Los Alamos NATIONAL LABORATORY EST. 1943 | Lawrence Livermore National Laboratory | OAK RIDGE National Laboratory | Sandia National Laboratories | STANFORD UNIVERSITY | California Institute of Technology

ExMatEx: Exascale Co-design Center for Materials in Extreme Environments

Timothy C. Germann

Los Alamos National Laboratory

SOS 18 Workshop

St. Moritz, Switzerland *March 17-20, 2014*

tcg@lanl.gov

U.S. DEPARTMENT OF ENERGY | Office of Science

Theme: Supercomputers as scientific instruments

The theme of the 2014 workshop is "Supercomputers as scientific instruments" that will visit the following topics:

- What and who should drive the development of supercomputers? The centres or domain science communities?
- Should these supercomputers be general purpose or appliances?
- Should they exist as standalone infrastructures or be integrated with experimental facilities or domain specific research centres?
- Furthermore, there is great interest in discussing the many dimensions of software and its sustainable development. Is the traditional split into application software vs. system software sustainable, or should we consider developing a new middleware layer for scientific computing? If yes, what should it look like and why would its development be more sustainable?

Abstract

Within the Exascale Co-design Center for Materials in Extreme Environments (ExMatEx), we have initiated an early and deep collaboration between domain (computational materials) scientists, applied mathematicians, computer scientists, and hardware architects, in order to establish the relationships between algorithms, software stacks, and architectures needed to enable exascale-ready materials science application codes within the next decade. We anticipate that we will be able to exploit hierarchical, heterogeneous architectures to achieve more realistic large-scale simulations with adaptive physics refinement, and are using tractable application scale-bridging proxy application testbeds to assess new approaches to resilience, OS/runtime and execution models, and power management. The current scale-bridging strategies accumulate (or recompute) a distributed response database from fine-scale calculations (tasks), in a top-down rather than bottom-up multiscale approach. I will demonstrate this approach and our initial assessments, using simplified proxies to encapsulate the expected scale-bridging workload and workflow.

Modeling and simulation is playing an increasing role in materials design and certification

- High-strength, light-weight structural materials are required for products from cars and airplanes to gas, wind, and jet turbine blades

- Materials Genome Initiative

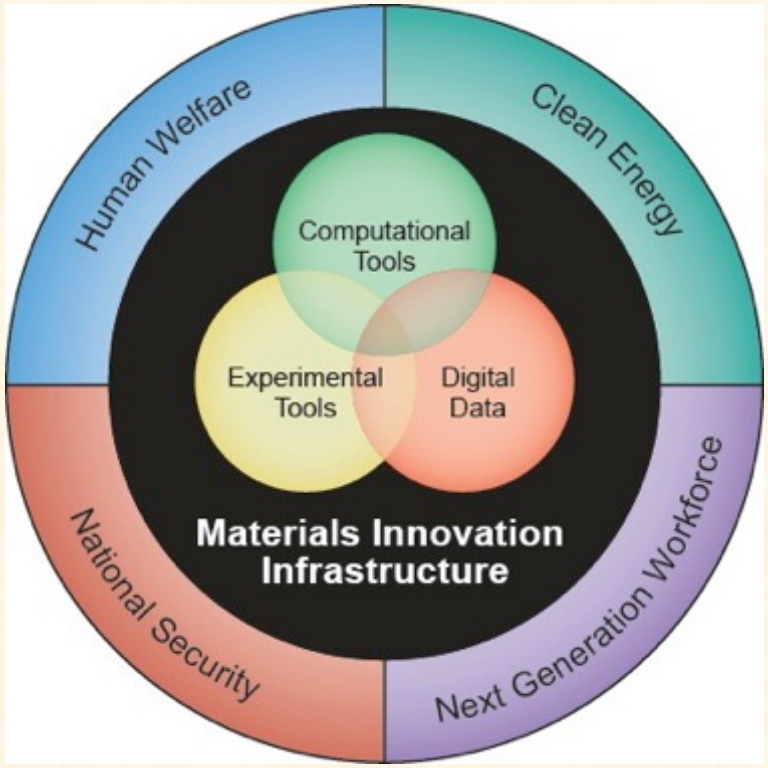

http://www.whitehouse.gov/MGI

Materials dynamics issues are important to NNSA's weapons mission.

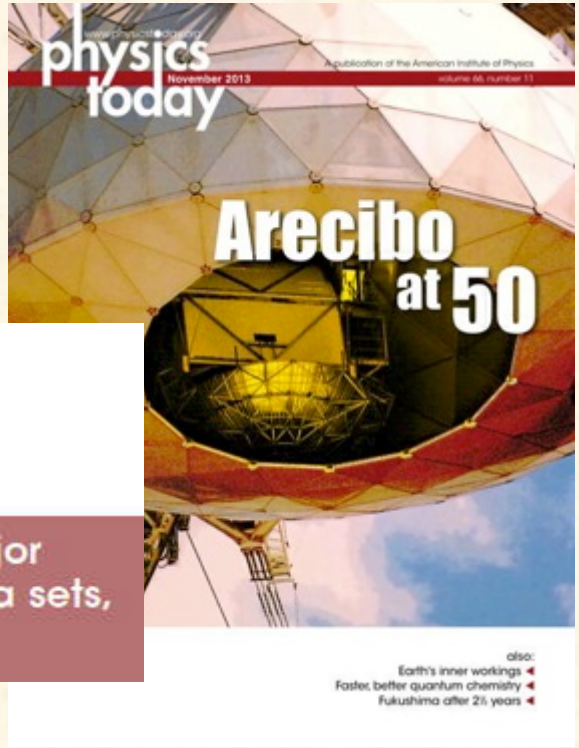

NNSA touts savings from supercomputing

Nuclear weaponeers look toward exascale computing, but major breakthroughs in power consumption, handling of massive data sets, and other areas are needed first.

"…if you look at where cracks develop in metals they always develop at the grain boundary. If you look at where corrosion occurs, it's at a grain boundary. If you look at the effect on materials of aging, you gather a lot of chemical contaminants at the grain boundary." – Donald Cook, NNSA Deputy Administrator for Defense Programs (*Physics Today*, Nov 2013)

Nov 2013

"First, the stress also serves as a **direct test of supercomputer simulations that model how metals behave. The better the data that goes in, the more reliable are the results that come out. That's important in trying to model the exact behavior of metals under stress, say the crash of a car or the impact of a bullet into armor.** And it's especially important for the Office of Science, since several of its labs are home to world-class supercomputers, which researchers are using for everything from simulating the 'subatomic soup' of the early universe to modeling air turbulence and thereby improving airplane performance.

Those better metal models could, in turn, lead to the design of even stronger and more durable materials. And those materials might come in handy for technologies that operate in extreme environments, such as shielding for satellites and space probes. They'll likely be useful in more everyday applications too."

Molecular dynamics simulations have revealed underlying unit processes during shock loading.

Microscopic View of Structural Phase Transitions Induced by Shock Waves

Kai Kadau,[1,2*] Timothy C. Germann,[3] Peter S. Lomdahl,[1] Brad Lee Holian[1]

Multimillion-atom molecular-dynamics simulations are used to investigate the shock-induced phase transformation of solid iron. Above a critical shock strength, many small close-packed grains nucleate in the shock-compressed body-centered cubic crystal growing on a picosecond time scale to form larger, energetically favored grains. A split two-wave shock structure is observed immediately above this threshold, with an elastic precursor ahead of the lagging transformation wave. For even higher shock strengths, a single, overdriven wave is obtained. The dynamics and orientation of the developing close-packed grains depend on the shock strength and especially on the crystallographic shock direction. Orientational relations between the unshocked and shocked regions are similar to those found for the temperature-driven martensitic transformation in iron and its alloys.

Science **296**, 1681 (2002)

Phys Rev B **72**, 064120 (2005)

Phys Rev Lett **98**, 135701 (2007)

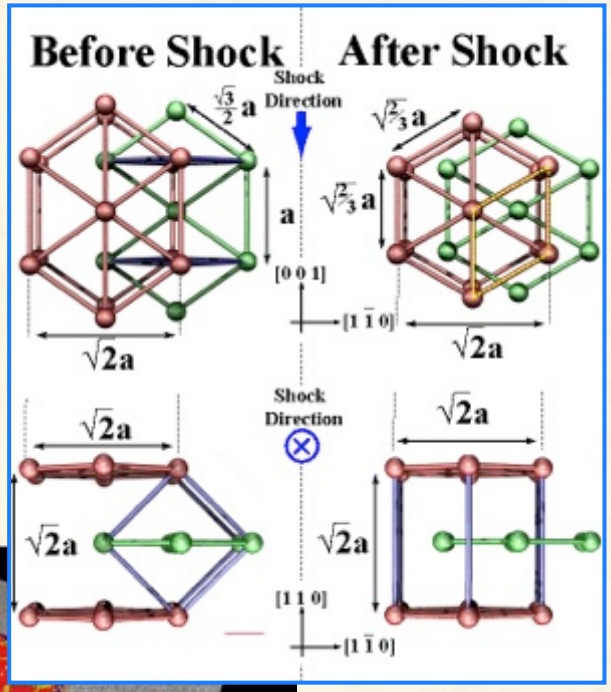

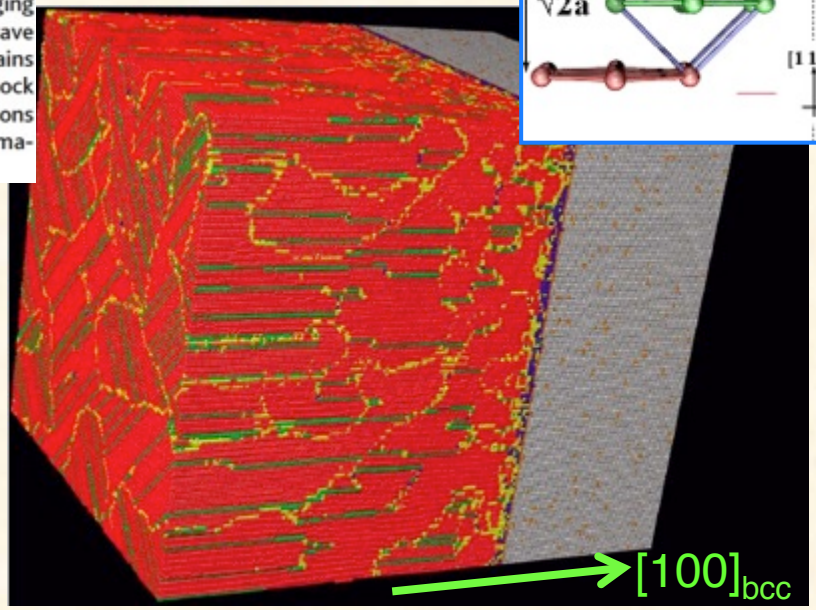

2005: First direct confirmation of the Fe α–ε transition via *in situ* X-ray diffraction

PRL 95, 075502 (2005) PHYSICAL REVIEW LETTERS week ending 12 AUGUST 2005

Direct Observation of the α-ε Transition in Shock-Compressed Iron via Nanosecond X-Ray Diffraction

D. H. Kalantar,[1] J. F. Belak,[1] G. W. Collins,[1] J. D. Colvin,[1] H. M. Davies,[2] J. H. Eggert,[1] T. C. Germann,[3] J. Hawreliak,[4] B. L. Holian,[3] K. Kadau,[3] P. S. Lomdahl,[3] H. E. Lorenzana,[1] M. A. Meyers,[5] K. Rosolankova,[4] M. S. Schneider,[5] J. Sheppard,[4] J. S. Stölken,[1] and J. S. Wark[4]

[1]*Lawrence Livermore National Laboratory, Livermore, California 94550, USA*
[2]*AWE Aldermaston, Reading, United Kingdom*
[3]*Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*
[4]*University of Oxford, Oxford, United Kingdom*
[5]*University of California, San Diego, La Jolla, California 92093, USA*

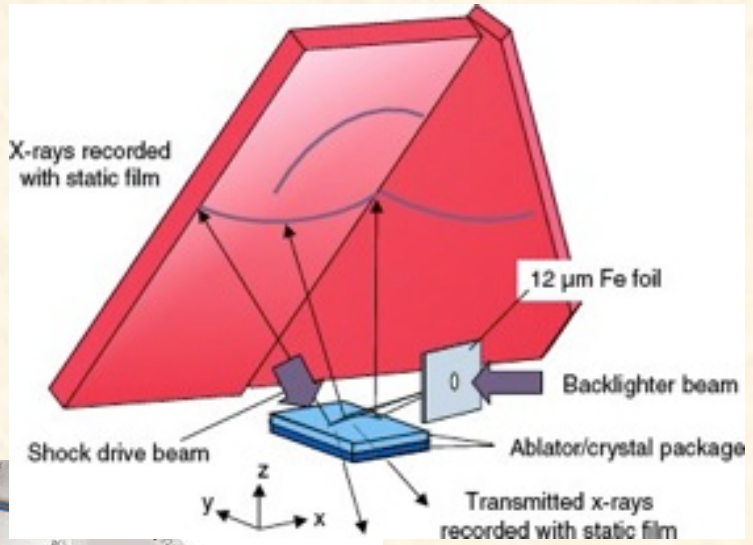

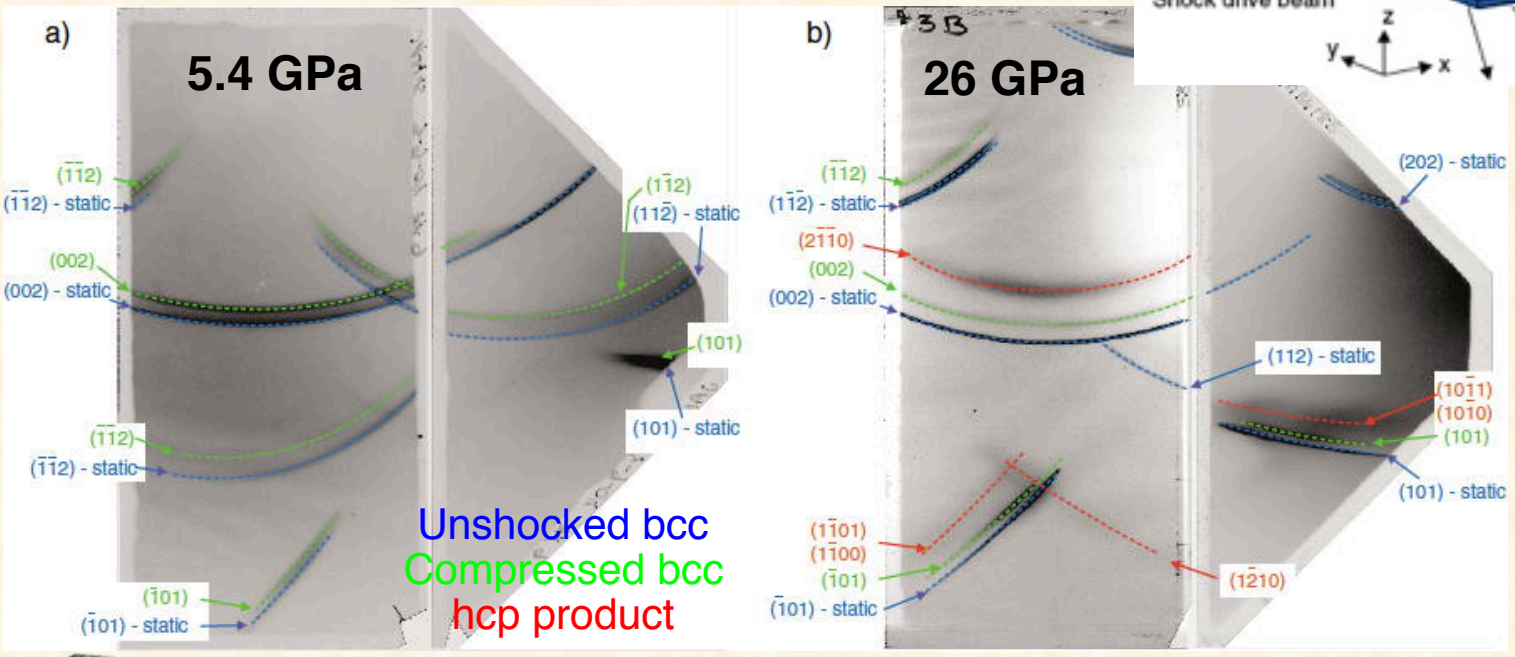

New light sources such as LCLS and APS are providing unprecedented spatio-temporal resolution.

Linac Coherent Light Source (LCLS) setup for shock experiments

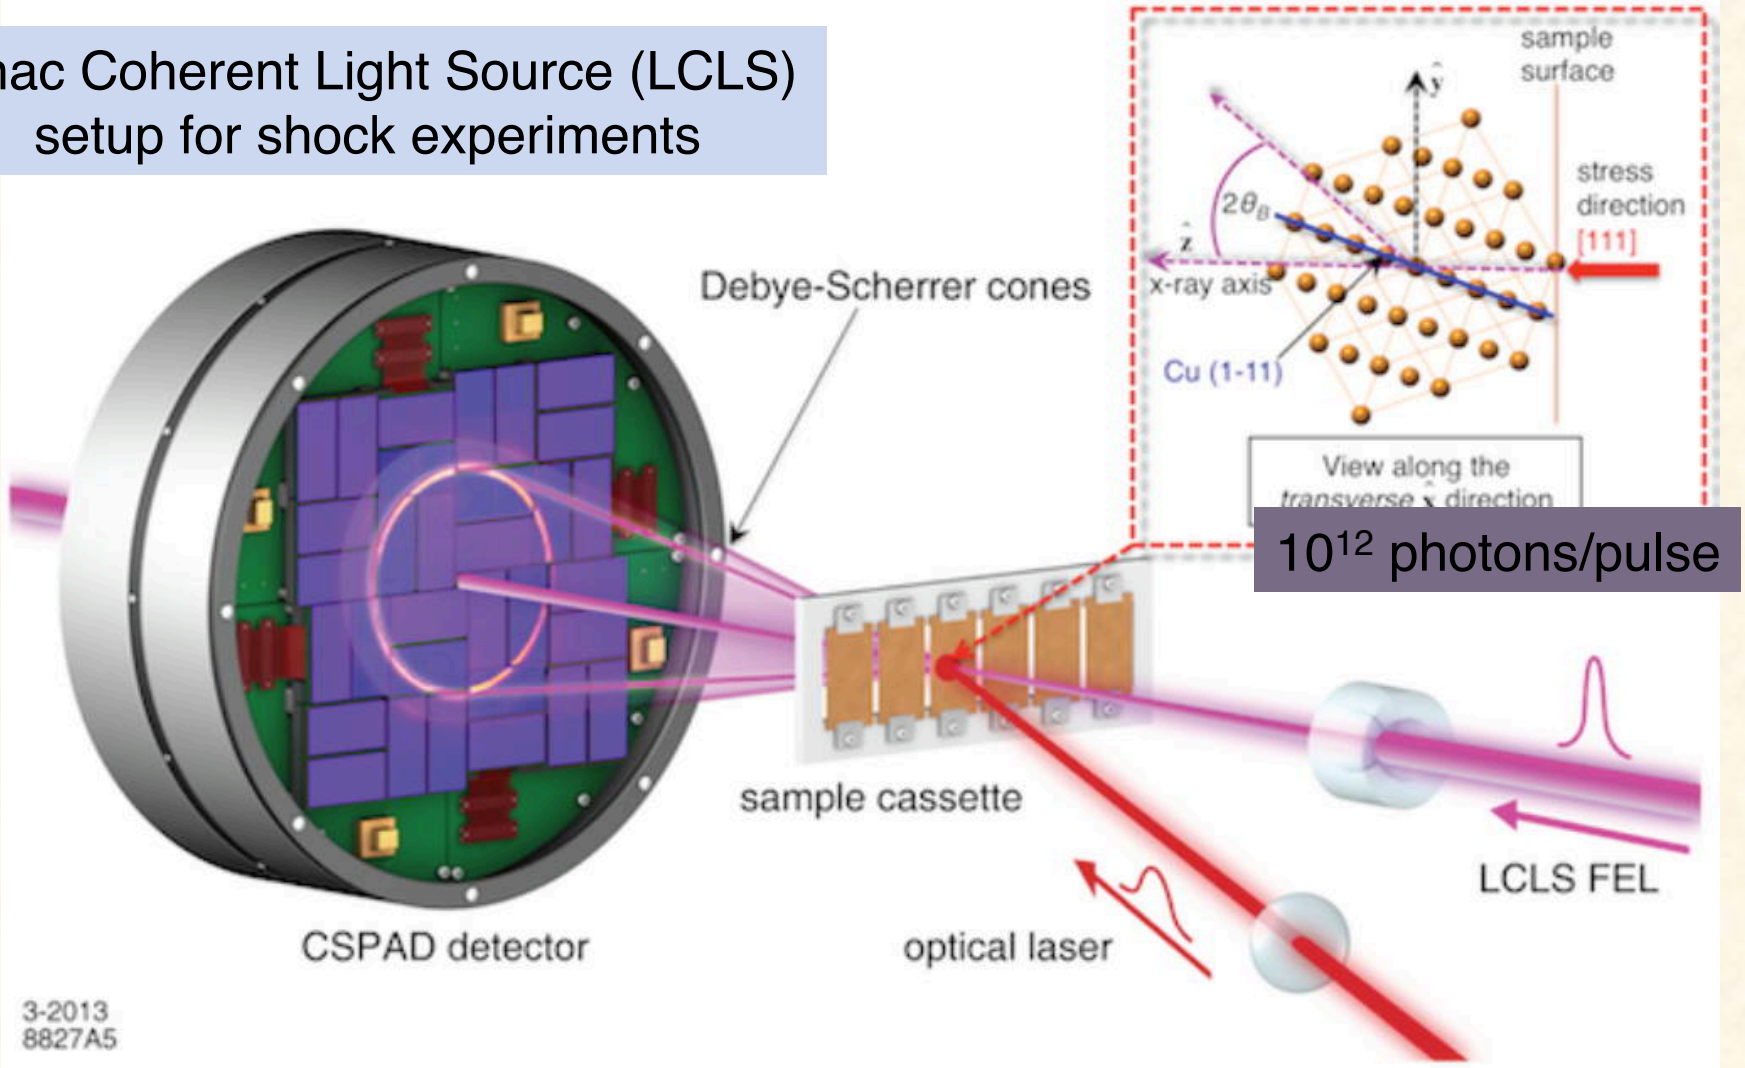

10^{12} photons/pulse

Despina Milathianaki *et al*, Science **342**, 220 (11 October 2013)

Initial shock experiments at LCLS exhibit excellent agreement with MD simulation predictions for Cu(111)

A peak compressive strain ~18% is reached before the onset of plastic flow

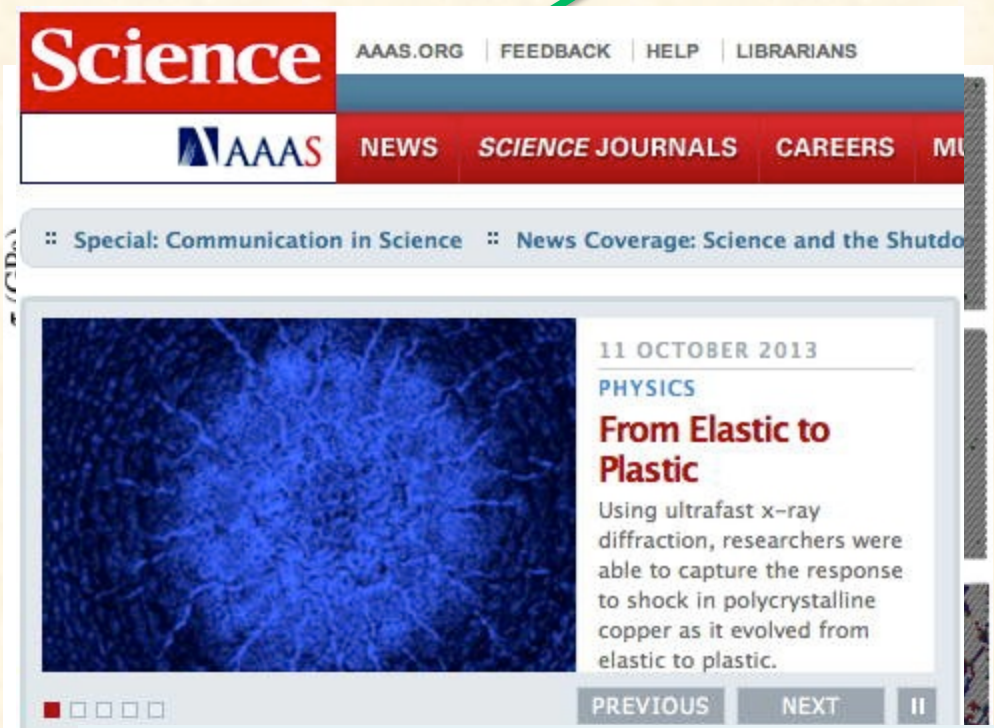

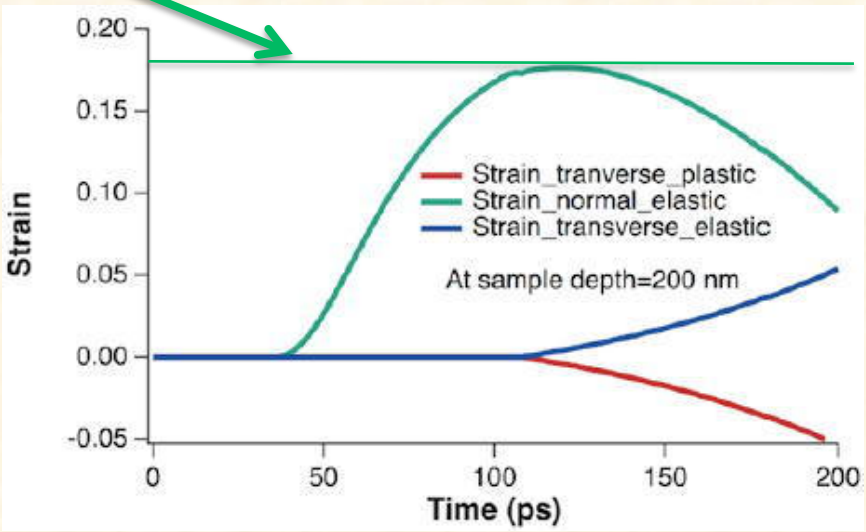

Milathianaki et al, *Science* **342**, 220 (2013)

(d)
Dup

"…the yield stress of the material is in excellent agreement with MD simulations in single crystal Cu at a strain rate of (10^9 s^{-1}) and for uniaxial compression along the [111] direction, thus confirming the considerably higher yield stress values predicted by simulations compared with those extracted from nanosecond shock experiments on samples of >>1-μm thickness."

Integration of experiment and simulation is too often only via the “viewgraph norm”

Interatomic potential

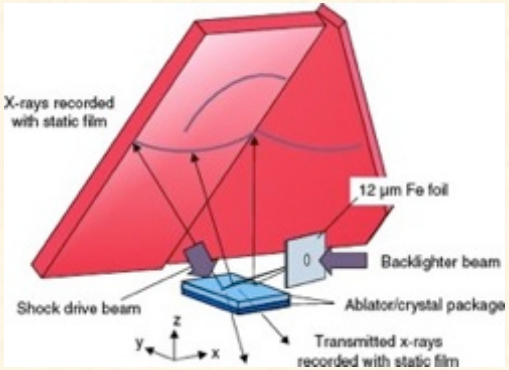

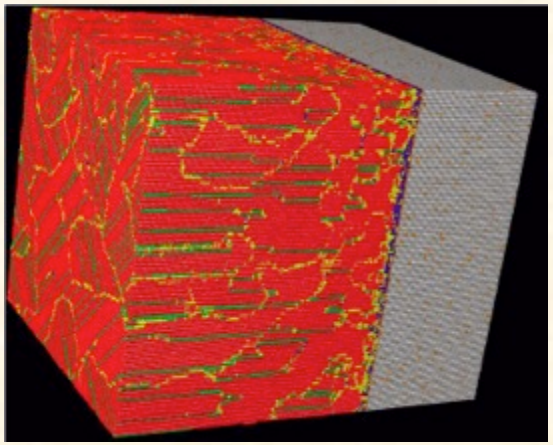

Non-equilibrium molecular dynamics simulation

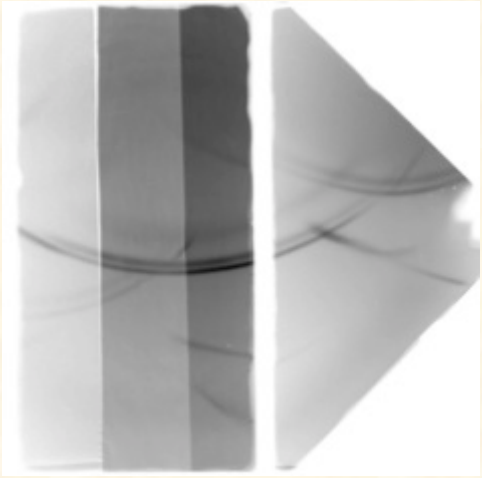

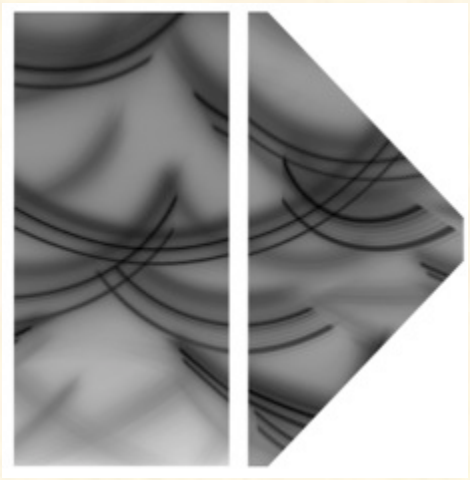

Simulated X-ray diffraction pattern

Data "challenges" presented by new light sources

- Data analysis *post situ*
- Volume, velocity, variety tolerable
- Veracity (noise) may be an issue

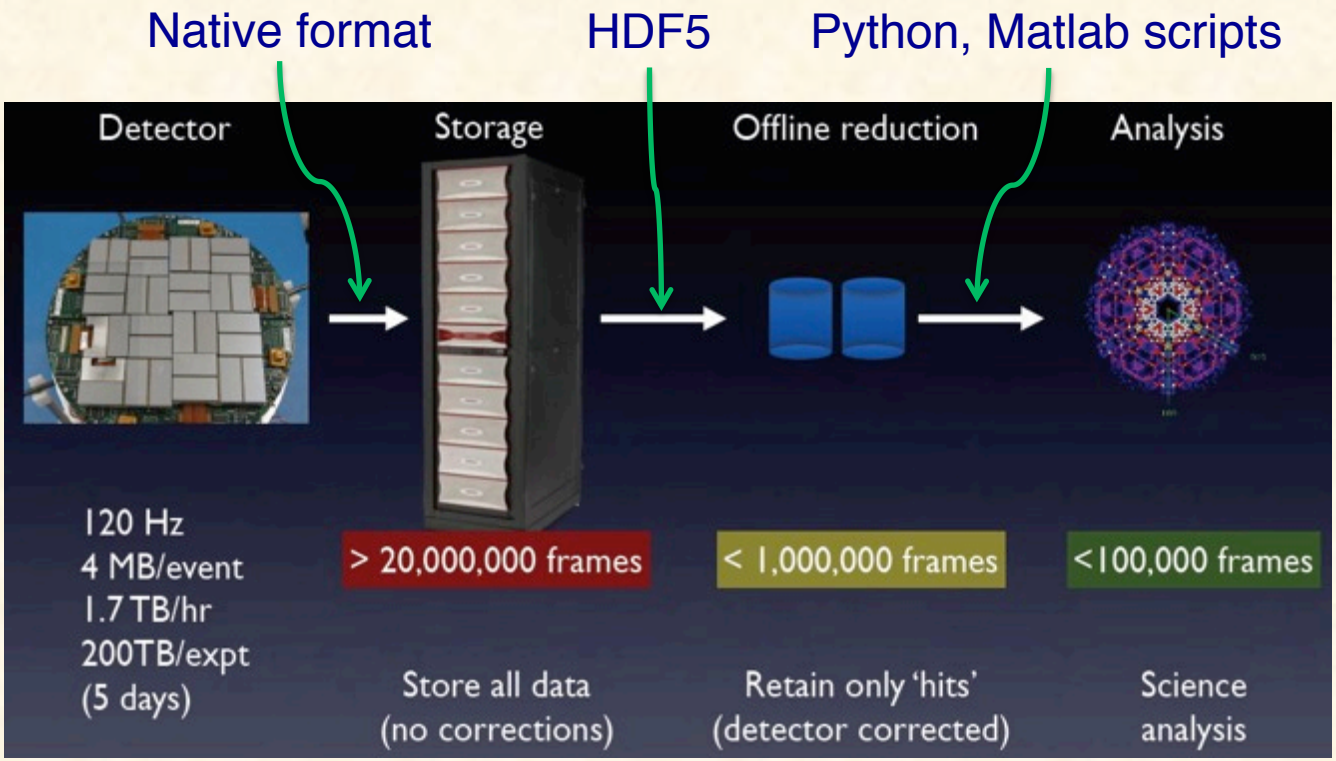

A. Barty, "Femtosecond serial imaging using fast integrating detectors," 8/19/2011

- Experiment & simulation *interact* via forward modeling, e.g. simulated diffraction patterns

Exascale use case: competing dislocation, twinning, and/or phase transitions under shock loading

Direct non-equilibrium molecular dynamics simulation matching time and length scales of LCLS experiments:

- *~1-2 μm thick nanocrystalline samples (Cu, Ti, Fe, Ta), ~400 nm grain size*
- *Laser drive: 10-20 ps rise time, 150 ps duration*
- *50 fs duration X-ray "snapshot" interrogation pulses at 10 ps intervals*

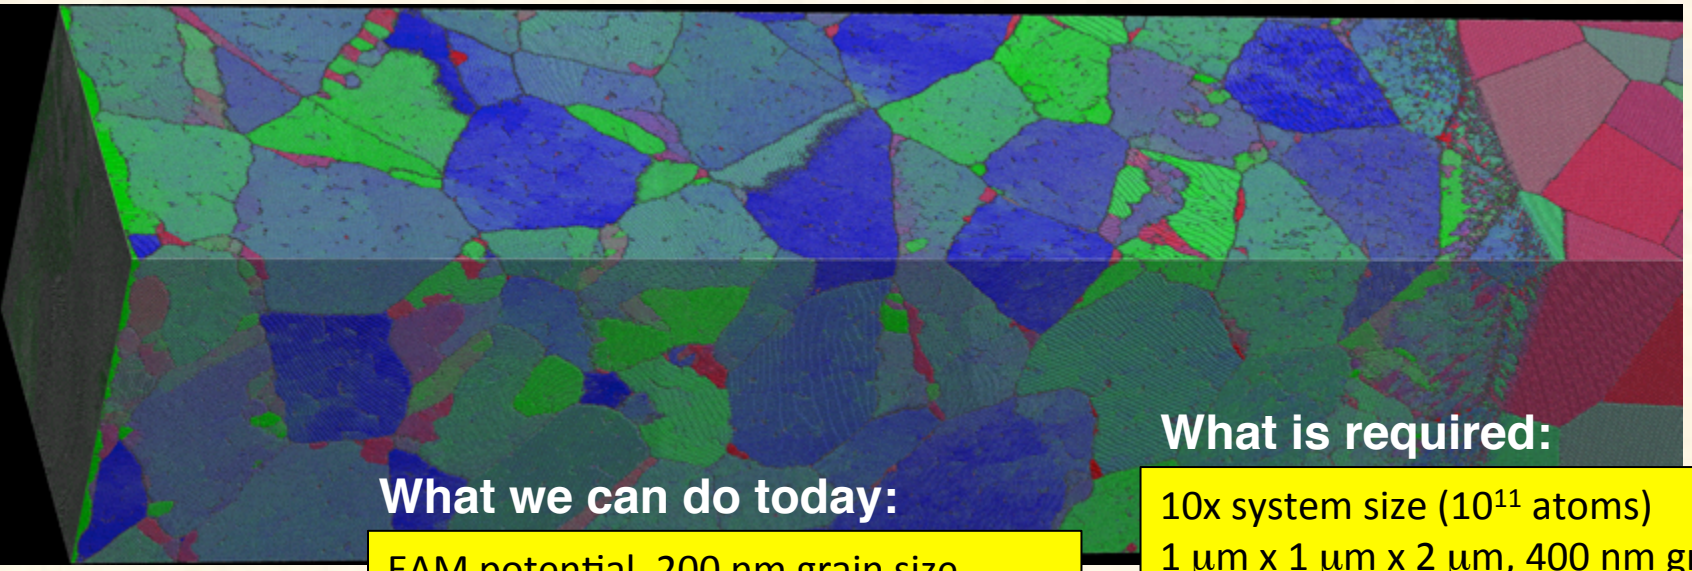

NEMD simulation of shocked nc-Ta on Cielito (R. Ravelo, LANL/ UTEP)

What we can do today:

EAM potential, 200 nm grain size
10^{10} atoms (0.5 μm x 0.5 μm x 1.5 μm)
Simulation time: 4 nsec (10^6 steps)
Wall clock: 2 days on Mira (½ Sequoia)

What is required:

10x system size (10^{11} atoms)
1 μm x 1 μm x 2 μm, 400 nm grain size

More accurate MGPT potential: 100x
3 weeks on exascale system

Traditional approach to subscale models: "sequential multiscale"

- Subscale models (e.g. interatomic potentials, equation of state and strength models) are developed from a combination of theory, experiment, and simulation.
  - *The specific combination depends on the developer, and may involve as much art as science.*

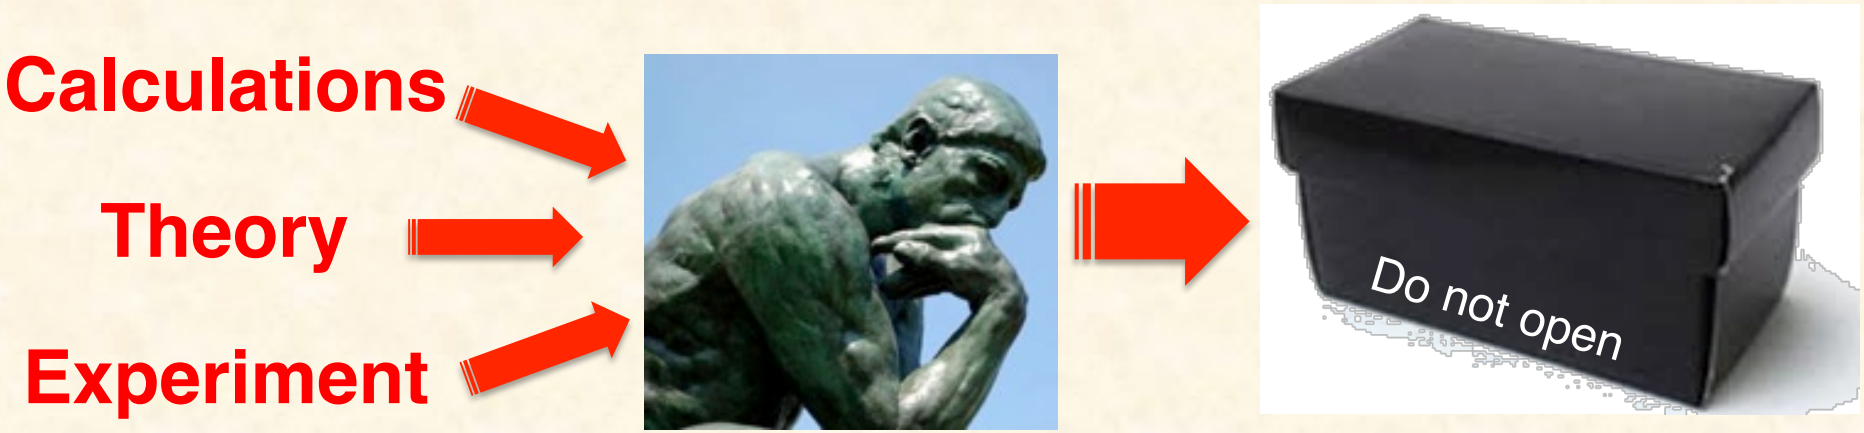

Computational materials science involves a hierarchy of length and time scales

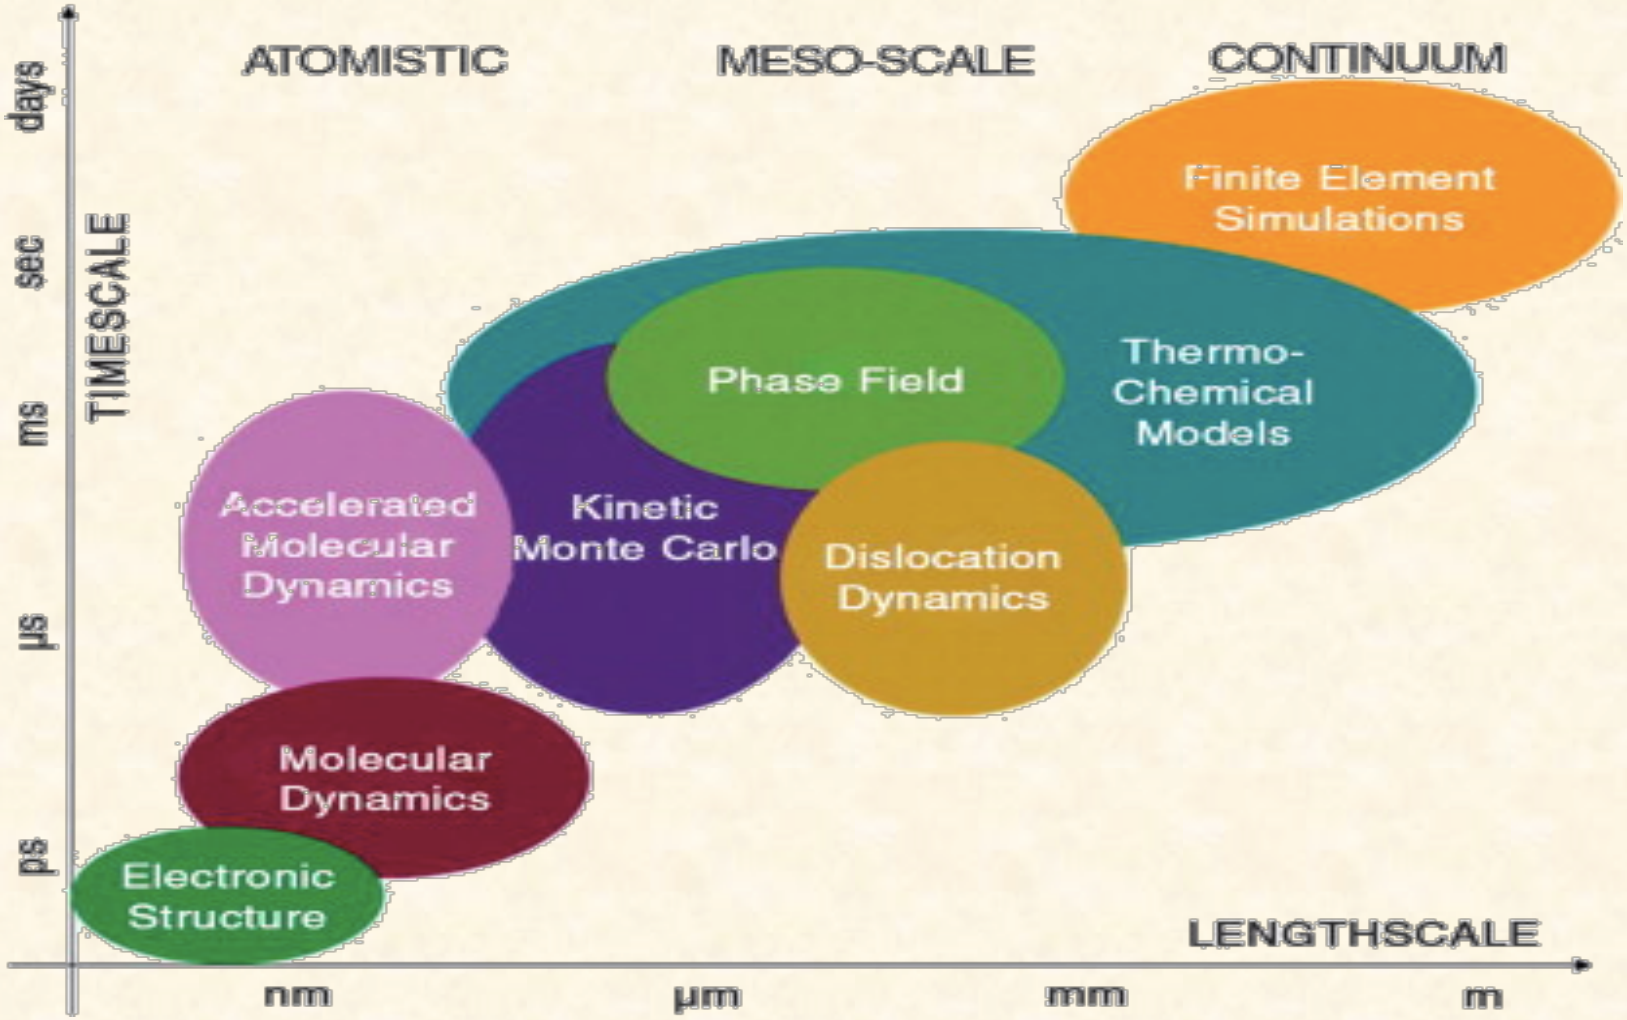

M. Stan, Materials Today **12**(11), 20 (2009)

Example: "sequential" multiscale strength model

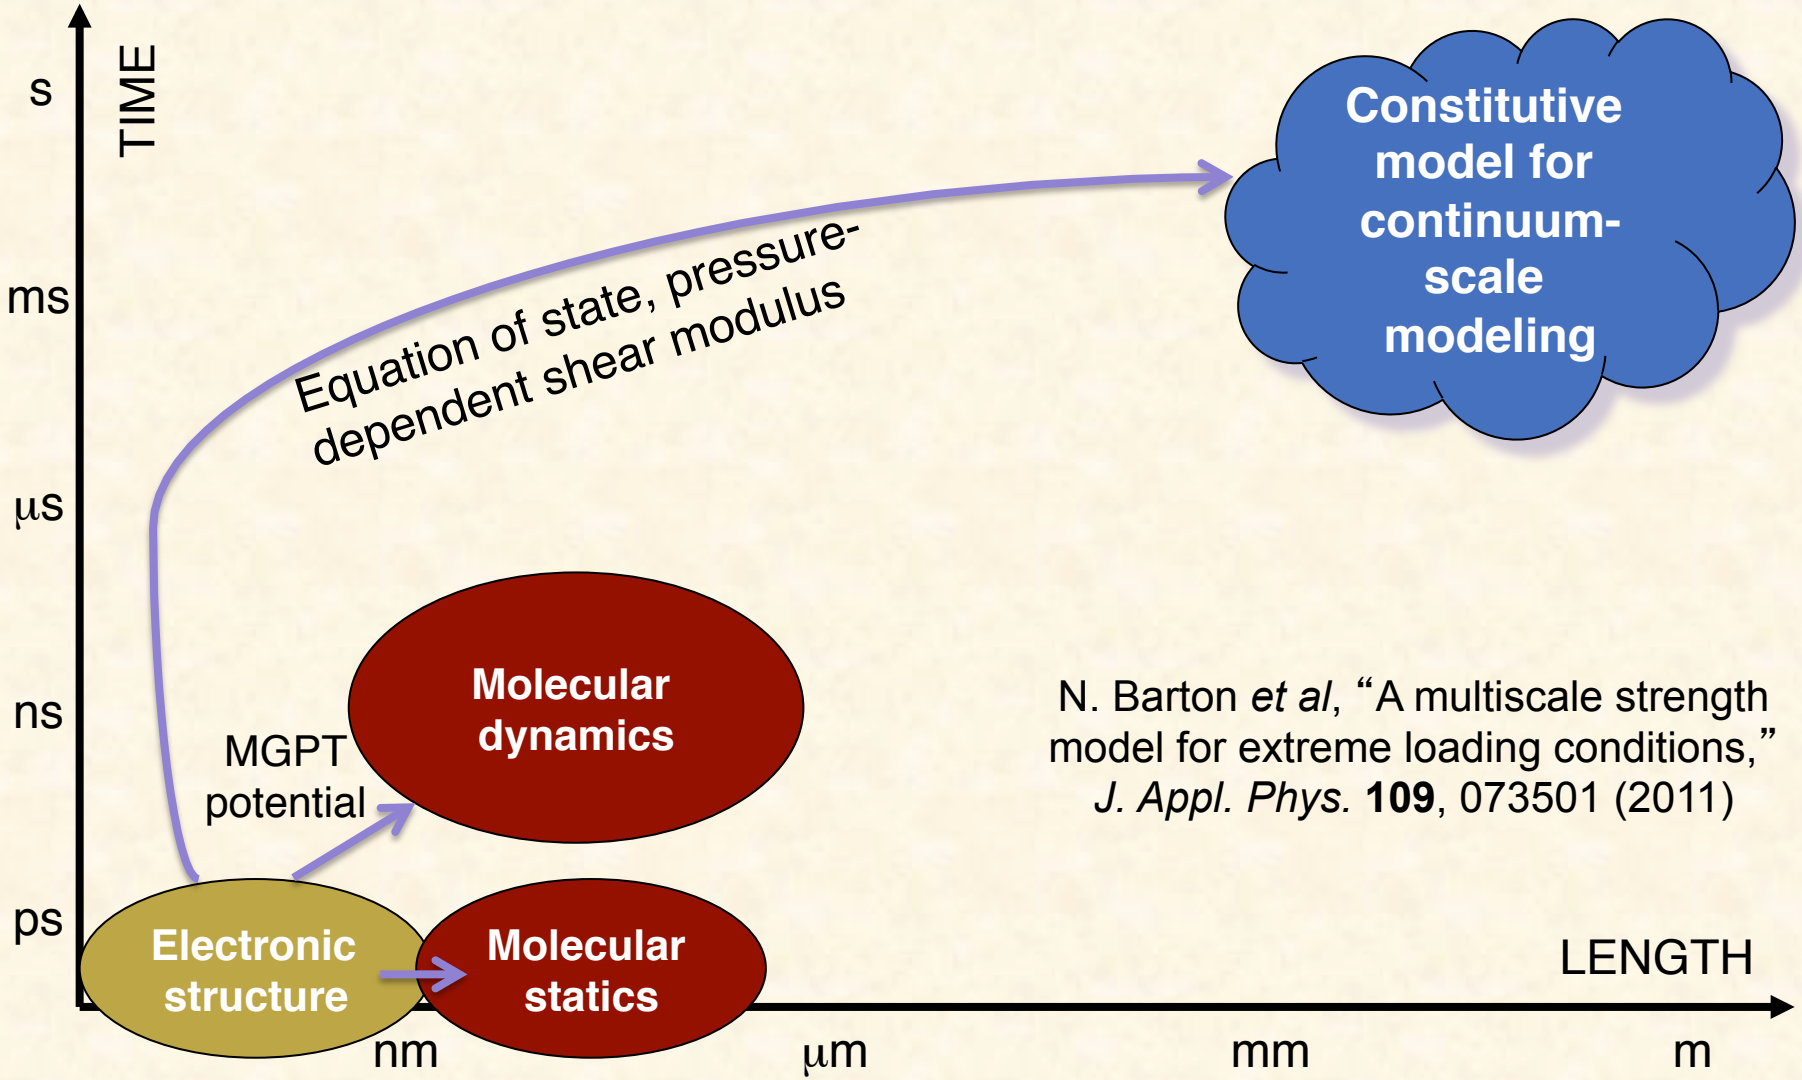

N. Barton *et al*, "A multiscale strength model for extreme loading conditions," *J. Appl. Phys.* **109**, 073501 (2011)

Example: "sequential" multiscale strength model

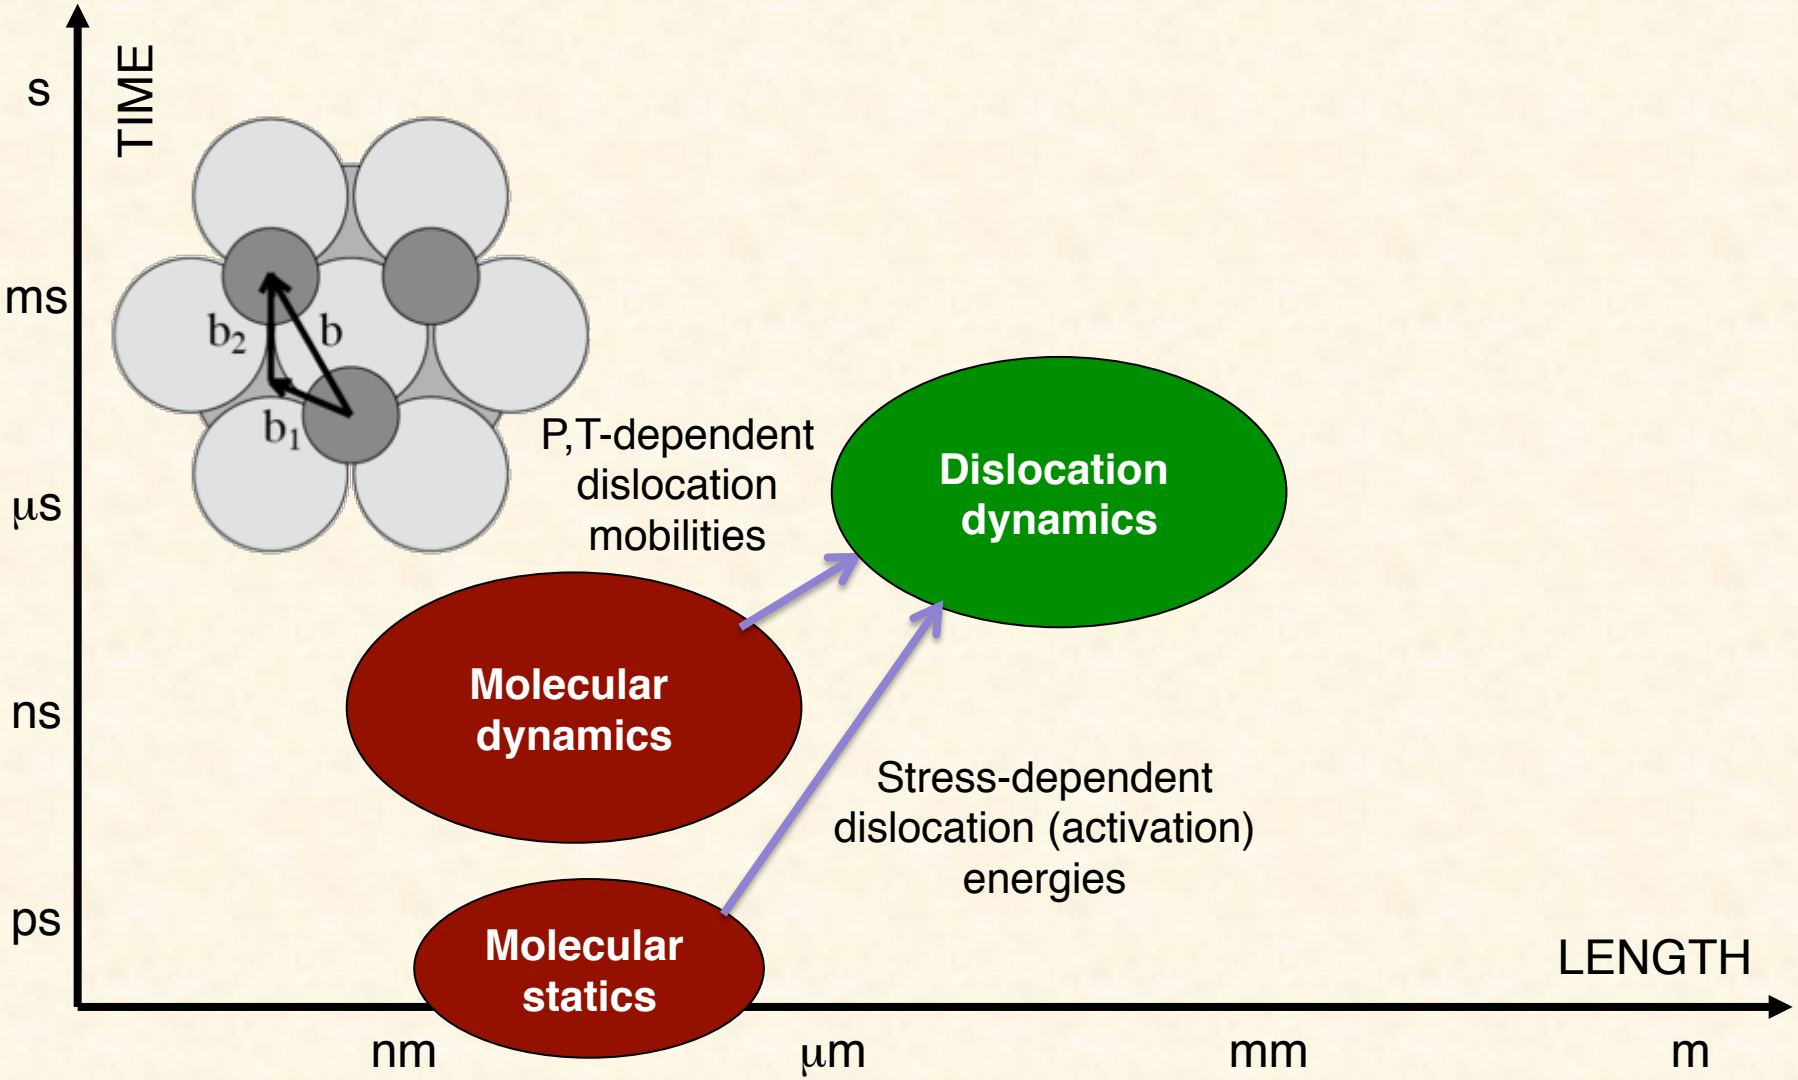

Example: “sequential” multiscale strength model

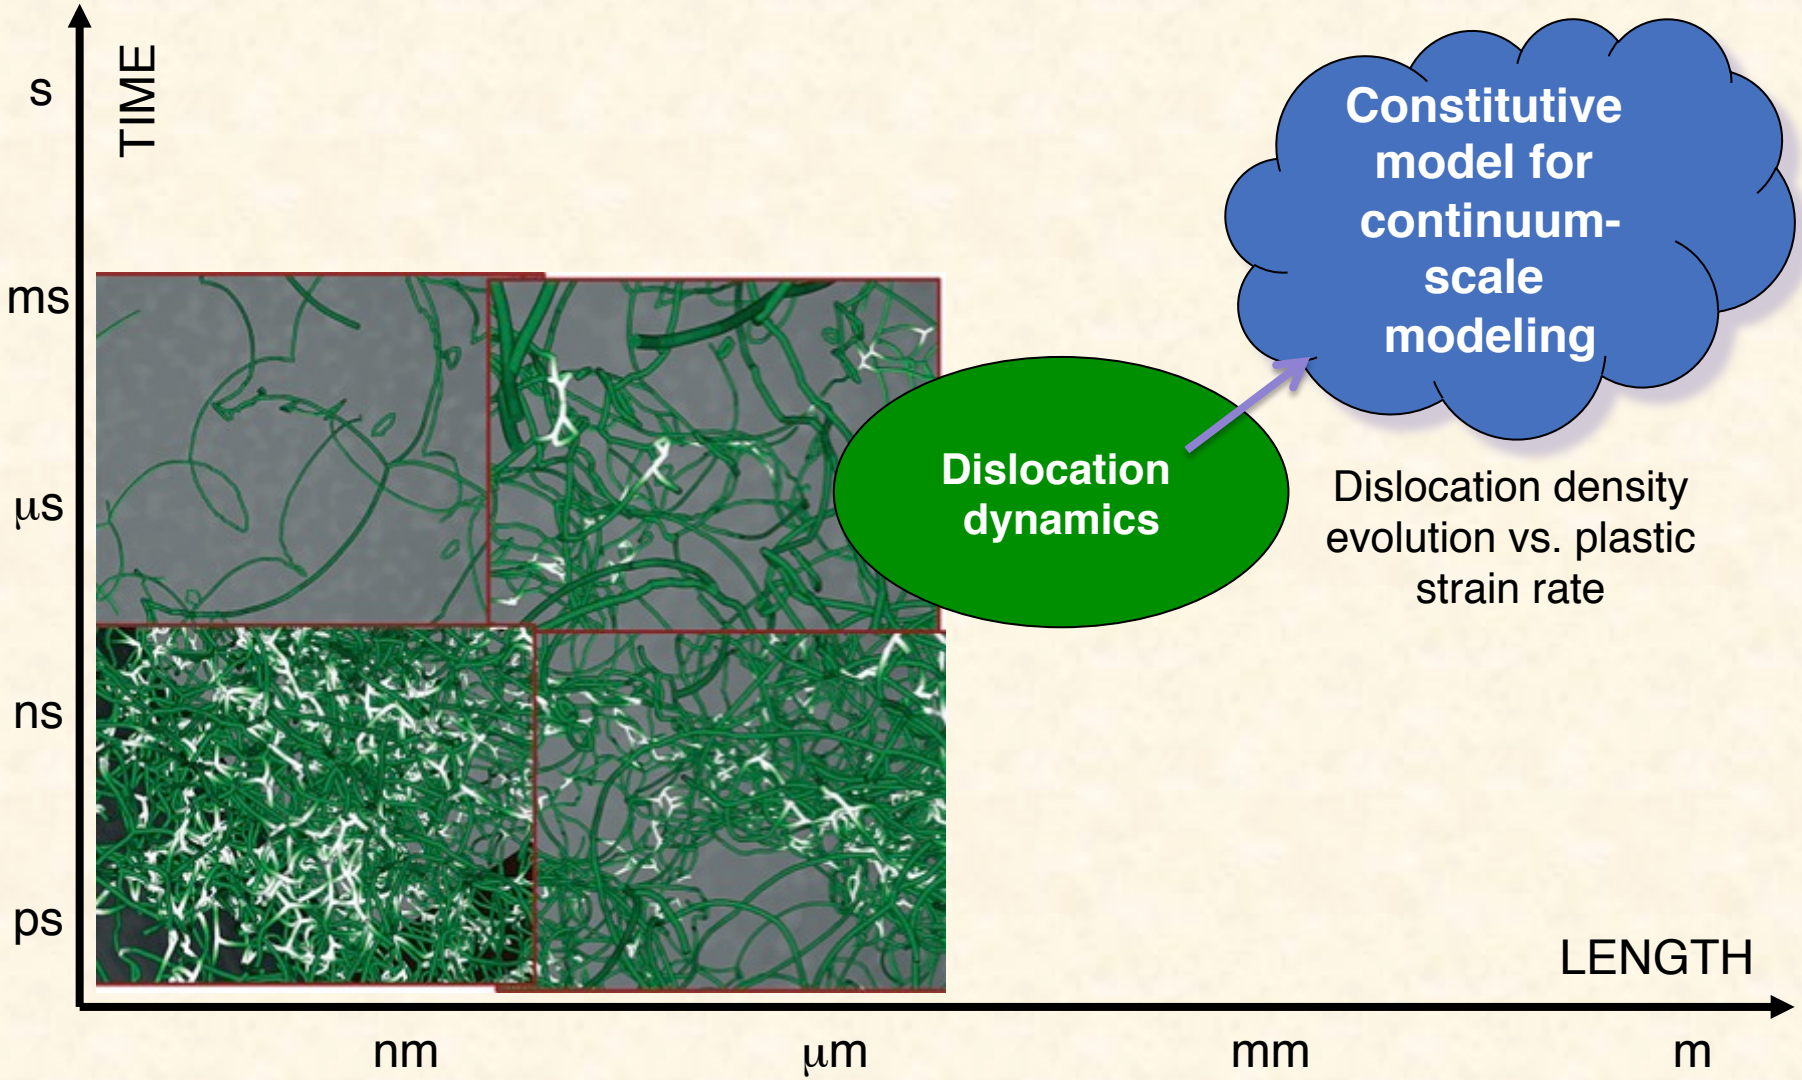

Example: "sequential" multiscale strength model

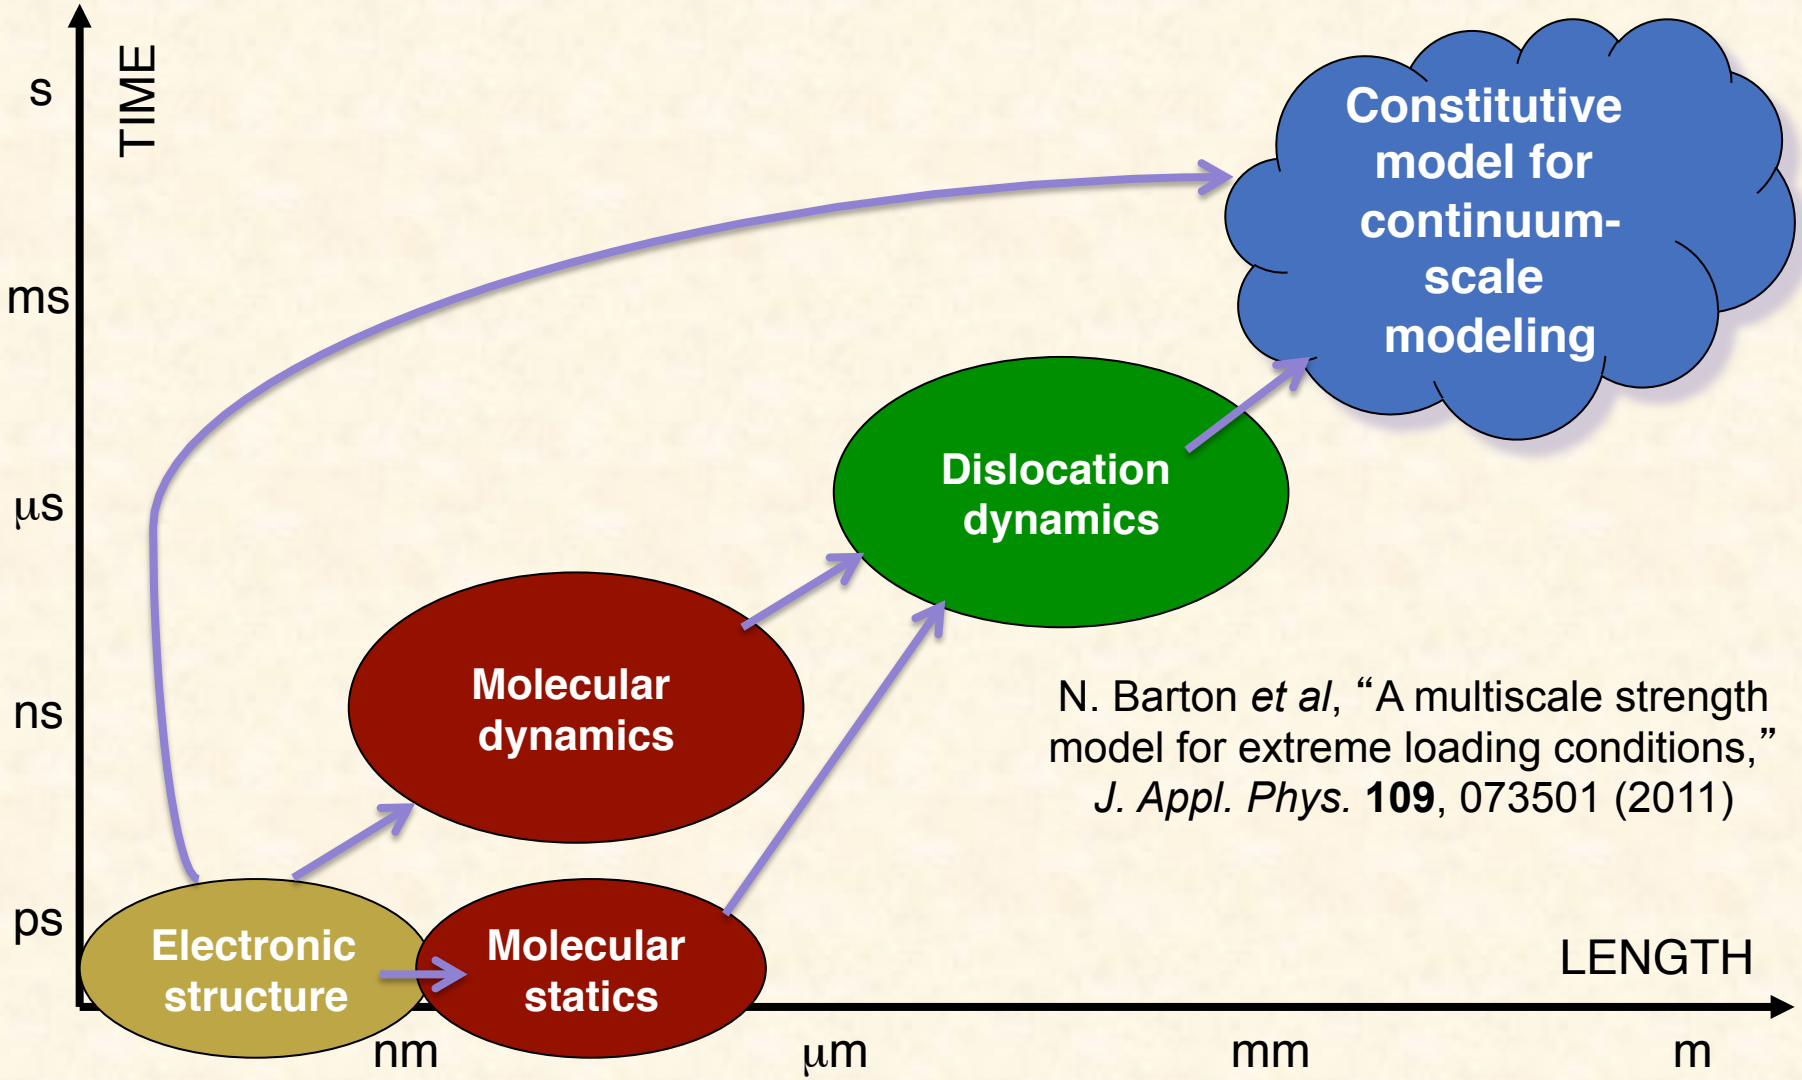

N. Barton *et al*, "A multiscale strength model for extreme loading conditions," *J. Appl. Phys.* **109**, 073501 (2011)

Challenges of a "sequential" multiscale approach

- Information is passed up a hierarchy of coupled length/time scales via a sequence of subscale models and parameters.
- This relies upon understanding how phenomena at shorter length/time scales control the behavior at longer length/time scales.
- Model complexity (and uncertainty) grows with each new physical mechanism.
  - *E.g. adding twinning and/or phase transformations to dislocation-based strength model*
  - *May need to account for coupling/competition between different physical processes*
  - *How does one include path (history) dependence (e.g., what is the strength of a material that has melted and then recrystallized?)*

Adaptive sampling techniques have been demonstrated under the LLNL "Petascale Initiative" LDRD.

- A coarse-scale model (e.g. FEM) calls a lower length-scale model (e.g. polycrystal plasticity) and stores the response obtained for a given microstructure, each time this model is interrogated.
- A microstructure-response database is thus populated.
- The fine-scale workload varies dramatically over the coarse-scale spatial and temporal domain.
- This requires dynamic workload balancing in a task parallel context.

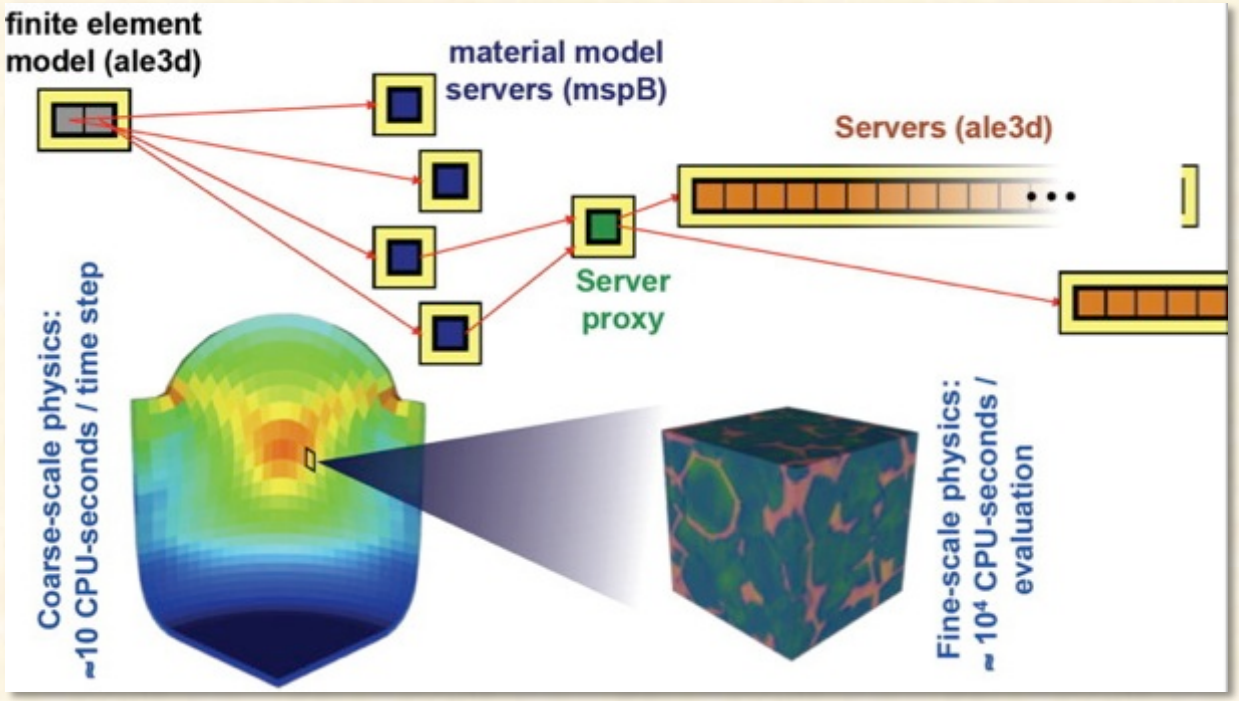

N. R. Barton, J. Knap, A. Arsenlis, R. Becker, R. D. Hornung, and D. R. Jefferson. Embedded polycrystal plasticity and adaptive sampling. *Int. J. Plast.* **24**, 242-266 (2008)

N. R. Barton *et.al.* A call to arms for task parallelism in multi-scale materials modeling. *Int. J. Numer. Meth. Engng* **86**, 744–764 (2011)

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)

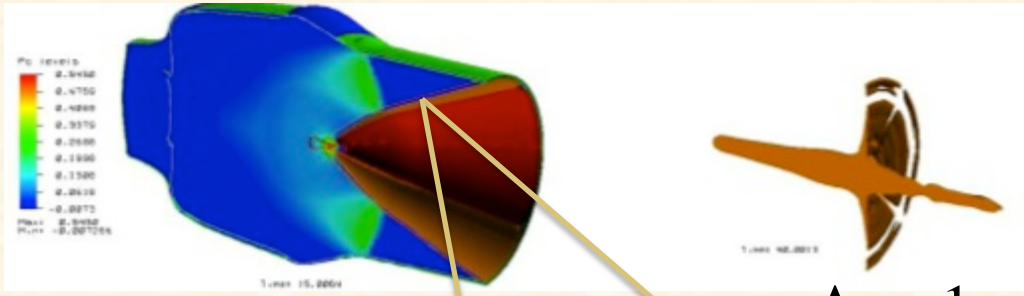

$\Delta\varepsilon \geq 1$

slow glide

$\Delta\varepsilon = 0.15$

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

Computational co-design is a process by which computer science, applied math, and domain science experts work together to enable scientific discovery

Hardware is changing dramatically

- *Increased concurrency*
- *Increased heterogeneity*
- *Reduced memory per core*
- *"Business as usual" is not going to work*

Algorithms and methods will have to be rethought / revisited for an era where:

- *Floating point operations ("flops") are (almost always) free*
- *Memory is at a premium*
- *Power is a constraint*
- *Resiliency is a challenge*

- Few domain scientists have the extended expertise "from hardware to application" to enable applications to run at exascale (10^{18} flops/sec).
- Success on the next generation(s) of machines will require extensive collaboration between domain scientists, applied mathematicians, computer scientists, and hardware manufacturers.

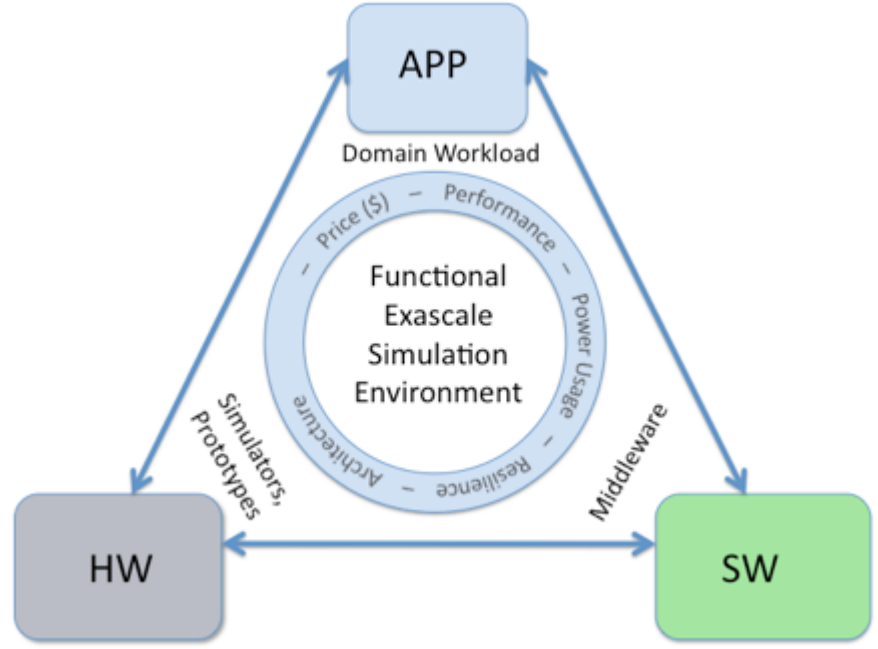

Exascale Co-Design Center for Materials in Extreme Environments

- ExMatEx is one of three* DOE/SC/ASCR co-design centers ($4M/year x 5 years) begun in August 2011

  Large scale collaborations between national labs, academia, vendors

  *Others are: CESAR (ANL/reactors), ExaCT (SNL-CA/combustion)

- Goal: to establish the relationships between algorithms, software stack, and architectures needed to enable exascale-ready science applications
- Coordinated with related NNSA/ASC tri-lab co-design efforts

Director: Tim Germann (LANL)

Deputy Director: Jim Belak (LLNL)

Exascale Computing
Proposed Timeline

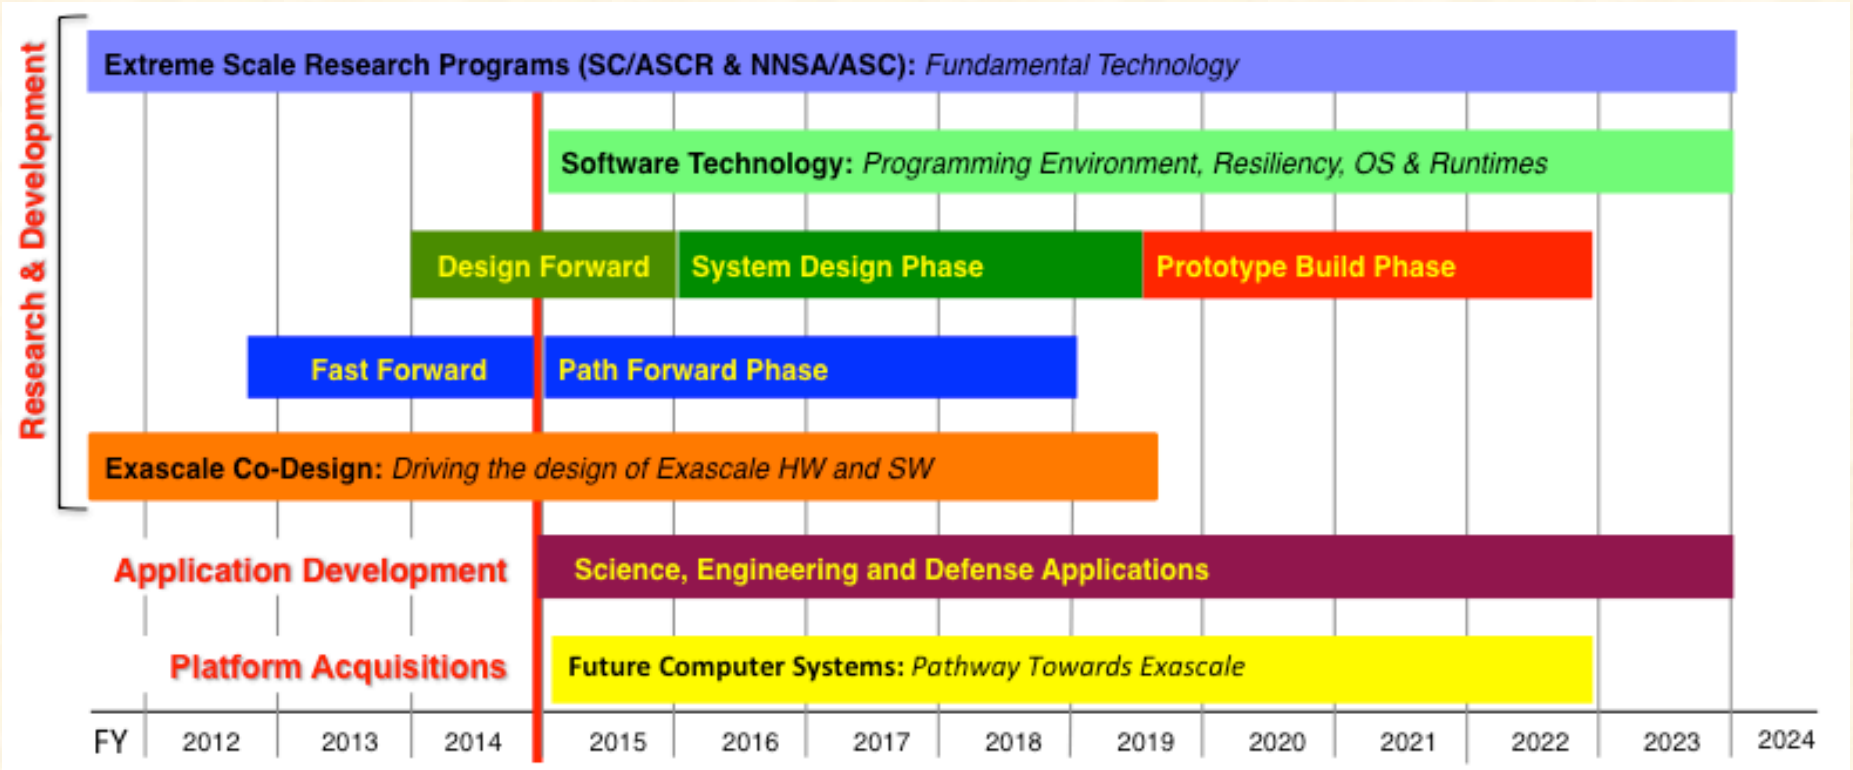

Fast Forward Projects

Fast Forward

- Jointly funded by SC & NNSA
- Two year contracts, started July 1, 2012

Project Goals & Objectives

- Initiate partnerships with multiple companies to accelerate the R&D of critical technologies needed for extreme-scale computing.
- Fund technologies targeted for productization in the 5–10 year timeframe.

Vendor	SCOPE	Value
AMD	Processor / Memory	$12,600,000
IBM	Memory	$10,476,714
Intel	Processor / Memory	$18,963,437
NVIDIA	Processor	$12,398,893
WhamCloud (Intel)	Storage & I/O	$7,996,053
Total		**$62,435,097**

Design Forward Projects

- **Design Forward**
  - Jointly funded by SC & NNSA
  - Two year contracts, started Fall 2013
  - $25.4 Million in Contracts
- **Project Goals & Objectives**
  - Initiate partnerships with multiple companies to accelerate the R&D of interconnect architectures for future extreme-scale computers.
  - Fund technologies targeted for productization in the 5–10 year timeframe.
- **Projects Funded**
  - **AMD:** interconnect architectures and associated execution models
  - **Cray:** open network protocol standards
  - **IBM:** energy-efficient interconnect architectures and messaging models
  - **Intel:** interconnect architectures and implementation approaches
  - **NVIDIA:** interconnect architectures for massively threaded processors.

Co-design requires an ongoing feedback between application, architecture, and middleware developers

- Coordinated efforts of Domain Scientists, Computer Scientists and Hardware Developers

Application Co-design Centers

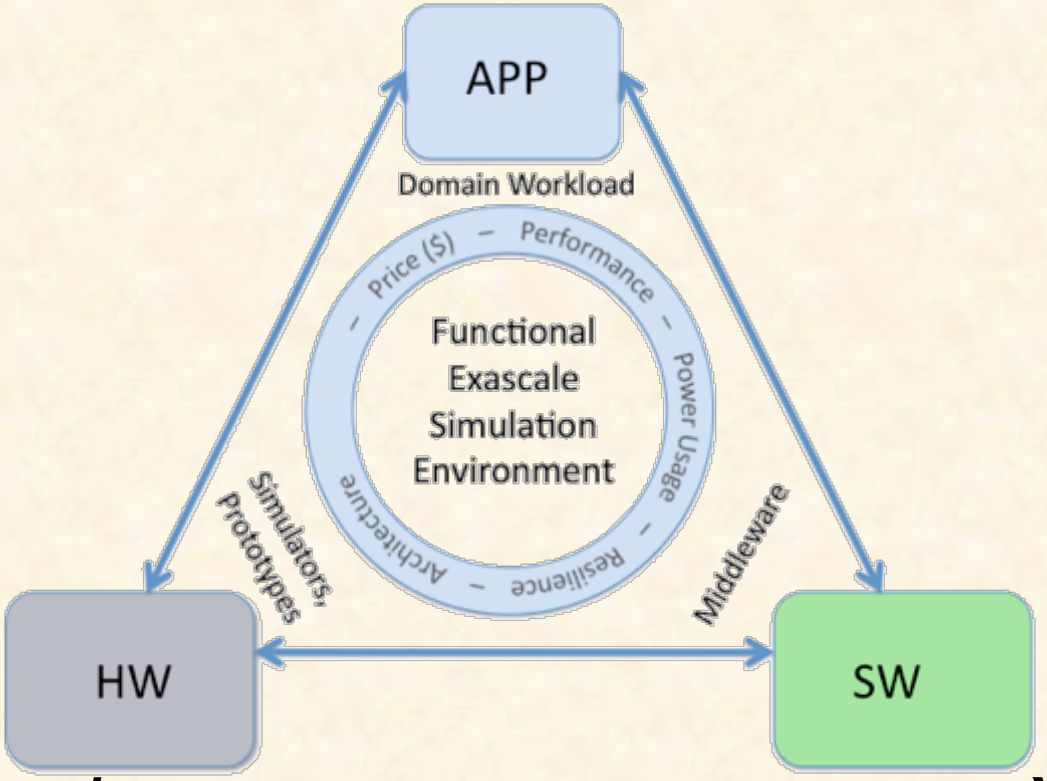

Fast Forward
Design Forward

X-Stack
Exascale OS/R

ExMatEx Co-Design Project Goal

- Our **goal** is to establish the interrelationship between hardware, middleware (software stack), programming models, and algorithms required to enable *a productive exascale environment* for multiphysics simulations of materials in extreme mechanical and radiation environments.
- We will exploit, rather than avoid, the greatly increased levels of concurrency, heterogeneity, and flop/byte ratios on the upcoming exascale platforms.

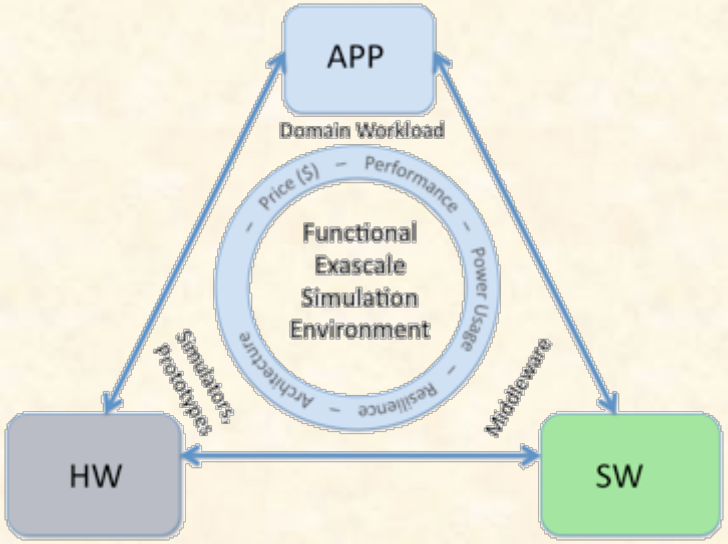

- The task-based *adaptive physics refinement* approach leverages the extensive concurrency and heterogeneity expected at exascale while enabling fault tolerance within applications.
- The programming models and approaches developed to achieve this will be broadly applicable to a variety of multiscale, multiphysics applications, including astrophysics, climate and weather prediction, structural engineering, plasma physics, and radiation hydrodynamics.

All ExMatEx activities are focused on the two ultimate objectives.

(1) Demonstrating and delivering a ***prototype*** scale-bridging materials science application based upon adaptive physics refinement.

(2) Identifying the ***requirements*** for the exascale ecosystem that are necessary to perform computational materials science simulations (both single- and multi-scale).

Seven pillars of computational materials science

Ab-initio	MD	Long-time	Phase Field	Dislocation	Crystal	Continuum
Inter-atomic forces, EOS	**Defects and interfaces, nucleation**	**Defects and defect structures**	**Meso-scale multi-phase evolution**	**Meso-scale strength**	**Meso-scale material response**	**Macro-scale material response**
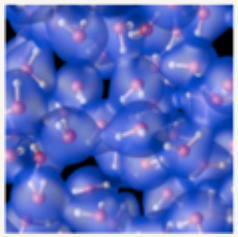	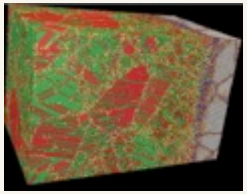	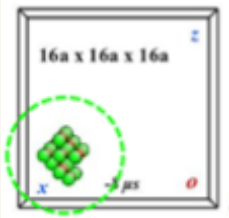	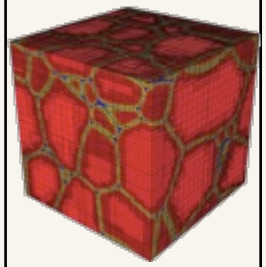	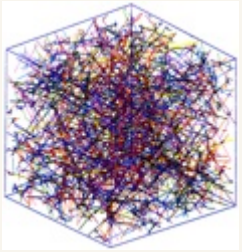	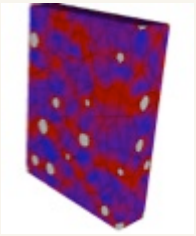	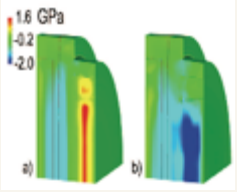
Code: Qbox/ LATTE Motif: Particles and wavefunctions, plane wave DFT, ScaLAPACK, BLACS, and custom parallel 3D FFTs Prog. Model: MPI + CUBLAS/ CUDA	Code: SPaSM/ ddcMD/CoMD Motif: Particles, explicit time integration, neighbor and linked lists, dynamic load balancing, parity error recovery, and *in situ* 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/ CoGL 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 arithmatic, 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

Co-design requires an ongoing feedback between application, architecture, and middleware developers

Workflow of Co-design between Application Co-design Centers, hardware vendors, and the broader research community

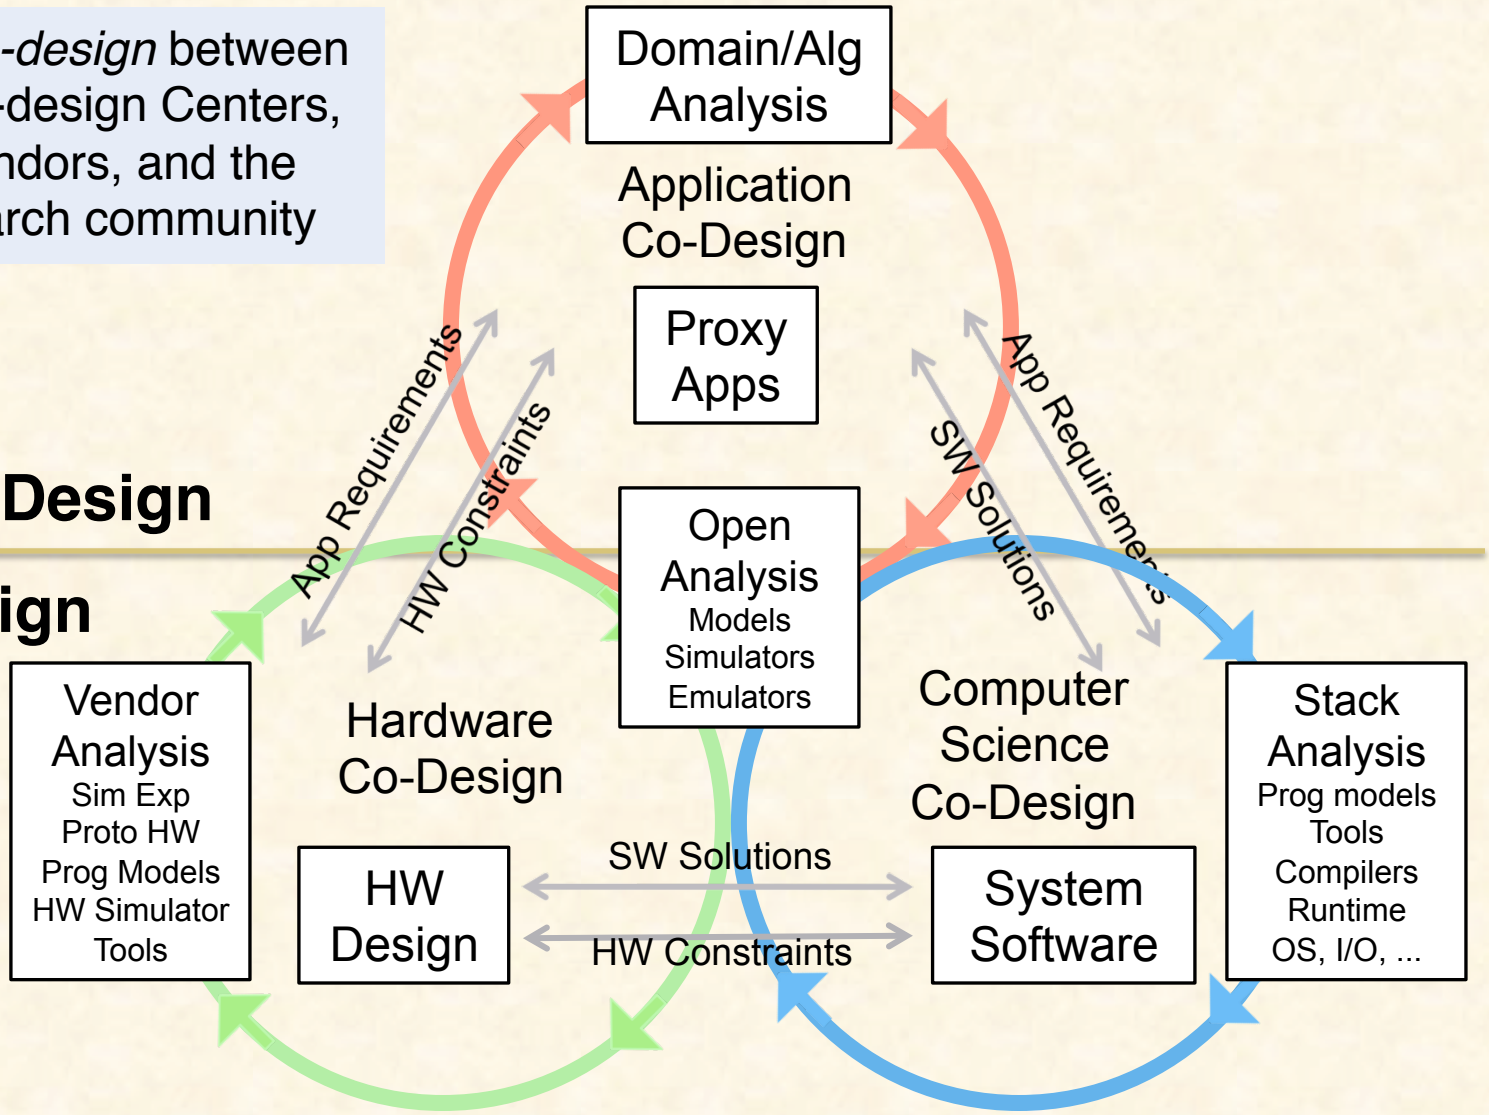

Proxy applications are central to co-design

- Proxy applications are a primary mechanism for collaboration between hardware architects, computer scientists, and domain scientists.
- Proxy apps representing the workflow have been an effective mechanism for:
  - *Identifying language/compiler weaknesses*
  - *Indicating bottlenecks that more complex computational workflows may have (vs. conventional benchmarks)*
  - *Providing tractable application testbeds for new approaches to resilience, OS/runtime/execution models, power management, …*
  - *Evaluating alternative programming models, e.g. task-based execution models & runtimes*
- Open-source Mantevo suite
  - *Sandia National Laboratories*

    + AWE, LANL, LLNL, NVIDIA

Our focus during the first 18 months was establishing the initial suite of single-scale SPMD proxy apps.

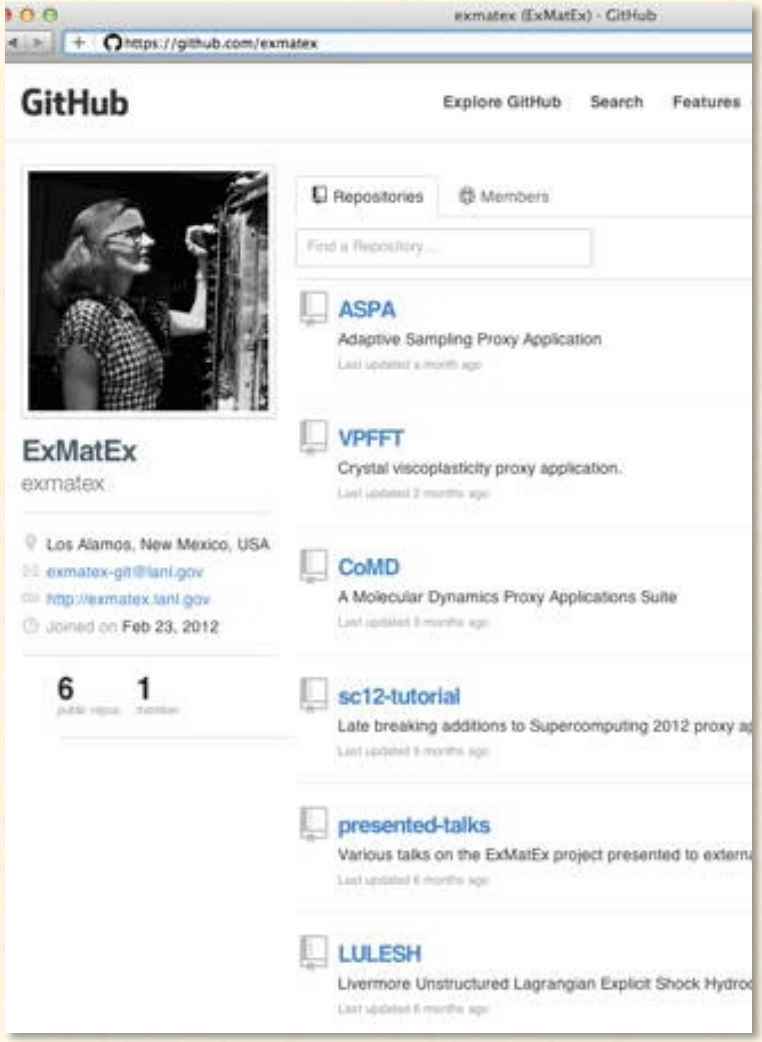

github.com/exmatex

- Single-scale proxies primarily address node-level SPMD issues:
  - *Microscale: CoMD*
    - » Molecular dynamics; particle-based
  - *Mesoscale: VPFFT, CoGL*
    - » Crystal plasticity, phase field; regular Eulerian grids (Fourier- & real-space alternatives)
  - *Macroscale: LULESH*
    - » Shock hydro; unstructured Lagrangian mesh
- CoMD and LULESH are two of the small set (~6) of compact applications that several of the vendor FastForward teams are focusing on as part of their projects.
- Several hackathons and deep dives have enhanced this collaboration.

Proxy apps are being used to identify critical features of programming models

The single-scale proxy apps developed in Year 1, primarily CoMD and LULESH, were used as the primary vehicle for the co-design process, notably several “hackathons” with vendor and X-stack partners.

From these activities, and exploration of various node and component-level programming models, several critical features were identified. Namely, they need to enable the developer to:

- Express control of workflow beyond communicating serial processes
- Express information (e.g. data dependencies) for higher-level dynamic control of workflow
- Express fine grain concurrency
- Express data locality / data layout
- Express asynchrony
- Express heterogeneity and hierarchy

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

Heterogeneous, hierarchical task-based MPMD algorithms:

- Escape the traditional bulk synchronous SPMD paradigm
- Map naturally to anticipated heterogeneous, hierarchical architectures
- Leverage concurrency and heterogeneity at exascale while enabling novel data models, power management, and fault tolerance strategies

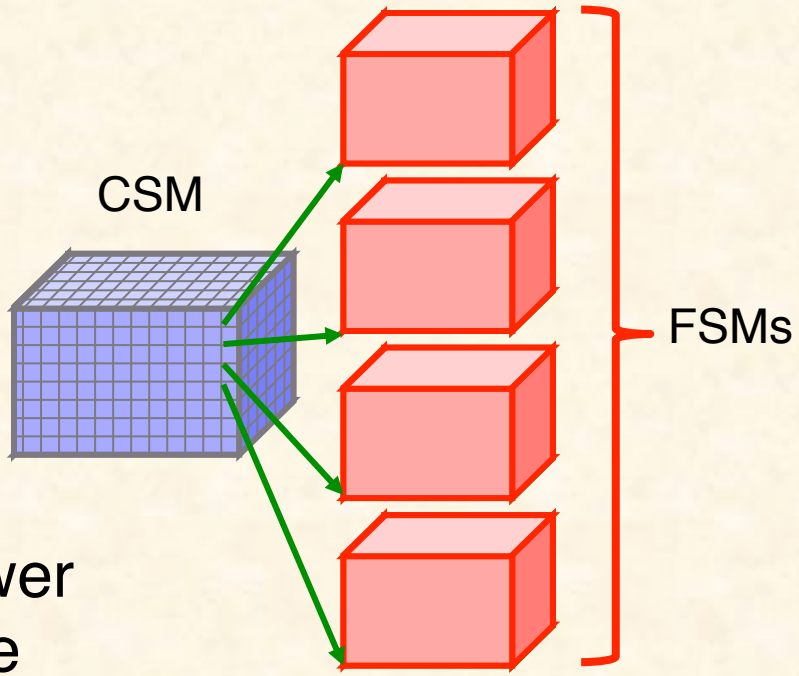

Adaptive sampling techniques have been demonstrated under the LLNL "Petascale Initiative" LDRD.

- A coarse-scale model (e.g. FEM) calls a lower length-scale model (e.g. polycrystal plasticity) and stores the response obtained for a given microstructure, each time this model is interrogated.
- A microstructure-response database is thus populated.
- The fine-scale workload varies dramatically over the coarse-scale spatial and temporal domain.
- This requires dynamic workload balancing in a task parallel context.

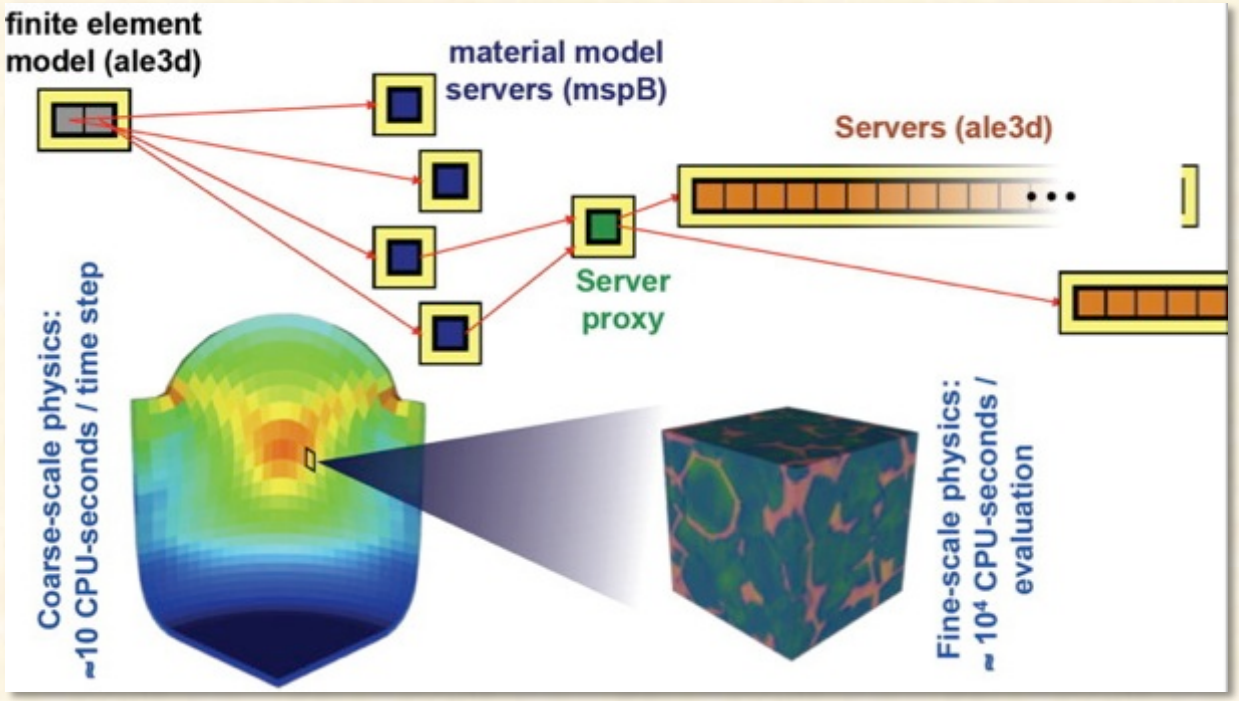

N. R. Barton, J. Knap, A. Arsenlis, R. Becker, R. D. Hornung, and D. R. Jefferson. Embedded polycrystal plasticity and adaptive sampling. *Int. J. Plast.* **24**, 242-266 (2008)

N. R. Barton *et.al.* A call to arms for task parallelism in multi-scale materials modeling. *Int. J. Numer. Meth. Engng* **86**, 744–764 (2011)

Kriging estimates are based on previously computed fine-scale responses.

Fine-scale responses accumulated in a database are interpolated (with error estimation) via a kriging algorithm.

• = fine scale evaluation
--- = linear regression model

Kriging model 1

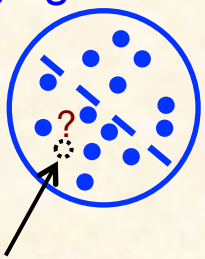

Sample point near existing model and satisfies tolerance:

- Just interpolate (saves fine-scale evaluation)

Kriging model 2

Sample point too far from existing models:

- Evaluate fine scale
- Create new model

Kriging model 3

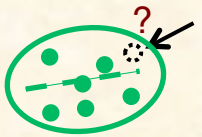

Sample point near existing model, but fails error tolerance:

- Evaluate fine scale
- Add to existing model

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

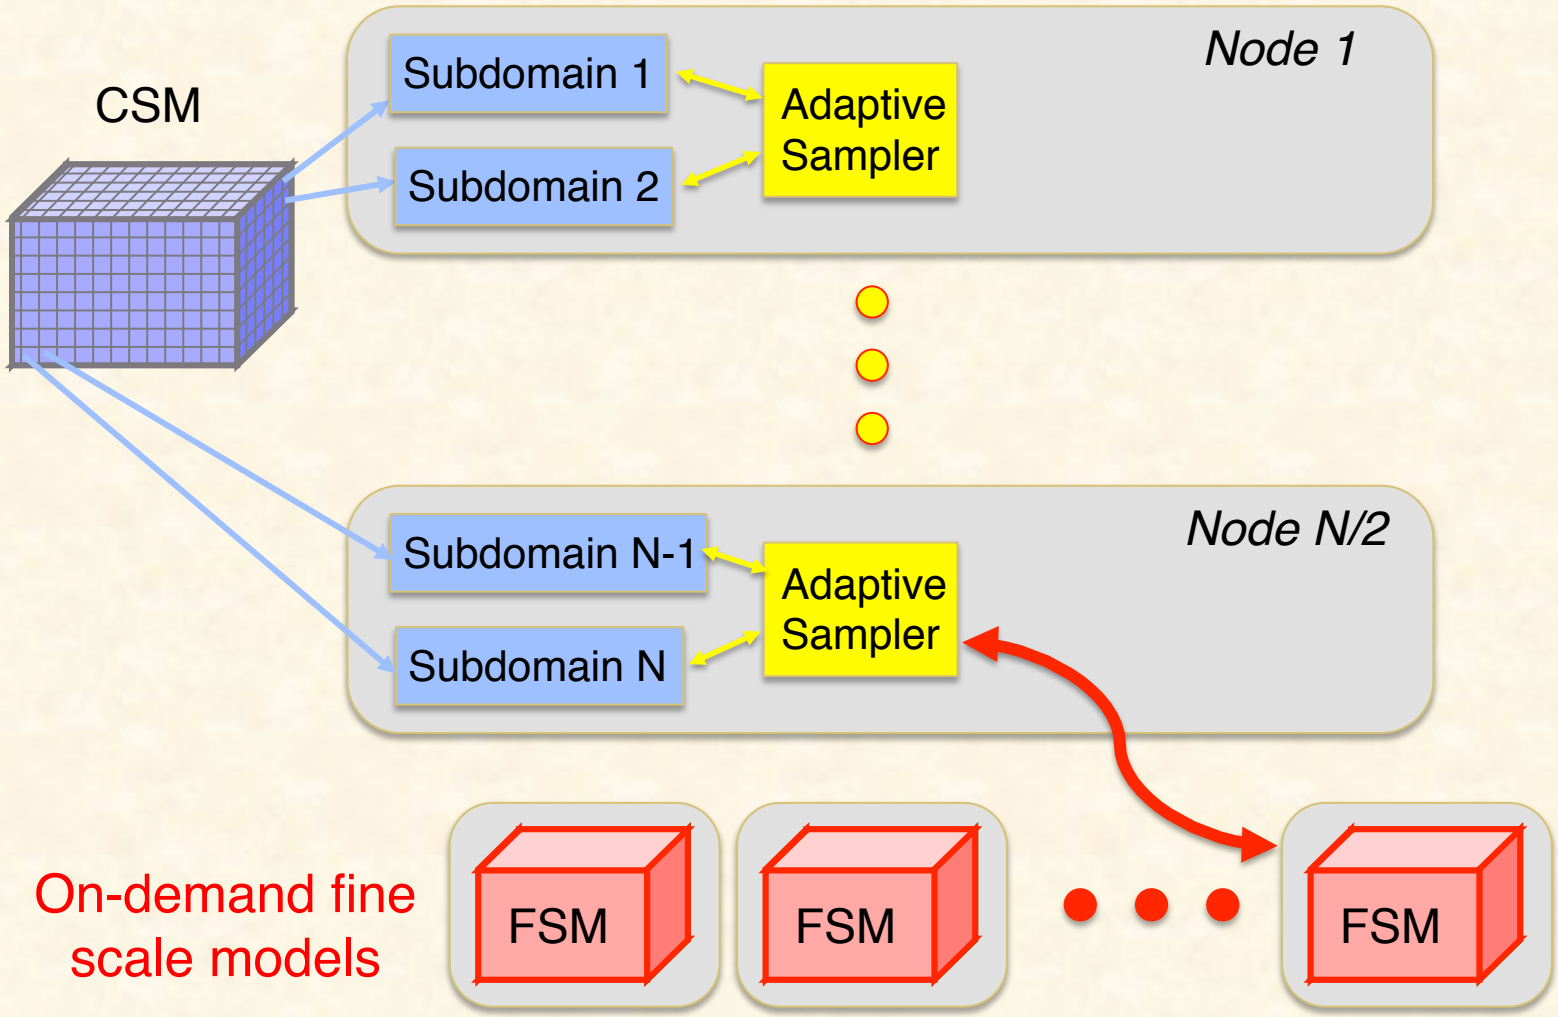

Tradeoff: re-use vs. re-computation of expensive fine-scale model results

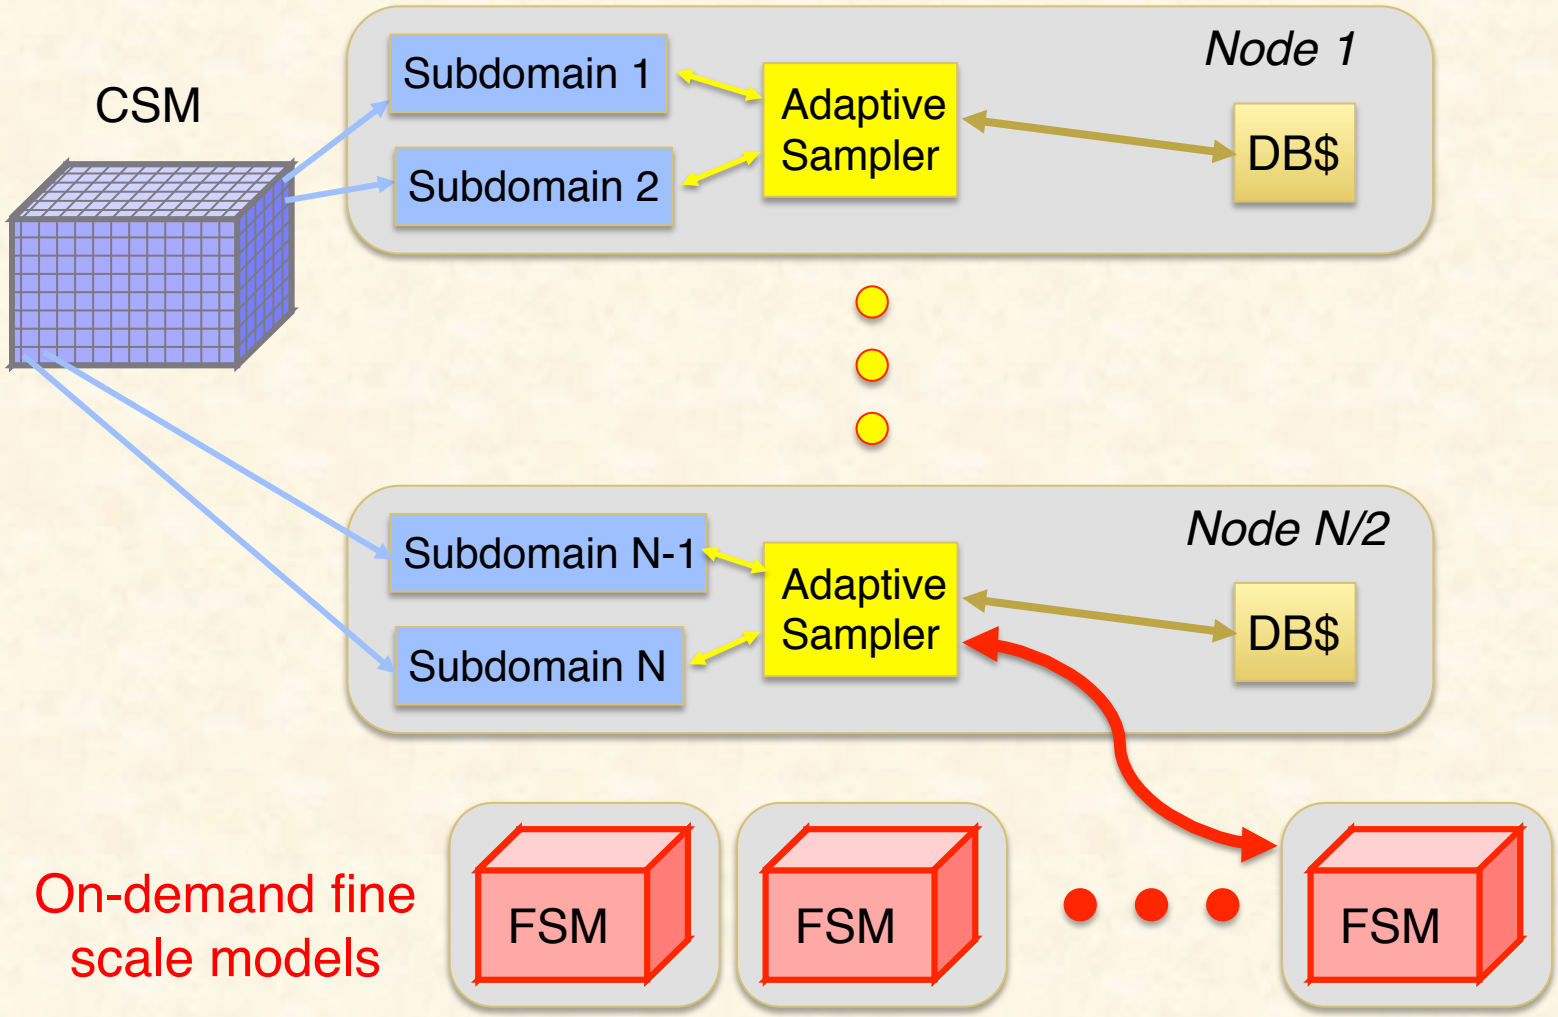

Tradeoff: re-use vs. re-computation of expensive fine-scale model results

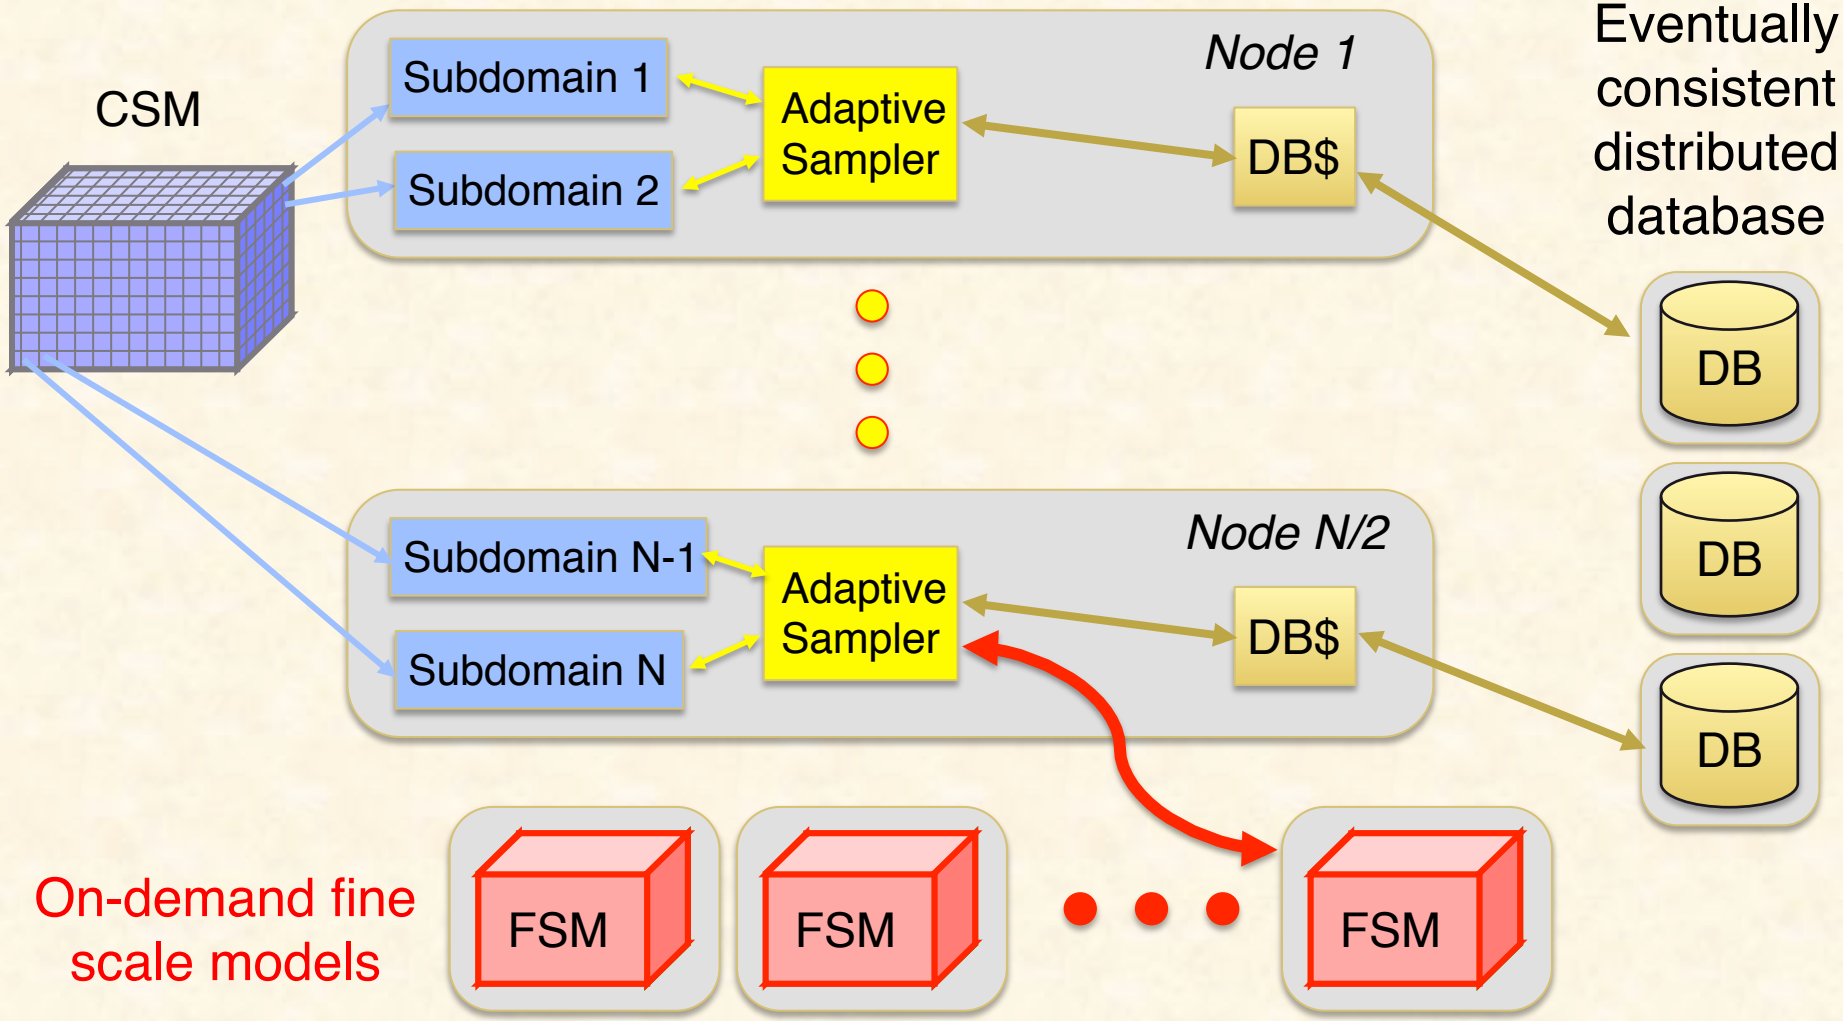

Concurrent scale-bridging approaches are being pursued in other materials science contexts

- Directly computing a potential surface from *ab initio* calculations
  - *Distinct from* ab initio *molecular dynamics*
  - *GAP: Gaussian approximation potentials*
    - » Bartók et al, PRL **104**, 136403 (2010)
  - *SNAP: Spectral neighbor analysis potentials*
    - » Aidan Thompson et al, SNL-NM
  - *Configuational database-driven dynamics*
    - » Jones and Shaughnessy, SNL-CA

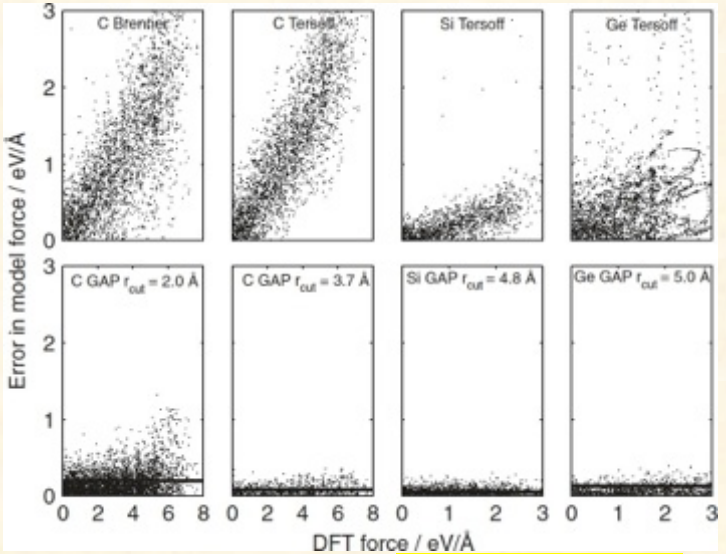

- On-the-fly kinetic Monte Carlo
  - *Henkelman and Jonsson, J. Chem. Phys. **115**, 9657 (2001)*
- Self-learning kinetic Monte Carlo
  - *Trushin et al, Phys. Rev. B* **72**, 115401 (2005)
- Self-evolving atomistic kinetic Monte Carlo
  - *Xu, Osetsky, and Stoller, Phys. Rev. B **84,** 132103 (2011)*

GAP/SNAP

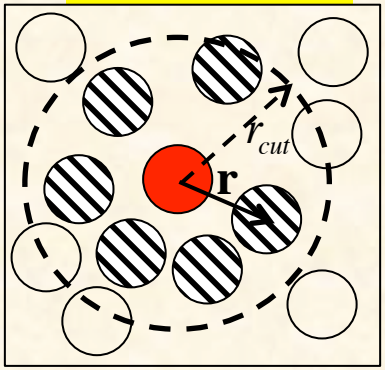

Our work on scale-bridging has followed two complementary paths.

- "Top-down"
  - We have developed an Adaptive Sampling Proxy App (ASPA) that represents the fine-scale query, database lookup, and kriging interpolation steps.
  - LULESH (coarse-scale) and VPFFT (fine-scale) proxies are coupled via ASPA to study the workflow for our target application problems.
    - » "Speeds & feeds"
    - » What are the frequency, number, and duration of fine-scale calculations?
    - » What size and type of data are communicated between scales?

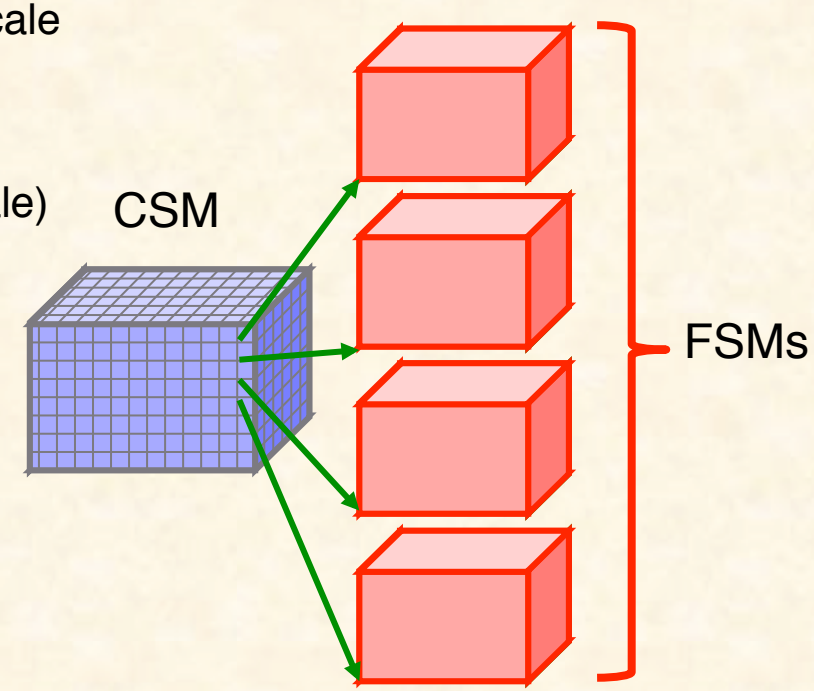

- "Bottom-up"
  - We have developed a tractable scale-bridging proxy (CoHMM) that represents the basic task-based modeling approach we are targeting.
  - It is being used to evaluate task-based OS/runtime requirements.

We are using the Heterogeneous Multiscale Method* as a scale-bridging prototype

- CoHMM presents the basic workflow requirements of a scale-bridging materials application.
- A full fine scale model (FSM, e.g. a crystal plasticity or molecular dynamics model) is run for every zone & time step of coarse scale model (CSM, e.g. an ALE code).
- It is being used to assess *basic* requirements for task-based runtime systems.
  - *The original HMM* is limited by its predictable, uniform workload pattern.*
  - *Adaptive coarsening provides a more dynamic and realistic workload.*

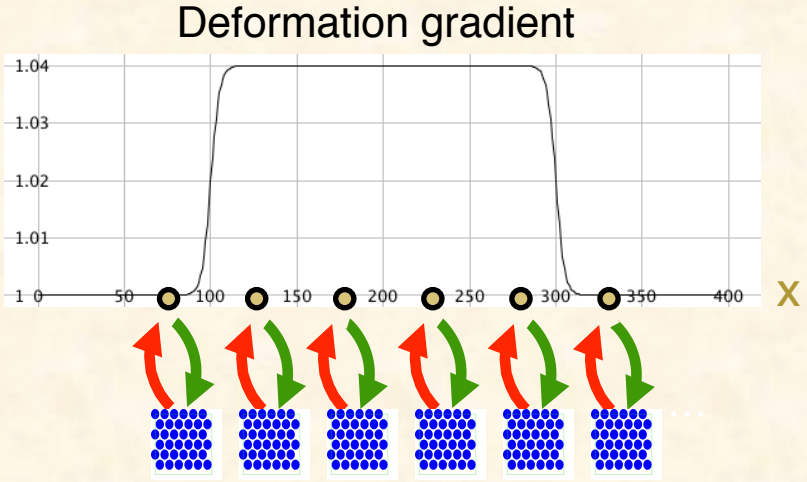

*Xiantao Li and Weinan E, "Multiscale modeling of the dynamics of solids at finite temperature," *J. Mech. Phys. Solids* **53**, 1650–1685 (2005)

Emerging approach to subscale models: "concurrent multiscale"

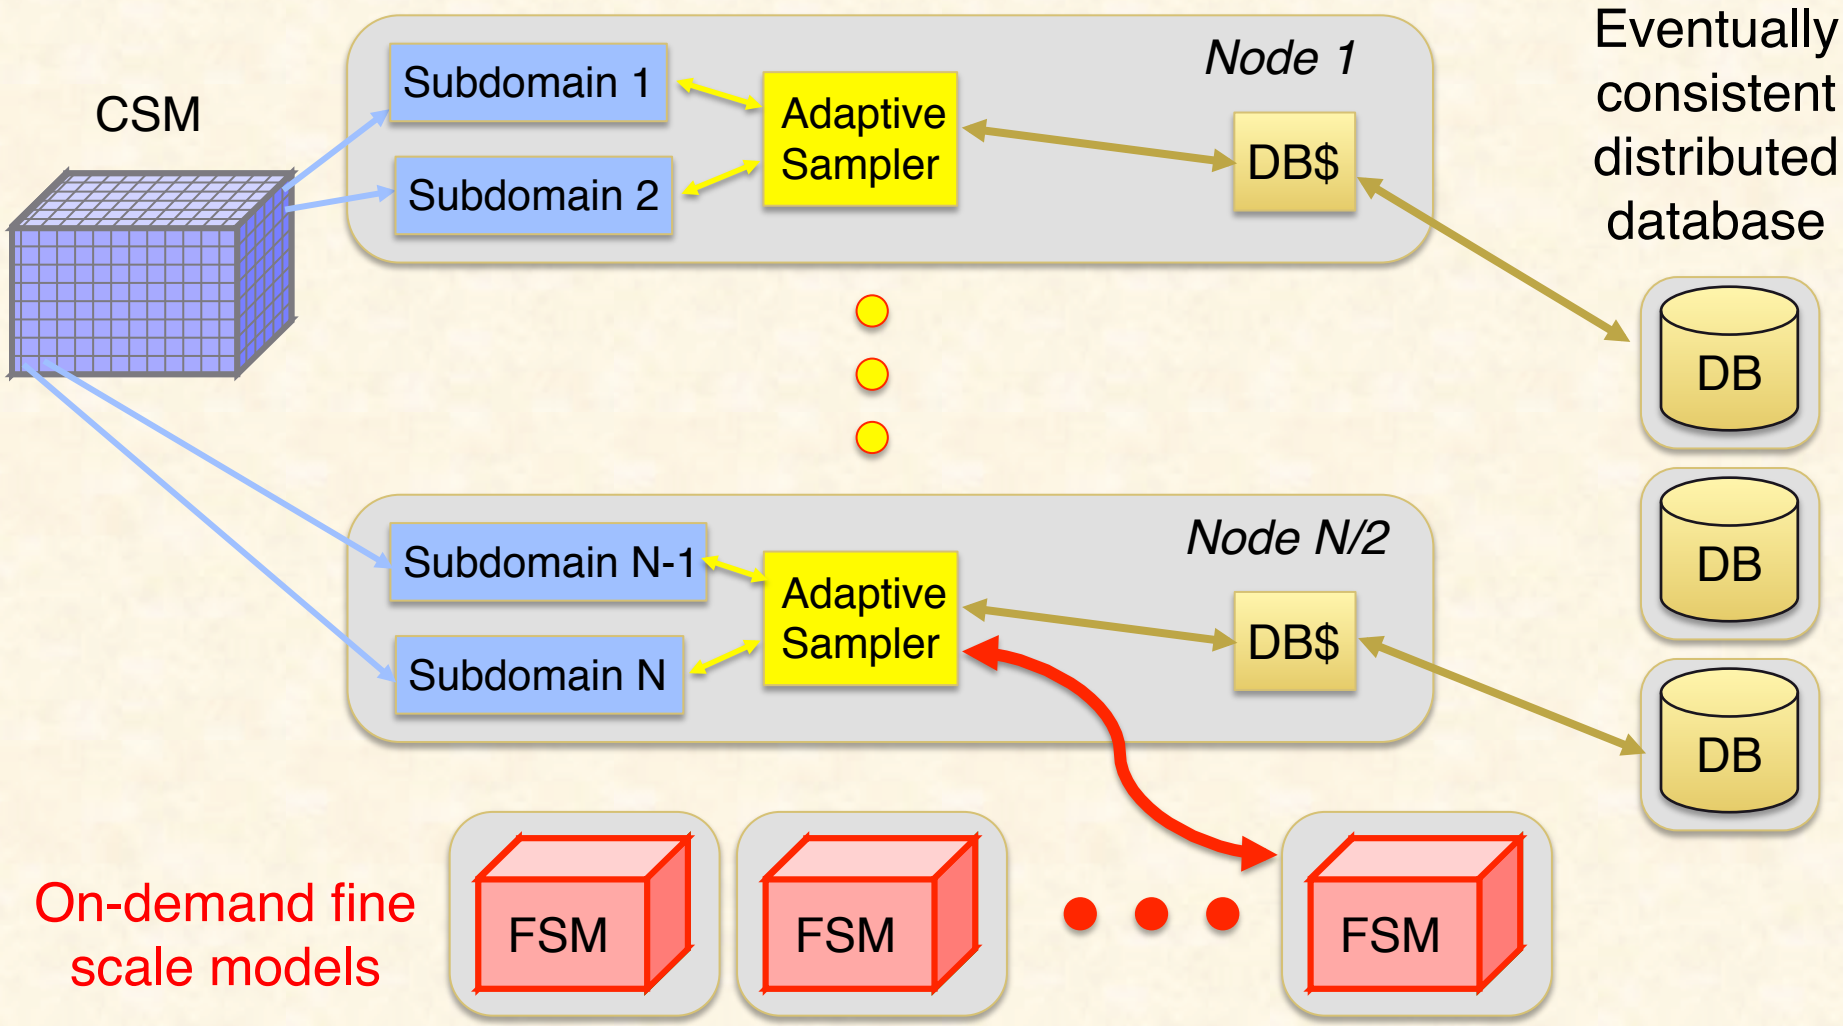

Traditional approach to subscale models: "sequential multiscale"

- Subscale models (e.g. interatomic potentials, equation of state and strength models) are developed from a combination of theory, experiment, and simulation.
  - *The specific combination depends on the developer, and may involve as much art as science.*

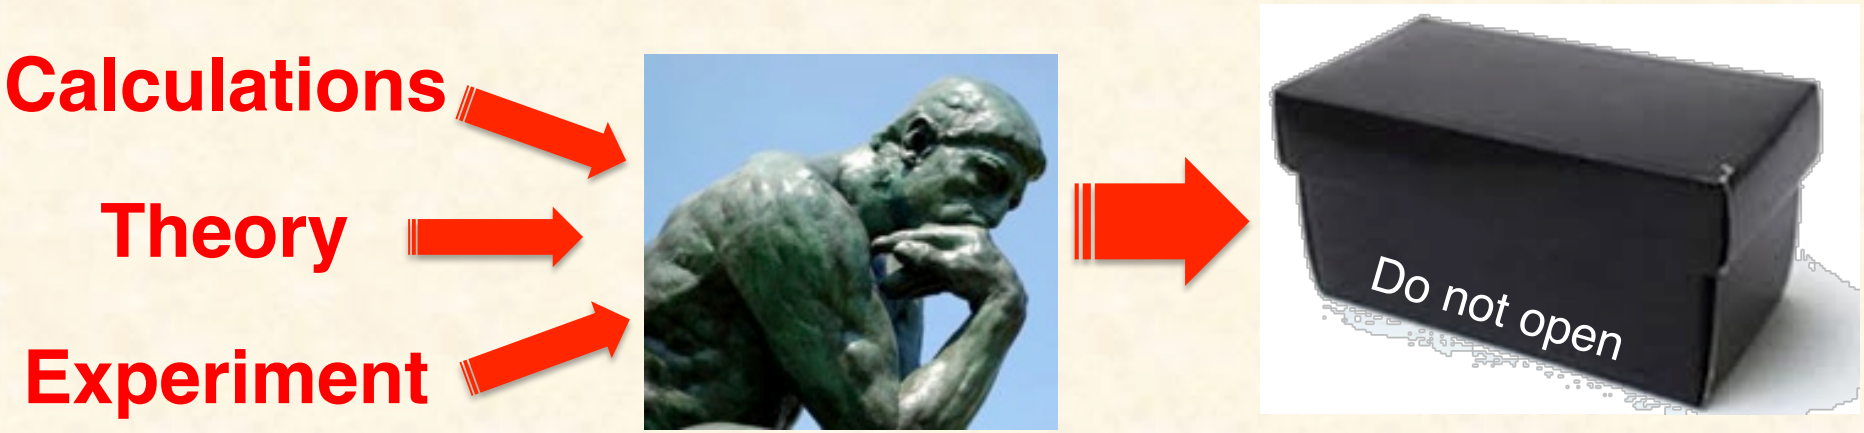

Integrating experimental data (and theory)

- Kriging interpolation models based on fine-scale calculations may be augmented by analytic models (e.g. near equilibrium) and/or small-scale experiments.

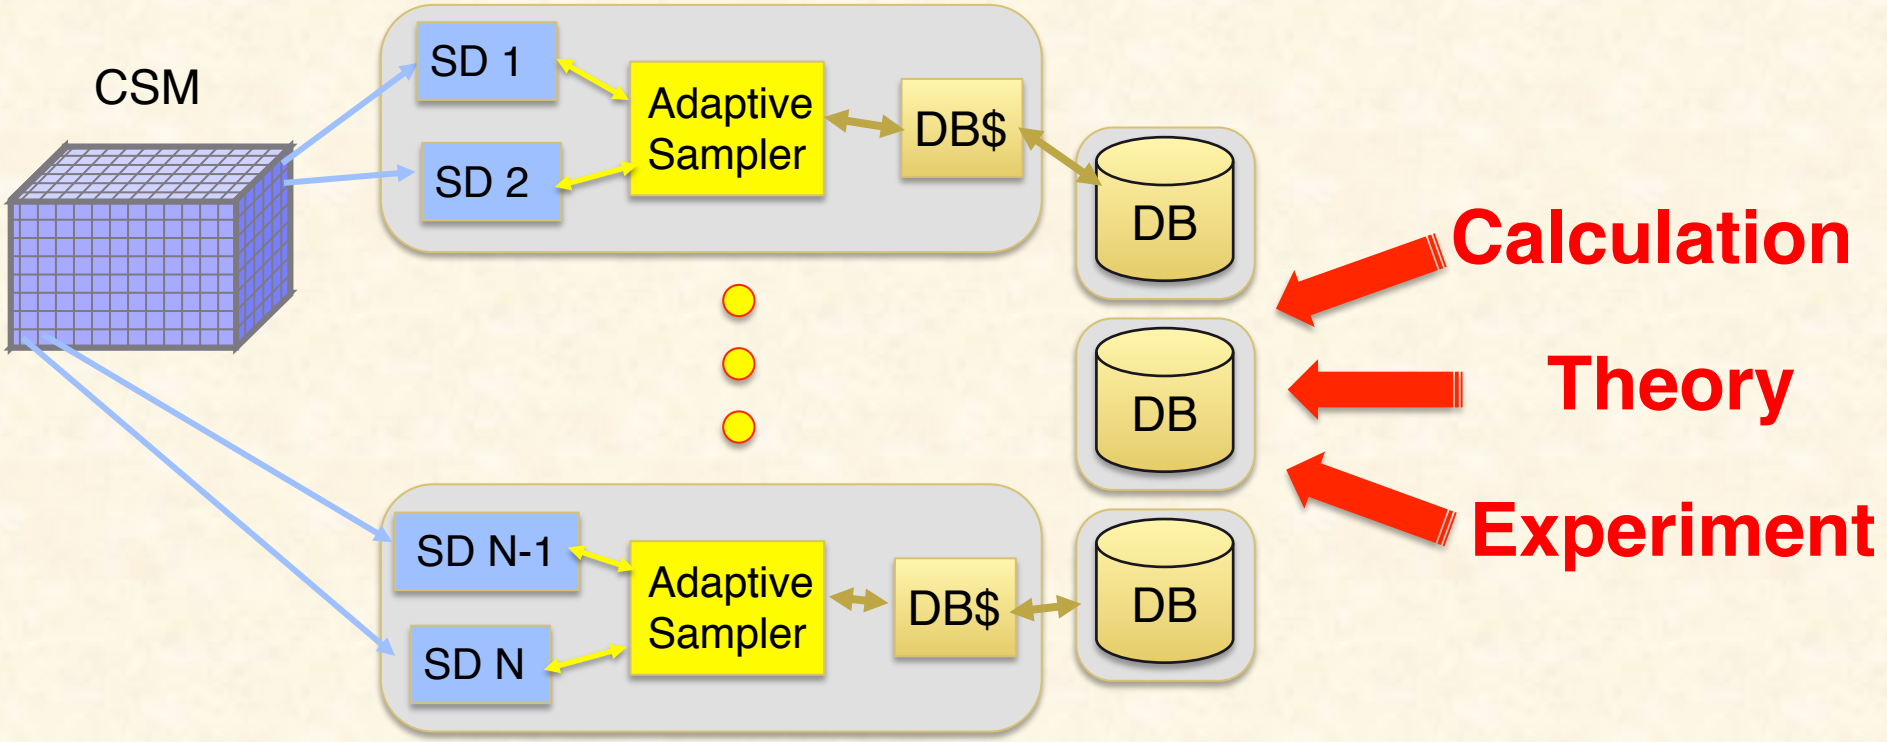

Summary

- Exascale-class materials simulation workloads will involve both large single-scale and scale-bridging simulations.
- The traditional bottom-up, sequential multiscale approach is being replaced by a top-down, on-demand scale-bridging one in a variety of materials simulation contexts.
  - *"All models are wrong, but some are useful" – G. Box*
  - *Caveat emptor : beware of GIGO*
- Computational co-design via proxy apps is establishing the relationships between algorithms, software stack, and architectures needed to enable exascale simulations.


ExMatEx