Exploiting High Performance Computing Resources to Drive Time-Parallelized Molecular Dynamics in a GPU-Dominated World

Peter Hatton[1], Danny Perez[2], Blas Uberuaga[1]

1 Materials Science and Technology Division, Los Alamos National Laboratory, Los Alamos, NM
2 Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM

CNLS Seminar
13th November 2023

Outline

1. Introduction to Molecular Dynamics (MD) and Spatial Parallelization

2. Introduction to Time-Parallelized (Accelerated) MD

3. An Example of Helium Bubbles in Tungsten

4. Increasing Efficiency of Time-Parallelized MD on GPUs

Molecular Dynamics (MD) - LAMMPS

Large-scale Atomic/Molecular Massively Parallel Simulation (LAMMPS)

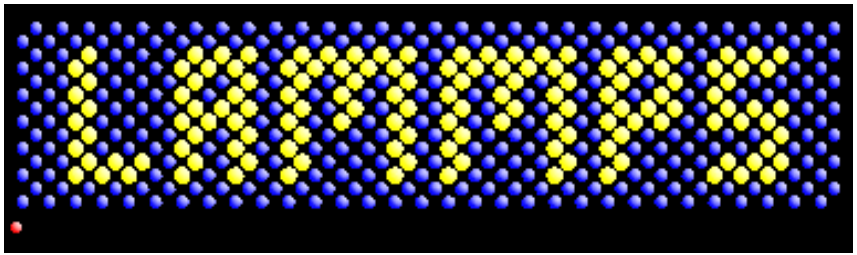

Molecular Dynamics (MD) - LAMMPS

Atomic Configuration

What is the initial positions of my atoms?

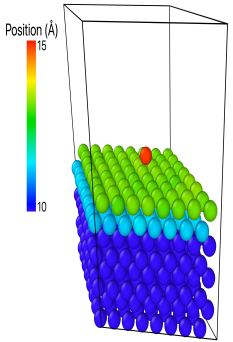

Large-scale Atomic/Molecular Massively Parallel Simulation (LAMMPS)

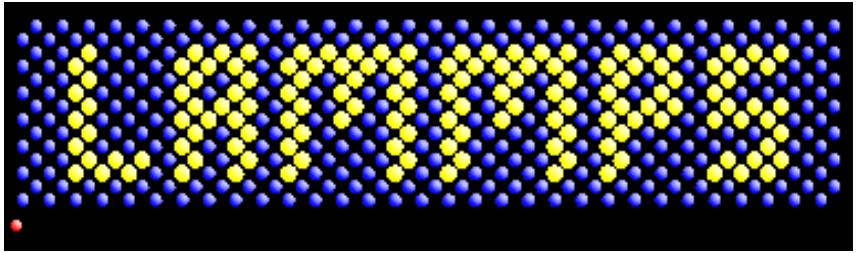

Molecular Dynamics (MD) - LAMMPS

Atomic Configuration

What is the initial positions of my atoms?

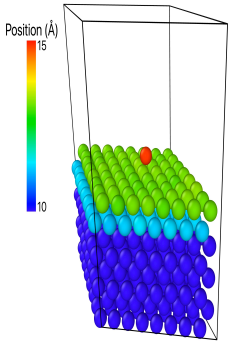

Large-scale Atomic/Molecular Massively Parallel Simulation (LAMMPS)

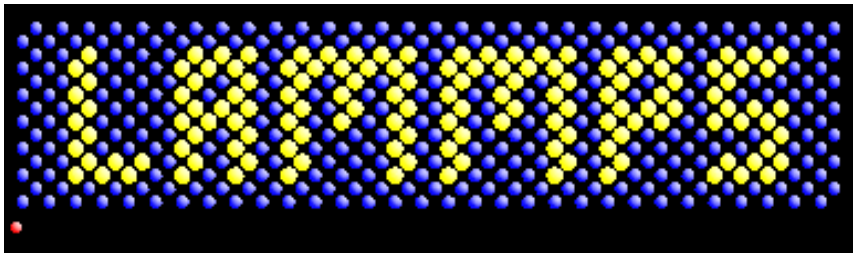

Interatomic Potential

How do the atoms interact?

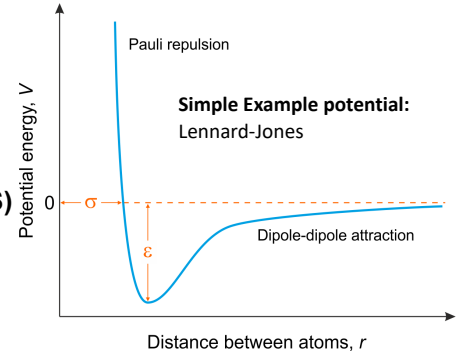

Molecular Dynamics (MD) - LAMMPS

Atomic Configuration

What is the initial positions of my atoms?

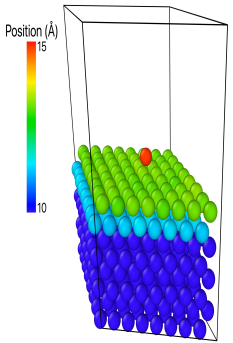

Interatomic Potential

How do the atoms interact?

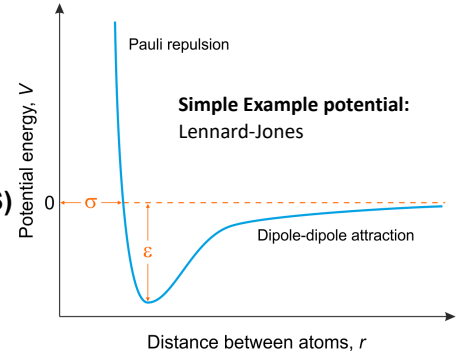

Large-scale Atomic/Molecular Massively Parallel Simulation (LAMMPS)

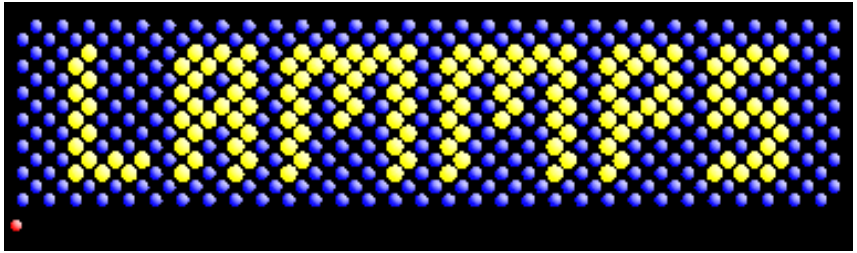

Simulation Parameters

Periodic Boundary Conditions? Fixed Atoms? …

Molecular Dynamics (MD) - LAMMPS

Atomic Configuration

What is the initial positions of my atoms?

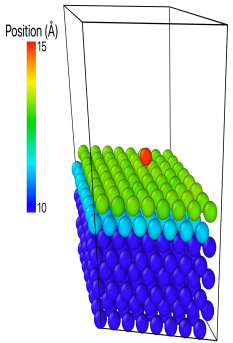

Interatomic Potential

How do the atoms interact?

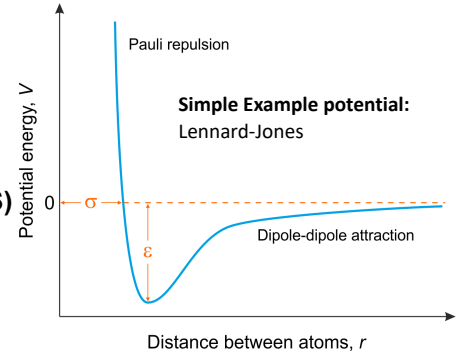

Large-scale Atomic/Molecular Massively Parallel Simulation (LAMMPS)

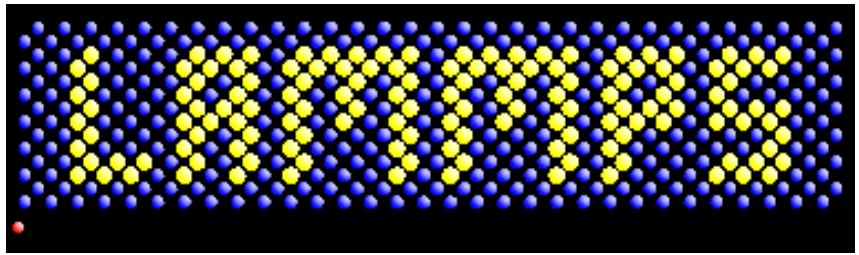

Temperature Thermostat

Thermostatics control the temperature of simulations

Simulation Parameters

Periodic Boundary Conditions? Fixed Atoms? …

Molecular Dynamics (MD) - LAMMPS

Atomic Configuration

What is the initial positions of my atoms?

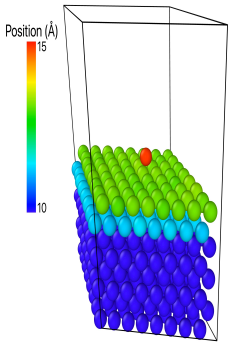

Interatomic Potential

How do the atoms interact?

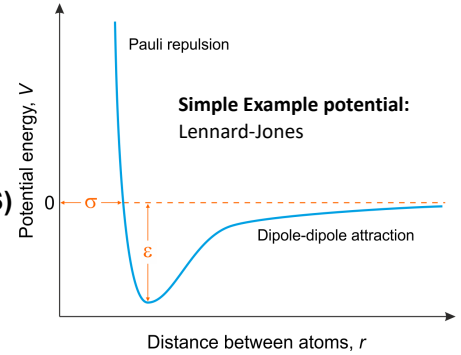

Large-scale Atomic/Molecular Massively Parallel Simulation (LAMMPS)

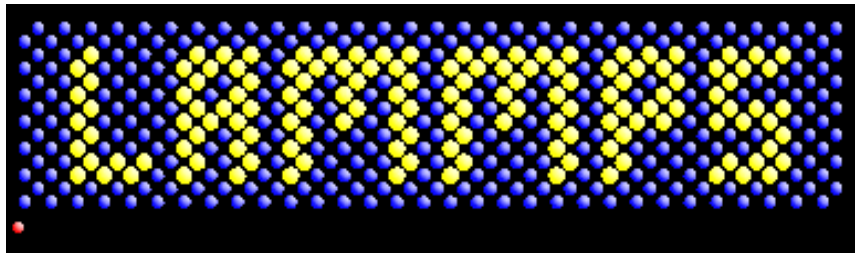

Temperature Thermostat

Thermostatics control the temperature of simulations

Dynamics

Simulation Parameters

Periodic Boundary Conditions? Fixed Atoms? …

Molecular Dynamics (MD) - LAMMPS

Atomic Configuration

What is the initial positions of my atoms?

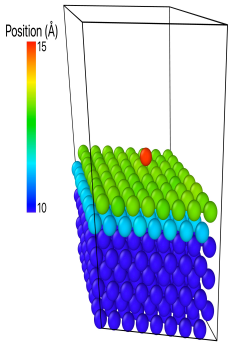

Interatomic Potential

How do the atoms interact?

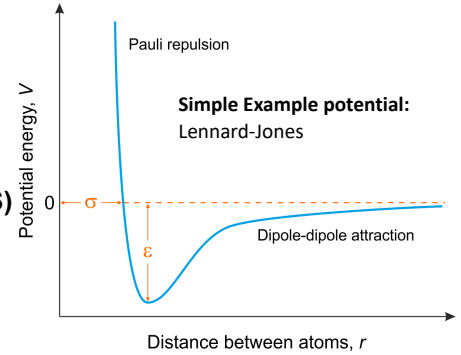

Large-scale Atomic/Molecular Massively Parallel Simulation (LAMMPS)

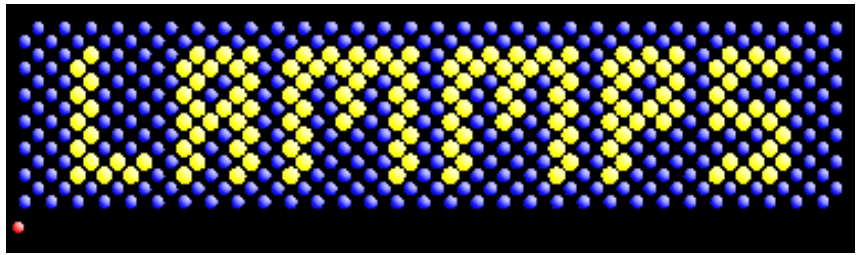

Temperature Thermostat

Thermostatics control the temperature of simulations

Dynamics

Simulation Parameters

Periodic Boundary Conditions? Fixed Atoms? …

Garbage

Gold

Molecular Dynamics (MD) – Parallelization

Standard Parallelization Method

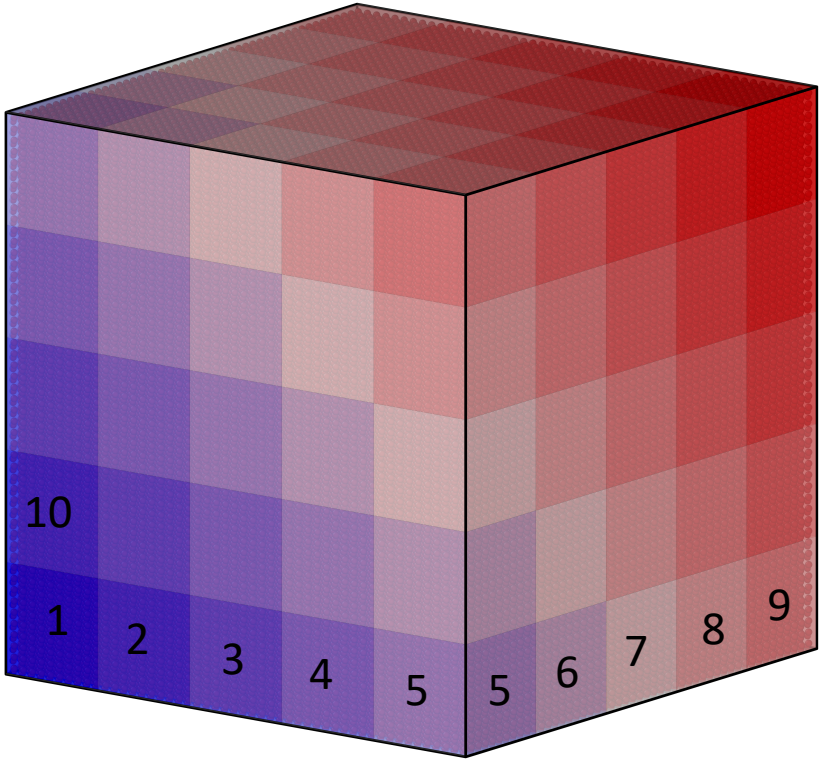

Molecular Dynamics (MD) – Parallelization

Standard Parallelization Method

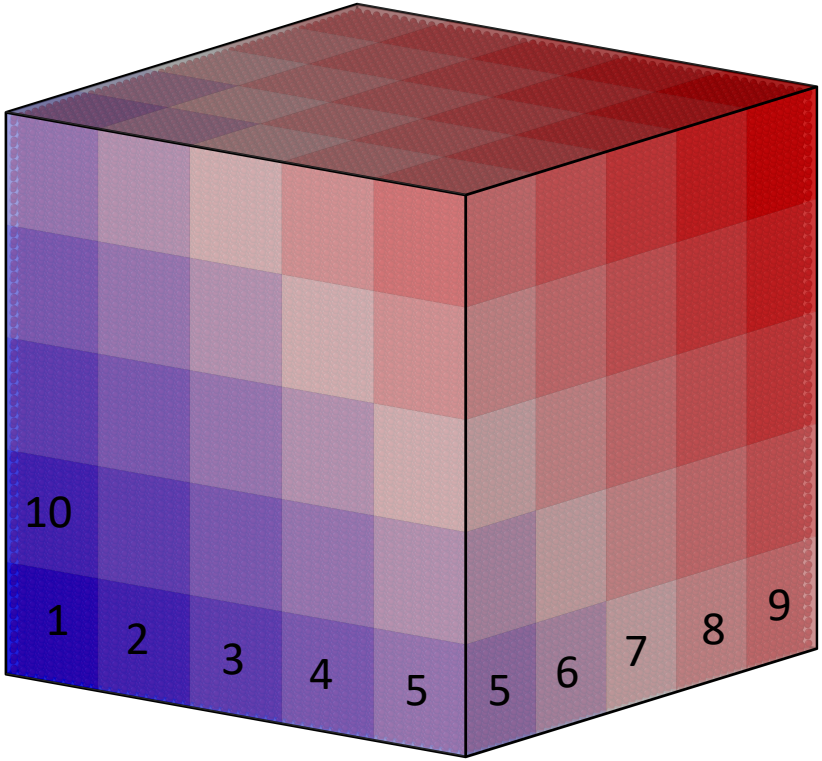

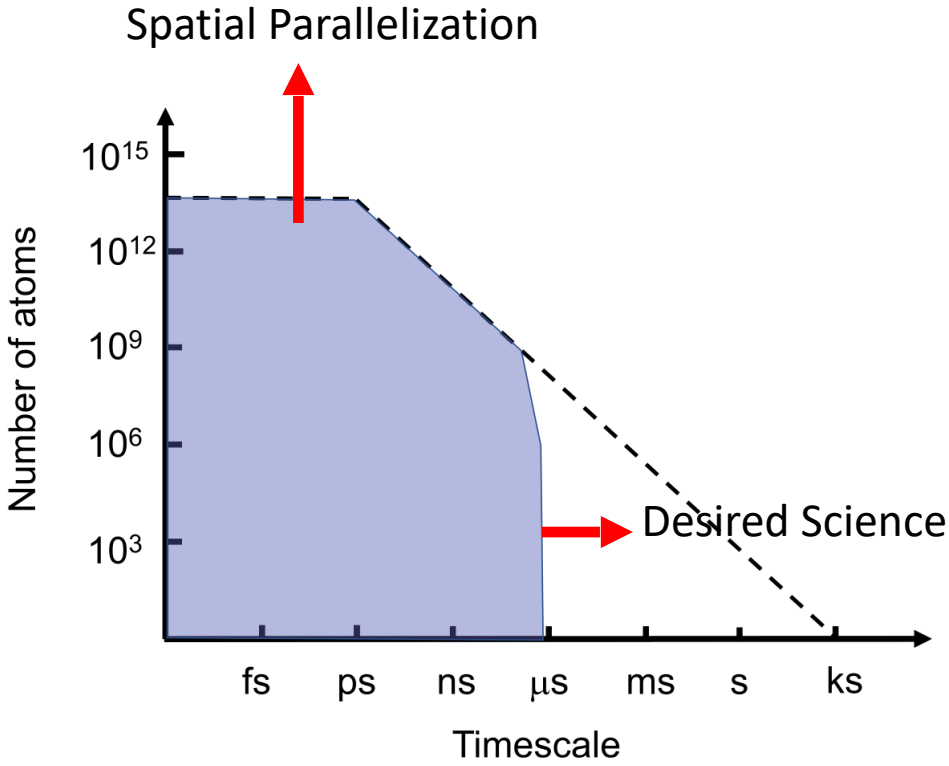

Time-Parallel Algorithms – The Idea

‘Continuous’ Picture

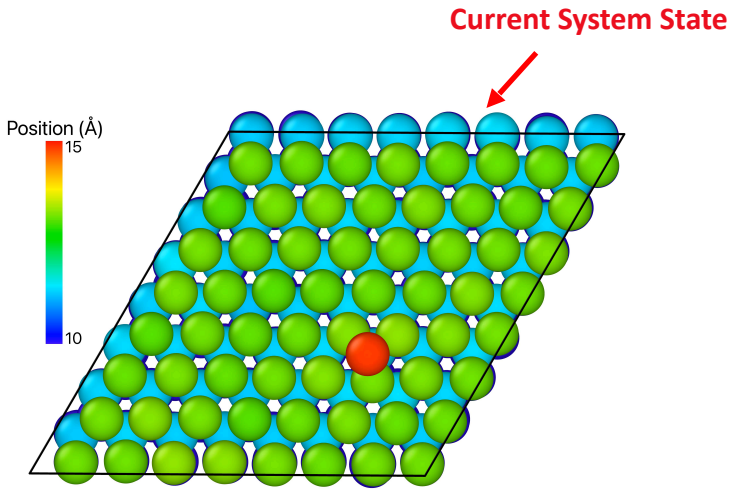

We don’t care about ‘in-state’ vibrations that don’t advance the dynamics

We only care when we ‘hop’ to an adjacent minima – these are mechanisms which advance the dynamics

Time-Parallel Algorithms – The Idea

‘Continuous’ Picture

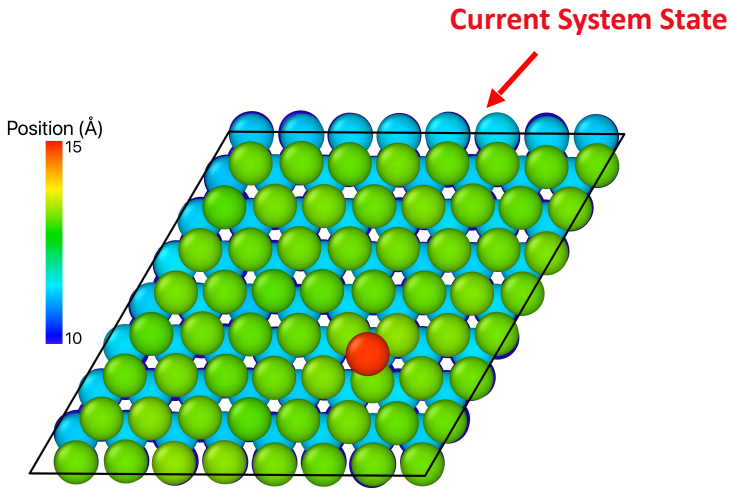

Discrete Picture

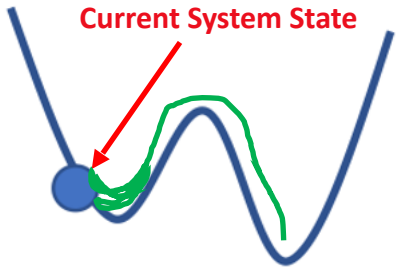

Energy Landscape - 2D Example

We don’t care about ‘in-state’ vibrations that don’t advance the dynamics

We only care when we ‘hop’ to an adjacent minima – these are mechanisms which advance the dynamics

Time-Parallel Algorithms – The Idea

In a regime where we are trapped in a minima configuration for a ‘long-time’ before we ’hop’ to a new state we can partition the trajectory that is just vibrating ’in-state’ onto many workers.

Time-Parallel Algorithms – The Idea

In a regime where we are trapped in a minima configuration for a ‘long-time’ before we ’hop’ to a new state we can partition the trajectory that is just vibrating ’in-state’ onto many workers.

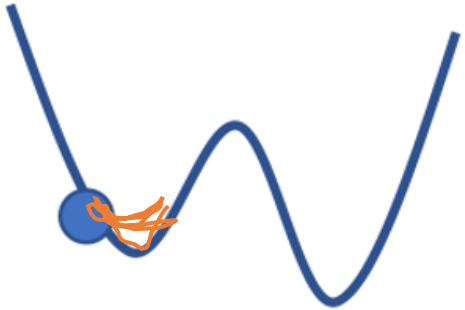

Time-Parallel Algorithms – The Idea

In a regime where we are trapped in a minima configuration for a ‘long-time’ before we ’hop’ to a new state we can partition the trajectory that is just vibrating ’in-state’ onto many workers.

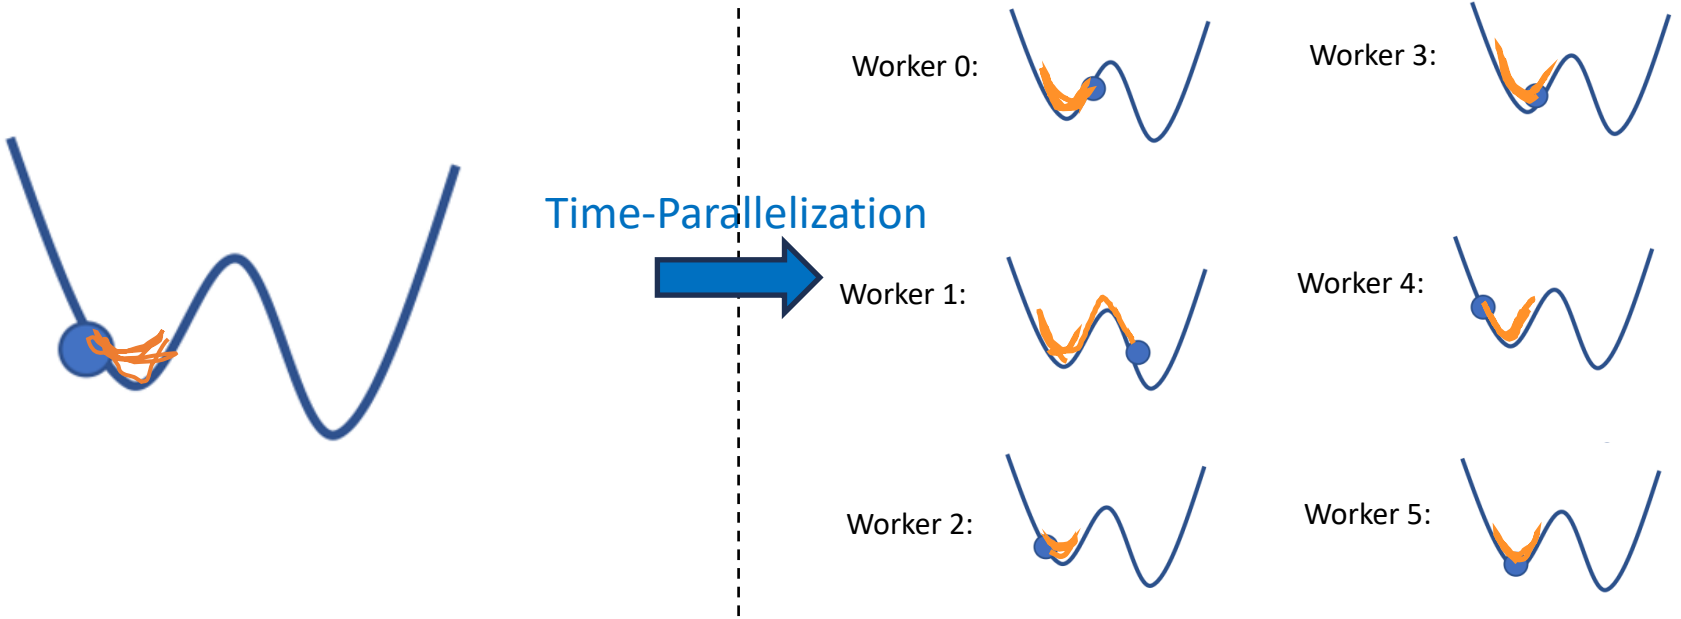

Time-Parallel Algorithms – The Idea

In a regime where we are trapped in a minima configuration for a ‘long-time’ before we ’hop’ to a new state we can partition the trajectory that is just vibrating ’in-state’ onto many workers.

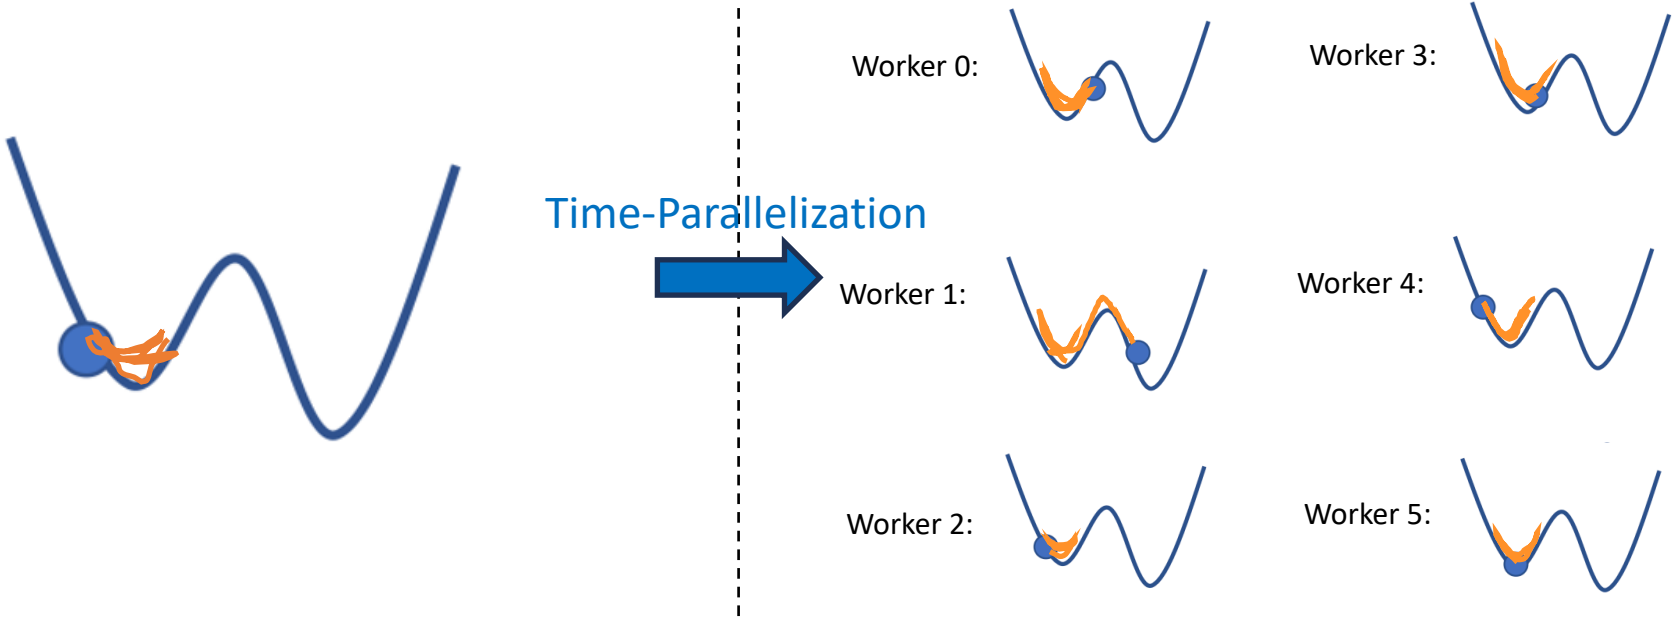

In this case we will find a transition to a new state ~6 times faster

Time-Parallel Algorithms - Example

Parallel Replica Dynamics (PRD):

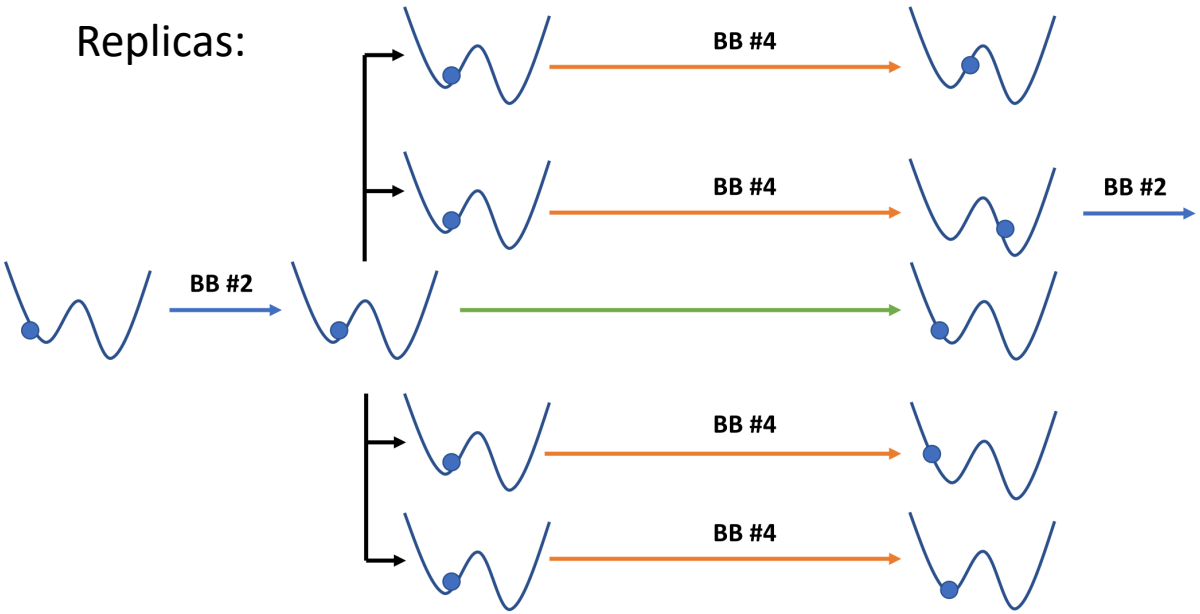

Time-Parallel Algorithms - Example

Parallel Replica Dynamics (PRD):

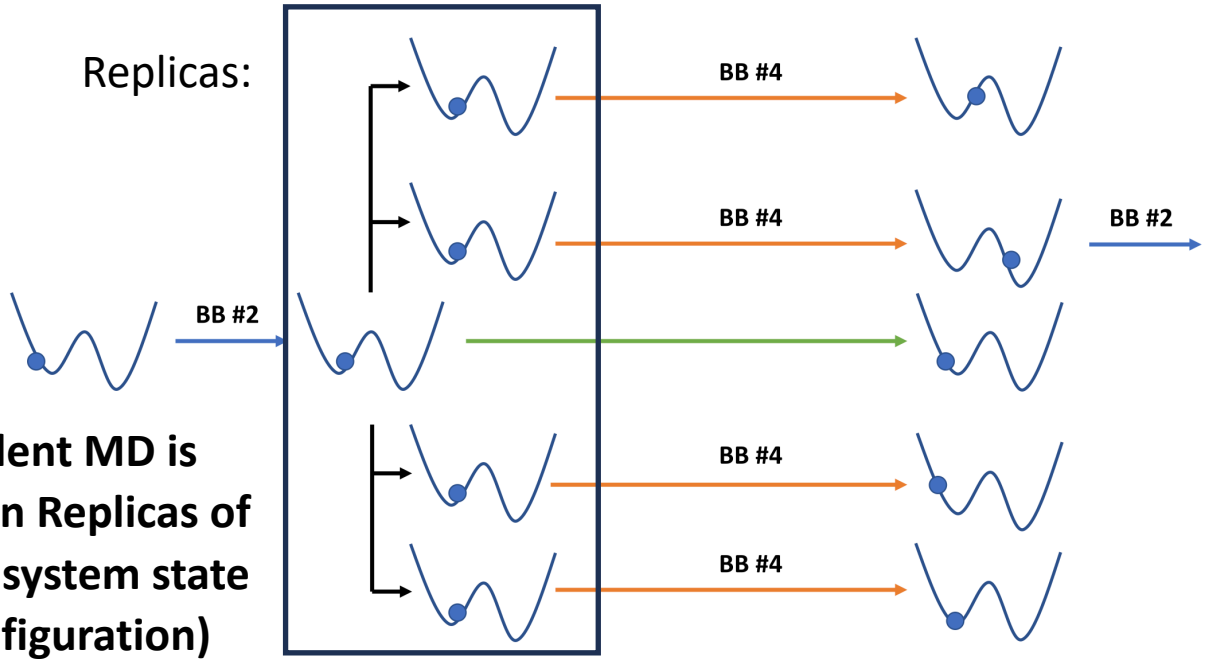

1: Independent MD is conducted in Replicas of the current system state (atomic configuration) on ‘workers’

Time-Parallel Algorithms - Example

Parallel Replica Dynamics (PRD):

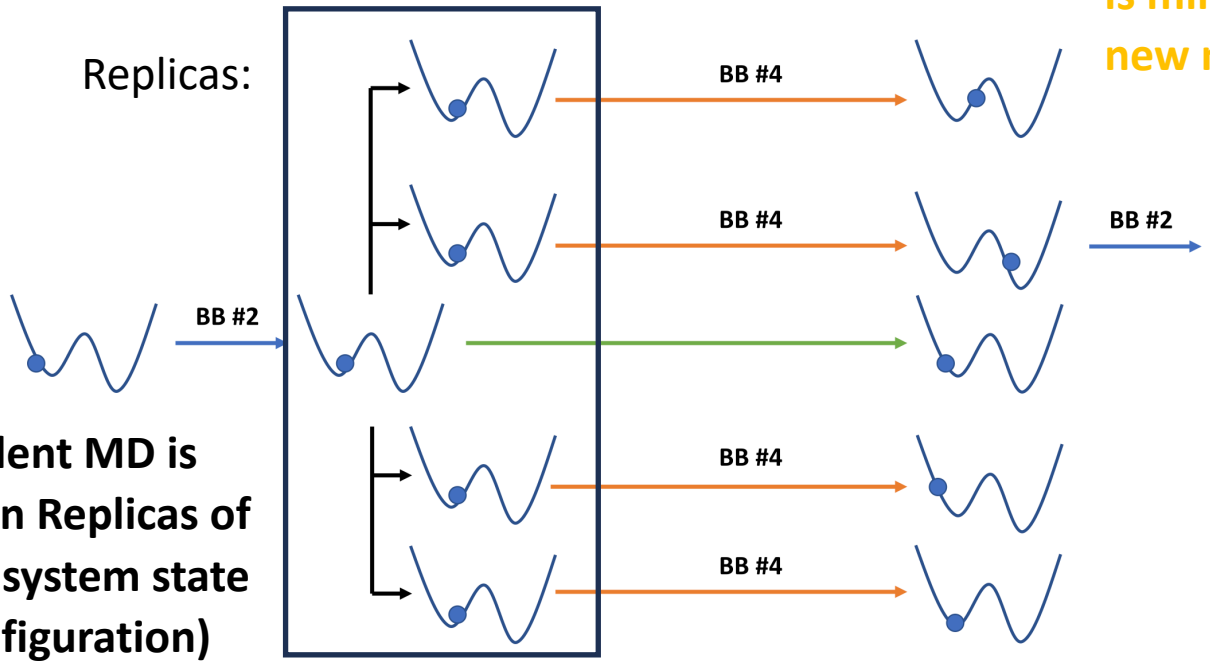

1: Independent MD is conducted in Replicas of the current system state (atomic configuration) on 'workers'

2: When a worker finds a transition, the system is minimized into that new minima.

Time-Parallel Algorithms - Example

Parallel Replica Dynamics (PRD):

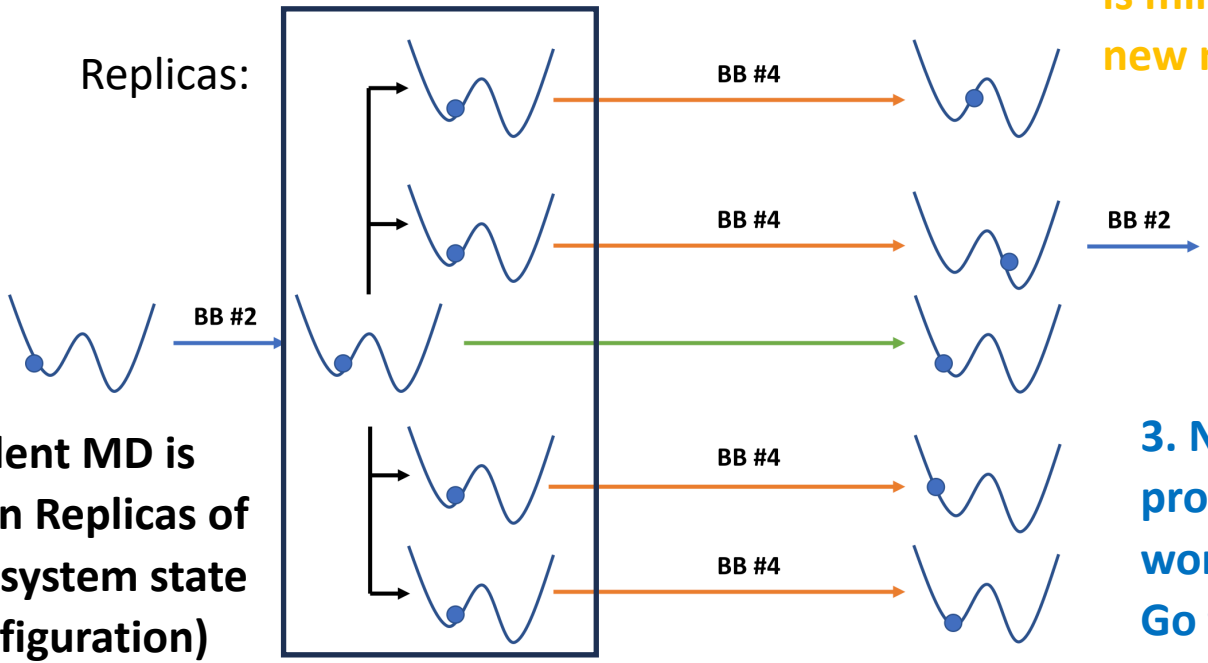

1: Independent MD is conducted in Replicas of the current system state (atomic configuration) on 'workers'

2: When a worker finds a transition, the system is minimized into that new minima.

3. New configuration is propagated to all of our workers.
Go to step 1.

Time-Parallel Algorithms

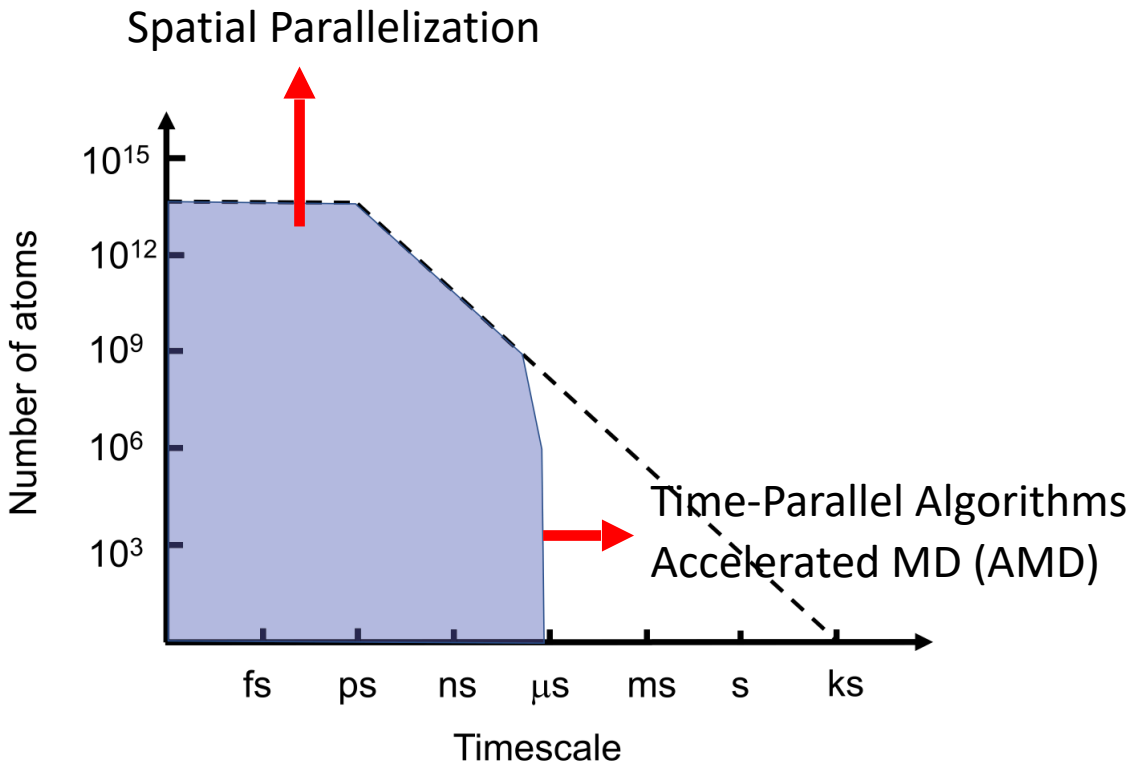

Los Alamos
NATIONAL LABORATORY

Some Materials Science…

Tungsten Divertor Damage

Tokamak

(**To**roidal **cha**mber with **ma**gnetic **c**oils)

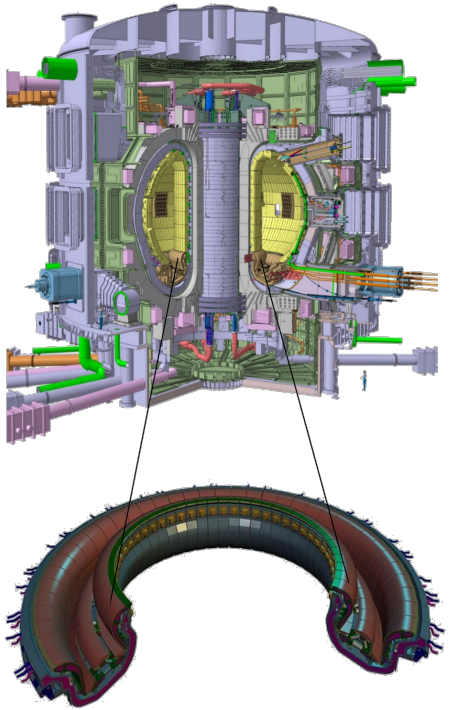

Divertor

Tungsten Divertor Damage

Tokamak

(**To**roidal **cha**mber with **ma**gnetic **c**oils)

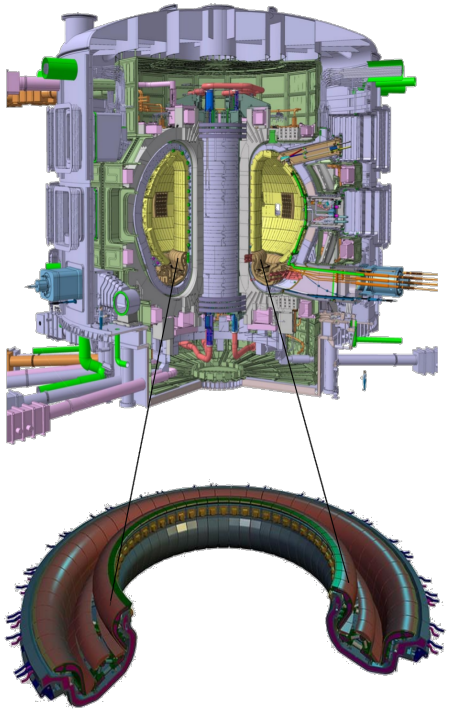

Divertor

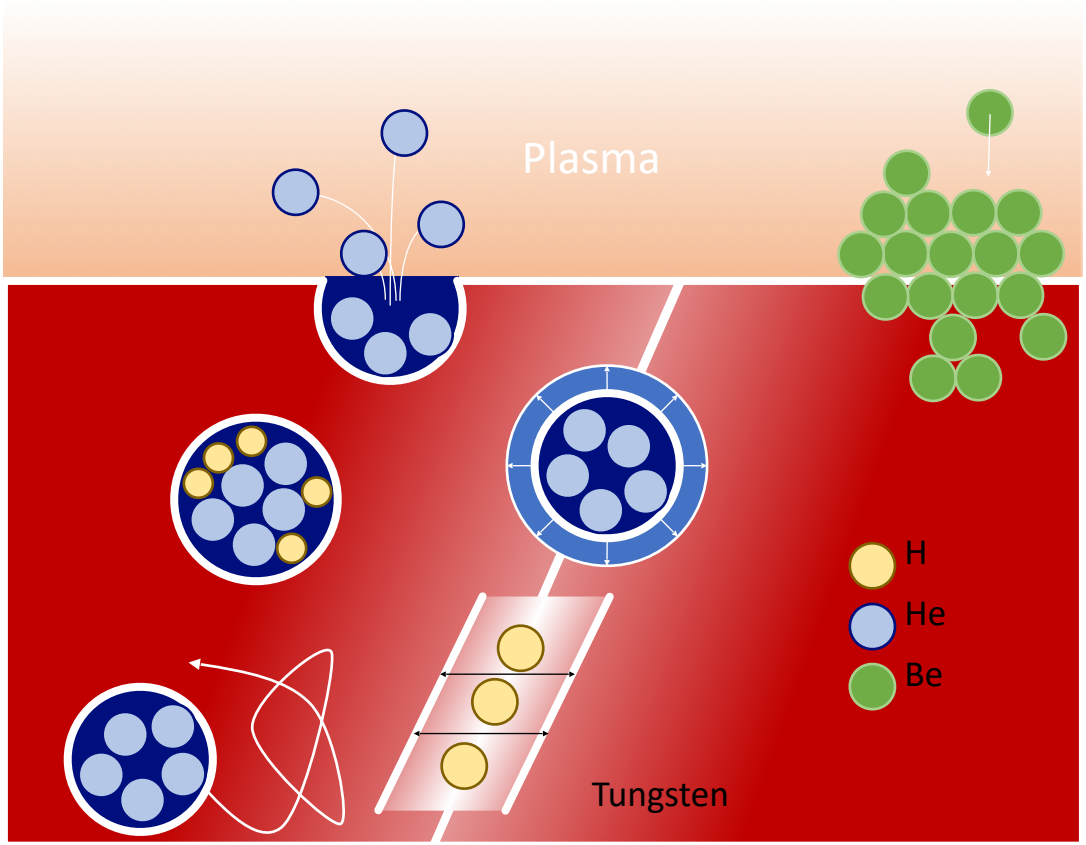

He Bubbles in Tungsten

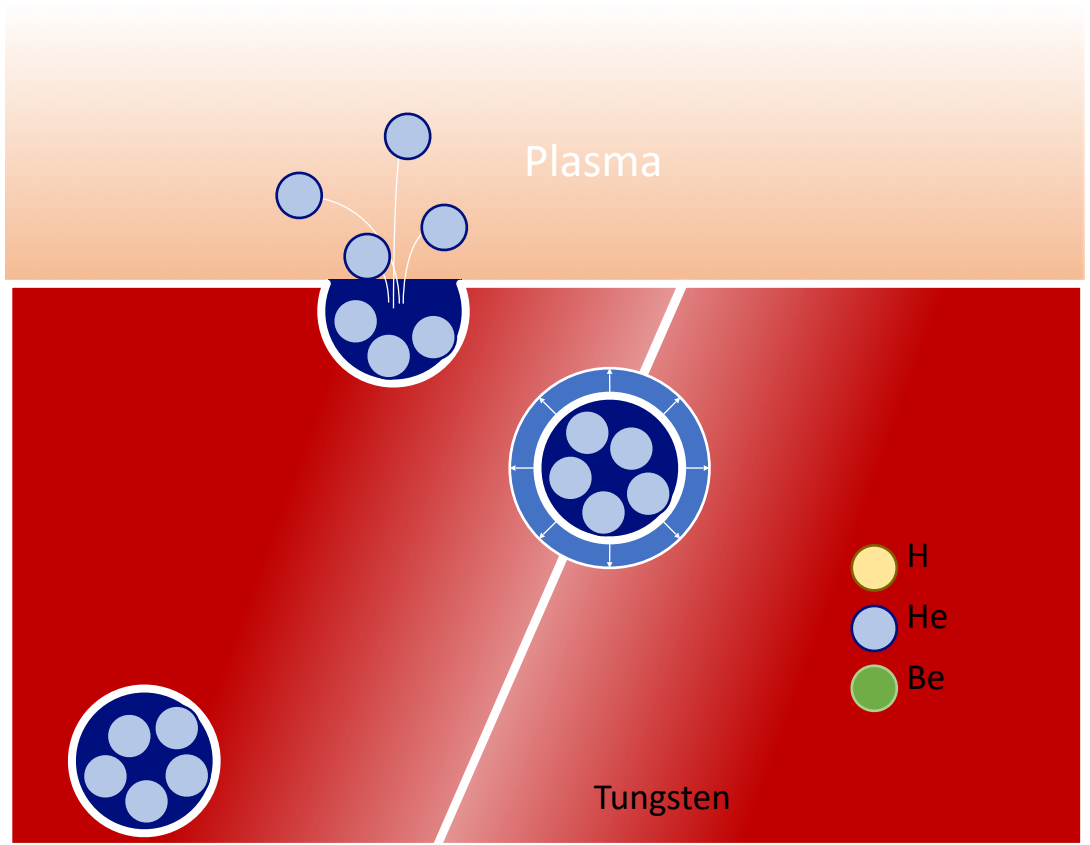

He Bubbles in Tungsten

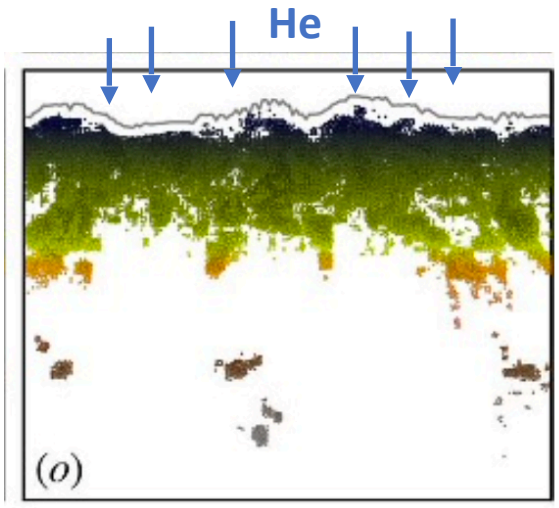

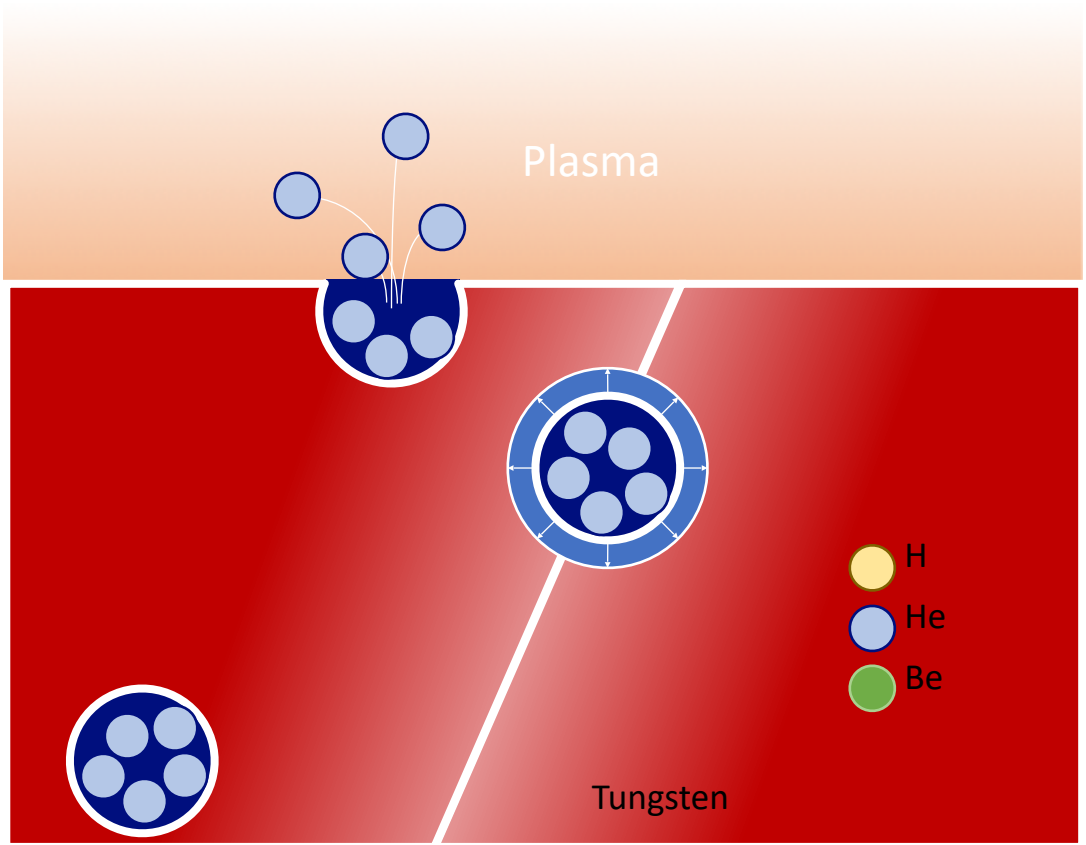

He Bubbles in Tungsten

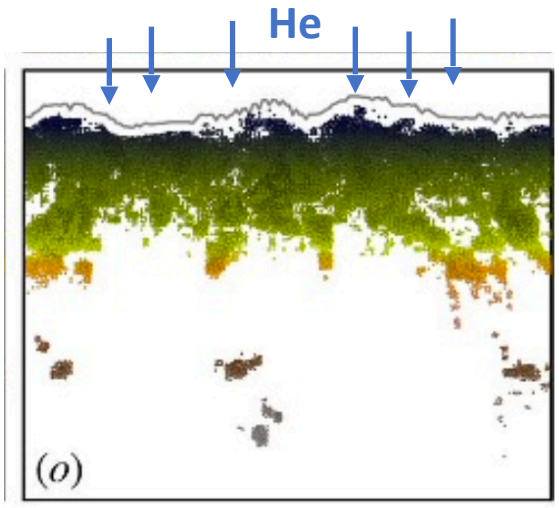

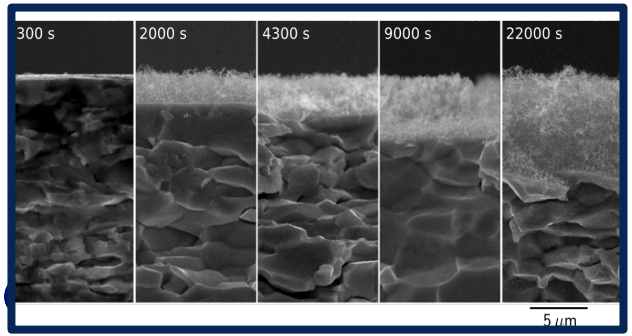

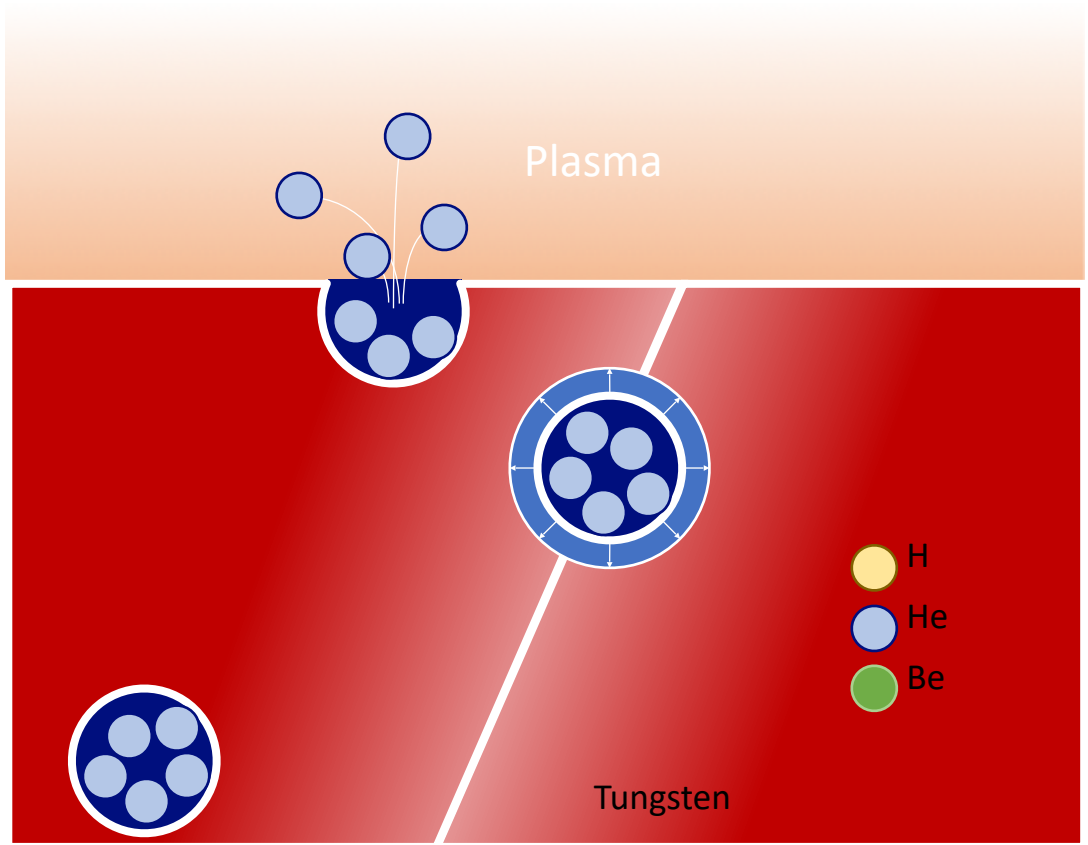

He Bubbles in Tungsten – Bulk W

Accelerated MD used to simulate He Bubble Growth in Tungsten.

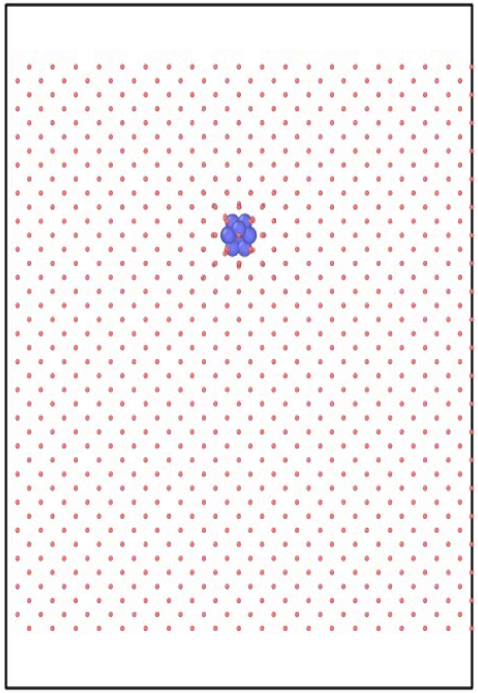

He Bubbles in Tungsten – Bulk W

Accelerated MD used to simulate He Bubble Growth in Tungsten.

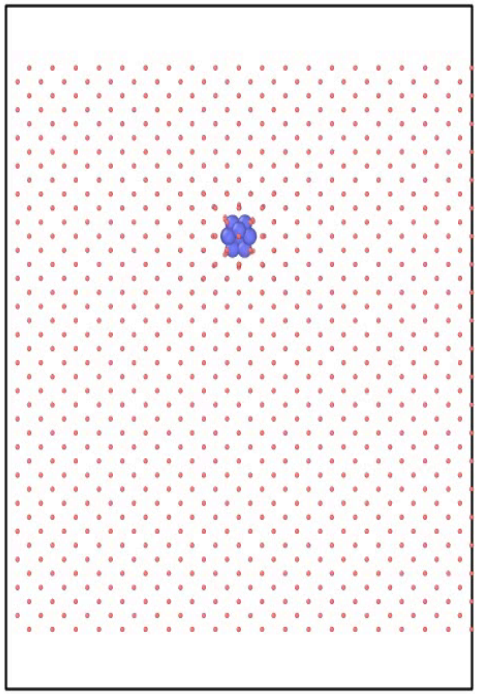

Bubble pressure controls loop punching and therefore controls surface morphology changes.

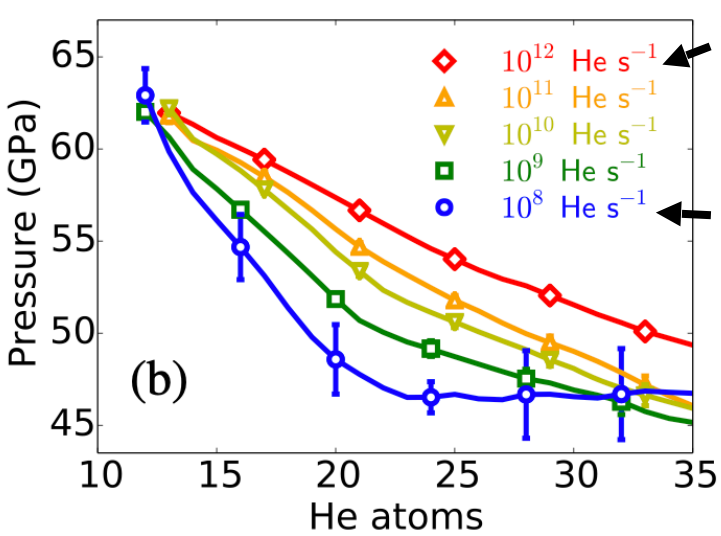

He Bubbles in Tungsten – W GB

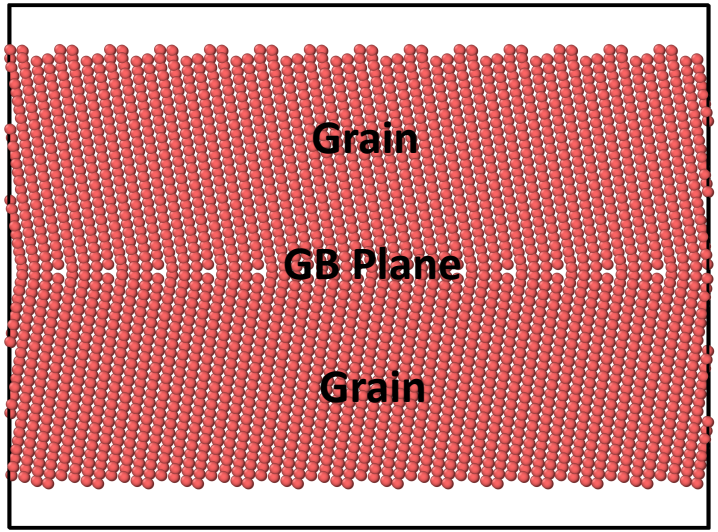

Grain Boundaries (GBs) are strong energetic sinks for defects.

He Bubbles are shown experimentally to concentrate at GBs.

P. Hatton *et al.*, MRS Comms. (2022)

He Bubbles in Tungsten – W GB

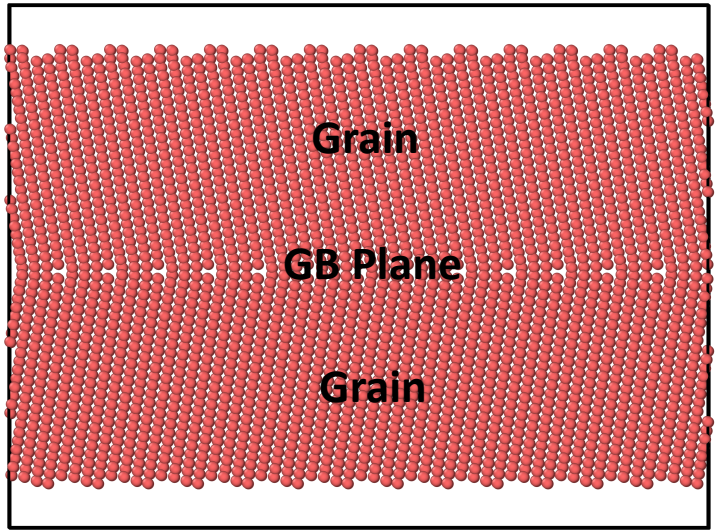

Blue: He, Green: W Interstitials, Orange: W vacancy.

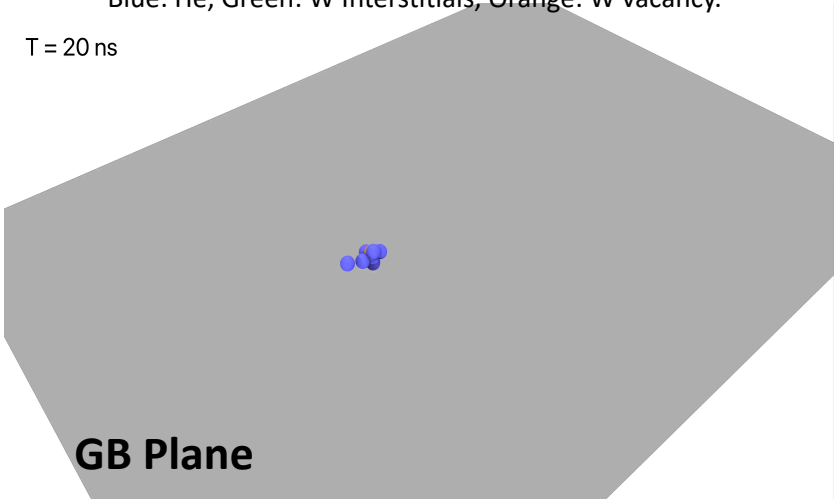

Grain Boundaries (GBs) are strong energetic sinks for defects.

He Bubbles are shown experimentally to concentrate at GBs.

50K atom supercell
400 A100 GPUs with one replica per GPU
He bubble manually grown at rate of 1 He / 10 ns.
With conventional MD we cap at 15 ns / day.
We want a bubble of size ~ 500 He atoms…

P. Hatton *et al.*, MRS Comms. (2022)

He Bubbles in Tungsten – W GB

t = 6.66 μs

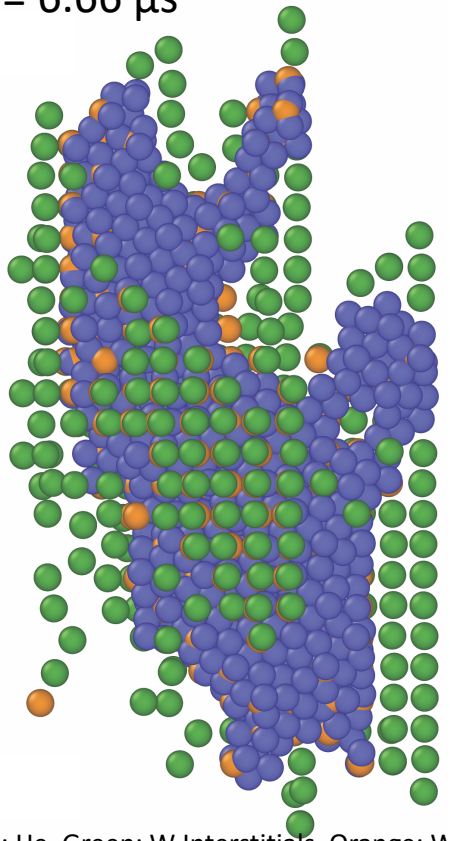

Blue: He, Green: W Interstitials, Orange: W vacancy.

He Bubbles in Tungsten – W GB

t = 6.66 μs

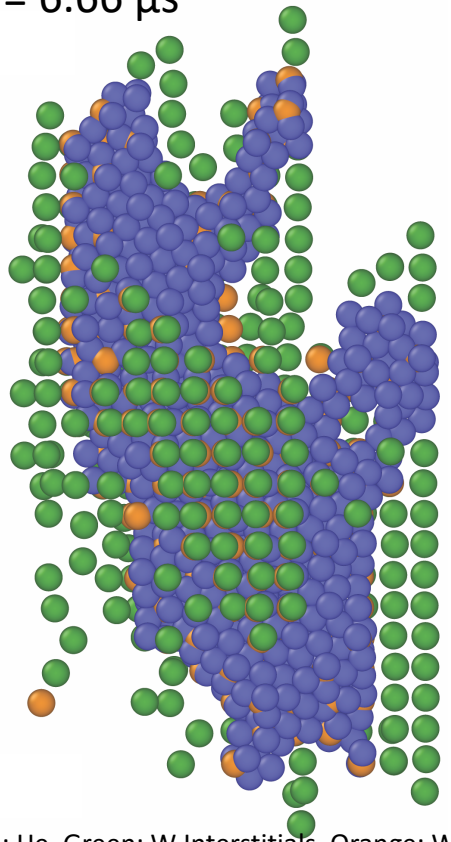

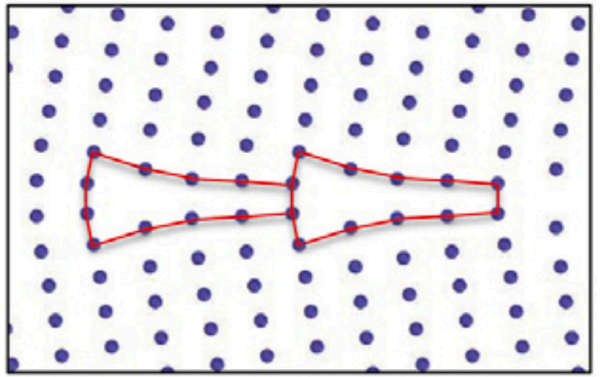

Blue: He, Green: W Interstitials, Orange: W vacancy.

He Bubbles in Tungsten – W GB

t = 6.66 μs

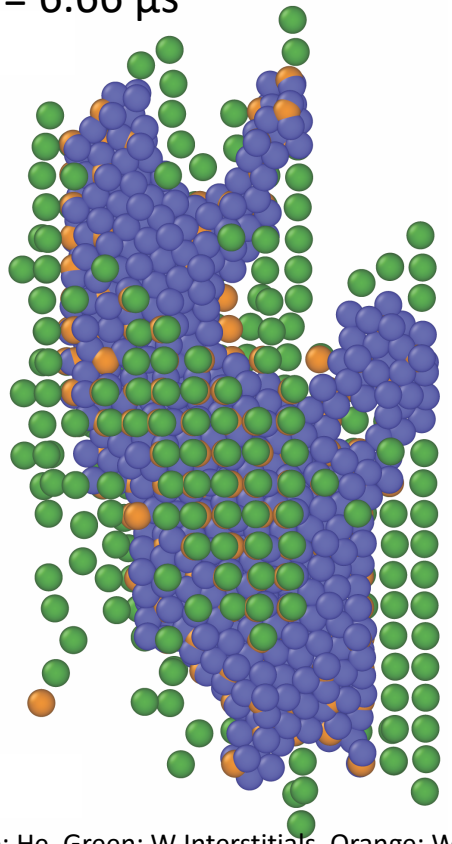

Blue: He, Green: W Interstitials, Orange: W vacancy.

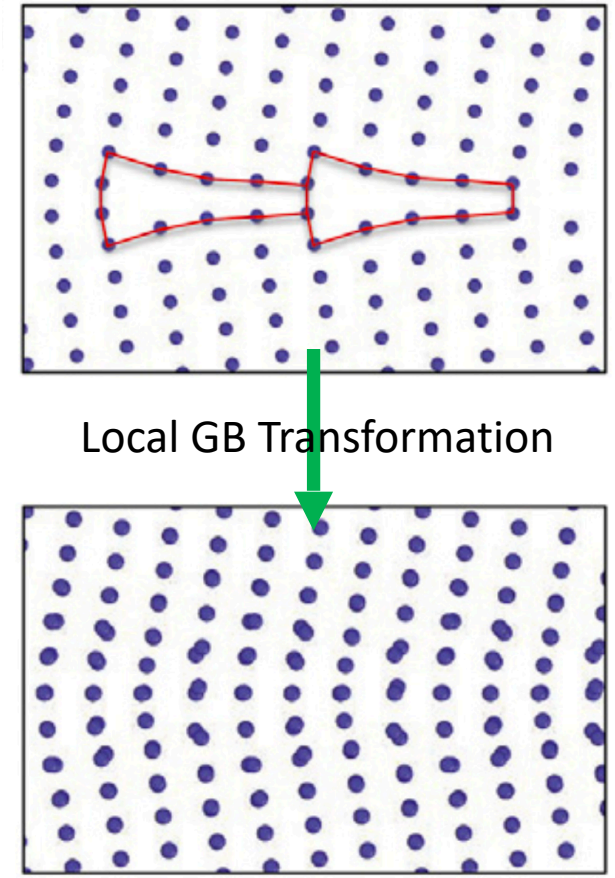

He Bubbles in Tungsten – W GB

t = 6.66 μs

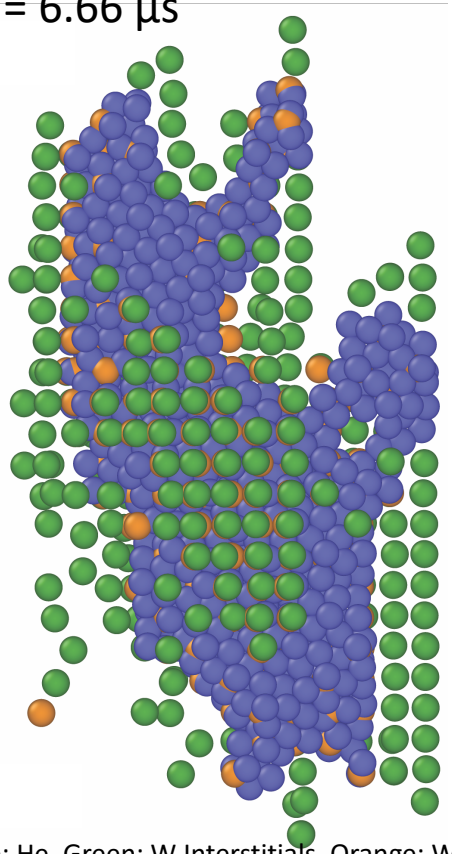

Blue: He, Green: W Interstitials, Orange: W vacancy.

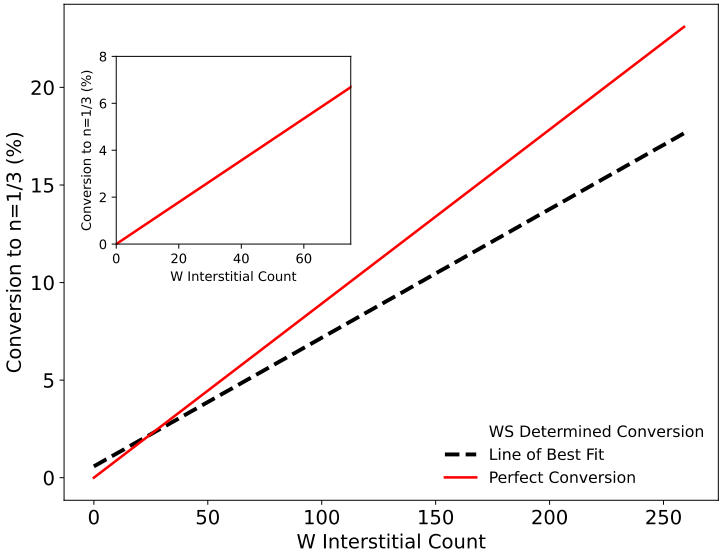

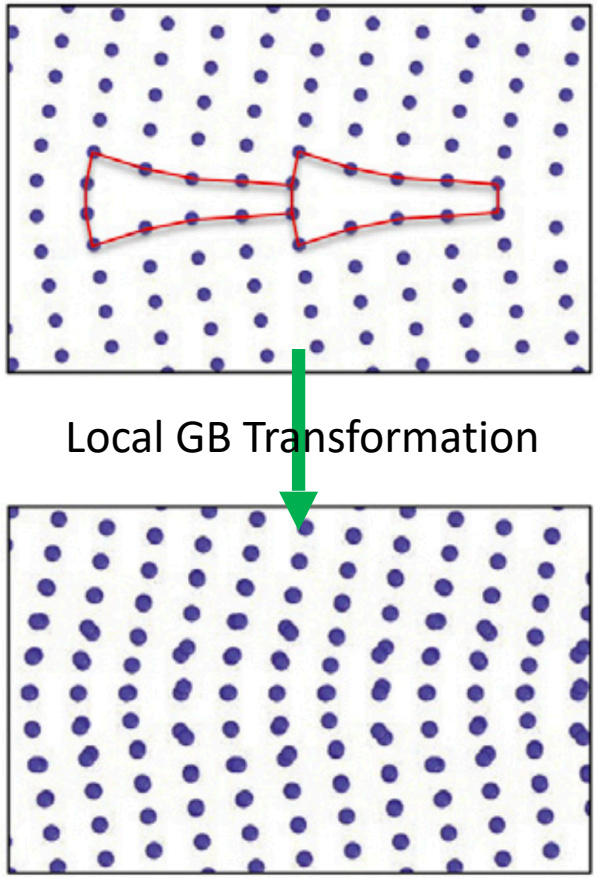

He Bubbles in Tungsten – W GB

t = 6.66 μs

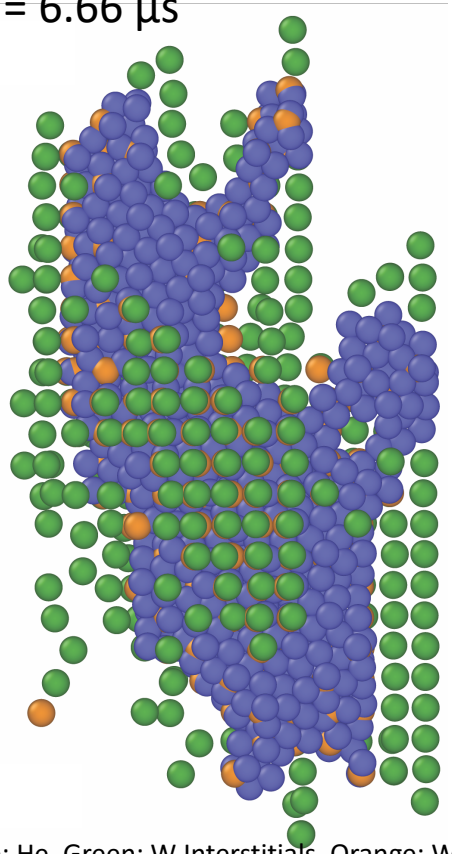

Blue: He, Green: W Interstitials, Orange: W vacancy.

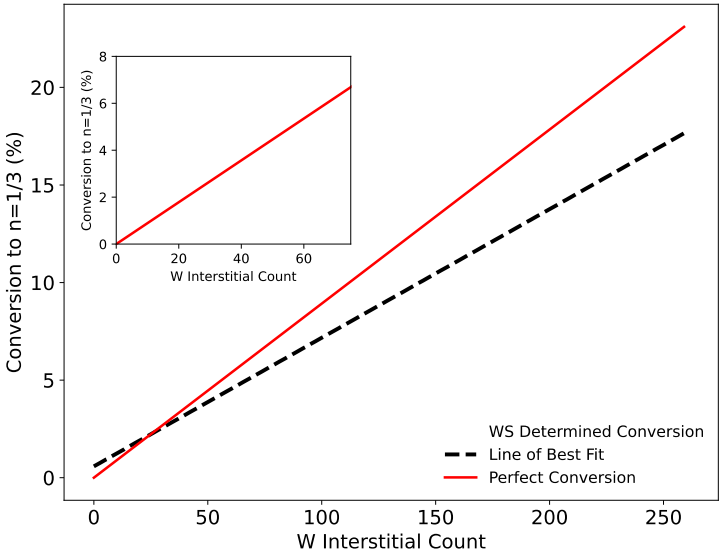

W Interstitials locally convert the GB to a higher density stable structure

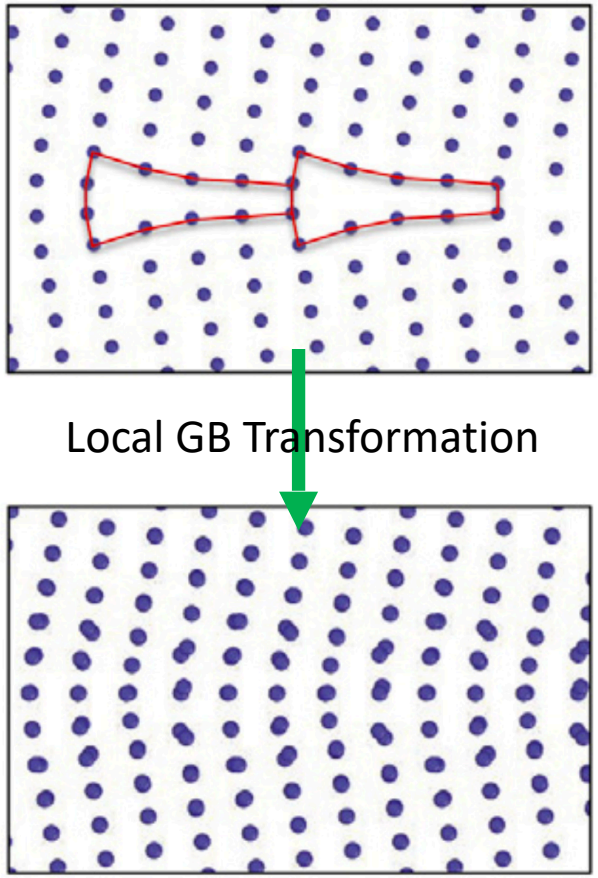

He Bubbles in Tungsten – W GB

1 He/ 10ns

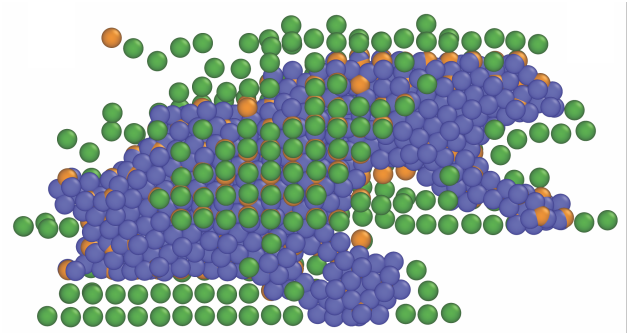

He Bubbles in Tungsten – W GB

1 He/ 10ns

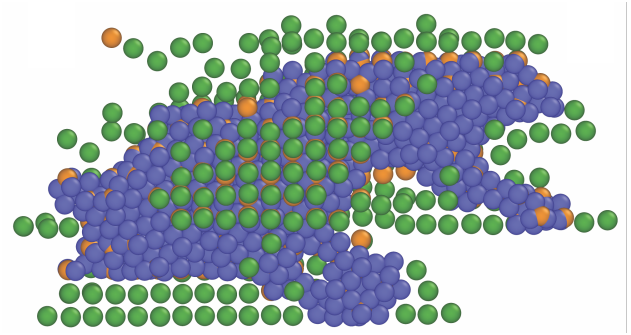

1 He/ 10ps

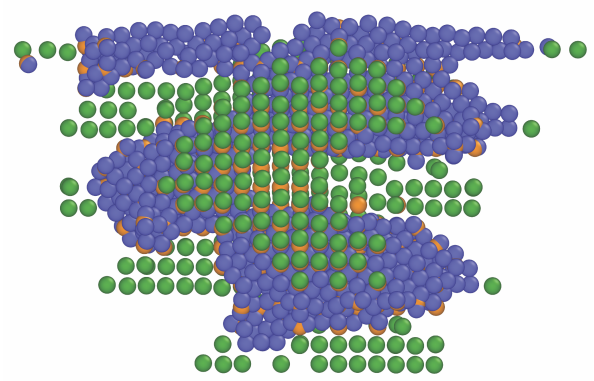

He Bubbles in Tungsten – W GB

1 He/ 10ns

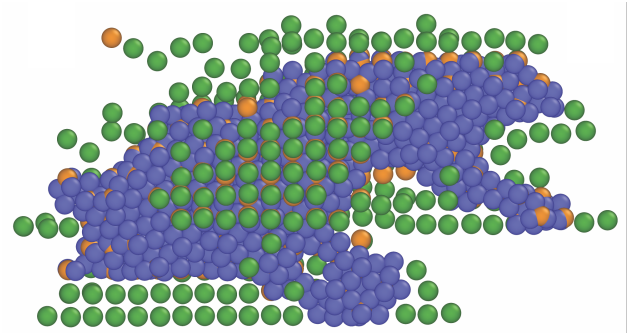

1 He/ 10ps

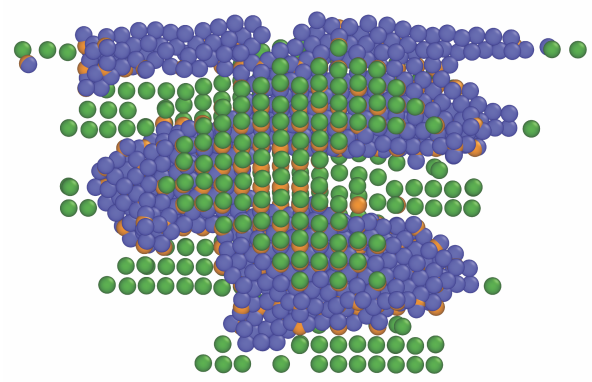

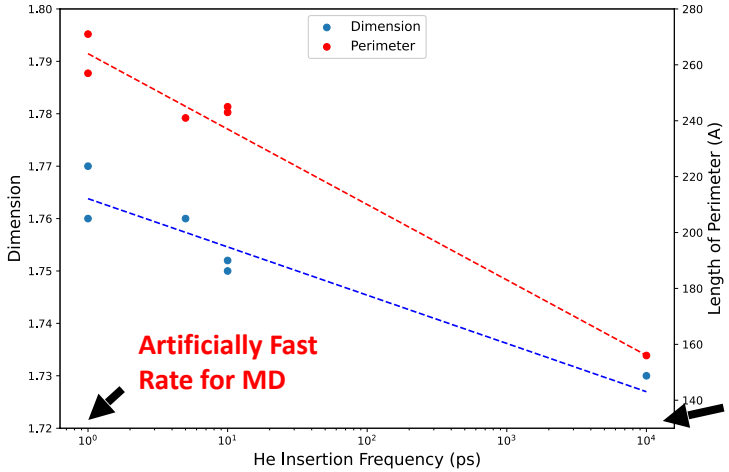

'Realistic' Slow Rate
unlocked by AMD

He Bubbles in Tungsten – W GB

1 He/ 10ns

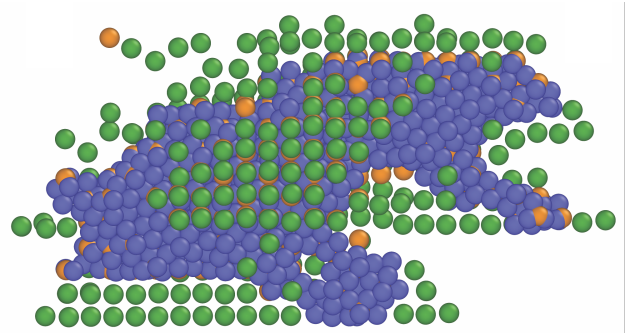

1 He/ 10ps

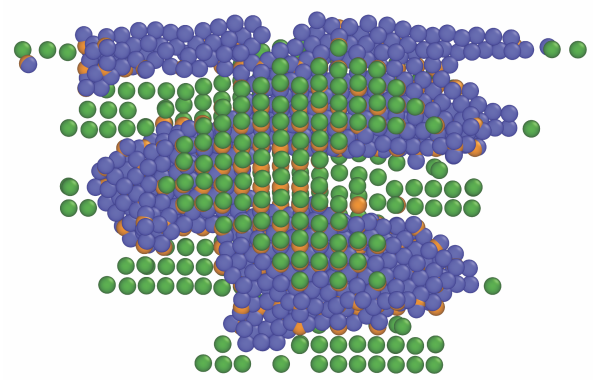

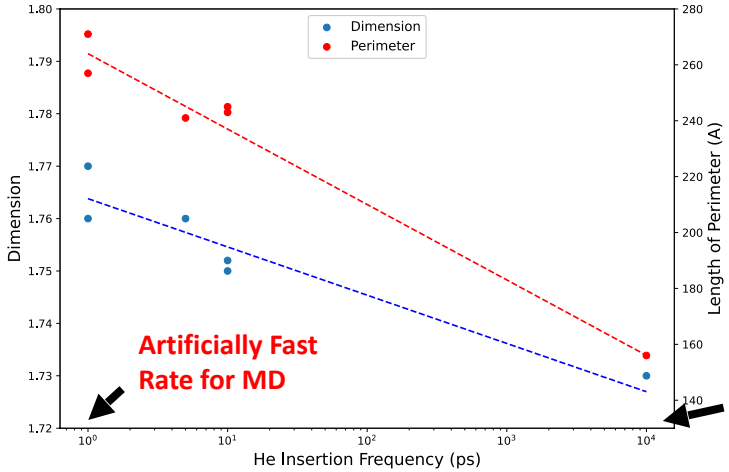

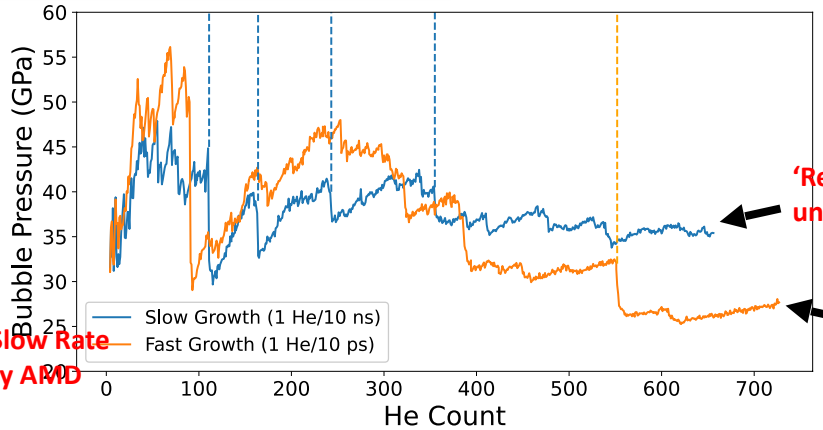

He Bubbles in Tungsten – W GB

1 He/ 10ns

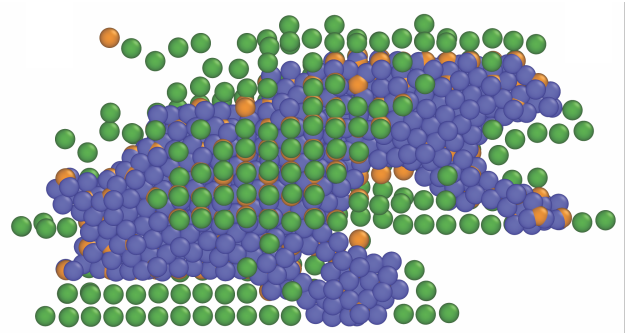

1 He/ 10ps

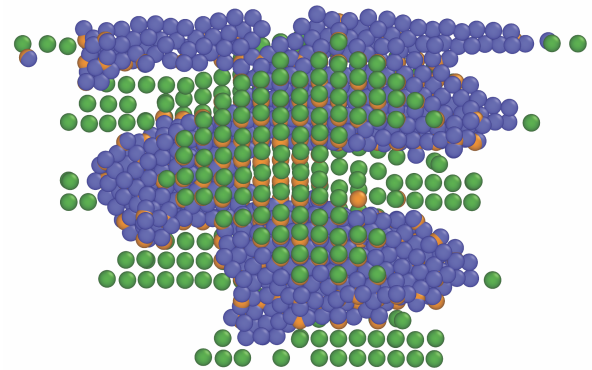

Bubbles grown at fast rates result in different growth behavior and visual properties.

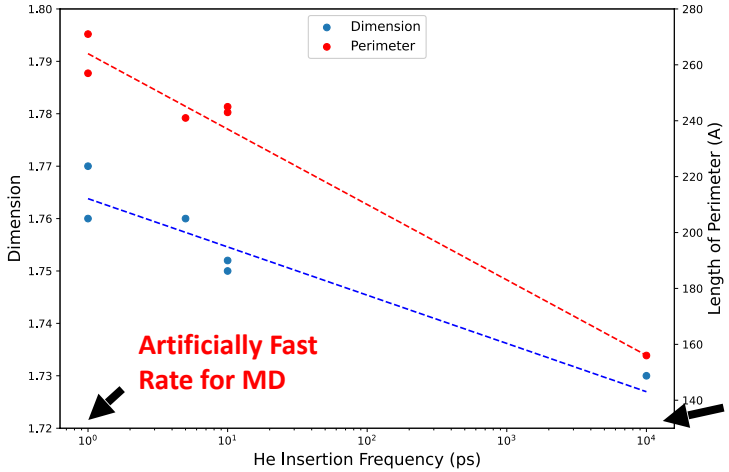

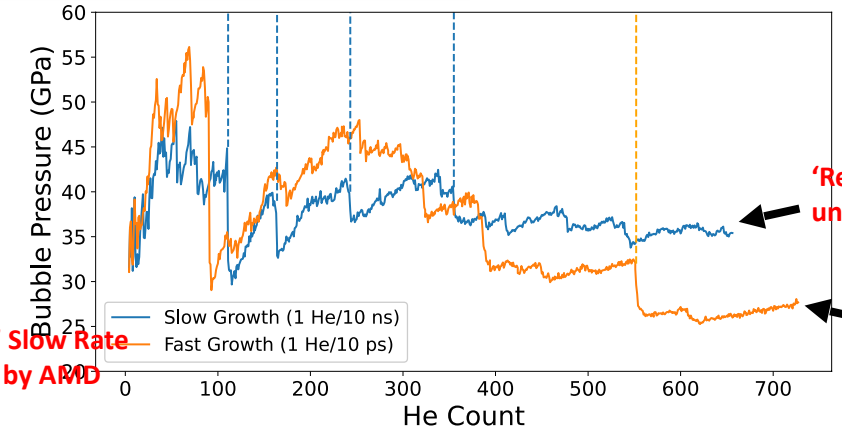

He Bubbles in Tungsten – W GB

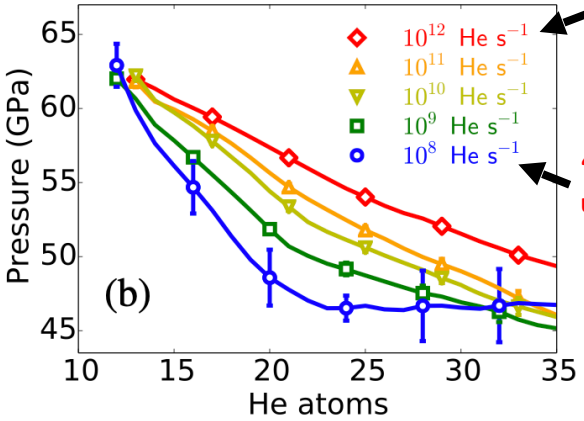

Artificially Fast Rate for MD

'Realistic' Slow Rate unlocked by AMD

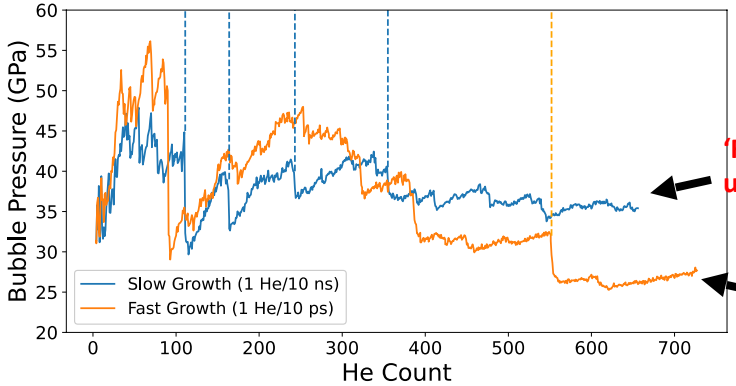

'Realistic' Slow Rate unlocked by AMD

Artificially Fast Rate for MD

He Bubbles in Tungsten – W GB

W Interstitials Generate Dislocation Loops – **Decreases Pressure**

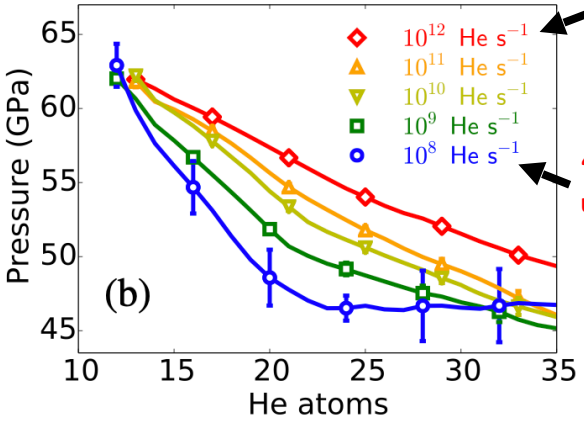

Artificially Fast Rate for MD

'Realistic' Slow Rate unlocked by AMD

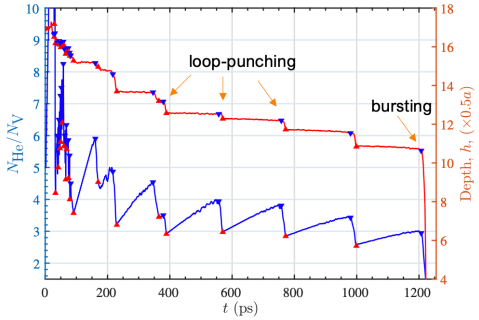

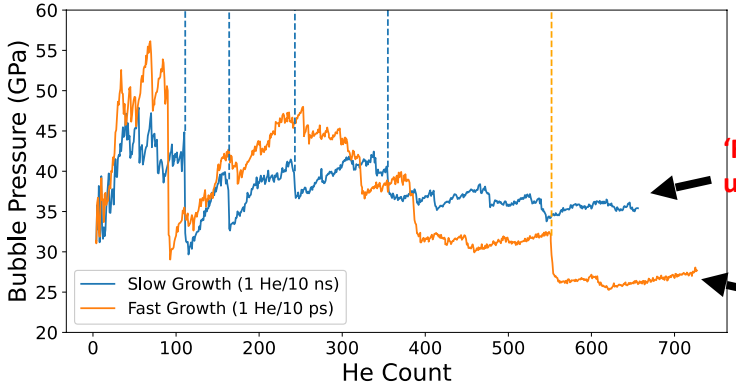

'Realistic' Slow Rate unlocked by AMD

Artificially Fast Rate for MD

He Bubbles in Tungsten – W GB

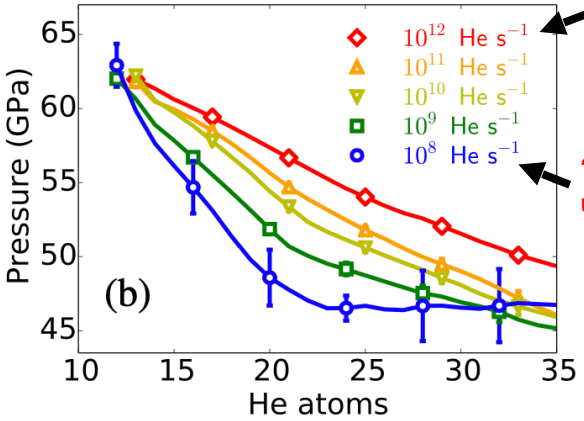

Artificially Fast Rate for MD

'Realistic' Slow Rate unlocked by AMD

W Interstitials Generate Dislocation Loops – **Decreases Pressure**

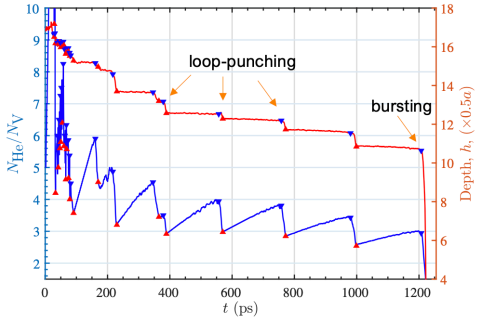

W Interstitial Restrict Bubble Growth – **Increases Pressure**

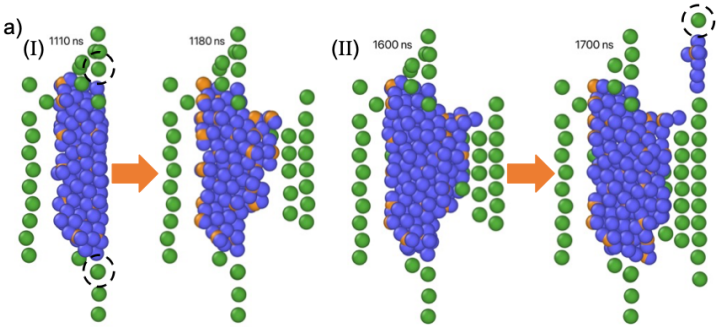

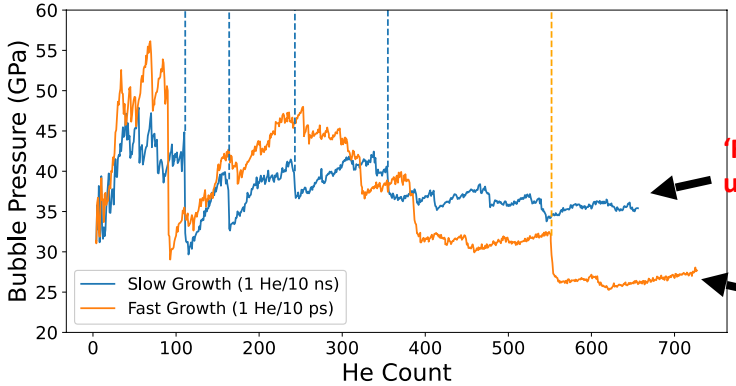

'Realistic' Slow Rate unlocked by AMD

Artificially Fast Rate for MD

He Bubbles in Tungsten – W GB

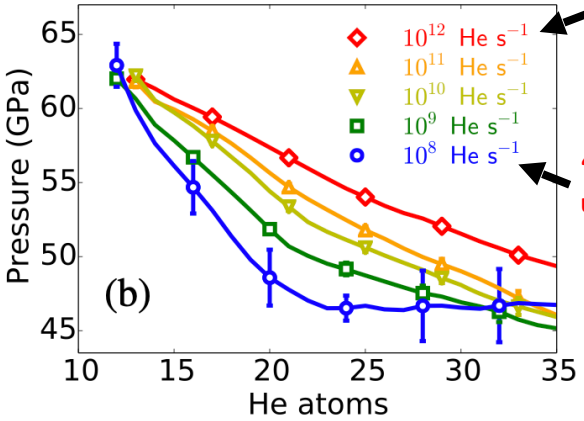

Artificially Fast Rate for MD

'Realistic' Slow Rate unlocked by AMD

W Interstitials Generate Dislocation Loops – **Decreases Pressure**

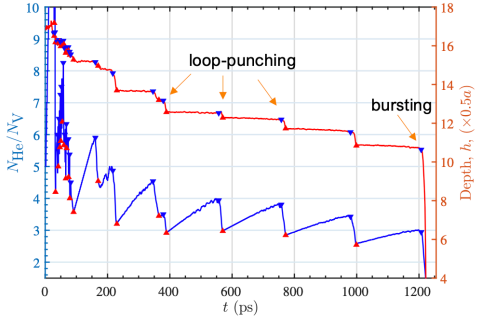

W Interstitial Restrict Bubble Growth – **Increases Pressure**

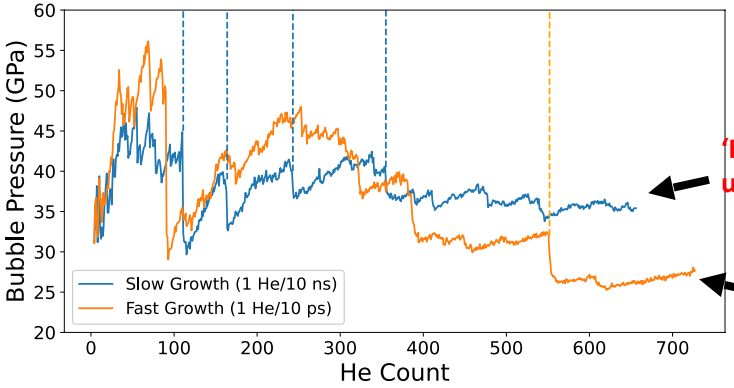

'Realistic' Slow Rate unlocked by AMD

Artificially Fast Rate for MD

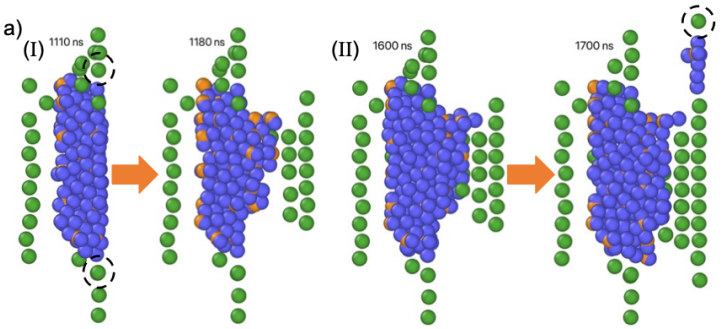

Highly non-intuitive rate effects that could not be guessed *a priori*

Los Alamos
NATIONAL LABORATORY

A New Challenge for Efficiency… GPUs

GPUs as Weak Scaling Monsters – The New Challenge

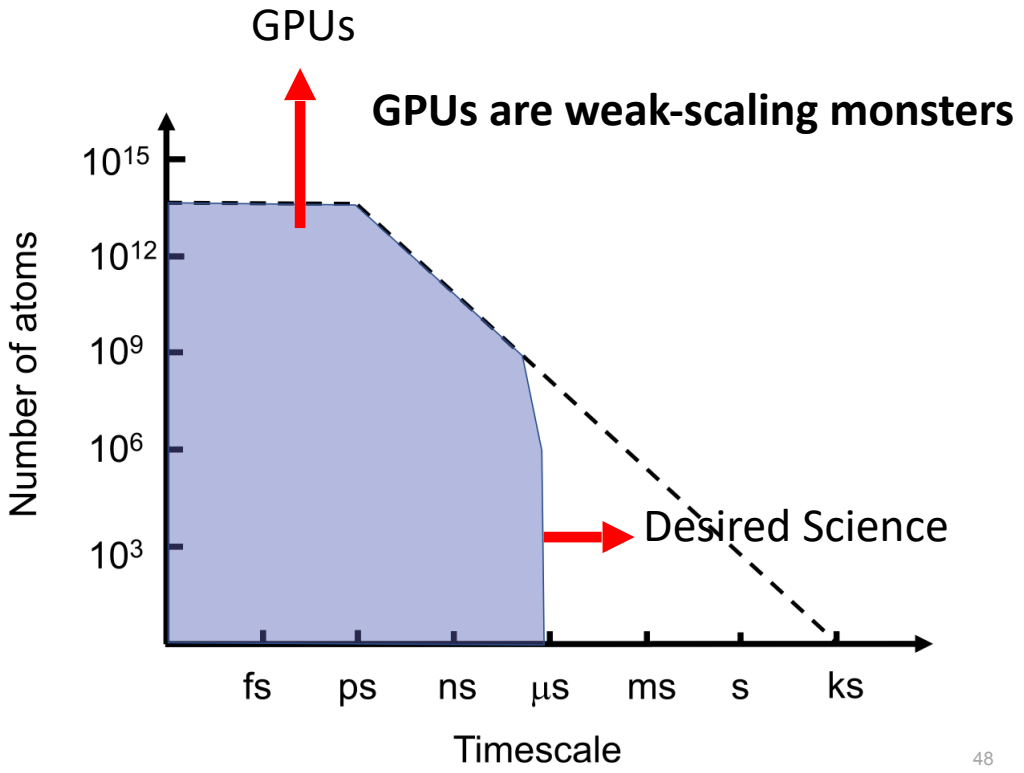

GPUs as Weak Scaling Monsters – The New Challenge

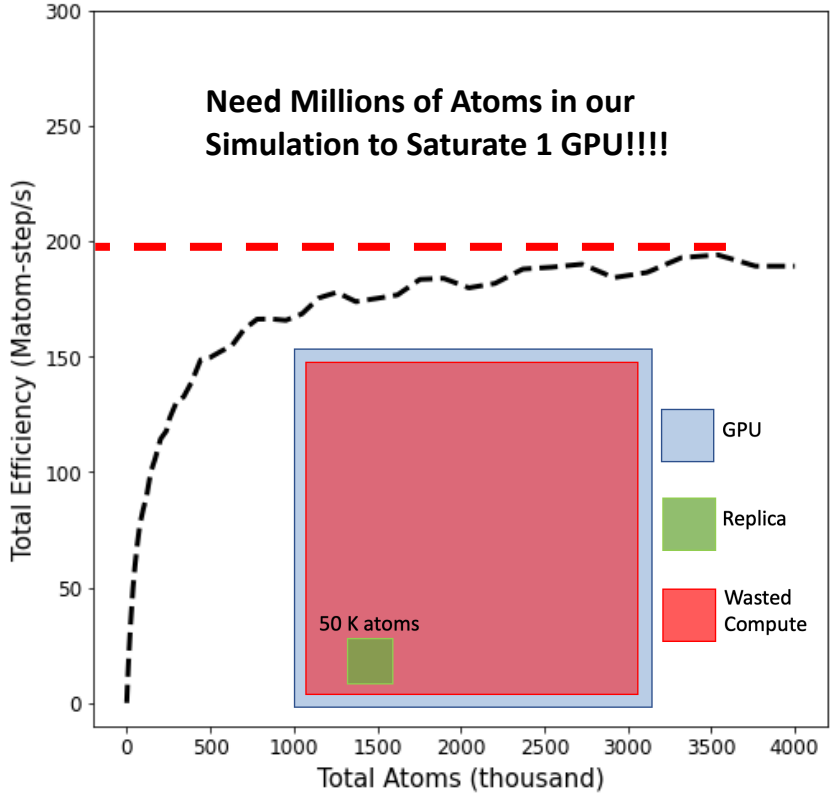

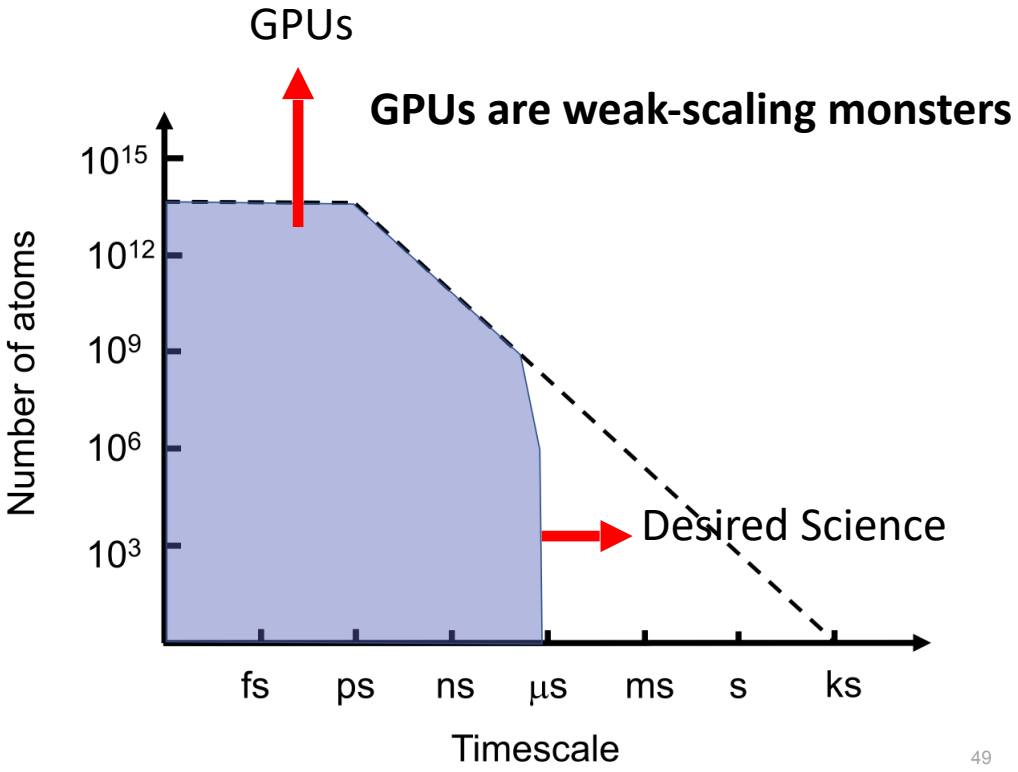

Los Alamos NATIONAL LABORATORY

3-Month ~~Vacation~~ Program at UCLA (IPAM)

institute for pure & applied mathematics

New Mathematics for the Exascale: Applications to Materials Science

March 13 - June 16, 2023

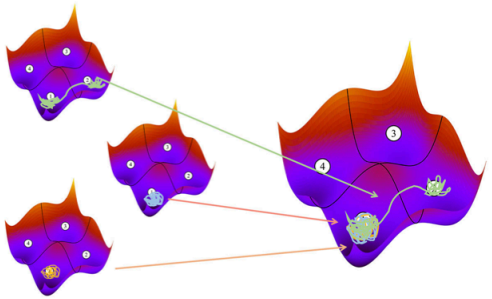

GPUs as Weak Scaling Monsters

THE DREAM

Turn these weak-scaling monsters into something stronger.

In some sense we would like to **Spatially Partition our GPUs** into multiple simulations.

GPUs as Weak Scaling Monsters

THE DREAM

Turn these weak-scaling monsters into something stronger.

In some sense we would like to **Spatially Partition our GPUs** into multiple simulations.

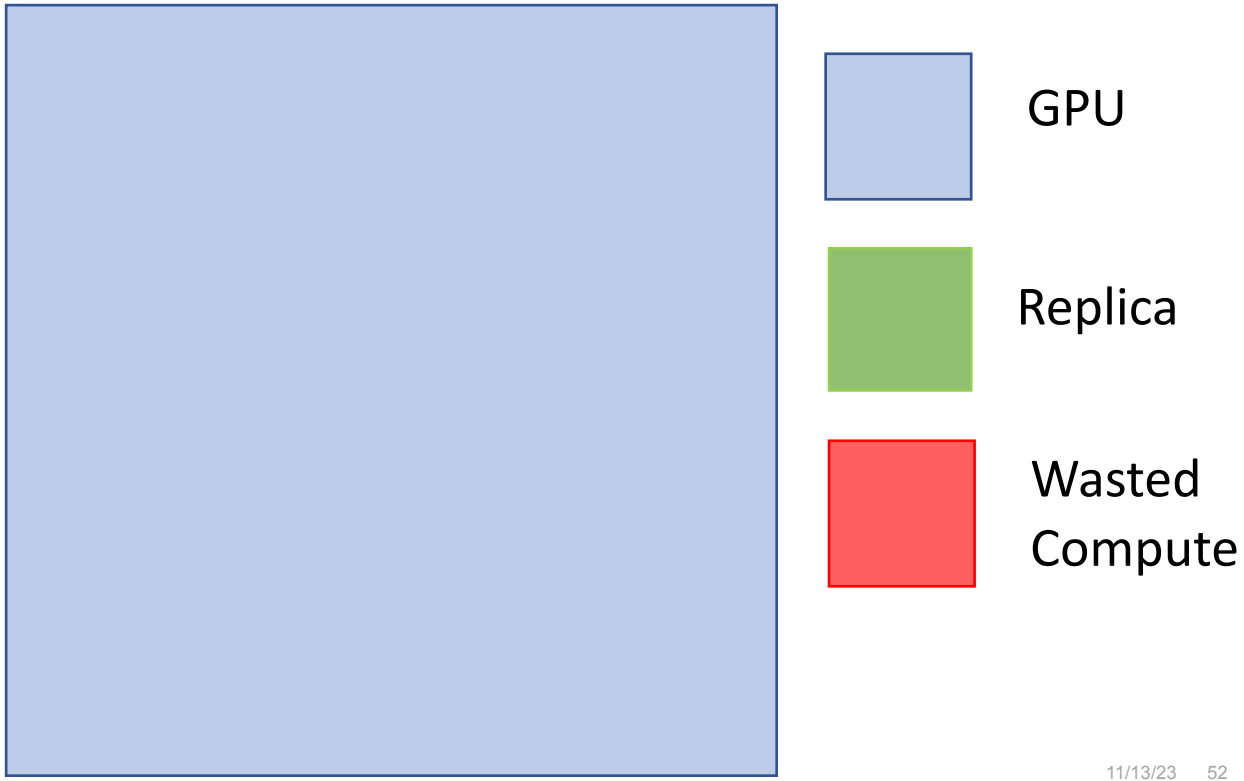

GPUs as Weak Scaling Monsters

THE DREAM

Turn these weak-scaling monsters into something stronger.

In some sense we would like to **Spatially Partition our GPUs** into multiple simulations.

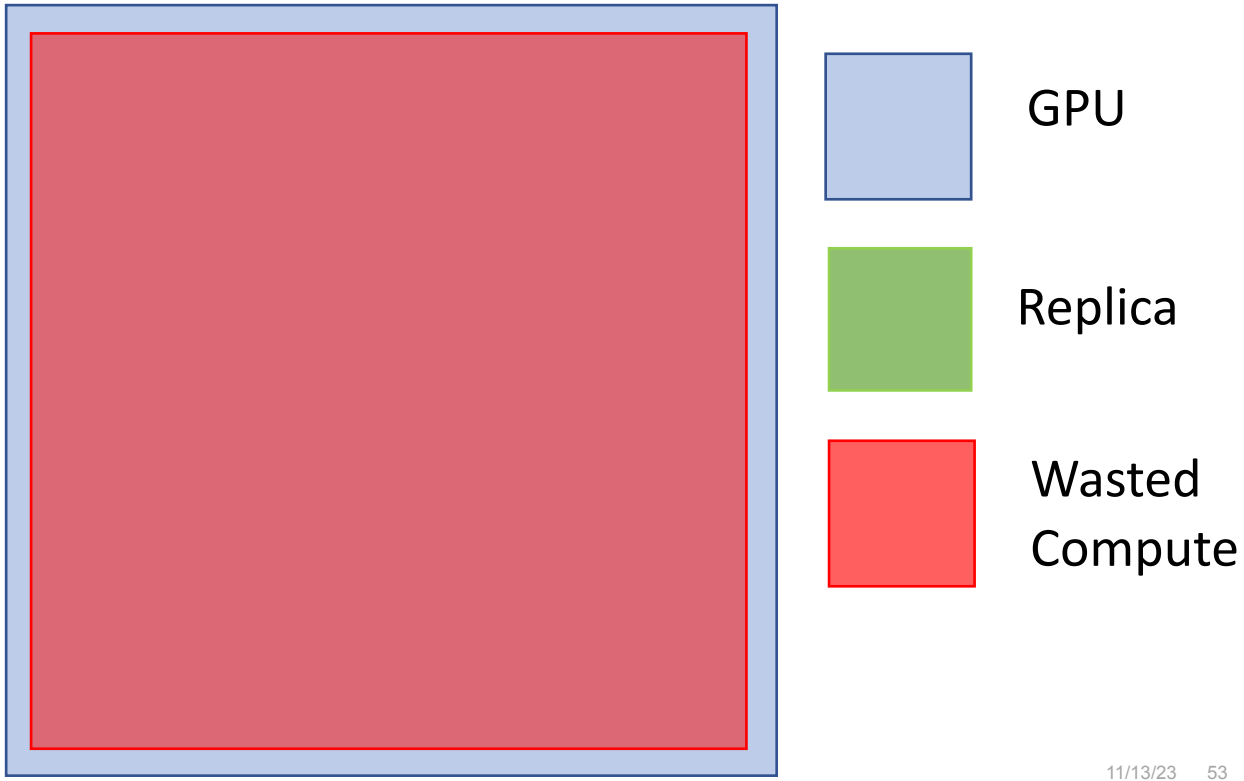

GPUs as Weak Scaling Monsters

THE DREAM

Turn these weak-scaling monsters into something stronger.

In some sense we would like to **Spatially Partition our GPUs** into multiple simulations.

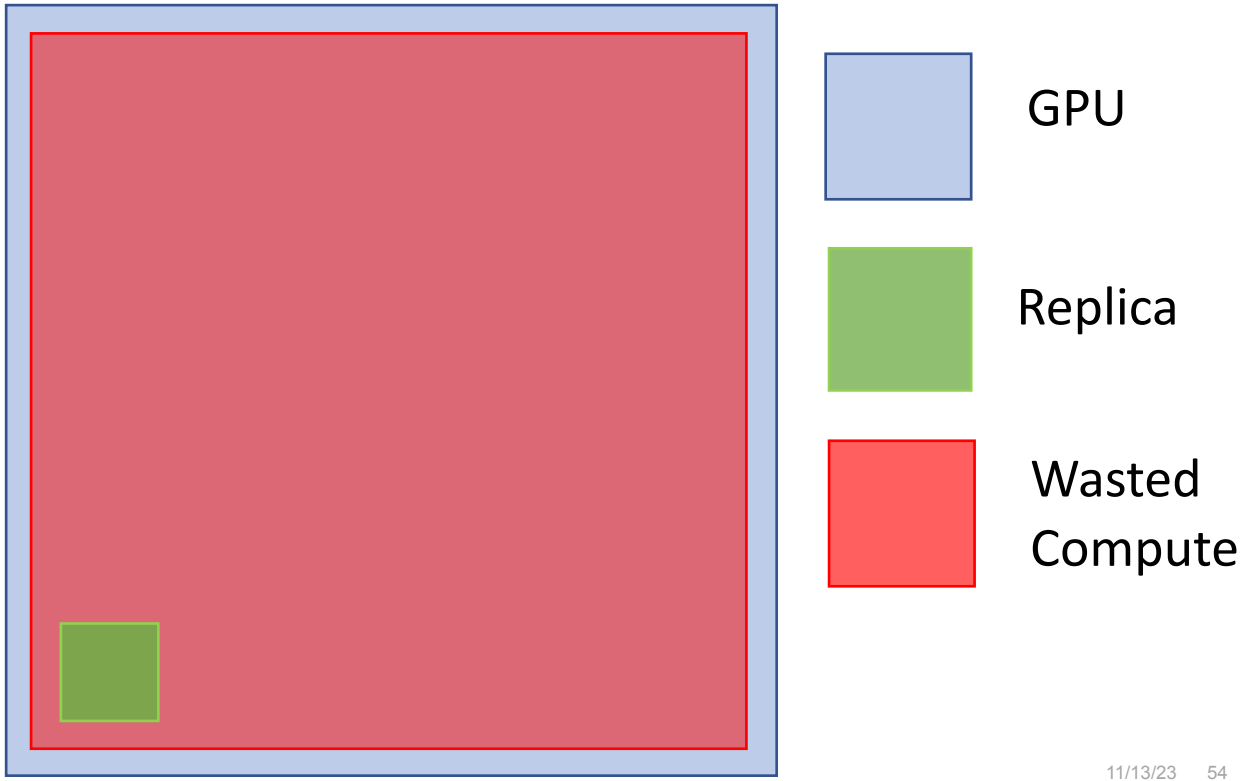

GPUs as Weak Scaling Monsters

THE DREAM

Turn these weak-scaling monsters into something stronger.

In some sense we would like to **Spatially Partition our GPUs** into multiple simulations.

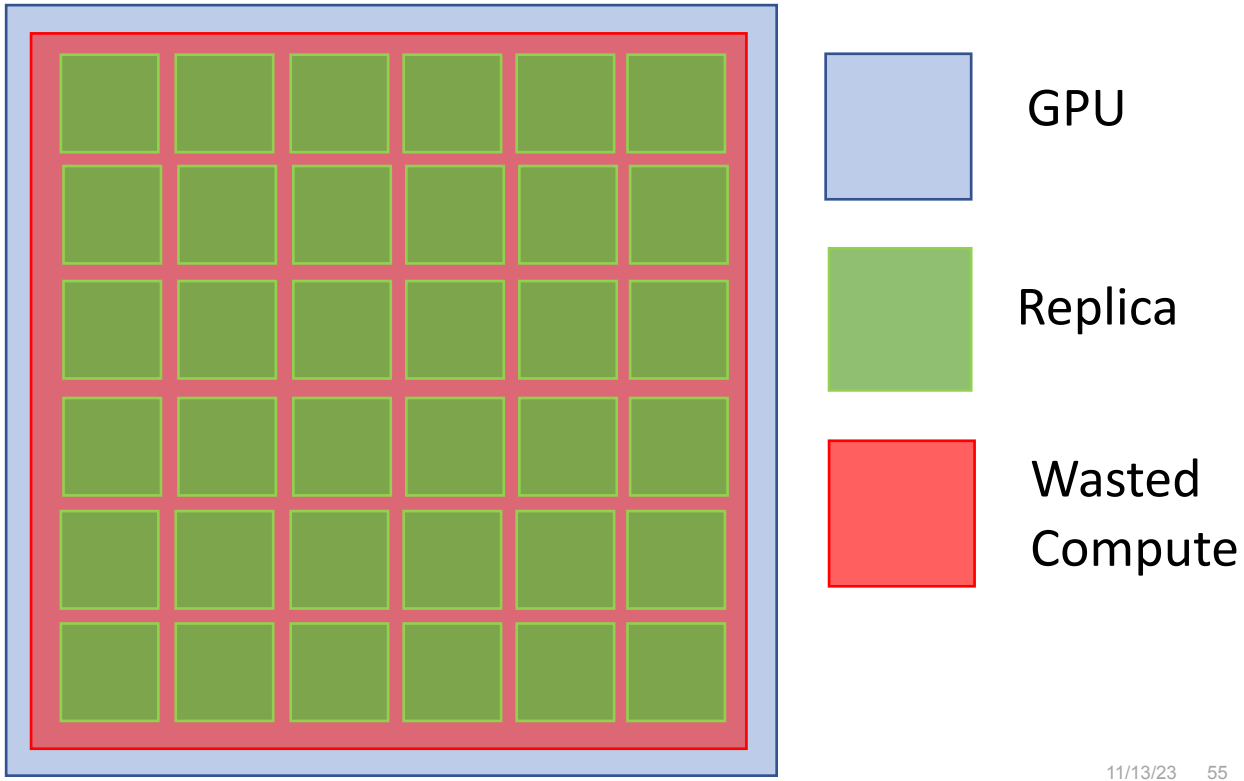

Making Oversubscription Cool

Making Oversubscription Cool

Multi-Instance GPU (MIG)

Physical subdivision of A100 GPU into partitions which can be used independently.

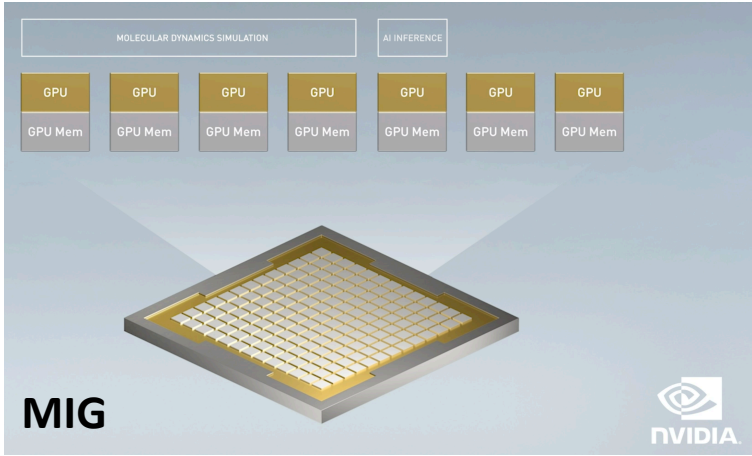

Making Oversubscription Cool

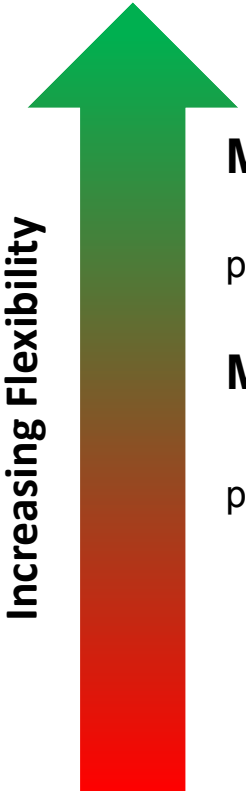

Multi-Instance GPU (MIG)

Physical subdivision of A100 GPU into partitions which can be used independently.

Multi-Process Service (MPS)

Logical subdivision of A100 GPU into partitions which can be used independently.

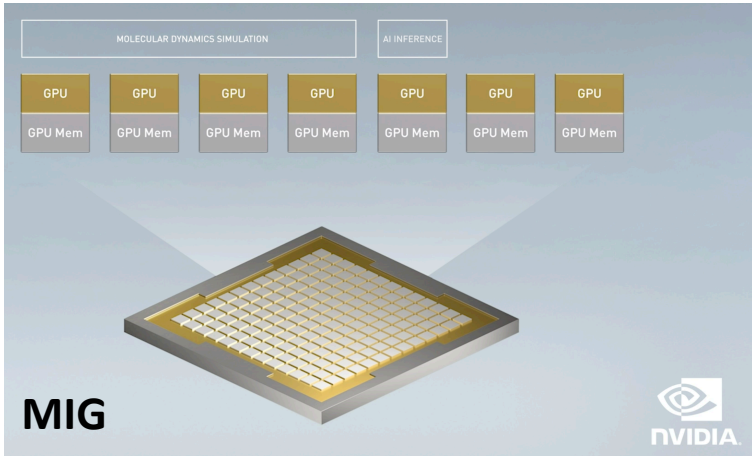

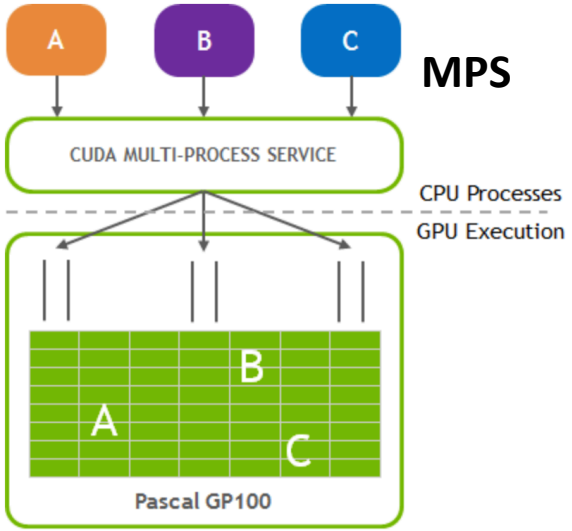

Los Alamos NATIONAL LABORATORY

Making Oversubscription Cool

Increasing Flexibility

Multi-Instance GPU (MIG)

Physical subdivision of A100 GPU into partitions which can be used independently.

Multi-Process Service (MPS)

Logical subdivision of A100 GPU into partitions which can be used independently.

LAMMPS Overlapping Simulations

Manipulate LAMMPS Neighbor list to run multiple replicas in the same simulation box.

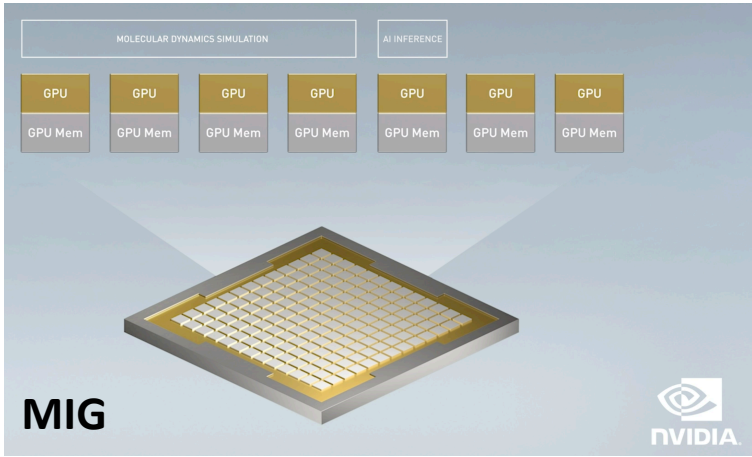

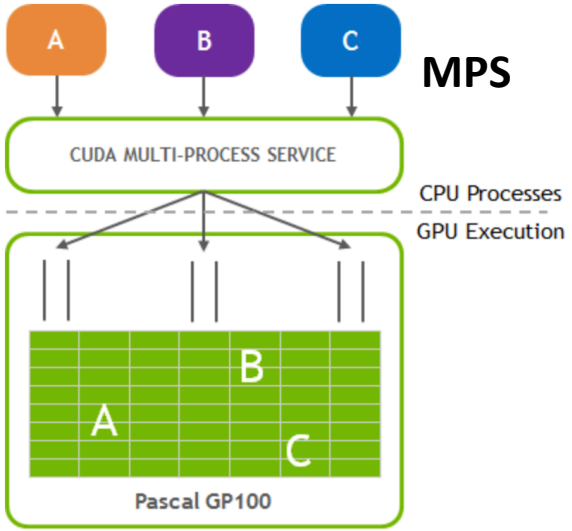

Making Oversubscription Cool

Increasing Flexibility

Multi-Instance GPU (MIG)

Physical subdivision of A100 GPU into partitions which can be used independently.

Multi-Process Service (MPS)

Logical subdivision of A100 GPU into partitions which can be used independently.

LAMMPS Overlapping Simulations

Manipulate LAMMPS Neighbor list to run multiple replicas in the same simulation box.

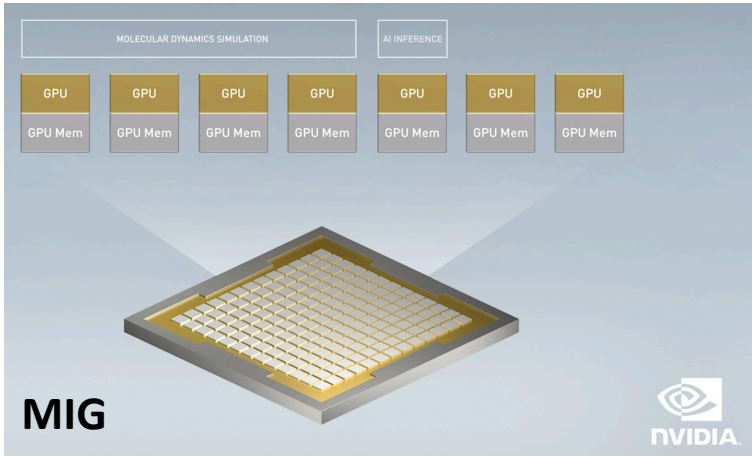

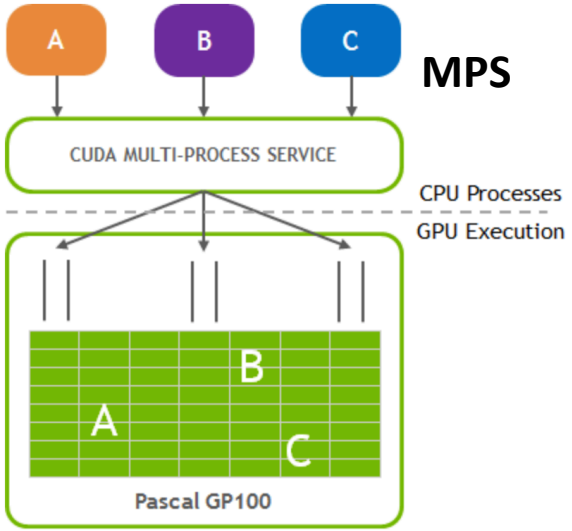

Making Oversubscription Cool - MPS

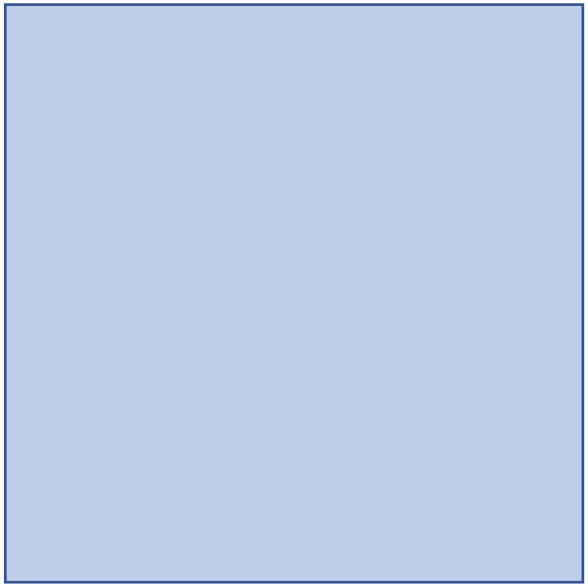

Multi-Process Service (MPS)

Supports up to 48 independent processes.

GPU MPS Instance LAMMPS Simulation

Making Oversubscription Cool - MPS

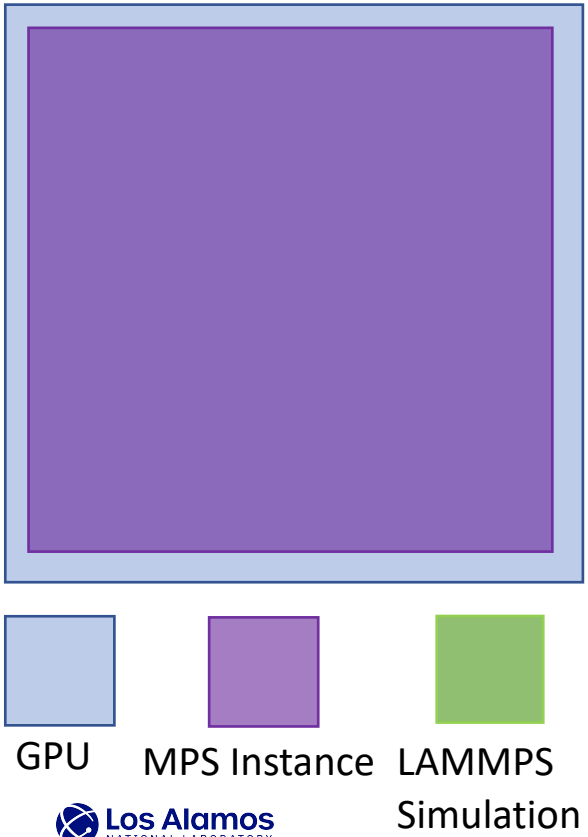

Multi-Process Service (MPS)

Supports up to 48
independent processes.

Making Oversubscription Cool - MPS

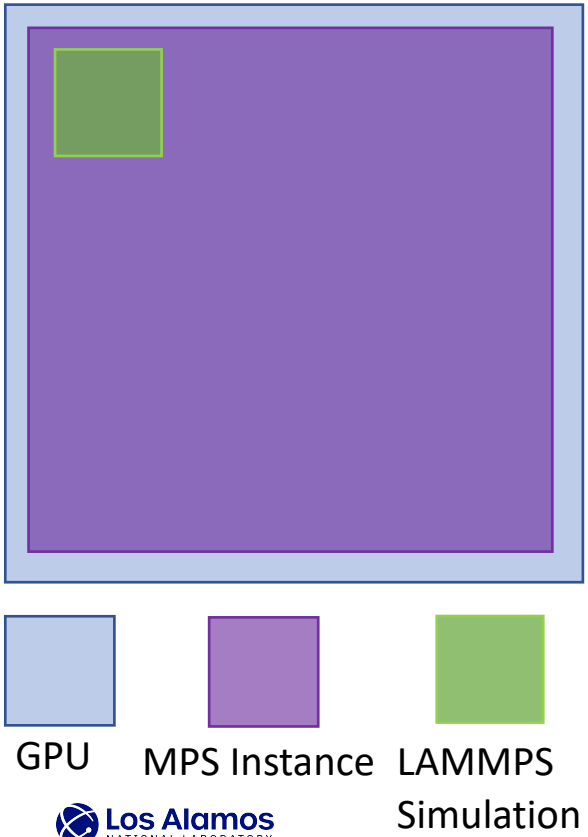

Multi-Process Service (MPS)

Supports up to 48
independent processes.

Making Oversubscription Cool - MPS

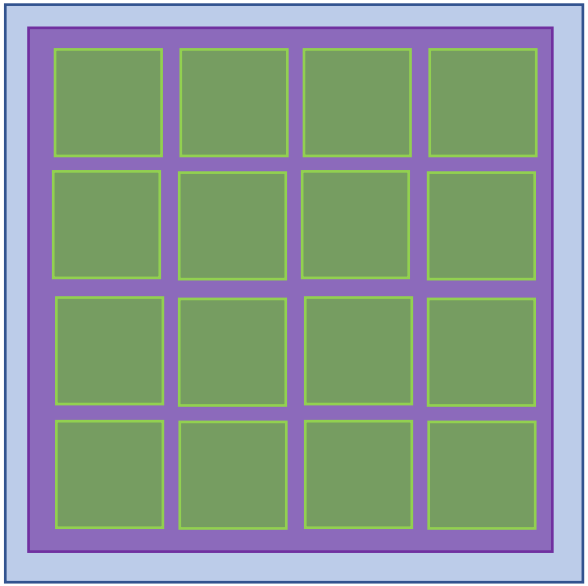

GPU MPS Instance LAMMPS Simulation

Multi-Process Service (MPS)

Supports up to 48 independent processes.

Making Oversubscription Cool - MPS

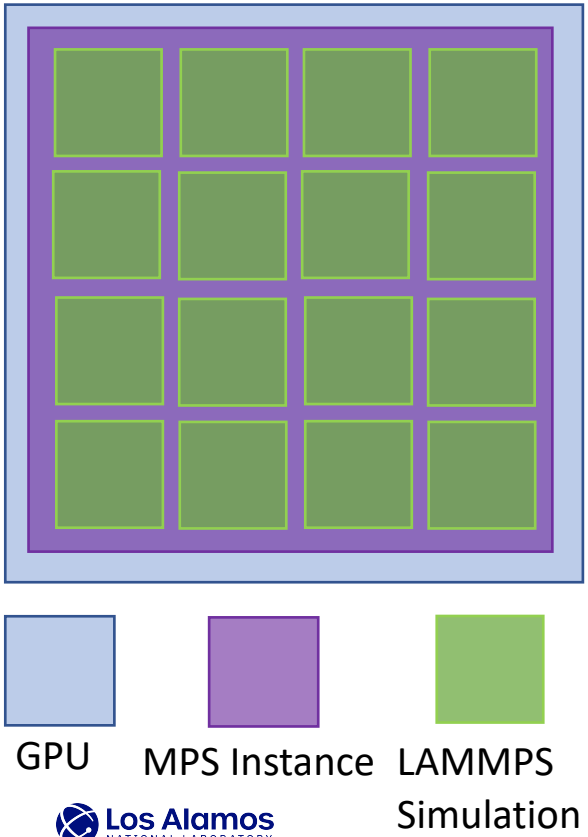

Multi-Process Service (MPS)

Supports up to 48 independent processes.

Poor total efficiency increase after ~ 5 MPS instances.

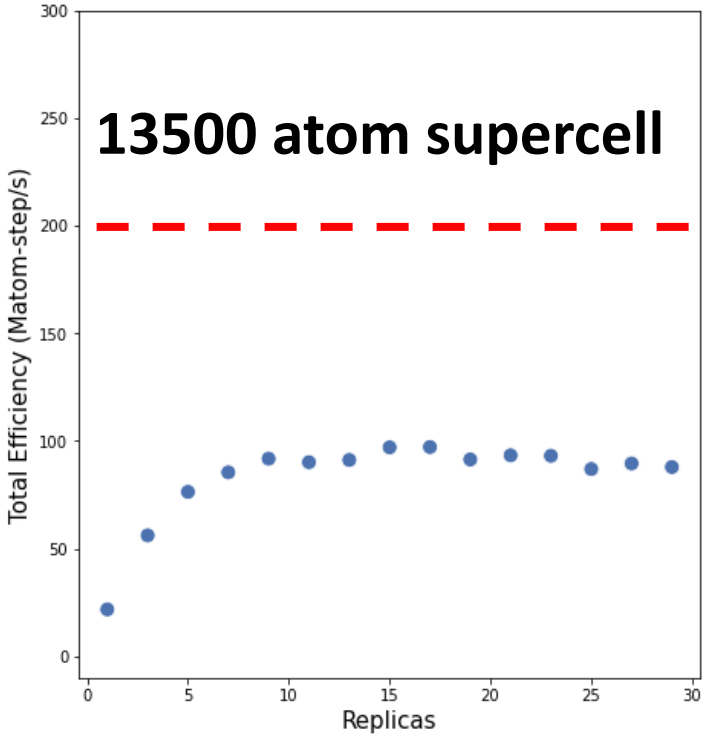

Los Alamos NATIONAL LABORATORY

Making Oversubscription Cool – Overlapping LAMMPS

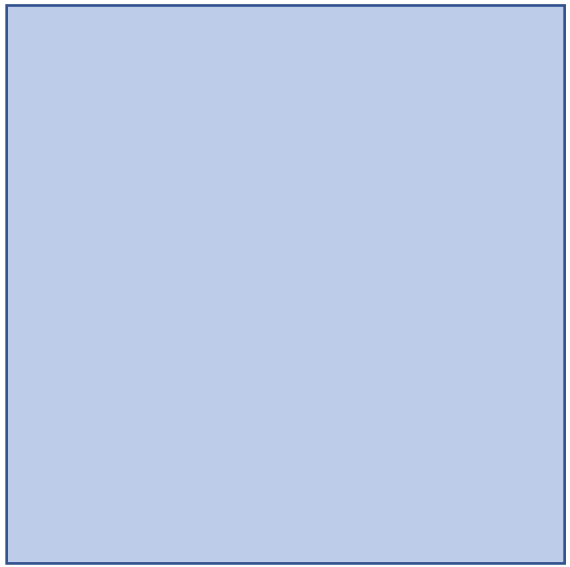

GPU

LAMMPS Instance

LAMMPS Simulation

Making Oversubscription Cool – Overlapping LAMMPS

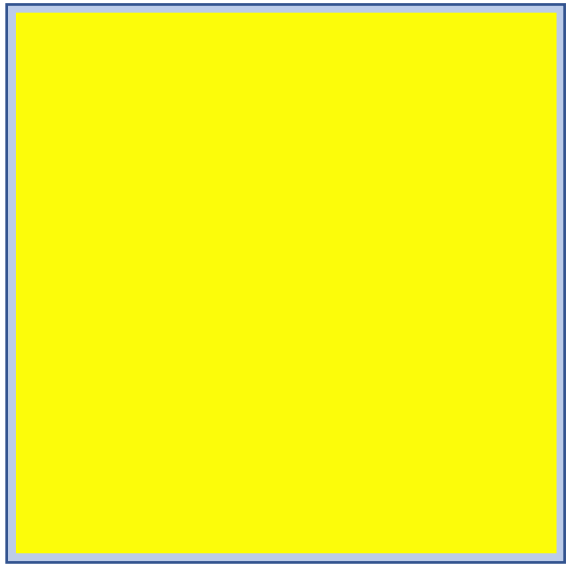

GPU

LAMMPS Instance

LAMMPS Simulation

Making Oversubscription Cool – Overlapping LAMMPS

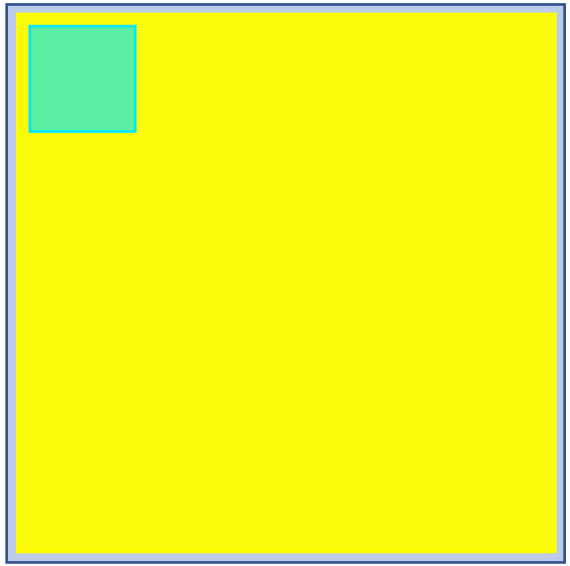

Making Oversubscription Cool – Overlapping LAMMPS

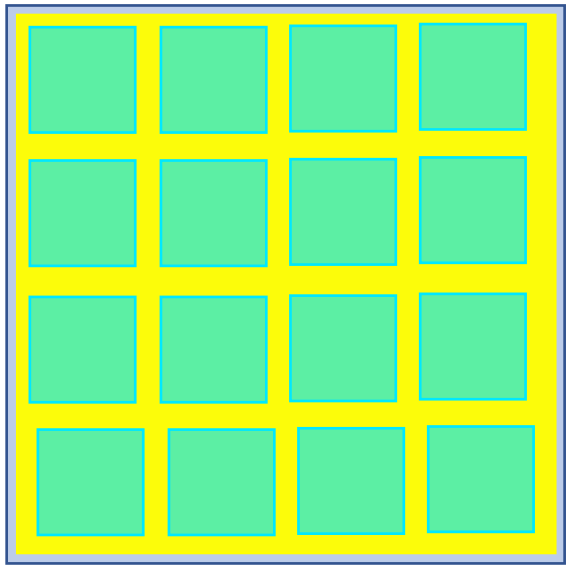

GPU

LAMMPS Instance

LAMMPS Simulation

Making Oversubscription Cool – Overlapping LAMMPS

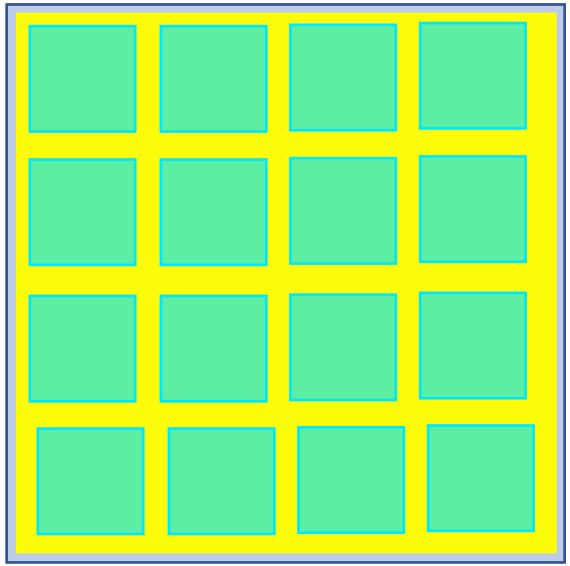

Supports up to 32 LAMMPS Simulations

Again, poor total efficiency increase after ~ 5 Replicas

GPU

LAMMPS Instance

LAMMPS Simulation

Making Oversubscription Cool – Overlapping LAMMPS

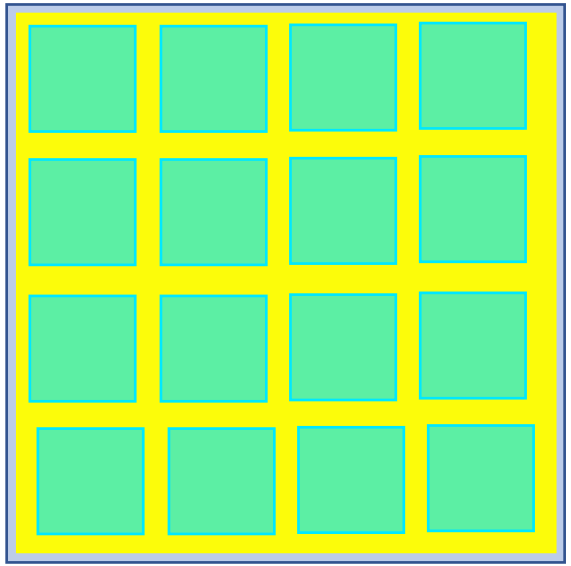

GPU

LAMMPS Instance

LAMMPS Simulation

Supports up to 32 LAMMPS Simulations

Again, poor total efficiency increase after ~ 5 Replicas

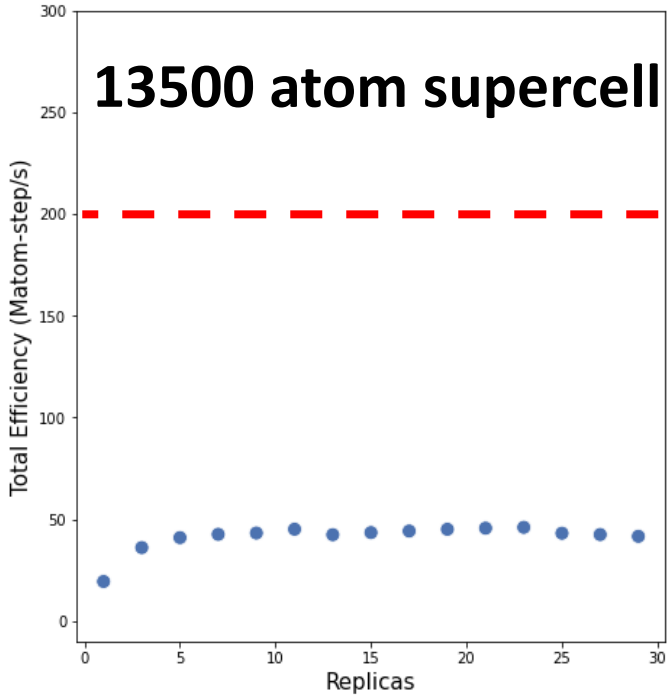

Making Oversubscription Cool – Combined Models

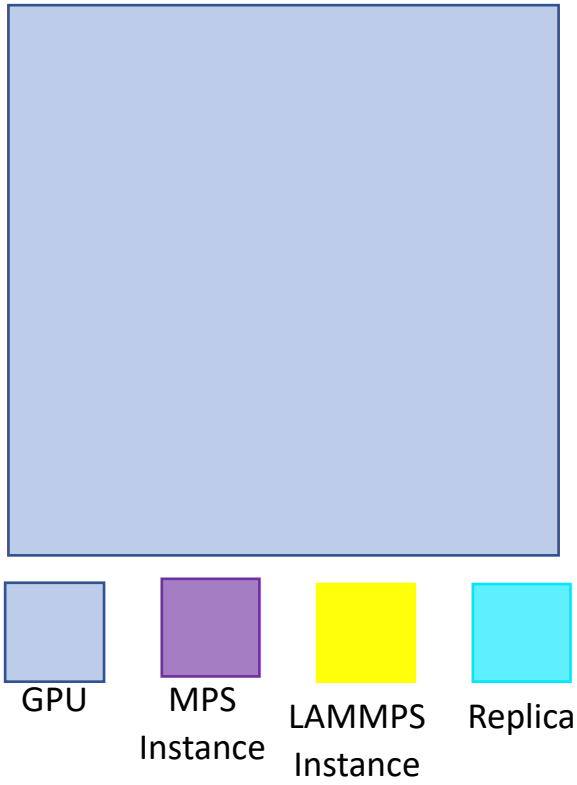

Making Oversubscription Cool – Combined Models

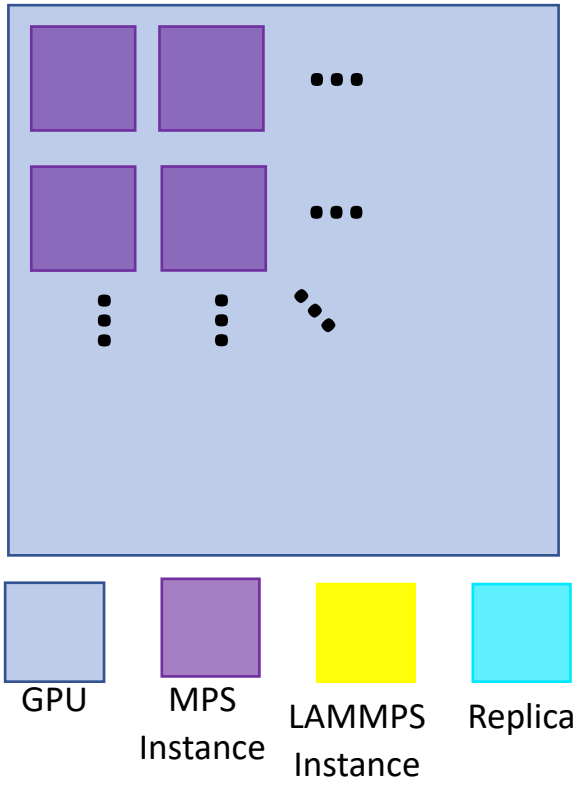

Making Oversubscription Cool – Combined Models

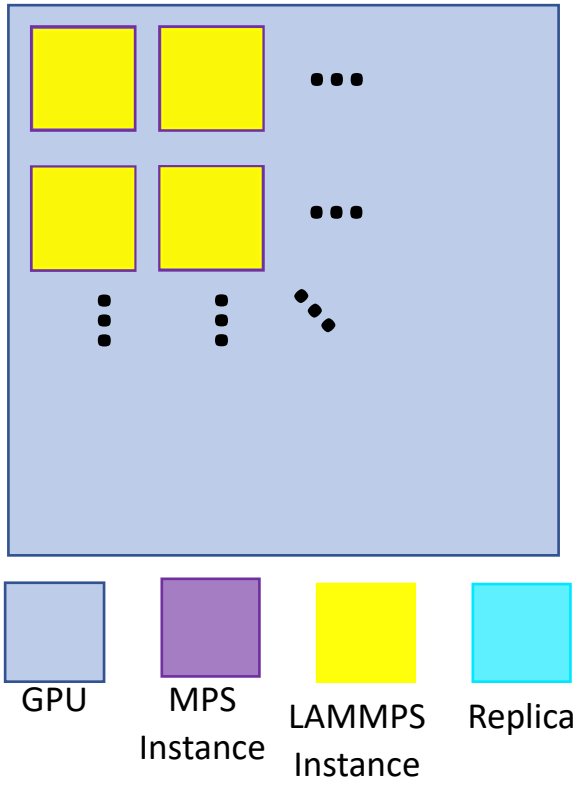

Making Oversubscription Cool – Combined Models

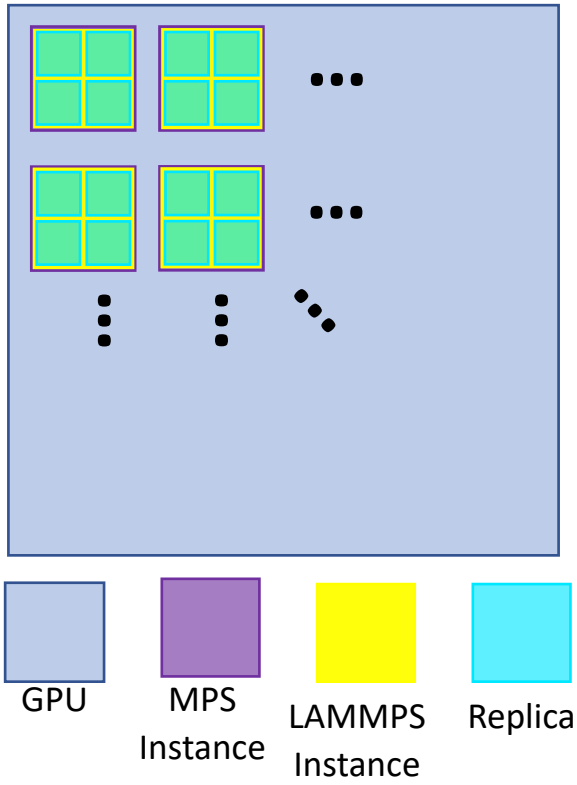

Making Oversubscription Cool – Combined Models

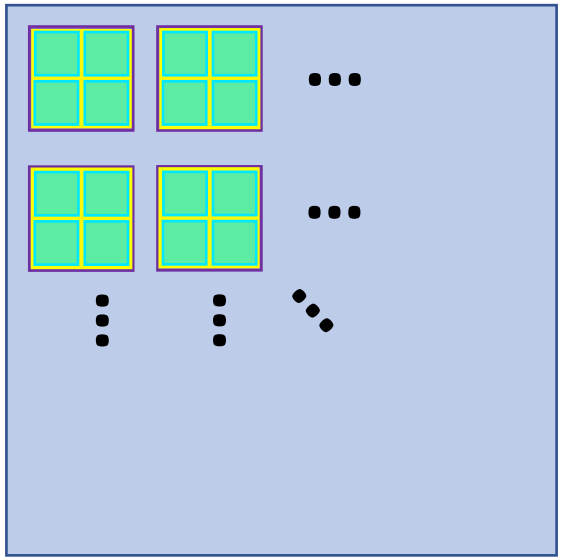

GPU

MPS Instance

LAMMPS Instance

Replica

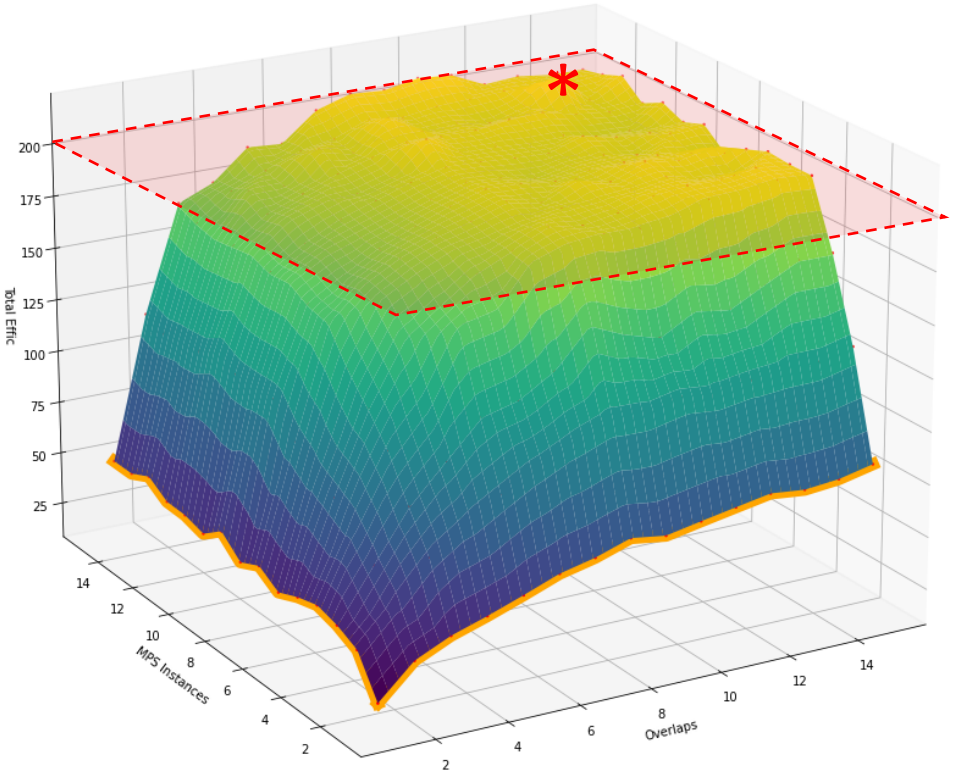

Making Oversubscription Cool – Example in Copper

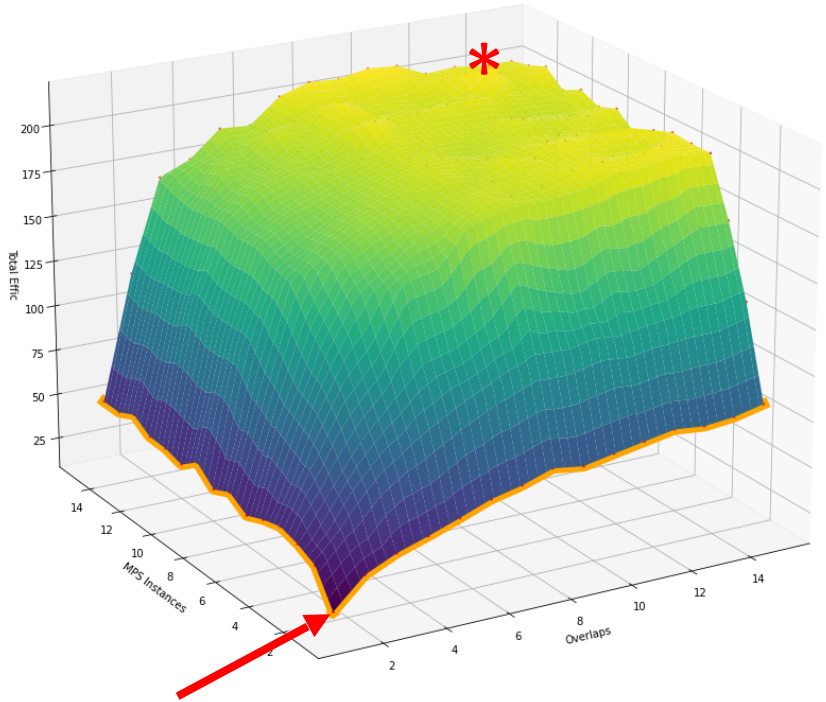

1 Replica per GPU (current method)

Making Oversubscription Cool – Example in Copper

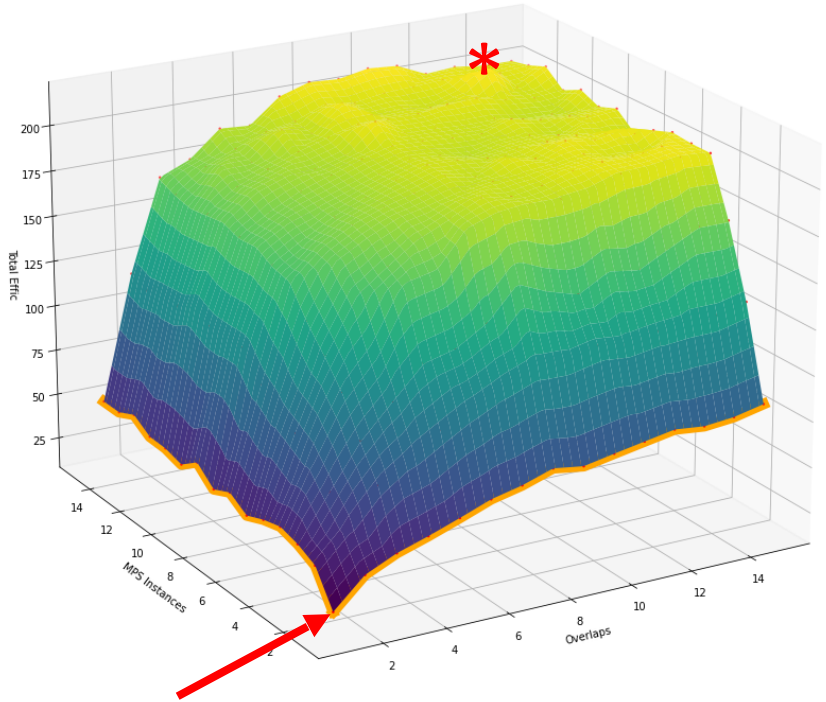

1 Replica per GPU (current method)

Example Setup:

13,500 atom supercell.

EAM interatomic potential.

Making Oversubscription Cool – Example in Copper

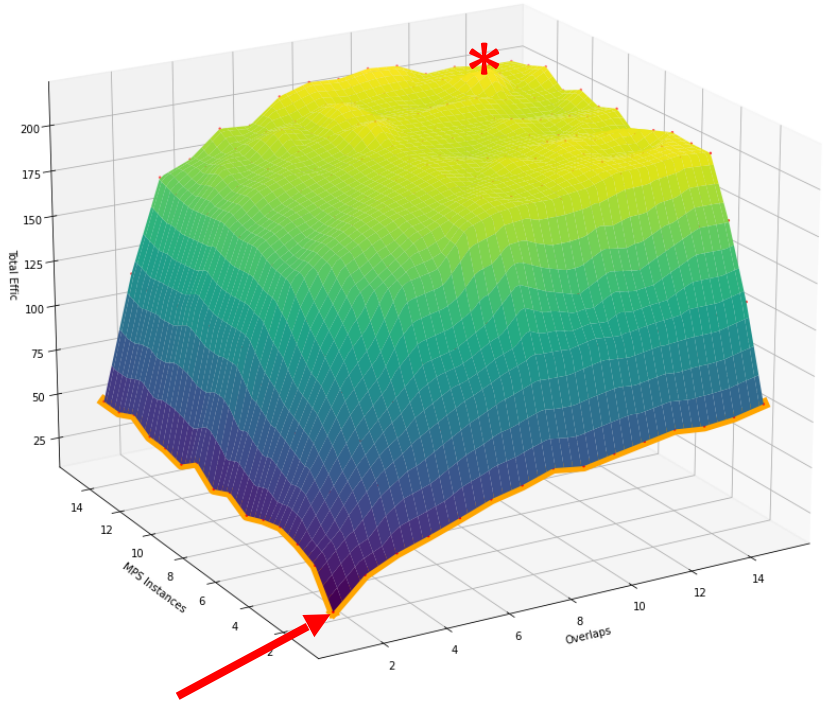

1 Replica per GPU (current method)

Example Setup:

13,500 atom supercell.

EAM interatomic potential.

Optimal Packing:

13 MPS instances

13 Overlapping Sims

Total = 169 Simulations on 1 GPU!

Making Oversubscription Cool – Example in Copper

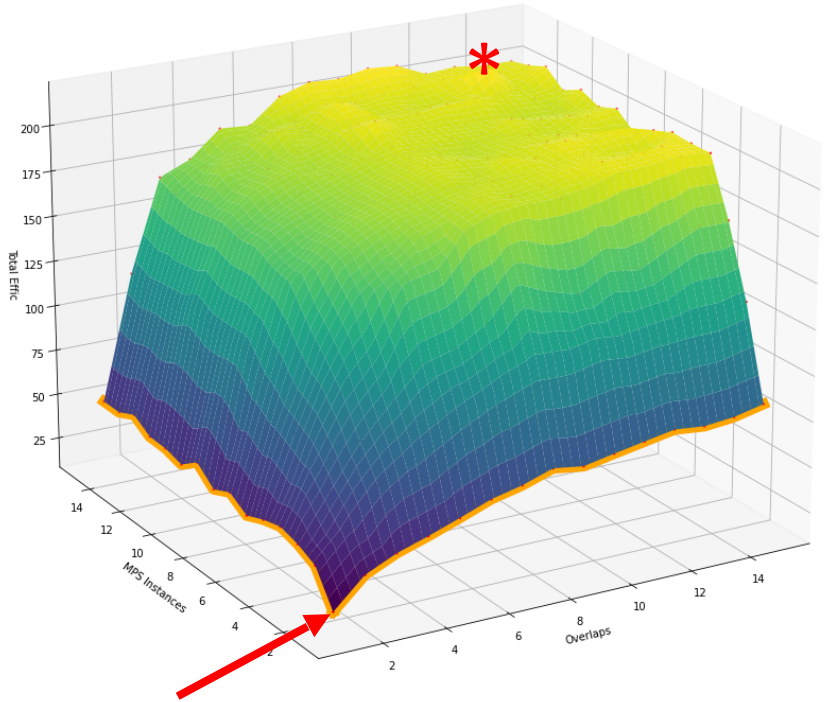

1 Replica per GPU (current method)

Example Setup:

13,500 atom supercell.

EAM interatomic potential.

Optimal Packing:

13 MPS instances

13 Overlapping Sims

Total = 169 Simulations on 1 GPU!

Total Efficiency:

190 Matom-step/s

5.577 μs/day

… on 1 GPU!

The Ongoing Battle for Efficiency

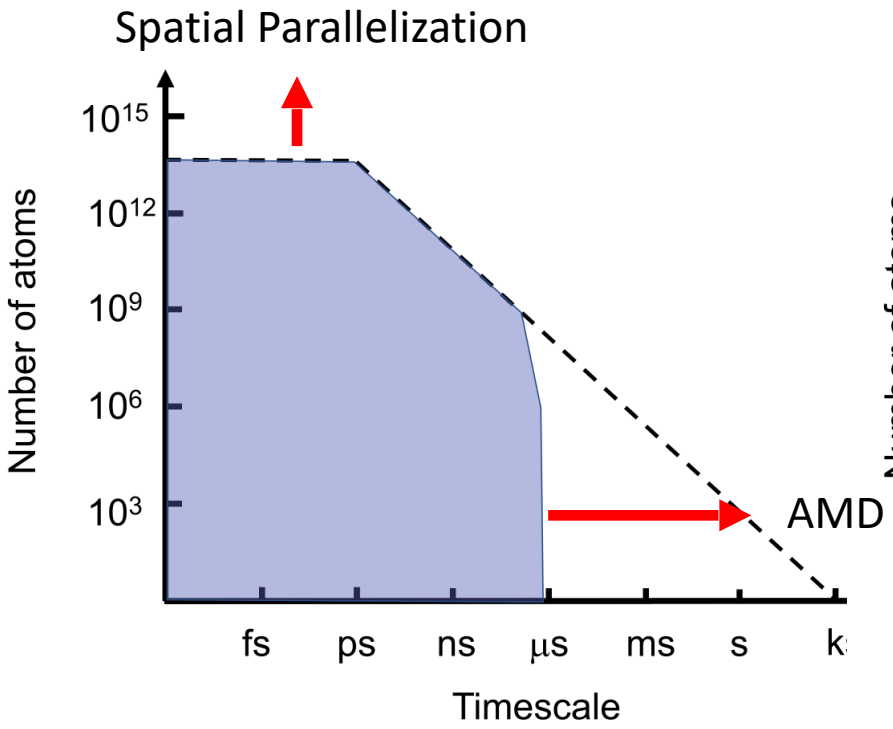

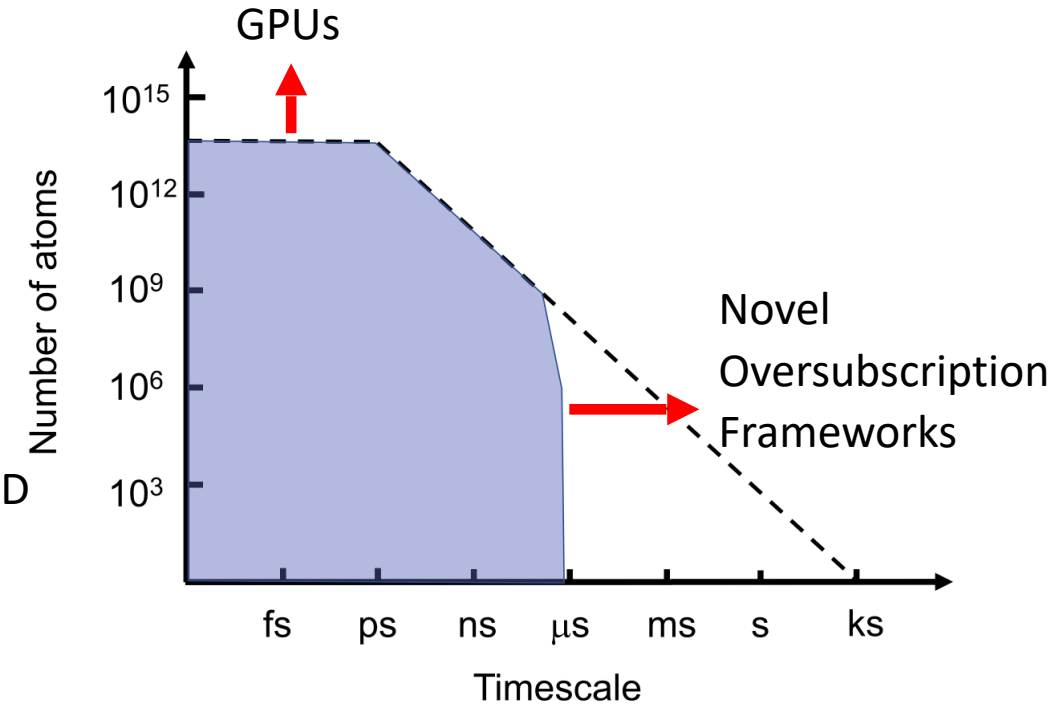

Conclusions

- **Software and hardware are inherently weak-scaling focused.**

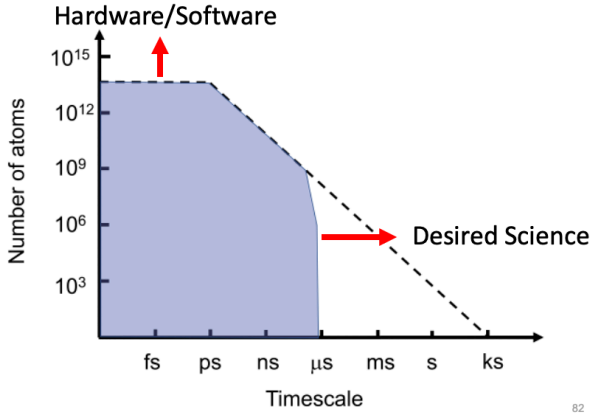

Conclusions

- **Software and hardware are inherently weak-scaling focused.**

- **Must devise ways to manipulate weak-scaling into strong(-ish) scaling to reach longer timescales.**

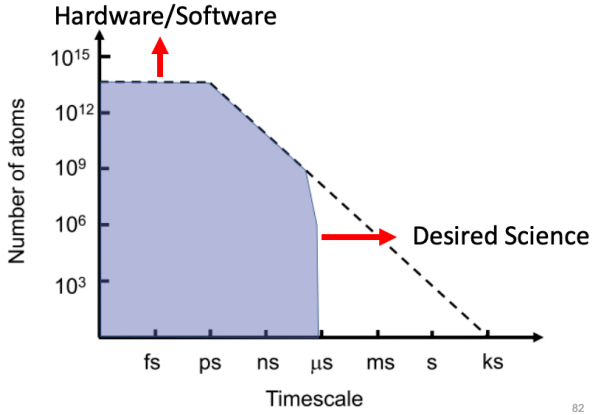

Conclusions

- **Software and hardware are inherently weak-scaling focused.**
- **Must devise ways to manipulate weak-scaling into strong(-ish) scaling to reach longer timescales.**
- **Novel oversubscription frameworks coupled with AMD allows us to reach new levels of computational efficiency on GPU-based HPCs.**
- **This can allow us to achieve scientific insight previously unattainable.**

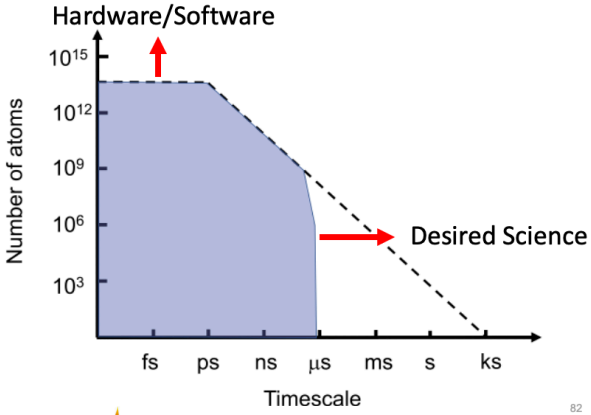

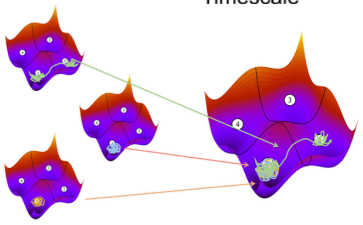

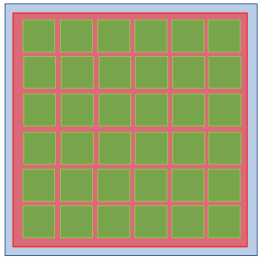

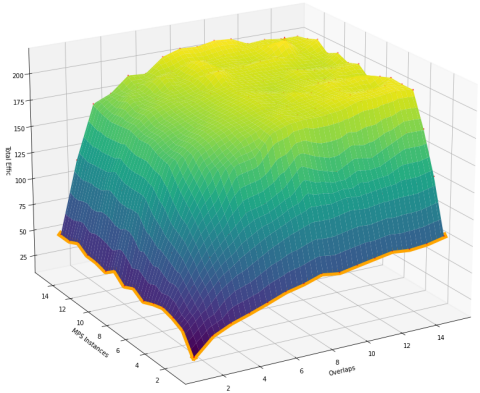

Acknowledgements

Funding:

Funding as part of the Scientific Discovery through Advanced Computing (SciDAC) program, which is jointly sponsored by the Fusion Energy Sciences (FES) and Advanced Scientific Computing Research (ASCR) programs within the US Department of Energy, Office of Science. Research supported by the US Department of Energy under DE-AC05-00OR22725.

More Scientific Details:

The work presented here is supported by publications:

P. Hatton *et al.*, *The importance of long-timescale simulations for driven systems: An example of He bubble growth at a W GB,* MRS Comms. (2022).

P. Hatton *et al.*, *He Bubble Induced Phase Transition of W Grain Boundaries using Accelerated Molecular Dynamics (In Production).*