Transforming Additive Manufacturing through Exascale Computing (ExaAM)

PI: John Turner (turnerja@ornl.gov)
Co-PI: **Jim Belak** (belak@llnl.gov)
www.ExaAM.org (no content yet)
Plus too many to name at ORNL(Babu), LLNL (King, Matthews), LANL(Henson, Carlson) and NIST (Levine)

SOS21, Davos Switzerland, 20-23 March, 02017

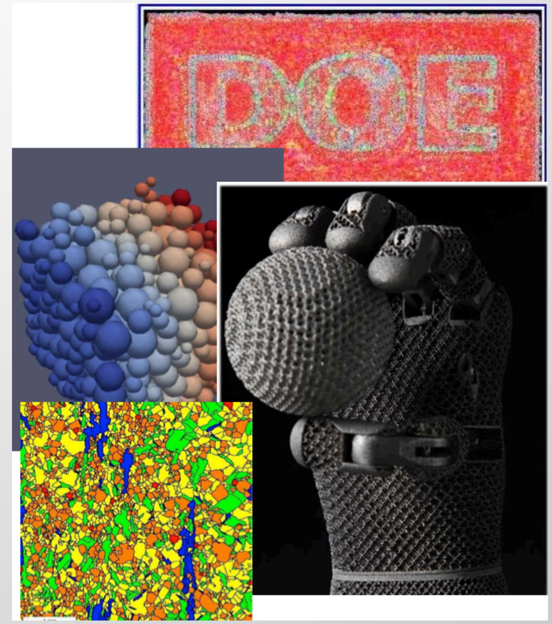

Oak Ridge National Laboratory
Lawrence Livermore National Laboratory
Los Alamos National Laboratory EST.1943
NIST National Institute of Standards and Technology Technology Administration, U.S. Department of Commerce

LLNL-PRES-XXXXXX
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

(from Dan Reed's talk) Something is missing. The first thing we do after making an observation is ask if it is consistent with our framework of understanding.

HPC Simulation of the Experiment (observational science versus experimental science)

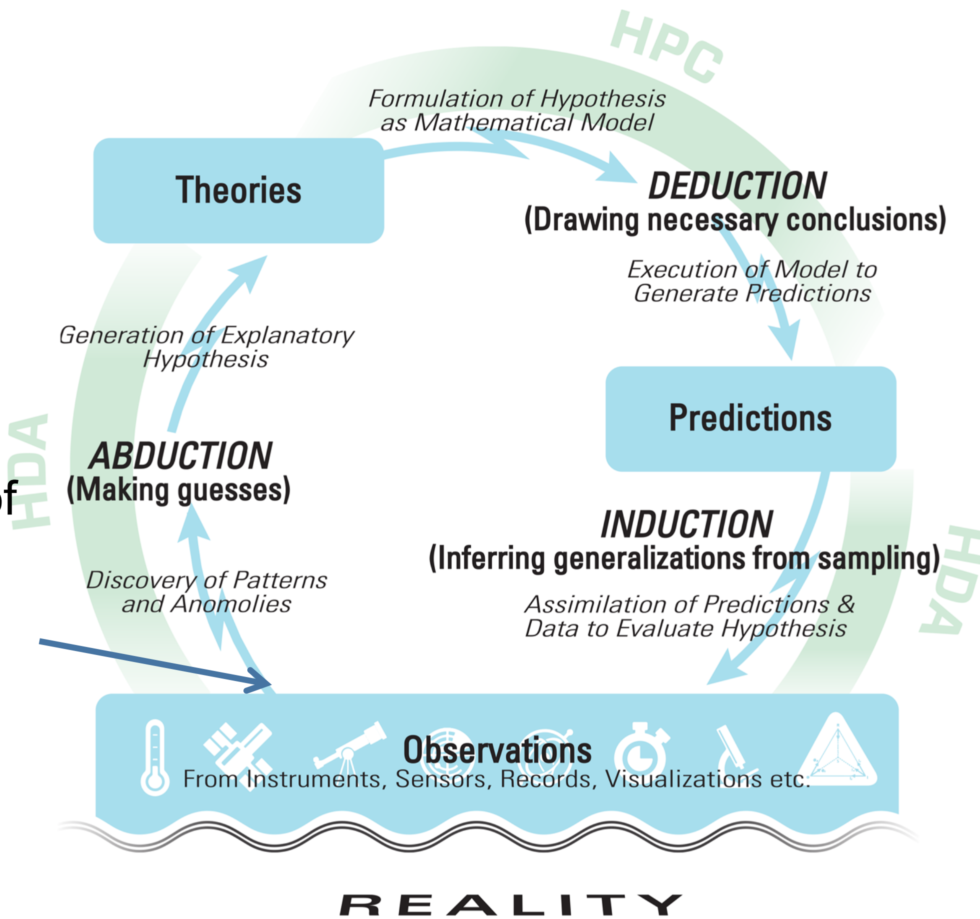

HPC Simulation of the Experiment (design of experiments)

BDEC doc

Eugene Brooks (1987) “Attack of the Killer Micros” (LLNL Massively Parallel Computing Initiative – 1990)
Fastforward 2010+ “Attack of the Killer Cell Phones”

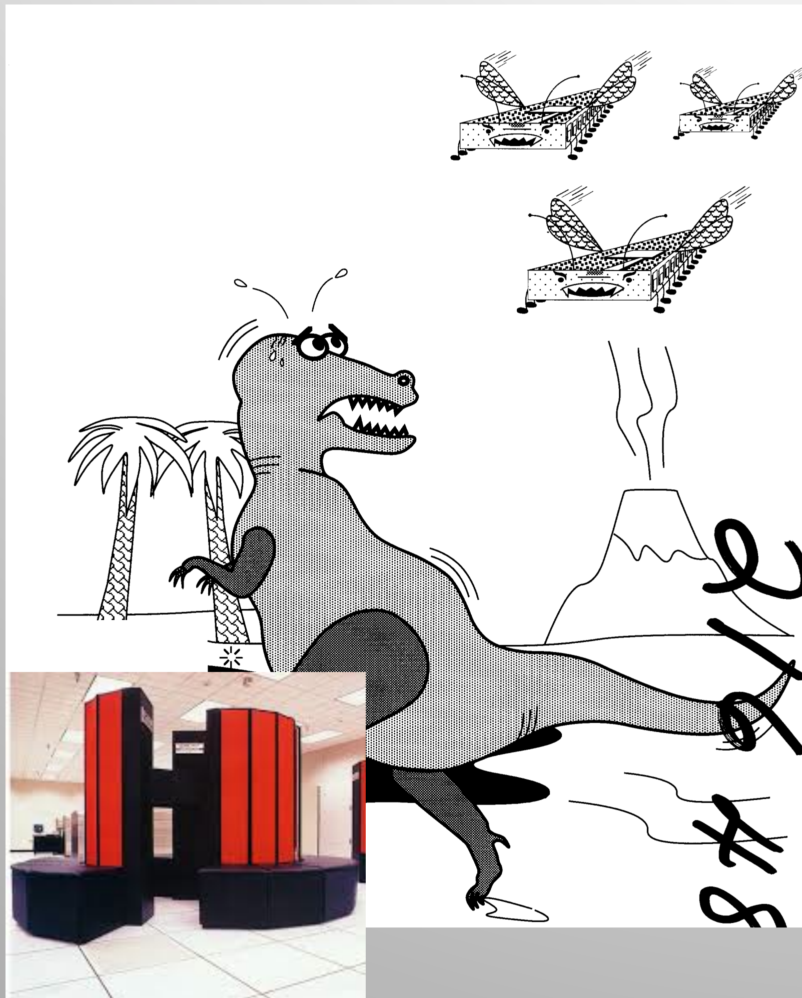

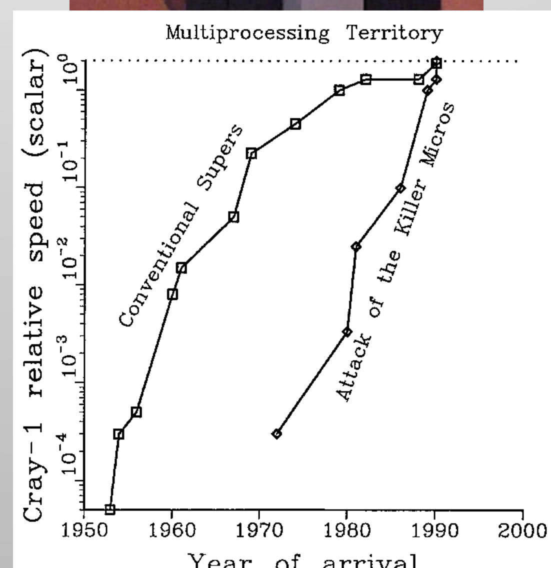

Game Processors

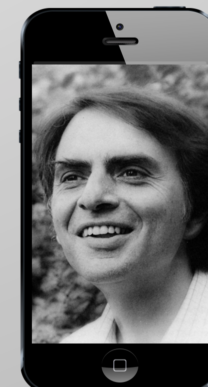

Billions and Billions of Cell Phone Processors

Apology to Mateo Valero (Director Barcelona Supercomputer Center)

ARM-Based Mont Blanc (BSC)

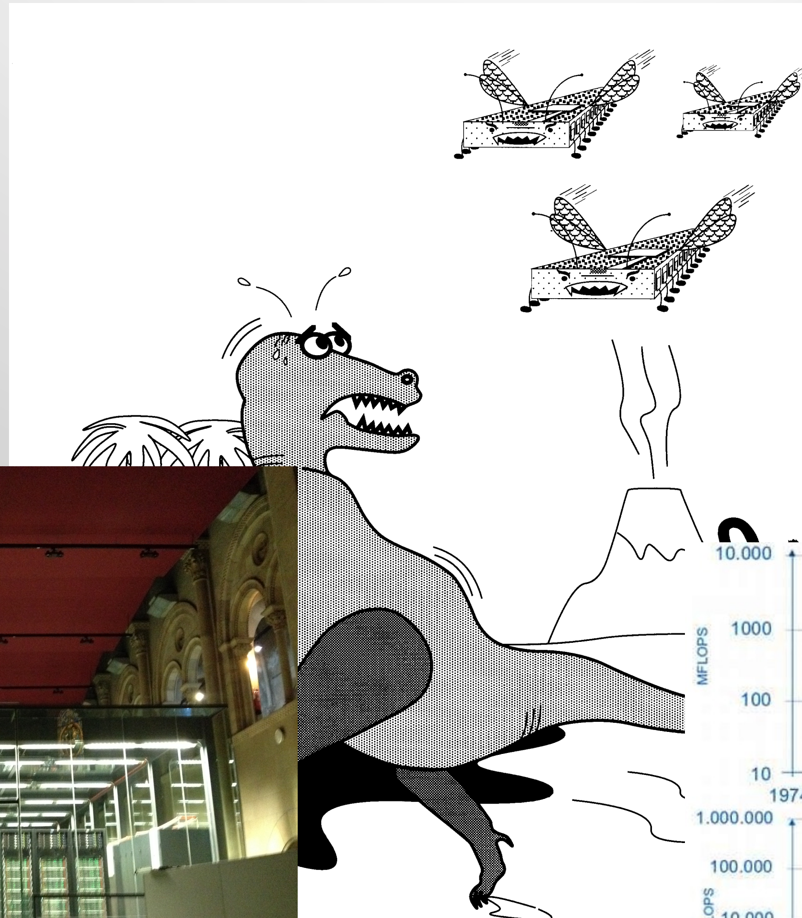

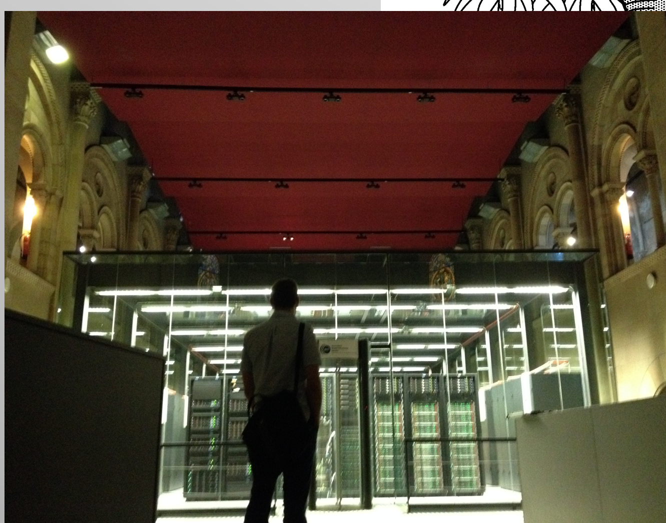

Mare Nostrum (BSC)

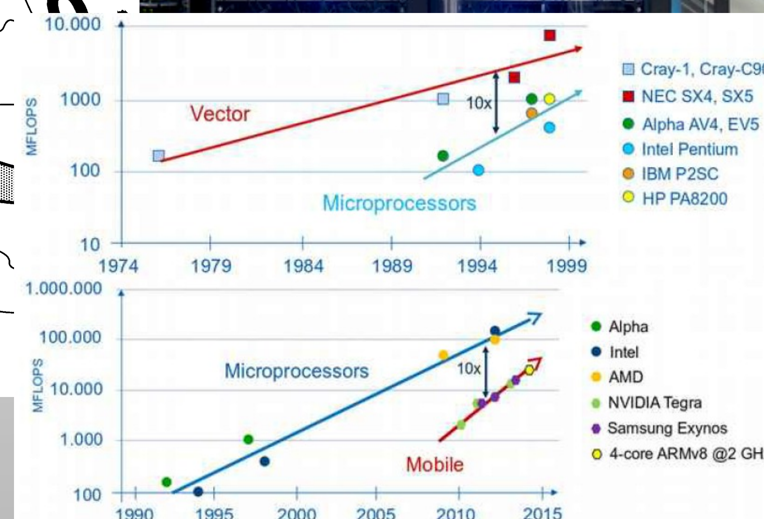

(Crisis) Commonly asked questions in exascale: What is the problem?

Power, energy, and heat dissipation are the central issues. Imagine a computer with **billions and billions** of cell phone processors (14MW) or **millions and millions** of throughput optimized cores, GPGPUs (20MW)

- How do you program it to work on one science problem?
- The architecture will be **heterogeneous** and **hierarchical**, with very high **flop/byte ratios**.
- Single program multiple data bulk synchronous parallelism will no longer be viable.

Data Movement will be expensive and computation will be cheap

- Need to present the physics so the computation occurs where the data is!
- Traditional global checkpoint/restart will be impractical: need local / micro checkpoint (flash memory?)

Simulation codes will need to become **fault tolerant and resilient**

- Recover from soft and hard errors, and anticipating faults
- Ability to drop or replace nodes and keep on running
- The curse of silent errors

What is the Exascale Computing Project (ECP)?

- Created in support of President Obama's National Strategic Computing initiative (NSCI)
- A collaborative effort of two US Dept of Energy (DOE) offices:
  – Office of Science (DOE-SC)
  – National Nuclear Security Administration (NNSA)
- A 10-year project to accelerate the development of a **capable** exascale ecosystem
  – 50x the performance of today's 20 PF/s systems
  – Operates in a power envelope of 20–30 MW
  – Is sufficiently resilient (average fault rate: ≤1/week)
  – Includes a software stack that meets the needs of a broad
    spectrum of applications and workloads
  – Led by DOE laboratories
  – Executed in collaboration with academia and industry

A **capable** exascale computing system will have a well-balanced ecosystem (software, hardware, applications)

ECP has formulated a holistic approach that uses co-design and integration to achieve capable exascale

Application Development	Software Technology	Hardware Technology	Exascale Systems
Science and mission applications	Scalable software stack	Hardware technology elements	Integrated exascale supercomputers

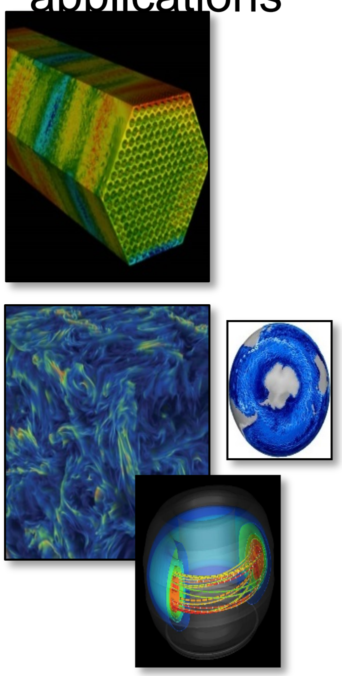

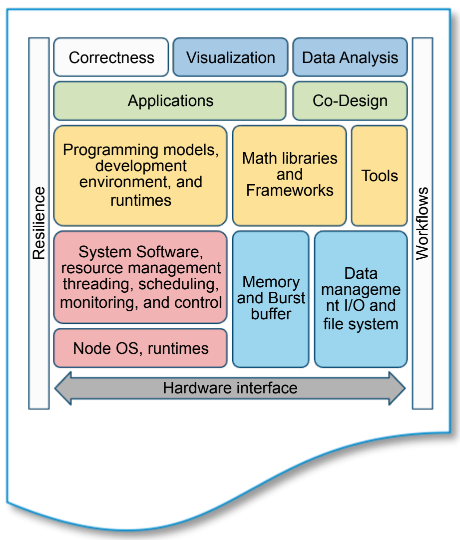

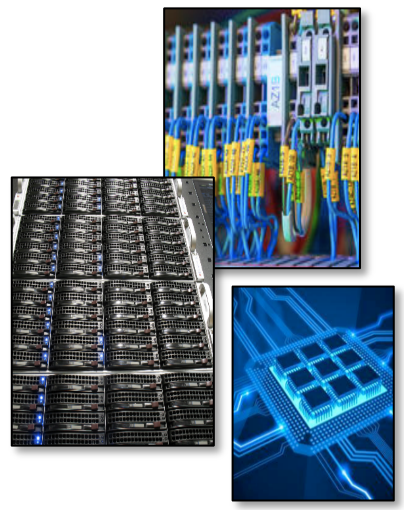

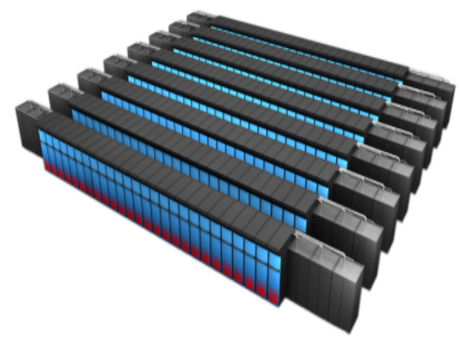

EXASCALE COMPUTING PROJECT

Creating Community Codes (1994-7) LAMMPS CRADA: Cray (Carpenter), Bristol-Myers (Stouch), Dupont (Lustig), LLNL (Belak), Sandia (Plimpton)

Three Essential Elements:

- Problem Specification
- Target Architecture
- Programming Model

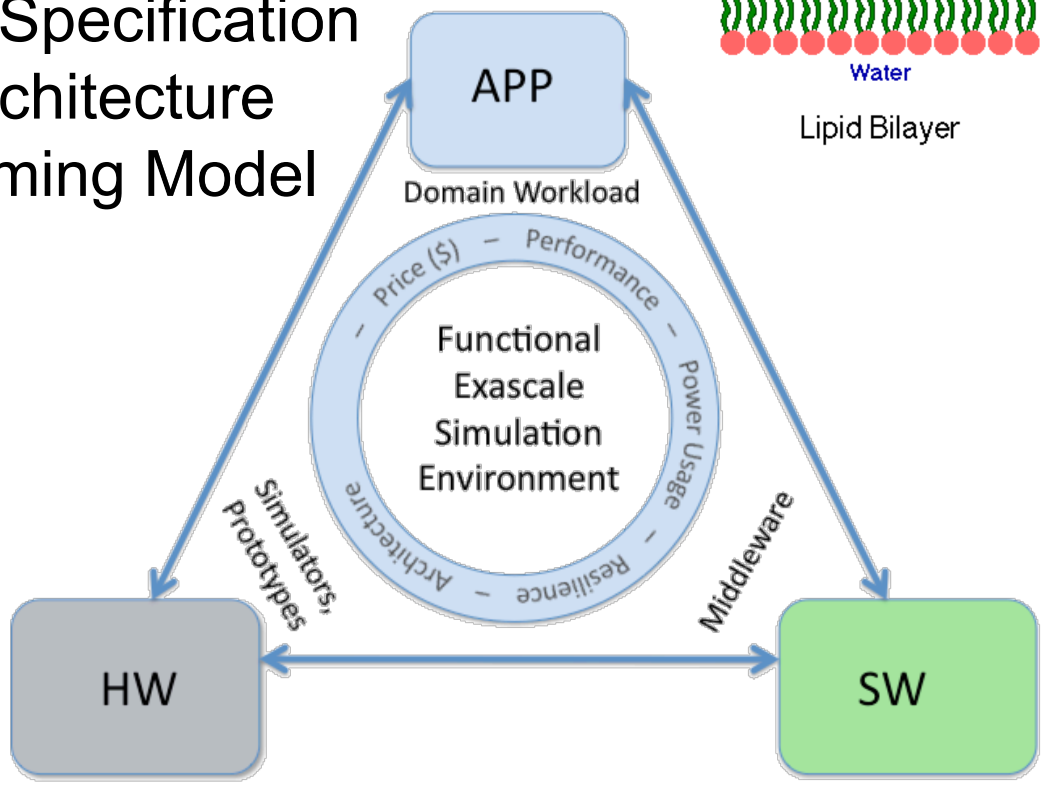

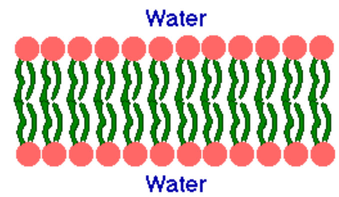

Lipid Bilayer

Diffusion of a small drug molecule through a bio-membrane

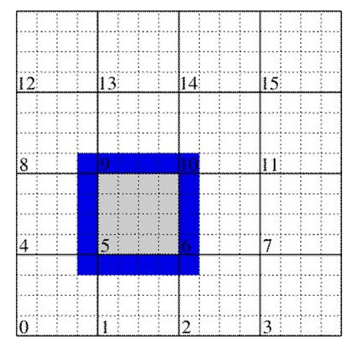

Cray T3D, Clock Speed: 0.15GHz, Date: 1994
Microprocessor Peak Teraflops: 0.02, Memory: 0.01TB
Processors: 128. Power Consumption: 41.40 kW

SPMD, MPI, Geometric domain decomposition

1997: LAMMPS became a successful community code because of Steve's idea to release it as open source, the ubiquitous Beowolf cluster and the stability of MPI.

- 20 Years Ago
  - Hardware: Linux Beowolf Cluster
  - Software Programming Model: MPI/SPMD
  - Application Code: LAMMPS
- ~5 Years in the Future
  - Hardware: Exascale? (Petascale in a Rack) Challenge: a scaled exascale computer with today's technology?
  - Software Programming Model: MPI+X, Q? Task-based with "Control" of Data Locality (User and Runtime)
  - Application Codes???

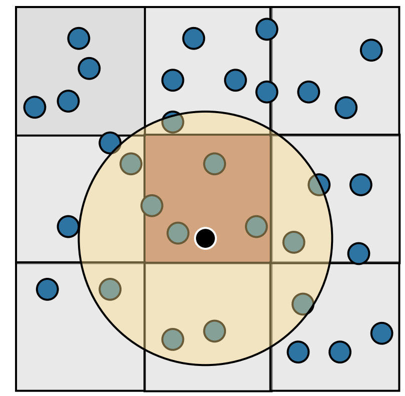

Worker

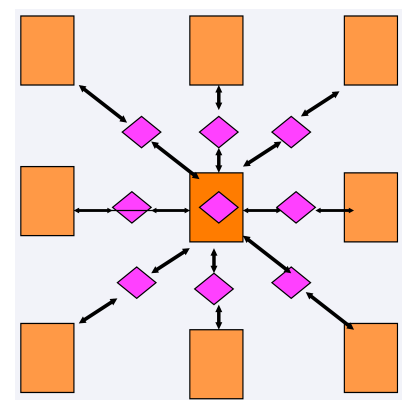

Work

Opportunity: Ubiquitous Peta-scale computing, in Industry, in Labs, in Academia.

Our Exascale AM Project is one of 15 initial ECP application development projects (plus 7 seeds)

Exascale Application Driver

Oak Ridge National Laboratory | Lawrence Livermore National Laboratory | Los Alamos National Laboratory | NIST National Institute of Standards and Technology, Technology Administration, U.S. Department of Commerce

Advanced Manufacturing

Gaps and Opportunities

- ✓ Advance quality, reliability, and application breadth of additive manufacturing (AM)
- ✓ Accelerate innovation in clean energy manufacturing institutes (NNMIs)
- ✓ Capture emerging manufacturing markets

Simulation Challenge Problems

- ✓ Continuum level predictions of non-uniform microstructure and its relationship to process parameters
- ✓ Predictive mesoscale models for dendritic solidification then scale-bridged to continuum

Prospective Outcomes and Impact

- ✓ Routine qualification of AM parts via process-aware design specs and reproducibility through process control
- ✓ Fabrication of metal parts with unique properties such as light weight strength and failure-proof joints and welds

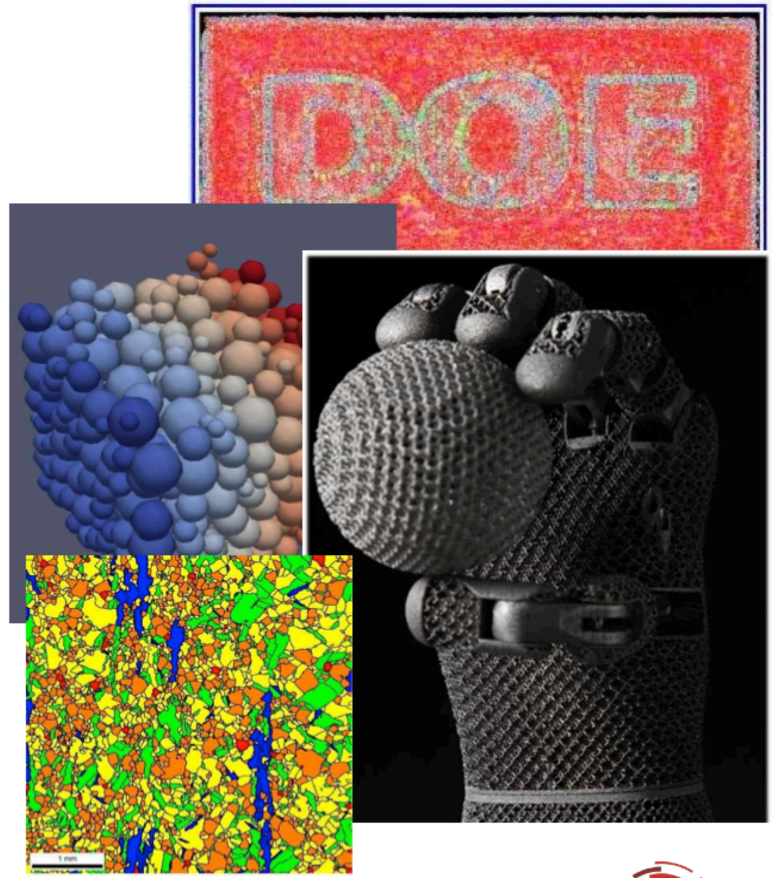

ECP EXASCALE COMPUTING PROJECT

Additive Manufacturing, a.k.a. 3D printing, is being used for metal as well as polymers

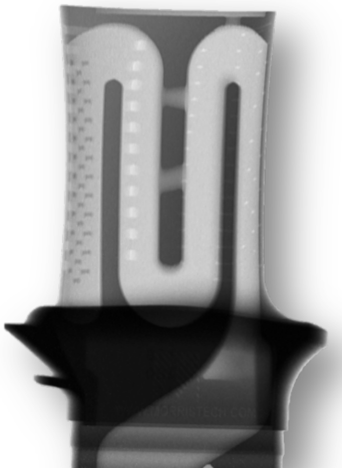

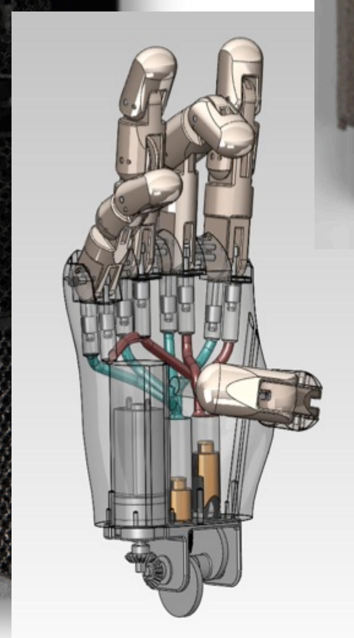

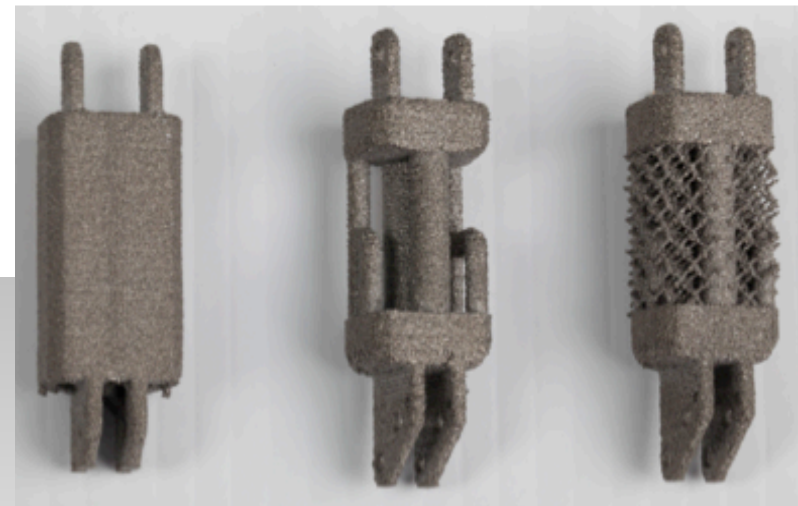

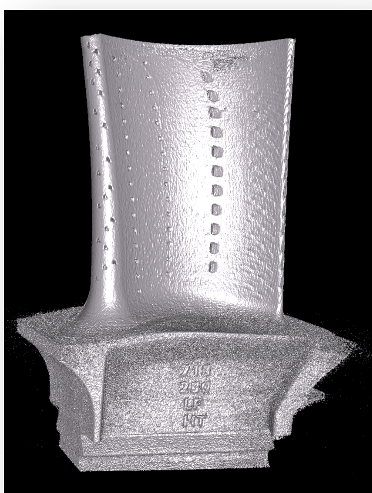

21.1 g   12.1 g   14.4 g

Multiple metal AM technologies

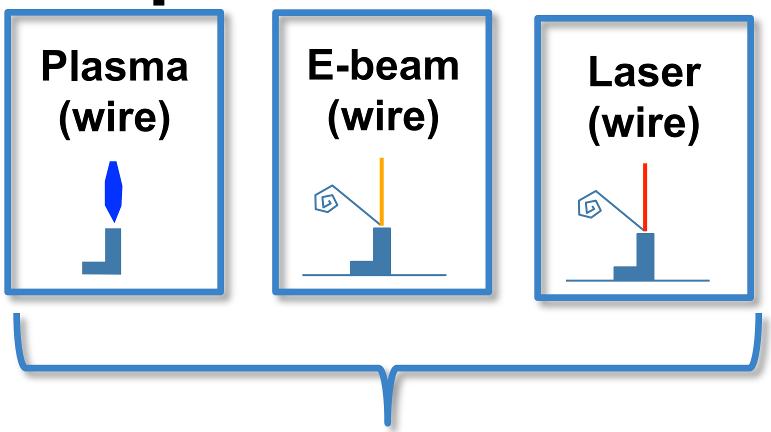

Large Melt Pool Technologies

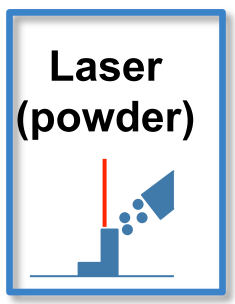

Direct Metal Deposition

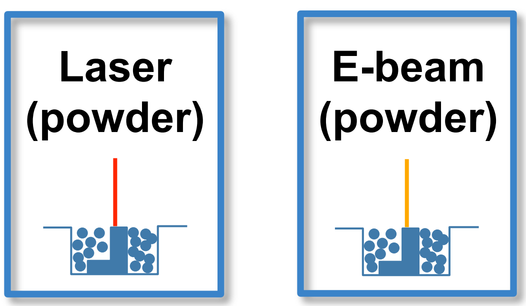

Powder Bed Technologies

Physical processes are similar

- Energy Deposition
- Melting & Powder Addition
- Evaporation & Condensation
- Heat & Mass Transfer
- Solidification
- Solid-State Phase Transformation
- Repeated Heating and Cooling
- Complex Geometries

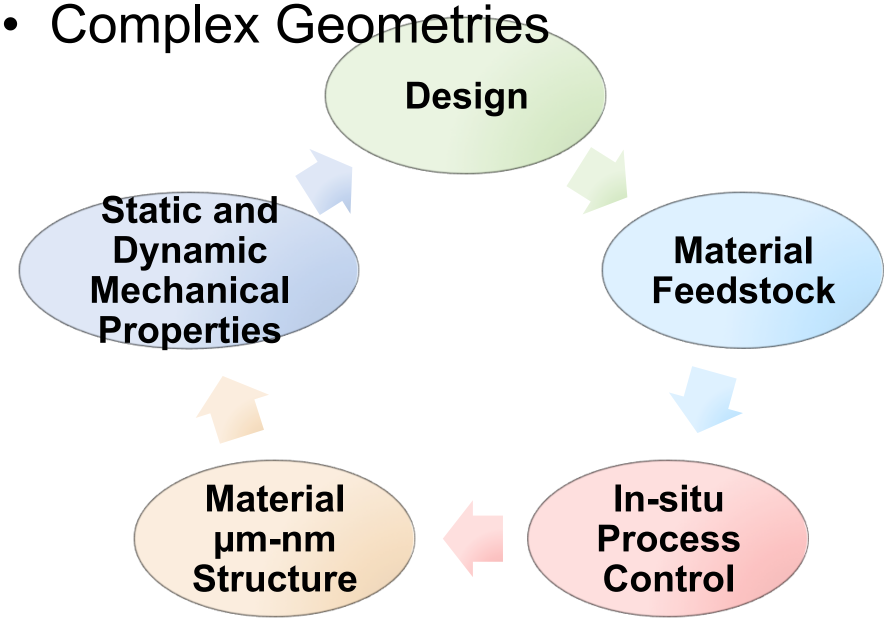

Multiple computational challenges must be addressed for AM

- 1 m^3 ~ 10^{12} particles ~ 10^9 m of "weld" line (assuming 50μm particles) and build times of hours
- Large temperature gradients, rapid heating and cooling
  - necessary / sufficient coupling between thermomechanics and melt/solidification
- Heterogeneous and multi-scale
  - resolution of energy sources and effective properties of powder for continuum simulations
- Path optimization
- Large number of parameters and missing understanding
  - key uncertainties and propagation of those uncertainties
- Validation is difficult as characterization is limited

How are we going to address the problem? Workflow for Design Optimization of Additively Manufactured Parts (Integrated Computational Engineering)

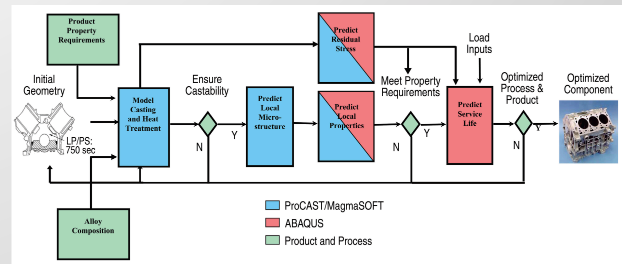

John Allison,JOM 2006

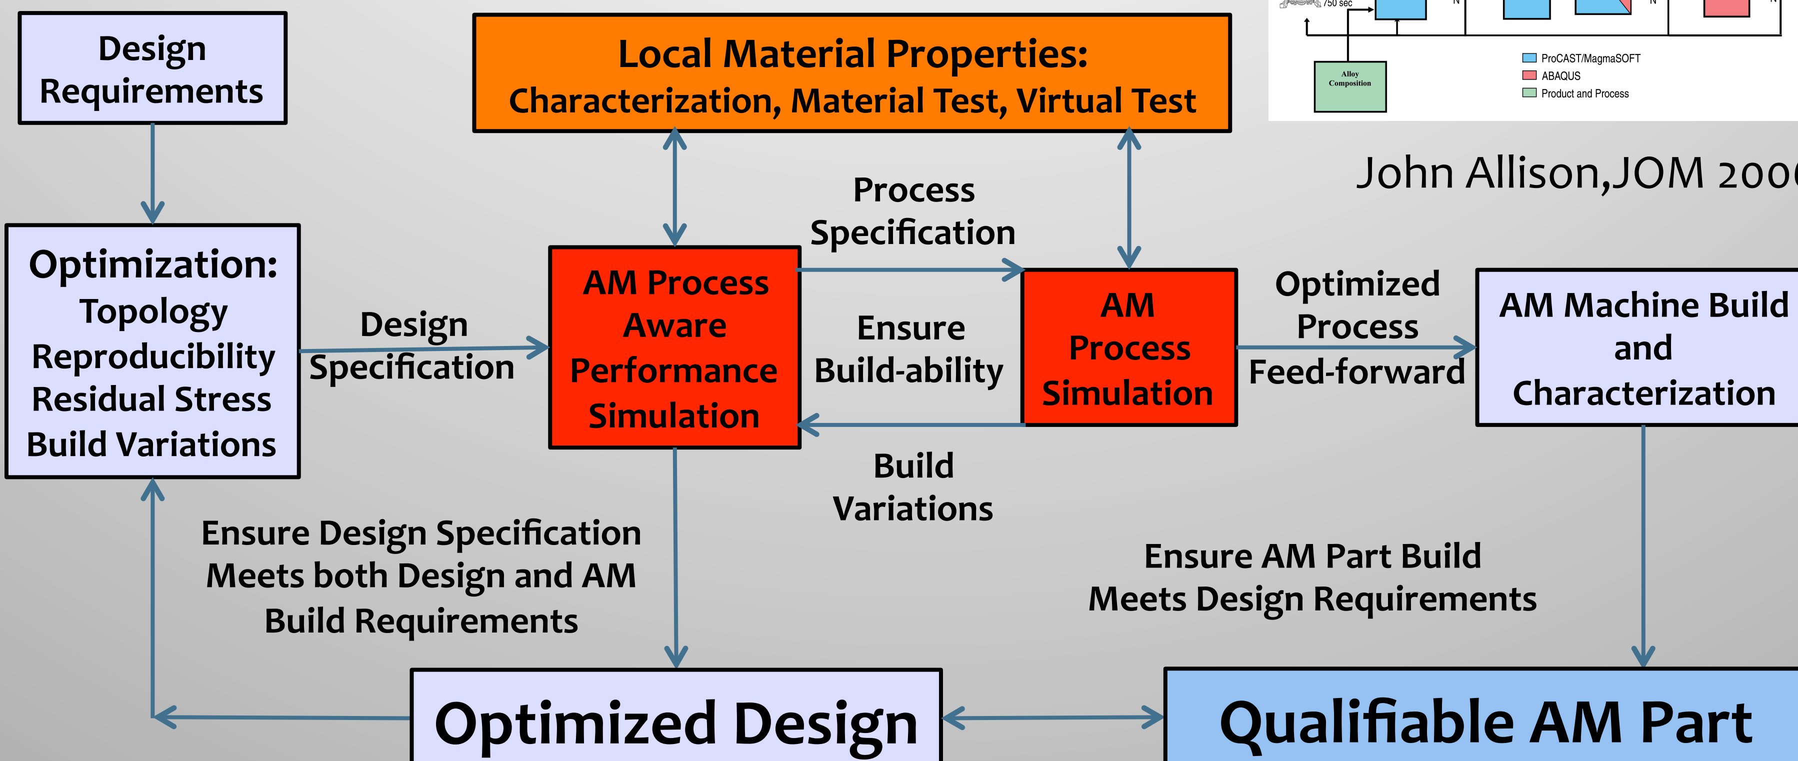

Intra-HPC Workflow (ExaAM)

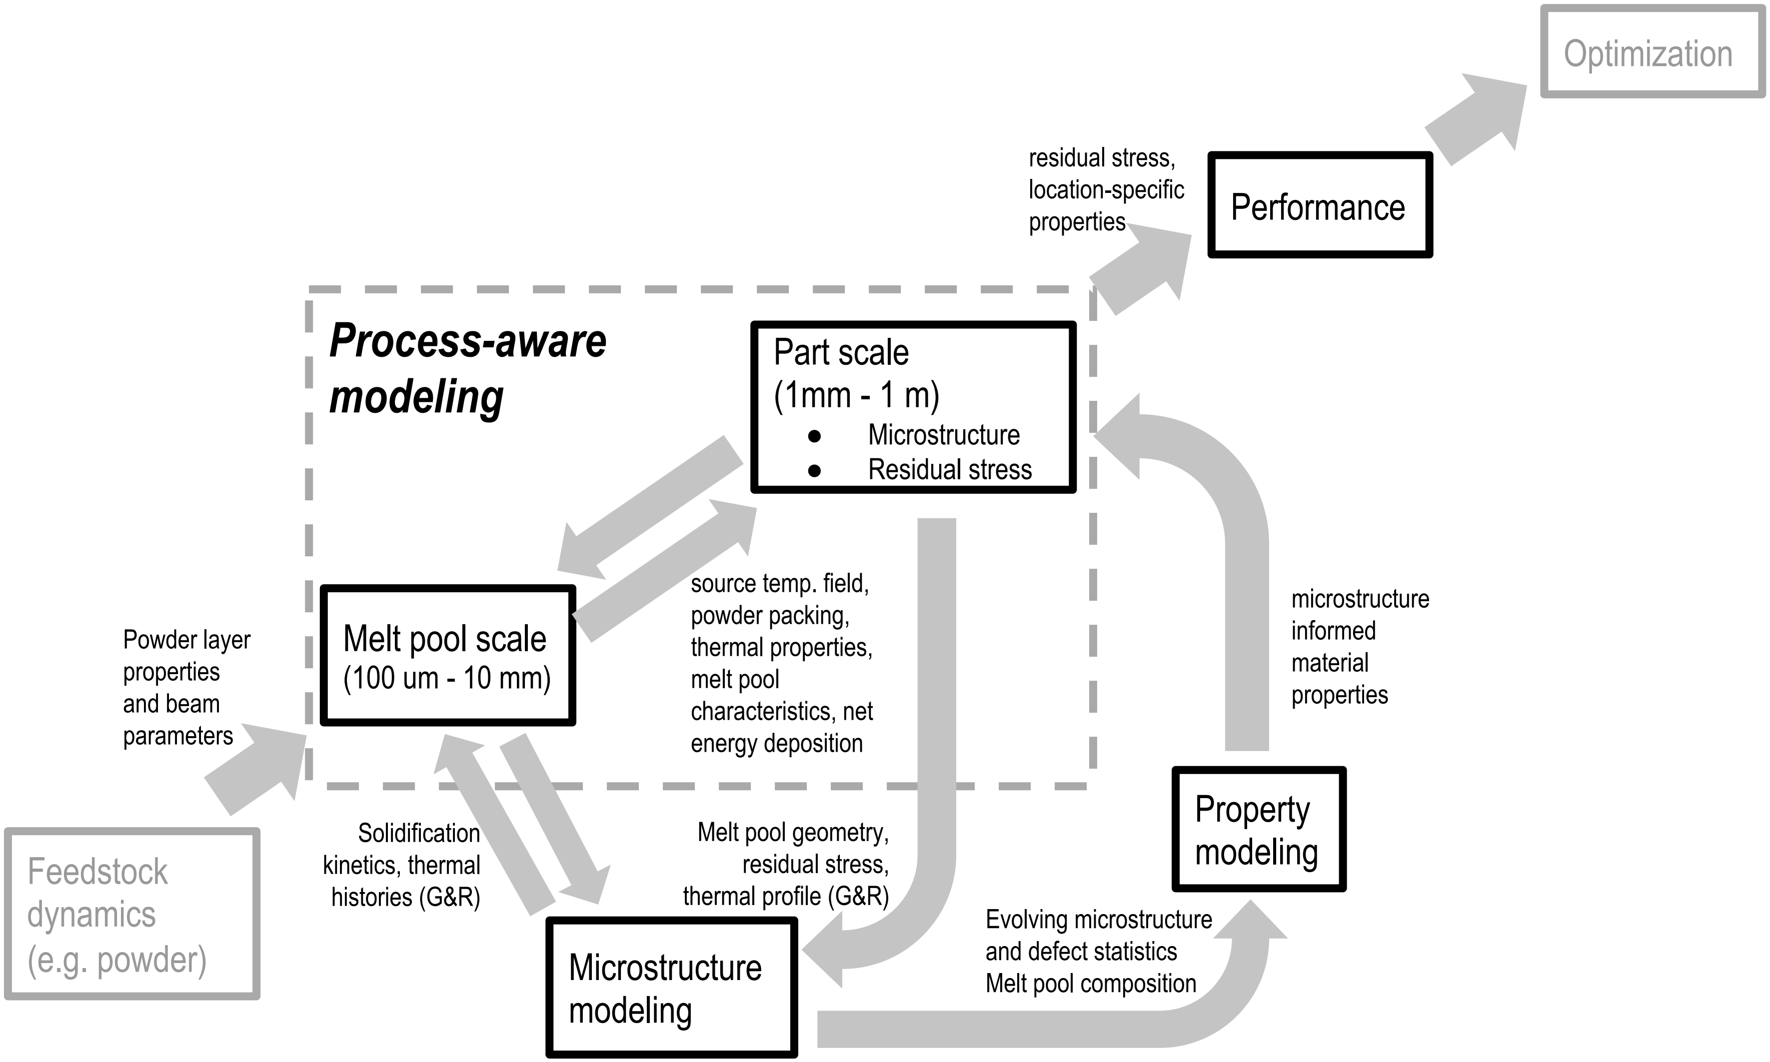

Mesoscopic 3D simulations provide insight into possible surface finish improvement options

Khairallah, S.A., Anderson, A., 2014.

Mesoscopic Simulation Model of Selective Laser Melting of Stainless Steel Powder.

Journal of Materials Processing Technology 214, 2627-2636 DOI. 10.1016/j.jmatprotec.2014.06.001.

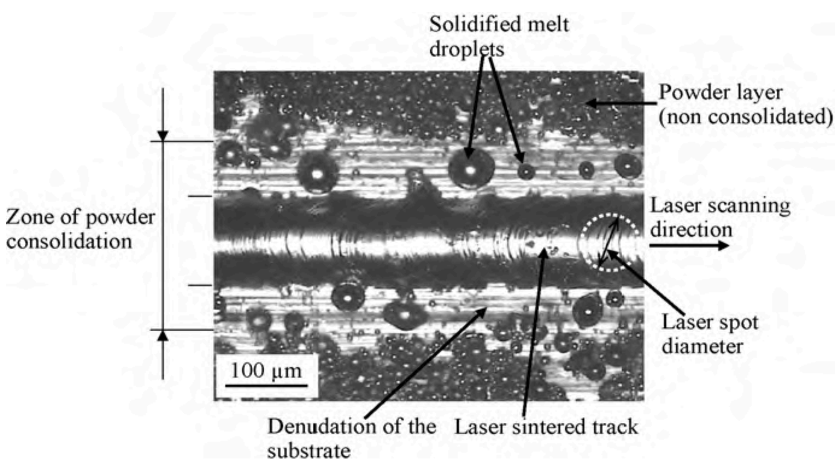

Yadroitsev, I., Gusarov, A., Yadroitsava, I., Smurov, I., 2010.
Single track formation in selective laser melting of metal powders.
Journal of Materials Processing Technology 210, 1624-1631.

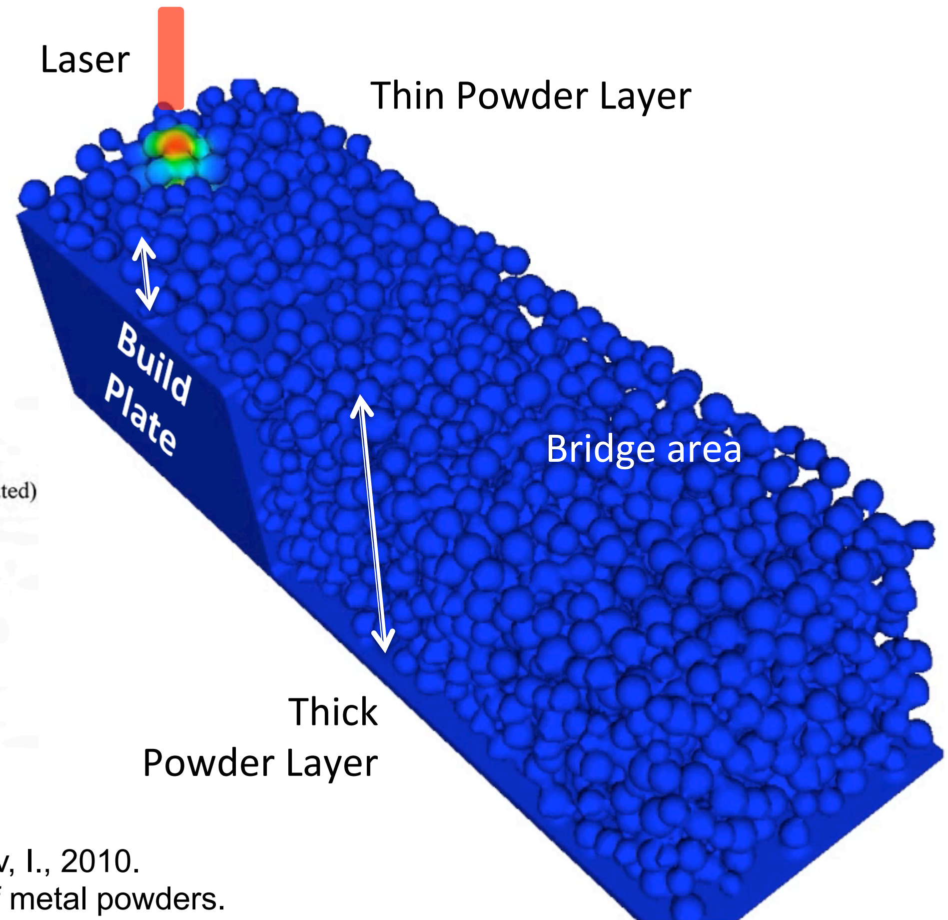

The effective medium model has been used to predict the development of residual stresses

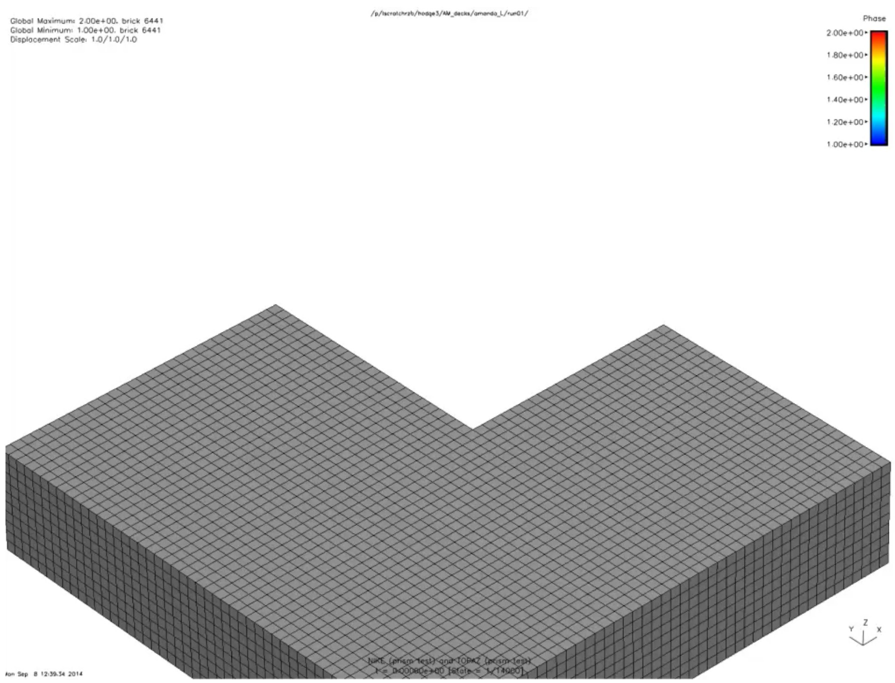

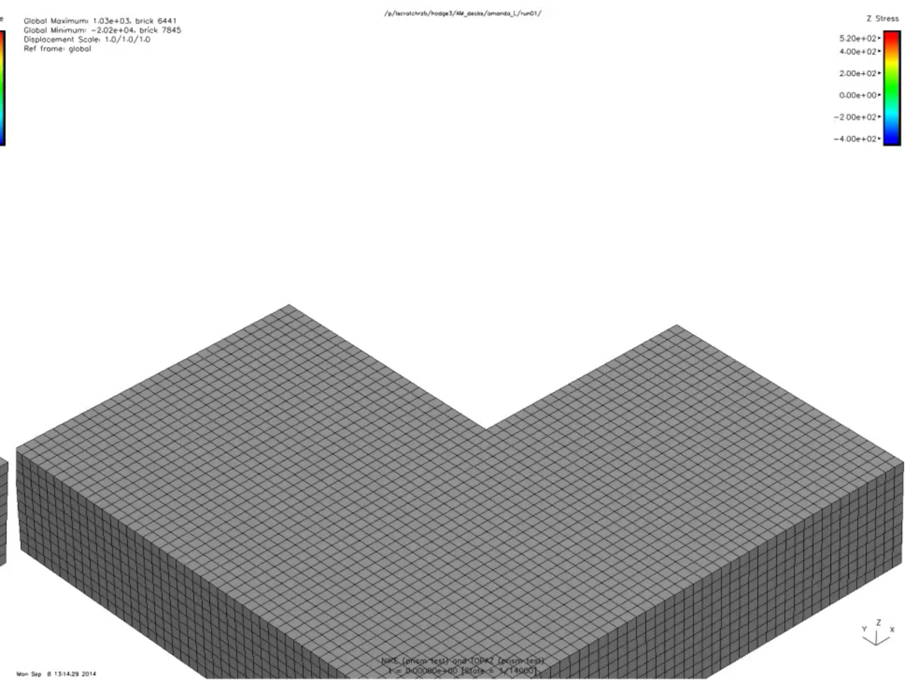

Hodge, N.E., Ferencz, R.M., Vignes, R.M., 2016. Experimental Comparison of Residual Stresses for a Thermomechanical Model for the Simulation of Selective Laser Melting. Additive Manufacturing DOI. http://dx.doi.org/10.1016/j.addma.2016.05.011.

Simulation compares well with digital image correlation combined with neutron diffraction

- Vertical stress assessed while on the build plate
- Perimeter values assessed from digital image correlation at horizontal cut

Experiments | **Diablo**

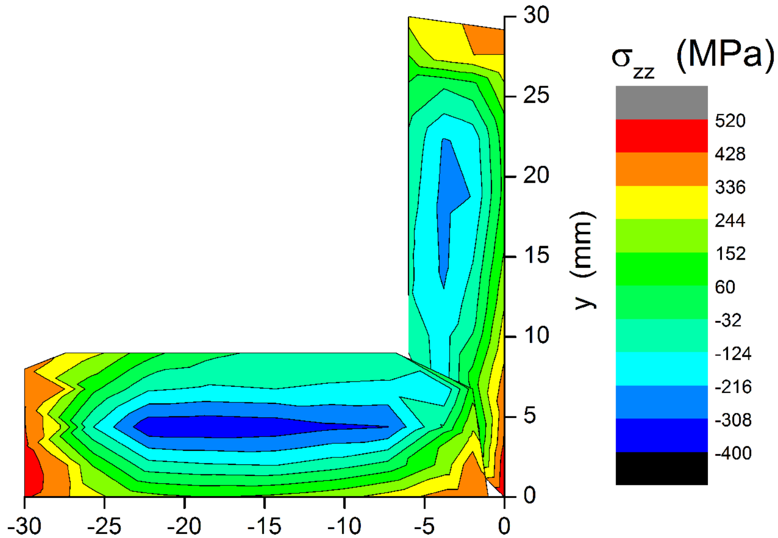

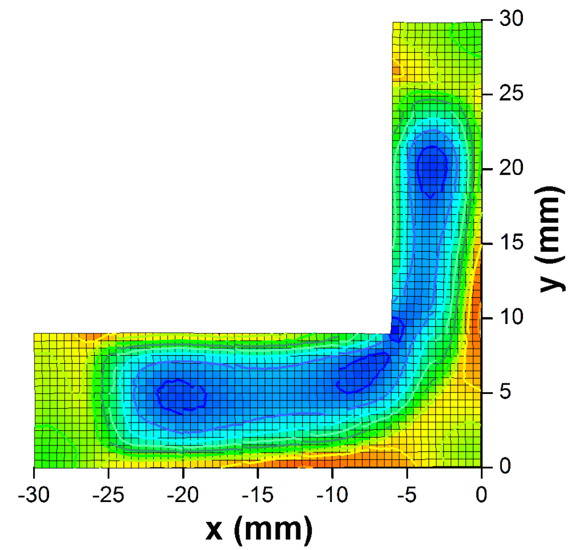

σ_{zz} normal to plane 15mm above base

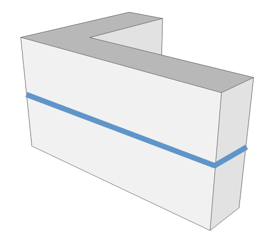

Hodge, N.E., Ferencz, R.M., Vignes, R.M., 2016. Experimental Comparison of Residual Stresses for a Thermomechanical Model for the Simulation of Selective Laser Melting. Additive Manufacturing DOI. http://dx.doi.org/10.1016/j.addma.2016.05.011.

We can model residual stresses and design for them

Exascale is about better Physics Fidelity: Coupling Atomistic with Microstructural Scales through Overlapping Simulations

Ab-initio	Atoms	Long-time	Microstructure	Dislocation	Crystal	Continuum
Inter-atomic forces, EOS, excited states	**Defects and interfaces, nucleation**	**Defects and defect structures**	**Meso-scale multi-phase, multi-grain evolution**	**Meso-scale strength**	**Meso-scale material response**	**Macro-scale material response**
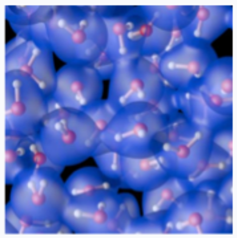	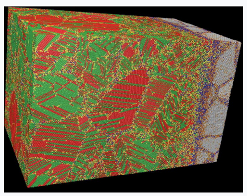	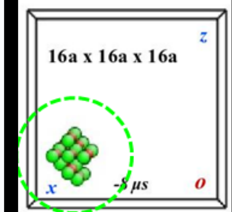	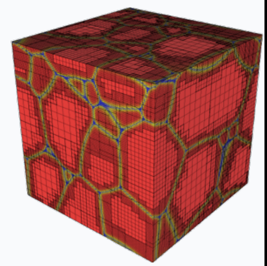	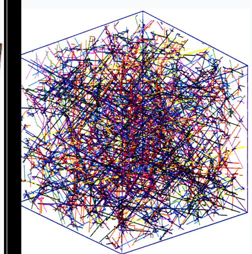	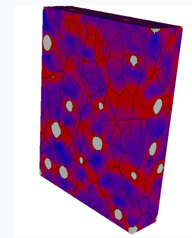	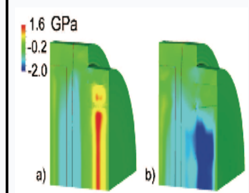
Code: Qbox/ LATTE Motif: Particles and wavefunctions, plane wave DFT, ScaLAPACK, BLACS, and custom parallel 3D FFTs Prog. Model: MPI + CUBLAS/CUDA	Code:SPaSM/ ddcMD/CoMD Motif: Particles, explicit time integration, neighbor and linked lists, dynamic load balancing, parity error recovery, and *in situ* visualization Prog. Model: MPI + Threads	Code: SEAKMC Motif: Particles and defects, explicit time integration, neighbor and linked lists, and *in situ* visualization Prog. Model: MPI + Threads	Code: AMPE/GL Motif: Regular and adaptive grids, implicit time integration, real-space and spectral methods, complex order parameter Prog. Model: MPI	Code: ParaDiS Motif: "segments" Regular mesh, implicit time integration, fast multipole method Prog. Model: MPI	Code: VP-FFT Motif: Regular grids, tensor arithmetic, meshless image processing, implicit time integration, 3D FFTs. Prog. Model: MPI + Threads	Code: ALE3D/ LULESH Motif: Regular and irregular grids, explicit and implicit time integration. Prog. Model: MPI + Threads

Rapid Solidification: Molecular Dynamics (MD) are now large enough to model the initiation of realistic microstructure

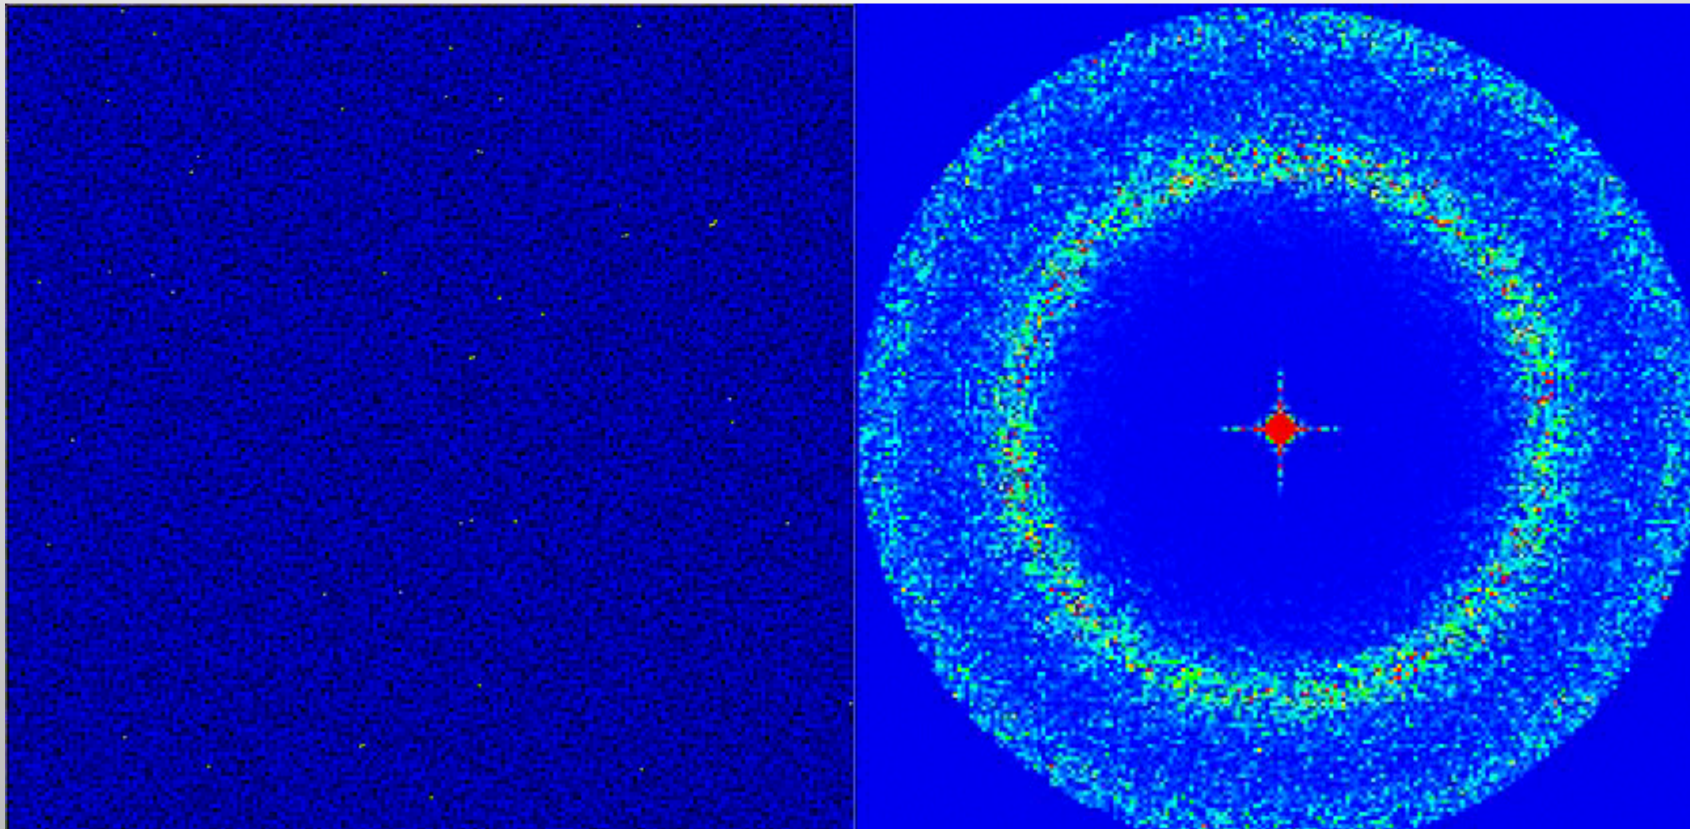

Simulations suggest novel in situ x-ray scattering experiments using emerging sources such as LCLS

Multi-scale paradigm: Phase-field model and MD simulations that overlap in space and time

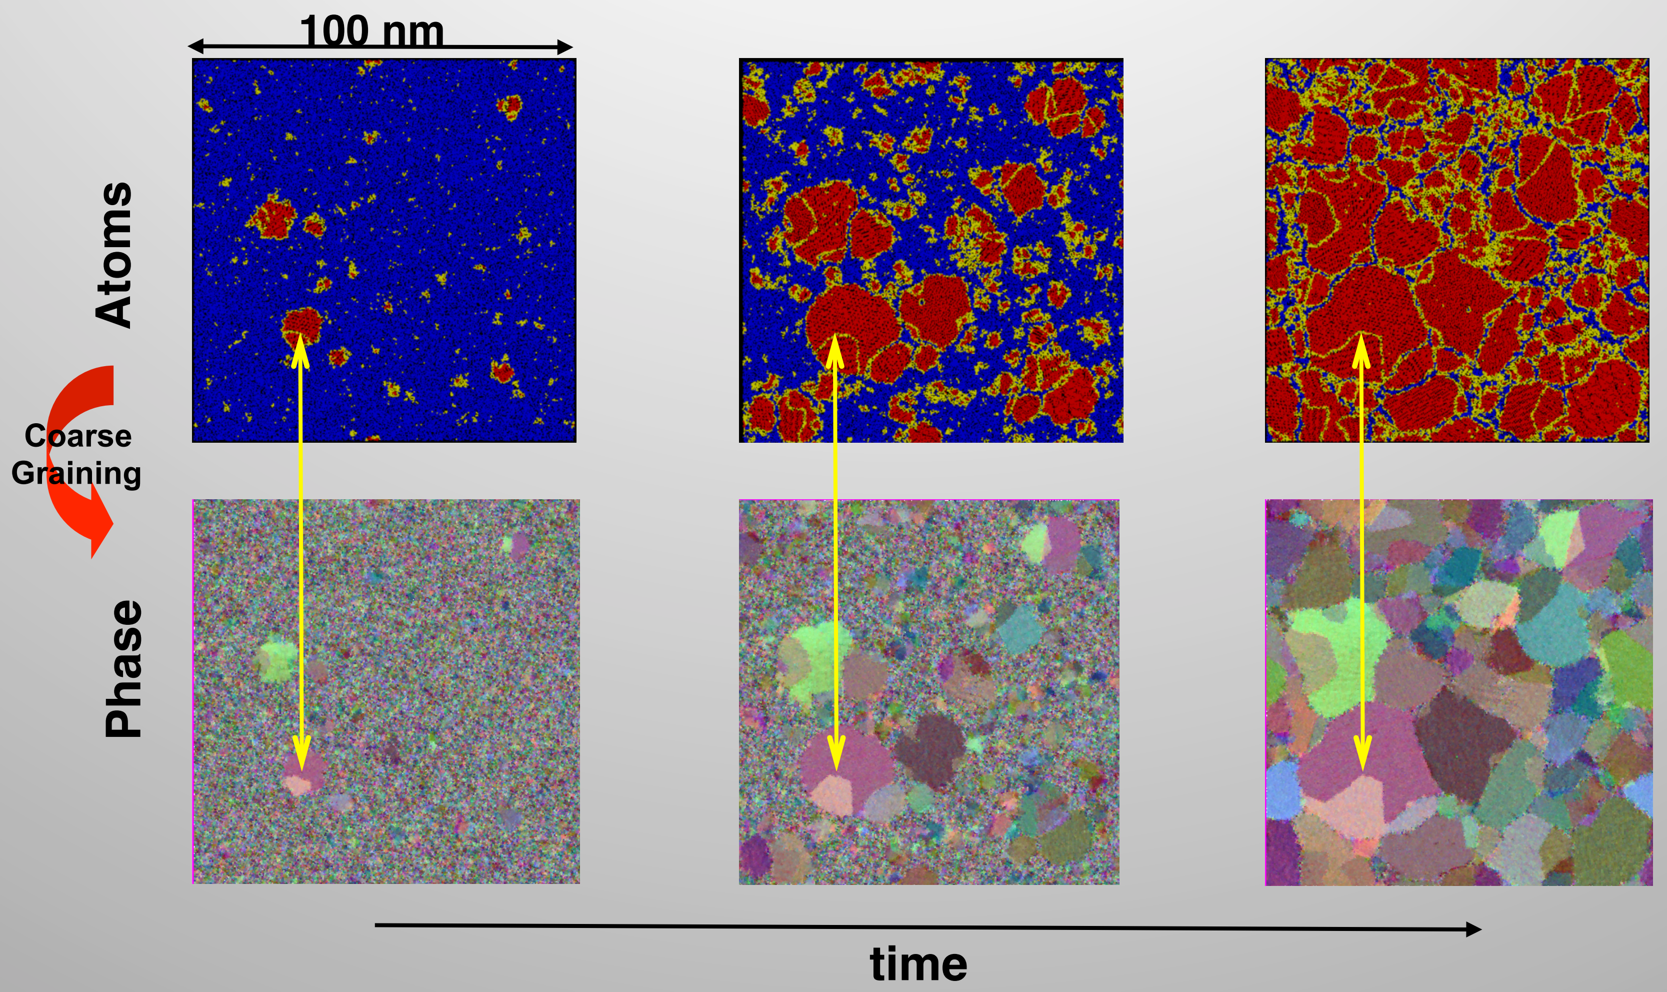

What is Phase Field Modeling (PFM)?

$\left|\vec{\phi}(r,t)\right| = liquid$

$\left|\vec{\phi}(r,t)\right| = solid$

$\vec{\phi}(r,t) = green$

$\vec{\phi}(r,t) = pink$

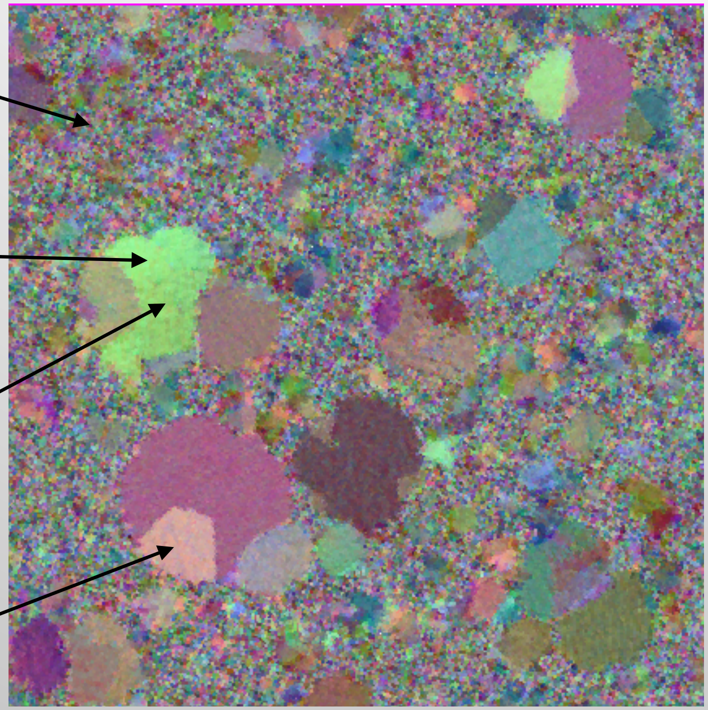

- **Each color represents a different value of the phase field $\vec{\phi}$ (solid orientation)**
- **Free energy describes how colors interact and evolve**
- **Accuracy depends on fidelity of physics in the equations**

Thermodynamic representation of phase (or "color") everywhere

Evolution Equations

$$F(P,T) = \int dx \left\{ \left|\nabla\vec{\phi}\right|^2 + f(\vec{\phi},P,T) + \ldots \right\}$$

$$\frac{\partial\vec{\phi}}{\partial t} = -\Gamma \frac{\delta F}{\delta\vec{\phi}} + noise$$

What does a crystallographic-aware phase-field model of polycrystal solidification look like?

Pusztai et al., have proposed a 3D quaternion-based phase-field model

- Represents crystal orientation with quaternion order parameter
- Quaternions are widely used to analyze crystallography of polycrystal interfaces
- Quaternion algebra is fast, efficient, avoids singularities, …

Free Energy

$$F = \int \left[\frac{\varepsilon_\varphi^2}{2} |\nabla \varphi|^2 + f(\varphi, c, T) + HT[1 - p(\varphi)] \left(\sum_i (\nabla q_i)^2 \right)^{1/2} \right] d^3 r$$

Evolution

$$\frac{\partial q_i}{\partial t} = -M_q \frac{\delta F}{\delta q_i} + \zeta_i = M_q \left[\nabla \cdot \left(D \frac{\nabla q_i}{|\nabla q_i|} \right) - 2\lambda q_i \right] + \zeta_i$$

Where q_i is the quaternion order parameter, M_q is the associated mobility and ζ is the fluctuation in q.

We have implemented the Pusztai model in our 3D AMR code

- Enhance energy functional to represent energetics of grain boundaries
- Crystal symmetry aware quaternion mathematics
- Extend energy functional to include elasticity and alloy concentration

Refs: T. Pusztai, G. Bortel, and L. Granasy, “Phase field theory of polycrystalline solidification in three dimensions,” Europhys. Lett, 71 (2005) 131-137; Dorr, M.R., Fattebert, J.-L., Wickett, M.E., Belak, J.F., and Turchi, P.E.A., “A Numerical Algorithm for the Solution of a Phase-field Model of Polycrystalline Materials,” J. Comp. Phys. Vol. 229, 626 (2010).

Representation of MD Data onto the AMR Grid Hierarchy using the SAMRAI AMR Library

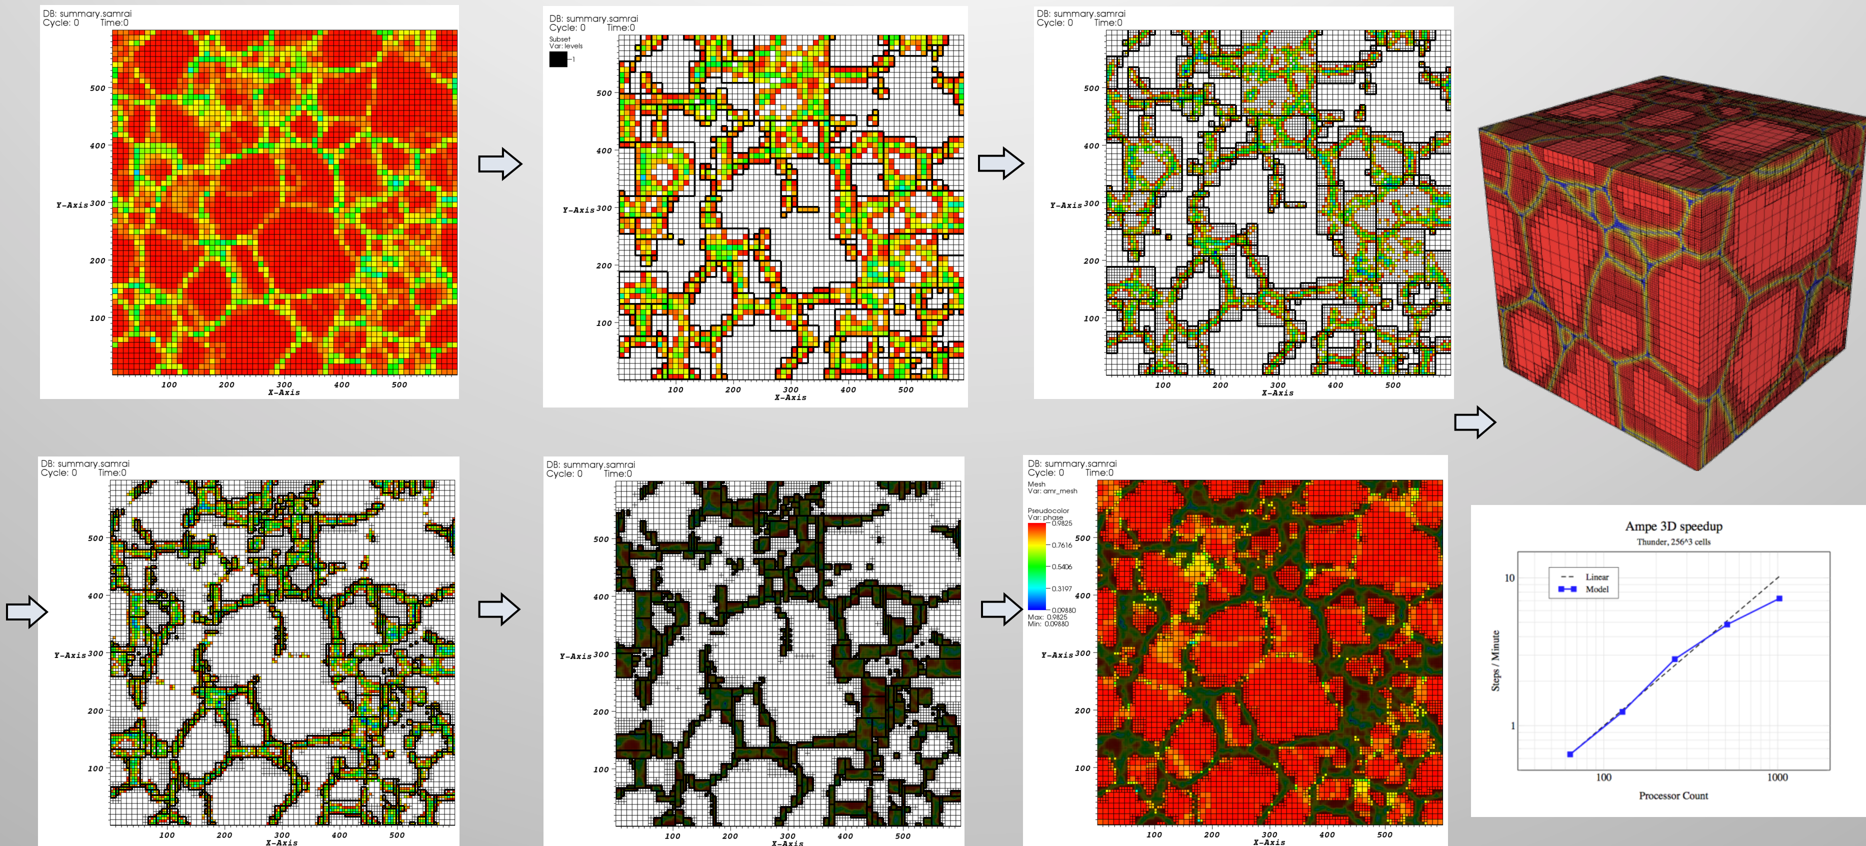

MD nucleated microstructure onto the micro-second hydro time-scale with the crystallographic quaternion model

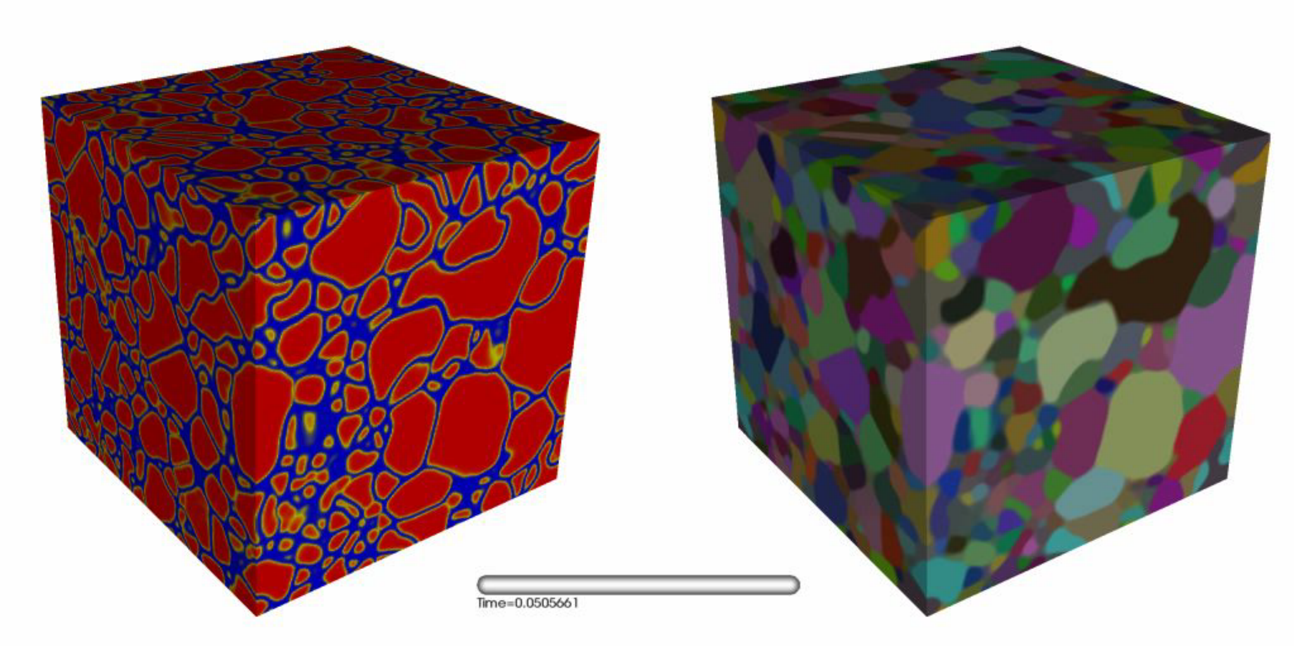

Phase Order Parameter

Quaternion Order Parameter

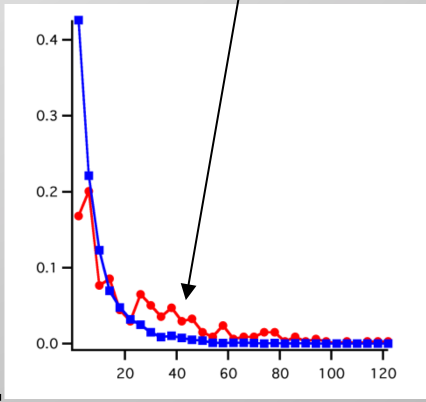

While significant grain coarsening has occurred on the microsecond scale, the microstructure is far from log-normal

ExMatEx and Materials Science Workflow: The 7 Pillars of Materials Science and Adaptive Physics Refinement

Ab-initio	Atoms	Long-time	Microstructure	Dislocation	Crystal	Continuum
Inter-atomic forces, EOS, excited states	**Defects and interfaces, nucleation**	**Defects and defect structures**	**Meso-scale multi-phase, multi-grain evolution**	**Meso-scale strength**	**Meso-scale material response**	**Macro-scale material response**
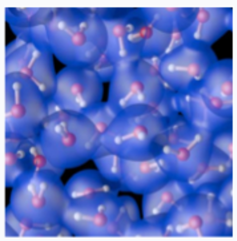	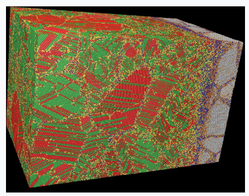	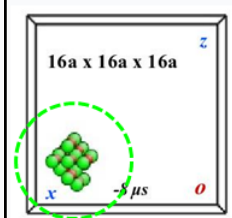	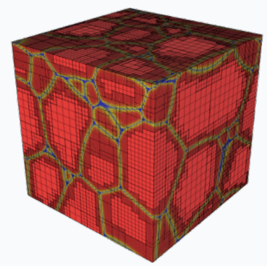	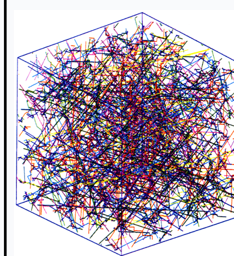	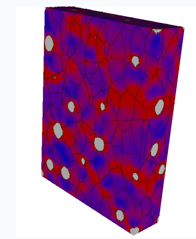	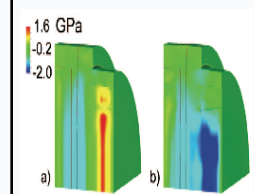
Code: Qbox/ LATTE Motif: Particles and wavefunctions, plane wave DFT, ScaLAPACK, BLACS, and custom parallel 3D FFTs Prog. Model: MPI + CUBLAS/CUDA	Code:SPaSM/ ddcMD/CoMD Motif: Particles, explicit time integration, neighbor and linked lists, dynamic load balancing, parity error recovery, and *in situ* visualization Prog. Model: MPI + Threads	Code: SEAKMC Motif: Particles and defects, explicit time integration, neighbor and linked lists, and *in situ* visualization Prog. Model: MPI + Threads	Code: AMPE/GL Motif: Regular and adaptive grids, implicit time integration, real-space and spectral methods, complex order parameter Prog. Model: MPI	Code: ParaDiS Motif: "segments" Regular mesh, implicit time integration, fast multipole method Prog. Model: MPI	Code: VP-FFT Motif: Regular grids, tensor arithmetic, meshless image processing, implicit time integration, 3D FFTs. Prog. Model: MPI + Threads	Code: ALE3D/ LULESH Motif: Regular and irregular grids, explicit and implicit time integration. Prog. Model: MPI + Threads

Use Case: Shaped-charge jets, breakup and 3D effects (e.g. spinning) require crystal plasticity and anisotropy

What is required:

Resolution: 10^{12} zones (10 cm cube)

Simulation time: 100 μsec (10^5 steps)

Strain rate: 10^6 /sec

Strain: 1-3

Using Small Strain Crystal Plasticity Model: ~10^4 sec (~3 h) wall clock on 10^9 cores

Large Strain Crystal Plasticity Model: 10x

Twinning / Scale Bridging Model: 100x

ALE3D simulation of shaped-charge jet (Rose McCallen, LLNL)

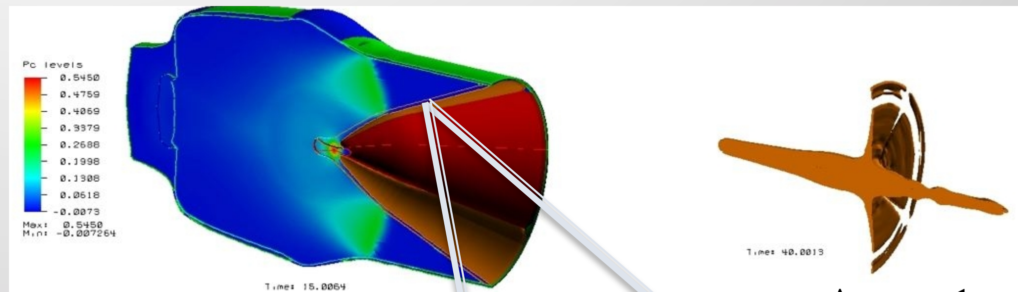

$\Delta\varepsilon \geq 1$

What we can do today:

Crystal plasticity simulation of high rate deformation (Nathan Barton, LLNL)

Model: Small Strain Crystal Plasticity

Number Zones: 10^7 (100 micron cube)

Simulation time: 10 μsec (10^4 steps)

Strain rate: 10^6 /sec

Strain: 0.15

Wall Clock: 1 day on 1/10 Cielo

$\Delta\varepsilon = 0.15$

Objective: Full utilization of exascale concurrency and locality

- Task-based embedded Scale-Bridging escapes the traditional synchronous SPMD paradigm and exploits the heterogeneity expected in exascale hardware.

- To achieve this, we are developing a UQ-driven *adaptive physics refinement* approach.
- Coarse-scale simulations dynamically spawn tightly coupled and self-consistent fine-scale simulations as needed.
- This *task-based* approach naturally maps to exascale heterogeneity, concurrency, and resiliency issues.

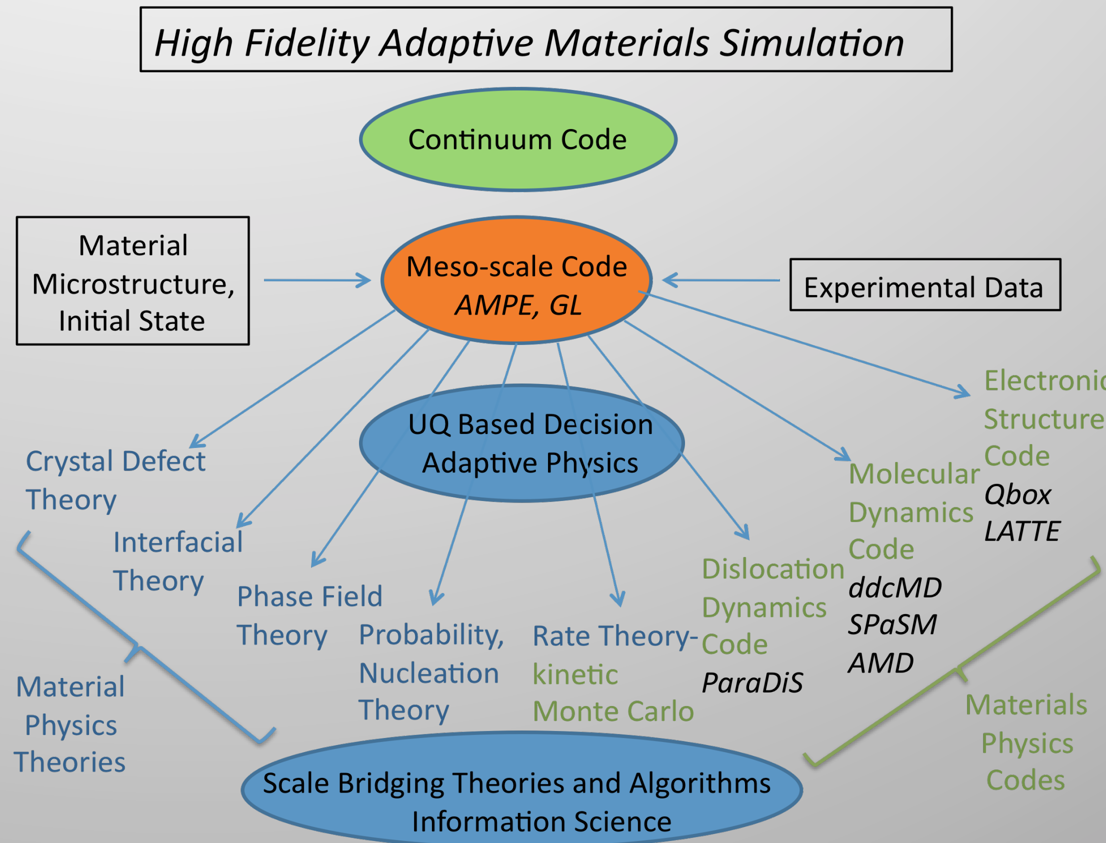

Direct multi-scale embedding requires full utilization of exascale concurrency and locality

- *Brute force multi-scale coupling:* Full fine scale model (FSM, e.g. a crystal plasticity model) run for every zone & time step of coarse scale mode (CSM, e.g. an ALE code)
- *Adaptive Sampling:*
  - Save FSM results in database
  - Before running another FSM, check database for FSM results similar enough to those needed that interpolation or extrapolation suffices
  - Only run full FSM when results in database not close enough

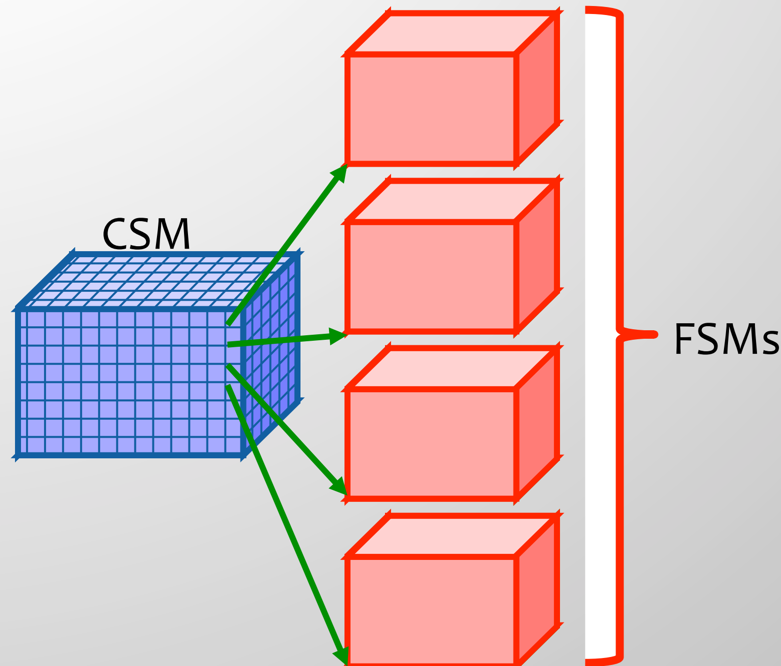

- Heterogeneous, hierarchical <u>MPMD</u> algorithms map naturally to anticipated heterogeneous, hierarchical architectures
- Escape the traditional bulk synchronous SPMD paradigm, improve <u>scalability</u> and reduce <u>scheduling</u>
- Task-based MPMD approach leverages <u>concurrency</u> and <u>heterogeneity</u> at exascale while enabling novel <u>data models</u>, <u>power management</u>, and <u>fault tolerance</u> strategies

Ref: Barton *et.al*, 'A call to arms for task parallelism in multi-scale materials modeling,' *Int. J. Numer. Meth. Engng* 2011; 86:744–764

Adaptive Sampling builds response on the fly

- Coarse scale model queries database for fine-scale material response
- If possible, approximate response from past evaluations
- Otherwise perform fine scale evaluation
- Fine-scale evaluations grow database

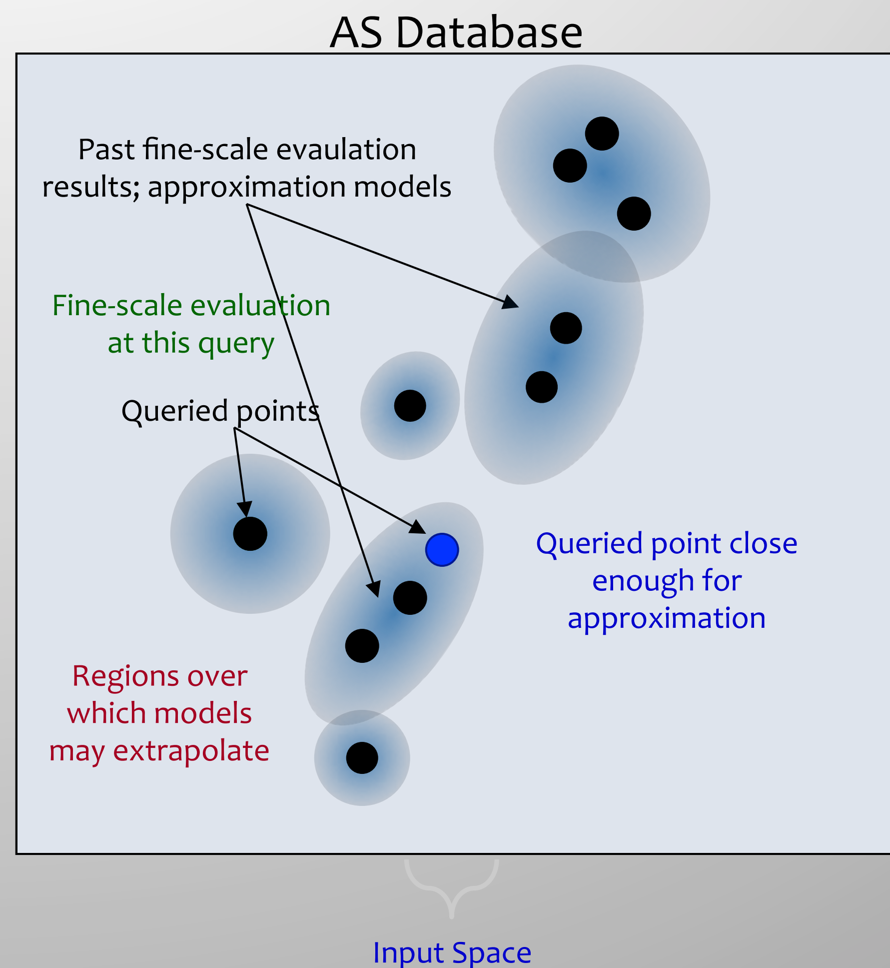

Tradeoff: re-use vs. re-computation of expensive fine-scale model results

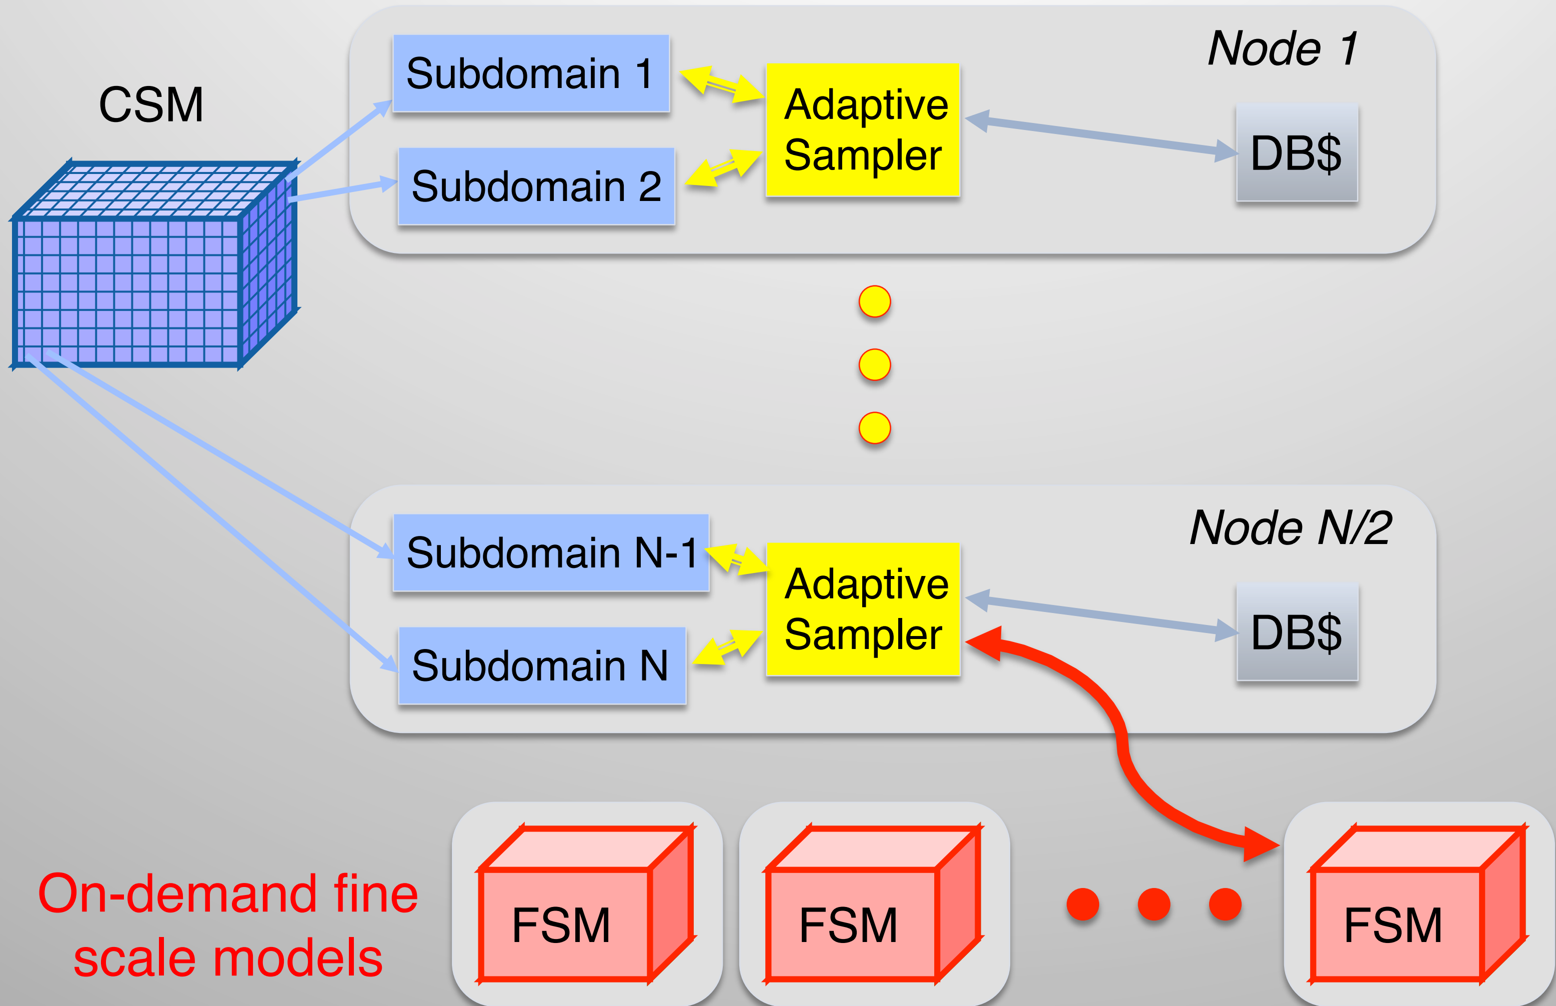

Tradeoff: re-use vs. re-computation of expensive fine-scale model results

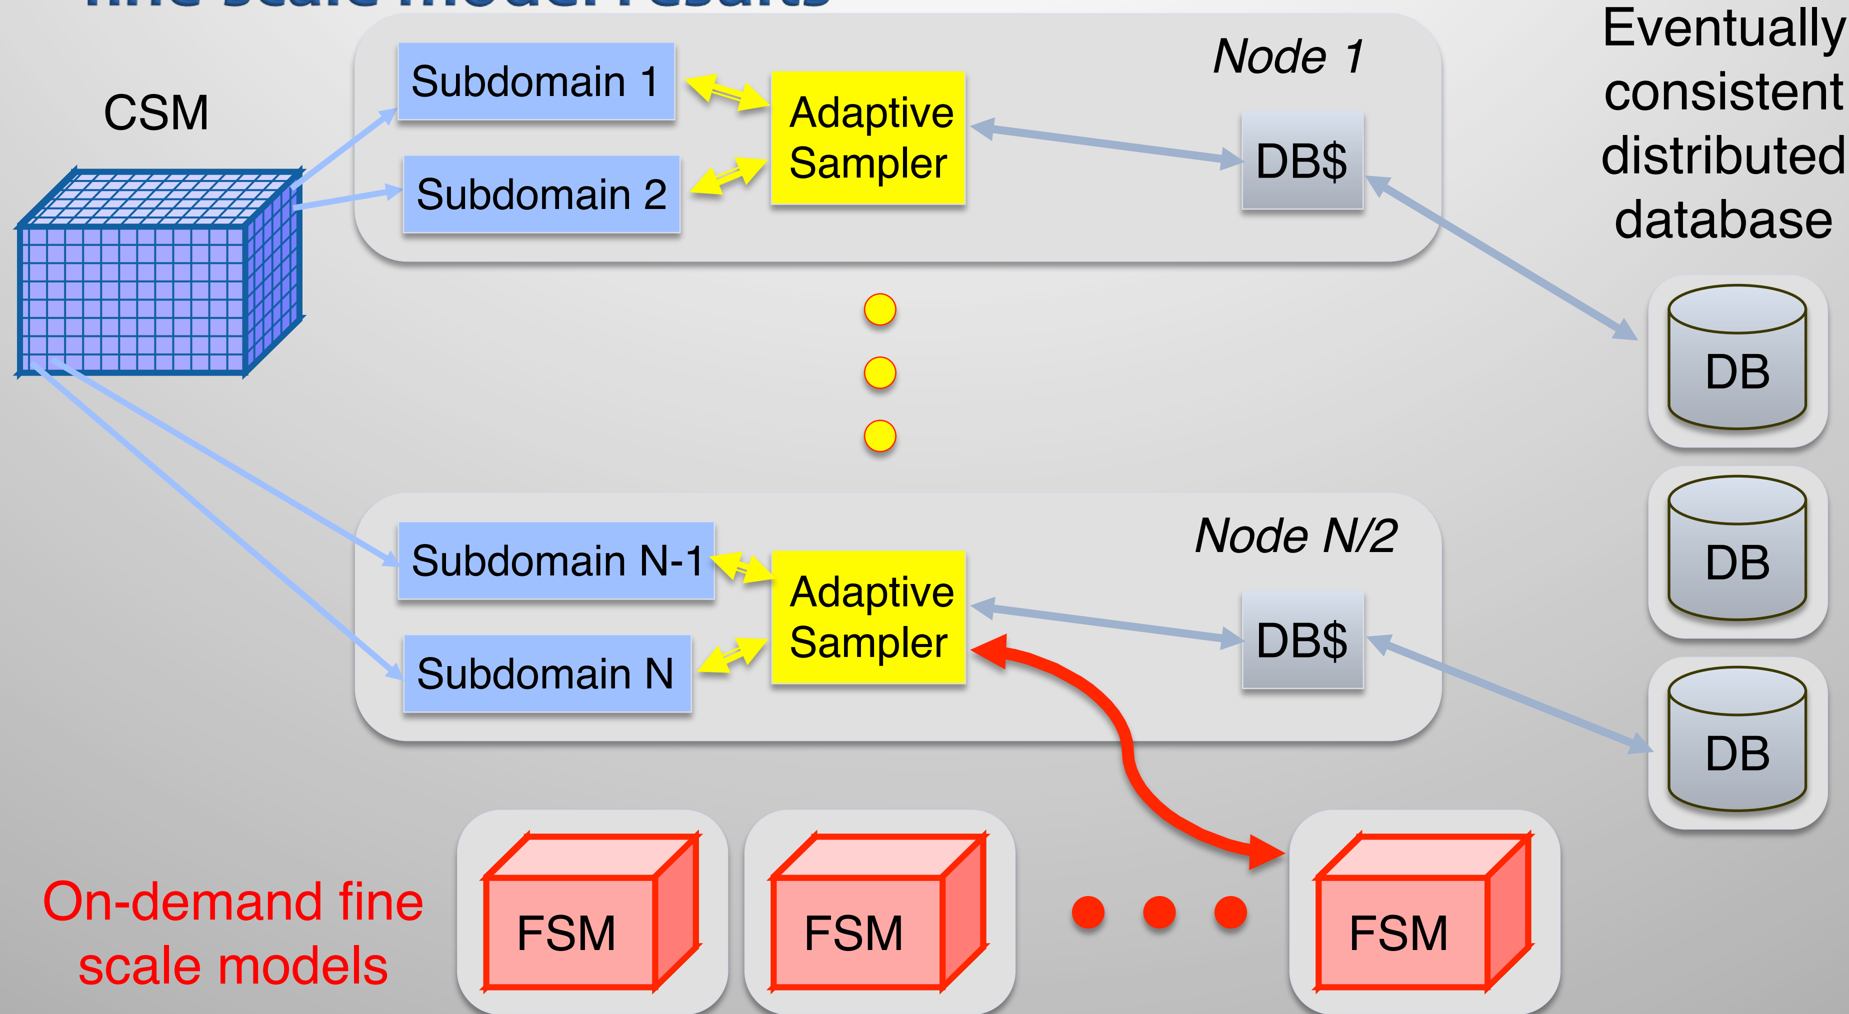

Several practical challenges must be addressed

- Interoperability between (often) legacy coarse and fine-scale models, and backing database
  - Scale-bridging implementations to date often "hacked" to work within MPI
  - "Monolithic" programming models (e.g. Charm++, Chapel) may require impractical rewrites of entire code base
  - Web/cloud technologies are often service-based, and emphasize programmer productivity, code agility, and code maintainability over performance
  - Component-based frameworks (e.g. Uintah, Pathos) have demonstrated success for multiphysics HPC applications
- Distributed database
  - Most implementations depend on conventional TCP/IP, which is often not supported on HPC software stacks (Infiniband)
  - Stanford's RAMCloud an interesting option

Seeds for Contemplation: Promote State (Data) be be the first class citizen.

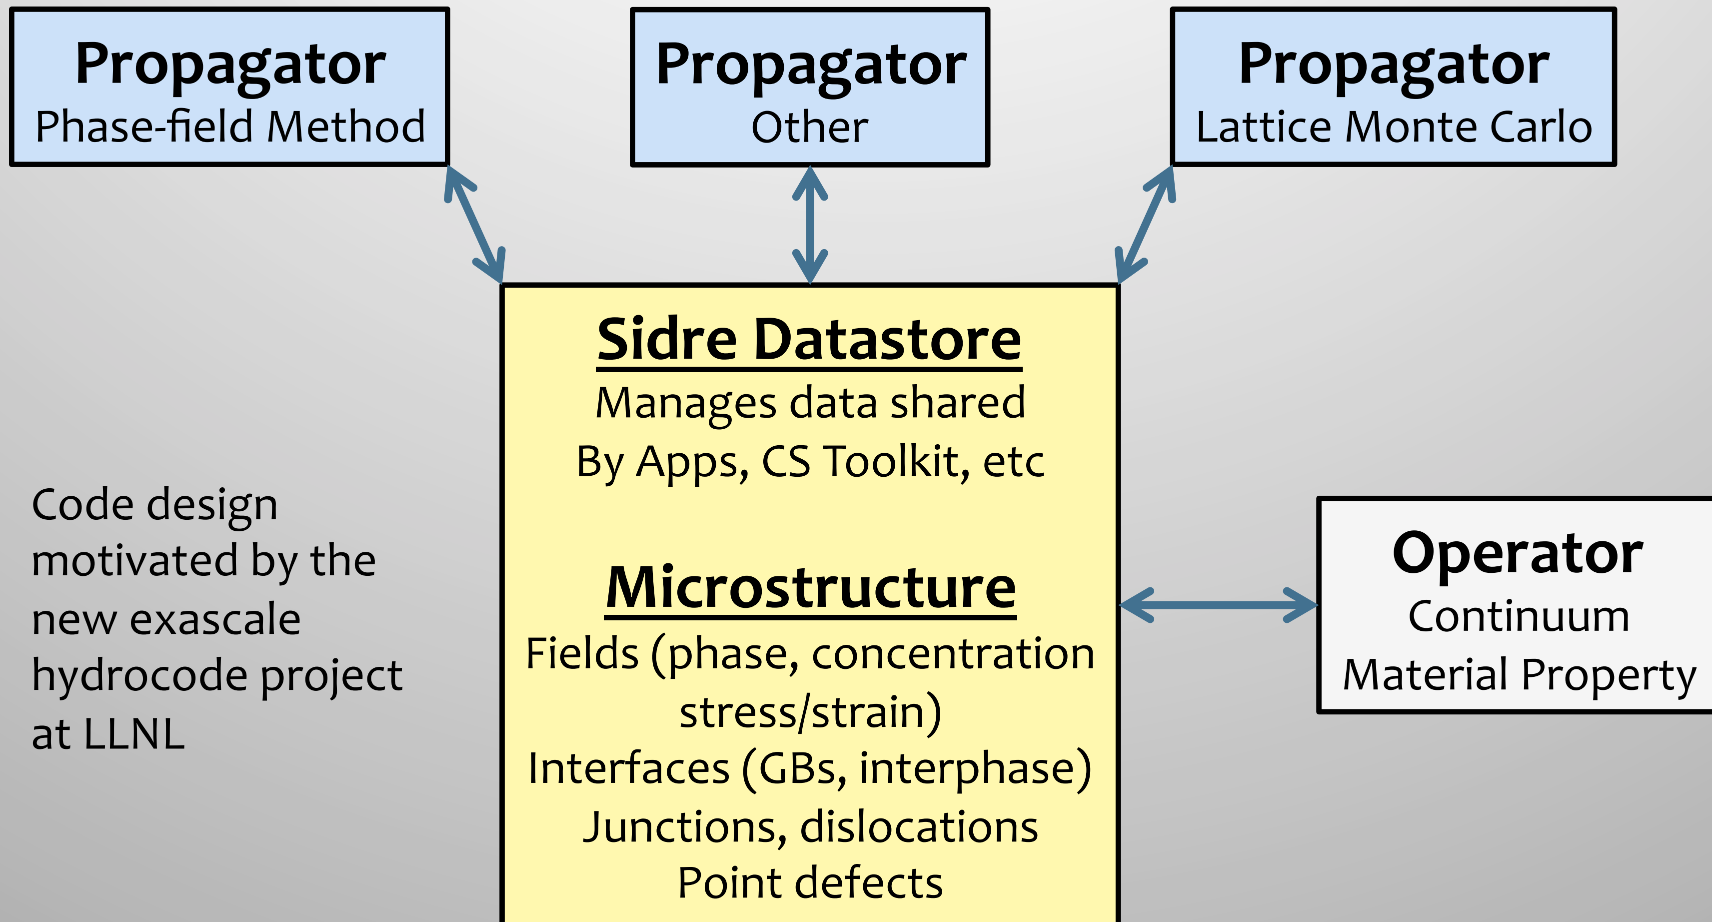

Code design motivated by the new exascale hydrocode project at LLNL

More Seeds for Contemplation

- **Workflow Manager:** Imagine imbedded chips (Neuromorphic? Vickie from I,Robot) that monitor and predict work flow during the simulation and moves tasks (both work and data) to optimal locations.
- **Machine Learning Constitutive Models:** the constitutive model is a response relation (e.g. stress-strain) that depends on the constitution of the material. It is the closure relation in the continuum conservation equations. Imagine using an ML model that learns from experimental data and direct numerical simulation.
- **Programming:** What does a scientific application look like when expressed in a programming framework designed for AI? e.g. TensorFlow code for O(N) electronic structure.