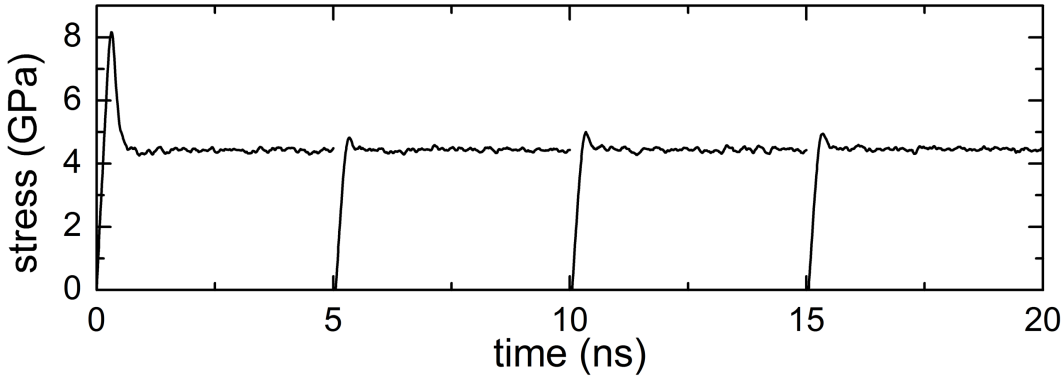

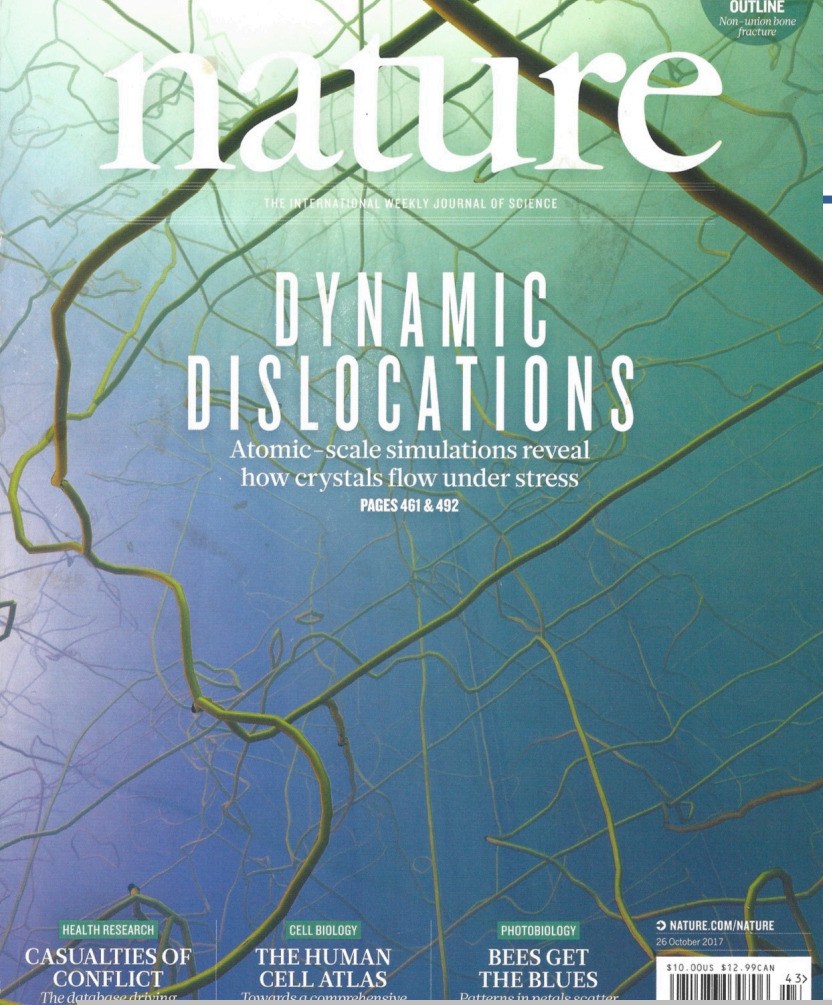

LETTER

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

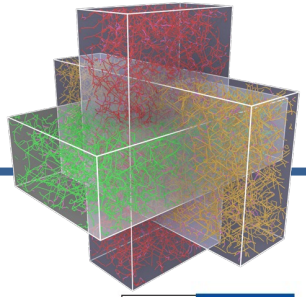

LETTER RESEARCH

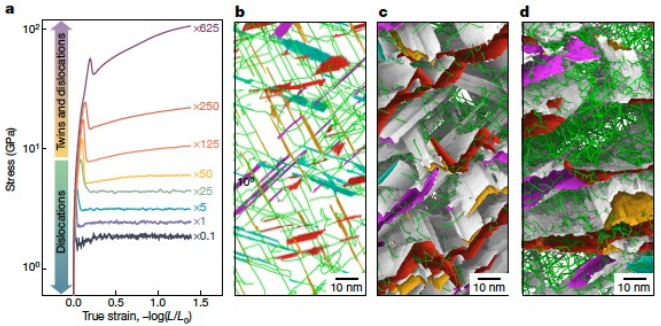

Figure 2 | Response to compression as a function of strain rate. a, Stress on a logarithmic scale as a function of true strain computed in molecular dynamics simulations of specimen compression at different strain rates (see colour-coded labels) and a temperature of 300 K. **b**, A snapshot of a simulation 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.

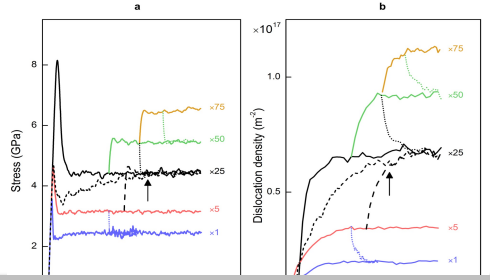

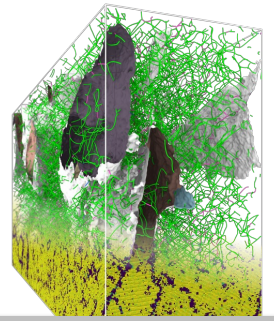

DEOS: under stationary straining, a metal attains a state of steady asymptotic plasticity

Direct MD simulations of crystal plasticity

Solving two ages-old conundrums of physical metallurgy

Staged hardening of metals

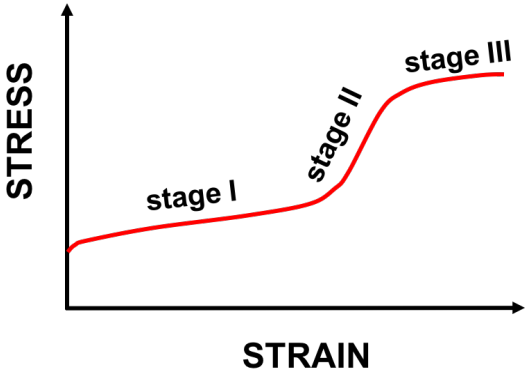

Work in progress

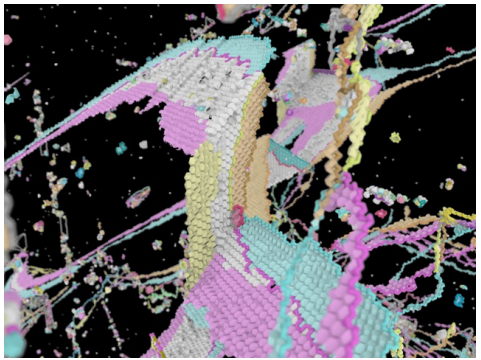

Dislocation patterns in crystals

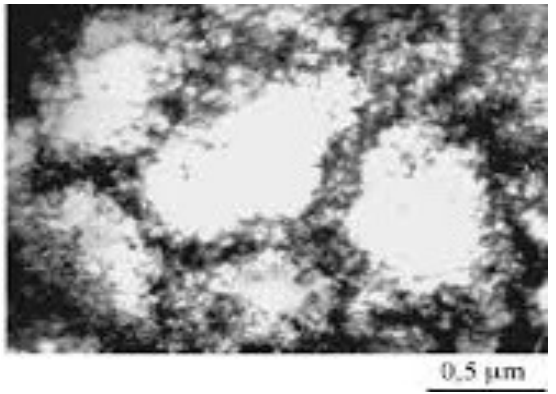

Origin of staged hardening

Experiment

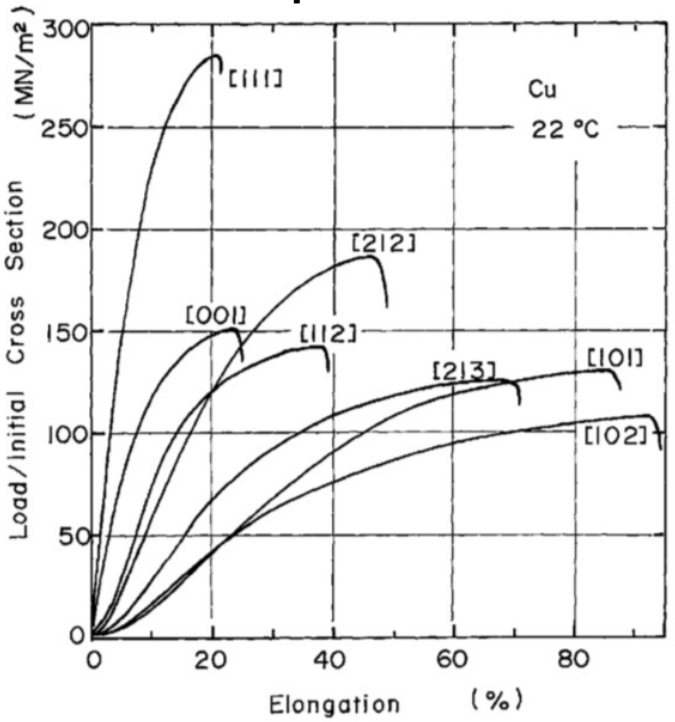

MD simulations

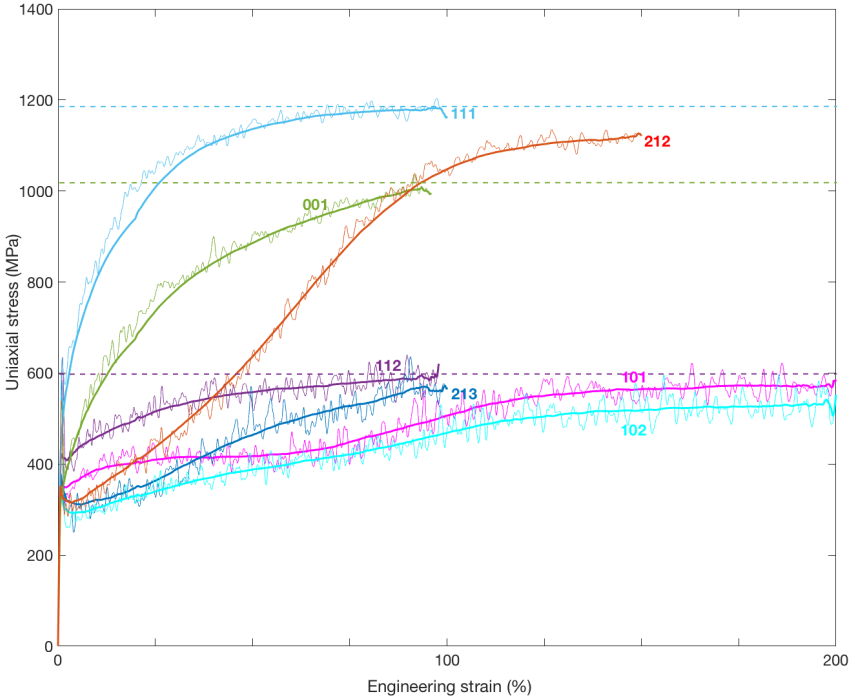

Summary: origin of 3-stage hardening

Initial orientation	Initial symmetry	Rotation observed?	Staged 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

Staged hardening is caused by crystal rotation during straining

MD simulations reveal details not accessible in experiments

In situ in-bulk microscopy

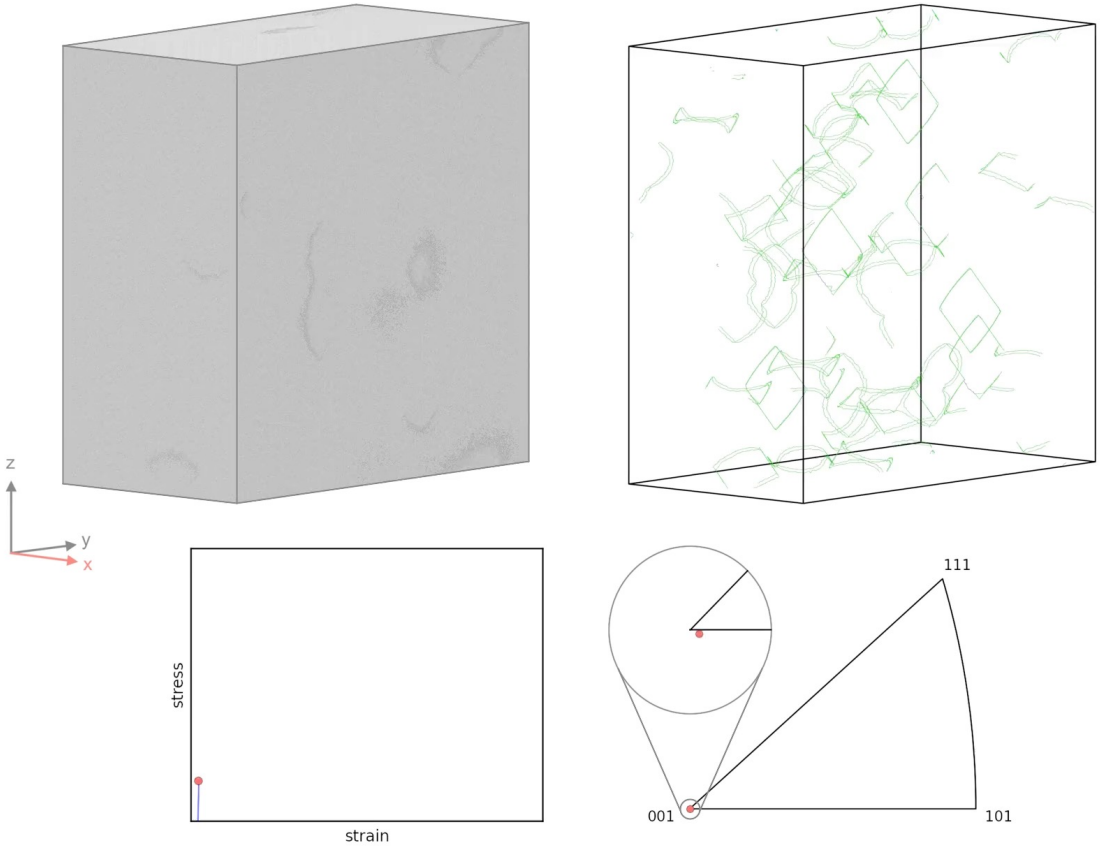

Straining along stable [001] axis

MD simulations reveal details not accessible in experiments

In silico in-bulk microscopy

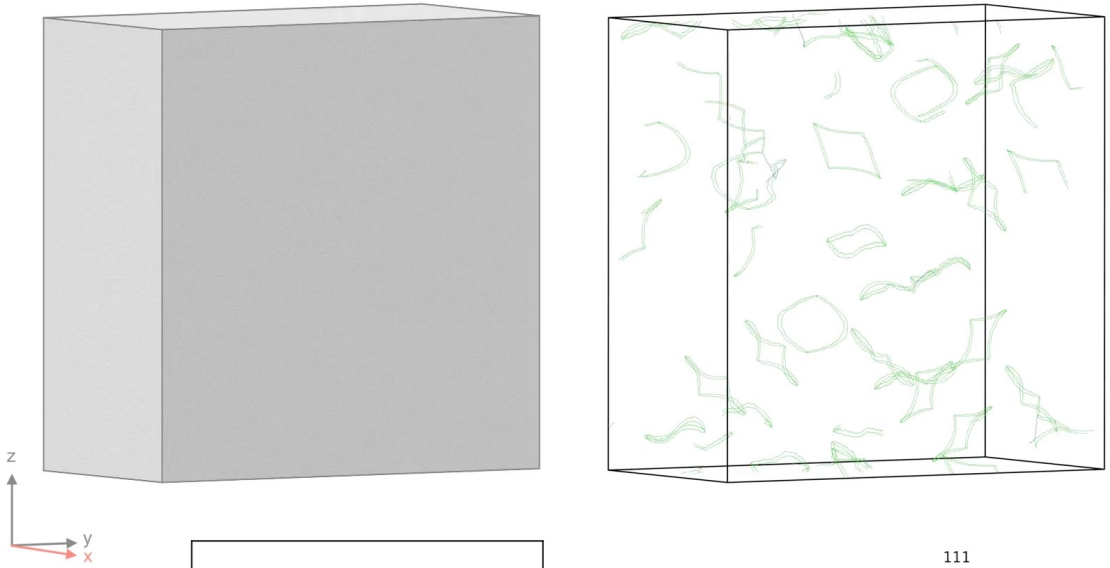

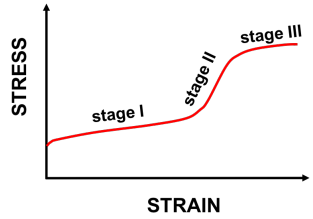

Straining along unstable [213] axis

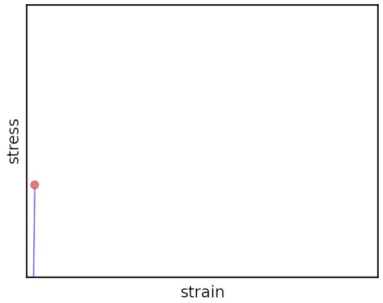

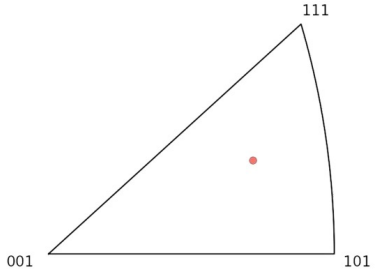

Origins of 3-stage hardening have been debated for over 60 years

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

Explanation of staged hardening has been a key aspiration for dislocation theorists, thousands of papers published.

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

nature materials

ARTICLES

https://doi.org/10.1038/s41563-020-00815-1

Check for updates

Atomistic insights into metal hardening

Luis A. Zepeda-Ruiz[1], Alexander Stukowski[2], Tomas Oppelstrup[1], Nicolas Bertin[1], Nathan R. Barton[1], Rodrigo Freitas[3,4] and Vasily V. Bulatov[1]✉

For millennia, humans have exploited the natural property of metals to get stronger or harden when mechanically deformed. Ultimately rooted in the motion of dislocations, mechanisms of metal hardening have remained in the cross-hairs of physical metallurgists for over a century. Here, we performed atomistic simulations at the limits of supercomputing that are sufficiently large to be statistically representative of macroscopic crystal plasticity yet fully resolved to examine the origins of metal hardening at its most fundamental level of atomic motion. We demonstrate that the notorious staged (inflection) hardening of met-

MD has become the workhorse method for predictive simulations of metal strength at LLNL. Is it worth to continue to develop the mesoscale DDD method?

Streamlined workflow for metal strength predictions

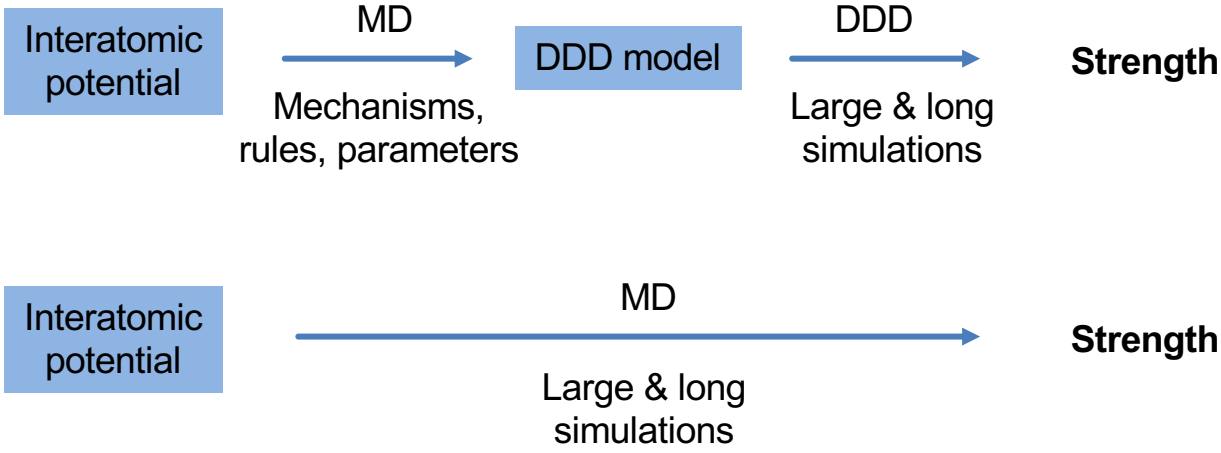

Benefits

- MD is a mature method, does not require further development.
- Reduces prediction uncertainties to interatomic potential.
- Shortens time from start to delivery, reduces overall cost.
- Unlimited scalability to exascale computing and beyond.

X-scale matching: learning from side-by-side comparison

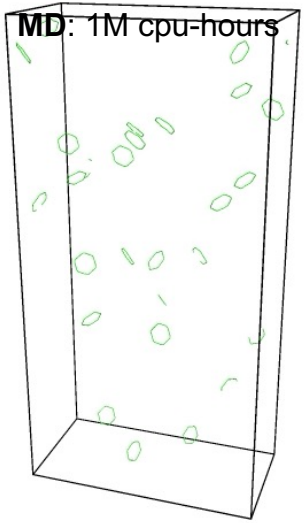

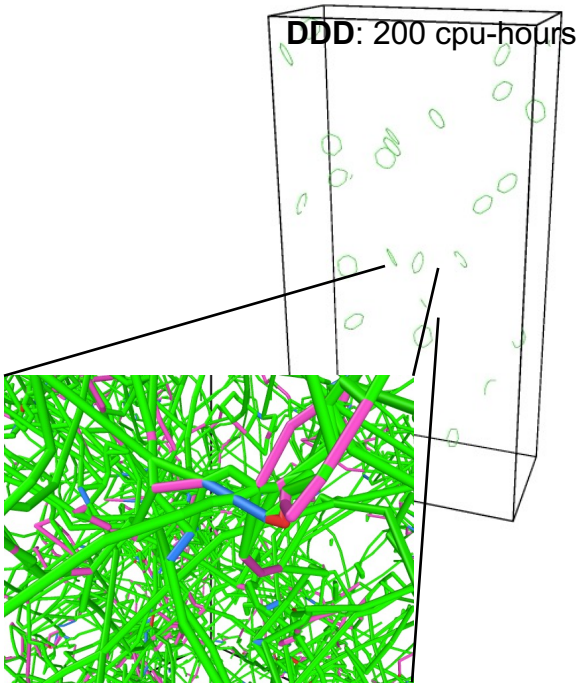

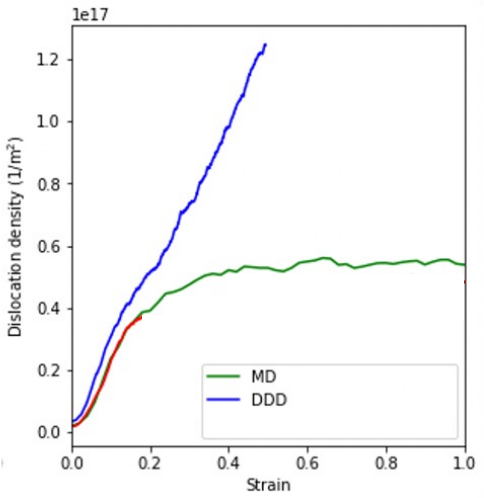

In DDD dislocation network nodes are “sticky” holding the tangles together, but not in MD

In MD dislocation network nodes are not sticky and move in 3D

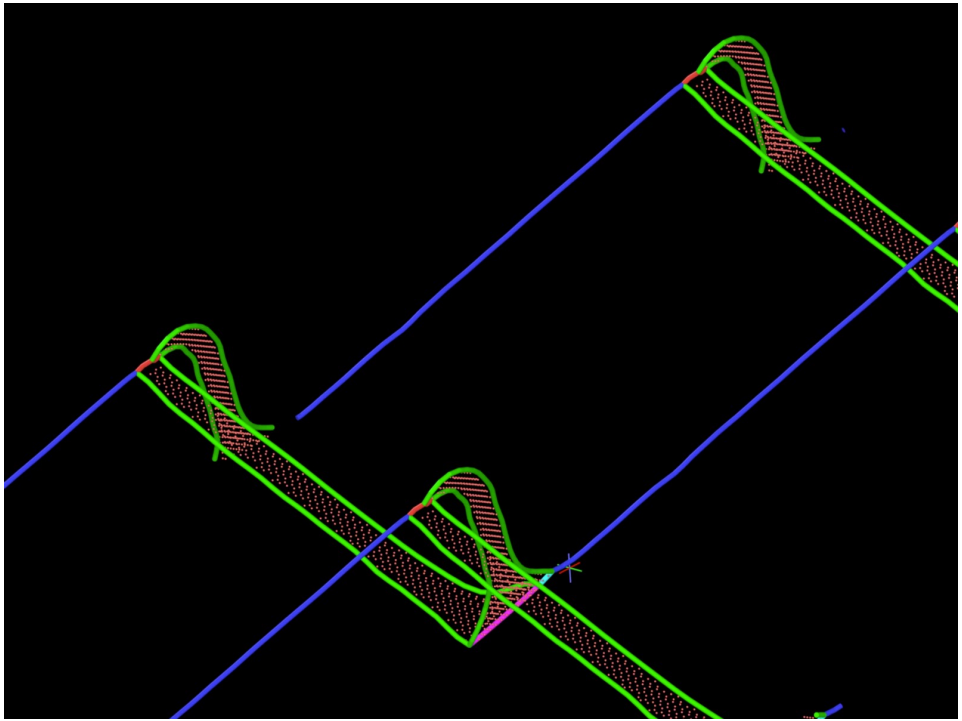

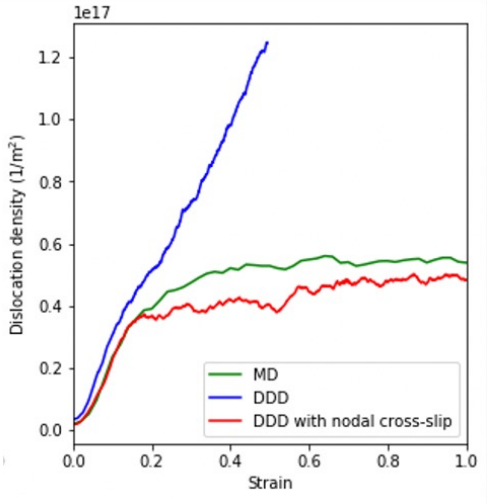

X-scale matching improves fidelity of DDD simulations

X-scale MD simulations

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

Ingredients for success

Accurate, transferrable and computationally efficient interatomic potentials.

Increasingly large and long MD simulations.

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

What is MD good for?

An engine to compute a number?

A fully resolved computational experiment to learn from?

68 years of Molecular Dynamics

Trajectory lines of atoms in a melting crystal from the very first MD simulation (64 atoms)

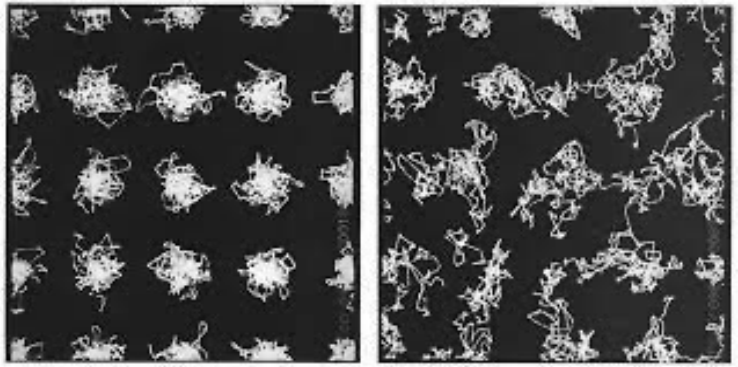

S. Campbell, B. Alder, T. Wainwright
LLNL (1956)

1957-1967: Seminal papers on melting (64 atoms) and “long-time tail” hydrodynamic interactions (108 atoms).

Livermore Big⊗Big simulation (LBB)

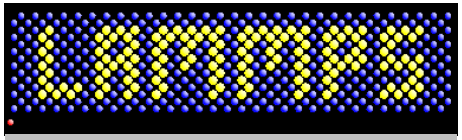

Developed at Sandia NL

LAMMPS on Sequoia

2017

20 days of full Sequoia worth of computing

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

Produced $9 \cdot 10^{19}$ bytes = 90 exabytes of recordable trajectory data

~ 6 googles of data (Google's worldwide storage capacity is ~ 15 exabytes)

Massively parallel machine at LLNL

Fate of LBB simulation data

5000 states saved at equal time intervals 10^5 time steps apart (~ 1 Pb of data)

Has taken about 2 months to transfer (htar) 1 Pb to tape storage

Has taken about 5 months to process 5000 snapshots to reveal dislocations and other defects

Hours spent watching the sequence/movie of defect configurations

Observed dynamics is highly intermittent (avalanche-like):

Nothing happens between most subsequent snapshots.

A whole lot happens between some snapshots, but dynamic details are irrevocably lost.

Few new insights gained using naïve approach to data management and processing

LBB simulation, if repeated on El Capitan (2023)

LBB would be completed in ~ 5 hrs of full El Cap

At its maximum efficiency parallel I/O can only write ~0.004 of concurrently generated MD trajectory

About 40% of the entire disk space will be filled

Would take some 60-80 years to transfer the partial trajectory to long-term storage and would exceed LLNL storage capacity (advertised as unlimited)

Would take a few hundred years to postprocess.

Even such an immense effort would not reveal relevant details of system dynamics

Growing disparity between data generation and data utilization rates

Is it hopeless?

Why do we want to run MD simulations on exa-scale and beyond?

The system needs to experience a statistically representative number of essential dynamic events

Exascale MD trajectory is grossly repetitive and redundant

How much of the LBB simulated trajectory was essential to retain?

90 Eb → < 1 Tb

~ 10^{-8} reduction

Essential data is of limited size but we don't know how to extract it from the trajectory

Desired specs on an exascale data management pipeline

- Intelligent data reduction: filter out redundant but retain essential data.
- Data reduction must be performed on the fly (or indiscriminately lost).
- Computational cost of data reduction should not exceed that of the simulation itself.
- It should be possible to process the reduced data into a form amenable to human analysis.

A workflow for reduction of exascale trajectory data is in the works

Summary

Direct MD simulations are reaching previously unthinkable scales and, where feasible, are superseding multiscale simulations.

X-scale matching: Fidelity of mesoscale simulation methods (DDD) can be improved by comparing to direct MD simulations performed on overlapping/mesoscopic scales.

Knowledge acquisition from MD simulation trajectories is becoming increasingly difficult. We are running into an “exa-scale data bottleneck”.