Highly-Scalable Discrete-Particle Simulations with Novel Coarse-Graining: Accessing the Microscale

Timothy I. Mattox (Engility Corporation, DoD HPCMP PETTT)

2017 LAMMPS Workshop and Symposium

August 2, 2017 at UNM in Albuquerque, NM

DOD HPC MODERNIZATION PROGRAM

ENGILITY

Co-Authors & Acknowledgements

U.S. Army Research Laboratory

James P. Larentzos†

John K. Brennan†

Brian Barnes

Computational Science and Engineering

Christopher P. Stone†

Engility Corporation

Timothy I. Mattox†

Ross Smith

Sandia National Laboratories

Stan G. Moore†

Daniel A. Ibanez†

Aidan P. Thompson†

Steven J. Plimpton†

Christian Trott

Czech Academy of Sciences

J.E. Purkinje University

Martin Lísal†

†**Co-Authors on our Gordon Bell Prize submission**

DoD Supercomputing Resource Centers (DSRCs):

Army Research Laboratory (ARL)

Air Force Research Laboratory (AFRL)

Army Engineer Research and Development Center (ERDC)

Los Alamos National Laboratory (LANL)

National Nuclear Security Administration (NNSA)

Texas Advanced Computing Center (TACC)

High Performance Computing Modernization Program (HPCMP)

This study was supported by the U.S. Dept. of Defense High Performance Computing Modernization Program (HPCMP) User Productivity Enhancement, Technology Transfer, and Training (PETTT) activity (GSA Contract No. GS04T09DBC0017 through Engility Corporation)

Motivation

Microstructure dependent shock response of Energetic Materials (EM)

- **EM composites are highly microstructured**
- **Sensitivity and performance believed to be due to localized "hot spots"**

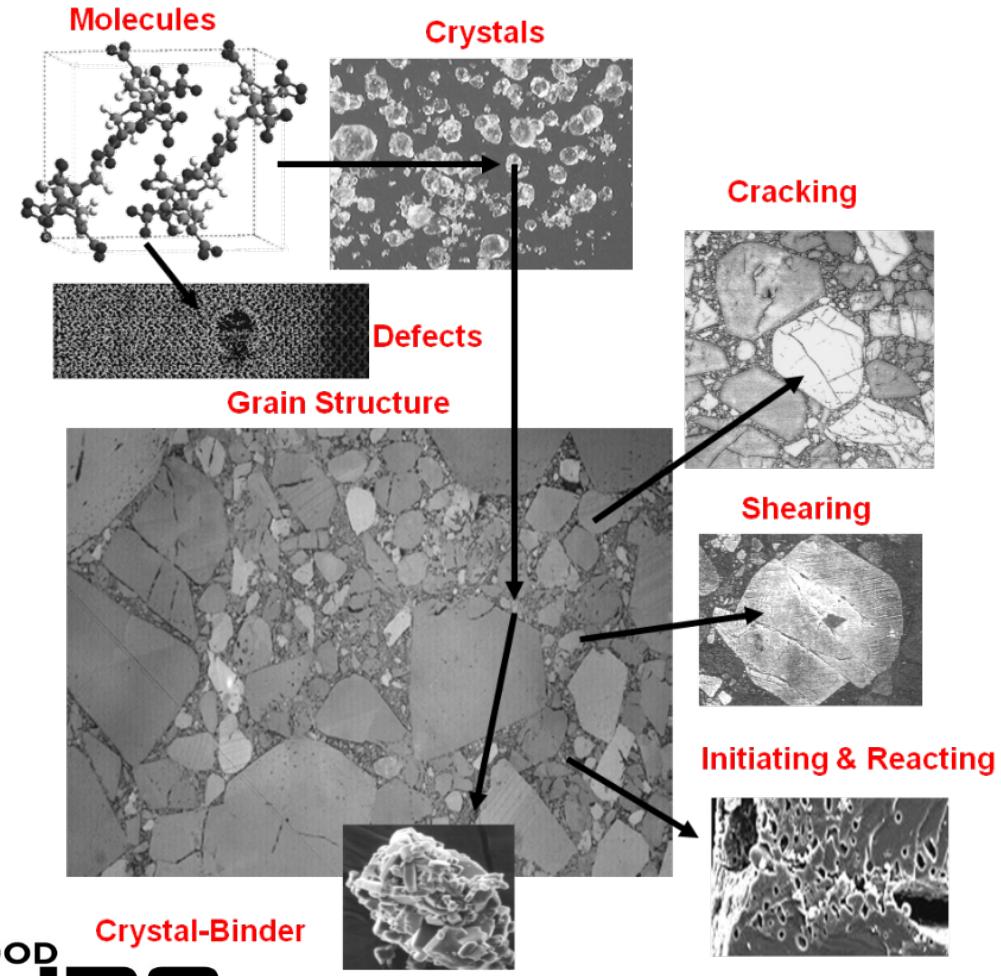

Hot spot mechanisms

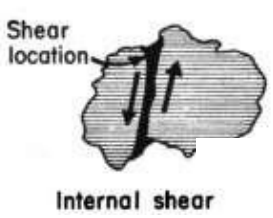

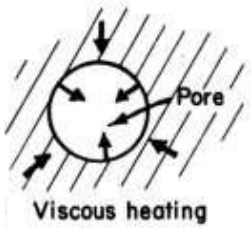

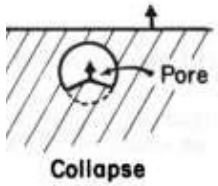

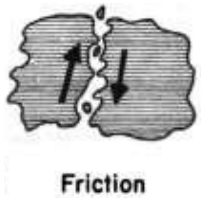

DOD HPC MODERNIZATION PROGRAM

ENGILITY

Microstructure Hot Spot Mechanisms

SINGLE IDEAL VOID: 10 nm

10 nm void – 3 km/s shock

- Collapsing void causes initiation
- Reactions at void location proceed farther and faster

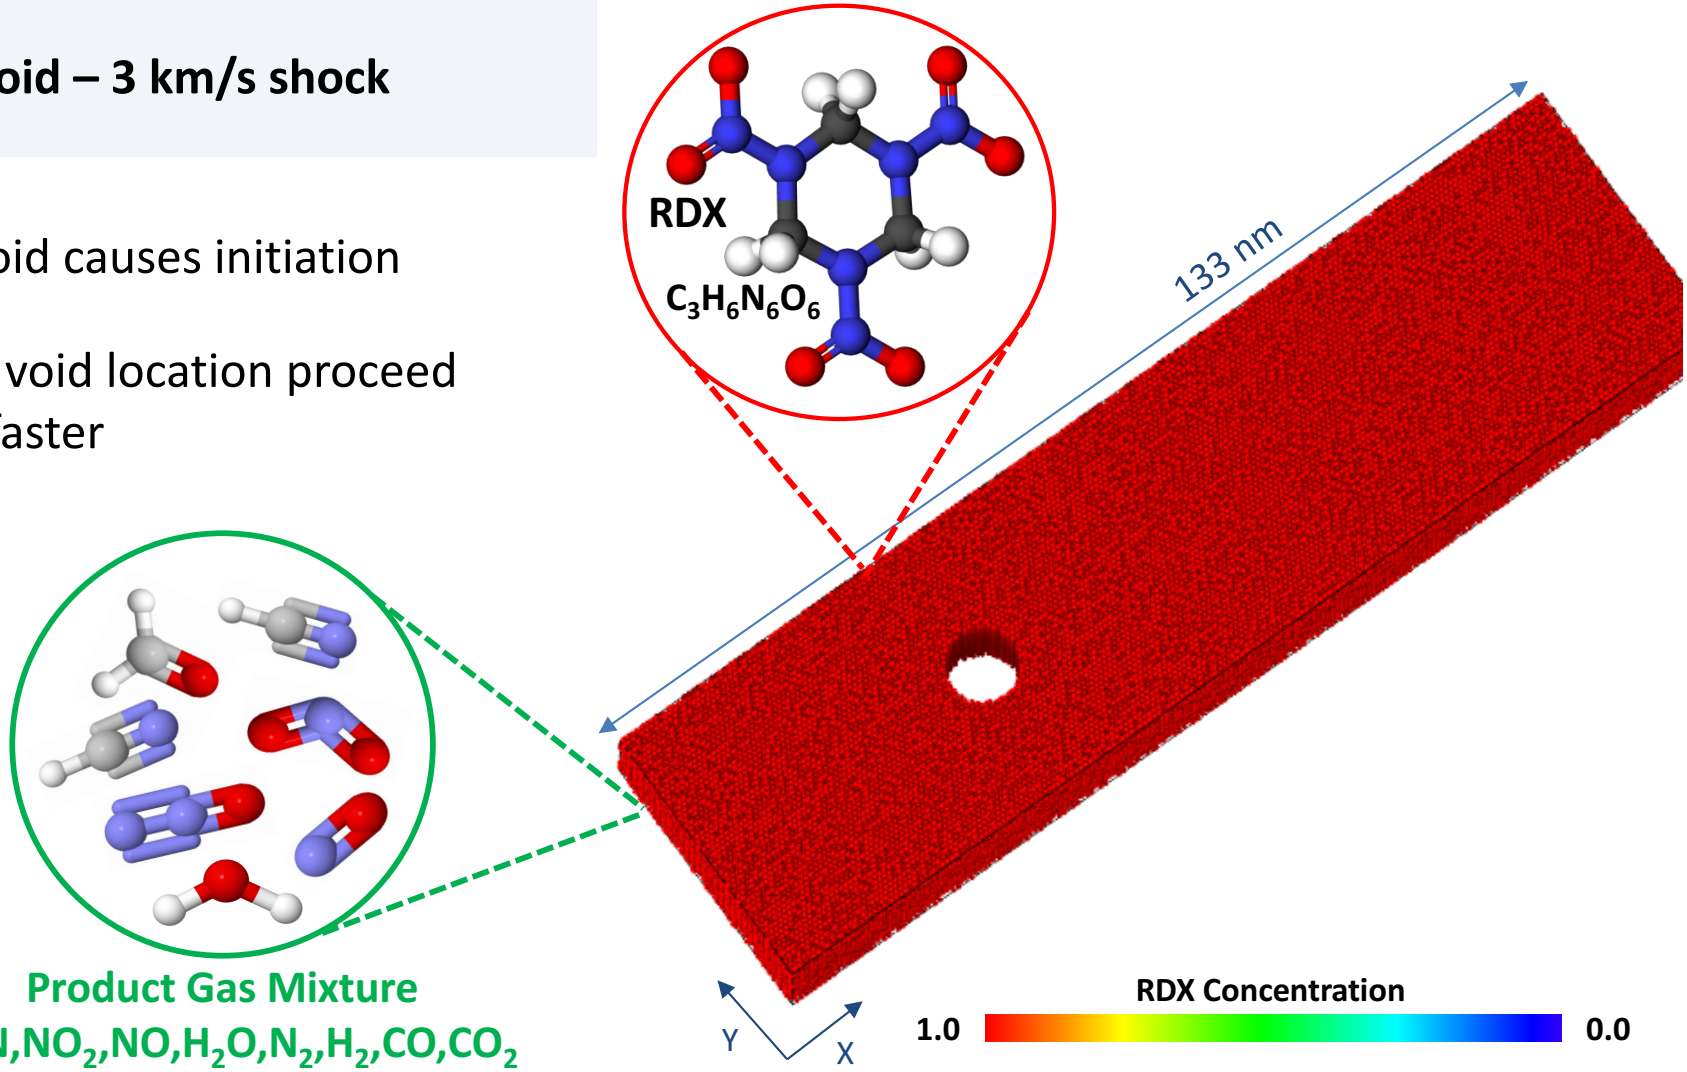

Goal: EM Composite Models

Approaching the scale of "dirty binder" regions

200.0 μm

2.5 μm

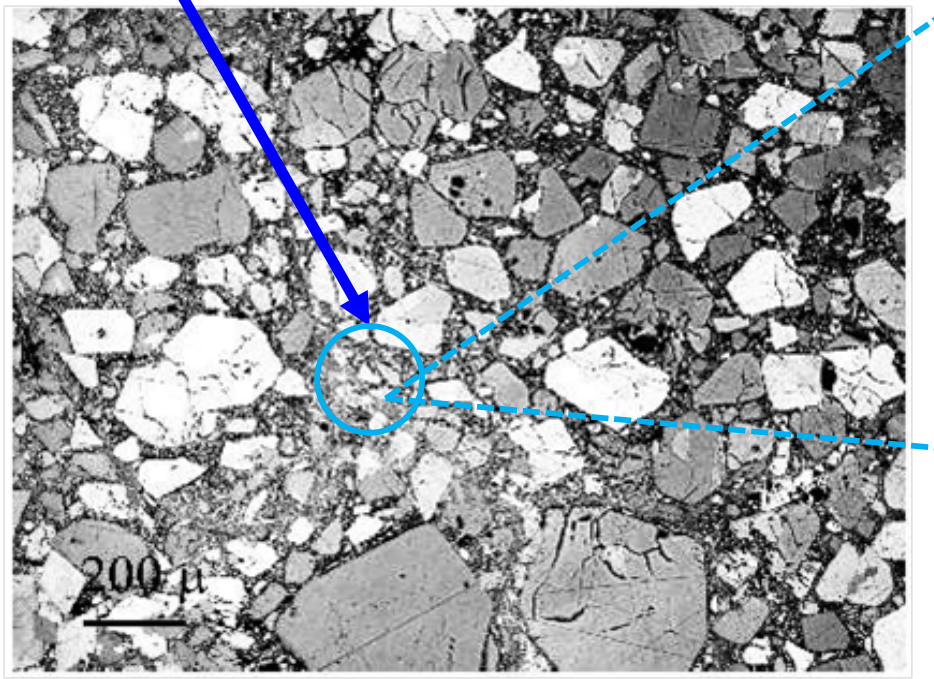

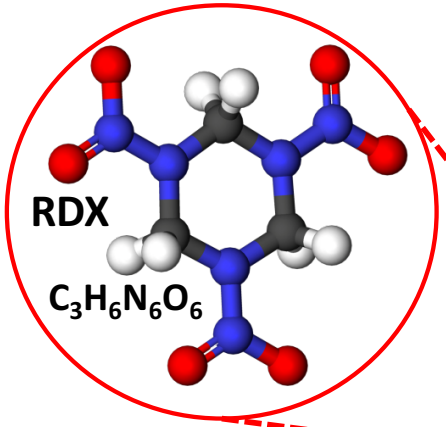

(yellow spheres, bonding not shown)

EM crystallite

polymer

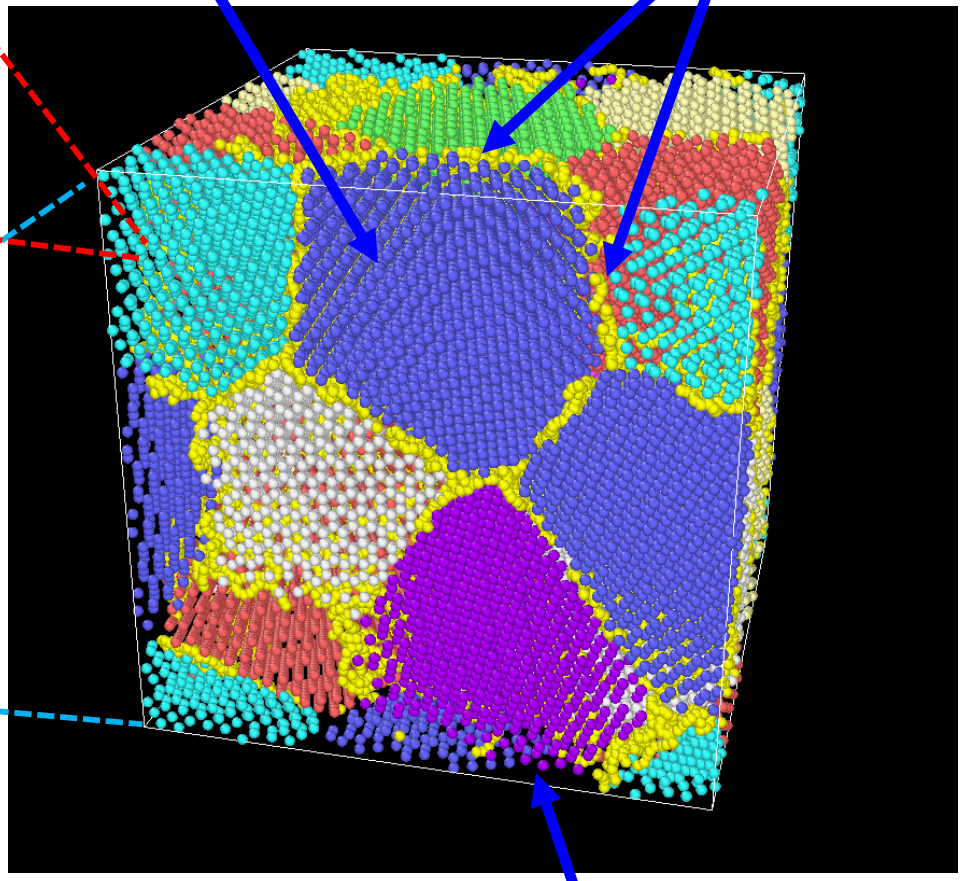

inter-granular void

Skidmore, C.B., Phillips, D.S., Howe, P.M., Mang, J.T., and J.A. Romero, 1998, 11th International Detonation Symposium, Snowmass Village, Colorado.

DOD HPC MODERNIZATION PROGRAM

ENGILITY

HPC Machines Used

System	Trinity Phase 2 (KNL)	Stampede-2	Thunder
Facility	Los Alamos National Laboratory (LANL)	Texas Advanced Computing Center (TACC)	US Airforce Research Laboratory (AFRL)
Peak Perf.	~27 PFLOP/s DP	~11 PFLOP/s DP	4.3 PFLOP/s DP
Approx. Size	~8,900 Nodes ~605,200 Cores	~3,600 Nodes ~244,800 Cores	3,216 Nodes 115,776 Cores
June 2017	est. **#6 or #7** (like Cori)	**#12** on Top500.org	**#36** on Top500.org
Kind	Cray XC40	Dell PowerEdge C6320P	HPE/SGI ICE X
Processor	Intel Xeon Phi 7250 KNL		Intel Xeon E5-2699v3 HSW
Node & Core Details	One 1.4 GHz CPU/node 68 cores/CPU (in 34 tiles) 4 Hardware Threads/core		Two 2.3 GHz CPUs/node 18 cores/CPU 1 Hardware Thread/core
Threads	Up to 272 threads/node		36 threads/node
Memory Hierarchy	32kB L1/core + 512kB L2/tile 16 GB MCDRAM/node 96 GB DDR4/node		32kB L1/core + 256kB L2/core 22.5 MB L3/CPU 128 GB DDR4/node

EM Shock Simulations at Scale: Inputs & Steps

Simulation Steps for a Shock applied to polycrystalline RDX:

1. **Load initial positions of all particles**
2. **Equilibrate at 325 Kelvin for 150 ps**
3. **Apply (fix wall/lj93) and Propagate Shock Wave for 350 ps**

Total of 500 ps simulated time

Model Name	Sample Size (nm^3)	Number of RDX Molecules	Number of Atoms	Grain Size Average (nm)	Number of Grains
Small	40 x 40 x 2,500	20,452,820	429,509,220	30	243
Medium	100 x 100 x 2,500	126,259,367	2,651,446,707	75	90
Large	300 x 300 x 2,500	1,126,926,339	23,665,453,119	225	31

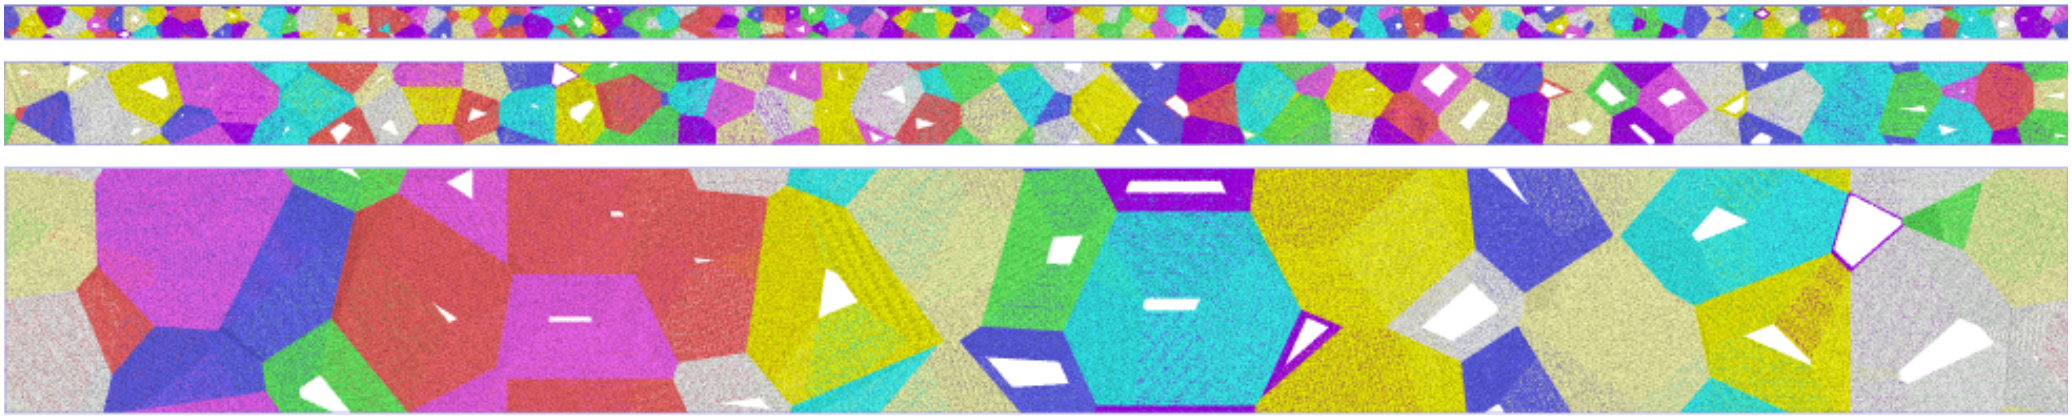

2.25 km/s plate induced shock

EM Shock Simulations at Scale: Atomistic Estimate

Estimating* Time-to-Solution with USER-REAXC in LAMMPS:

1. **Real runs have typical Computational Intensity of 500 to 2000 atoms/core**
2. **At 600 atoms per KNL core, achieve approximately 1.5 time steps/second**
3. **A 0.1 fs time step for a 500 ps simulation gives 5 x 10^6 time steps**
4. **On 8,820 nodes of Trinity Phase 2 (using 564,480 KNL cores):**

Model Name	Sample Size (nm^3)	Number of RDX Molecules	Number of Atoms	Atoms/ KNL Core	Wall Time (estimated*)
Small	40 x 40 x 2,500	20,452,820	429,509,220	761	7 weeks
Medium	100 x 100 x 2,500	126,259,367	2,651,446,707	4,697	10 months
Large	300 x 300 x 2,500	1,126,926,339	23,665,453,119	41,924	**7.4 years**

**Estimates are ideal scaling extrapolations by Aidan Thompson & Stan Moore (SNL) based on smaller EM simulation runs on Trinity with the USER-REAXC package in LAMMPS.*

There has to be a better way!

A ~~Better~~ Faster Way: Use Coarse-Graining (CG)

Why Coarse-Grain Simulation?

- ✓ Highly heterogeneous materials
- ✓ Dynamic responses over wide range of spatial and temporal scales
- ✓ All-atom simulation is too slow

energetic material composite

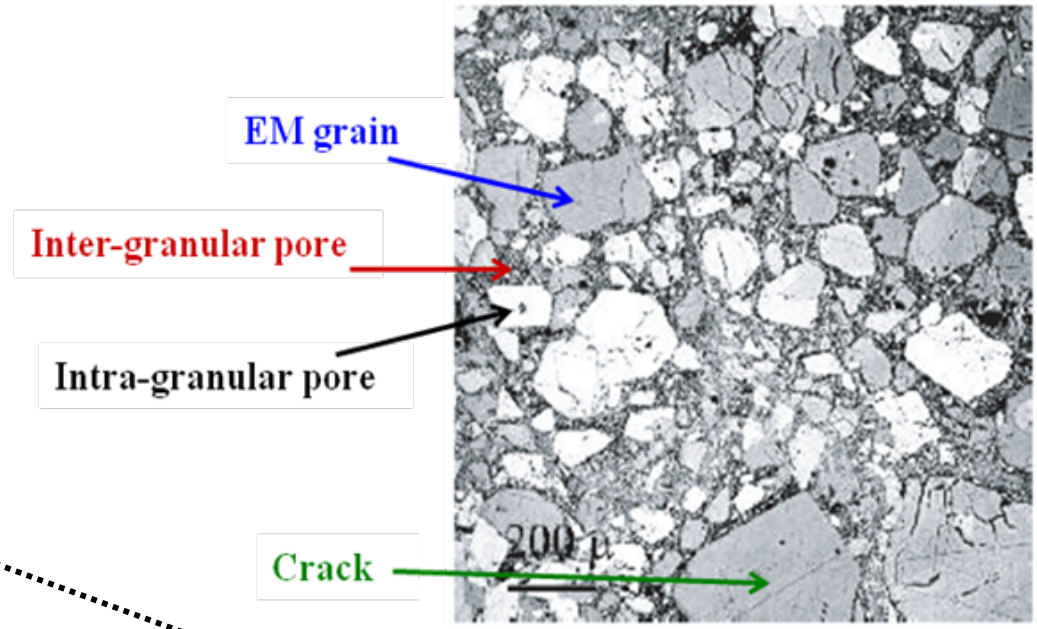

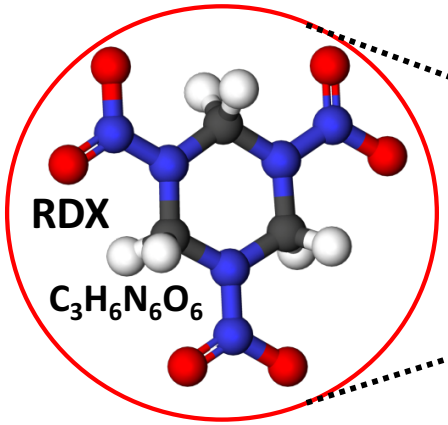

21 atoms

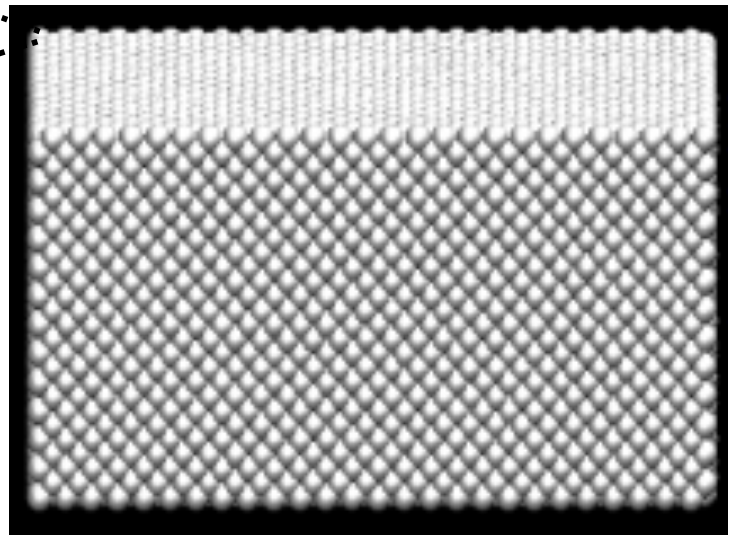

- *sacrifice atomistic detail*
- *gain computational speed*

DOD HPC MODERNIZATION PROGRAM

ENGILITY

1-site CG Model of RDX

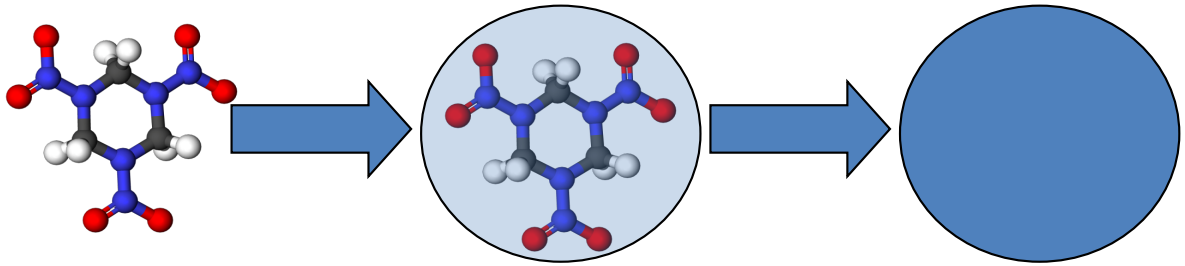

Force-Matching

Minimize $\chi^2 = \left\langle \left| F^{Atomistic} - F^{CG} \right|^2 \right\rangle$

CG potential fit by ***force matching***
CG and atomistic model

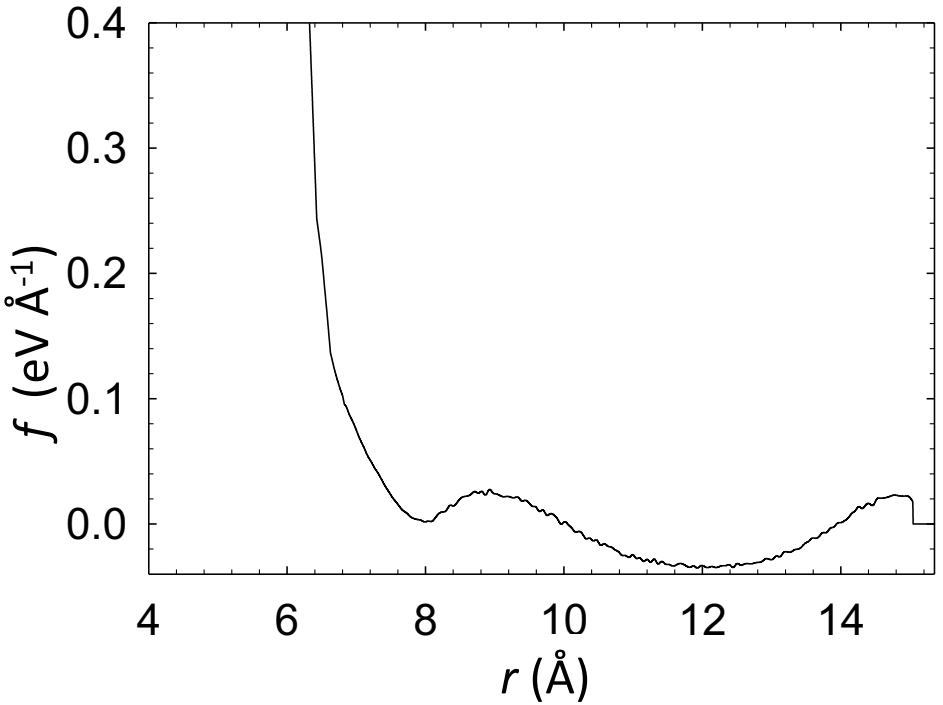

Pros:

- ❑Good standard state properties and Hugoniot
- ❑Good mechanical properties vs pressure

Cons:

- ❑Does not capture shear band formation
- ❑Melting point is too high at high pressures

J.D. Moore, B.C. Barnes, S. Izvekov et al., J. Chem. Phys., 144, 104501 (2016).

CG Method: Dissipative Particle Dynamics (DPD)

More specifically, DPD with Reactions (DPD-RX)

- ***Constant-E (or H) DPD helps*** **recapture lost degrees-of-freedom from CG-ing**
- **Thermal behavior can be recovered**

Each CG Particle Modeled as a Batch Reactor

RDX Decomposition

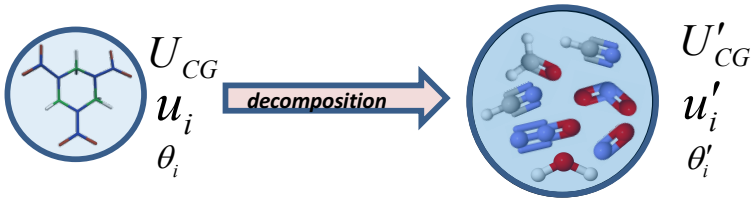

RDX

Product Gas Mixture
($HCN, NO_2, NO, H_20, N_2, H_2, CO, CO_2$)

Reduced Reaction Model

$RDX \rightarrow 3HCN + 3/2(NO_2 + NO + H_2O)$	endothermic
$HCN + NO_2 \rightarrow NO + ½(N_2 + H_2) + CO$ $HCN + NO \rightarrow CO + N_2 + ½H_2$ $NO + CO \rightarrow ½N_2 + CO_2$	exothermic

Brennan, J.K., Lísal, M., Moore, J.D., Izvekov, S., Schweigert, I.V., Larentzos, J.P. **Coarse-Grain Model Simulations of Nonequilibrium Dynamics in Heterogeneous Materials** *The Journal of Physical Chemistry Letters* **2014** *5* (12)

A simple matter of coding…

Several years to implement previous slides in an ARL customized LAMMPS:

- **Constant Energy Dissipative Particle Dynamics (DPD-E) Method**
- **Serial and Parallel Stochastic Integration Scheme: Shardlow Splitting Algorithm (SSA)†**

  †Note: The SSA allows ~2 orders of magnitude larger time-step size than Velocity-Verlet for DPD-E.
- **Reaction Kinetics Solver for "Batch Reactors" (DPD-RX)**

Dr. Larentzos (with Dr. Barnes) presented above at LAMMPS Workshop August 7, 2015:

"*Recent Extensions of Dissipative Particle Dynamics Methods and Application to Hierarchical Multiscale Simulation*"

Last two years (~17,000 changed/added lines of code since Jan 2017):

- **Merged USER-DPD into mainline LAMMPS**
- **Made the neighbor list code more modular to support SSA, etc.**
- **Optimized the code for performance on modern architectures**
- **Improved scalability and portability via a Kokkos implementation**

A few problems on the way…

- **Stochastic… means Random**
  - Changes in particle force evaluation order, *will* change the end trajectories!
  - Had to develop sophisticated validation scripts and methods
- **Simulating a billion particles will overflow 32-bit integers… someplace…**
- **Creating input files of a billion particles is hard.**
- **Reading those input files has to scale too…**
- **Don't talk to me (yet) about saving/storing full trajectories of a billion particles!**
- **Many parallel computations can't be easily described as a one-dimensional loop over independent coordinates or units of work.**
  - To expose parallelism within the SSA required a new "coloring" scheme, expressed as a new neighbor "list" in LAMMPS
  - Developed a way to compactly represent a work plan for the parallel evaluation of non-interfering particle pairs in the SSA
- **Automatic Vectorization flags for compilers are not magical**
- **The KNL based supercomputers were in "Early Access/Pioneer" periods**
  - Our code caused physical failure of KNL nodes in odd ways… not seen by others

We did some runs…

Tested the code using tiny versions of the shock simulations on:

- **ARL KNL testbed system**
- **TACC KNL testbed system**
- **Haswell production systems at the various DoD DSRCs**

Made MPI geometry selections for three Computational Intensities

- **High, ~5000 particles per KNL core**
- **Moderate, ~2000 particles per KNL core**
- **Low, ~570 particles per KNL core**

Ran the simulations for a Gordon Bell Prize submission on:

Trinity/KNL


Stampede-2

Thunder

DOD HPC MODERNIZATION PROGRAM

USER-DPD Time-to-solution on Trinity Phase 2

Energetic Material Simulation Milestones:

- **Scaled to over 0.5 million cores**
- **Ran 1.1 billion DPD-RX particles for 0.5 ns of simulated time**

Science results in hours, not years!

Computational Intensity	Model Name	Trinity KNL Nodes	KNL Cores Used	Particles per Core	Atoms per Core	ReaxFF Estimated Wall Time	USER-DPD Measured Wall Time
High	Small	63	4,032	5,073	106,525	18.8 years	8.93 hours
	Medium	390	24,960	5,058	106,228	18.7 years	9.50 hours
	Large	3,465	221,760	5,082	106,717	18.8 years	11.02 hours
Moderate	Small	162	10,368	1,973	41,426	7.3 years	3.75 hours
	Medium	988	63,232	1,997	41,932	7.4 years	4.34 hours
	Large	8,820	**564,480**	1,996	41,924	**7.4** years	**5.59** hours
Low	Small	561	35904	570	11,963	2.1 years	1.67 hours
	Medium	3,465	221,760	569	11,956	2.1 years	1.64 hours

DOD HPC MODERNIZATION PROGRAM

ENGILITY

USER-DPD Scaling on Trinity/KNL & Stampede 2

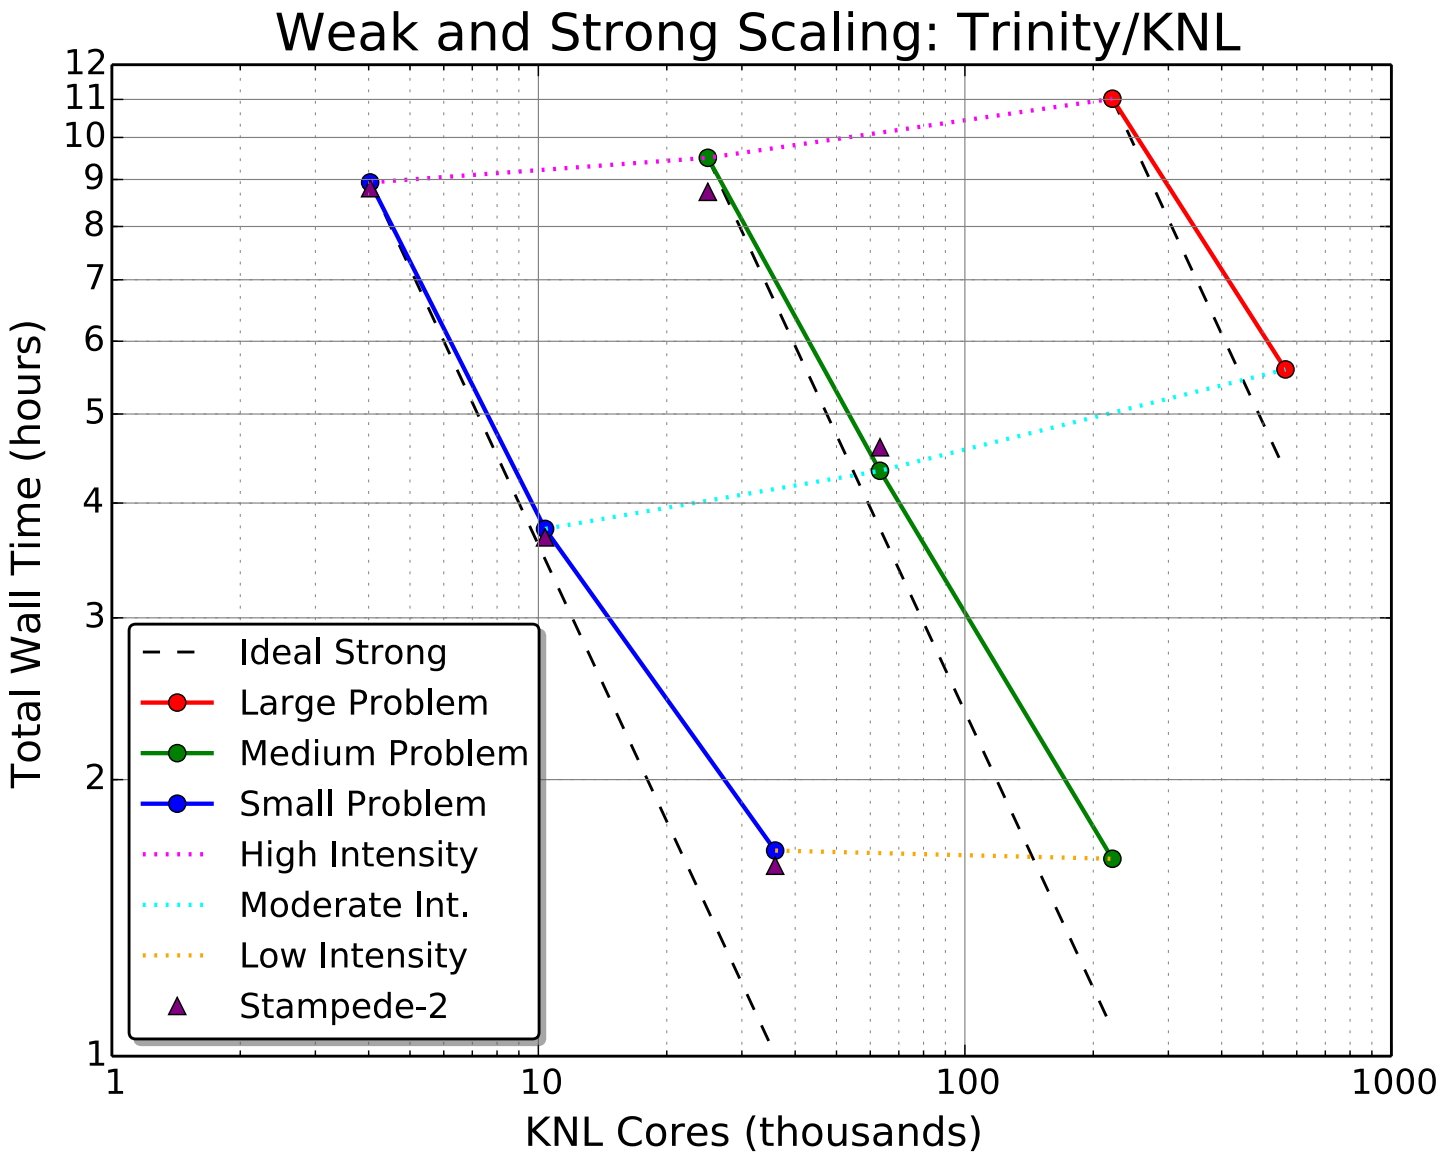

USER-DPD Scaling on Trinity/KNL & Stampede 2

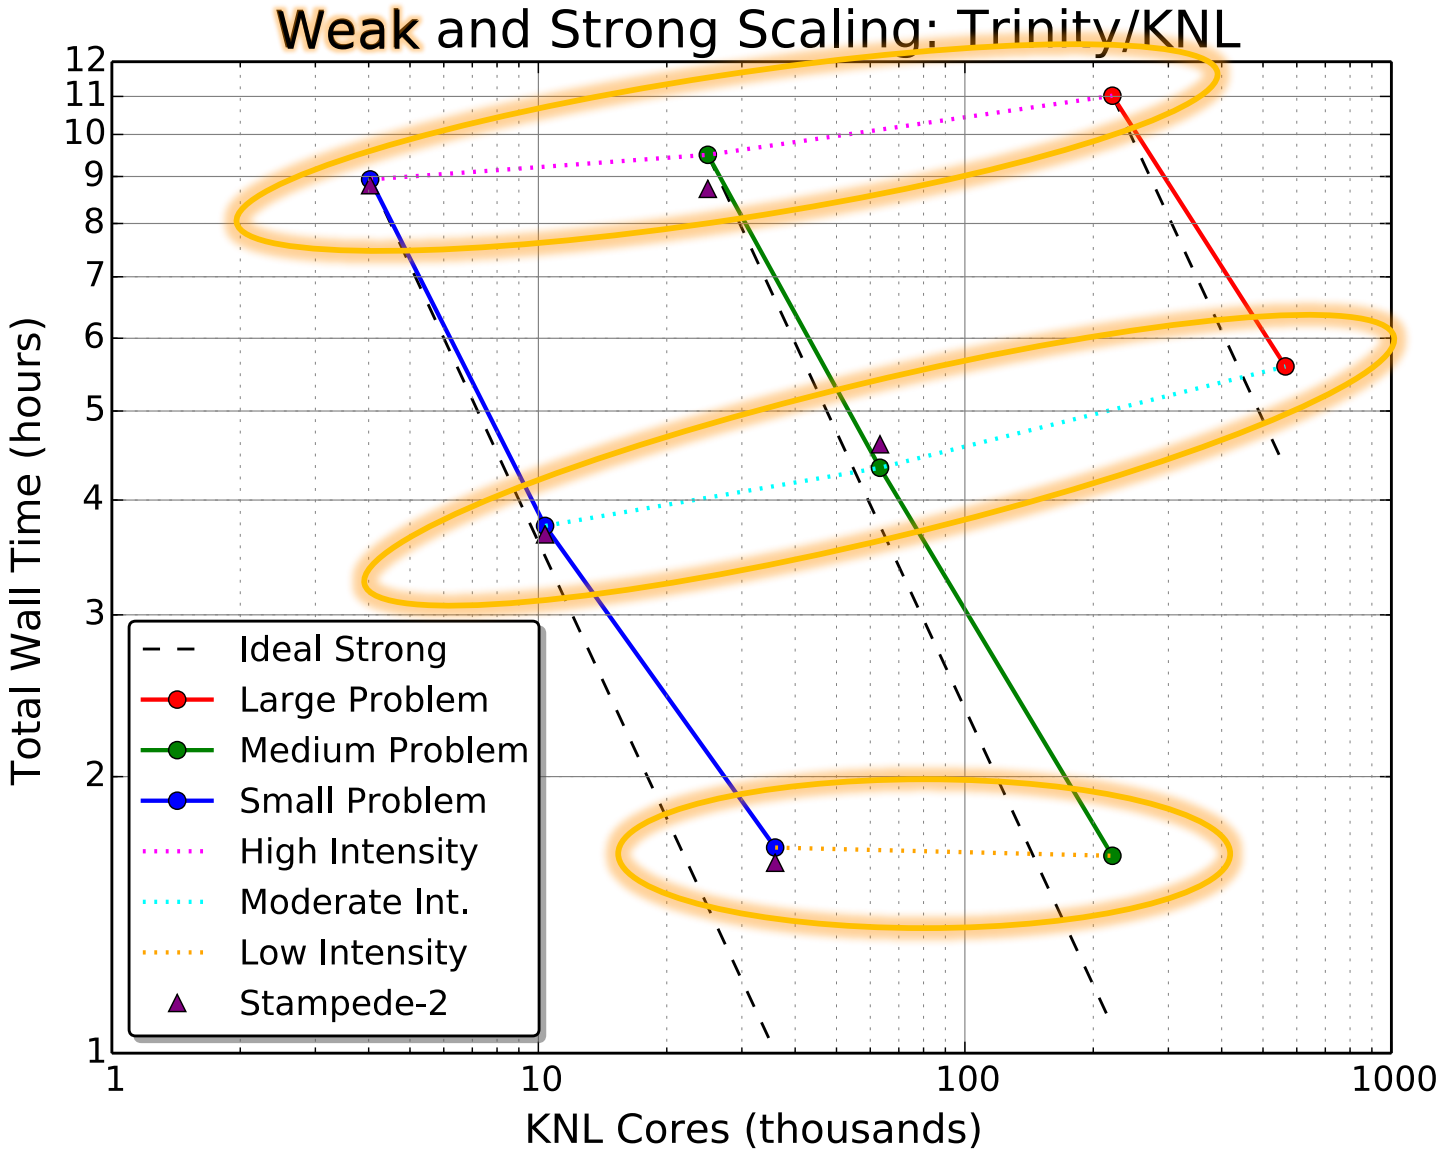

USER-DPD Scaling on Trinity/KNL & Stampede 2

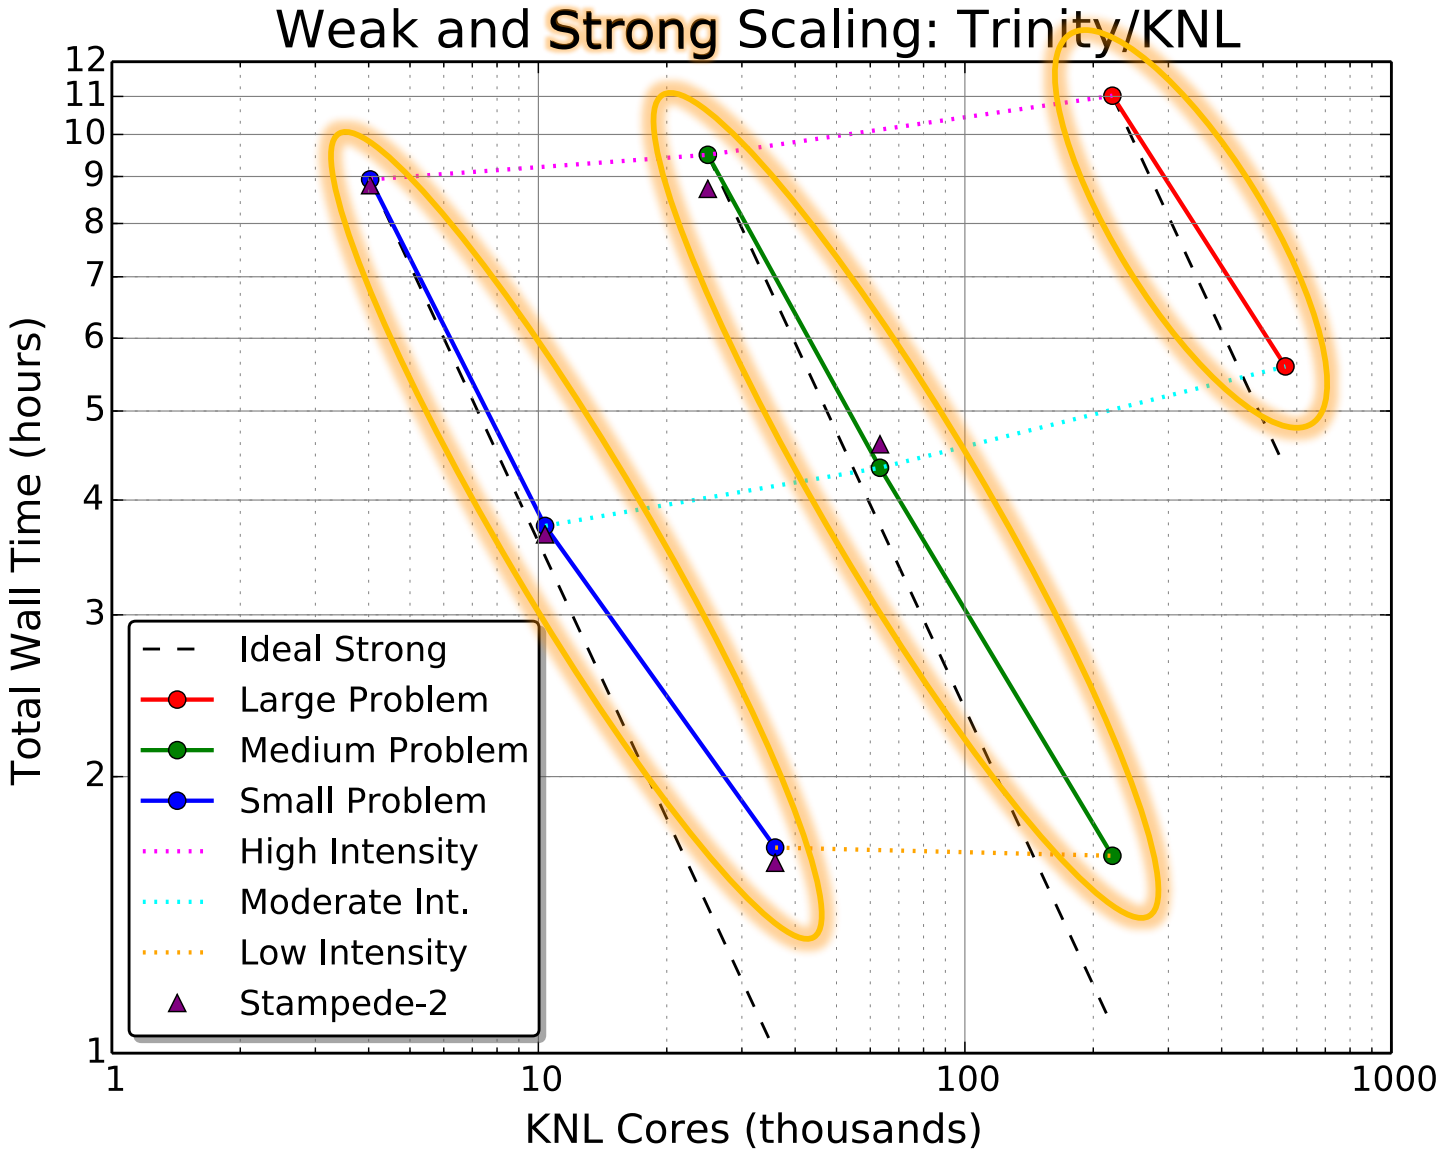

USER-DPD Scaling on Trinity/KNL & Stampede 2

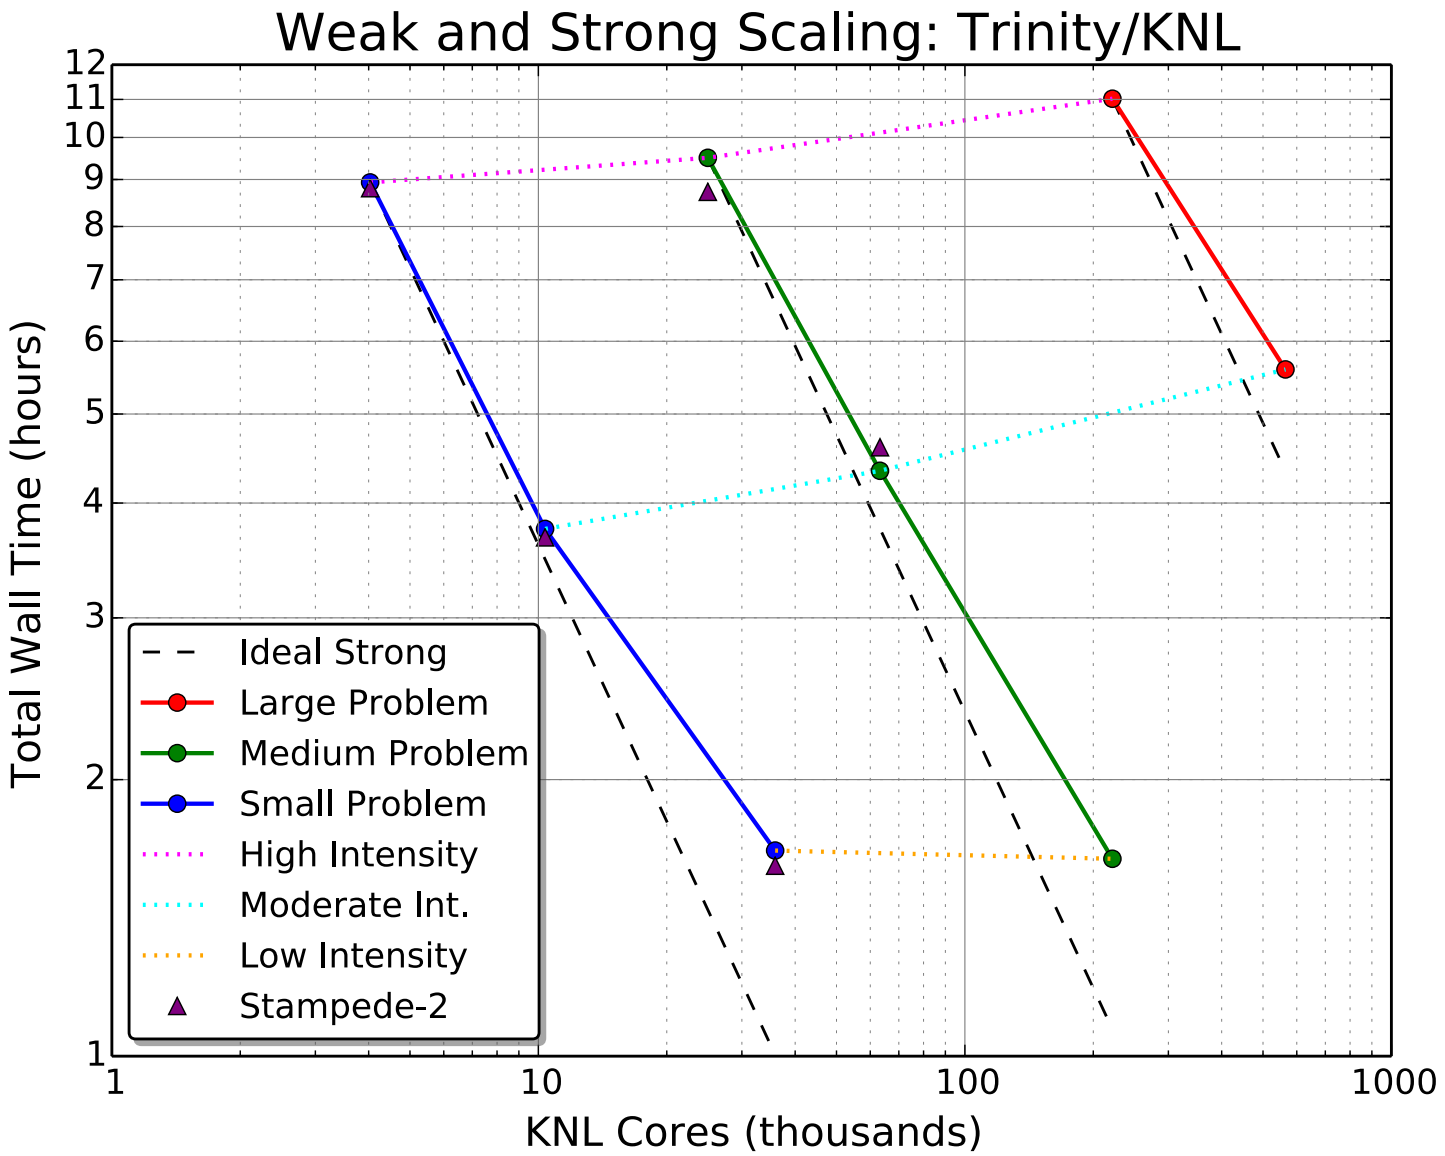

A 1.1 Billion Particle Polycrystal Shock Simulation

~225 nm grains

2500x300x300 nm^3

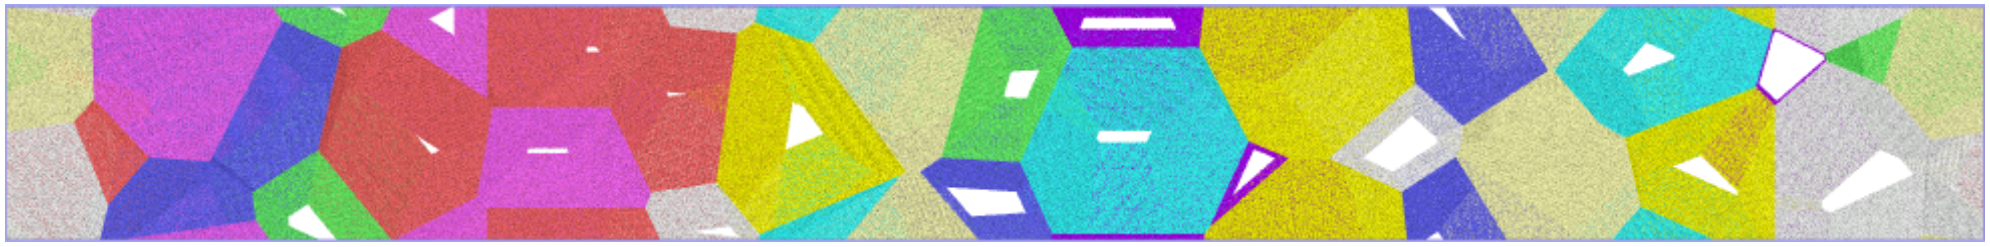

- **snapshots after 100 ps**

RDX Decomposition

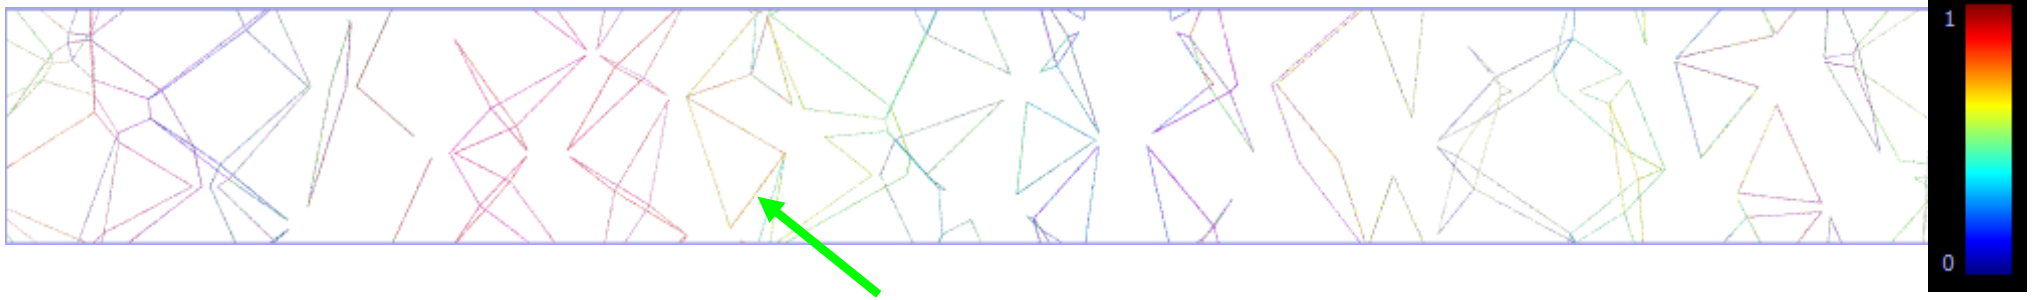

particles at grain interface(s)

u_p = 2.25 km/s

- **Particles highlighted if at least 10% of RDX has decomposed**

DOD HPC MODERNIZATION PROGRAM

ENGILITY

A 1.1 Billion Particle Polycrystal Shock Simulation

~225 nm grains

2500x300x300 nm^3

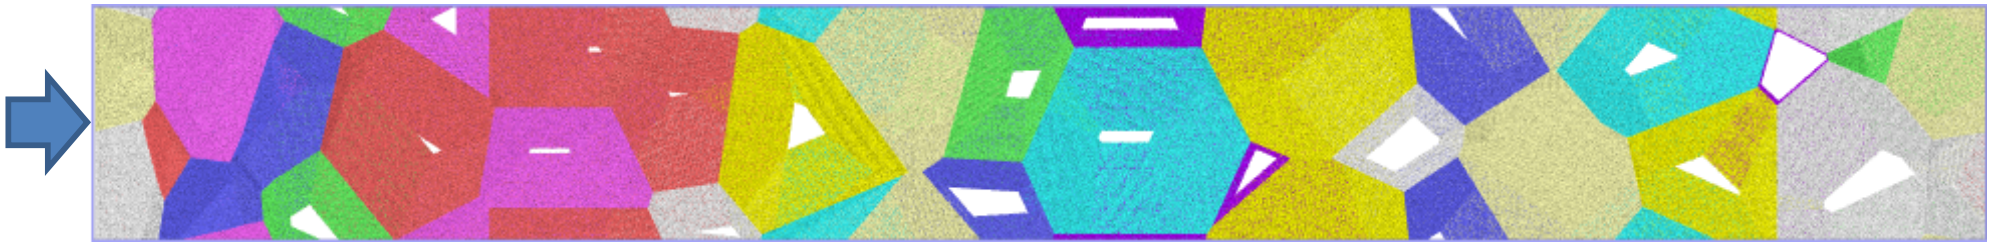

- **snapshots after 200 ps**

RDX Decomposition

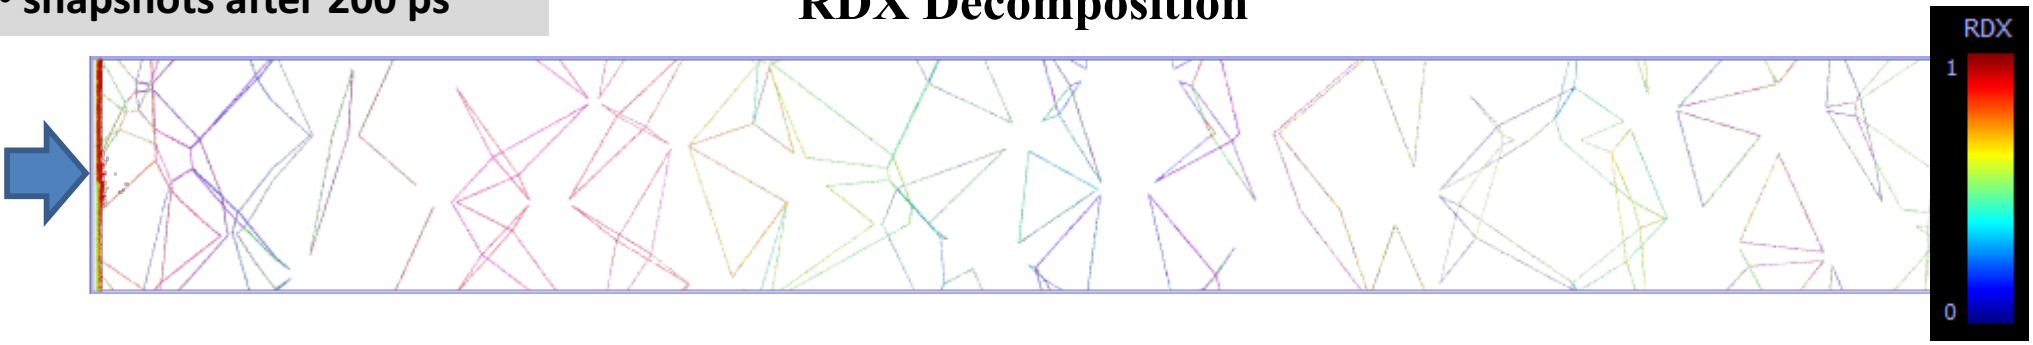

u_p = 2.25 km/s

- **Particles highlighted if at least 10% of RDX has decomposed**

DOD HPC MODERNIZATION PROGRAM

ENGILITY

A 1.1 Billion Particle Polycrystal Shock Simulation

~225 nm grains

2500x300x300 nm^3

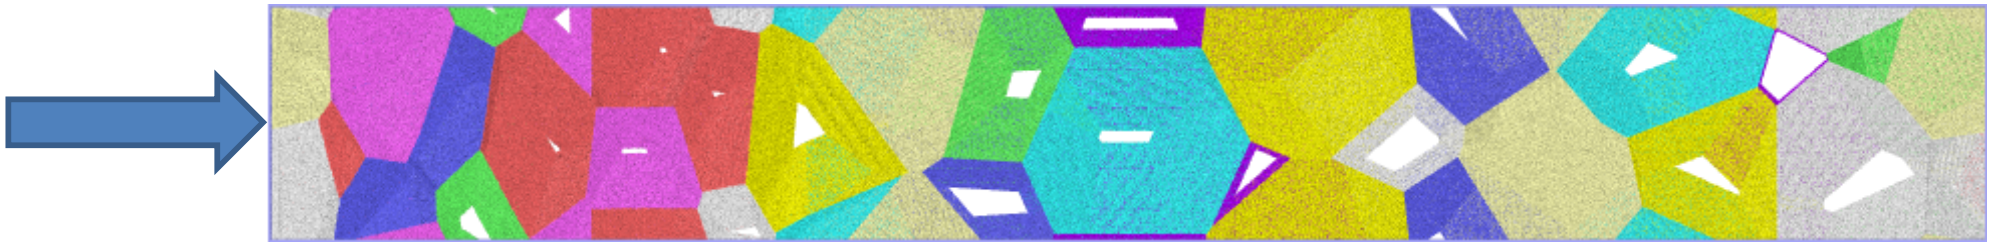

- **snapshots after 300 ps**

RDX Decomposition

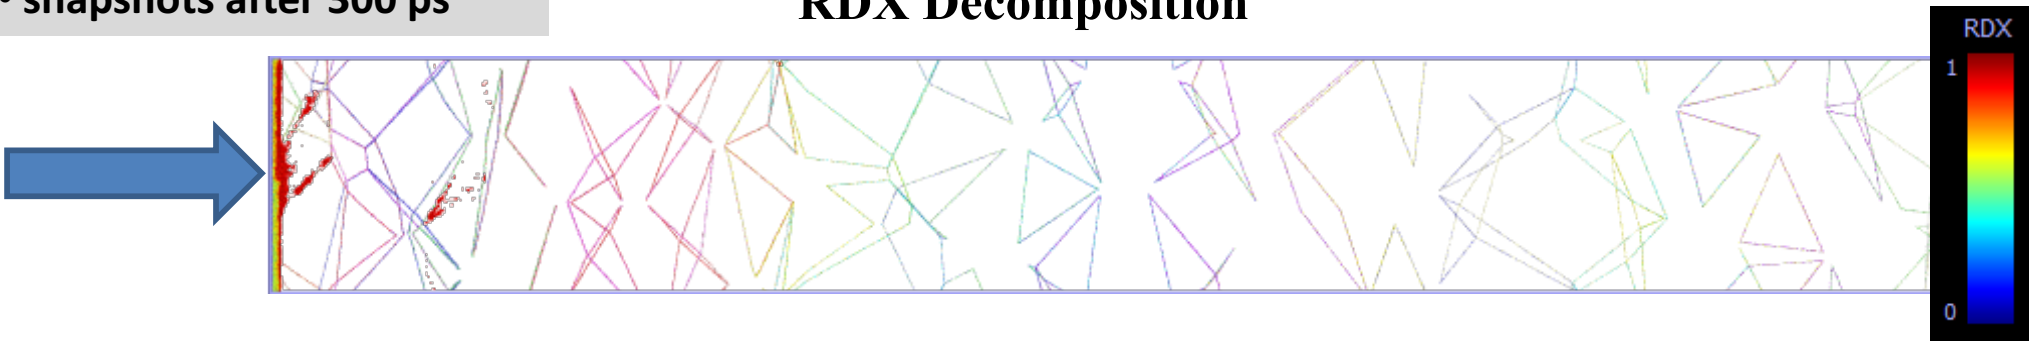

u_p = 2.25 km/s

- **Particles highlighted if at least 10% of RDX has decomposed**

DOD HPC MODERNIZATION PROGRAM

ENGILITY

A 1.1 Billion Particle Polycrystal Shock Simulation

~225 nm grains

2500x300x300 nm^3

- **snapshots after 400 ps**

RDX Decomposition

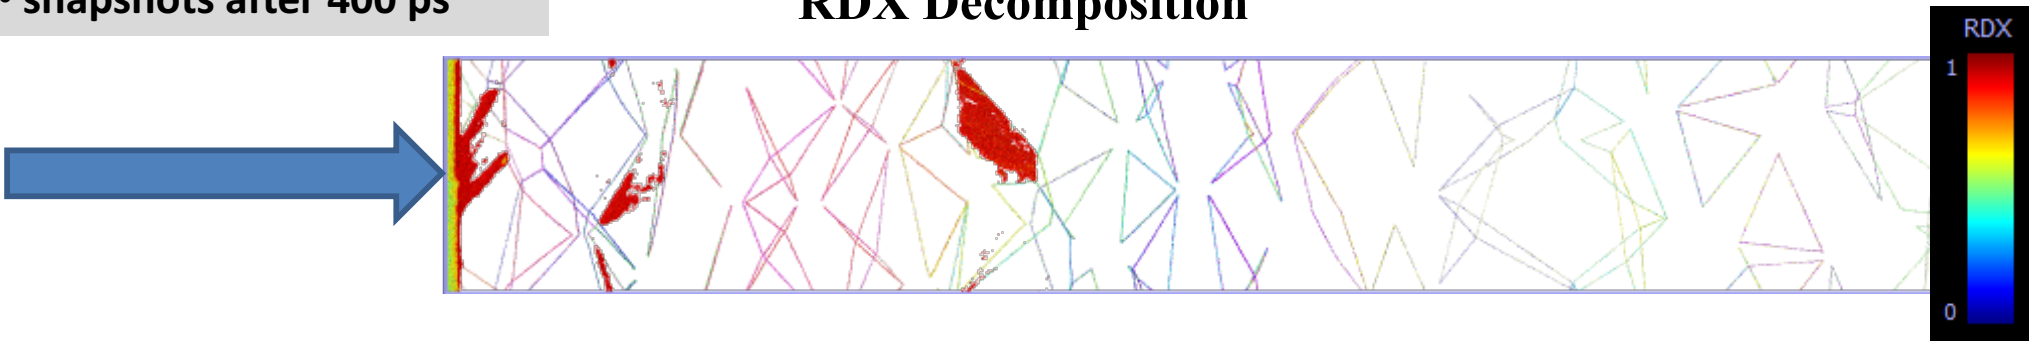

u_p = 2.25 km/s

- **Particles highlighted if at least 10% of RDX has decomposed**

DOD HPC MODERNIZATION PROGRAM

ENGILITY

A 1.1 Billion Particle Polycrystal Shock Simulation

~225 nm grains

2500x300x300 nm^3

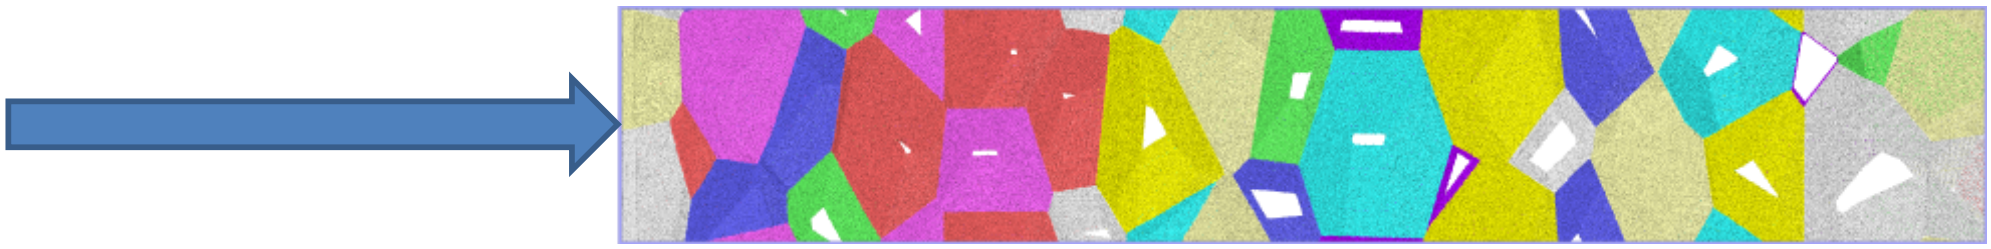

- **snapshots after 500 ps**

RDX Decomposition

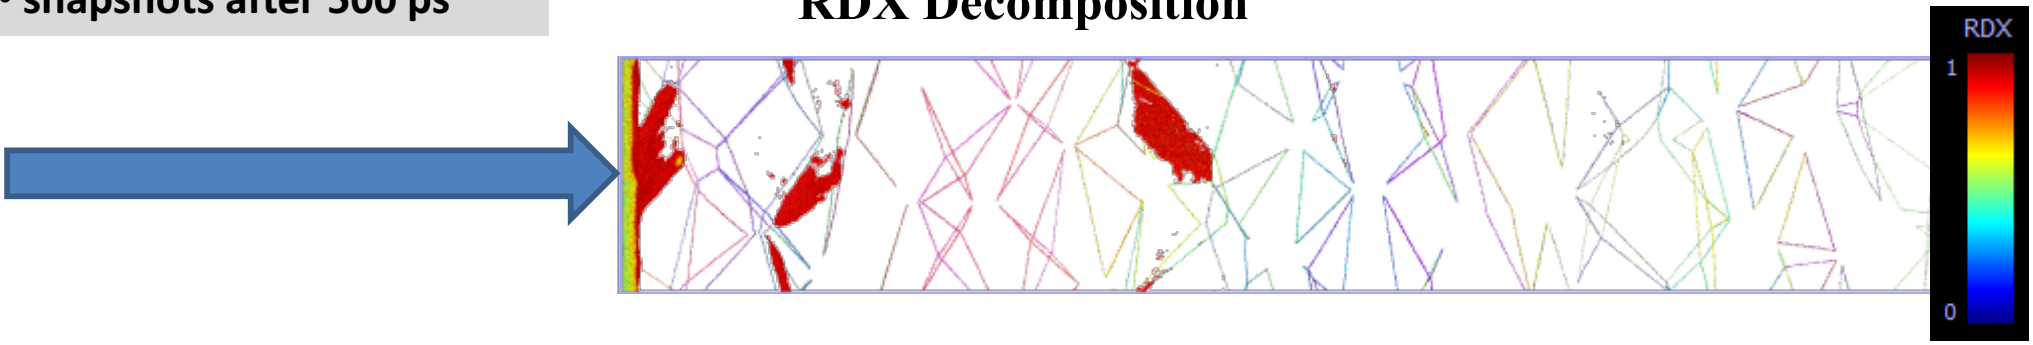

u_p = 2.25 km/s

- **Particles highlighted if at least 10% of RDX has decomposed**

DOD HPC MODERNIZATION PROGRAM

ENGILITY

A 1.1 Billion Particle Polycrystal Shock Simulation

2500x300x300 nm^3 RDX Polycrystal, ~225 nm grains

RDX Decomposition

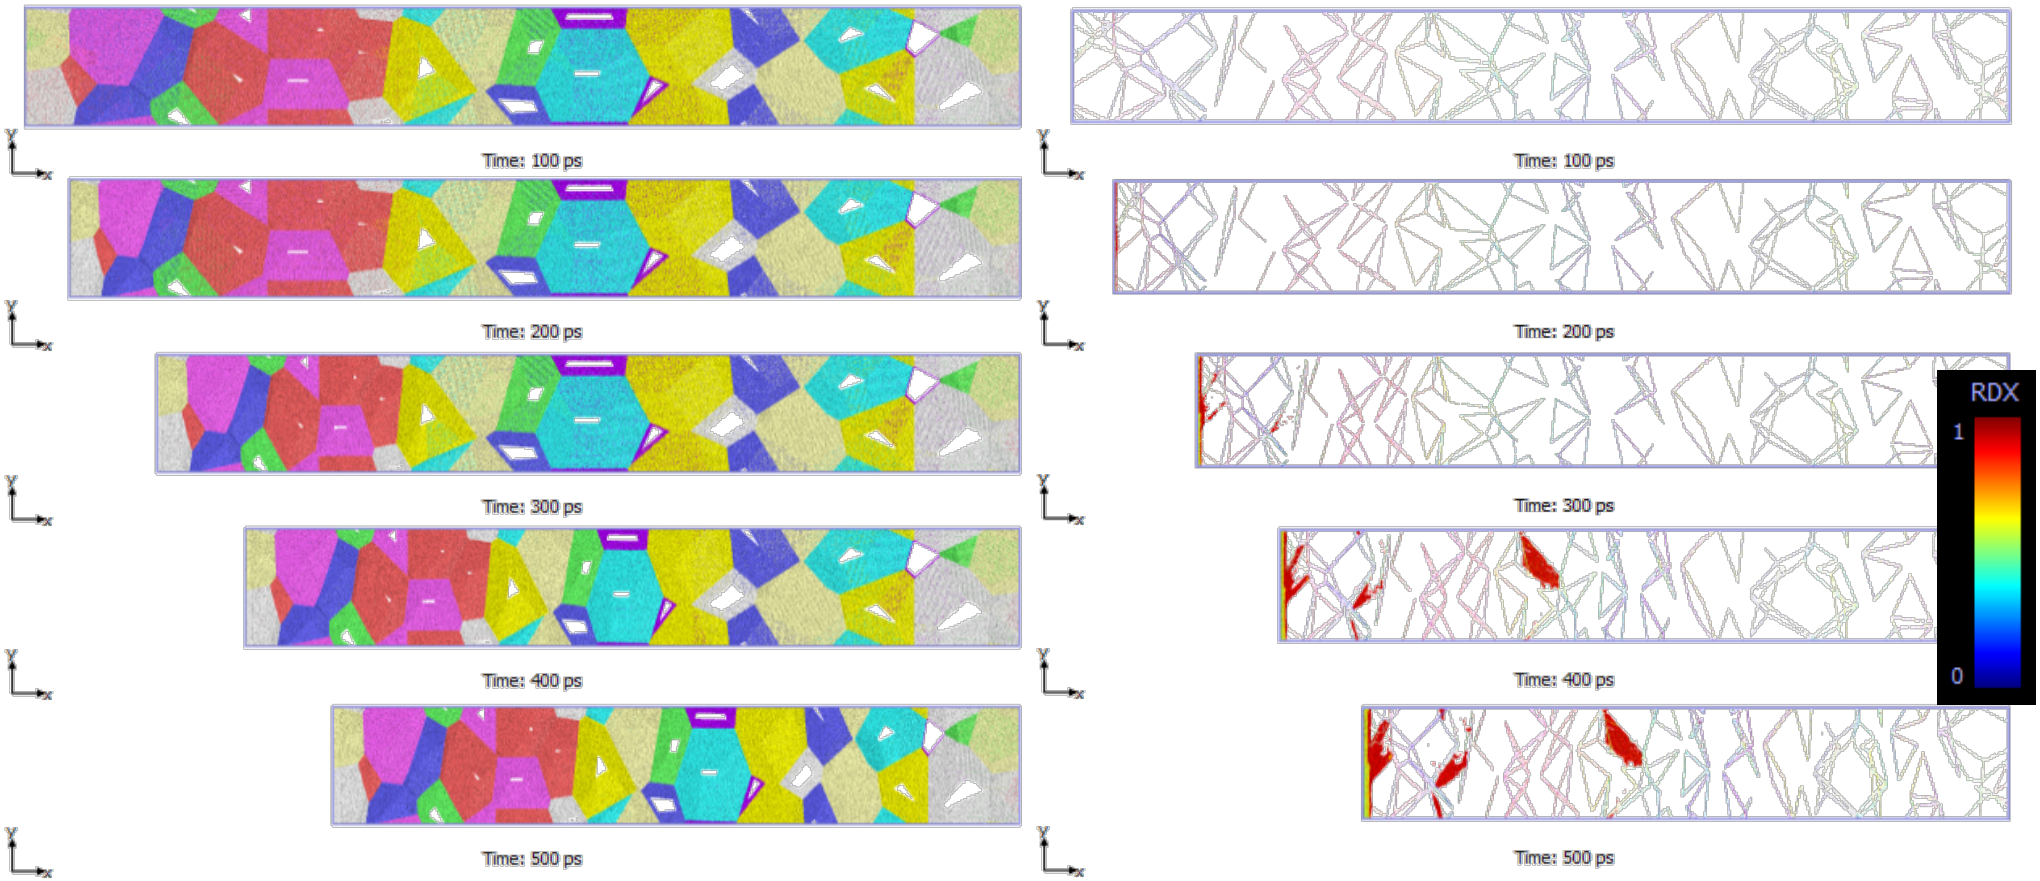

u_p = 2.25 km/s

- **Particles highlighted if at least 10% of RDX has decomposed**

DOD HPC MODERNIZATION PROGRAM

ENGILITY

Conclusions

Material Science:

- O(billion) particle DPD-RX simulations are now feasible
- Science results in hours, not years!
- Studies of microstructure effects on EM sensitivity are ongoing

Computational Science:

- Scaling Simulations to >500,000 cores takes a village
- Even with help of Kokkos, intra-node parallelism is tough
- Scaling the simulation itself isn't enough
  - How to generate the input sets at scale?
  - How to save the results without overloading the IO system?
  - How to do the analysis of the now much larger results?
  - How to archive the results?

Future & Collaborative Work

- Hierarchical modeling techniques to reach larger scales
- Studying alternatives/improvements to Kokkos
- DPD method improvements

Questions? timothy.mattox@engilitycorp.com

Again, “Thank you!” to all the many people, agencies, and understanding spouses that made this work possible!

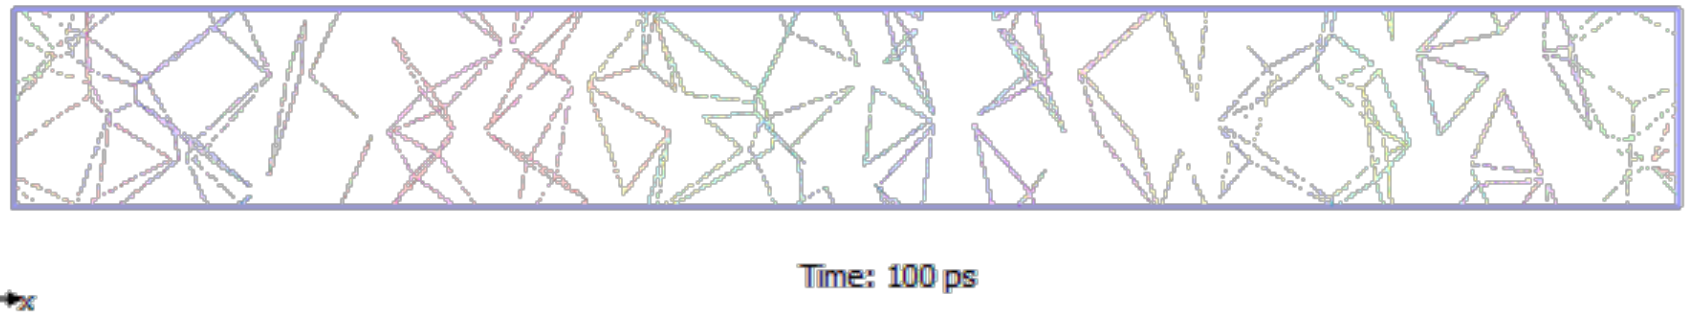

Backup Slides

USER-DPD Time-to-solution: Three Machines Data

Machine	Model Name	Nodes	Cores Used	Particles per Core	Atoms per Core	ReaxFF Estimated Wall Time	USER-DPD Measured Wall Time	Estimated Speedup
Thunder (HSW)	Small	51	1,836	11,140	233,937	18.5 years	8.64 hrs	1.9×10^4
		114	4,104	4,984	104,656	8.3 years	3.87 hrs	1.9×10^4
		452	16,272	1,257	26,396	2.1 years	1.54 hrs	1.2×10^4
	Medium	312	11,232	11,241	236,062	18.7 years	9.71 hrs	1.7×10^4
		702	25,272	4,996	104,916	8.3 years	4.87 hrs	1.5×10^4
Stampede2 (KNL)	Small	63	4,032	5,073	106,525	18.8 years	8.81 hrs	1.9×10^4
		162	10,368	1,973	41,426	7.3 years	3.67 hrs	1.7×10^4
		561	35904	570	11,963	2.1 years	1.61 hrs	1.1×10^4
	Medium	390	24,960	5,058	106,228	18.7 years	8.73 hrs	1.9×10^4
		988	63,232	1,997	41,932	7.4 years	4.60 hrs	1.4×10^4
Trinity (KNL)	Small	63	4,032	5,073	106,525	18.8 years	8.93 hrs	1.8×10^4
		162	10,368	1,973	41,426	7.3 years	3.75 hrs	1.7×10^4
		561	35904	570	11,963	2.1 years	1.67 hrs	1.1×10^4
	Medium	390	24,960	5,058	106,228	18.7 years	9.50 hrs	1.7×10^4
		988	63,232	1,997	41,932	7.4 years	4.34 hrs	1.5×10^4
		3,465	221,760	569	11,956	2.1 years	1.64 hrs	1.1×10^4
	Large	3,465	221,760	5,082	106,717	18.8 years	11.02 hrs	1.5×10^4
		8,820	**564,480**	1,996	41,924	7.4 years	5.59 hrs	1.2×10^4

Comparison of Microstructure Types

- ~1/10th of sample is shown after 50 ps
- Particles shown if at least 30% of RDX has decomposed

u_p = 2.25 km/s

2500x40x40 nm^3

Particle Internal Temperature

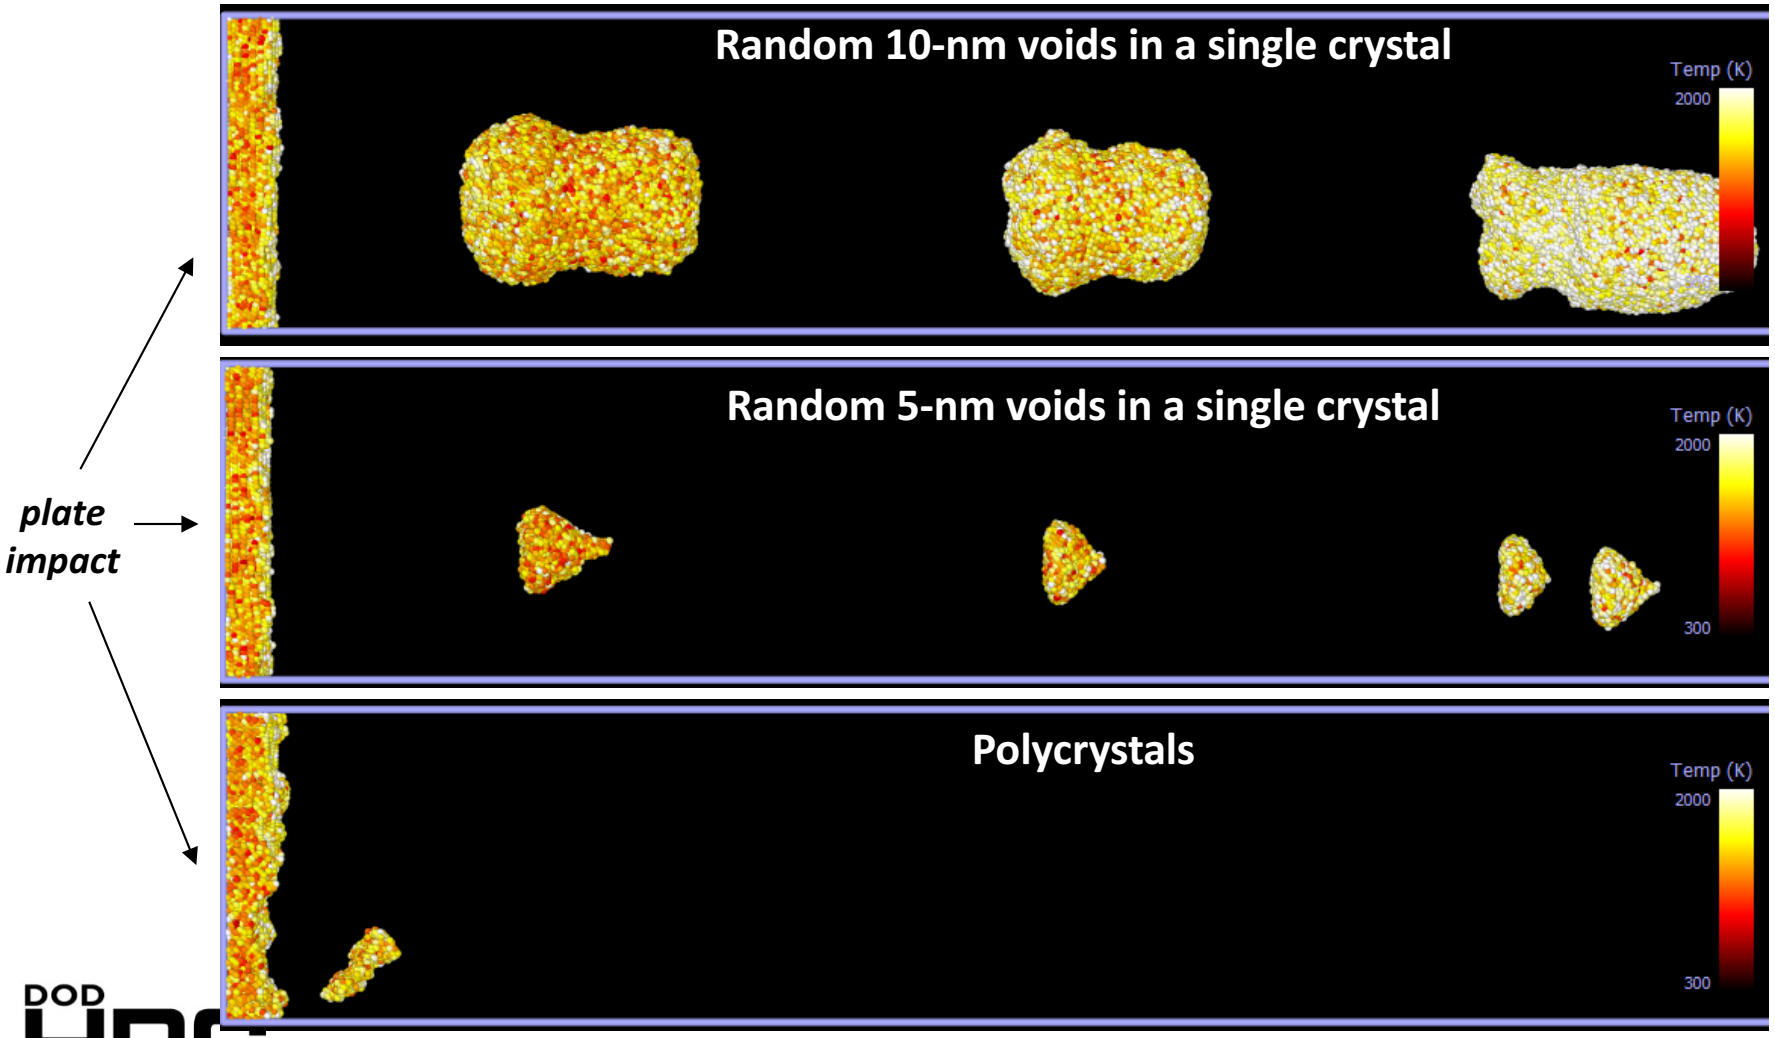

Comparison of Microstructure Types

- **~1/10th of sample is shown after 50 ps**
- **Particles shown if at least 30% of RDX has decomposed**

u_p = 2.25 km/s

2500x40x40 nm^3

Particle Internal Temperature

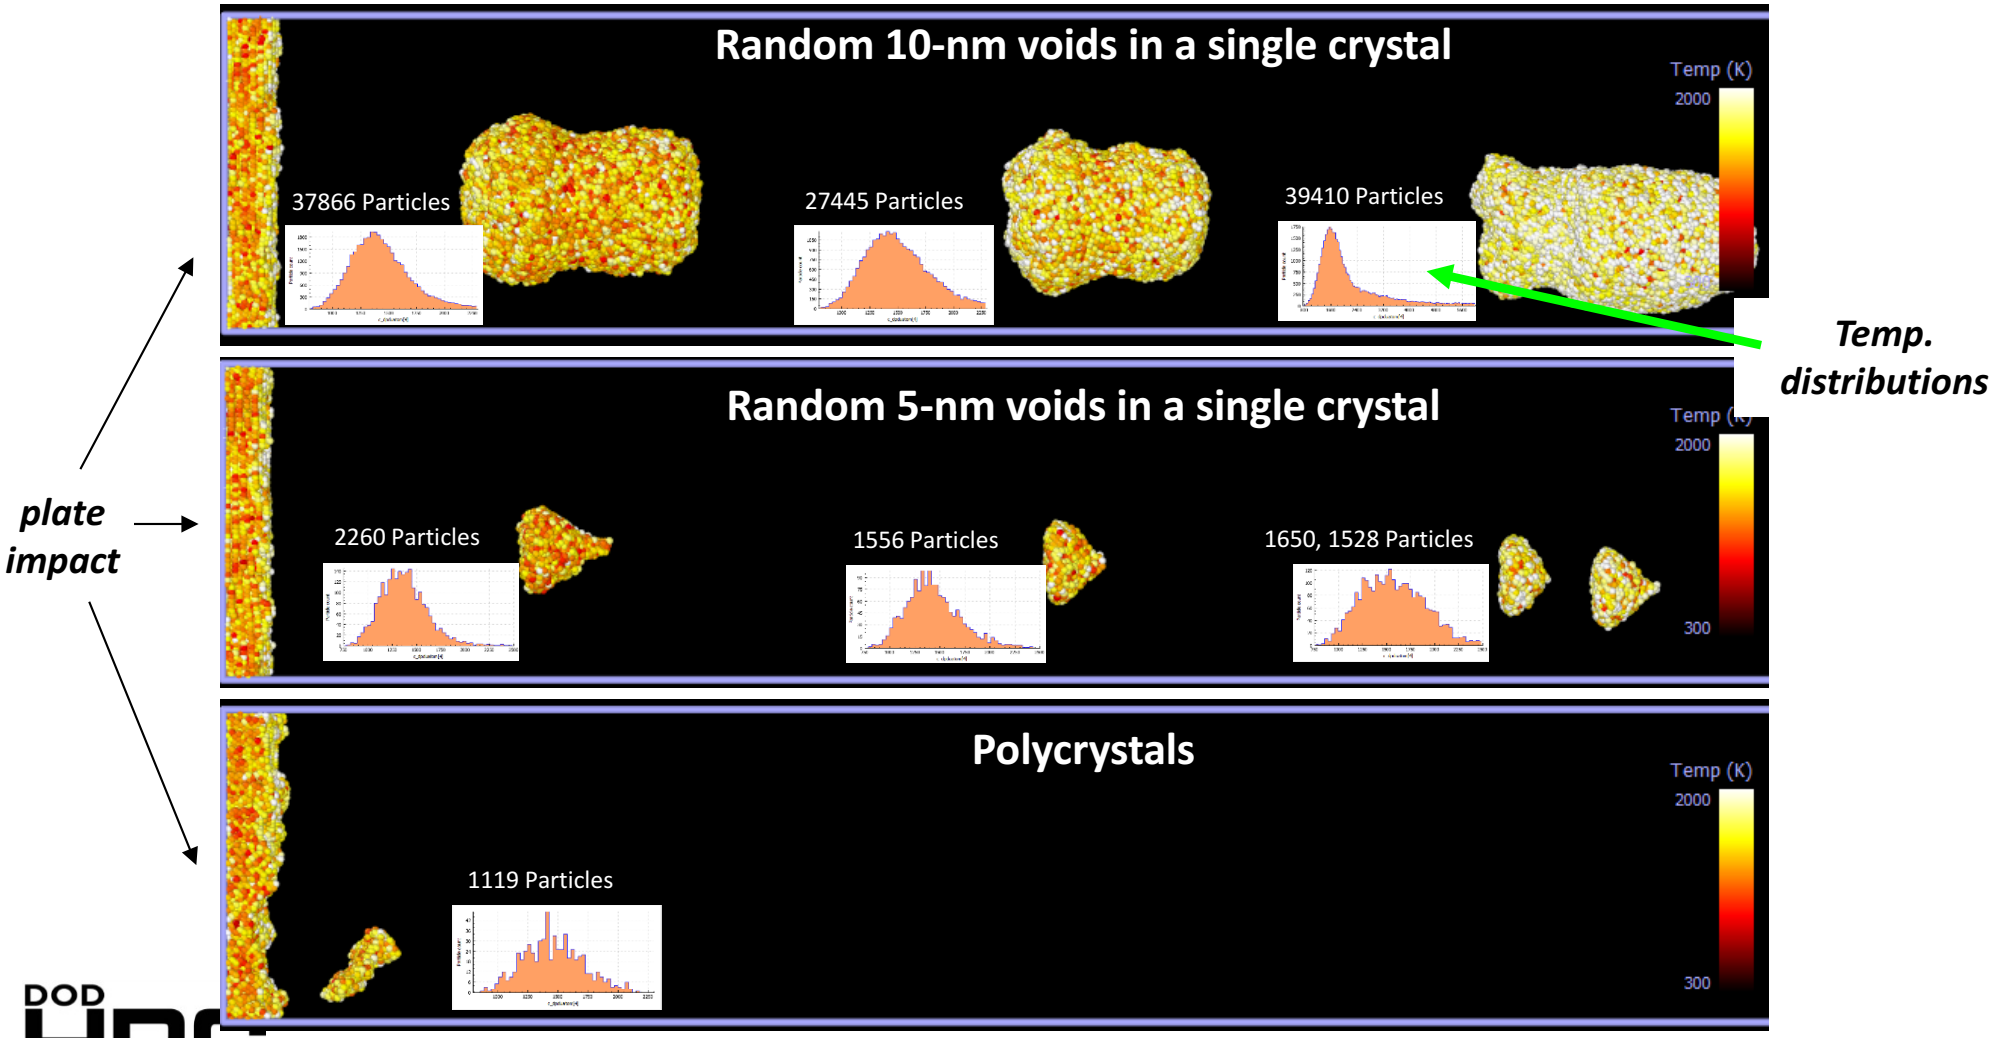

USER-DPD Scaling on Thunder