2005 Fall Creek Falls Conference: *Computational Science at Scale*

Materials modeling on BlueGene/L:

Resolidification of molten metals using ddcMD

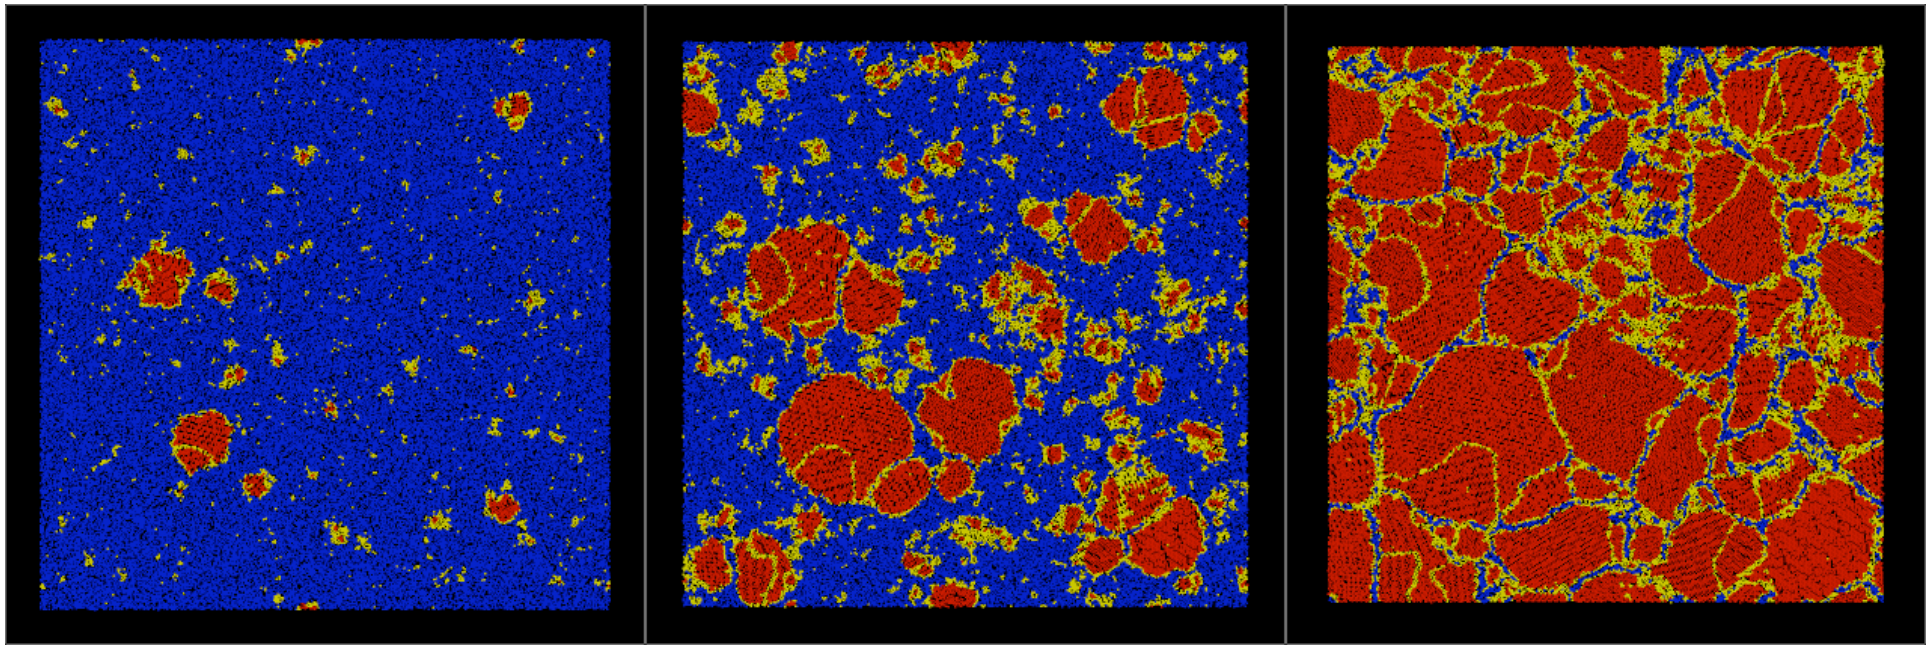

Frederick H. Streitz, James N. Glosli, Mehul V. Patel (AX-Div)
Bor Chan (CASC) and John Gunnels (IBM)
Lawrence Livermore National Laboratory
Visualization: Liam Krauss

UCRL-PRES-216478

Physics & Advanced Technologies

Investigation of rapid resolidification

Use molecular dynamics to uncover details of the solidification process

- Determine structure and stability of metals under dynamic loading conditions
- Identify relevant time scales (from ps to μs) for solidification
- Locate non-equilibrium phase boundaries
- Describe rate and path dependence of approach to final structure

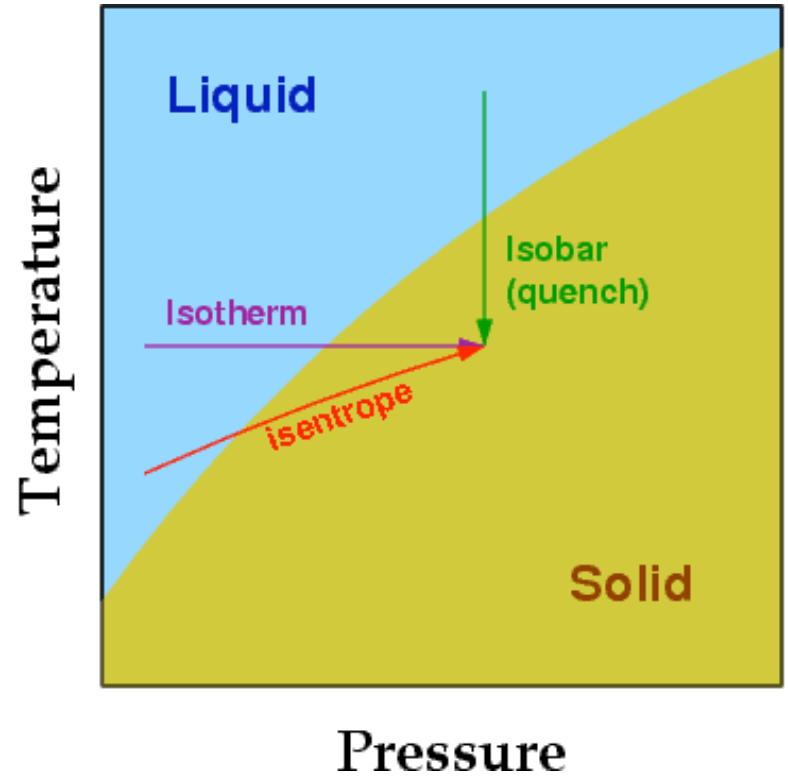

⇨ Corroborate details with experiments where possible

⇨ Condense results into phenomenological models which can be incorporated into hydrocodes

ddcMD (domain decomposition + Molecular Dynamics): designed for flexibility

Initial design requirements:

- Run efficiently on arbitrary number of processors (not just N^3)
- Excellent strong scaling to extend MD time scale
- Excellent weak scaling to extend size of simulation
- Robust restart capability

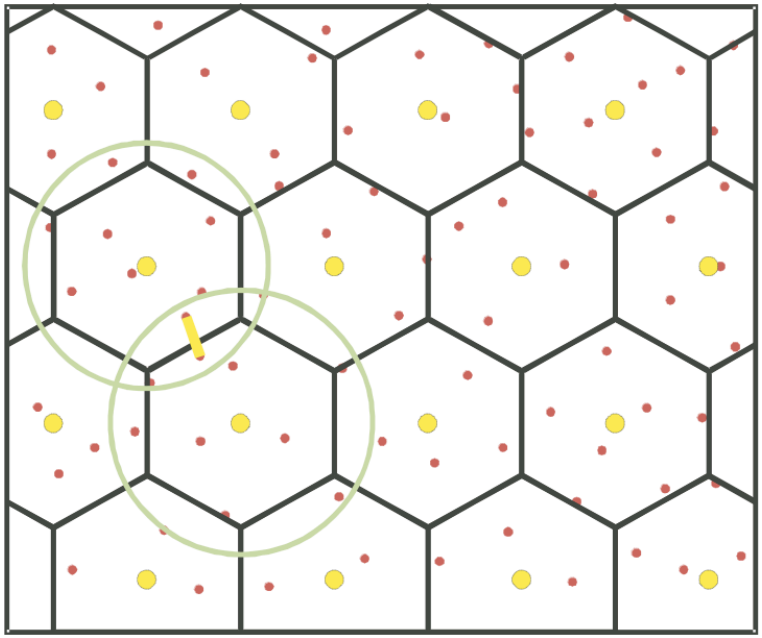

Solution:

- Particle-based domain decomposition - processors own particles, not regions
- Maintain minimum communication list for given decomposition
- Utilize separate random number stream per particle

Scaling ddcMD up to 131,072 CPUs

- Massively parallel system has some gotchas …
  - Algorithm inefficiencies for thousands of processors
  - Quantity of data files created difficulties for filesystem
  - Early time MTBF is not very long compared to run times

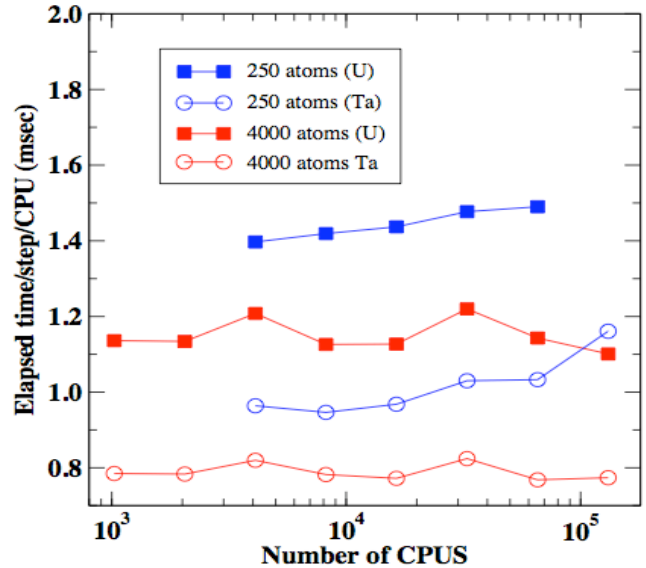

- … but allows unprecedented scaling of size or time
  - Weak scaling is virtually flat across the entire machine
  - Strong scaling shows speedup down to 8 atoms/CPU

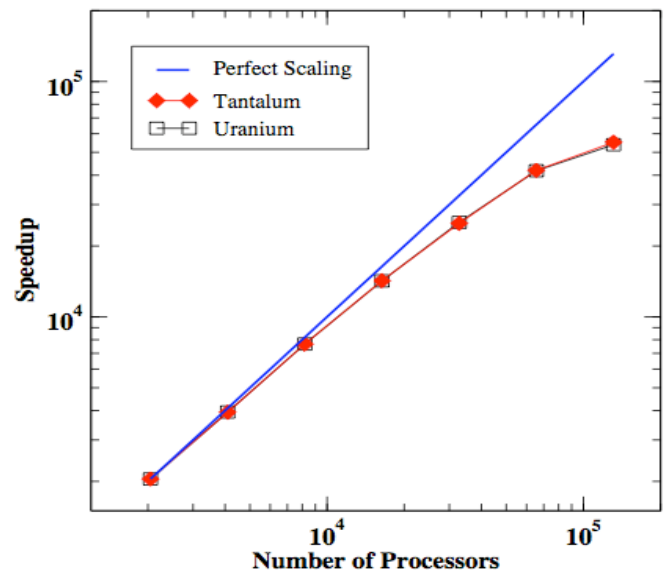

Performance of ddcMD on BlueGene/L

Weak scaling results on BlueGene/L:

Simulated quench of MGPT uranium

CPUs	Total atoms	Local/ Remote		Elapsed time	Tflop/s
1,024	4,096,000	0.95	bcc	914.90	0.84
2,048	8,192,000	0.94	fcc	914.38	1.67
4,096	16,384,000	0.85	sc	969.20	3.09
8,192	32,768,000	0.95	bcc	906.39	6.62
16,384	65,536,000	0.94	fcc	906.22	13.23
32,768	131,072,000	0.85	sc	985.33	24.41
65,536	262,144,000	0.95	bcc	916.81	52.23

Performance of ddcMD on BlueGene/L

MD		ddc	
Force	86%	Computation	1%
I/O	4%	Communication	5%
		Barrier	4%
Total	90%	Total	10%

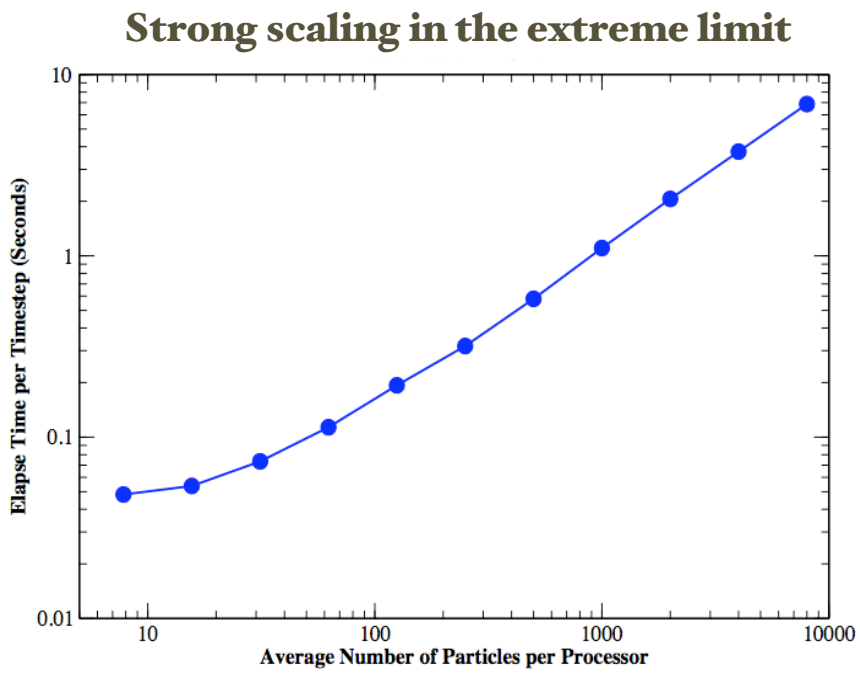

- Sustained performance (over a 7 hour run) at better than a TF/$M
- Science runs totaling hundreds of (wall clock) hours with up to 512M atoms across 128k CPUs
- Code/compiler/kernel upgrades in progress promise even better performance

Molecular dynamics simulations

- Model resolidification of tantalum (*d*-electron, bcc metal with no structural phase transitions)
- Constant NVT simulations with imposed strain rate
- Stochastic (Langevin) thermostat
- Liquid condensed from random (hot gas) positions
- Isothermally compressed on exponential ramp
- Interatomic interaction modeled using MGPT potentials

MGPT Interaction Potentials:

$$E_{MGPT} = E_0(\Omega) + \frac{1}{2}\sum_{ij} v_2(ij,\Omega) + \frac{1}{6}\sum_{ijk} v_3(ijk,\Omega) + \frac{1}{24}\sum_{ijkl} v_4(ijkl,\Omega) + \ldots$$

where the $\upsilon_n(\Omega)$ represent many-body terms which are volume dependent but structure independent and thus rigorously transferable

⇨ About 20× slower than EAM

Small system size has large effect on modeling of solidification

MGPT tantalum modeled using ddcMD

64,000 atoms (tiled image) **16,384,000 atoms**

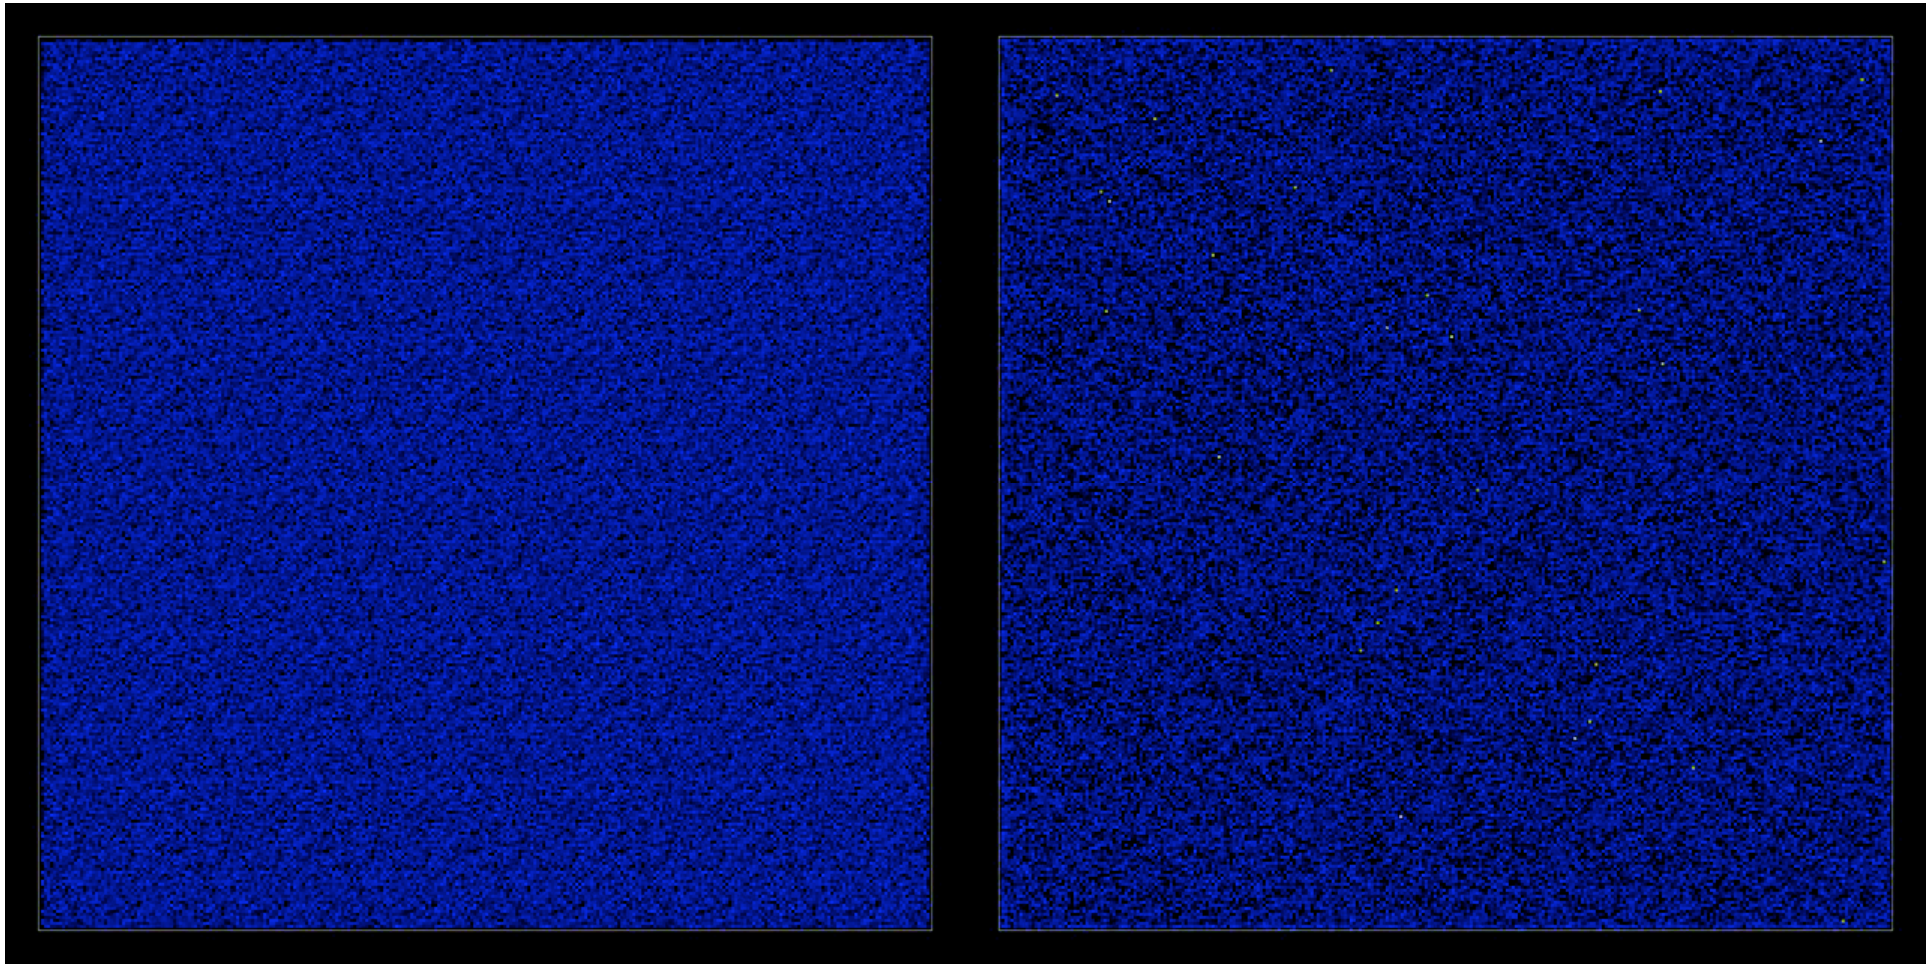

Finite size scaling theory can be applied to cluster sizes at coalescence

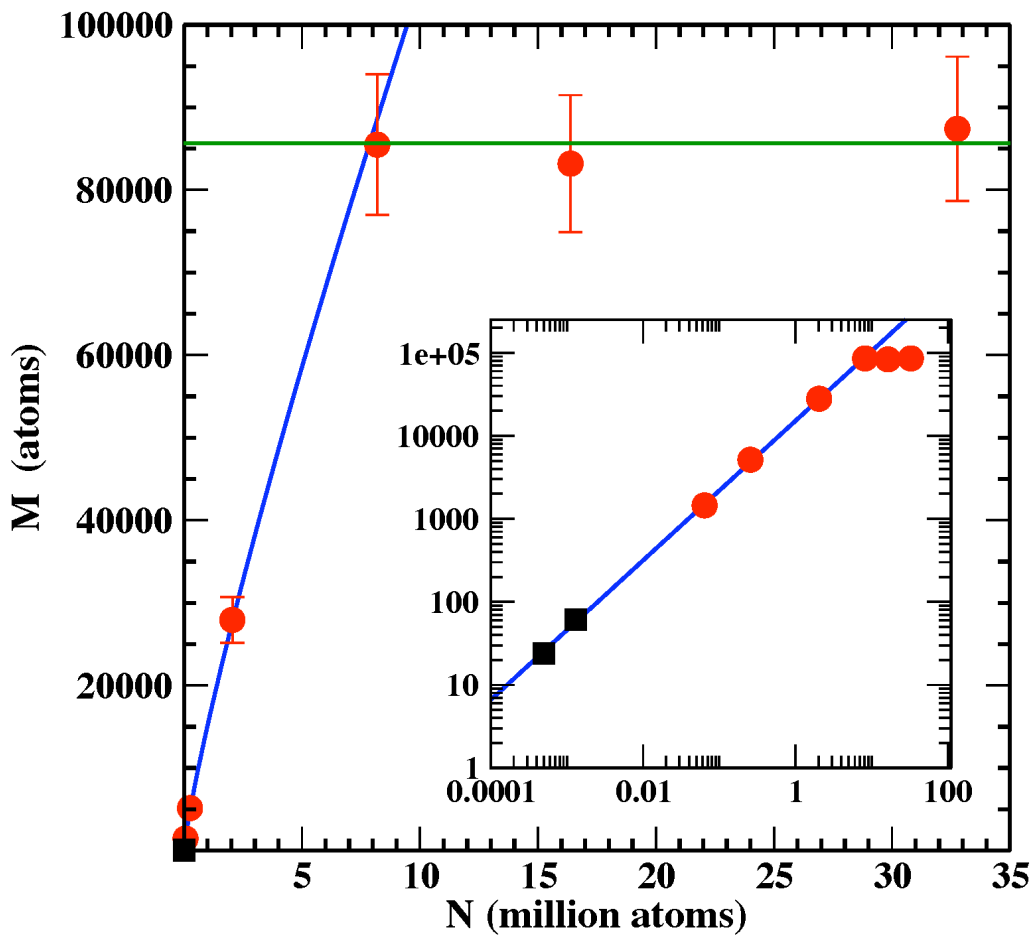

Finite size scaling theory at percolation:

$$M \propto N^{d_f/d}$$

$$d_f = d - \beta/\nu \cong 2.54$$

$$d_f/d \cong 0.84$$

with β=0.41 and ν=0.88 in d=3 dimensions.

By 16M-atoms, cluster size has crossed over to system-size independent value

Summary

- Utilizing up to 128k processors on BlueGene/L, we have successfully resolidified 64k, 256k, 2M, 8M, 16M and 32M-atom simulations during an isothermal compression at 5000 K to a final pressure of approximately 240 GPa.
  - **Our results clearly demonstrate the need for samples larger than several million atoms**
  - **Largest simulations required many CPU-centuries of compute time - we are in a new regime of computational physics!**

- ➠ Weak scaling performance is near perfect - allows simulation of up to several billion MGPT atoms across 64 racks
- ➠ Strong scaling performance is exemplary - enables very long simulations of millions of MGPT atoms
- ➠ We attain overall price/performance of substantially better than 1 TFlop/$M with production code