Addressing National Challenge Problems with Exascale Applications
Application Development Plans in the Exascale Computing Project (ECP)

Douglas B. Kothe
Deputy Associate Laboratory Director
Computing and Computational Sciences Directorate
Oak Ridge National Laboratory

Presentation to
Argonne Training Program on Extreme-Scale Computing
Aug 1, 2016

EXASCALE COMPUTING PROJECT

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science

ECP WBS

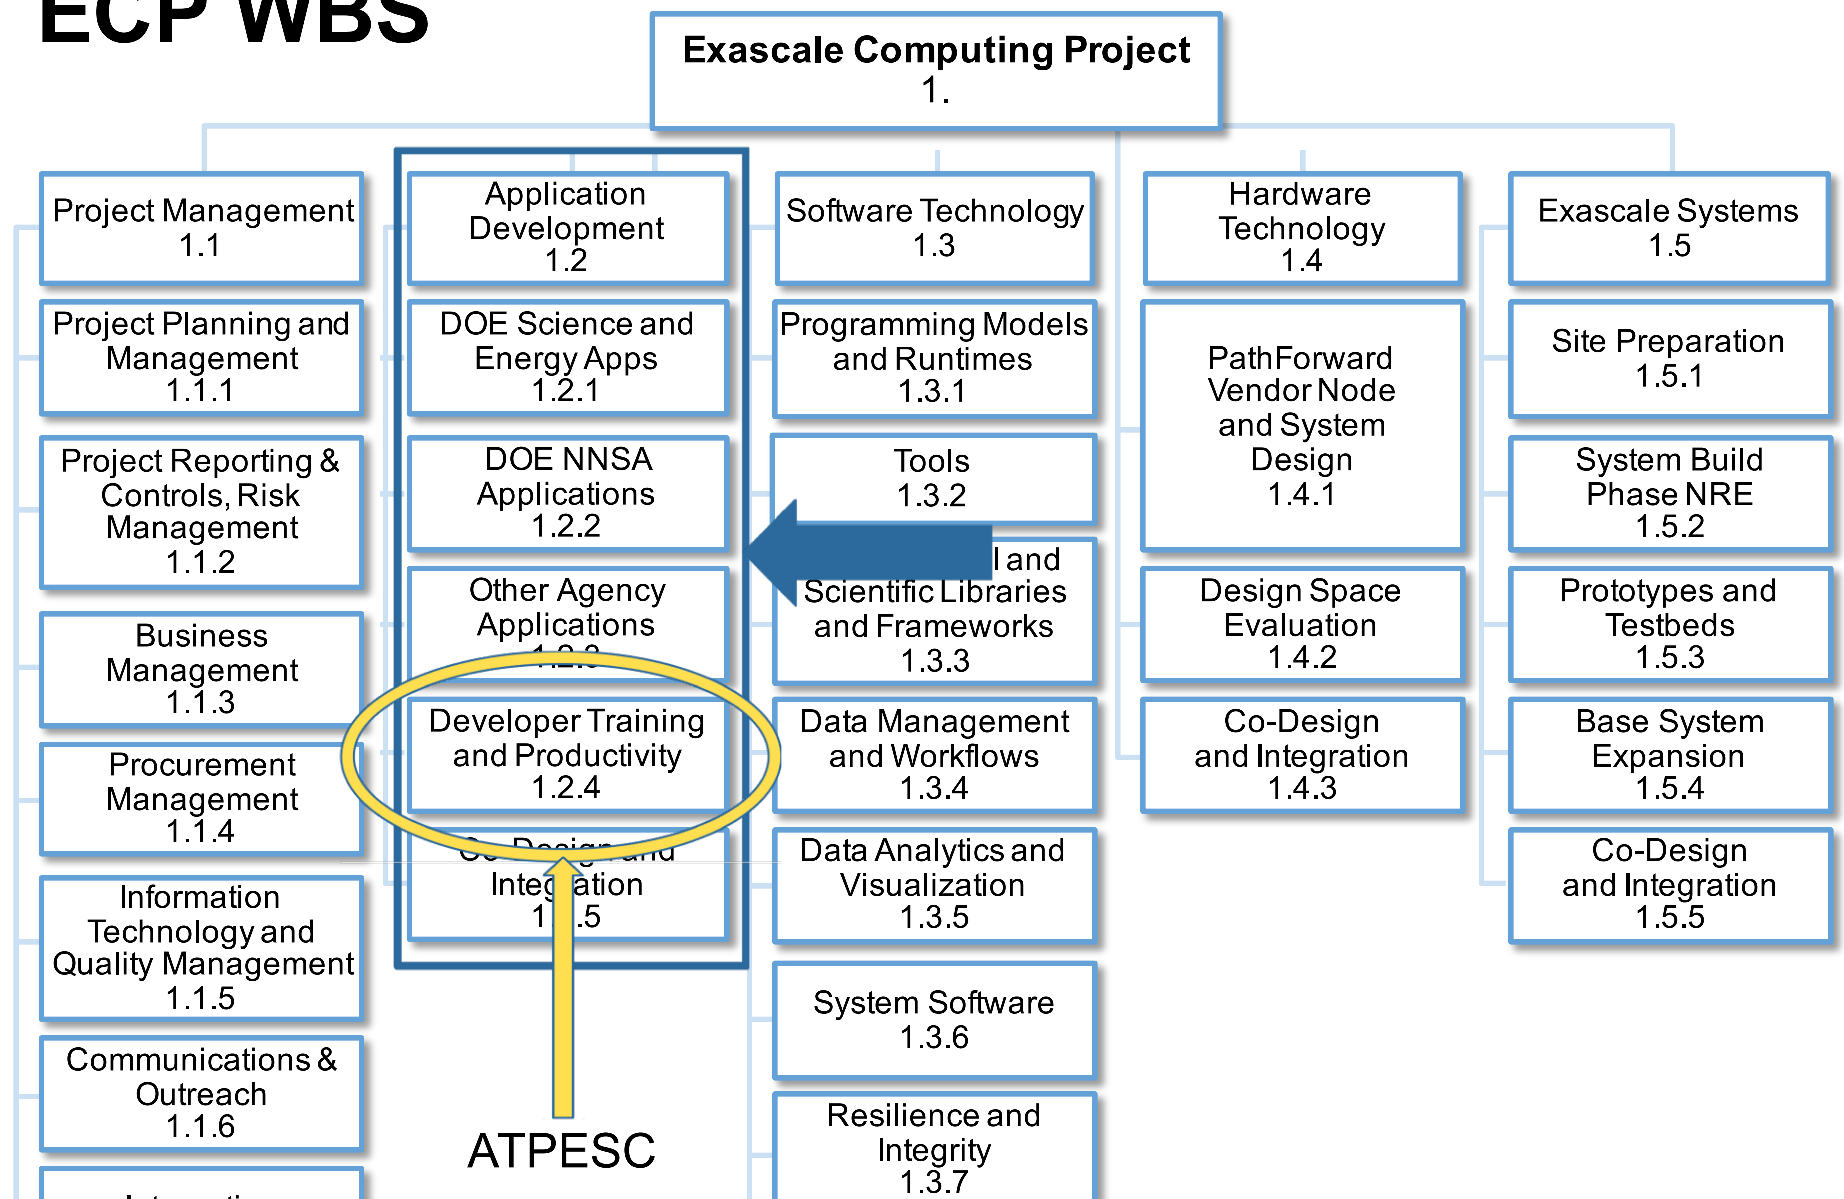

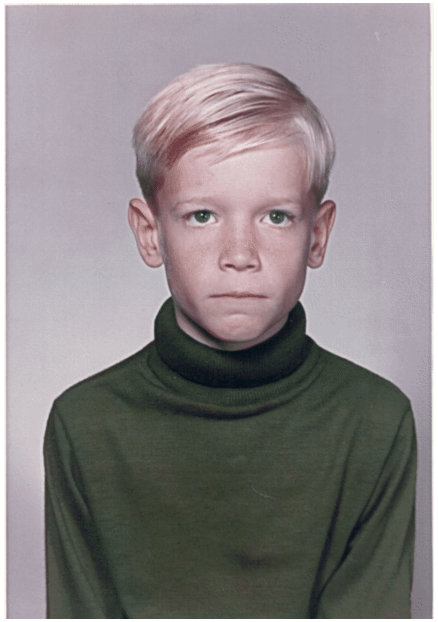

Before Photo: Me prior to agreeing to lead the ECP Application Development focus area

Credits: The “Cult of the Code”

- FLIP-PHD
  - 2D imploding inertial confinement fusion (ICF) targets (DOE)
- NASA-VOF2D, NASA-VOF3D
  - 2D/3D micro-gravity free surface flows (NASA micro-gravity)
- RIPPLE
  - 2D free surface flows (NASA micro-gravity; Xerox inkjet, …)
- CFDLIB
  - 2D/3D multiphase flows in fluidized catalytic crackers (Exxon); diaper making
- PAGOSA
  - Armor/anti-armor program (Army)
- POP
  - Global ocean circulation
- TRUCHAS/TELLURIDE
  - Casting/welding processes; spray forming; coastal hydrodynamics; Corporate Lethality Program (MDA)
- VERA (Virtual Environment for Reactor Applications)
  - Nuclear reactors
- Others !TBN

Imploding ICF Target Design

Fluid-Implicit-Particle (FLIP) particle-in-cell (PIC) method, now a base technology for Disney Animation Studios!

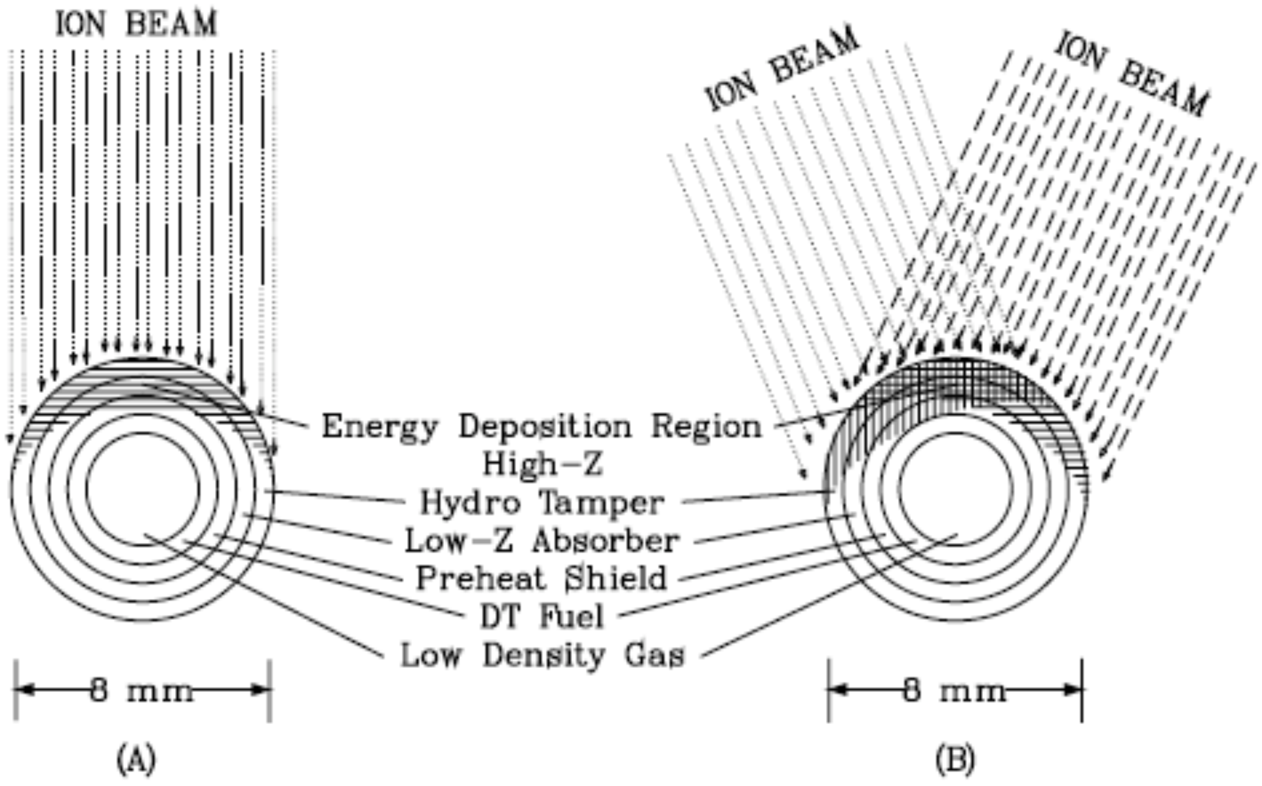

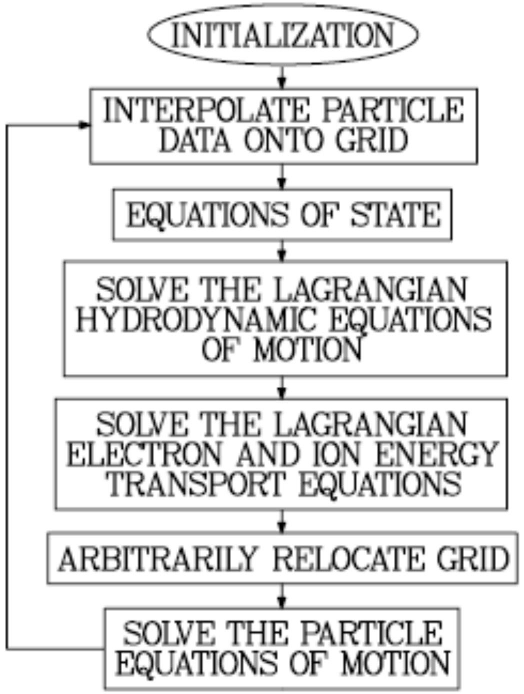

Unique characteristics of numerical method

- 2nd order time/space PIC method (FLIP) with inherent stability and fluid-like collisionality
- Adaptive grid (*not* AMR – it was the pre-AMR days)
- Discrete ray-tracing for ion beam penetration & energy deposition
- Natural ability to track interfaces via particle identity
- Innate sensitivity to unstable hydrodynamics (particle/grid Eulerian/Lagrangian duality)

Imploding ICF Target Design

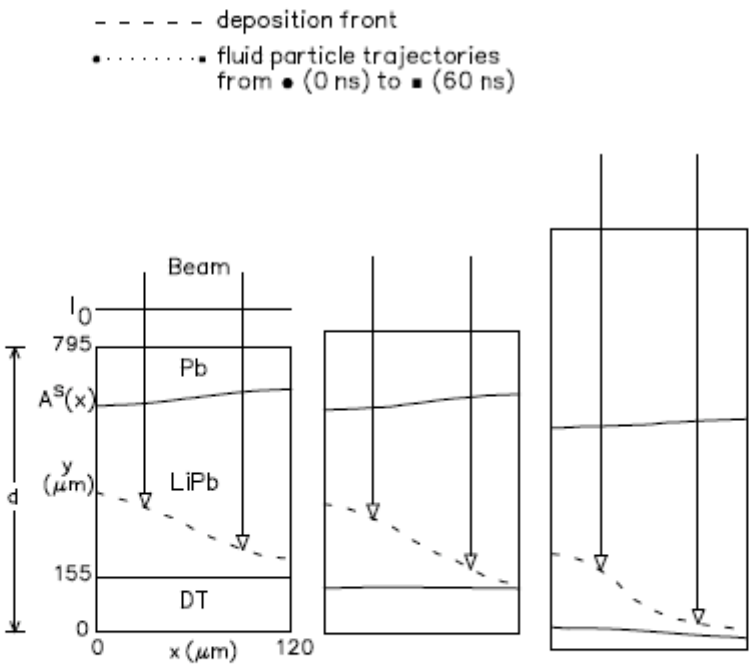

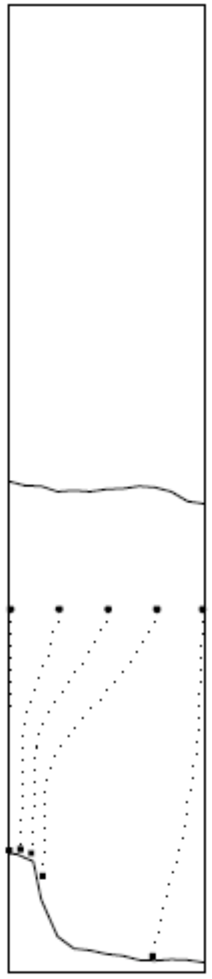

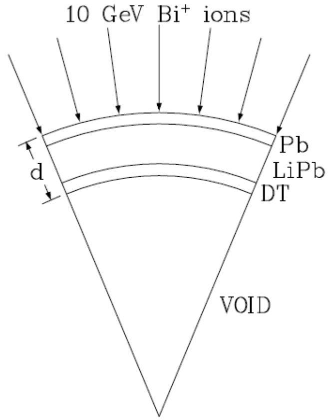

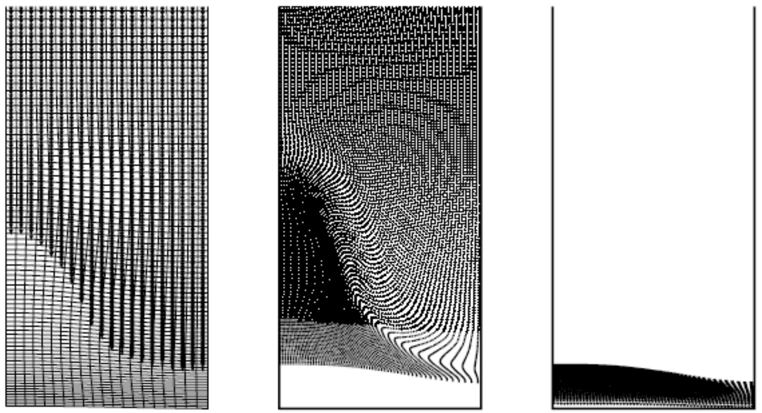

(A) (B) (C)

Characteristics of Implementation

- Fortran 77 (some inlined CAL)
- Linked lists for particles
- Kershaw's ICCG
- A memory/CPU hog
- CDC 7600, Cray XMP

Material Point Methods

Ref: Joseph Teran (UCLA) et al.

Snow

Material Point Methods

Ref: Joseph Teran (UCLA) et al.

Heat transfer/phase change

A New Model for Surface Tension ("CSF")

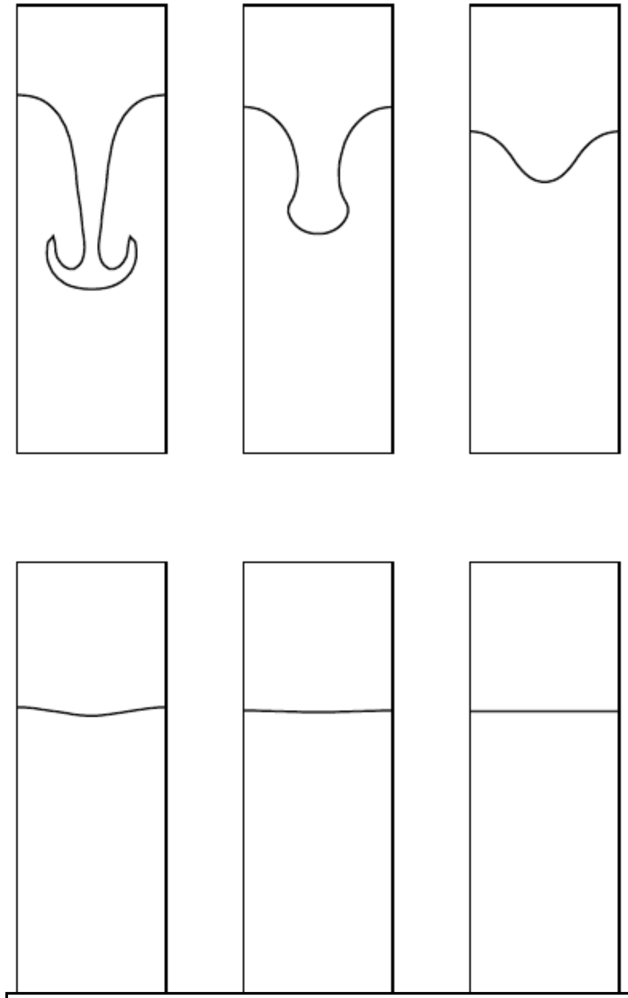

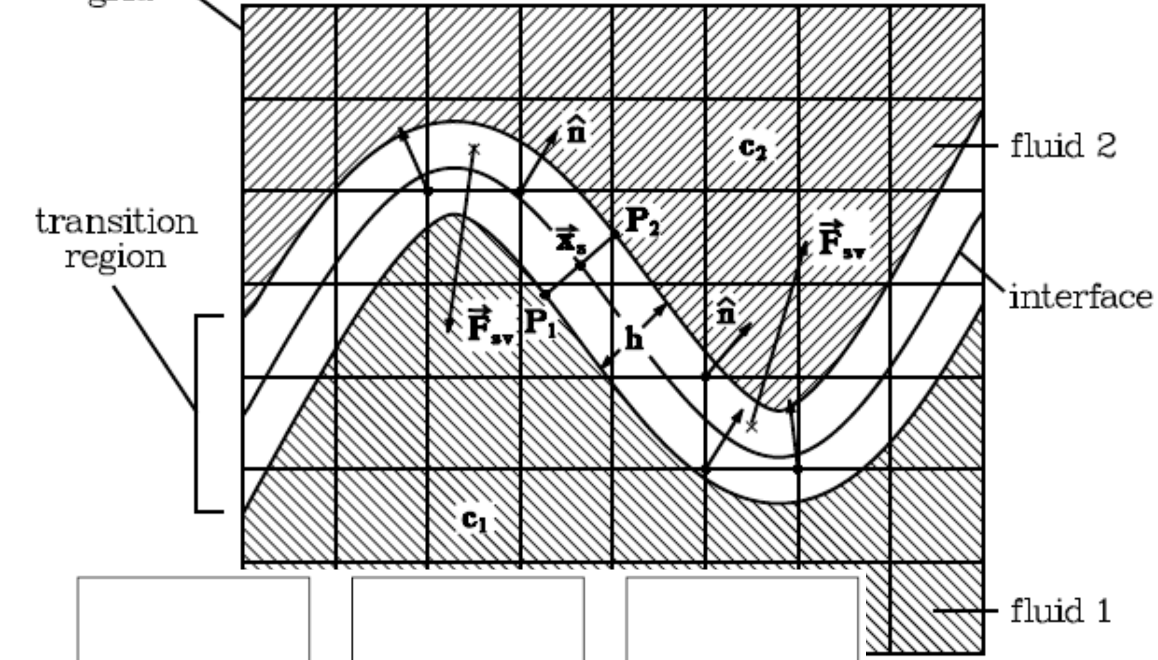

Recognize in models for interface kinematics and dynamics that the interface has a finite thickness for h>0

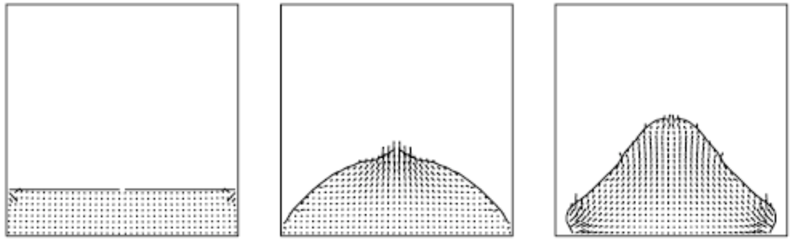

Imposition of a simple boundary condition gives an interesting model for wetting / de-wetting

The CSF model has withstood the test of time

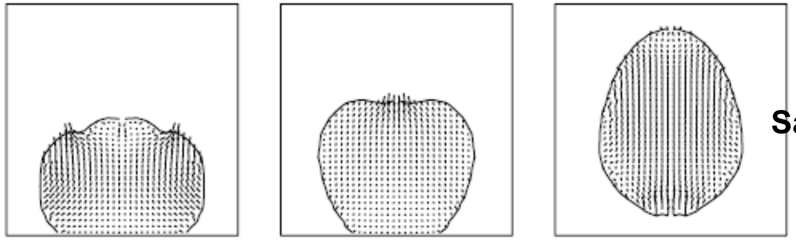

Sally Ride

Micro-Gravity Free Surface Flows for NASA
Development of "Ripple"

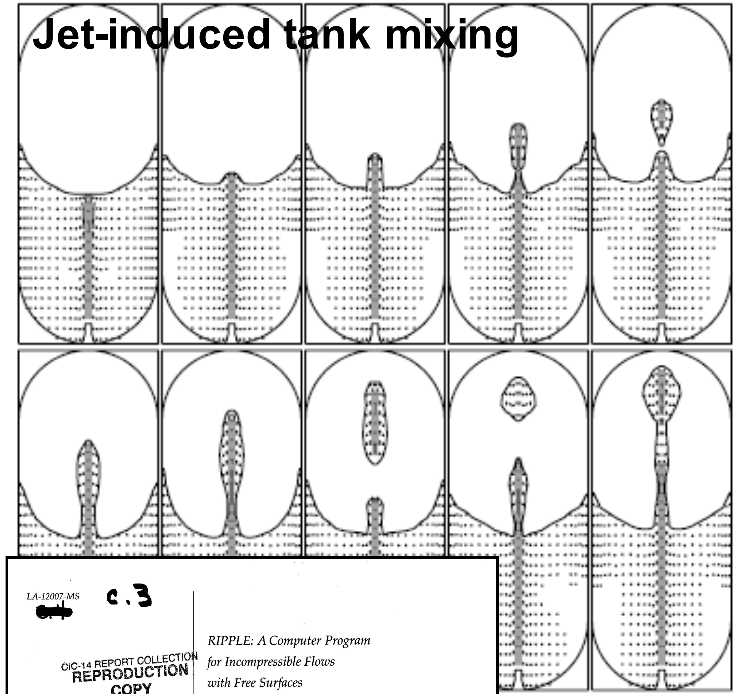

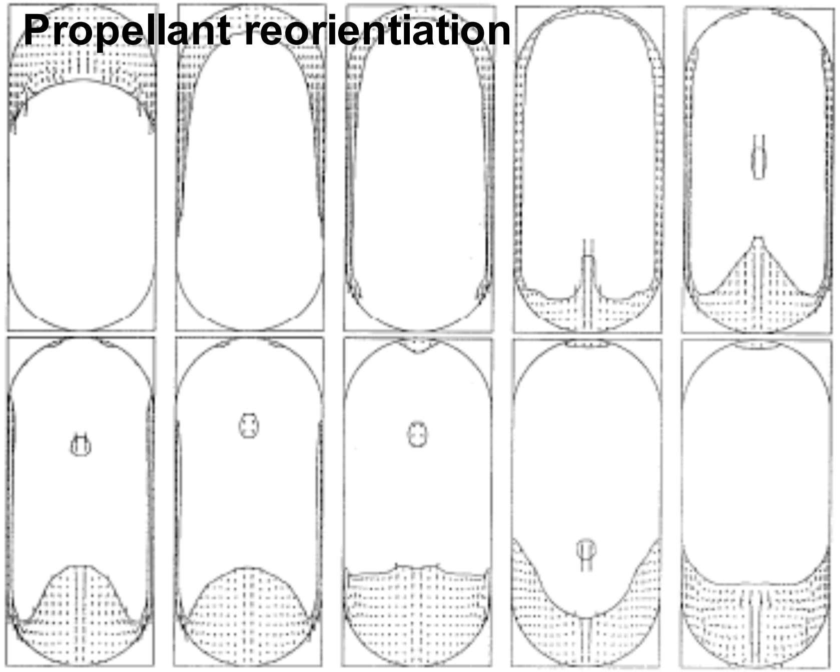

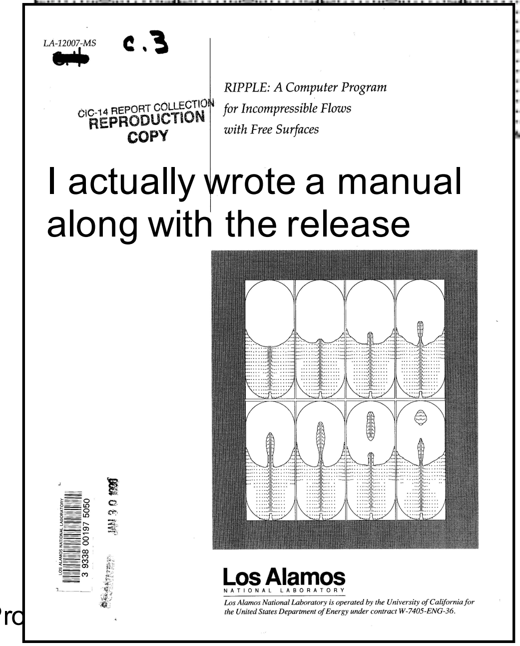
LA-12007-MS
c.3

RIPPLE: A Computer Program for Incompressible Flows with Free Surfaces

Los Alamos
NATIONAL LABORATORY

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36.

I actually wrote a manual along with the release

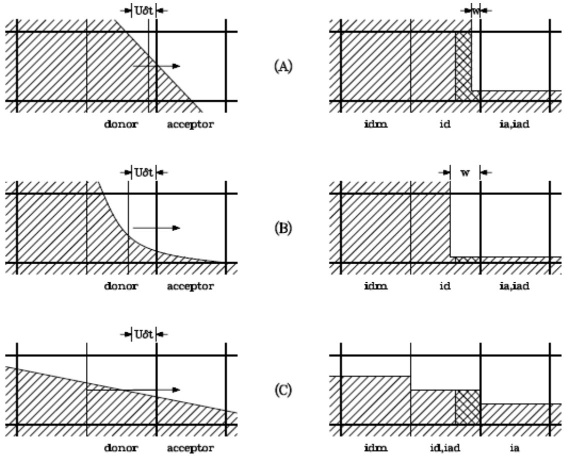

- Formal public release of software in 1991 at the ESTSC (in Oak Ridge!)
- Top 5 downloaded package for many years
- OS: COS 1.15 (Cray X-MP)

Build It and Application Developers Will Come

Connection Machines at LANL in Early 1990s

CM-2, CM-200 (1989)

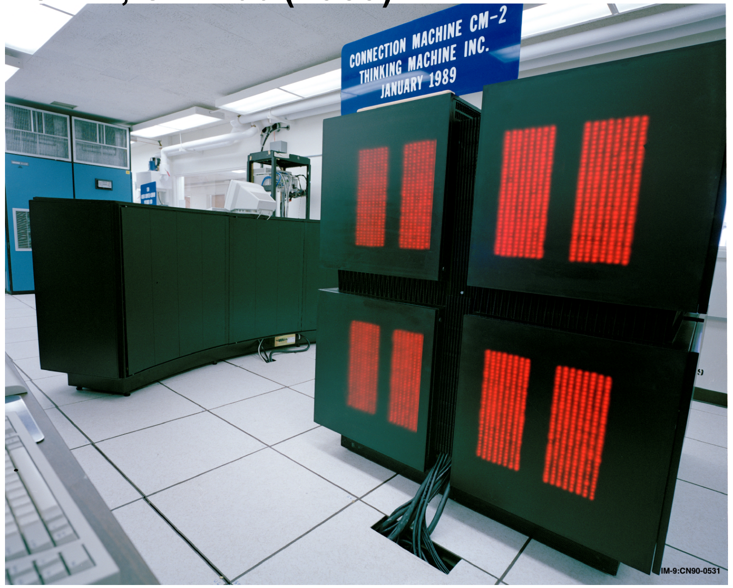

CM-5 (1992)

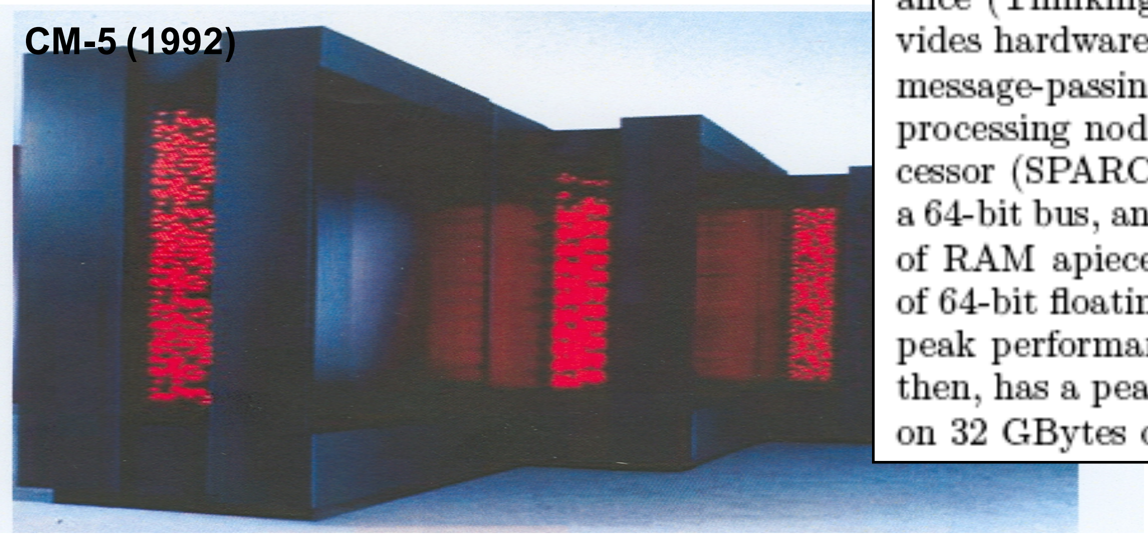

The CM-200 at LANL is a SIMD parallel supercomputer that has 2^{16} bit-serial processors and 2^{11} 64-bit Weitek floating-point units (FPUs) connected as a hypercube. Each bit-serial processor has 2^{10} Kbits of random access memory (RAM), providing a total memory of 8 GBytes. The Weitek FPUs have a 10-MHz clock, giving the CM-200 a theoretical peak speed of 40.9 GFlops (Flops = floating-point operations per second).

The CM-5 is a more flexible parallel supercomputer combining the attractive features of existing parallel architectures, including fine- and coarse-grained concurrence, MIMD and SIMD control, and fault tolerance (Thinking Machines Corporation 1991). It provides hardware support for both the data parallel and message-passing programs. The LANL CM-5 has 1024 processing nodes (PNs), each with a RISC microprocessor (SPARC technology), a network interface chip, a 64-bit bus, and 4 vector units (VUs) having 8 MBytes of RAM apiece. The VUs are capable of 32 MFlops of 64-bit floating-point performance, giving each PN a peak performance of 128 MFlops. A 1024-PN CM-5, then, has a peak performance of 131 GFlops operating on 32 GBytes of memory.

Memories: NEWS, Prism, CMSSL, CSHIFT, 10 GB DV, serial/parallel dimensions, compiler data layout

A Multi-Material Hydrodynamics Model for 3D High-Speed Flow and High-Rate Material Deformation

The PAGOSA code

PAGOSA oil well perforator simulation presented at SC91 (Albq, NM)

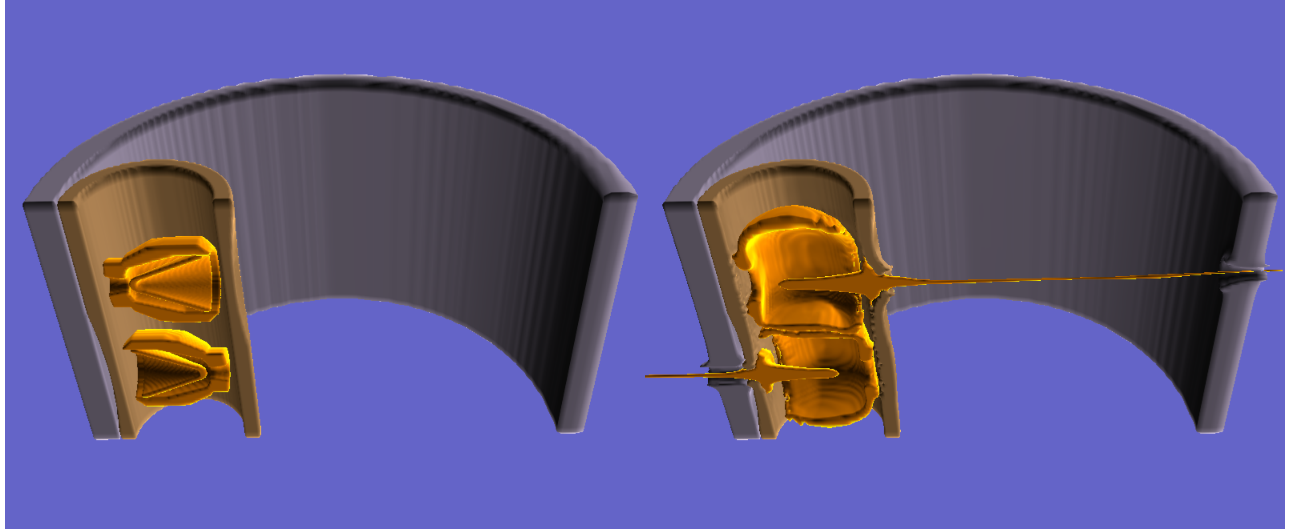

Simulation Details
- Steel carrier tube holding 2 oil well perforators inside of a steel oil well casing
- 10.3 CPU hours on 512 nodes of the CM-5
- 0.5 mm mesh size
- 1.9M cells
- 3 GB total memory

Key Features
- Developed from scratch on CM-200, CM-5 (in CMFortran)
- Finite difference discretization on structured orthogonal (hex) meshes
- Continuum mechanical conversation equations solved in Eulerian frame with a Lagrangian/remap algorithm
- 2nd-order accurate predictor-corrector method for Lagrangian phase; 3rd-order van Leer upwind scheme for advection
- Piecewise linear (“Youngs”) method for tracking material interfaces
- Ported later by SNL to nCUBE2 and Intel Paragon XP/S (SAND97-2551)
- Funded in part by Joint DoD/DOE Munitions Technology Development Program

ECP EXASCALE COMPUTING PROJECT

Application Example
Gravity Casting Mold Fill

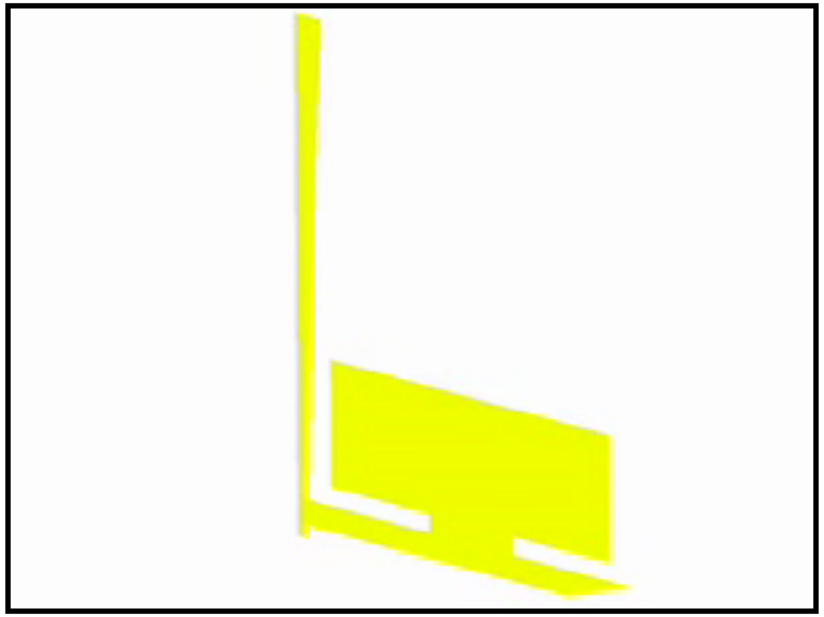

- International casting mold-fill benchmark put together by manufacturing industries and R&D agencies
- Molten aluminum into air
  - Turbulence free surface flow of Navier-Stokes fluid
- Blind test tried by *many* codes – and many failed
- We “cheated” – did not do our first simulation until years after the benchmark was published

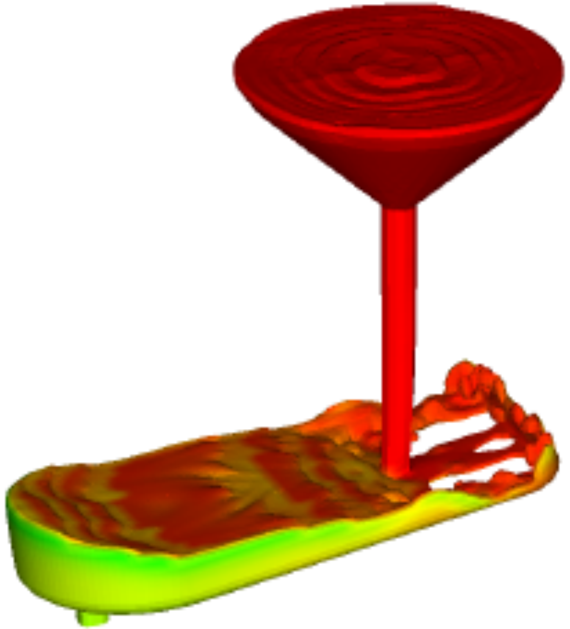

Crucible draining into a “launder” of an actual gravity-pour casting process

EXASCALE COMPUTING PROJECT

Application Example
Dam-Break Bore Interacting with Square Cylinder

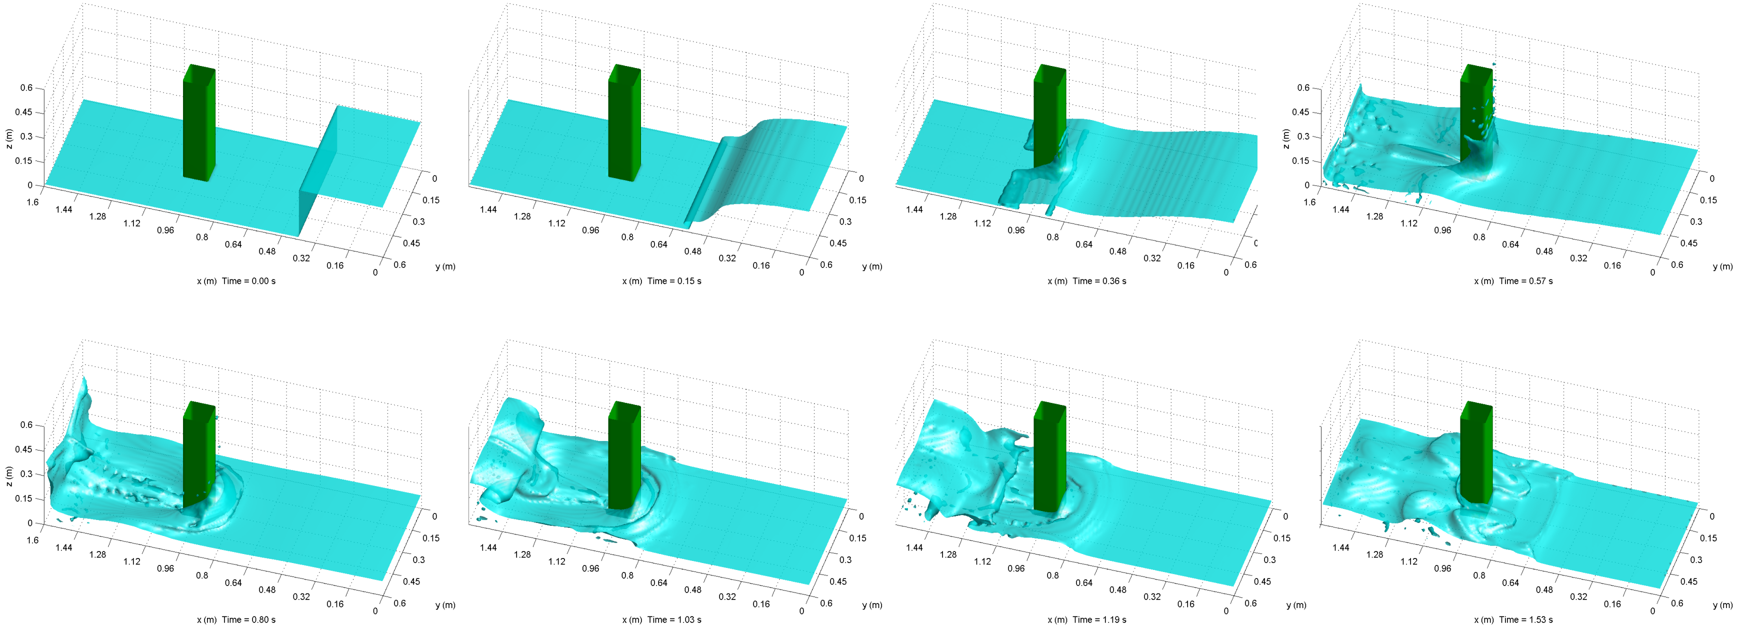

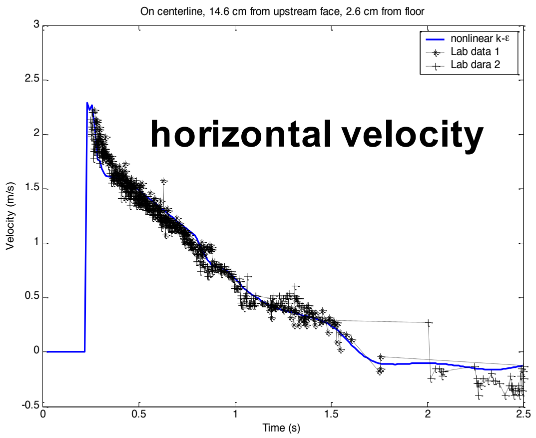

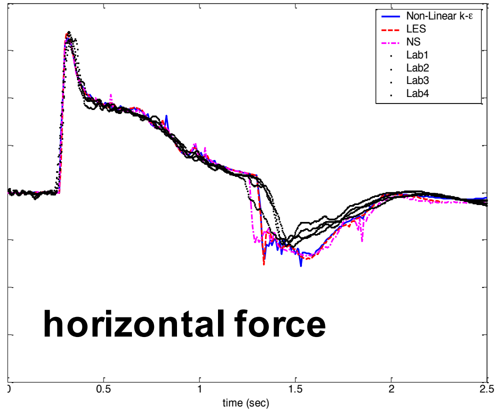

- W. Mo, K. Irschik, H. Oumeraci, and P. L.-F. Liu, *A 3D numerical model for computing non-breaking wave forces on slender piles,* J Eng Math 58:19–30 (2007)

EXASCALE COMPUTING PROJECT

U.S. Strategic Advantage: NSCI Critical Applications

NSCI vision is supported by a national tech base, integrated capabilities, and R&D foundation

President issued an executive order creating the NSCI -- National Strategic Computing Initiative

NSCI has 5 strategic themes:

- Create systems that can apply Exaflops of computing power to Exabytes of data (Exa = 10^{18})
- Keep the United States at the forefront of High Performance Computing (HPC) capabilities
- Improve HPC application developer productivity
- Make HPC readily available
- Establish hardware technology for future HPC systems

National Security

- Stockpile Stewardship
- Decision Support
- Battlefield Command
- Counter-Terrorism
- Secure Communication
- Cyber Defense
- Signals Intelligence

Economic Competitiveness

- Energy Production
- Advanced Manufacturing
- Digital Engineering
- Drug Design
- Personalized Medicine
- Health Care
- Business Analytics
- Financial Services
- E-Commerce
- Social Networking

Scientific Discovery

- Climate Science
- Fusion Science
- Materials Genome
- Particle Physics
- Neuroscience
- Weather Prediction
- Genomic Discovery

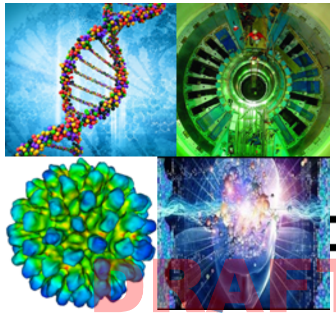

ECP Application Strategy: Statement of Mission Need

- **Science Needs: Support Six Science Programs**
  - Discovery and characterization of next-generation *materials*
  - Systematic understanding and improvement of *chemical processes*
  - Analysis of the extremely large datasets resulting from the next generation of *particle physics* experiments
  - Extraction of knowledge from *systems-biology* studies of the microbiome.
  - Advances in *applied energy* technologies, notably whole-device modeling in plasma-based fusion systems
- **National Security Needs**
  - Stockpile Stewardship Annual Assessment and Significant Finding Investigations
  - Robust UQ techniques in support of lifetime extension programs
  - Understanding evolving nuclear threats posed by adversaries and in developing policies to mitigate these threats
- **Key Science and Technology Challenges to be Addressed with Exascale**

Materials Discovery and Design
Nuclear Energy
Large-Data Applications
National Security

Climate Science
Combustion Science
Fusion Energy
Additive Manufacturing

ECP Application Strategy

- **Precipitate, foster, and grow a national push in applications**
  - Deliver to ten stakeholder DOE program offices: 5 SC, 4 Energy, NNSA DP, …
  - Deliver to NSCI ”deployment agencies”: NIH, NOAA, NASA, DHS, FBI
  - Deliver on next generation stockpile stewardship applications (3 NNSA Labs)
  - Require incubation of “non-traditional” applications (data science, experimental work flow)
  - Reassert international leadership and grow the gap between 1st (U.S.) and the followers
- **Applications should support key national initiatives**
  - National Security, Energy Security, Economic Security, Scientific Discovery, Climate/Environmental Science, Healthcare
- **Applications must adequately drive a broad set of requirements for the software ecosystem**
- **Support application teams in Tiers and seek/secure programmatic cost sharing**
  - Tiers based on complexity, maturity, strategic importance, ranging from “startup” to teams of 5, 10, and 15+
  - Startup crucial to address promising applications currently viewed as high risk based on their technology starting point and technical plans
  - Not all application teams led by DOE (e.g., could be led by Other Federal Agencies, Universities, Industry)

DEPARTMENT OF ENERGY

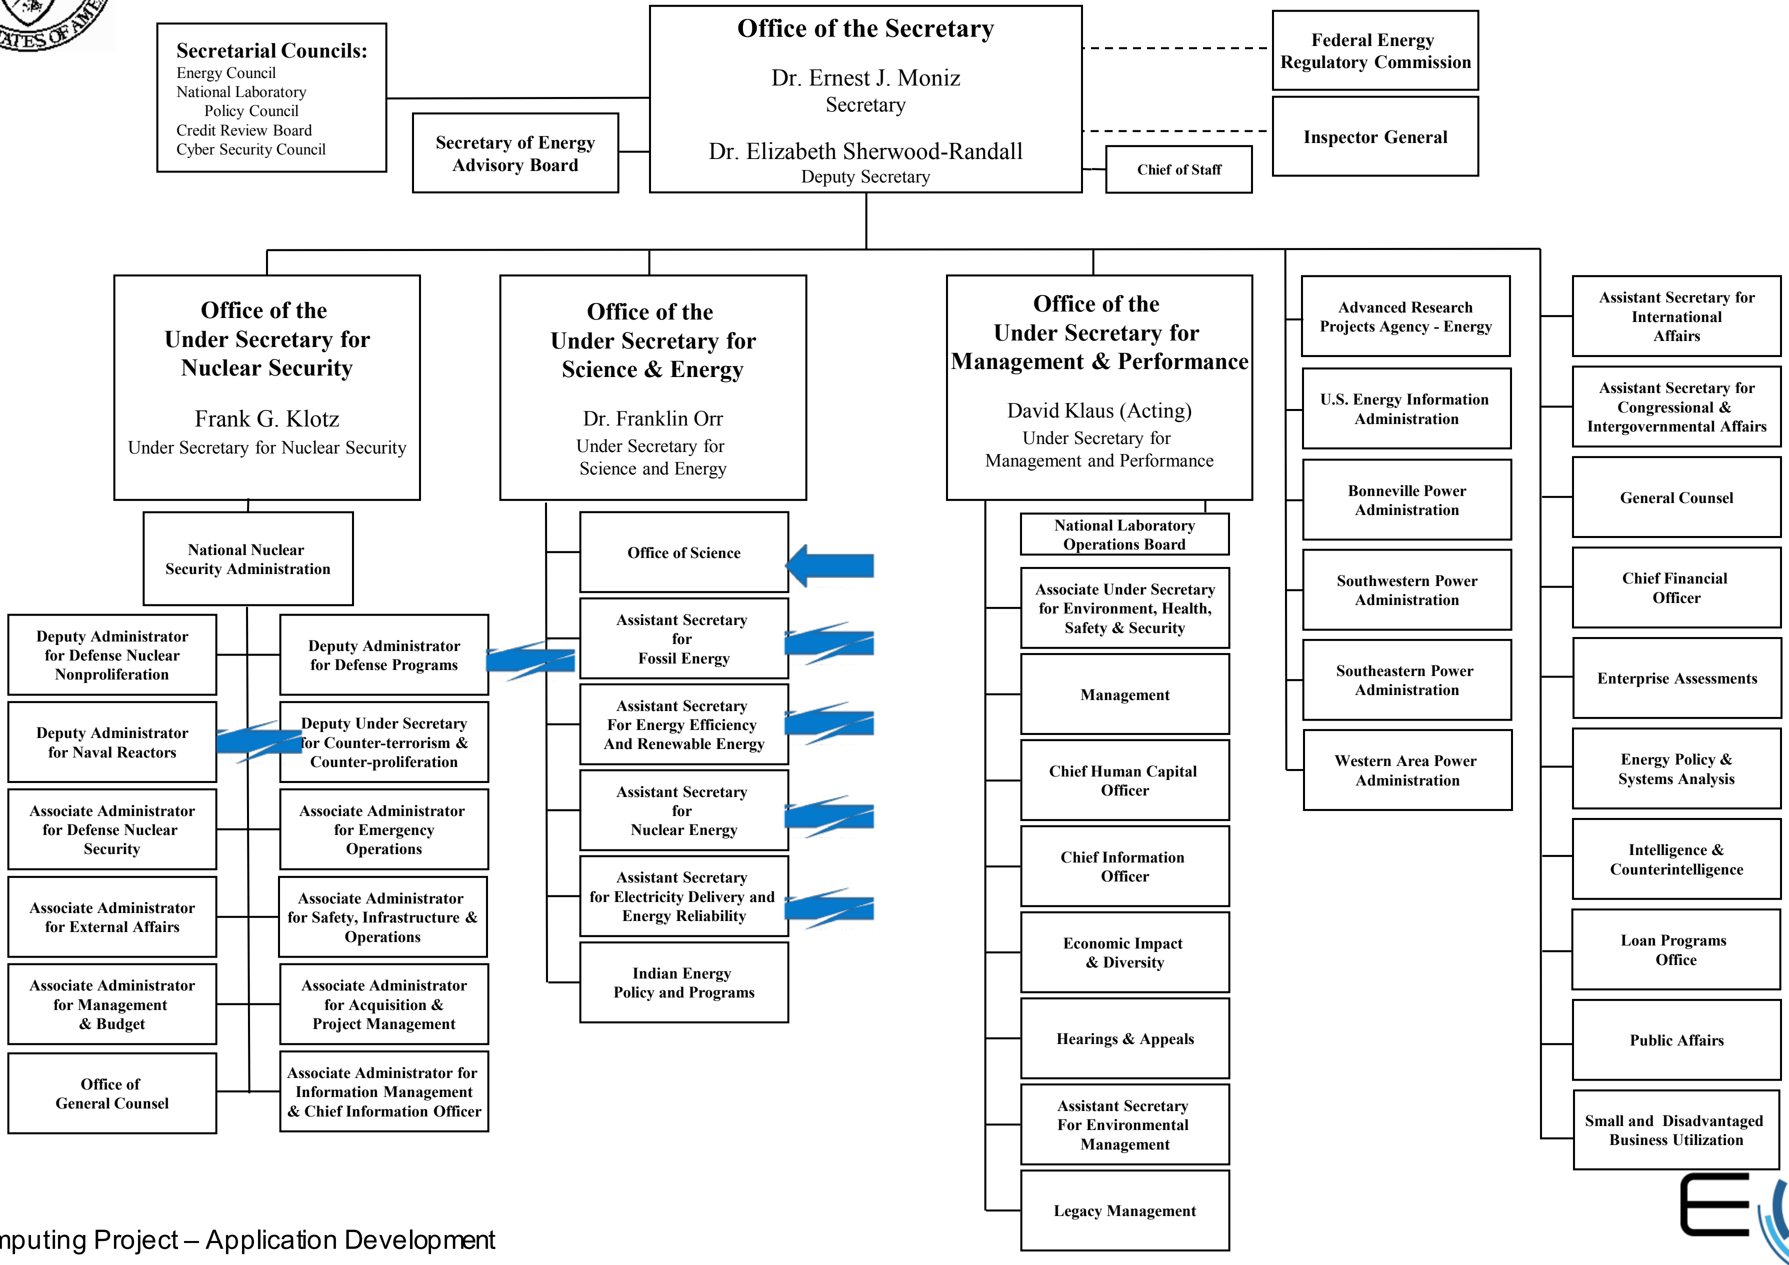

U.S. Strategic Advantage: Strengthened with ECP Applications

Realized through targeted and aggressive application development fostered by ECP

National Security

- Stockpile Stewardship
- Hypersonic Scramjet-Glide Vehicle Design

Energy Security

- Turbine Wind Plant Efficiency
- Design/Commercialization of SMRs
- Nuclear Fission and Fusion Reactor Materials Design
- Subsurface Use for Carbon Capture, Petro Extraction, Waste Disposal
- High-Efficiency, Low-Emission Combustion Engine and Gas Turbine Design
- Carbon Capture and Sequestration Scaleup
- Biofuel Catalyst Design
- Advanced IC Engine Design
- Energy Conversion and Storage Materials Design
- Lignocellulosic Biomass Deconstruction
- Materials for Batteries, Solar Cells, Optoelectronics

Economic Security

- Additive Manufacturing of Qualifiable Metal Parts
- Urban Planning
- Reliable and Efficient Planning of the Power Grid
- Earthquake Hazard and Risk Assessment
- Metal Fatigue Prediction/Control

Scientific Discovery

- Cosmological Probe of the Standard Model (SM) of Particle Physics
- Validate Fundamental Laws of Nature (SM)
- Plasma Wakefield Accelerator Design
- Light Source-Enabled Analysis of Protein and Molecular Structure and Design
- Find, Predict, and Control Materials and Properties
- Predict and Control Stable ITER Operational Performance
- Demystify Origin of Chemical Elements
- Formation of Heavy Elements in the Universe

Climate and Environmental Science

- Accurate Regional Impact Assessment of Climate Change
- Stress-Resistant Crop Analysis and Catalytic Conversion of Biomass-Derived Alcohols
- Metegenomics for Analysis of Biogeochemical Cycles, Climate Change, Environ Remediation
- Systems Biology Model of Environmental Bioorganisms

Healthcare

- Accelerate and Translate Cancer Research

ECP EXASCALE COMPUTING PROJECT

ECP Application Development
Key Assumptions

- Each application addresses a Challenge Problem simulation target
  - Challenge Problem targets span a spectrum of technical risks/challenges
  - All are impactful, mission critical, and possess an appropriate success likelihood
- Application characteristics (hence funding support) will vary
  - Some may have co-investment (e.g., DOE program offices or agencies), some not
  - Some start from mature code bases; some only as promising prototypes
  - Some explore new problems and paradigms (e.g., data science) in non-traditional HPC areas
- Application development team size and quality key to success
  - Led by experienced computational scientists with proven track records
  - Remain moderate in size (5-10); larger teams can have negative productivity
  - Expected to adopt agile project management approaches
  - Some mature/complex apps will need larger teams to meet required targets, but must prove they have a structure that scales

EXASCALE COMPUTING PROJECT

Exascale Applications Will Address National Challenges

Nuclear Energy (NE)

Accelerate design and commercialization of next-generation small modular reactors

Climate Action Plan; SMR licensing support; GAIN

Climate (BER)

Accurate regional impact assessment of climate change

Climate Action Plan

Carbon Capture and Storage (FE)

Scaling carbon capture/storage laboratory designs of multiphase reactors to industrial size

Climate Action Plan; SunShot; 2020 greenhouse gas/2030 carbon emission goals

Wind Energy (EERE)

Increase efficiency and reduce cost of turbine wind plants sited in complex terrains

Climate Action Plan

Combustion (BES)

Design high-efficiency, low-emission combustion engines and gas turbines

2020 greenhouse gas and 2030 carbon emission goals

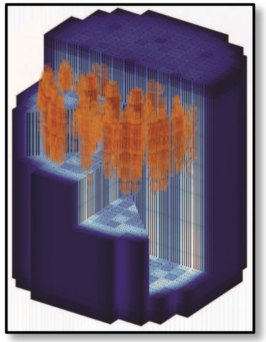

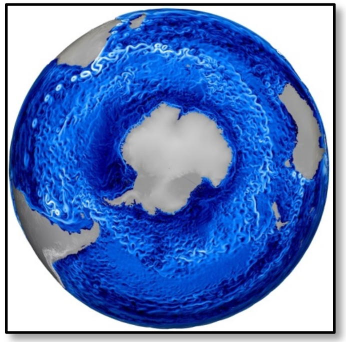

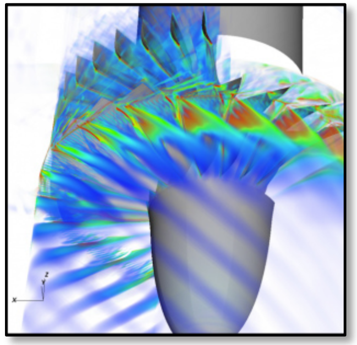

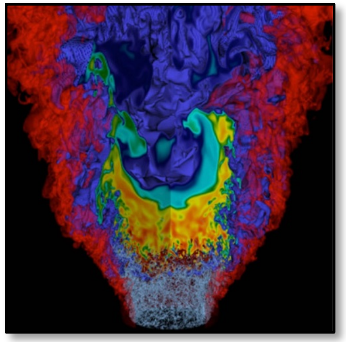

Exascale Application Development

Science and Energy Driver

Nuclear Energy

Gaps and Opportunities

- ✓ Understanding and predicting fuel failure and core damage in severe reactor accidents
- ✓ Near real-time load-following core simulator supporting onsite operating plant decisions
- ✓ Engineering scale predictions of nuclear fuel performance and barriers to higher burnup

Simulation Challenge Problems

- ✓ Core-wide multi-physics: Monte Carlo neutronics, fluids, fuels, chemistry, material, local wear and contacts, structural response
- ✓ Core-wide multi-scale fluids performance with two-way DNS-to-LES CFD coupling

Prospective Outcomes and Impact

- ✓ Accelerate and innovate designs for small and advanced reactors
- ✓ Virtual test reactor for advanced fuel designs outside of principal test base
- ✓ Improved efficiency, economics, and safety of existing reactor fleet

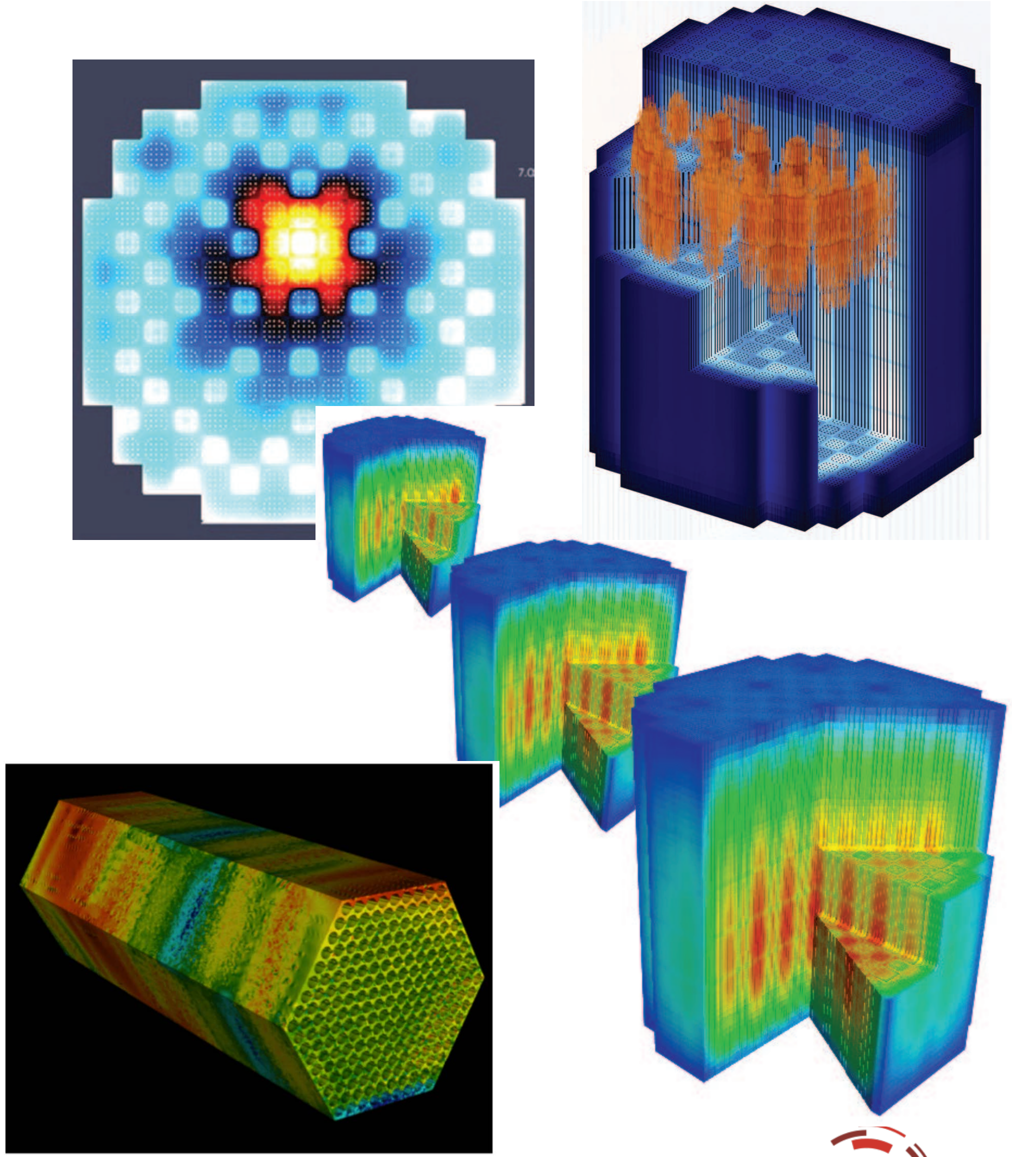

EXASCALE COMPUTING PROJECT

CASL Is Developing a High-Fidelity Virtual Environment for Reactor Applications (VERA)

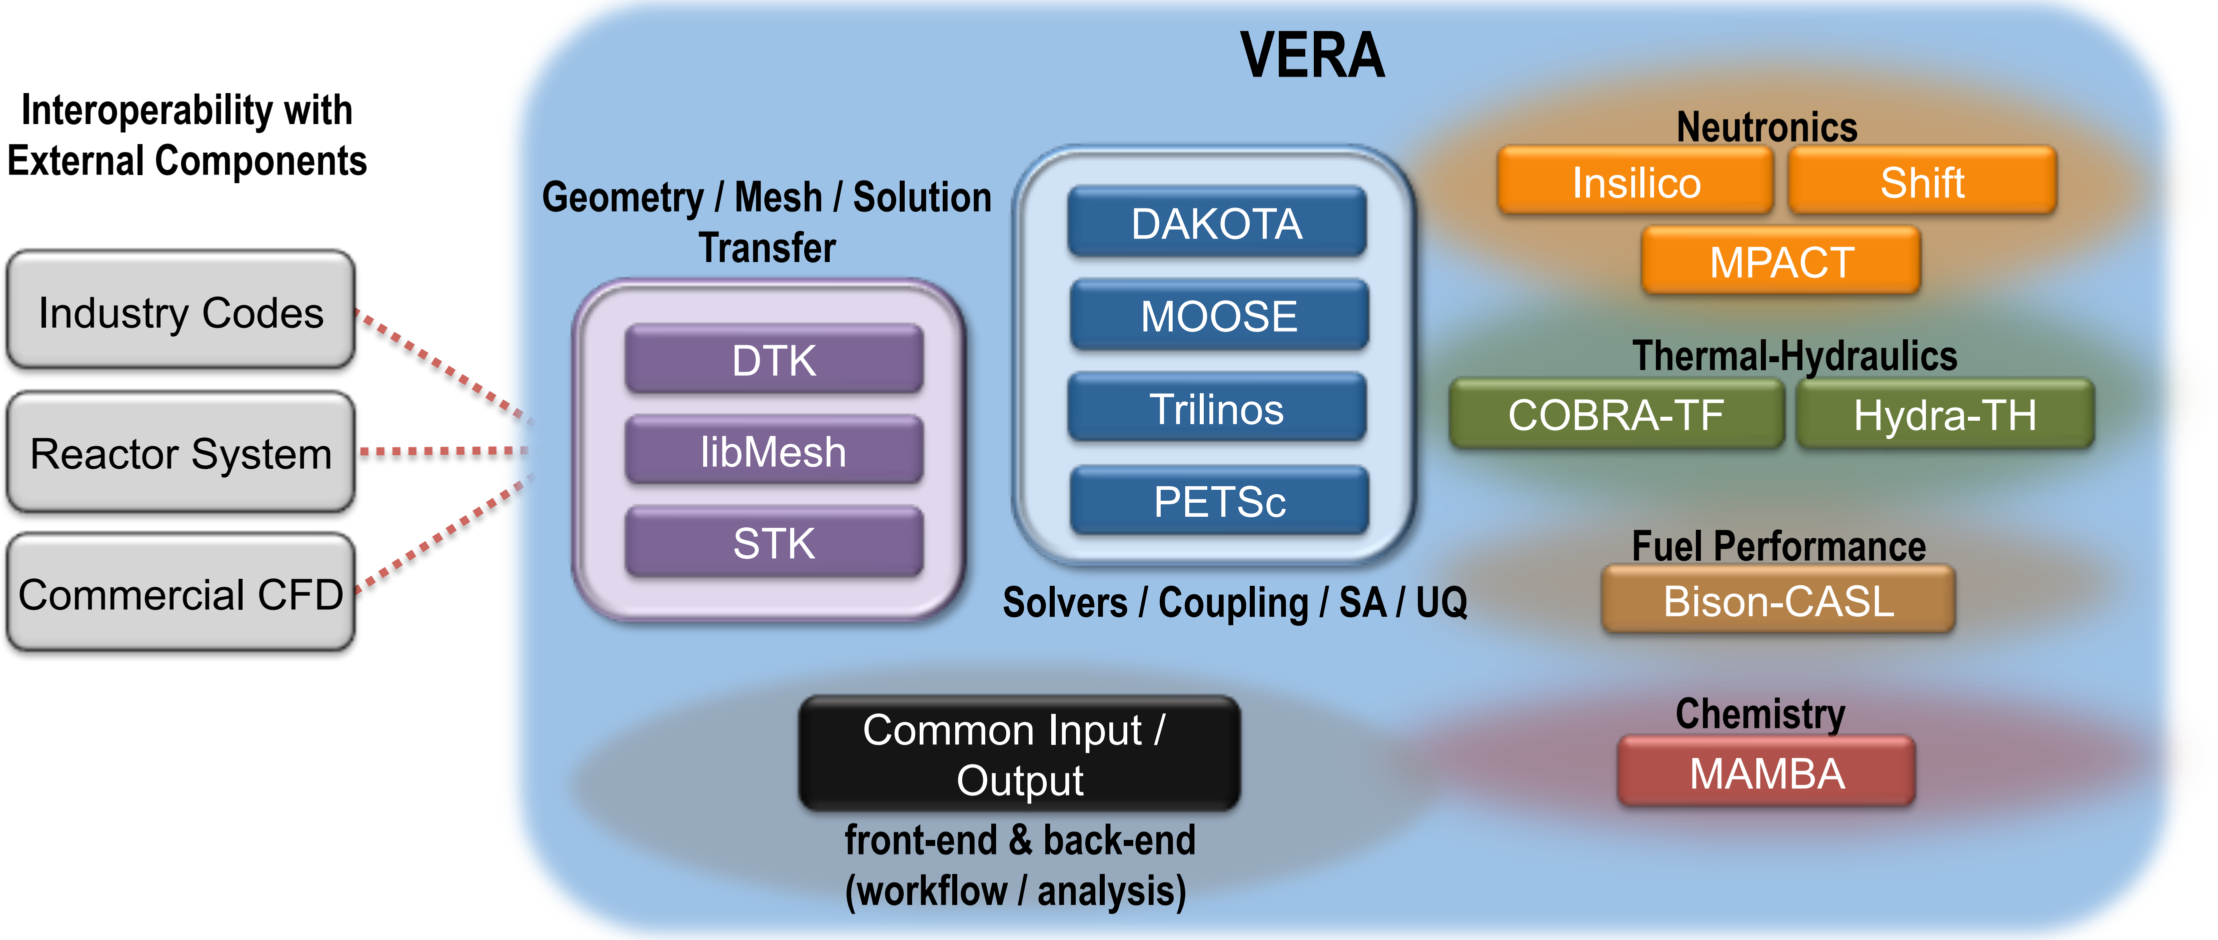

Exascale Horizon Opens the Door to Increased Predictability

Terascale Platforms standard today	Petascale Platforms possible today	Exascale Platforms possible soon?
***Engineering** Analysis*	***High-Fidelity** Core Analysis*	***Extreme-Fidelity** Analysis*
• Criticality and safety set-points • Core power predictions • Cycle fuel depletions • Transient safety analysis • Core loading optimization • Operator-assist predictions • Real-time operator training simulators	• Criticality and safety set-points • Core pin power predictions • Cycle isotopic fuel depletions • Localized sub-channel feedback • Assembly or full core structural models	• Azimuthal/radial intra-pellet isotopics • Rim effects in high burnup fuel pins • Localized CRUD deposition//corrosion • Fluid/structure vibrations/wear • Physics-based DNBR predictions • Vessel flow asymmetry and instabilities • Fully coupled TH/structural full core

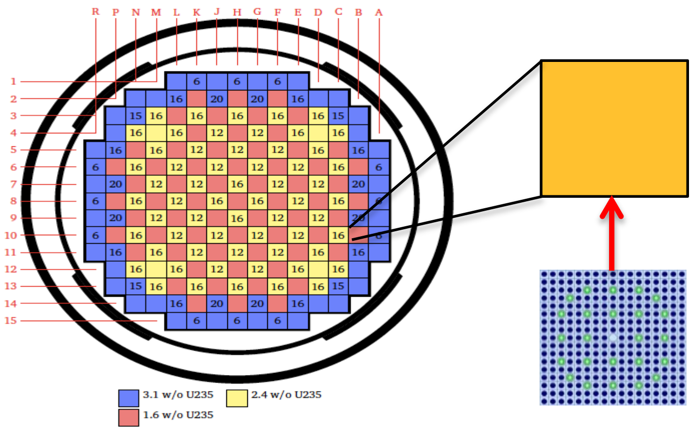

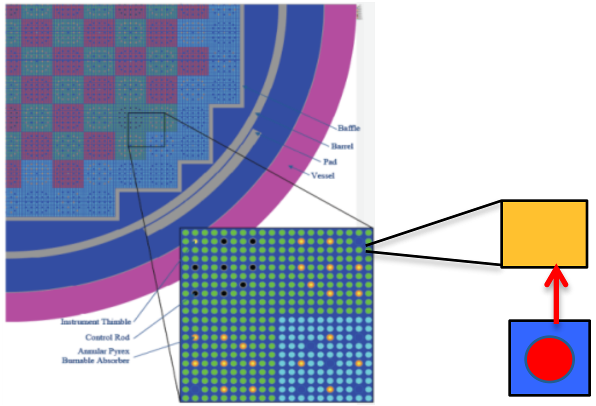

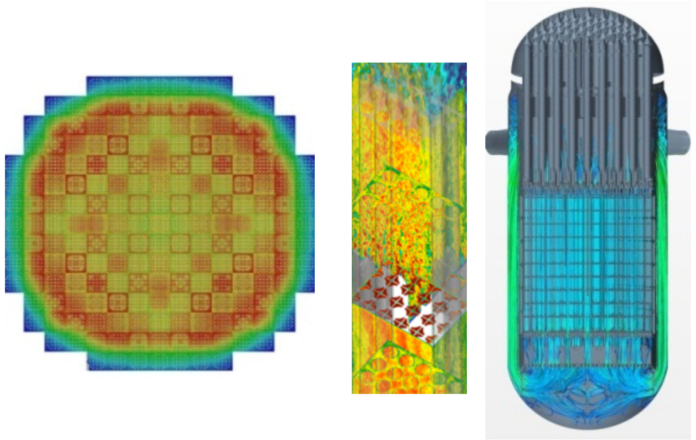

*Homogenized Fuel **Assemblies***	*Homogenized Fuel **Pin-Cells***	***Explicit** Fuel/Clad/Fluids &Vessel*
• Pre-computed assembly data tables • Few-group nodal diffusion neutronics • Characteristic-channel fuel pin • Characteristic-channel thermal fluids • Macroscopic fuel assembly depletion • Lumped-parameter closure relations	• Pre-computed pin-cell data tables • Multi-group transport neutronics • Simplified explicit-pin fuel mechanics • Sub-channel and CFD thermal fluids • Microscopic fuel pin depletion • Simplified-physics closure relations	• No pre-computed data tables • Continuous-energy Monte Carlo • Meso/macro fuel pin mechanics • CFD and DNS thermal fluids • Intra-pellet isotopics in fuel depletion • Physics-based closure relations

Exascale Application Development

Science and Energy Driver

Climate

Gaps and Opportunities

- ✓ Higher confidence projections of global and regional climate system changes (rise in sea level and rainfall events) thru increased realism, number and reliability of model-based climate predictions
- ✓ Quantification of uncertainty in climate model prediction
- ✓ More accurate explicit simulation of local to global weather phenomena, including extreme events, resulting in improved probabilistic prediction across time and space scales.

Simulation Challenge Problems

- ✓ 1 km resolved scales in global simulations with feature-aware models tracking 3D dynamics similar to weather models
- ✓ High resolution coupling of ocean (<10 km) to atmosphere (<5 km), with even higher resolution under ice shelves
- ✓ Fully coupled atmosphere-ocean-land-land ice-sea ice model executing at 5 simulated years/day with 1 km atmosphere in critical regions, 100 m land/land ice, and 10 km ocean and sea ice

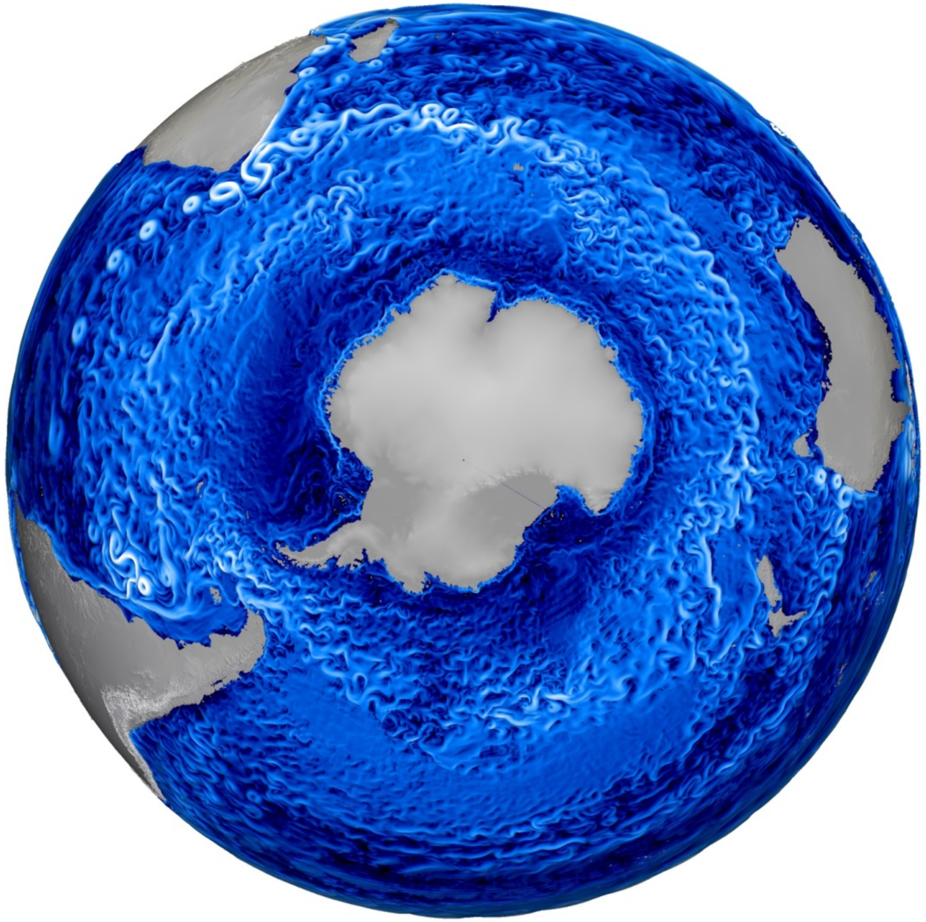

ACME
Accelerated Climate Modeling for Energy

Exascale Application Development

Science and Energy Driver

Climate

Prospective Outcomes and Impact

- ✓ Forecast water resources with increased confidence
- ✓ Project future changes in very severe weather with increased confidence
- ✓ Address food supply changes with defensible estimations of the impact on productivity of crop, tree, and grass species
- ✓ Provide informed input on regional climate change impacts
- ✓ Reduce quantified uncertainty in the timing and extent of sea level rise
- ✓ Simulate small-scale climate change impacts relevant to the resiliency of energy and public health infrastructure

Support the President's 2013 Climate Action Plan and the global climate deal adopted at the 2015 United Nations Conference on Climate Change

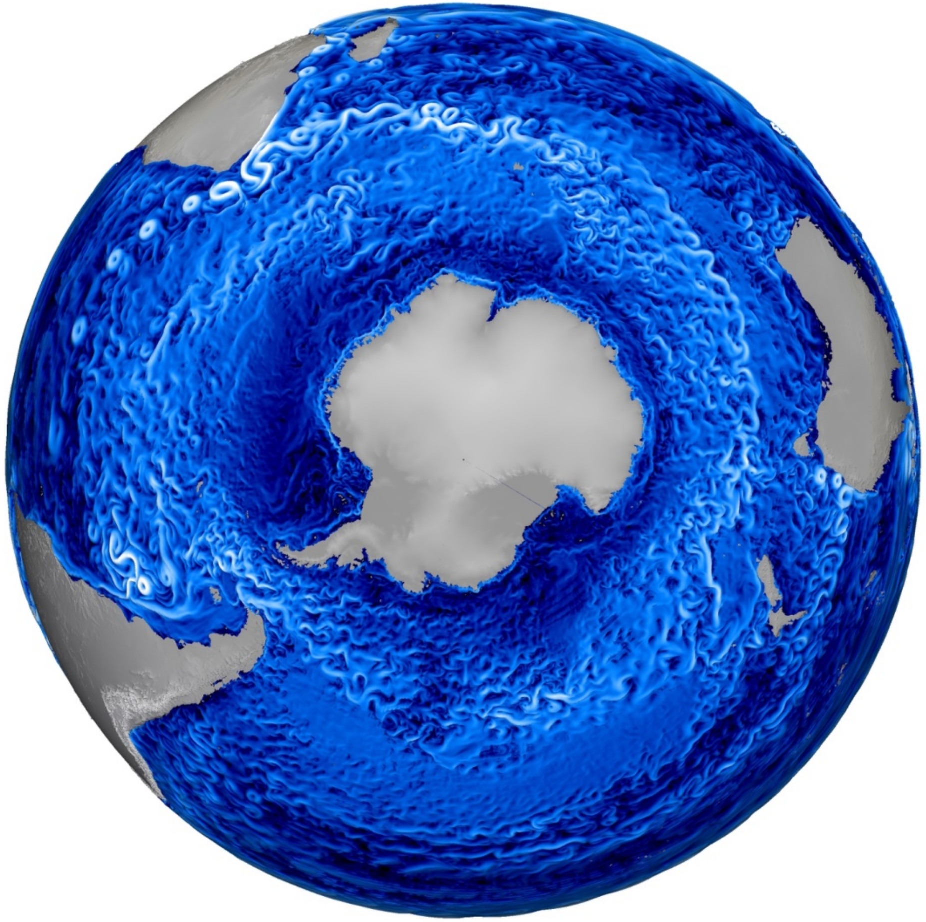

Exascale Application Development

Science and Energy Driver

Carbon Capture and Storage

Gaps and Opportunities

- ✓ Ability to conclusively determine best conceivable mixed-matrix membrane technology for post-combustion CO_2 capture
- ✓ Predictive simulation tools for carbon storage, monitoring and verification, risk assessment, and use/reuse

Simulation Challenge Problems

- ✓ Screen millions of membrane candidates via MD and atomistic grand canonical MC simulations for CO_2 adsorption
- ✓ Discrete element model for billions of reacting particles or bubbles in multiphase flow systems for validation of two-fluid models of CO_2 absorbers
- ✓ 100M node carbon storage reservoir-to-receptor simulation over a 1000-year period

Prospective Outcomes and Impact

- ✓ Reduce CO_2 capture costs (e.g., by \$5-\$10 per ton), helping to promote wide-scale adoption
- ✓ Safe and permanent storage and/or utilization of CO_2 captured from point sources

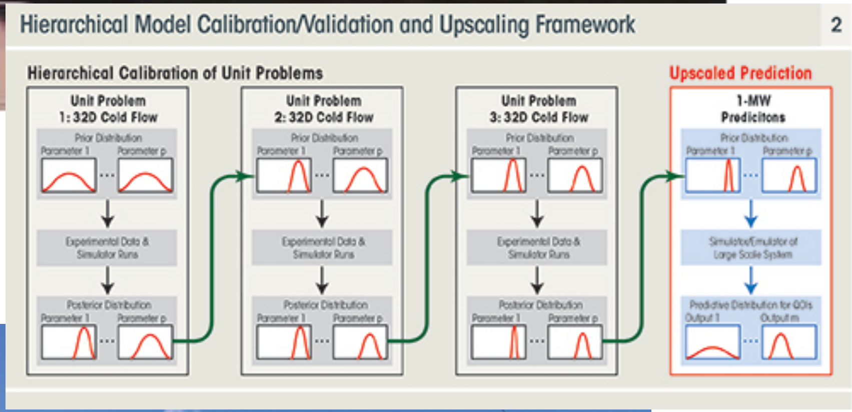

EXASCALE COMPUTING PROJECT

Exascale Application Development

Science and Energy Driver

Wind Energy

Gaps and Opportunities

- ✓ Wide-scale deployment of unsubsidized wind plants hampered by large plant-level energy losses, currently at ~20%

Simulation Challenge Problems

- ✓ Wind turbine blade-resolved CFD detached eddy simulation of a ~50-turbine wind plant over a useful operating period (hours)
- ✓ Predict plant flow physics: starting from the atmospheric boundary layer interaction with the wind plant, down to turbine-turbine interactions, the response of individual turbines, and impact of complex terrain

Prospective Outcomes and Impact

- ✓ Harden wind plant design and layout against energy loss susceptibility
- ✓ 1% wind plant performance improvement translates to >$100M annual cost savings
- ✓ Higher penetration of wind energy (~20-30%) in U.S. electrical supply supports President's Climate Action Plan

ECP EXASCALE COMPUTING PROJECT

Exascale Application Development

Science and Energy Driver

Combustion Science and Technology

Gaps and Opportunities

- ✓ Combustion-based systems will dominate marketplace for decades; must be optimized for energy efficiency and reduced emissions
- ✓ Current cut-and-try approaches for combustion system design take too long; cannot evaluate design parameter space for optimized results

Simulation Challenge Problems

- ✓ Fully coupled multi-scale and multi-physics LES treatment of complex combustion processes in propulsion and power devices, capturing combined effects of geometry, heat transfer, and multiphase reacting flow
- ✓ DNS of chemical mechanisms for kinetics of complex hydrocarbons and alternative fuels and related turbulence-chemistry interactions

Prospective Outcomes and Impact

- ✓ Address outstanding challenges for advanced gas turbines and reciprocating engines: flame stabilization, flashback, thermo-acoustics, pollutant formation in gas turbines; effects of fuel composition and spray parameters on ignition and soot formation in engines

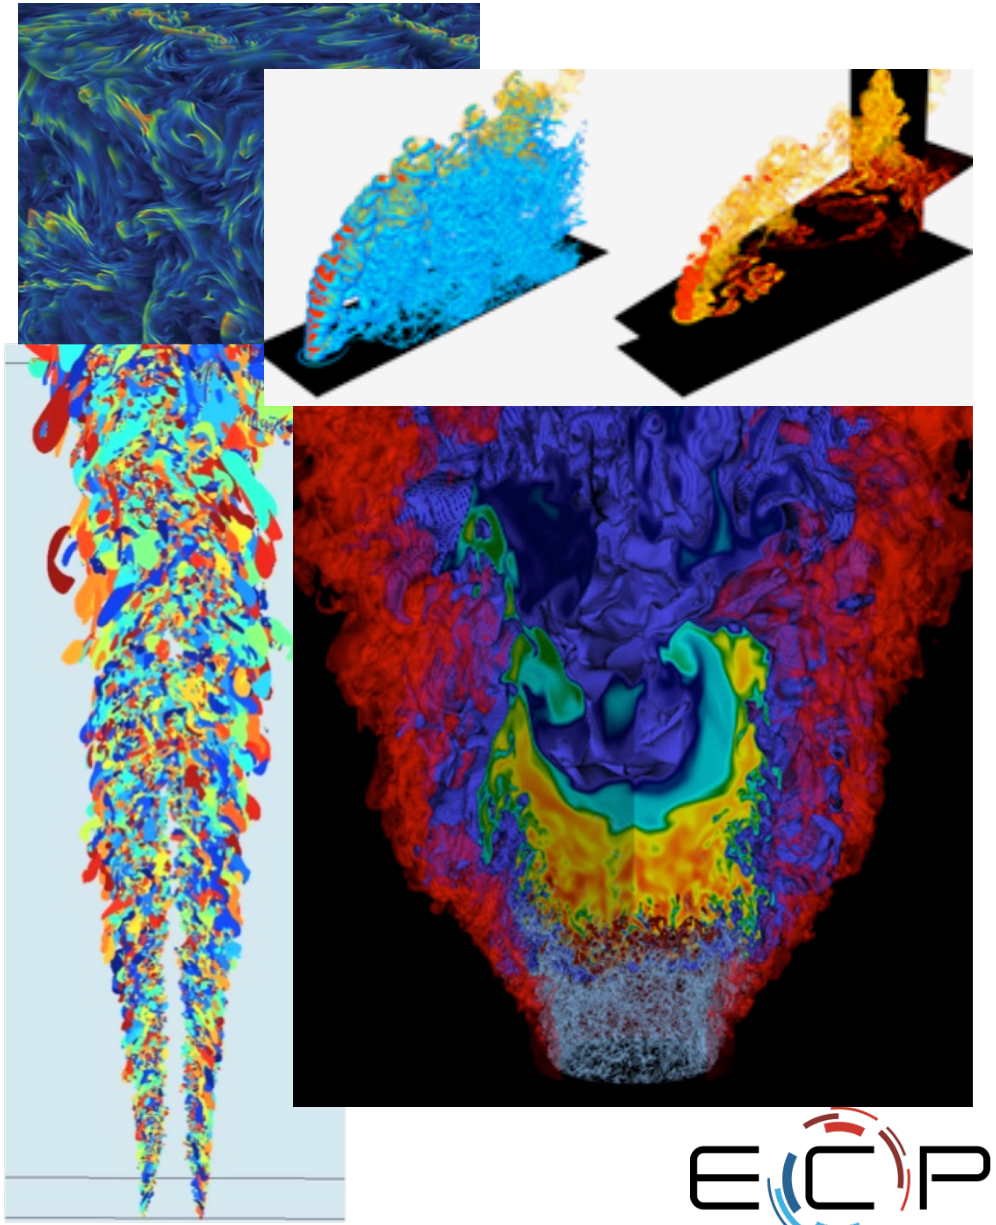

ECP EXASCALE COMPUTING PROJECT

Exascale Applications Will Address National Challenges

Magnetic Fusion Energy (FES)

Predict and guide stable ITER operational performance with an integrated whole device model

ITER; fusion experiments: NSTX, DIII-D, Alcator C-Mod

Advanced Manufacturing (EERE)

Additive manufacturing process design for qualifiable metal components

NNMIs; Clean Energy Manufacturing Initiative

Chemical Science (BES)

Design catalysts for conversion of cellulosic-based chemicals into fuels, bioproducts

Climate Action Plan; SunShot Initiative; MGI

Precision Medicine for Cancer (NIH)

Accelerate and translate cancer research in RAS pathways, drug responses, treatment strategies

Precision Medicine in Oncology; Cancer Moonshot

Urban Systems Science (EERE)

Retrofit and improve urban districts with new technologies, knowledge, and tools

Energy-Water Nexus; Smart Cities Initiative

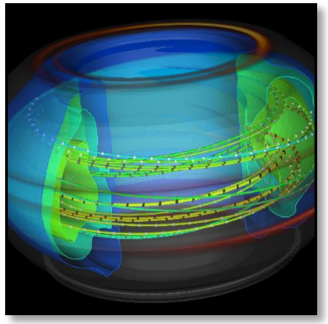

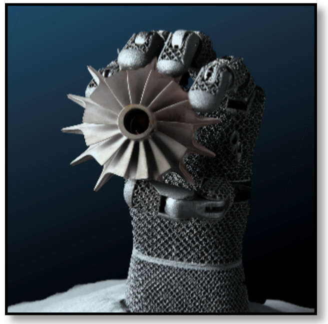

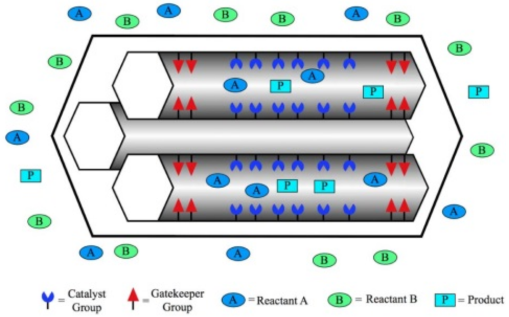

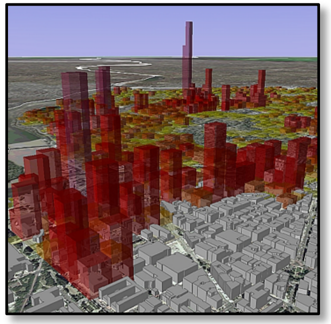

Exascale Application Development

Science and Energy Driver

Magnetic Fusion Energy

Gaps and Opportunities

- ✓ Prepare for and exploit ITER and other coming international major experiments such as JET-DT, JT-60SA, and W 7-X

Simulation Challenge Problems

- ✓ Whole-device fusion reactor simulation: tightly-coupled full-f/delta-f models and loosely coupled source/boundary models, including electron/ion kinetics, MHD, and energetic particles in core and edge regions
- ✓ Simulate and characterize tokamak disruptions and mitigations, incorporating kinetics, MHD, and fast particles
- ✓ Plasma boundary region analysis: edge kinetic effects, material interaction, radiation and detachment, power and particle exhaust

Prospective Outcomes and Impact

- ✓ Prepare for and fully simulate ITER experiments and increase ROI of validation data and understanding
- ✓ Prepare for next step beyond-ITER devices such as nuclear science facilities and DEMO

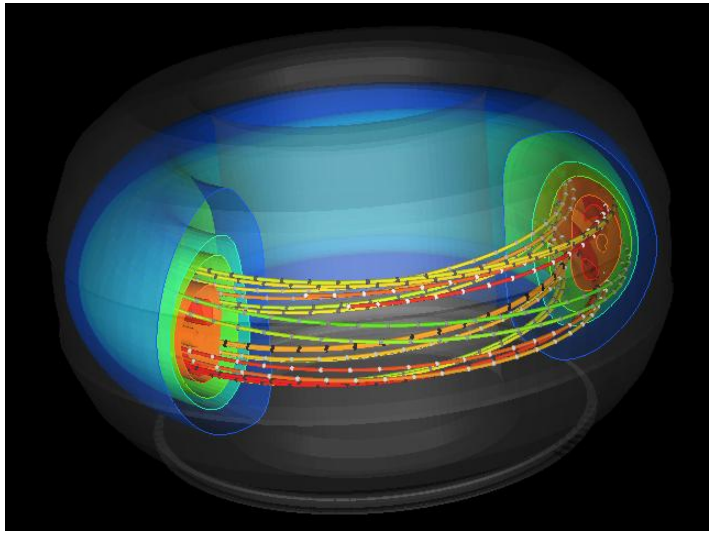

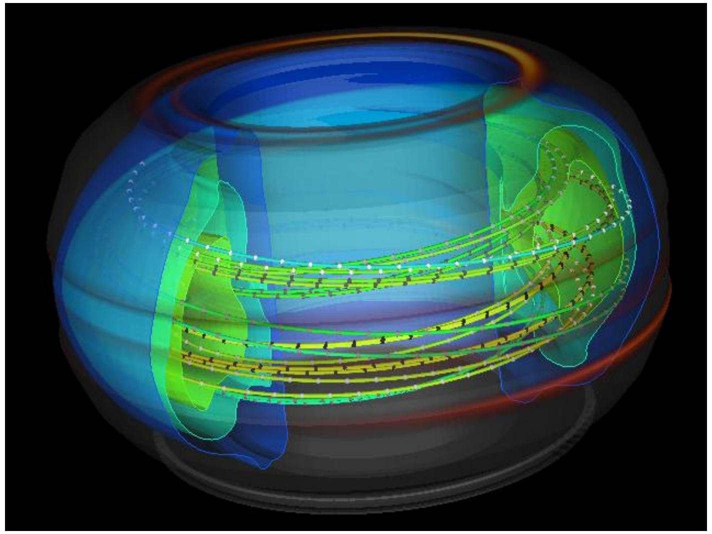

EXASCALE COMPUTING PROJECT

Exascale Application Development

Science and Energy Driver

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

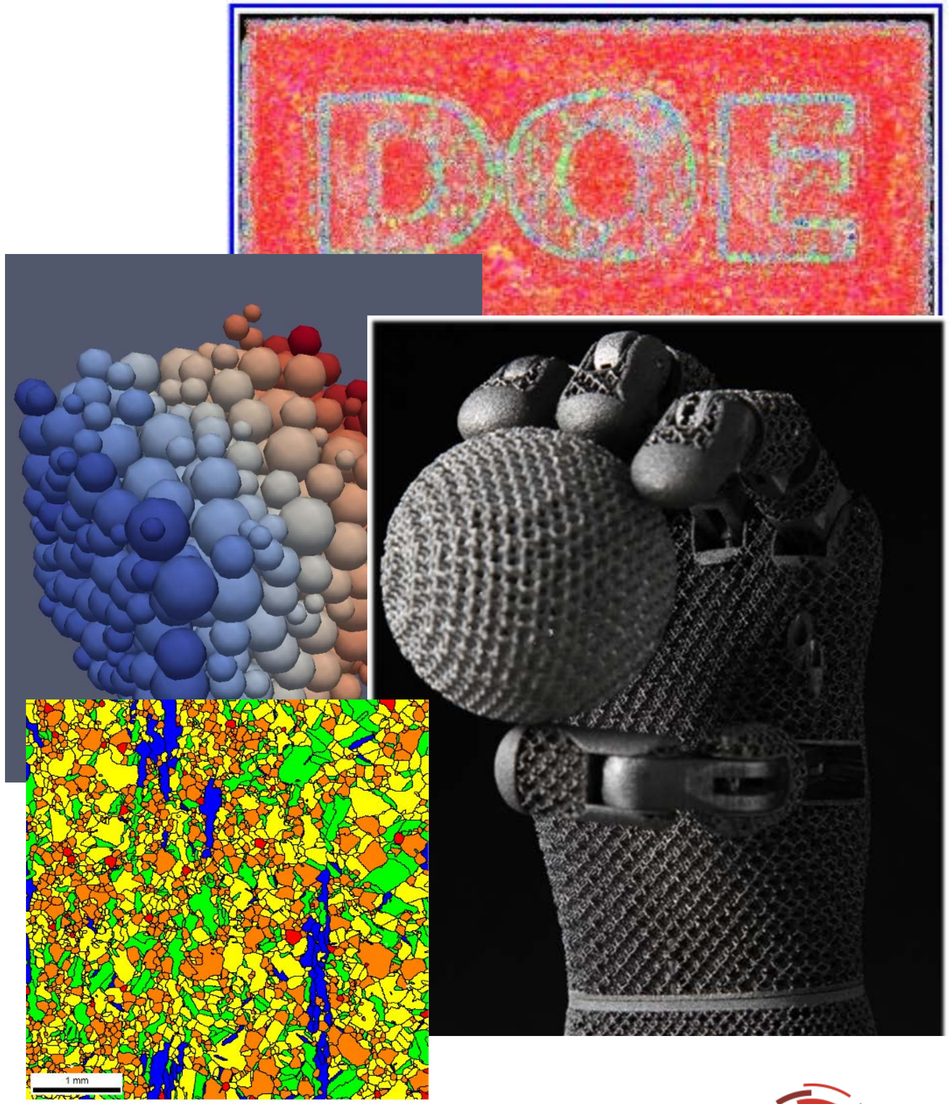

Exascale Applications Will Address National Challenges

Materials Science (BES)	Combustion Engine Design (BES, EERE)	Cosmology (HEP)	Materials Science (BES)	Nuclear Physics (NP)
Understanding fatigue in poly-crystal metals as guided by real-time experimental steering	Predict soot and NOx emission for selected fuels in low temperature internal combustion engine designs	Cosmological probe of standard model (SM) of particle physics: Inflation, dark matter, dark energy	Tailor materials for energy conversion/ storage, thermal management, thermoelectricity	QCD-based elucidation of fundamental laws of nature: SM validation and beyond SM discoveries
MGI; new energy and transportation system technologies	2020 greenhouse gas and 2030 carbon emission goals	Particle Physics Project Prioritization Panel (P5)	MGI, Climate Action Plan; SunShot Initiative; Nanotechnology-Inspired Grand Challenge for Future Computing	2015 Long Range Plan for Nuclear Science; RHIC, CEBAF, FRIB

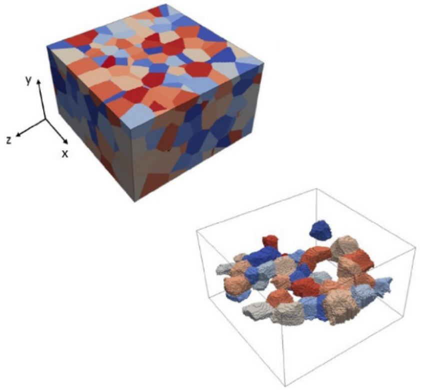

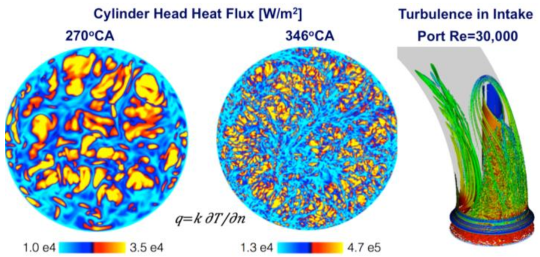

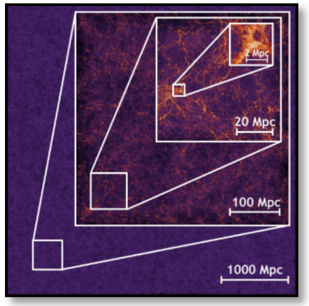

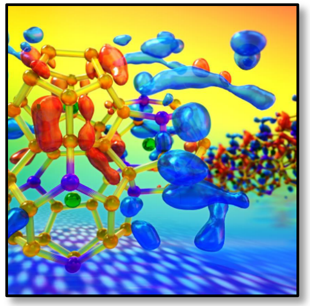

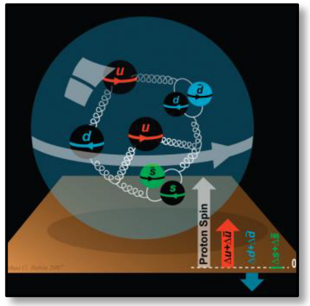

EXASCALE COMPUTING PROJECT

Exascale Applications Will Address National Challenges

Geoscience (BES, BER, EERE, FE, NE)

Safe and efficient use of subsurface for carbon capture and storage, petroleum extraction, geothermal energy, nuclear waste

EERE Forge; FE NRAP; Energy-Water Nexus; SubTER Crosscut

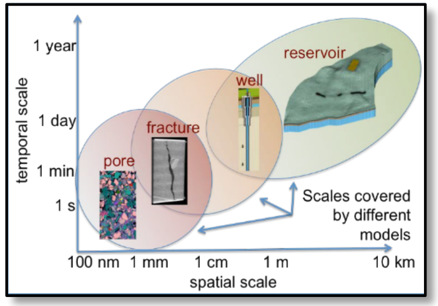

Metagenomics (BER)

Leveraging microbial diversity in metagenomic datasets for new products and life forms

Climate Action Plan; Human Microbiome Project; Marine Microbiome Initiative

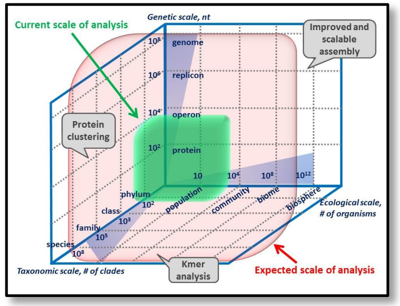

Accelerator Physics (HEP)

Practical economic design of 1 TeV electron-positron high-energy collider with plasma wakefield acceleration

>30k accelerators today in industry, security, energy, environment, medicine

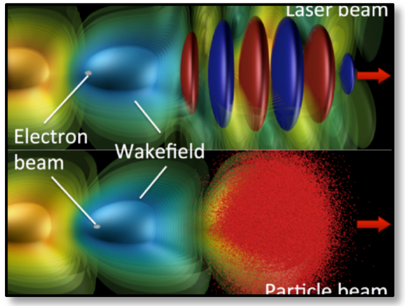

Materials Science (BES)

Efficiency and performance characteristics of materials for batteries, solar cells, and optoelectronics

MGI; Climate Action Plan

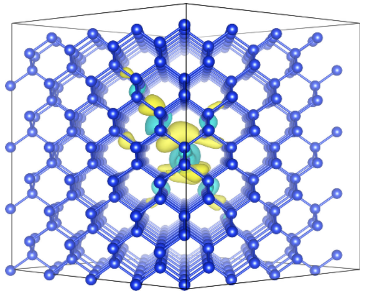

Astrophysics (NP)

Demystify origin of chemical elements (> Fe); confirm LIGO gravitational wave and DUNE neutrino signatures

2015 Long Range Plan for Nuclear Science; origin of universe and nuclear matter in universe

Exascale Applications Will Address National Challenges

Seismic (EERE, NE, NNSA)

Reliable earthquake hazard and risk assessment in relevant frequency ranges

DOE Critical Facilities Risk Assessment; urban area risk assessment; treaty verification

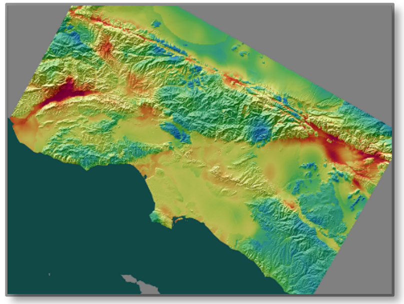

Scramjet Design (DoD, NASA)

Maneuverable hypersonic scramjet-glide vehicle design for national security and commercial space access

NASA 2030 Vision; 90-min global travel; defense technologies

Nuclear Materials (BES, NE, FES)

Extend nuclear reactor fuel burnup and develop fusion reactor plasma-facing materials

Climate Action Plan; MGI; Light Water Reactor Sustainability; ITER; Stockpile Stewardship Program

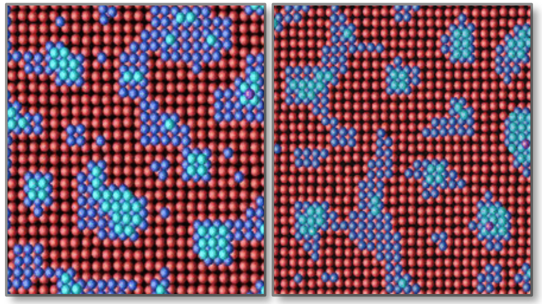

Nuclear Physics (NP)

Nuclear binding and formation of heavy elements in the universe; neutrino oscillation

2015 Long Range Plan for Nuclear Science; FRIB, DUNE, Majorana, EXO experiments

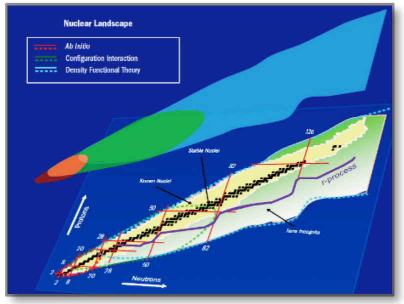

Bioenergy (BER, BES, EERE)

Predictive understanding of lignocellulosic biomass deconstruction into biofuels and bioproducts

MGI; Climate Action Plan

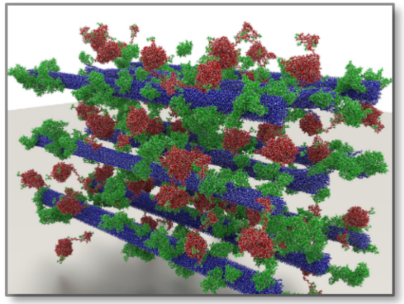

Exascale Applications Will Address National Challenges

Materials Science (BES)

Find, predict, and control materials and properties: property change due to hetero-interfaces and complex structures

MGI

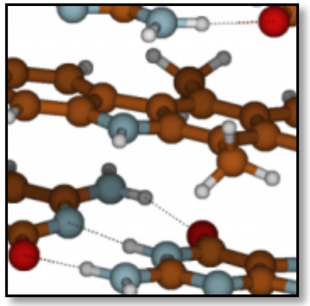

Systems Biology (BER)

Systems model of biological data and molecular mechanisms responsible for complex biological systems

Cancer Moonshot; BRAIN Initiative; Climate Action Plan

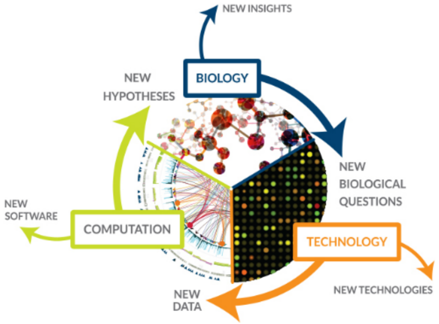

Chemical Science (BES, BER)

Biofuel catalysts design; stress-resistant crops

Climate Action Plan; MGI

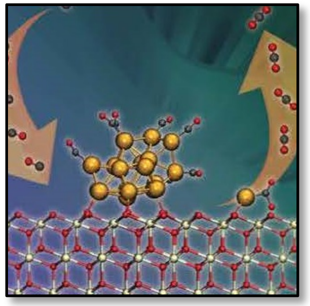

Power Grid (EERE, OE)

Reliably and efficiently planning our nation's grid for societal drivers: rapidly increasing renewable energy penetration, more active consumers

Grid Modernization Initiative; Climate Action Plan

Materials Science (BES)

Protein structure and dynamics; 3D molecular structure design of engineering functional properties

MGI; LCLS-II 2025 Path Forward

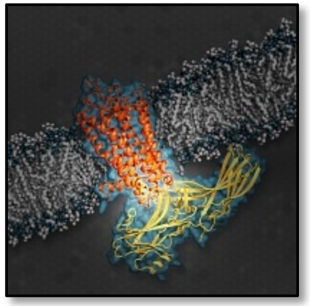

Application Co-Design (CD)
Scope

- Preferentially target crosscutting algorithmic methods that capture the most common patterns of computation and communication (referred to as *motifs*) in ECP applications
- Key application motifs will be addressed within the CD activity, with the goal being to develop software components that characterize the motifs' functionality. These co-designed components will then be integrated into the respective application software environments for testing, use, and requirements feedback
- The CD activity will be executed by a number of CD Centers, with each center focusing on a unique collection of algorithmic motifs needed by multiple (ideally two or more) applications of strategic interest to (e.g., within the scope of) the ECP

Application Motifs*

Algorithmic methods that capture a common pattern of computation and communication

1. **Dense Linear Algebra**
   – Dense matrices or vectors (e.g., BLAS Level 1/2/3)
2. **Sparse Linear Algebra**
   – Many zeros, usually stored in compressed matrices to access nonzero values (e.g., Krylov solvers)
3. **Spectral Methods**
   – Frequency domain, combining multiply-add with specific patterns of data permutation with all-to-all for some stages (e.g., 3D FFT)
4. **N-Body Methods (Particles)**
   – Interaction between many discrete points, with variations being particle-particle or hierarchical particle methods (e.g., PIC, SPH, PME)
5. **Structured Grids**
   – Regular grid with points on a grid conceptually updated together with high spatial locality (e.g., FDM-based PDE solvers)
6. **Unstructured Grids**
   – Irregular grid with data locations determined by app and connectivity to neighboring points provided (e.g., FEM-based PDE solvers)
7. **Monte Carlo**
   – Calculations depend upon statistical results of repeated random trials
8. **Combinational Logic**
   – Simple operations on large amounts of data, often exploiting bit-level parallelism (e.g., Cyclic Redundancy Codes or RSA encryption)
9. **Graph Traversal**
   – Traversing objects and examining their characteristics, e.g., for searches, often with indirect table lookups and little computation
10. **Graphical Models**
   – Graphs representing random variables as nodes and dependencies as edges (e.g., Bayesian networks, Hidden Markov Models)
11. **Finite State Machines**
   – Interconnected set of states (e.g., for parsing); often decomposed into multiple simultaneously active state machines that can act in parallel
12. **Dynamic Programming**
   – Computes solutions by solving simpler overlapping subproblems, e.g., for optimization solutions derived from optimal subproblem results
13. **Backtrack and Branch-and-Bound**
   – Solving search and global optimization problems for intractably large spaces where regions of the search space with no interesting solutions are ruled out. Use the divide and conquer principle: subdivide the search space into smaller subregions (“branching”), and bounds are found on solutions contained in each subregion under consideration

The Landscape of Parallel Computing Research: A View from Berkeley, Technical Report No. UCB/EECS-2006-183 (Dec 2006).

Survey of Application Motifs

Application	Monte Carlo	Particles	Sparse Linear Algebra	Dense Linear Algebra	Spectral Methods	Unstructured Grid	Structured Grid	Comb. Logic	Graph Traversal	Dynamical Program	Backtrack & Branch and Bound	Graphical Models	Finite State Machine
Cosmology		■	■		■		■		■				
Subsurface			■			■	■						
Materials (QMC)	■	■	■	■	■					■			
Additive Manufacturing		■	■	■	■	■							
Chemistry for Catalysts & Plants	■		■	■	■		■		■				
Climate Science		■	■			■	■						
Precision Medicine Machine Learning			■	■				■	■	■	■	■	
QCD for Standard Model Validation	■		■	■	■		■						
Accelerator Physics		■	■		■		■						
Nuclear Binding and Heavy Elements	■		■	■									
MD for Materials Discovery & Design		■	■	■	■		■						
Magnetically Confined Fusion		■	■		■	■	■						

Application "Predictivity": A Step Change at Exascale?

Predictability Capability Maturity Model . . . developed by SNL over the past 15 years . . .

Increasing completeness and rigor →

Content ↓

MATURITY / ELEMENT	Maturity Level 0 Low Consequence, Minimal M&S Impact, e.g. Scoping Studies	Maturity Level 1 Moderate Consequence, Some M&S Impact, e.g. Design Support	Maturity Level 2 High-Consequence, High M&S Impact, e.g. Qualification Support	Maturity Level 3 High-Consequence, Decision-Making Based on M&S, e.g. Qualification or Certification
Representation and Geometric Fidelity What features are neglected because of simplifications or stylizations?	• Judgment only • Little or no representational or geometric fidelity for the system and BCs	• Significant simplification or stylization of the system and BCs • Geometry or representation of major components is defined	• Limited simplification or stylization of major components and BCs • Geometry or representation is well defined for major components and some minor components • Some peer review conducted	• Essentially no simplification or stylization of components in the system and BCs • Geometry or representation of all components is at the detail of "as built", e.g., gaps, material interfaces, fasteners • Independent peer review conducted
Physics and Material Model Fidelity How fundamental are the physics and material models and what is the level of model calibration?	• Judgment only • Model forms are either unknown or fully empirical • Few, if any, physics-informed models • No coupling of models	• Some models are physics based and are calibrated using data from related systems • Minimal or ad hoc coupling of models	• Physics-based models for all important processes • Significant calibration needed using separate effects tests (SETs) and integral effects tests (IETs) • One-way coupling of models • Some peer review conducted	• All models are physics based • Minimal need for calibration using SETs and IETs • Sound physical basis for extrapolation and coupling of models • Full, two-way coupling of models • Independent peer review conducted
Code Verification Are algorithm deficiencies, software errors, and poor SQE practices corrupting the simulation results?	• Judgment only • Minimal testing of any software elements • Little or no SQE procedures specified or followed	• Code is managed by SQE procedures • Unit and regression testing conducted • Some comparisons made with benchmarks	• Some algorithms are tested to determine the observed order of numerical convergence • Some features & capabilities (F&C) are tested with benchmark solutions • Some peer review conducted	• All important algorithms are tested to determine the observed order of numerical convergence • All important F&Cs are tested with rigorous benchmark solutions • Independent peer review conducted
Solution Verification Are numerical solution errors and human procedural errors corrupting the simulation results?	• Judgment only • Numerical errors have an unknown or large effect on simulation results	• Numerical effects on relevant SRQs are qualitatively estimated • Input/output (I/O) verified only by the analysts	• Numerical effects are quantitatively estimated to be small on some SRQs • I/O independently verified • Some peer review conducted	• Numerical effects are determined to be small on all important SRQs • Important simulations are independently reproduced • Independent peer review conducted
Model Validation How carefully is the accuracy of the simulation and experimental results assessed at various tiers in a validation hierarchy?	• Judgment only • Few, if any, comparisons with measurements from similar systems or applications	• Quantitative assessment of accuracy of SRQs not directly relevant to the application of interest • Large or unknown exper-imental uncertainties	• Quantitative assessment of predictive accuracy for some key SRQs from IETs and SETs • Experimental uncertainties are well characterized for most SETs, but poorly known for IETs • Some peer review conducted	• Quantitative assessment of predictive accuracy for all important SRQs from IETs and SETs at conditions/geometries directly relevant to the application • Experimental uncertainties are well characterized for all IETs and SETs • Independent peer review conducted
Uncertainty Quantification and Sensitivity Analysis How thoroughly are uncertainties and sensitivities characterized and propagated?	• Judgment only • Only deterministic analyses are conducted • Uncertainties and sensitivities are not addressed	• Aleatory and epistemic (A&E) uncertainties propagated, but without distinction • Informal sensitivity studies conducted • Many strong UQ/SA assumptions made	• A&E uncertainties segregated, propagated and identified in SRQs • Quantitative sensitivity analyses conducted for most parameters • Numerical propagation errors are estimated and their effect known • Some strong assumptions made • Some peer review conducted	• A&E uncertainties comprehensively treated and properly interpreted • Comprehensive sensitivity analyses conducted for parameters and models • Numerical propagation errors are demonstrated to be small • No significant UQ/SA assumptions made • Independent peer review conducted

ECP Application Development

To create or enhance high-performance computing applications and train researchers to effectively use Exascale systems

Mission need

Create and/or enhance important DOE applications through development of models, algorithms, and methods; integration of software and hardware using co-design methodologies; systematic improvement of exascale system readiness and utilization; and demonstration and assessment of effective software and hardware integration

Objective

Deliver a broad array of comprehensive science-based computational applications that effectively exploit exascale HPC technology to provide breakthrough modeling and simulation solutions, yielding high-confidence insights and answers to the nation's most critical problems and challenges in scientific discovery, energy assurance, economic competitiveness, health enhancement, and national security

Questions?

EXASCALE COMPUTING PROJECT