LA-UR- 12-00143

Approved for public release; distribution is unlimited.

Title: Computational ~~and~~ & Computer Science at Los Alamos National Laboratory

Author(s): Alan Bishop, PADSTE

Intended for: GAO Review, January 12, 2012

Los Alamos NATIONAL LABORATORY EST. 1943

Computational ~~and~~ & Computer Science at Los Alamos National Laboratory

By

Alan Bishop, Principal Associate Director for Science, Technology and Engineering

Presented to

Harvard University Engineering School

January 13, 2012

Abstract

Dr Bishop presents an overview of computational and computer science activities at Los Alamos National Laboratory. This briefing discusses new trends in computing science (co-design, exascale, petaflop delivery, MaRIE, earth models, and the future of energy technology) under development at LANL and applications for this area of technical expertise to the broader mission at Los Alamos. Detailed information on the new ExMatEx co-design project is also provided.

Computational & Computer Science at Los Alamos National Laboratory

Alan Bishop, Principal Associate Director
Science, Technology & Engineering

January 13, 2012

Computational Co-Design: High-Performance Computing at a Critical Time of Challenges and Opportunities

High expectations (from White House, DOE, industry…) as fundamental enabling capability: Accelerating the discovery and prediction to impact and commercialization cycles to meet (inter-)national imperatives!

Major (disruptive) technology transitions

- Hybrid/multicore…parallelism not clock speed
- Cooling, power, resiliency, concurrency, file systems, networking, agile programming environments for broad S,T&E portfolio… managing the complexity trade-offs

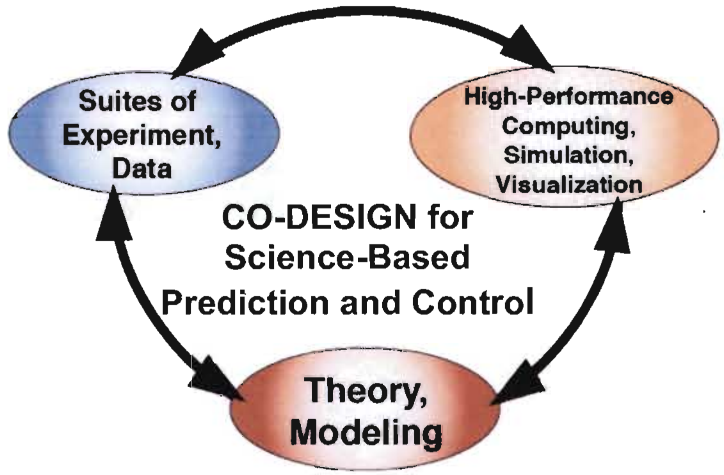

Multiscale/Multiphysics: V&V, UQ

Compute-and-store AND data-intensive frontiers

Training the next generation of leaders

The DOE road to "Exascale"
Multi-Lab associations, including vendor partners
Co-Design Centers (Applications, Algorithms, Architectures)

High-Performance/Advanced/Supercomputing and Visualization

A "NEW-ly" mature scientific capability for the scientific method

Many leaders emerged from traditional disciplines (Physics, Mathematics….); increasingly an identified discipline at universities (Computer Science…)

~1950

Interdisciplinary Teaming & Partnering

Co-Design of Assets at Scale

2020

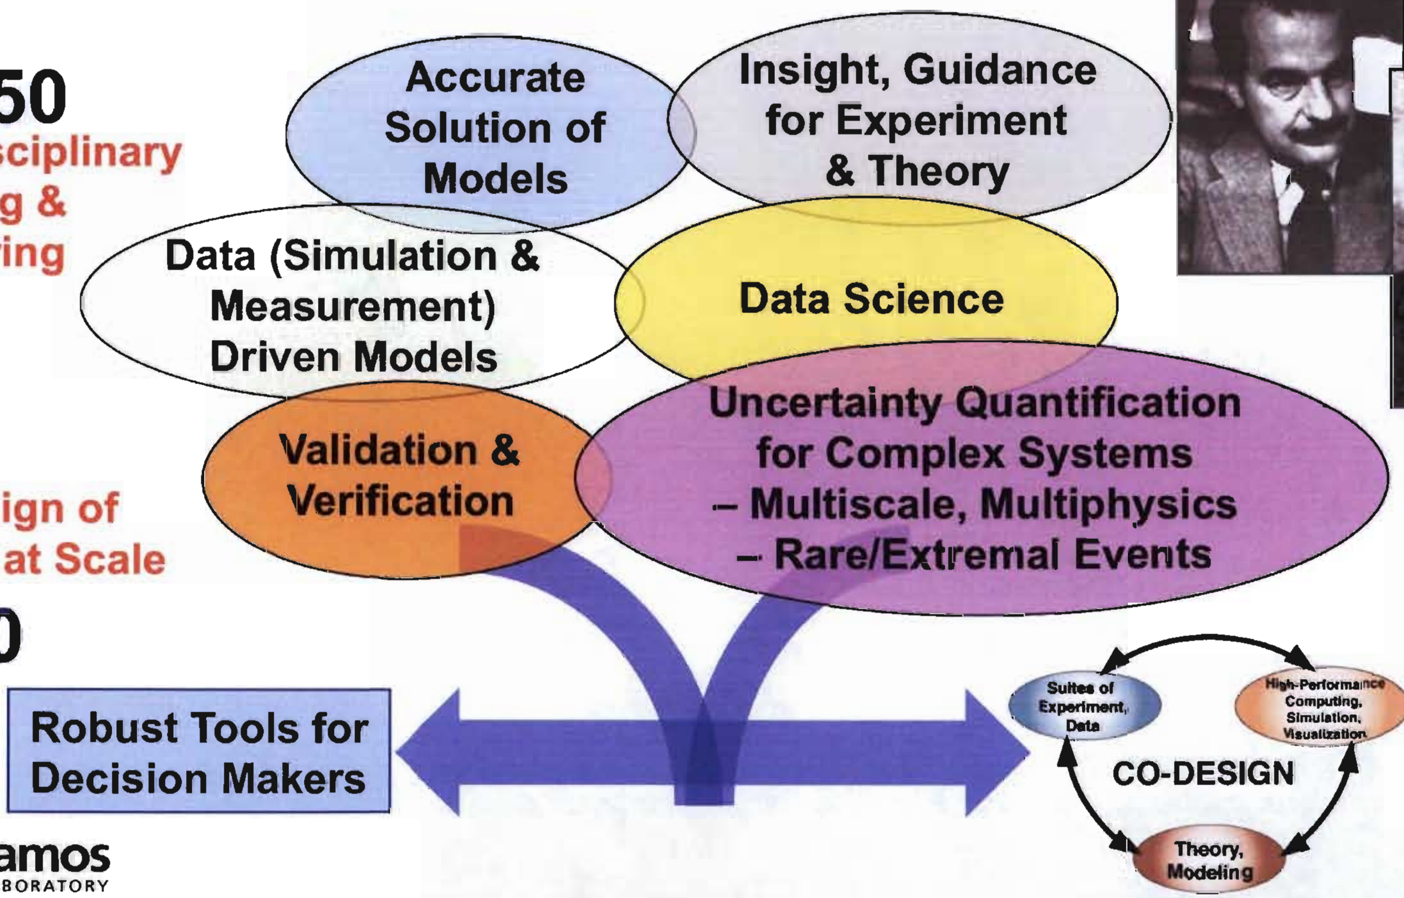

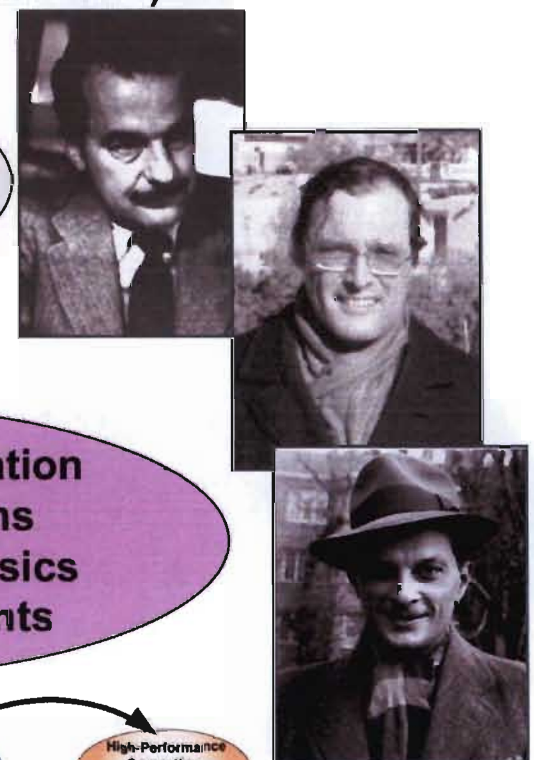

Science at Los Alamos National Laboratory: Broad and Deep

Stockpile Stewardship

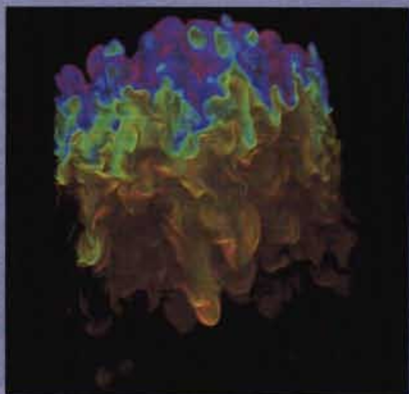

Hydrodynamics
Turbulence

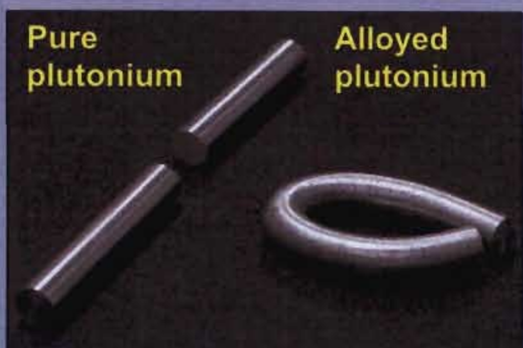

Plutonium Science
Metallurgy

Global Security

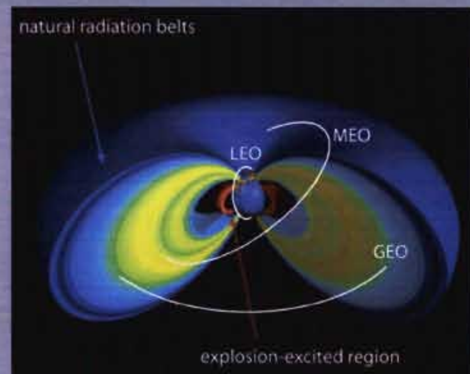

Threats from Space
Dynamic Radiation Environment Assimilation Model

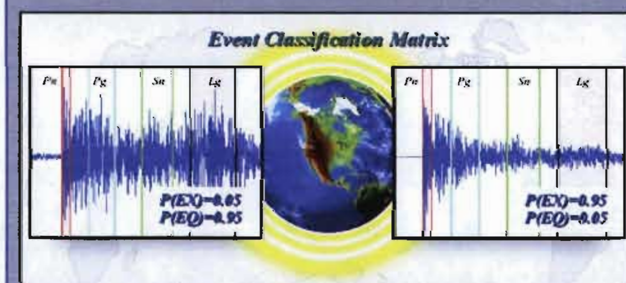

Seismic Detection of Nuclear Explosions

Energy Security

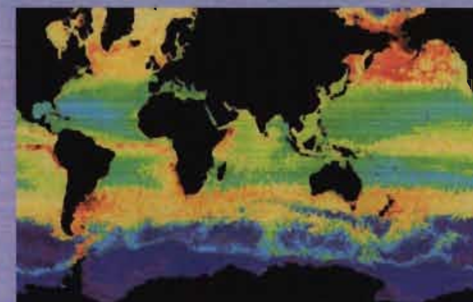

Climate/Energy Impacts
Measurement, simulation, prediction

Materials
Energy Generation and Transmission

"The Century of Complexity" (S. Hawking)

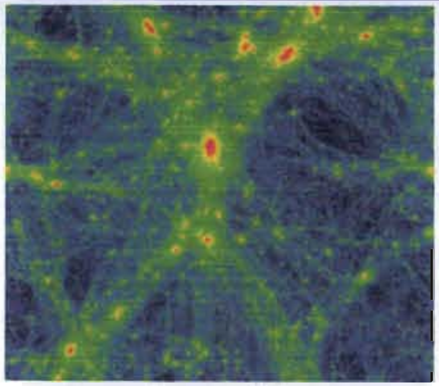

Cosmology: Filaments, Clusters, and Voids

Science @ Scale

Systems of connected functional scales space, time; Emergent functions; Extreme conditions

Enabled by huge advances in Data, Simulation, Nonlinear Science…

BUT….

? Origins, Measures, Consequences ?

Multiscale Modeling, Simulating, Measuring ≠…at Multiple Scales: Need IS&T

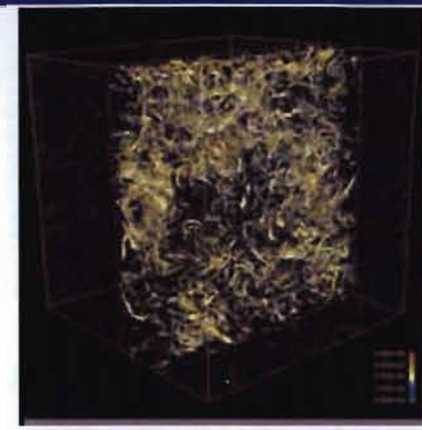

Fluid Turbulence

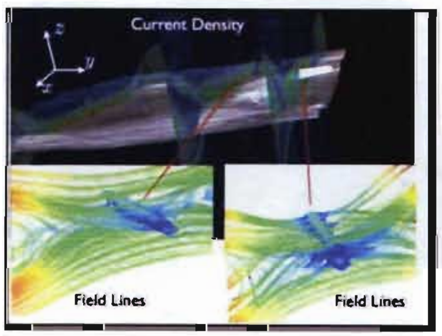

Magnetic Reconnection

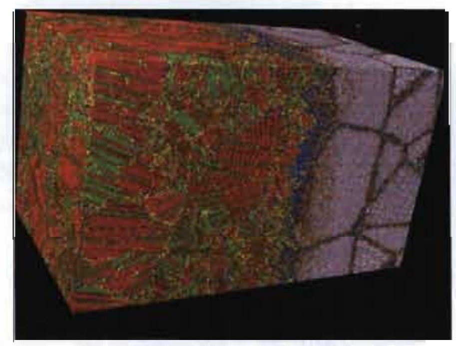

Shocked Metals

Communication Networks

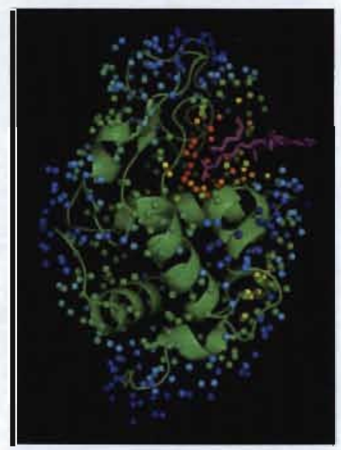

Protein Dynamics

Our vision for high-performance computing embraces both multi- and unit-physics codes

Lay the groundwork for advanced architectures that can significantly increase simulation performance in the future (Milagro, Sweep3D)

Multi-scale, multi-physics codes

Complex & varied physics models

100x slower than Unit-physics codes

under-resolved

Provide resources necessary for predictive <u>science at scale</u> simulations (VPIC, SPaSM)

Experiments (NIF, ZR ...)

- **Complex & varied physics**
- **Difficult @ high-energy-density**
- **Under-resolved diagnostics**
- **Tend to be integral**

Unit-physics codes

- **Few physics assumptions**
- **100x faster than Multi-physics**
- **Fully resolved**
- **Simple, repetitive manipulations**

The Promise and Challenge of S&T for Complex Networks

Isolating complicated phenomena to "understand" them is not always sufficient

science of prediction **&** **uncertainty quantification**

for complex systems/networks

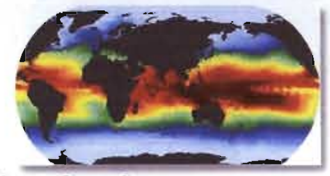

NW
Nuclear stewardship, non-proliferation

global security
Cyber, IND, borders, infrastructure

energy & environment
Climate, carbon seq, grids, EM, nuclear

complex materials
Nano S&T, extreme conditions

biological systems
Disease, vaccination, cognition

Quantitative tools for decision makers/risk assessment

- (coupled) socio, economic, humanities, physical sciences …
- from observation to prediction and uncertainty quantification

"Co-Design": A future with maximum impact

A national S,T,E management challenge: How we do business

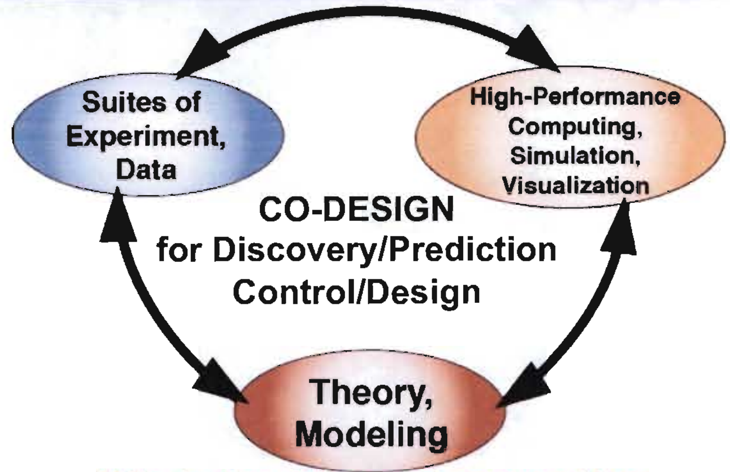

What should a next-generation user facility look like?

- **A new framework for transformational S,T, E at Science & Mission Frontiers**
  Integration and collaboration (DOE: EFRCs, SciDAC, Hubs, Co-Design Centers…)
  (OSTP: Adv. Manufacturing, Materials Genome…Initiatives)

- **LANL opportunities being developed: NW predictive capability framework, energy-climate, informatics, environmental management, cyber, … MaRIE**

DOE (SC, NNSA, App. Energy) has a full spectrum of assets for this future
Integrating National Assets for Discovery, Prediction, Control, Design

LANL must provide the infrastructure for "Connecting the Dots"

"Quantitative X"

Finding the Dots

Understanding the Systems

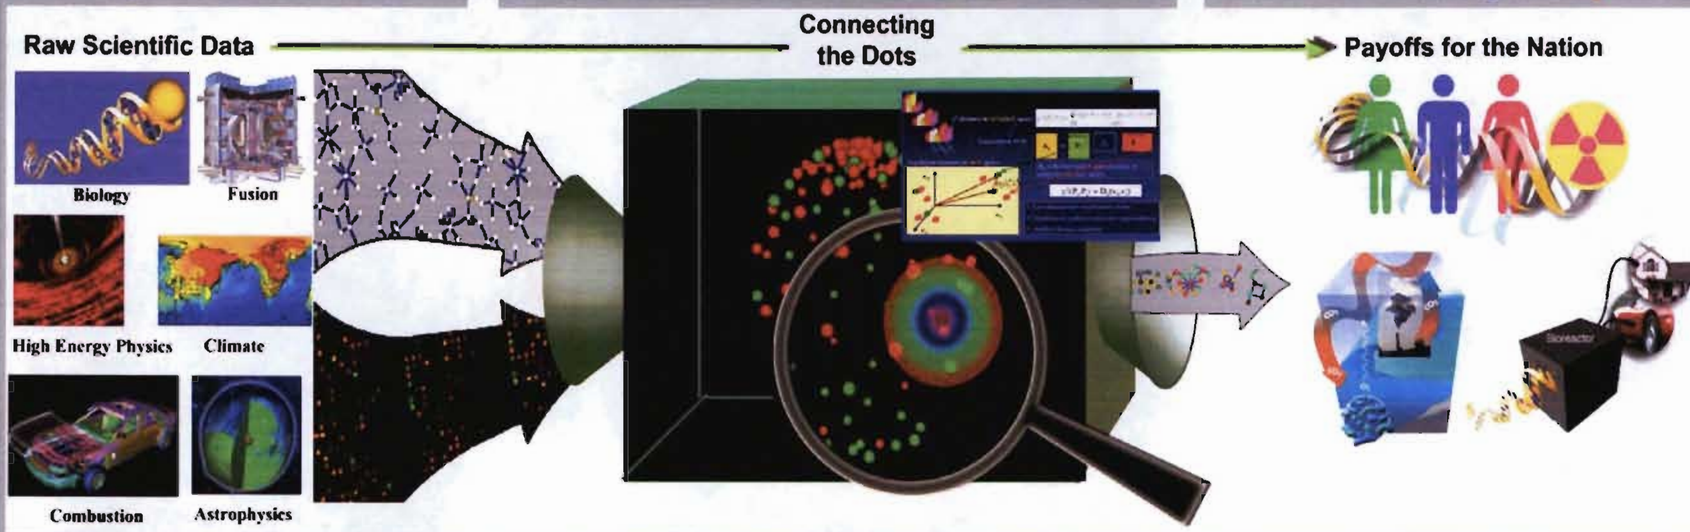

New Worlds of High-Resolution, High-Precision Data

Advanced Mathematics and Algorithms

- Huge dimensional space
- Combinatorial challenge
- Complicated by noisy data
- Requires advanced CS, HPC, theory, modeling, uncertainty quantification, data curation

Providing Predictive Understanding

- Nuclear Security
- Produce hydrogen-based energy
- Stabilize carbon dioxide
- Clean/dispose toxic waste
- Cyber Security

Office of Science
U.S. DEPARTMENT OF ENERGY

Adapted from: Raymond L. Orbach, DOE Undersecretary for Science
2006 AAAS Annual Meeting

IS&T at Los Alamos connects to and integrates many fields and activities across the Laboratory

- **Weapons stockpile stewardship**
  Assessment and integration of simulation and experimental data, experiment design, QMU&UQ, penetrating imaging, weapons data and knowledge ontology and preservation, ...
- **Global security and Intelligence analysis**
  Ubiquitous sensing, fusing, correlating and streaming of images, signals, sensors, and databases; Situational awareness; Cyber security, Cognition, ...
- **Energy security**
  Climate-Energy-Infrastructure; Smart Grids; Nuclear Fuel Cycle...
- **Materials science**
  Matter-Radiation Interactions in Extremes signature facility (MaRIE, ...); Materials-by-design
- **Biothreat response**
  Surveillance of emergent infectious diseases, Bioinformatics, ...
- **Digital libraries**

Research and mission frontiers: Must accelerate synergy and exchange of technical progress for agile, innovative response to current and emerging national security challenges

Applying IS&T to Accelerate Predictive Capability

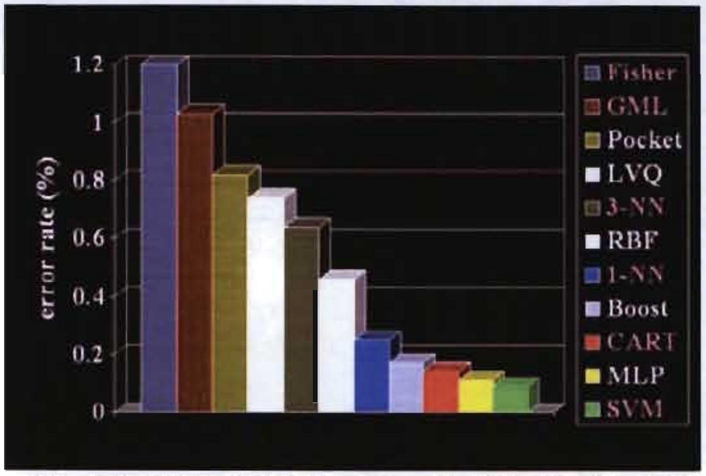

False Alarm rate for cybersecurity applications

IS&T developments are needed to identify "significant" events in real time and extract information-rich data for statistical interrogation, anomaly detection, and classification.

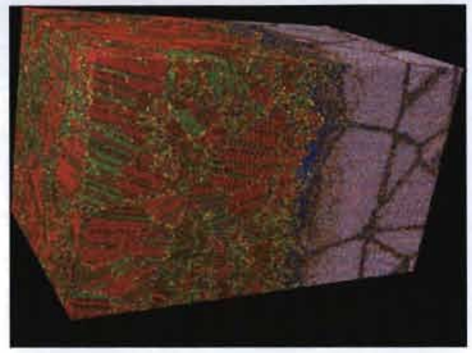

Model evaluation techniques are needed to assess sensitivity, quantify uncertainty, and facilitate development of realistic potential functions, defect distributions, grain orientations, and boundaries.

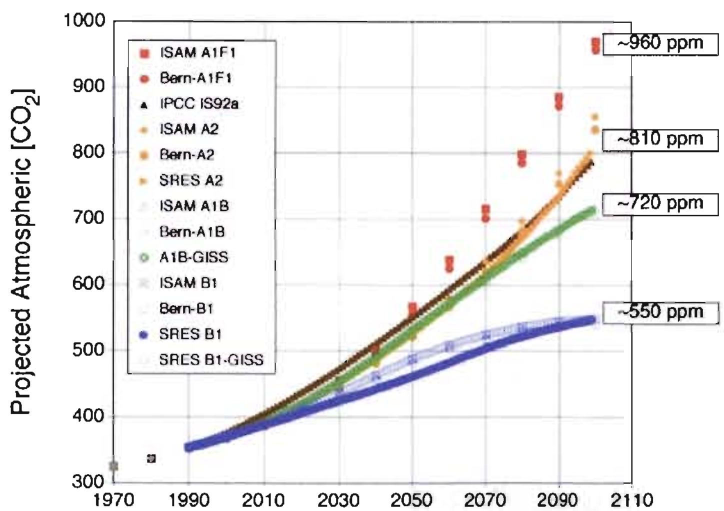

Uncertainty quantification for climate predictions

IS&T provides the framework to efficiently integrate information from experiment and theory, & accelerate model development

Ikonos, 23 May 2000 | Quickbird, 26 March 2006

Change Detection in Images

Many Data-Rich Science Frontiers for Global Security: Intel, Non-proliferation, Space . . .

Space Situational Awareness: Threats from Space

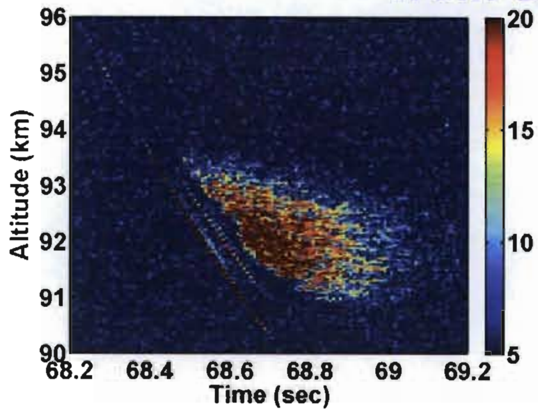

Large interstellar dust

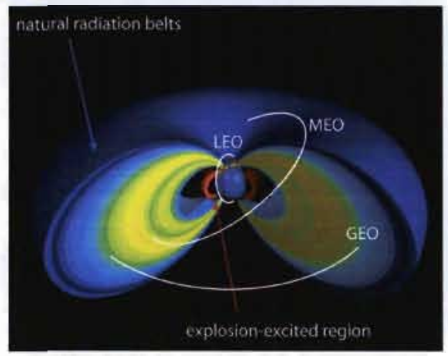

DREAM: Dynamic Radiation Environment Assimilation Model

Nuclear Forensics: Identification & Attribution

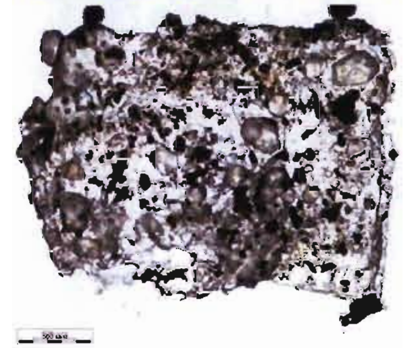

Actinide particle in soil

Electrical and Optical Control of Materials

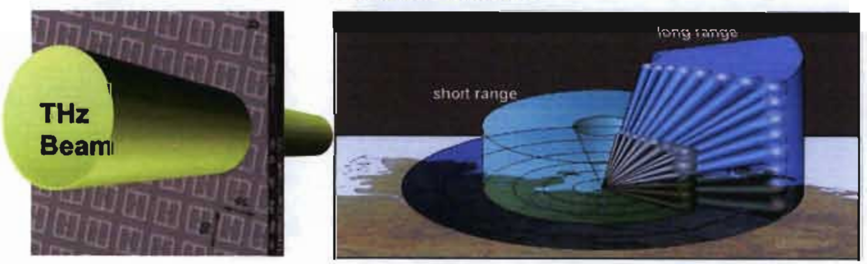

Terahertz metamaterials that modulate

Seismic Detection of Nuclear Explosions

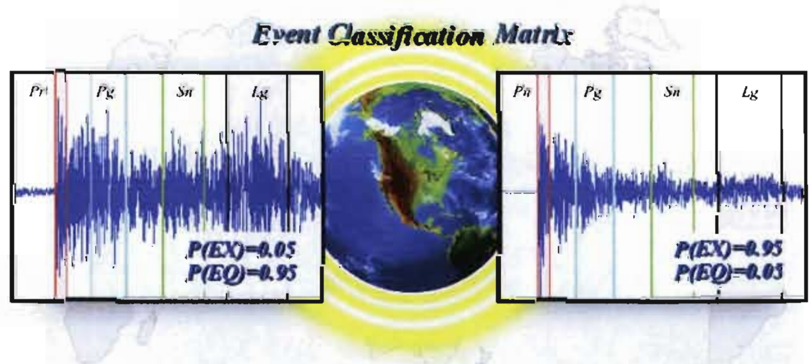

Data Science at Scale: Visualization and analysis of massive (including streaming) data

- **LANL exploring "Middle Ways" between numerically-intensive and data-intensive supercomputing**
  - Need for interactive scientific visualization of massive data quantities
- **Developing novel ways to use emerging computer hardware to enable real-time visualization and analysis of massive streaming datasets**
  - Use active storage and networks
  - Examples: situational awareness, cyber, space, infrastructure, HPC . . .
- **Will enable a system that provides real-time**
  - Processing (correlation) of incoming measurements
  - Analysis of correlated data to identify events of interest, their storage and use

Immersion Visualization (CAVE)

LANL *Data-Intensive Supercomputing* Testbed (G. Grider et al.)

The Problem

A 2-min look at 1/50,000 of the haystack

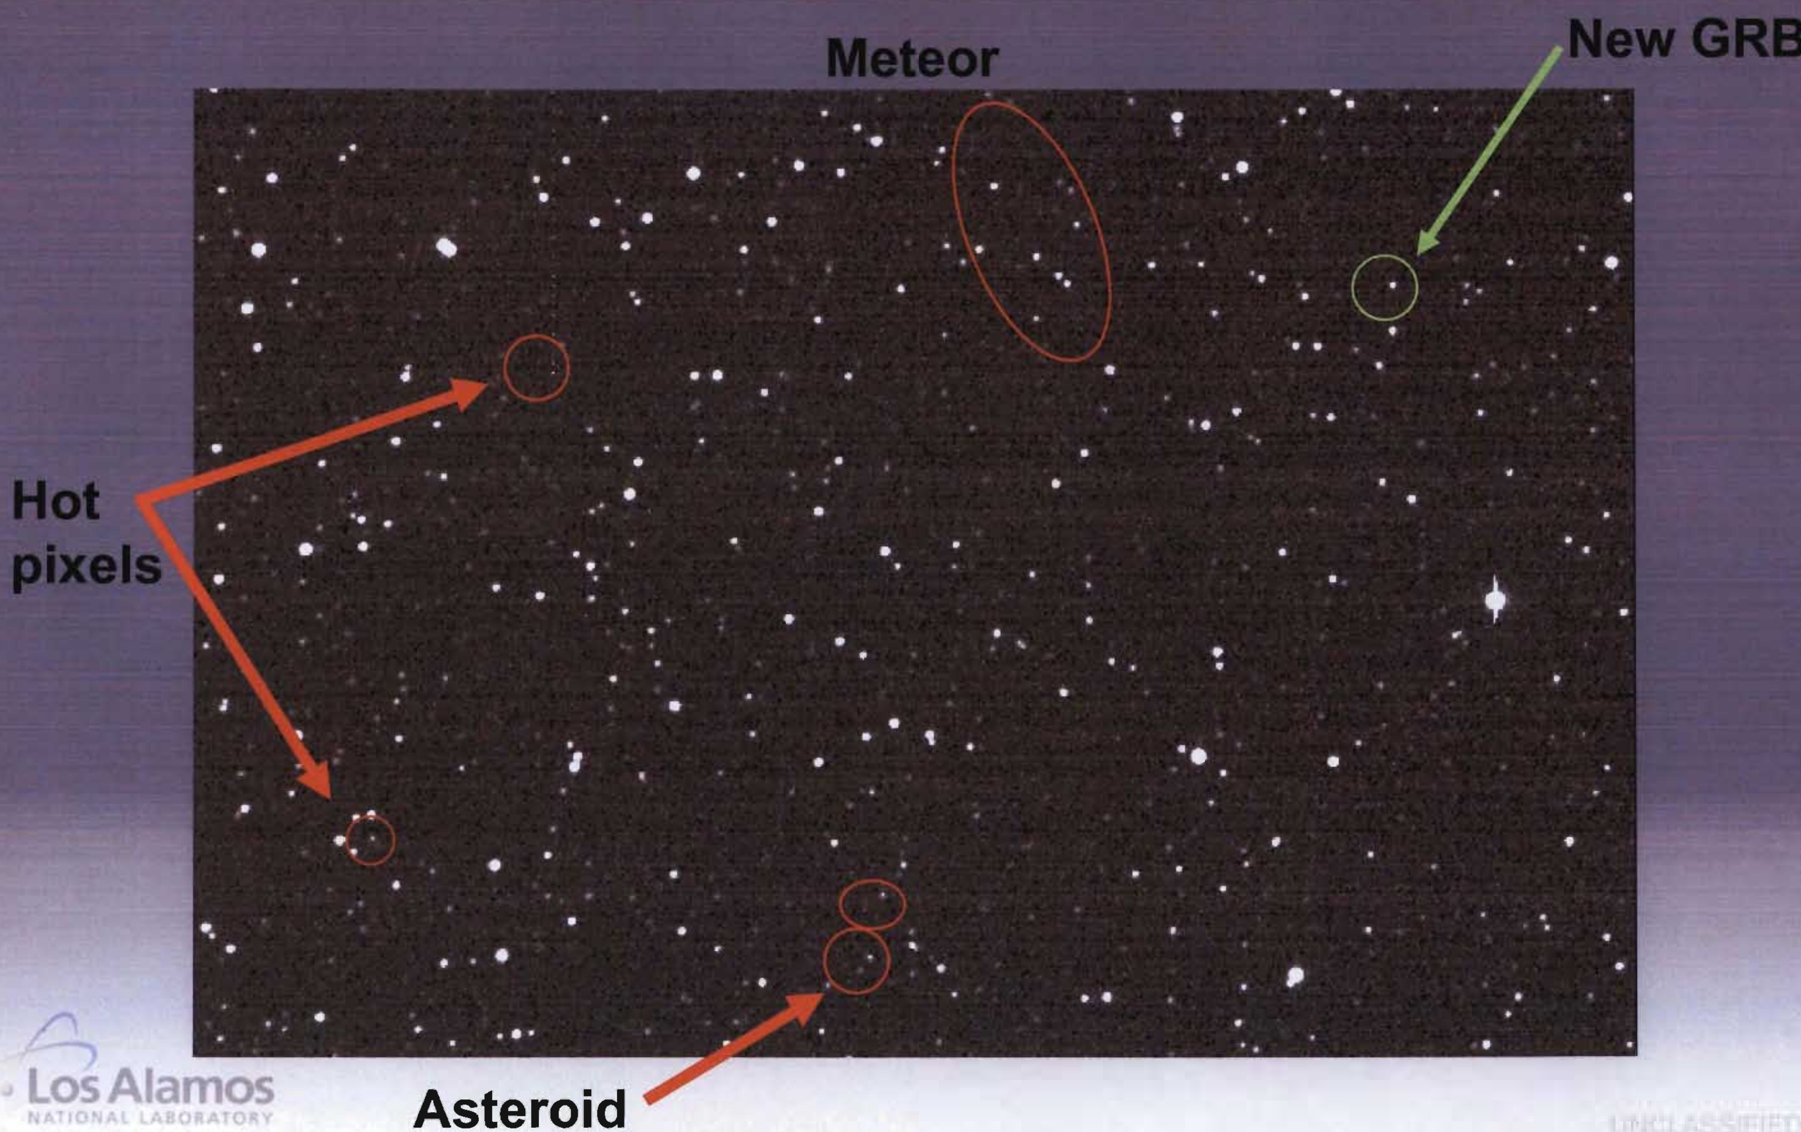

Integrating, interdisciplinary Cyber Science and Technology for real-time distributed network analysis: detecting & mitigating risk

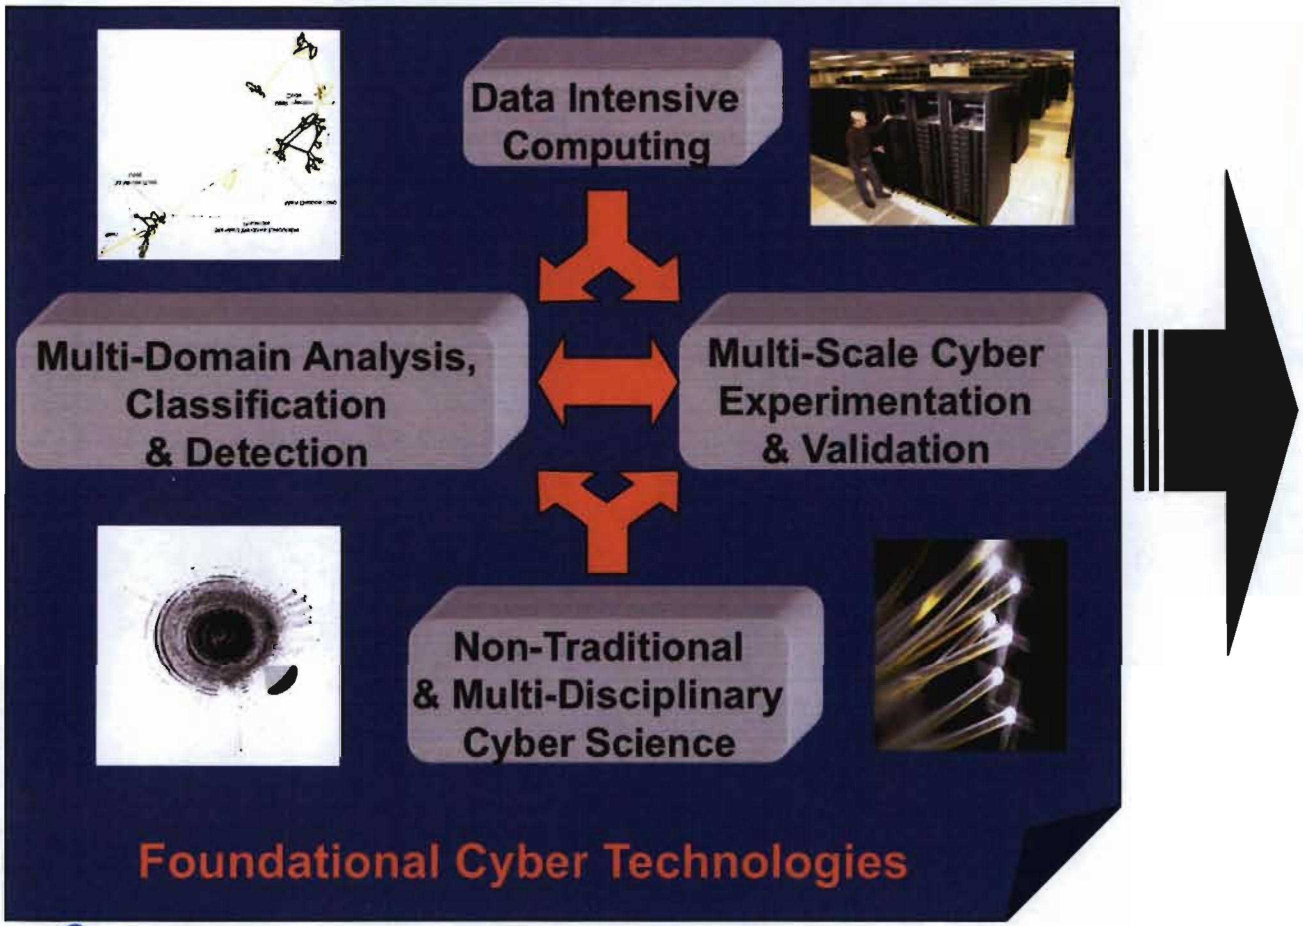

- **Advanced threat analysis, detection, and response tools**
- **Mobile Cyber platform solutions**
- **Cyber operation & decision support frameworks**
- **Experiential Cyber training**

Beyond "Perimeter Defense"

The Future of Materials Science: Control science via integration

Accelerated materials discovery and design

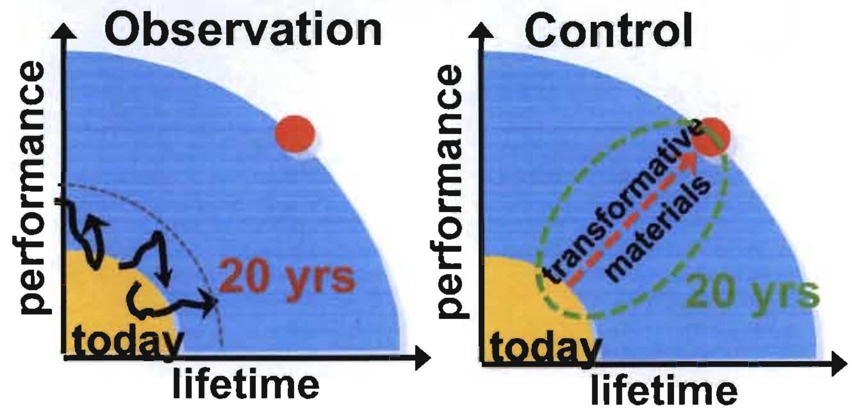

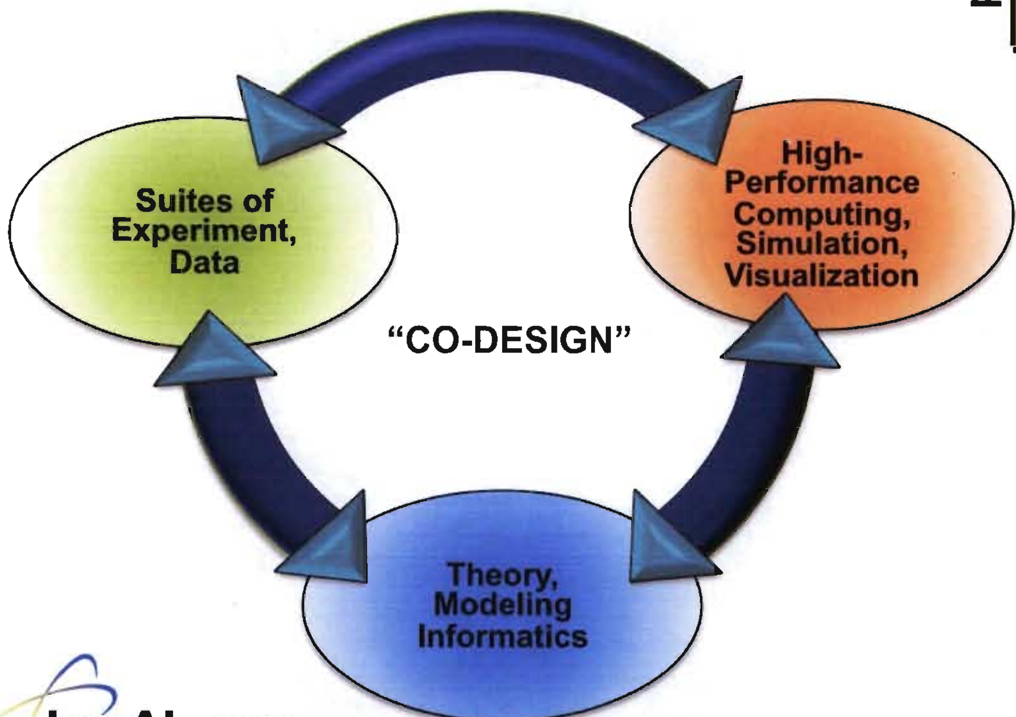

Integration

Key to prediction of material properties

- Theory and models that take function to structure
- Synthetic control of defects and interfaces
- Characterization of the evolution of defects and interfaces in multiple extremes to provide feedback

Exascale Co-Design Center for Materials in Extreme Environments

ExMatEx

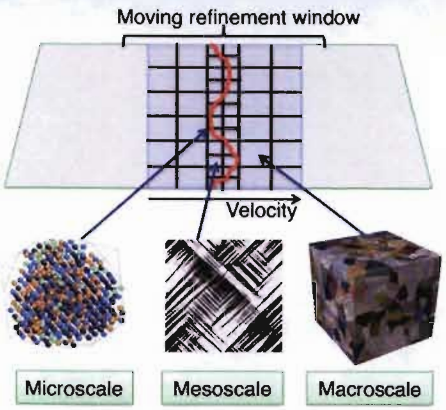

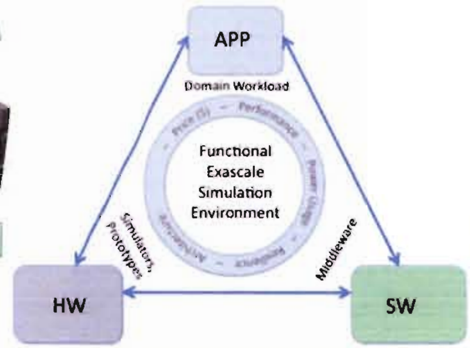

Novel Ideas

Embedded Scale-Bridging Materials Science

Adaptive physics refinement

Asynchronous task-based approach

Agile Development of Proxy Application Suite

Single-scale apps target node-level issues

Scale-bridging apps target system-level issues

Co-optimization for P^{3}R (Price, Performance, Power, and Resiliency)

ASPEN, SST models & simulators

GREMLIN emulator for stress-testing

Impact and Champions

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.

The design and development of extreme environment tolerant advanced materials by manipulating microstructure and interfaces, at the grain scale, depends on such predictive capabilities.

Director: Tim Germann (LANL)

Deputy Director: Jim Belak (LLNL)

Milestones/Dates/Status

Kickoff Workshop	AUG 2011
Initial molecular dynamics (MD) SPMD proxy app(s)	DEC 2011
Initial scale-bridging MPMD proxy app(s)	MAR 2012
Prototype MD DSL	SEP 2012
Assessment of data/resource sharing requirements, both for scale-bridging and *in situ* visualization/analysis	2013
Demonstrate scale-bridging on 10+ PF-class platform	2015

A Decadal Opportunity for Materials Science: Removing key scientific barriers to discovery, prediction, and control

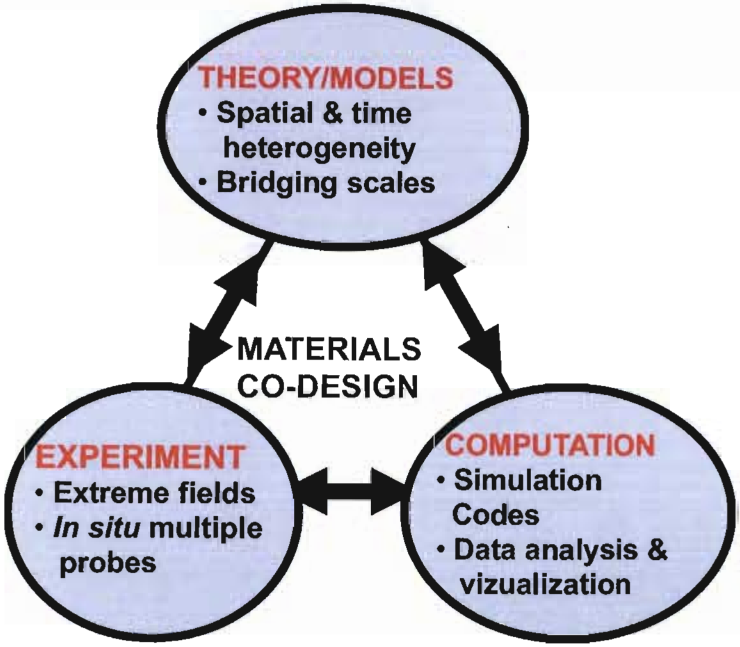

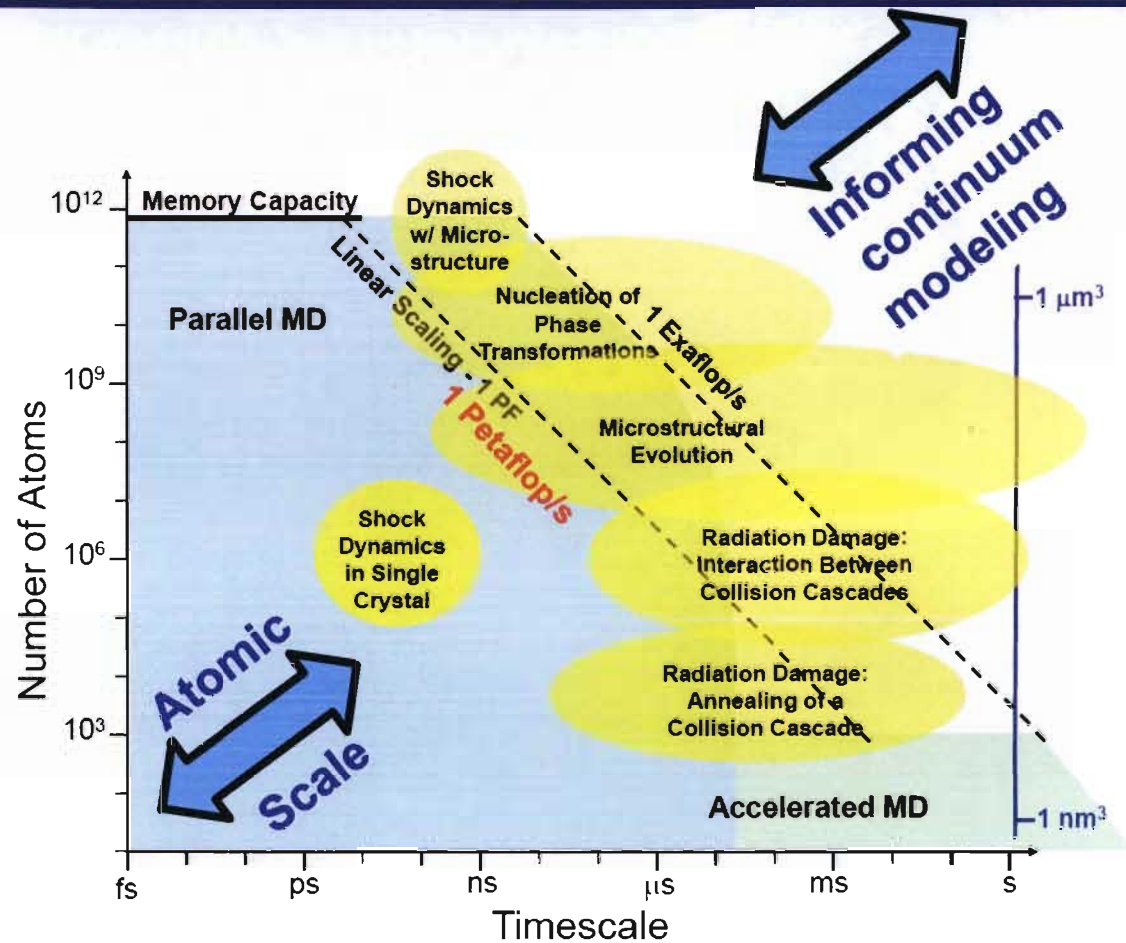

Anticipated advances in modeling, petaflop–exaflop computing, and experimental tools with unprecedented resolution, will allow access to rate-limiting phenomena at the meso (micron) scale

MaRIE: Matter-Radiation Interactions in Extremes
LANL Signature Facility plan (NW, NE, Fission-Fusion…)

Co-design: Seamless integration of domain disciplines to optimally solve advanced, strategic problems of national interest

Experiment
- X-ray
- P-rad
- VISAR
- TEM

Theory
- Multiscale Models
- Nonequilibrium Physics
- Fluctuations
- Disorder/ Inhomogeneities

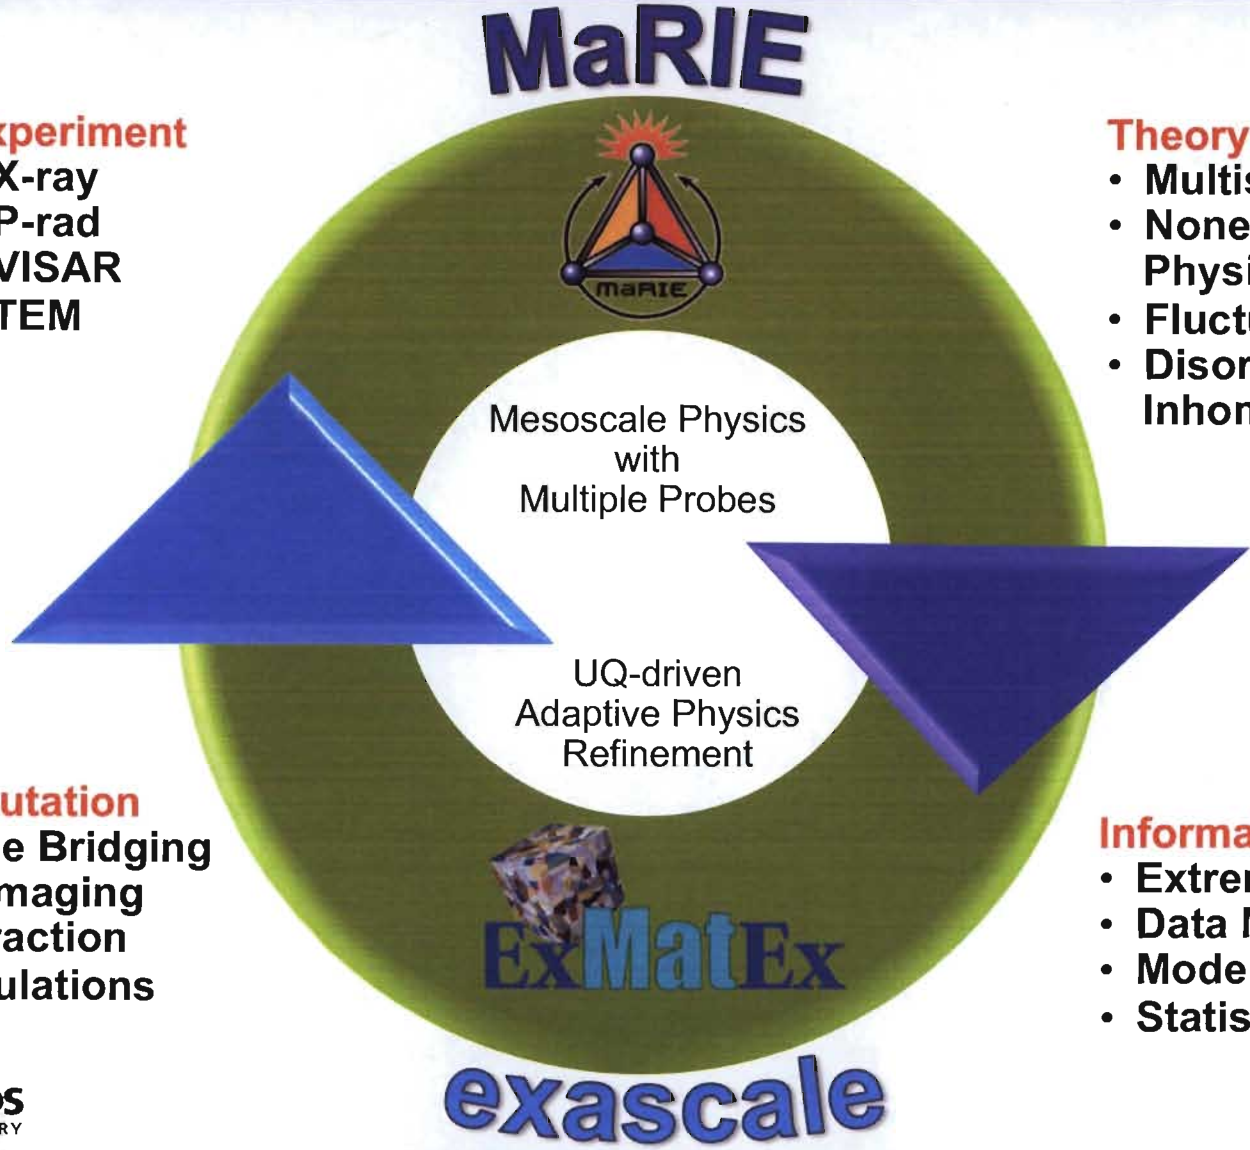

Computation
- Scale Bridging
- 3D Imaging
- Diffraction
- Simulations

Information Science
- Extremal Statistics
- Data Mining
- Model Selection
- Statistical Inference

Codesign: Integrated capabilities to impact mission challenges

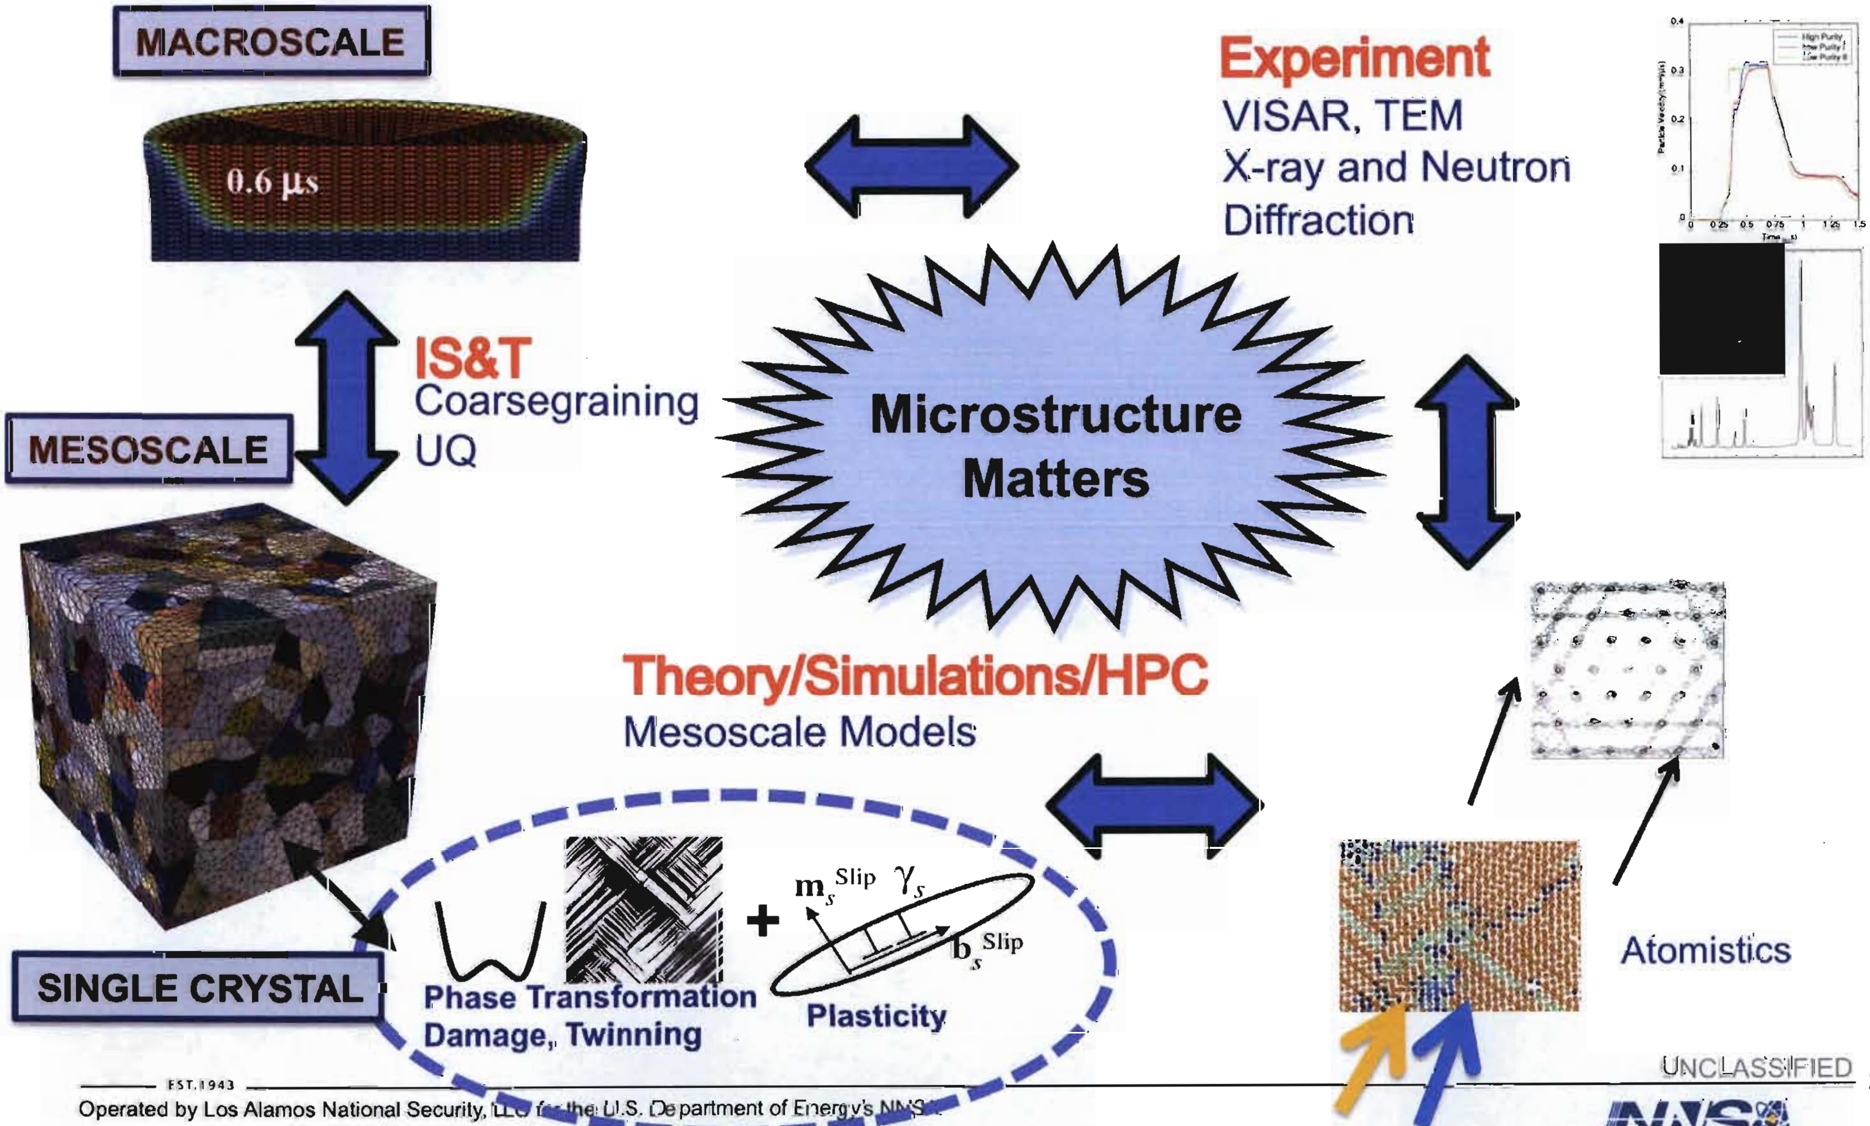

The Consortium for Advanced Simulation of Light Water Reactors
CASL: Unique partnership and the first DOE "HUB"

Building on long-standing, productive relationships and collaborations to forge a close, cohesive, and interdependent team that is fully committed to a well-defined plan of action

Core partners

- Oak Ridge National Laboratory
- Electric Power Research Institute
- Idaho National Laboratory
- Los Alamos National Laboratory
- Massachusetts Institute of Technology
- North Carolina State University
- Sandia National Laboratories
- Tennessee Valley Authority
- University of Michigan
- Westinghouse Electric Company

Individual contributors

- ASCOMP GmbH
- CD-adapco, Inc.
- City University of New York
- Florida State University
- Imperial College London
- Rensselaer Polytechnic Institute
- Southern States Energy Board
- Texas A&M University
- University of Florida
- University of Tennessee
- University of Wisconsin
- Worcester Polytechnic Institute

Using Roadrunner cell architecture for viral phylogenetics: Inheritance vs. Adaptation? HIV vaccine design?

Seeking correlation patterns of HIV viruses that are immunologically potent; infer "family tree" of >4,000 HIV sequences (10^4 nucleotides each)

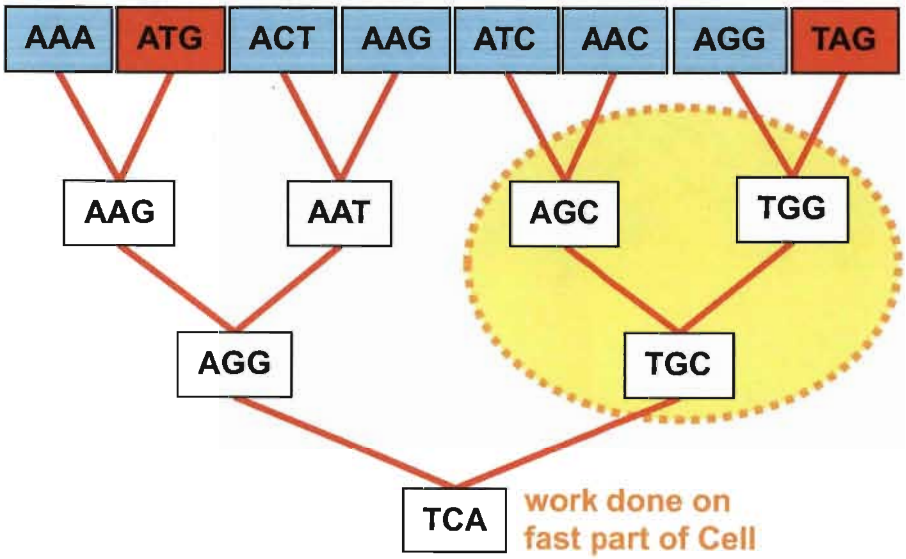

>10^5 CPU hours on Roadrunner
Huge # matrix inversions

HIV Vaccine Implications?

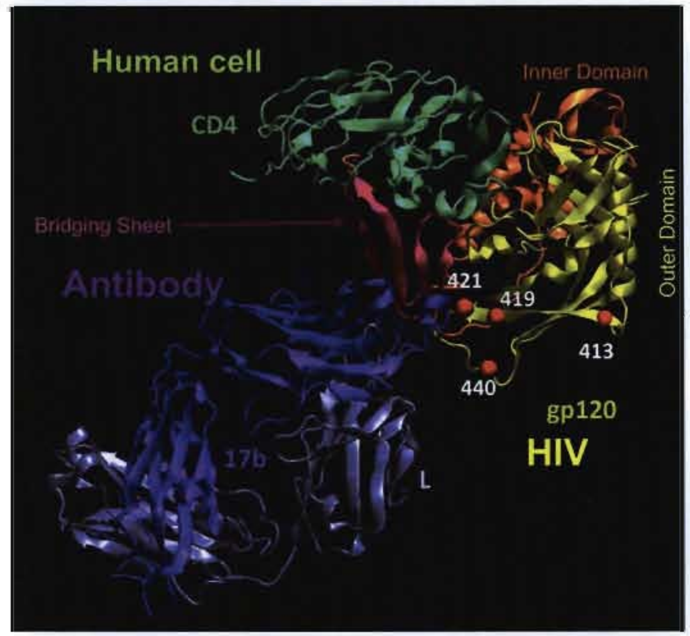

- **CD4i region of viral envelope important for good immune response (inducing antibodies)**
- **Experiments underway to verify**

BILL & MELINDA GATES foundation

CHAVI
CENTER FOR HIV/AIDS VACCINE IMMUNOLOGY

(T. Bhattacharya et al.)

Coupled Earth System Model

- **DOE SciDAC effort to help create CCSM model with complete carbon, sulfur and nitrogen cycles**
- **Community Climate System Model**
  - Fully interactive physical models of atmosphere, ocean, land, sea ice
  - Collaboration with CCSM biogeochemistry working group
- **Coupled chemistry/biogeochemistry**
  - Atmosphere, ocean and land exchanging CO_2 fluxes
  - Atmosphere: 90-100 chemical compounds and over 200 interactions
  - Ocean: 26 species and tracers using Doney-Moore-Lindsay ecosystem (plankton, nutrients, detritus…) and LANL trace gas model including dimethyl sulfide
  - Land: many carbon pools

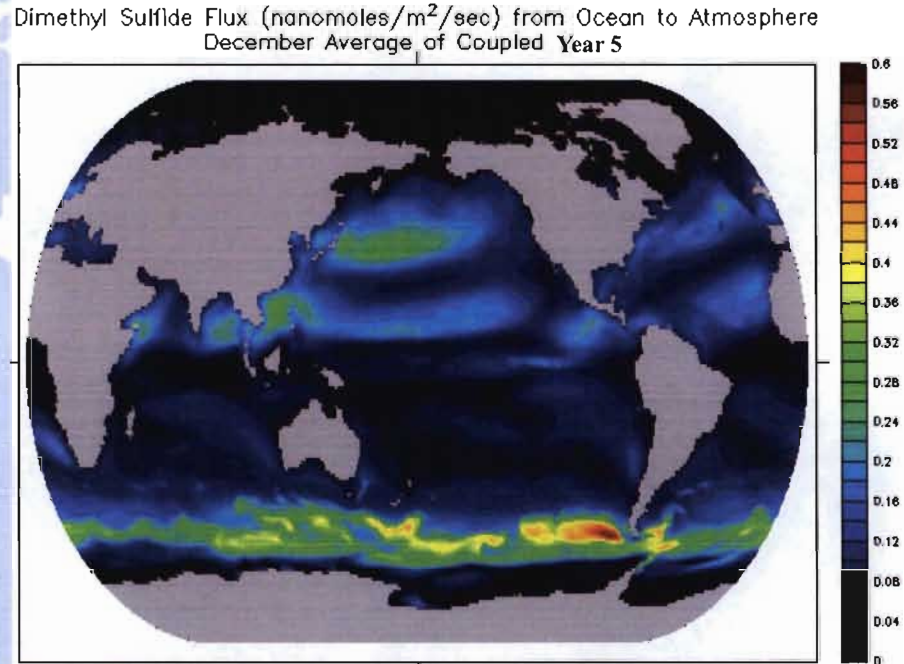

Ocean Dimethyl Sulfide Flux: important natural source of aerosols

Co-Design for Energy and Climate Science: Bringing key assets together to address Impacts and Mitigations

Measurements & Analysis (Microbes to Satellites!)

Climate and Modeling

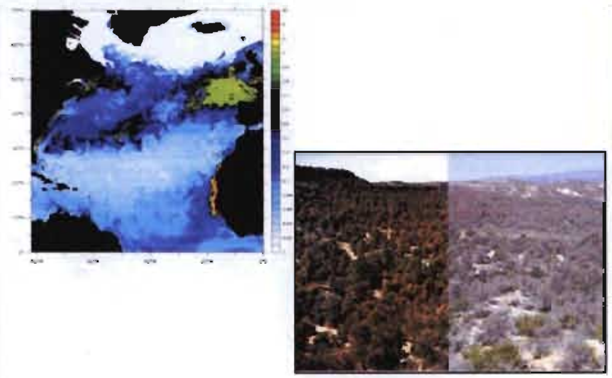

Uncertainty Quantification

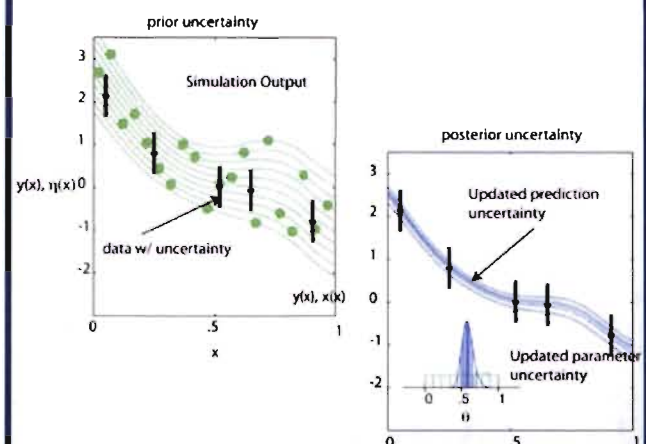

Social, Energy, and Infrastructure Modeling

ASCEM National Laboratory Consortium

The *Advanced Simulation Capability for Environmental Management* (ASCEM) program

- Developing a state-of-the-art scientific tool and approach for understanding and predicting contaminant fate and transport in natural and engineered systems
- Will facilitate standardized development of performance and risk assessments for EM clean-up/closure activities
- Multiple Labs: LANL, LBNL, PNNL, ORNL, SRNL (LANL, ANL, INL)
- Bridging Communities: Applied Mathematics, Computational Scientists, Geoscientists
- Building on advances from ASC, ASCR Applied Math, and SciDAC

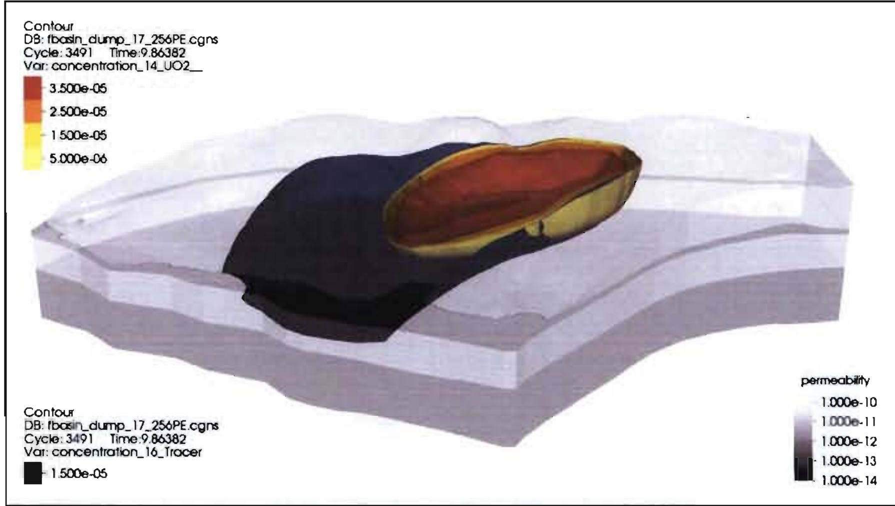

Iso-surfaces of the Uranium plume (yellow and red) and a non-reactive tracer (blue) are shown for a simplified model of the F-Area seepage basins at Savannah River. Results computed with the new parallel open-source ASCEM Multi-Process HPC simulator, Amanzi.

Science-Based Prediction of Natural System Performance
requires system-level probabilities based on process level phenomena

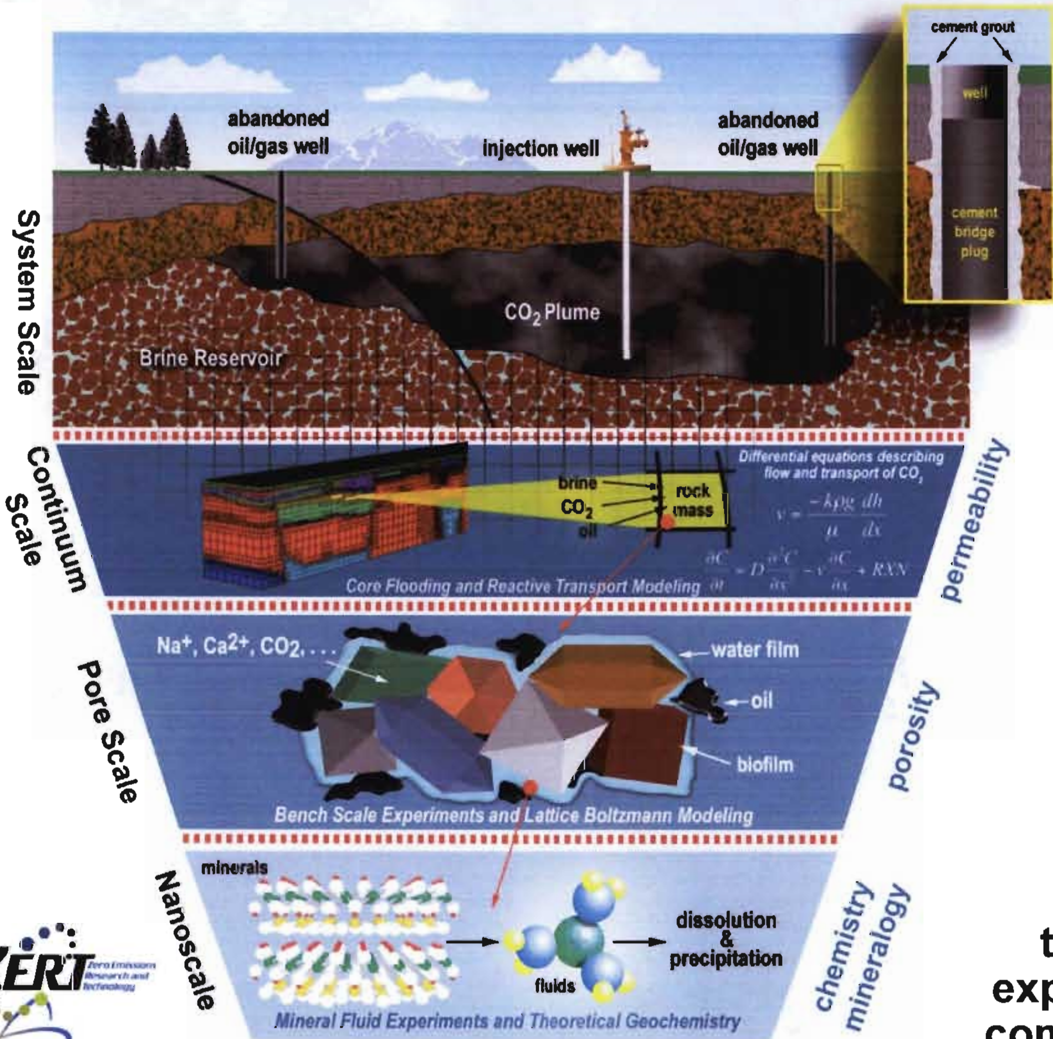

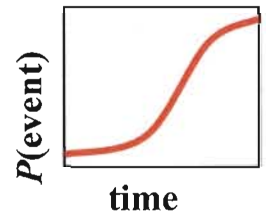

system model (probabilistic)

theory, experiment, computation

observation (analog sites)

Smart Grid: LANL Science and Programs

R&D Problems for Smart Grids

A future grid—in which modern sensors, communication links, and computational power are used to improve efficiency, stability, and flexibility—has become known as the "smart grid."

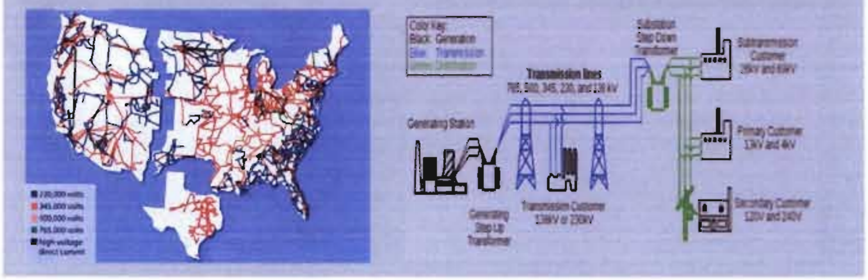

R&D Methodology: Road Map for Smart Grids

Driven by emerging technologies (e.g., renewables, storage, and meters), specifies the technical challenges in *Grid Planning*, *Grid Control*, and *Grid Stability* and requires scientific advances in

- Analysis & Control
- Scalability/Reliability Mosaics
- State Estimation
- Data Aggregation & Assimilation
- Middleware for the Grid
- Modeling Consumer Response

Grid Planning

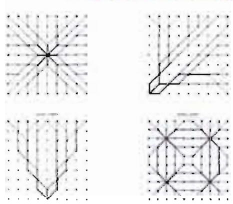

Bent, Bercheid, Toole; 09-11, *Generation and Transmission Extension Planning*

Johnson, Chertkov; 10-11, *Network Optimization Approach to Planning*

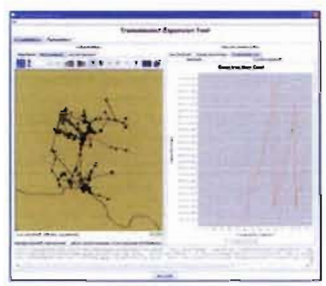

Grid Stability

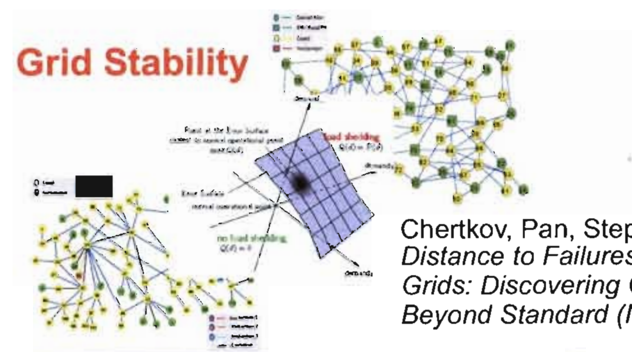

Chertkov, Pan, Stepanov ; 10-11, *Distance to Failures in Transmission Grids: Discovering Contingency Beyond Standard (N-1 criterium)*

Grid Control

Turitsyn, Sinitsyn, Backhaus, Chertkov; 10-11, *Control of Electric Vehicle Queuing*

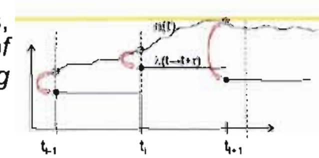

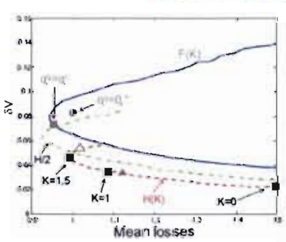

Sulc, Turitsyn, Backhaus, Chertkov; 10-11, *Reactive Control of Photovoltaic-Rich Distribution System*

LANL contributing to the White House's Advanced Manufacturing Partnership

"What parent wouldn't want indestructible military-grade diapers?"

President Barack Obama
June 24, 2011

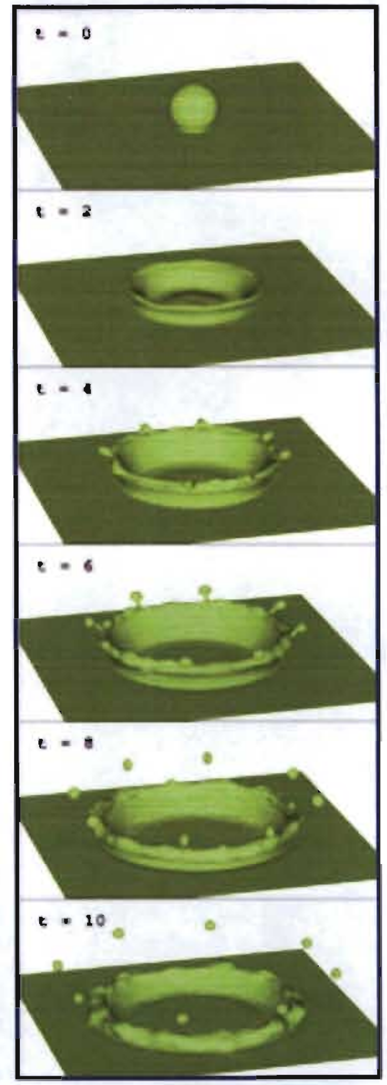

For nearly 20 years, Procter & Gamble and Los Alamos National Laboratory have collaborated to incorporate computational and modeling technologies developed for national security into cutting-edge tools for manufacturing

Supplementary

Strategic Direction: LANL leads the Exascale Co-Design Center for Materials in Extreme Environments (ExMatEx)

- Co-design is a new paradigm in which the exascale hardware, system software, and application codes are concurrently designed to create the exascale simulation environment.
- The goal is to achieve more realistic large-scale simulations of materials in extreme mechanical and radiation environments.
- A predictive understanding of the response of materials to extreme conditions underpins DOE and NNSA missions and Laboratory programs.

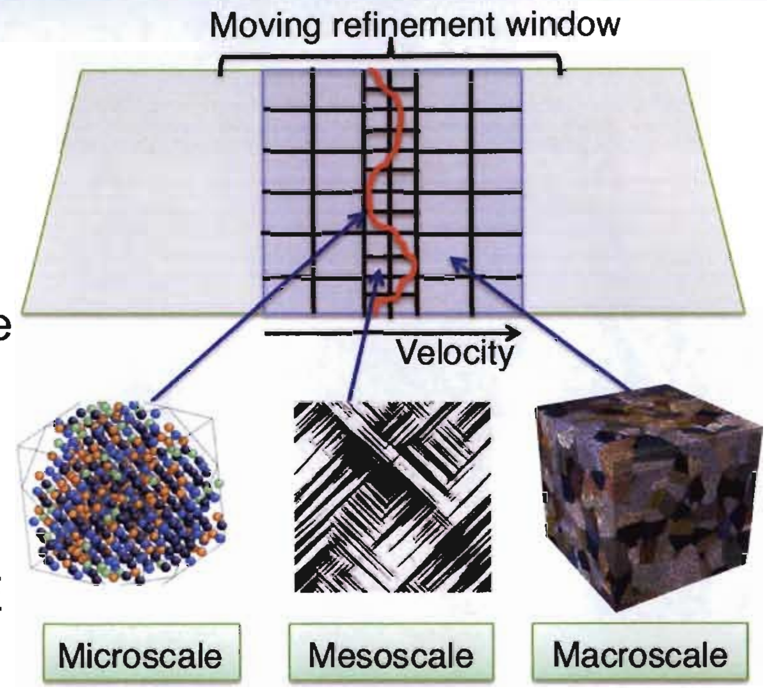

Los Alamos NATIONAL LABORATORY EST. 1943

Sandia National Laboratories

OAK RIDGE National Laboratory

ExMatEx

Massachusetts Institute of Technology

California Institute of Technology 1891

PRINCETON UNIVERSITY

Lawrence Livermore National Laboratory

CRAY THE SUPERCOMPUTER COMPANY

intel

STANFORD UNIVERSITY

Understanding the role of microstructure-based heterogeneity evolution in material damage

Example: Plutonium and stockpile metals strength and damage

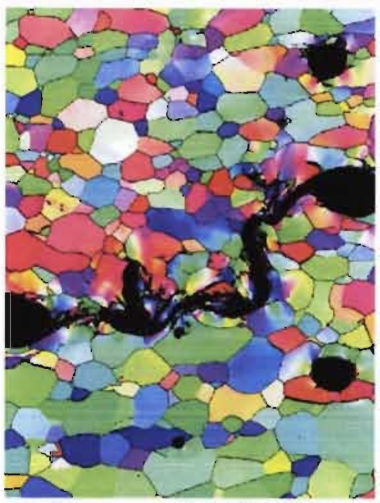

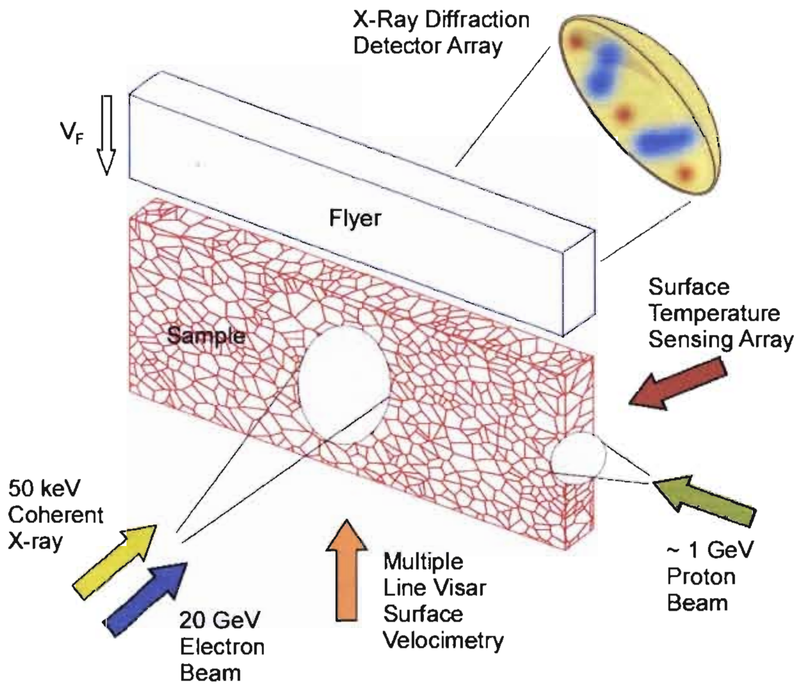

Sub-µm resolution
100s–1000s µm samples

Sub-ns resolution
~30 frames in 1-µs duration

+

The goal
Predict dynamic microstructure and damage evolution

The first experiment
Multiple, simultaneous dynamic *in situ* diagnostics with resolution at the scale of nucleation sites (< 1 µm; ps – ns)

The model
Accurate sub-grain models of microstructure evolution coupled to molecular dynamics

Team includes Curt Bronkhorst et al. (LANL, UK AWE, BYU, CalTech, Ohio State, …)

LANL always an early adopter of transformational High-Performance Computing technology

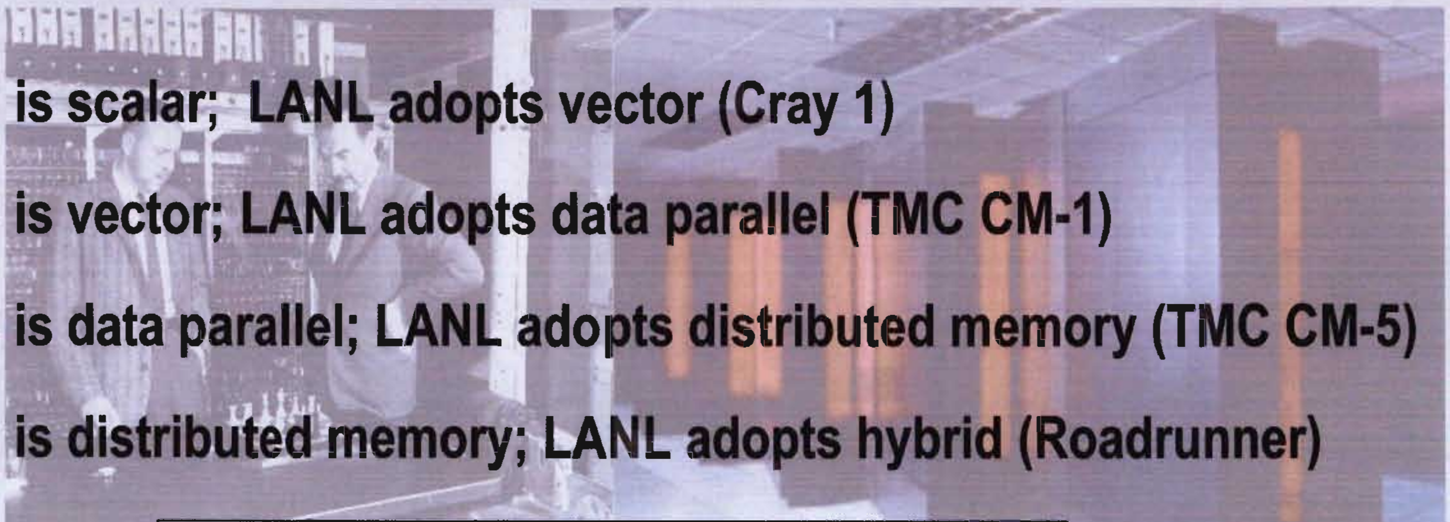

1970s: HPC is scalar; LANL adopts vector (Cray 1)

1980s: HPC is vector; LANL adopts data parallel (TMC CM-1)

1990s: HPC is data parallel; LANL adopts distributed memory (TMC CM-5)

2000s: HPC is distributed memory; LANL adopts hybrid (Roadrunner)

The decision to commit to new hybrid technology was formed by our application and platform experience

Roadrunner embodies many key architectural trends, each in moderation

Roadrunner has multicore, short-vector SIMD, threads, heterogeneous instruction sets, local stores instead of caches, on-chip CPU/memory networks, remote accelerators, and cluster computing.

- **You can use any of these features as needed, without needing to go to extremes in any one of them.**
- **_Roadrunner_'s scale and flexibility makes it an ideal base from which to explore the changing landscape of HPC**
- **And it provides immediate benefits!**

Open Science Projects

Lessons-Learned

Roadrunner architecture is cell-accelerated

- Huge speed-up of some dominant compute elements
- But very new code, I/O, memory strategies needed

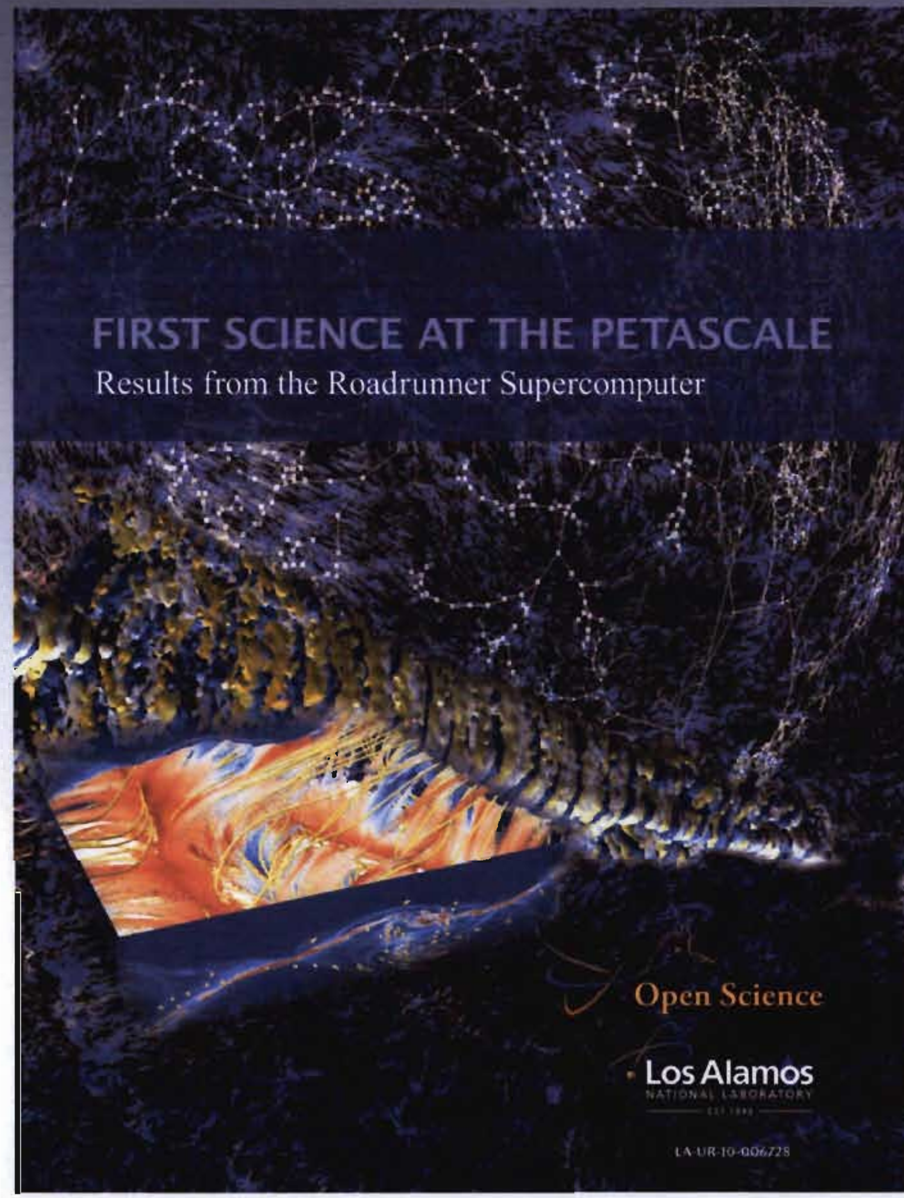

Los Alamos National Laboratory EST. 1943 | Lawrence Livermore National Laboratory | Oak Ridge National Laboratory | Sandia National Laboratories | Stanford University | California Institute of Technology

Exascale Co-Design Center for Materials in Extreme Environments

ExMatEx

Timothy C. Germann

Exascale Research PI Meeting
Annapolis, MD
13 October 2011

A predictive understanding of the response of materials to extreme conditions (mechanical and/or irradiation) underpins many DOE missions.

ADVANCED SIMULATION & COMPUTING™

CMIME
Center for Materials at Irradiation and Mechanical Extremes

CASL

"With the advent of exascale computing, the possibility exists to achieve predictive capabilities to manipulate microstructure and interfaces, at the grain scale, to enable the design and development of extreme environment tolerant advanced materials."
– *Scientific Grand Challenges for National Security* report

Computational materials science involves a hierarchy of length and time scales

Ab-initio Methods	Molecular Dynamics	Phase-Field Modeling	Continuum Methods
Inter-atomic force model, equation of state	Defect and interface mobility, nucleation	Direct numerical simulation of multi-phase evolution	Multi-phase material response, experimental observables
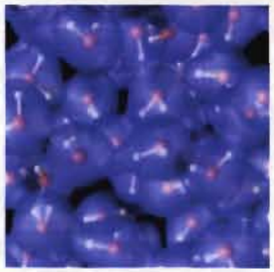			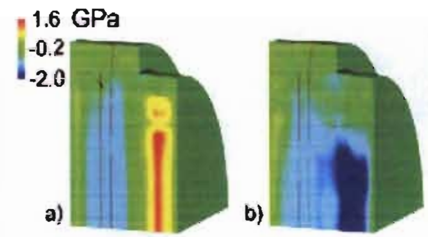
Length/time: nm, ps Codes: Qbox/LATTE Motif: Particles and wavefunctions, plane wave DFT with nonlocal norm-conserving, ScaLAPACK, BLACS, and custom parallel 3D FFTs Prog. Model: MPI	Length/time: μm, ns Codes: SPaSM/ddcMD Motif: Particles, domain decomposition, explicit time integration, neighbor and linked lists, dynamic load balancing, parity error recovery, and *in situ* visualization Prog. Model: MPI + Threads	Length/time: 100 μm, μs Codes: AMPE/GL Motif: Regular and adaptive grids, implicit time integration, real-space and spectral methods, complex order parameter (phase, crystal, species) Prog. Model: MPI	Length/time: cm, ms Codes: VP-FFT/ALE3d Motif: Regular and irregular grids, implicit time integration, 3D FFTs, polycrystal and single crystal plasticity, Prog. Model: MPI

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

ExMatEx Co-Design Project Goals

- 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.

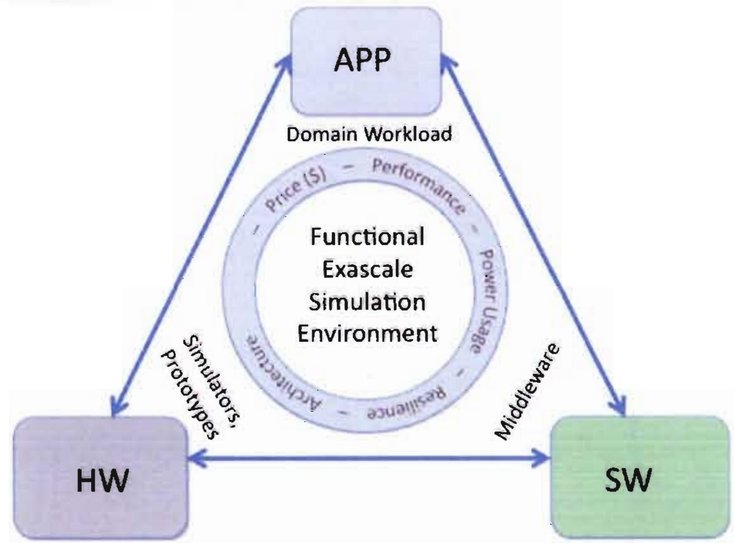

- This task-based approach leverages the extensive concurrency and heterogeneity expected at exascale while enabling fault tolerance within applications.
- The programming models and approaches developed to achieve this will be broadly applicable to a variety of multiscale, multiphysics applications,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

Executive Advisory Board

Alan Bishop (LANL)
Tomás Díaz de la Rubia (LLNL)
Rick Stevens (ANL)
Kathy Yelick (LBNL)
Steve Zinkle (ORNL)

Exascale Co-Design Center for Materials in Extreme Environments

Center Director: Tim Germann (LANL)
Deputy Director: Jim Belak (LLNL)

SC/ASCR

Exascale Ecosystem

Advanced Algorithms & Co-design "Code-Team"
Lead: David Richards (LLNL) Erik Draeger (LLNL), Jamal Mohd-Yusof (LANL), Danny Perez (LANL)

Computer Science

Lead: Allen McPherson (LANL)
Co-lead: Scott Futral (LLNL)

Programming Models
Lead: Allen McPherson (LANL)
Pat Hanrahan (Stanford)
David Jefferson (LLNL)

Data/Resource Sharing
Lead: Jim Ahrens (LANL)
Chris Sewell (LANL)

Analysis Tools At Scale
Lead: Martin Schulz (LLNL)

Performance Modeling
Lead: Jeff Vetter (ORNL)
Jim Ang (SNL)
Arun Rodrigues (SNL)

Software Stack Engagement
Jim Ahrens (LANL)
Martin Schulz (LLNL)

Vendor Engagement
Matt Leininger (LLNL)
Pat McCormick (LANL)

Applied Math

Lead: Milo Dorr (LLNL)
Co-lead: Dana Knoll (LANL)

Scale-Bridging Algorithms
Lead: Dana Knoll (LANL)
Frank Alexander (LANL)
Milo Dorr (LLNL)
Jean-Luc Fattebert (LLNL)

V&V+UQ
Lead: Houman Owhadi (CalTech)
Clint Scovel (LANL)

Computational Materials Science

Lead: Turab Lookman (LANL)

High Strain-Rate Applications
Lead: Turab Lookman (LANL)
Frank Addessio (LANL)
Nathan Barton (LLNL)
Curt Bronkhorst (LANL)
Ricardo Lebensohn (LANL)
Michael Ortiz (CalTech)

Irradiation Applications
Lead: Roger Stoller (ORNL)
Yuri Osetskiy (ORNL)
Art Voter (LANL)

Agile proxy application development

- Petascale single-scale SPMD and scale-bridging MPMD proxy apps will be used to explore algorithm and programming model design space with domain experts, hardware architects and system software developers.
- Proxy apps communicate the application workload to the hardware architects and system software developers, and are used in models/simulators/emulators to assess performance, power, and resiliency.
- These proxy applications will not be "toy models", but will realistically encapsulate the workload, data flow and mathematical algorithms of the full applications.

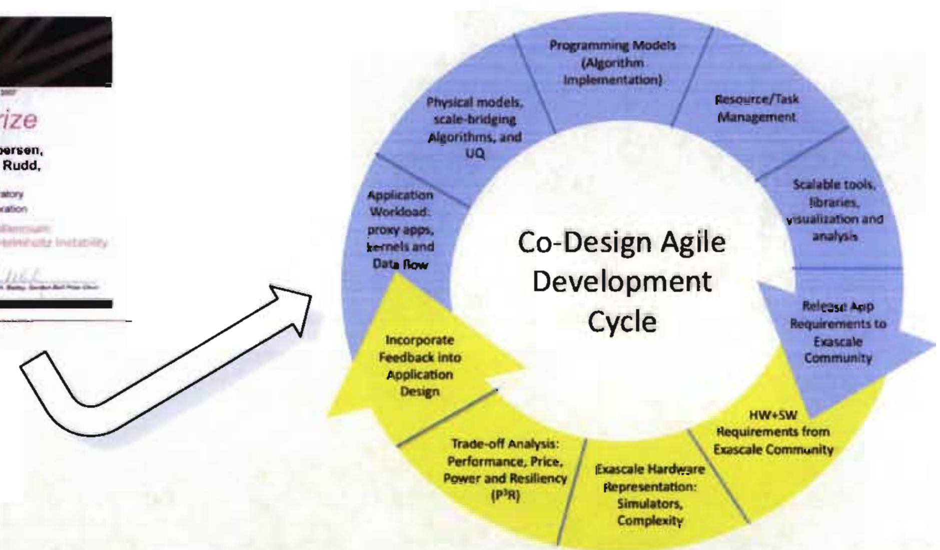

Agile proxy application development

- Proxy apps for single-scale SPMD applications (e.g. molecular dynamics) will be used to assess node-level issues including:
  - *Data structures*
  - *Hierarchical memory storage and access*
  - *Power management strategies*
  - *Node-level performance*
- The asynchronous task-based MPMD scale-bridging proxy apps will be used to optimize:
  - *System-level data movement*
  - *Resilience (fault management)*
  - *Load balancing techniques*
  - *Performance scalability*
- These proxy apps are **not** static entities, but the central mechanism for our co-design process.
- Application, software, and hardware communities analyze and respond to trade-offs with new requirements and capabilities, both from and to the application.

Embedded Scale-Bridging Algorithms

- Our goal is to introduce more detailed physics into computational materials science applications in a way which:
  - escapes the traditional synchronous SPMD paradigm
  - exploits the heterogeneity and concurrency expected in exascale hardware
  - improves data locality
  - reduces I/O burden
- 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 heterogeneous, hierarchical MPMD task-based approach naturally maps to anticipated heterogeneous, hierarchical architectures and concurrency and resiliency issues.

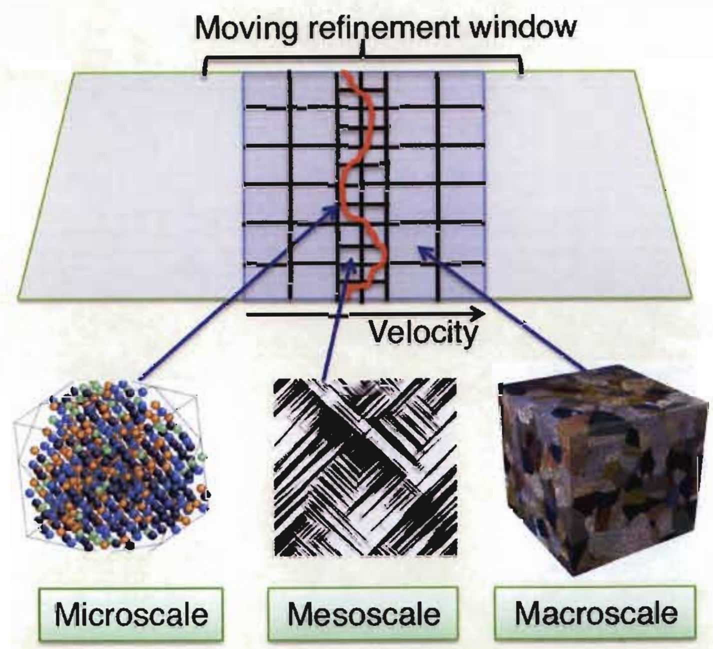

Embedded Scale-Bridging Algorithms

- Scale-bridging algorithms require a consistent two-way algorithmic coupling between temporally evolving distinct spatial levels; they are not "modeling", and not one-way information flow.
- Our focus is on coupling between macro (coarse-scale model) and meso (fine-scale model) scales with all unit physics being deterministic.
- We begin by building off of our adaptive sampling success, but move to the use of temporally evolving mesoscale and spatial adaption.
- Similar concepts apply in the time domain, e.g. using *ab initio* techniques to compute activation energies for a rate theory or kinetic Monte Carlo model ("on-the-fly kMC") applied to radiation damage modeling.

Uncertainty quantification

- On-the-fly spawning of finer scale models (adaptive sampling already has some uncertainty information)
- In P^3R optimization process, to prioritize measurements on leadership-class platforms

Application friendly programming models

- A hierarchy of programming models exposes and exploits the architectural heterogeneity, while providing a transparent layer of abstraction that insulates the application programmer from the continuous flux and complexity of the underlying hardware.
- We implement abstractions…
  - *…on top of the technologies proven by the proxies*
    - » Abstractions "compile" to underlying technology
  - *…as libraries, run-time systems, domain-specific languages*
    - » Domain science abstractions, e.g. meshes, particles, halo communication, etc.
    - » System-level abstractions, e.g. interoperability (resource sharing) for scale-bridging, in situ viz, tools, scheduling, load-balancing, fault tolerance, etc.
- Well defined abstractions enable the 3 P's
  - *Portable…as the domain*
  - *Productive…more so than low-level APIs (e.g. Cuda, pthreads, etc.)*
  - *Performance…optimization enabled by domain knowledge*
- The hierarchical programming models and approaches developed to achieve our scale-bridging materials application will be broadly applicable to a variety of multiscale, multiphysics applications, and will escape the traditional bulk synchronous parallel paradigm.

Hierarchical programming models

- This hierarchy will replace the traditional bulk synchronous parallel paradigm:
  - *On-node task parallelism* will allow us to couple multiple tightly coupled application components or segments while exploiting on-node resources to their full extent.
  - *Inter-node cooperative parallelism* will provide the necessary capabilities to execute scalable, dynamically structured MPMD applications.
  - *Domain specific languages* aim to encapsulate these levels, enable programmer productivity, and bridge disparate architectures.

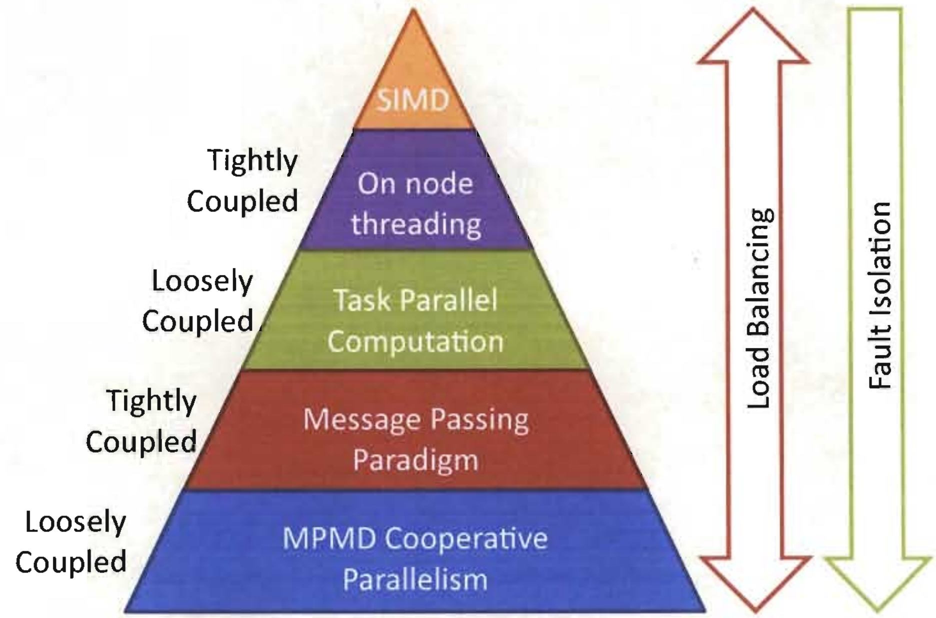

- This task-based MPMD approach leverages the concurrency and heterogeneity at exascale, while enabling novel data models, power management, and fault tolerance strategies.

Holistic analysis and optimization

- A hierarchy of performance models, simulators, and emulators are used to explore algorithm, programming model, and hardware design space and perform trade-off analysis between competing requirements and capabilities in a tightly coupled optimization loop.
- Node- to system-level models and simulators
  - ➢ *ASPEN: Rapid exploration of design space using application skeletons*
  - ➢ *SST: Detailed simulation of data flow, performance and energy/power cost*
- Exascale emulation layer (GREMLIN) to introduce perturbations similar to those expected on future architectures
- Performance analysis on hardware prototypes and leadership-class machines

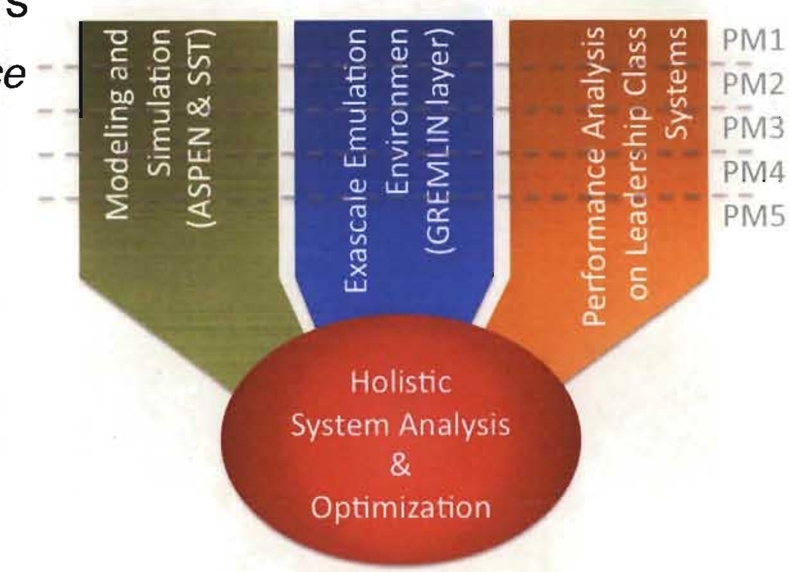

- Trade-off analysis leads to a co-optimization of algorithms and architectures for price, performance, power (chiefly memory and data movement), and resilience (P^3R)

Collaboration with the exascale ecosystem

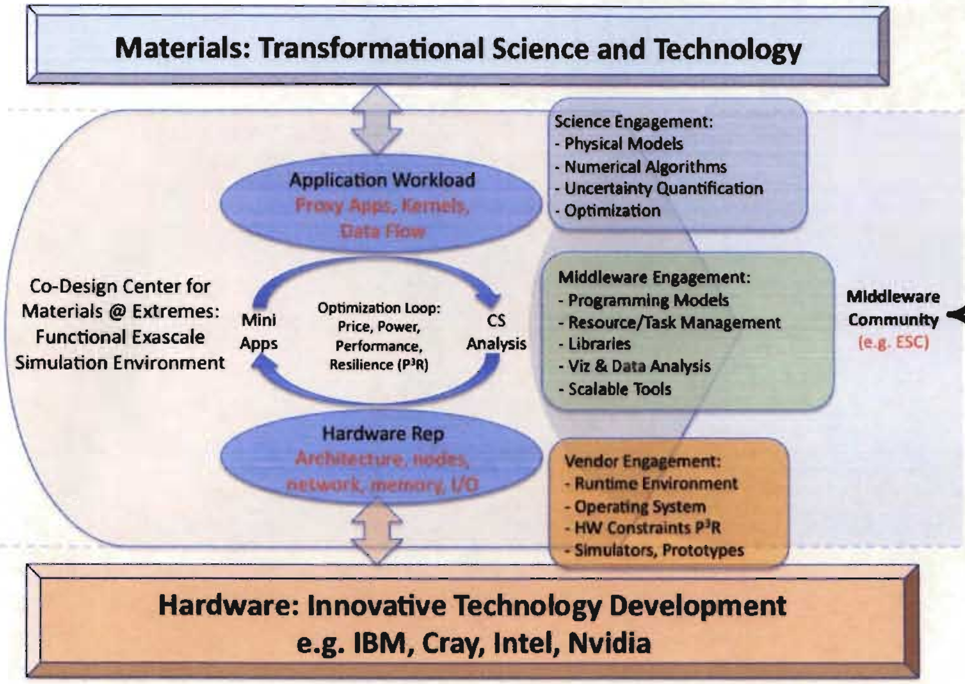

Scientific Data Management & Analysis

- In situ visualization and analysis techniques exploiting task based approach
- Data analysis and visualization techniques for multi-scale, multi-physics simulations

X-stack software

- Task based programming models, frameworks, compilers, and scalable tools
- Runtime systems that enable resource allocation, sharing, and management
- Co-design of fault tolerance and power management strategies

Advanced Architectures

- Mapping single-scale SPMD and scale-bridging MPMD proxy applications to evaluate performance and cost
- Co-design of fault tolerance and power management strategies

These are a few examples from several potential areas of synergy with ASCR supported exascale computer science research

Summary

- Our objective is to establish the interrelationship between algorithms, system software, and hardware required to develop a multiphysics exascale simulation framework for modeling materials subjected to extreme mechanical and radiation environments.
- This effort is focused in four areas:
  - *Scale-bridging algorithms*
    - » UQ-driven adaptive physics refinement
  - *Programming models*
    - » Task-based MPMD approaches to leverage concurrency and heterogeneity at exascale while enabling fault tolerance
  - *Proxy applications*
    - » Communicate the application workload to the hardware architects and system software developers, and used in performance models/simulators/emulators
  - *Co-design analysis and optimization*
    - » Optimization of algorithms and architectures for performance, memory and data movement, power, and resiliency