CASL: The Consortium for Advanced Simulation of Light Water Reactors

A U.S. Department of Energy Innovation Hub for Modeling and Simulation of Nuclear Reactors

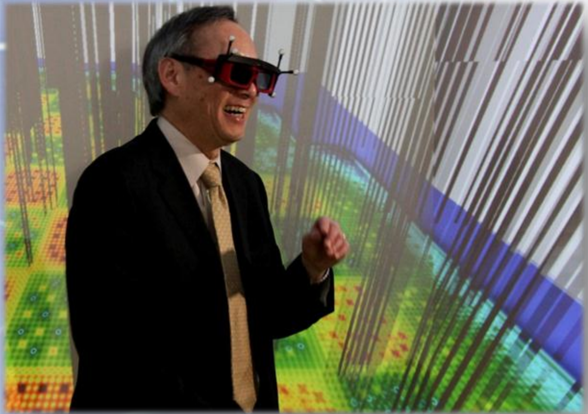

John A. Turner
Oak Ridge National Laboratory
Computer Science & Math Division
Virtual Reactor Integration Lead, CASL

NVIDIA GPU Tech. Conf.
San Jose, CA
May 15, 2012

U.S. DEPARTMENT OF ENERGY | Nuclear Energy

What is a DOE Innovation Hub?

- **04/06/2009**: Secretary Chu proposes 8 Energy Innovation Hubs (idea pre-dates Chu)
  - modeled after research entities like the Manhattan Project (nuclear weapons), Lincoln Lab at MIT (radar), and AT&T Bell Labs (transistor)
    - highly-integrated & collaborative teams - solve priority technology challenges to national climate and energy goals
    - problems that have proven the most resistant to solution via the normal R&D enterprise
    - focused, spanning spectrum from basic research through engineering development to partnering with industry in commercialization
    - bring together expertise across the R&D enterprise (gov, academia, industry, non-profits)
  - $25M per yr for 5 years, with possible 5-yr extension
- **06/25/2009**: House bill did not approve any of the proposed Hubs
  - $35M in Basic Energy Sciences for the Secretary to select one Hub
- **07/09/2009**: Senate approves 3 of the proposed hubs, but at $22M
  - Fuels from sunlight (in EERE)
  - Energy efficient building systems (in EERE)
  - Modeling and simulation for nuclear energy systems (in NE)
- **10/01/2009**: Final bill out of conference matches Senate bill
- **01/20/2010**: FOA released, proposals due 03/08/2010
- **05/27/2010**: CASL selected, first funding arrived 07/01/2010

CASL

The Consortium for Advanced Simulation of Light Water Reactors (CASL)

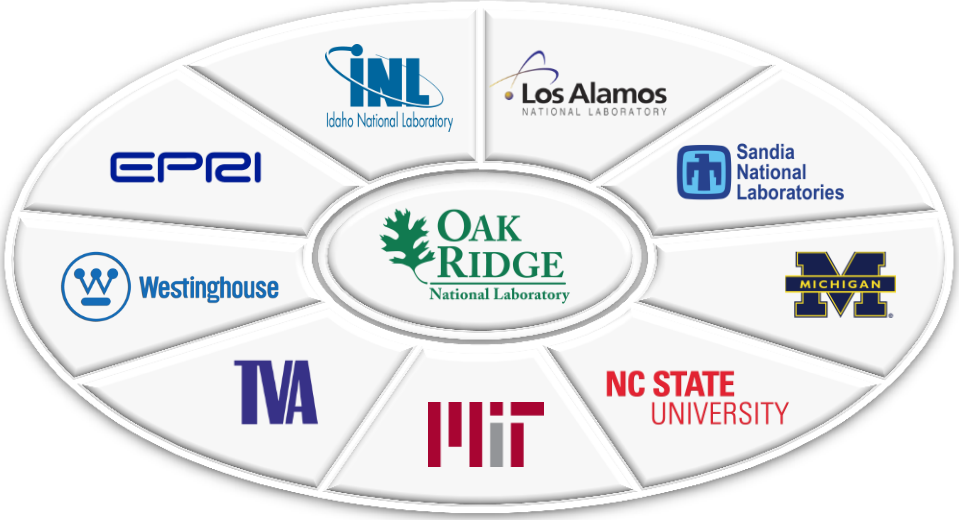

The Consortium for Advanced Simulation of Light Water Reactors (CASL)

Core partners

Oak Ridge National Laboratory
Idaho National Laboratory
Sandia National Laboratories
Los Alamos National Laboratory

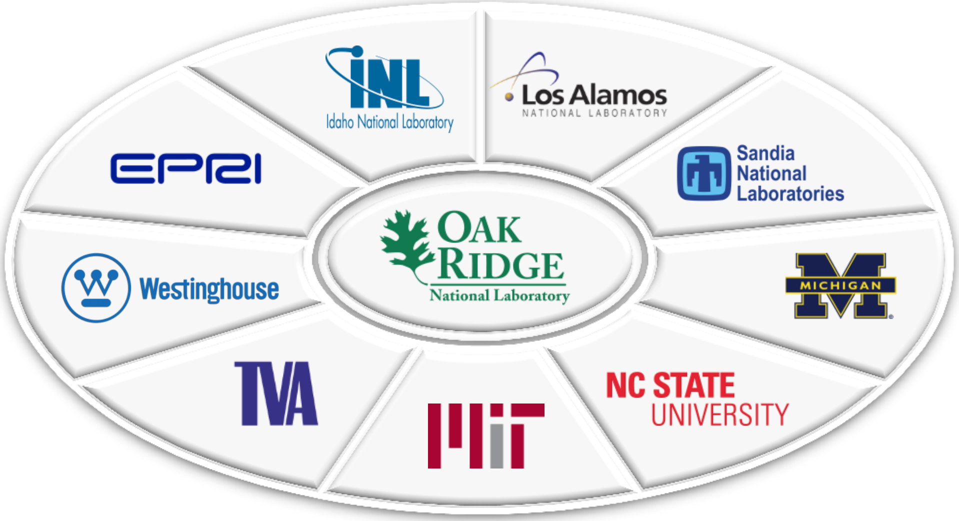

The Consortium for Advanced Simulation of Light Water Reactors (CASL)

Core partners

Oak Ridge National Laboratory

Idaho National Laboratory

Sandia National Laboratories

Los Alamos National Laboratory

University of Michigan

North Carolina State University

Massachusetts Institute of Technology

The Consortium for Advanced Simulation of Light Water Reactors (CASL)

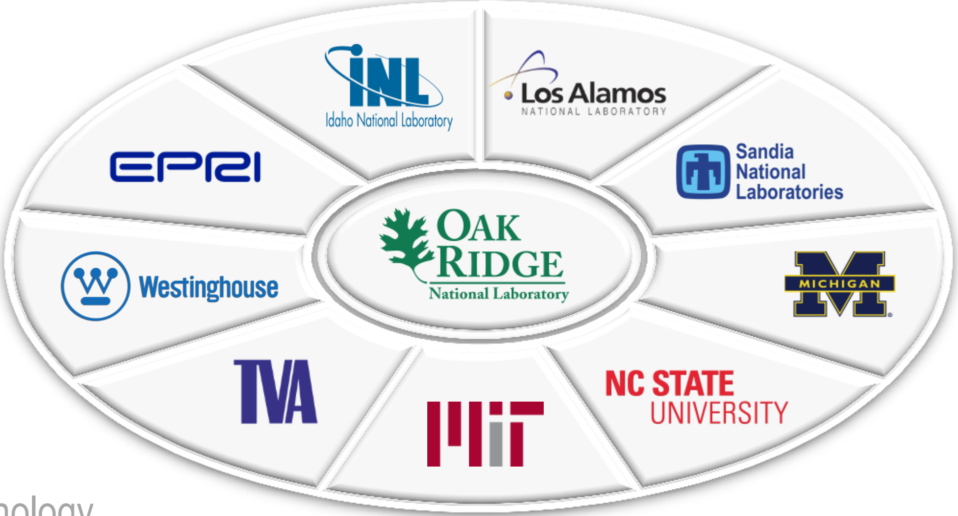

Core partners

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

The Consortium for Advanced Simulation of Light Water Reactors (CASL)

Core partners

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

The Consortium for Advanced Simulation of Light Water Reactors (CASL)

Core partners

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

Challenges

- High visibility
- Geographically-dispersed
- Diversity of experience
- Wide range of motivation / priorities
- Proprietary codes and data
- Role of commercial codes
- Export control

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

Nuclear Energy Overview

Source: Nuclear Energy Institute (NEI)

- World nuclear power generating capacity
  - 439 plants (U.S.- 104 plants in 31 states)
  - 373 GWe (U.S.- 100.7 GWe, 798.7 TWh in 2009)
  - ~90% capacity factor (>6 GWe added to grid)
- U.S. electricity from nuclear: 20.2%
  - One uranium fuel pellet provides as much energy as:
    - one ton of coal
    - 149 gallons of oil
    - 17,000 cubic feet of natural gas
- U.S. electricity demand projected to grow 25% by 2030
  - 2007: 3.99 TWh
  - 2030: 4.97 TWh
- nuclear accounts for 73% of emission-free electricity in US

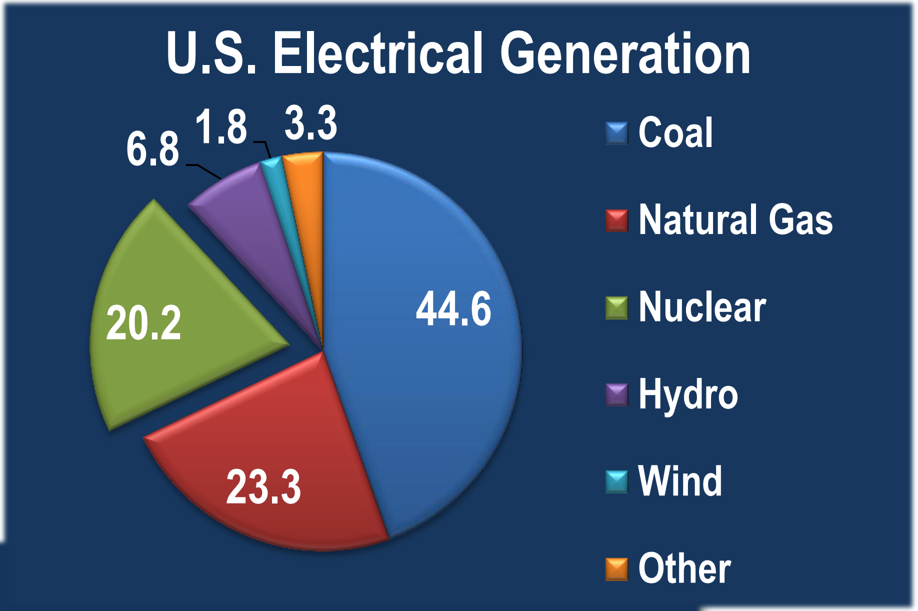

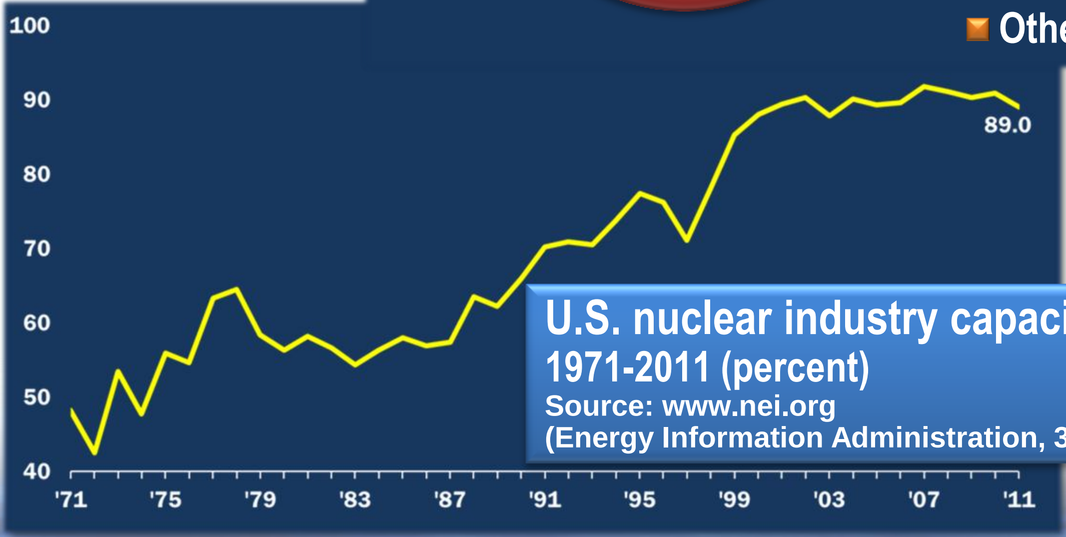

U.S. nuclear industry capacity factors 1971-2011 (percent)
Source: www.nei.org
(Energy Information Administration, 3/12)

Anatomy of a Nuclear Reactor

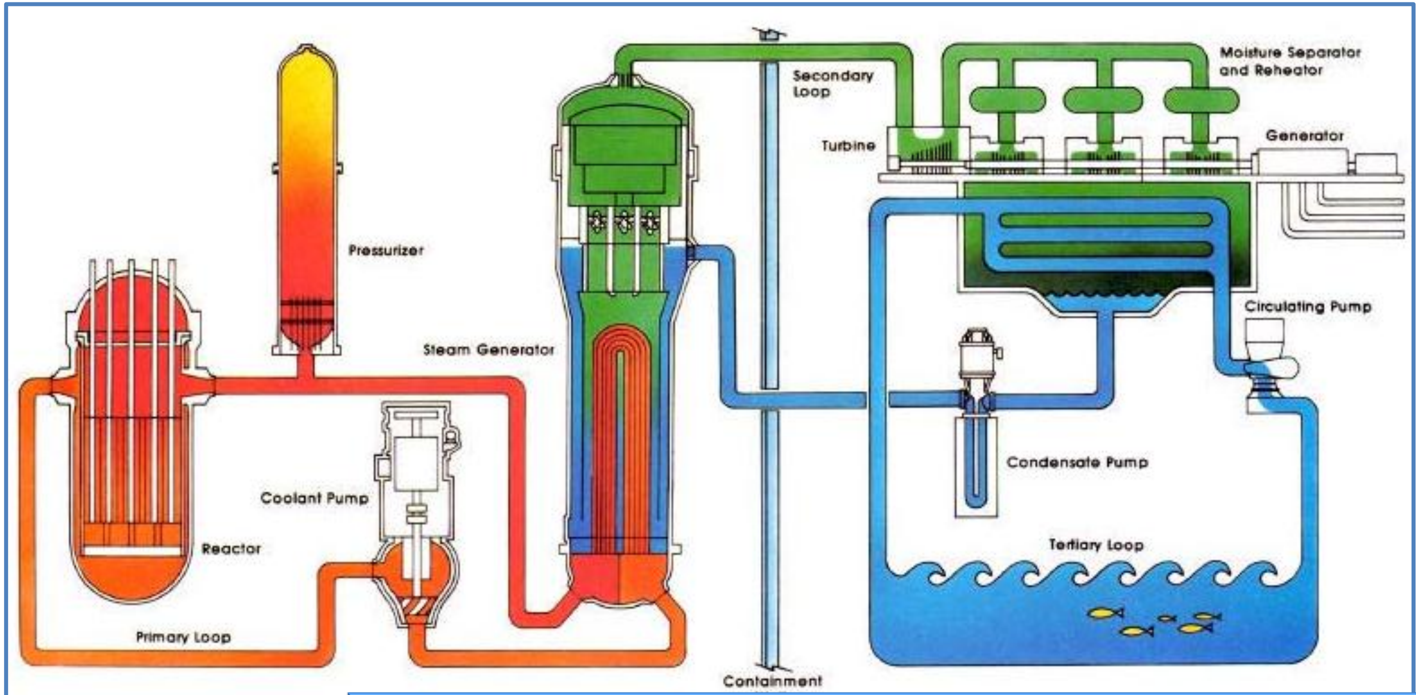

Example: Westinghouse 4-Loop Pressurized Water Reactor (PWR)

Power: ~1170 MWe (~3400 MWth)
Containment Building: 115’ diameter x 156’ high steel / concrete
Pressure Vessel: 14.4’ diameter x 41.3’ high x 0.72’ thick alloy steel
Coolant: pressurized water (2250 psia), T_{in} ~ 545°F, T_{out} ~ 610°F, 134M lb/h (4 pumps)

CASL

Anatomy of a Nuclear Reactor

Example: Westinghouse 4-Loop Pressurized Water Reactor (PWR)

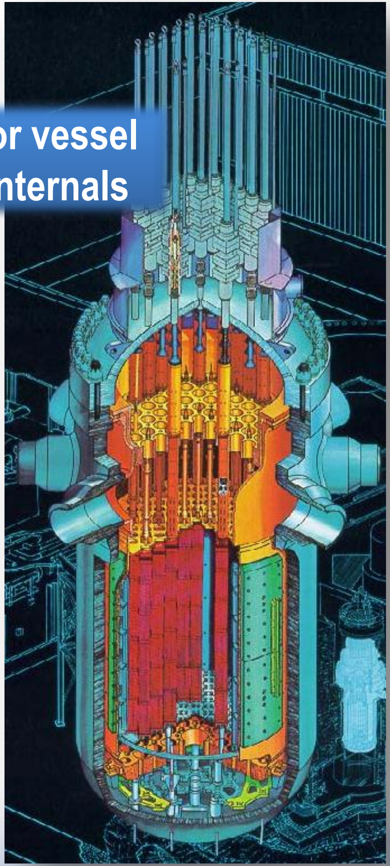

reactor vessel and internals

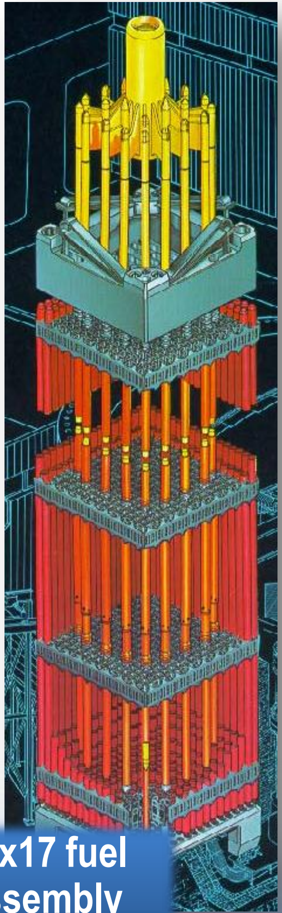

17x17 fuel assembly

Core
- 11.1’ diameter x 12’ high
- 193 fuel assemblies
- 107.7 tons of UO_2 (~3-5% U_{235})

Fuel Assemblies
- 17x17 pin lattice (14.3 mm pitch)
- 204 pins per assembly

Fuel Pins
- ~300-400 pellets stacked within 12’ high x 0.61 mm thick Zr-4 cladding tube

Fuel Pellets
- 9.29 mm diameter x ~10.0 mm high

Fuel Temperatures
- 4140° F (max centerline)
- 657° F (max clad surface)

~51,000 fuel pins and over 16M fuel pellets in the core of a PWR

CASL mission is to improve reactor performance
(initially currently-operating LWRs)

Power uprates

- 5–7 GWe delivered at ~20% of new reactor cost
- Advances in M&S needed to enable further uprates (up to 20 GWe)
- **Key concerns:**
  - Damage to structures, systems, and components (SSC)
  - Fuel and steam generator integrity
  - Violation of safety limits

Lifetime extension

- Reduces cost of electricity
- Essentially expands existing nuclear power fleet
- Requires ability to predict structures, systems, and components aging and life-cycle management
- **Key concerns:**
  - Effects of increased radiation and aging on integrity of reactor vessel and internals
  - Ex-vessel performance (effects of aging on containment and piping)
  - Significant financial decisions to support operation beyond 60 years must be made in ~5 yrs

Higher burnup

- Supports reduction in amount of used nuclear fuel
- Supports uprates by avoiding need for additional fuel
- **Key concerns:**
  - Cladding integrity
  - Fretting
  - Corrosion/ CRUD
  - Hydriding
  - Creep
  - Fuel-cladding mechanical interactions

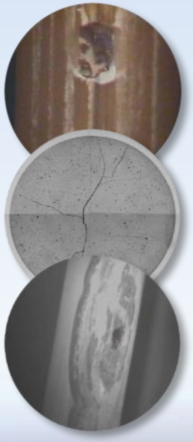

CASL

CASL Challenge Problems

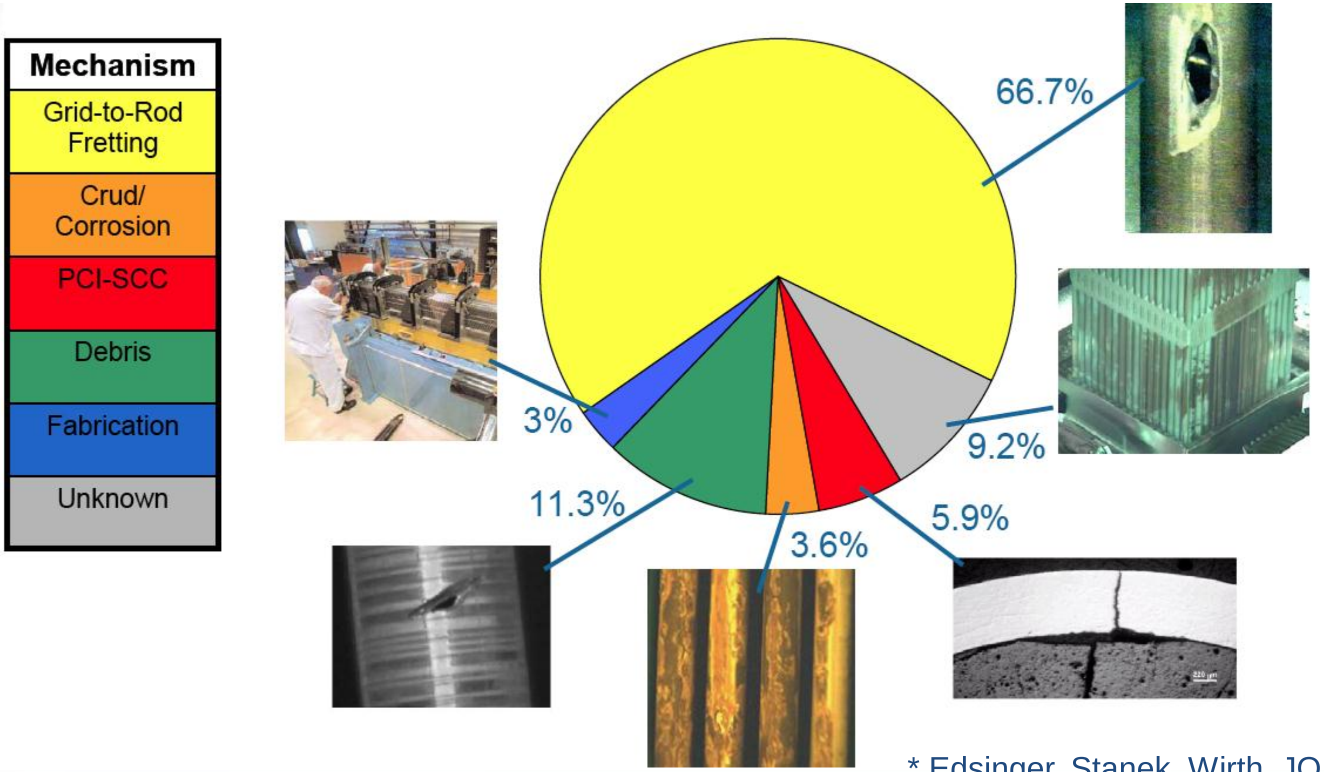

Summary of US fuel failure mechanisms (2000-2008)

* Edsinger, Stanek, Wirth, JOM **63, no. 8** (2011)

Fuel failure modes provide motivation for CASL activities

Grid-to-Rod-Fretting (GTRF)

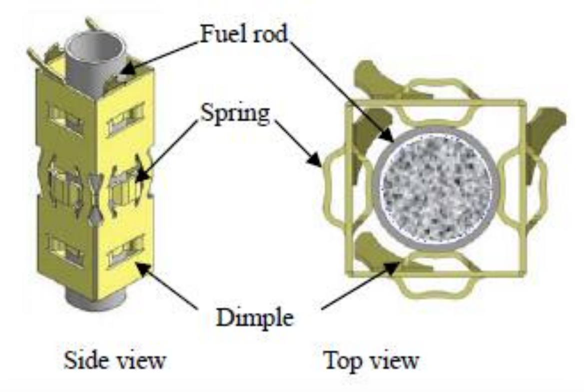

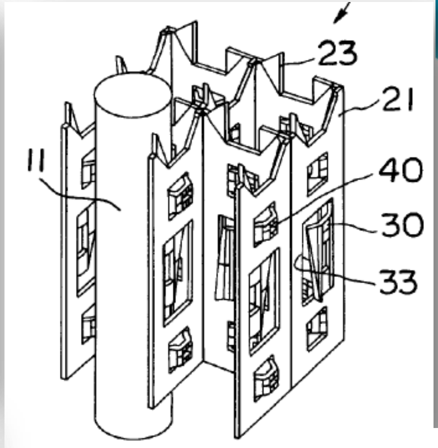

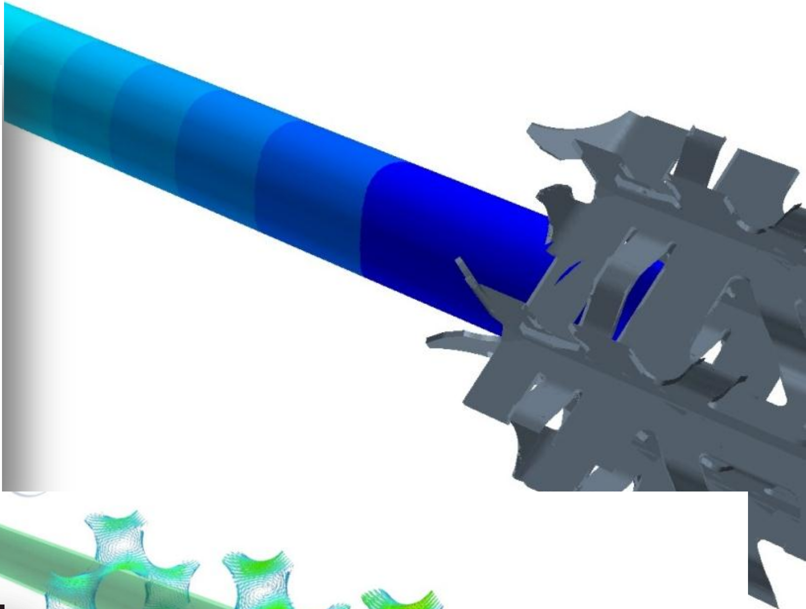

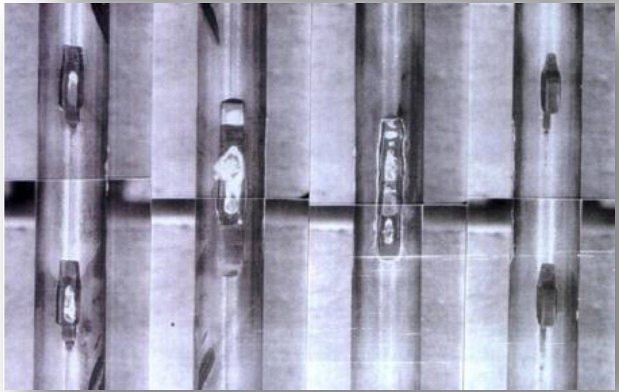

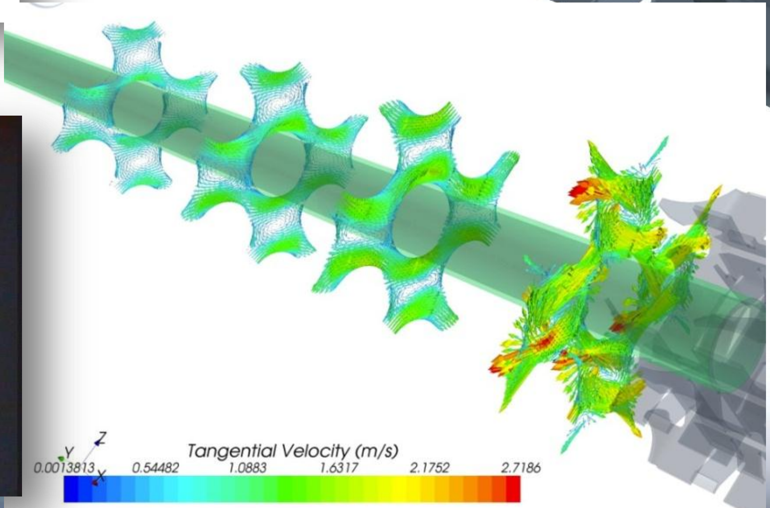

CRUD-induced power shift (CIPS)

- deviation in axial power shape
  - Cause: boron uptake in CRUD deposits
    in high power density regions with subcooled boiling
  - affects fuel management and thermal margin in many plants
- power uprates will increase potential for CRUD growth

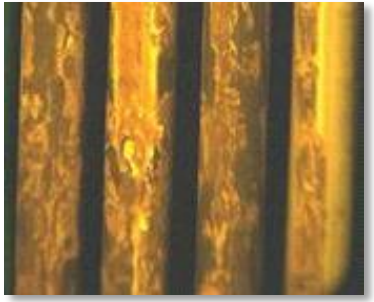

CRUD deposits

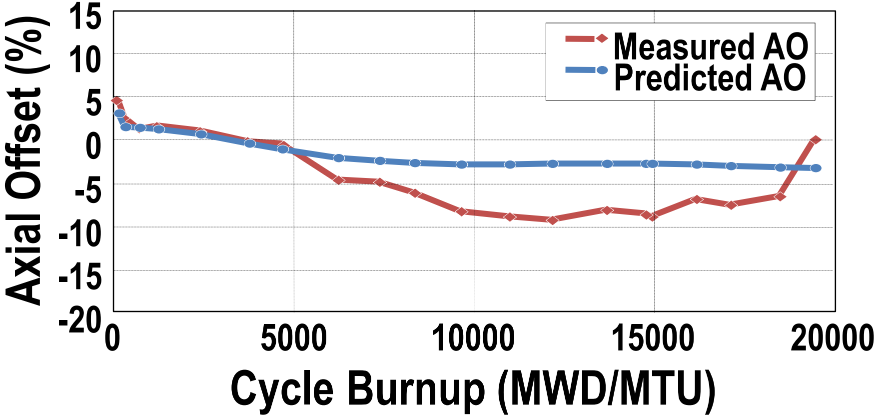

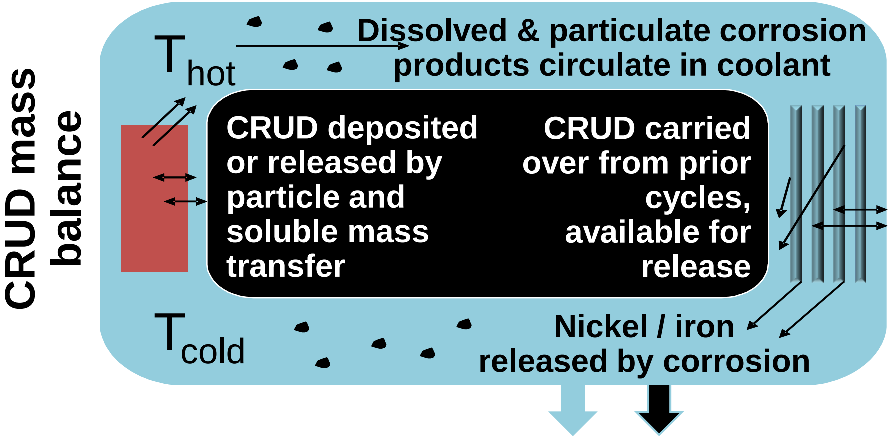

Need: Multi-physics chemistry, flow, and neutronics model to predict CRUD growth

CASL

Virtual Environment for Reactor Applications (VERA)

A suite of tools for scalable simulation of nuclear reactor core behavior

- Flexible coupling of physics components
- Toolkit of components
  - Not a single executable
  - Both legacy and new capability
  - Both proprietary and distributable

- Attention to usability
- Rigorous software processes
- Fundamental focus on V&V and UQ

- Development guided by relevant challenge problems
- Broad applicability

- Scalable from high-end workstations to existing and future HPC platforms
  - Diversity of models, approximations, algorithms
  - Architecture-aware implementations

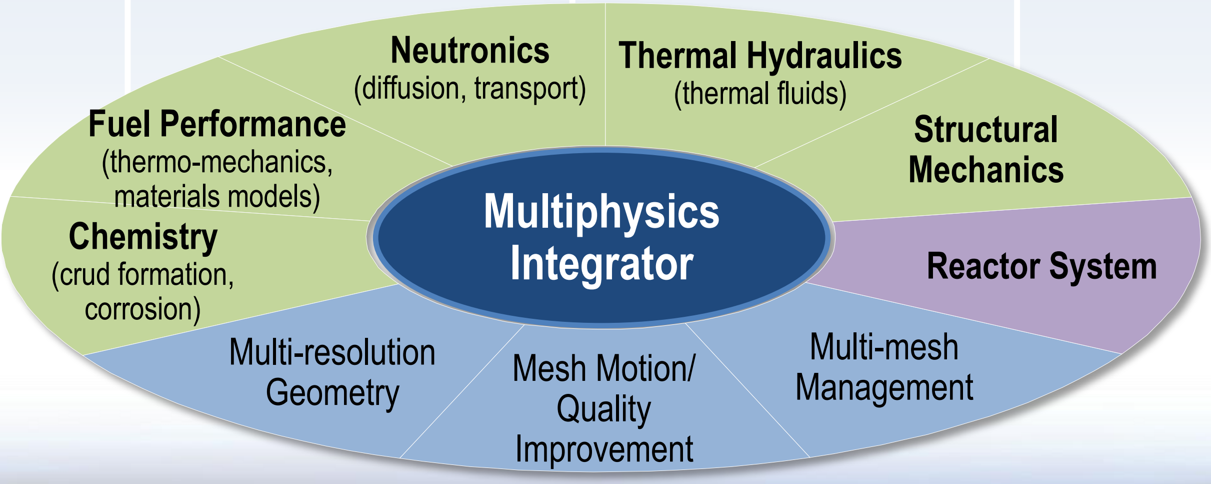

Lightweight Integrating Multiphysics Environment (LIME)

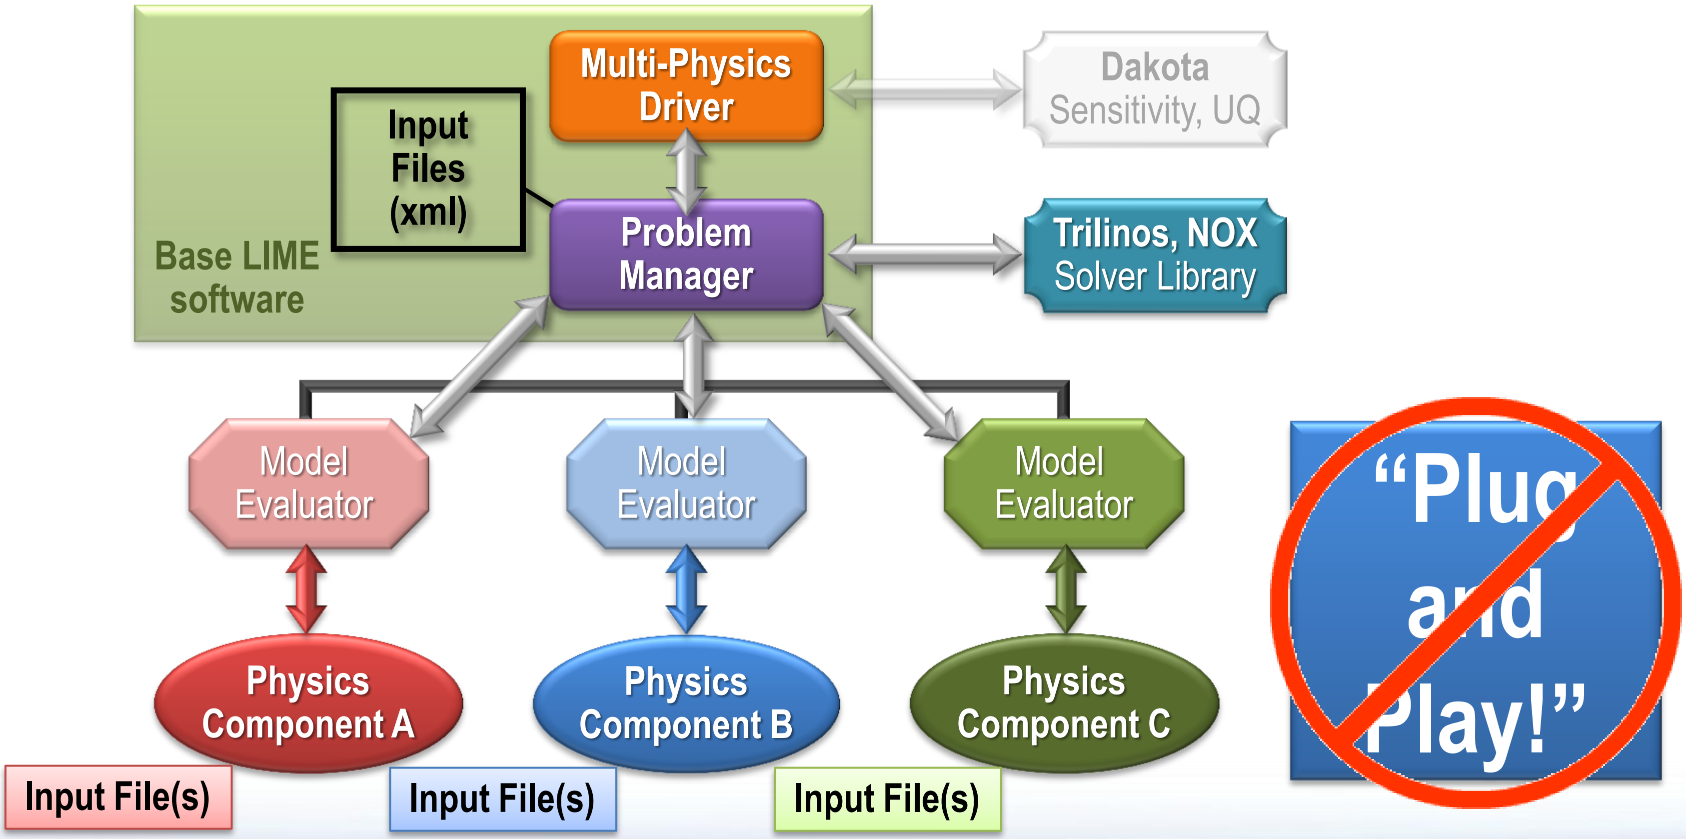

https://sourceforge.net/projects/lime1/

Writing software is easy

- "Writing songs is easy. Writing great songs is hard."
  - Bono (? couldn't verify)
- Writing software is easy. Writing great software is hard.

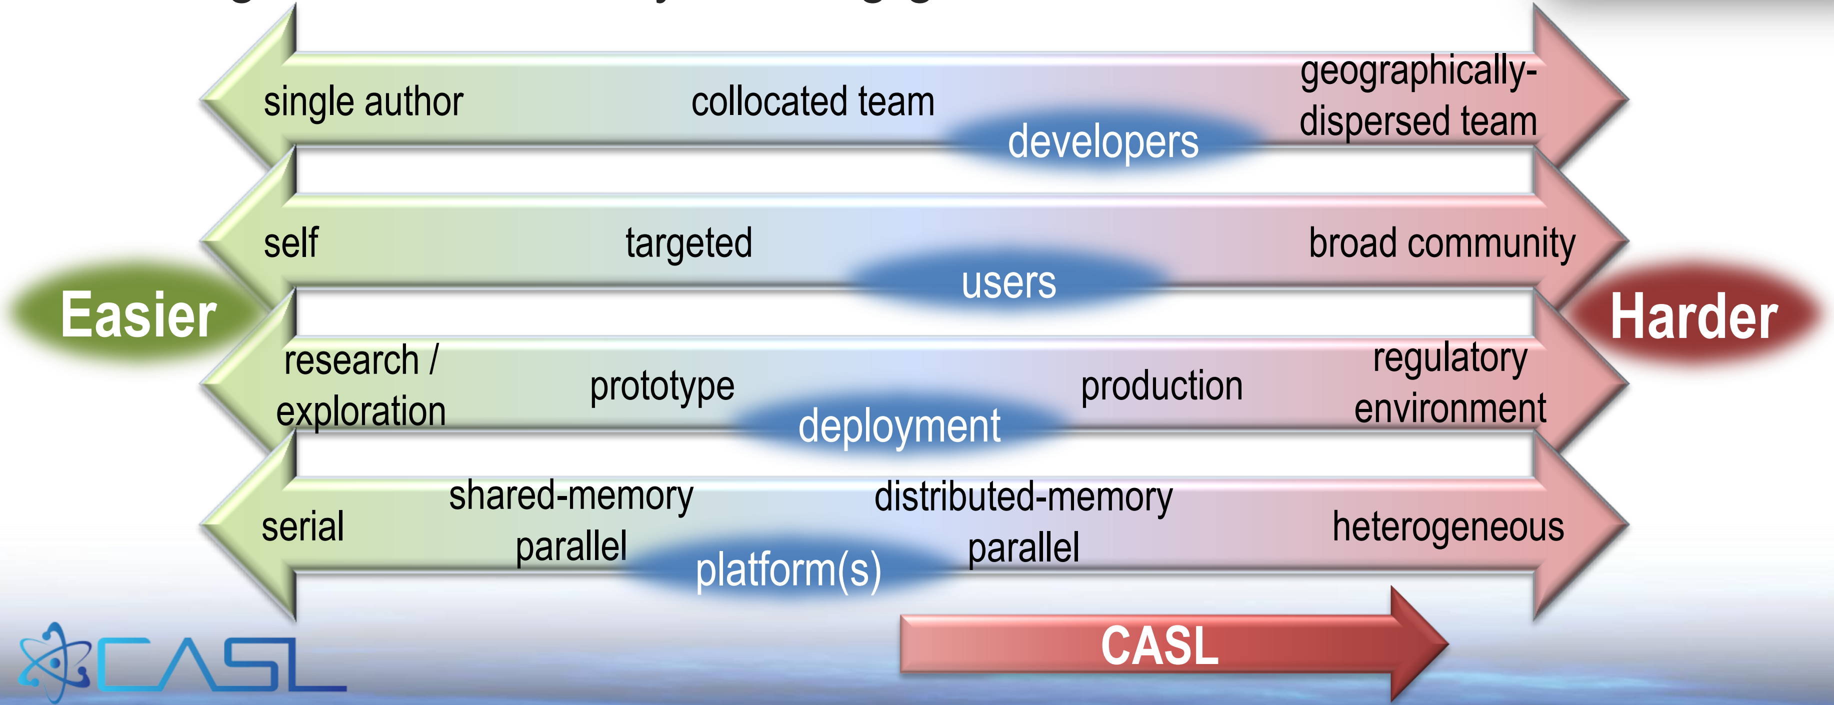

CFD is required for several challenge problems (GTRF, CRUD/CIPS) - remainder of presentation focuses on neutronics…

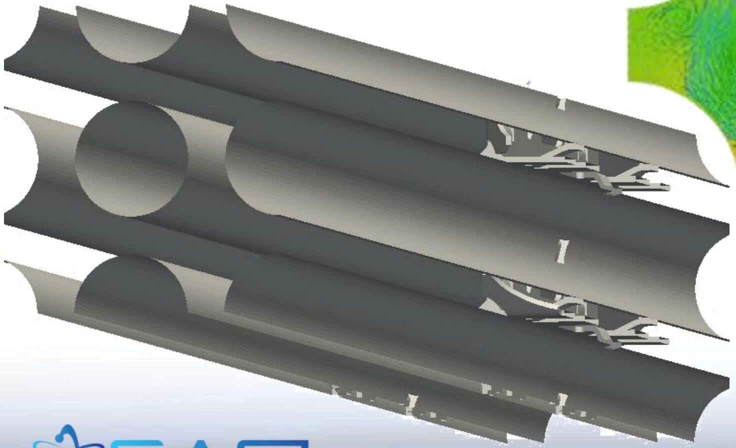

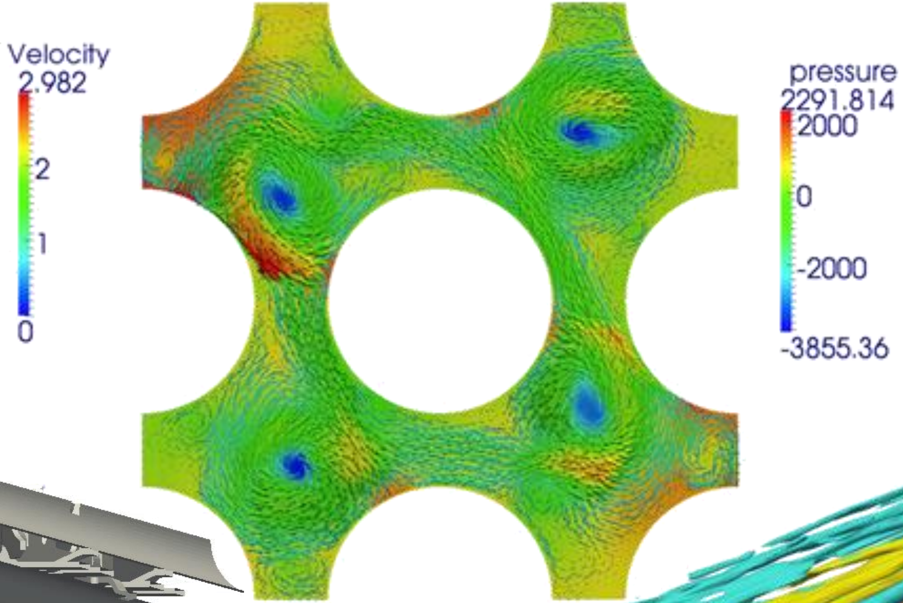

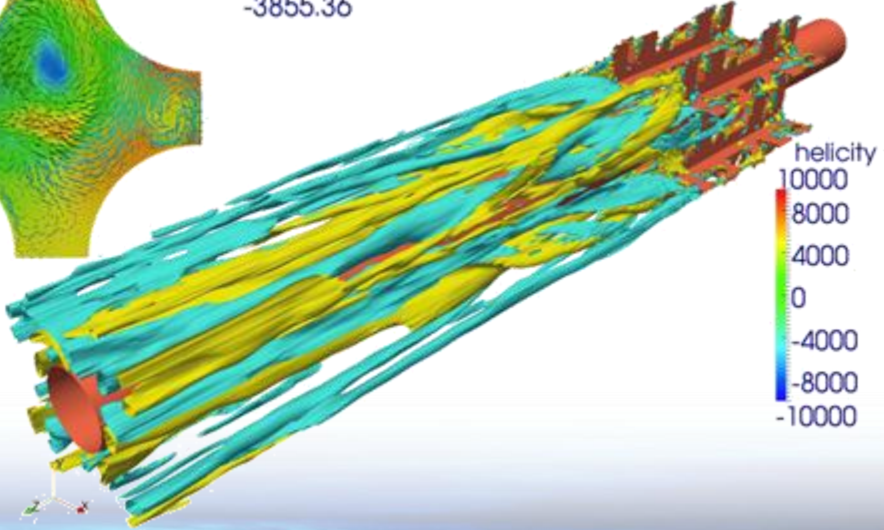

Discrete Ordinates Methods for Neutron Behavior

- We solve the first-order form of the transport equation:
  - Eigenvalue form for multiplying media (fission):

$$\hat{\mathbf{\Omega}} \cdot \nabla\psi(\mathbf{r},\mathbf{\Omega},E) + \Sigma(\mathbf{r},E,T)\psi(\mathbf{r},\mathbf{\Omega},E) = \int dE' \int_{4\pi} d\mathbf{\Omega}'\, \Sigma_{\mathrm{s}}(\mathbf{r},\hat{\mathbf{\Omega}}'\cdot\hat{\mathbf{\Omega}},E' \rightarrow E,T)\psi(\mathbf{r},\mathbf{\Omega}',E') + \frac{1}{k}\frac{\chi(E)}{4\pi}\int dE' \int_{4\pi} d\mathbf{\Omega}'\, \nu\Sigma_{\mathrm{f}}(\mathbf{r},E',T)\psi(\mathbf{r},\mathbf{\Omega}',E')$$

  - T-H coupling comes through the temperature-dependent material cross sections
- Total number of unknowns in solve:
  - $\mathrm{unknowns} = N_g \times N_c \times N_u \times N_a \times N_m$
- An ideal (conservative) estimate.
  - (238) x (1x10^9) x (4) x (288) x (16)

unknowns > 4 x 10^{15}

Eigenvalue Solvers
Power iteration
Arnoldi
Shifted-inverse

↓

Multigroup Solvers
Gauss-Seidel
Residual Krylov
Gauss-Seidel + Krylov

↓

Within-group Solvers
Krylov
Residual Krylov
Source iteration

Current State-of-the-Art in Reactor Neutronics

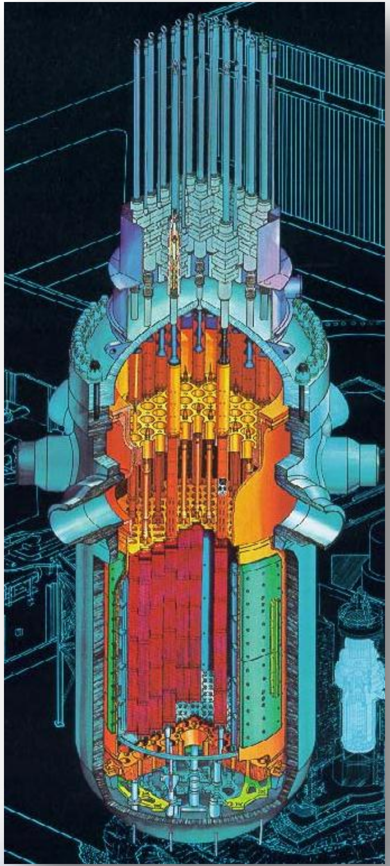

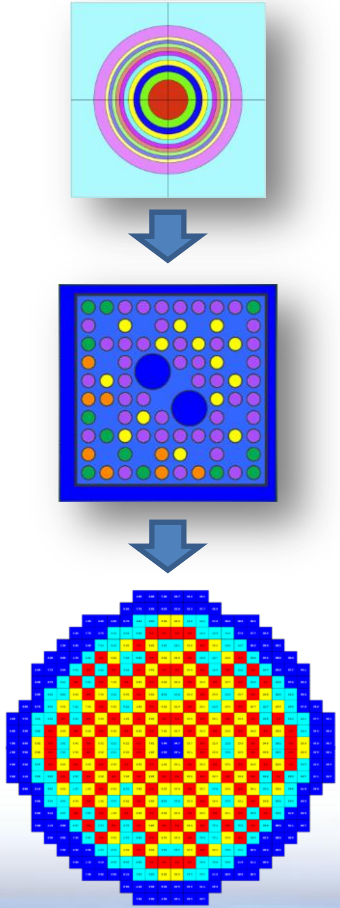

Pin cell (single fuel pin)
- 0/1-D transport
- high energy fidelity ($10^{2\text{-}5}$ unknowns)
- approximate state and BCs

Lattice cell (single assembly)
- 2-D transport
- moderate energy fidelity (7-102 groups)
- approximate state and BCs
- depletion with spectral corrections
- space-energy homogenization

Full core
- 3-D diffusion
- low energy fidelity (2-4 groups)
- homogeneous lattice cells
- heterogeneous flux reconstruction
- coupled physics

Can we approach resolution/fidelity of current 2D analysis in 3D for full core analysis?

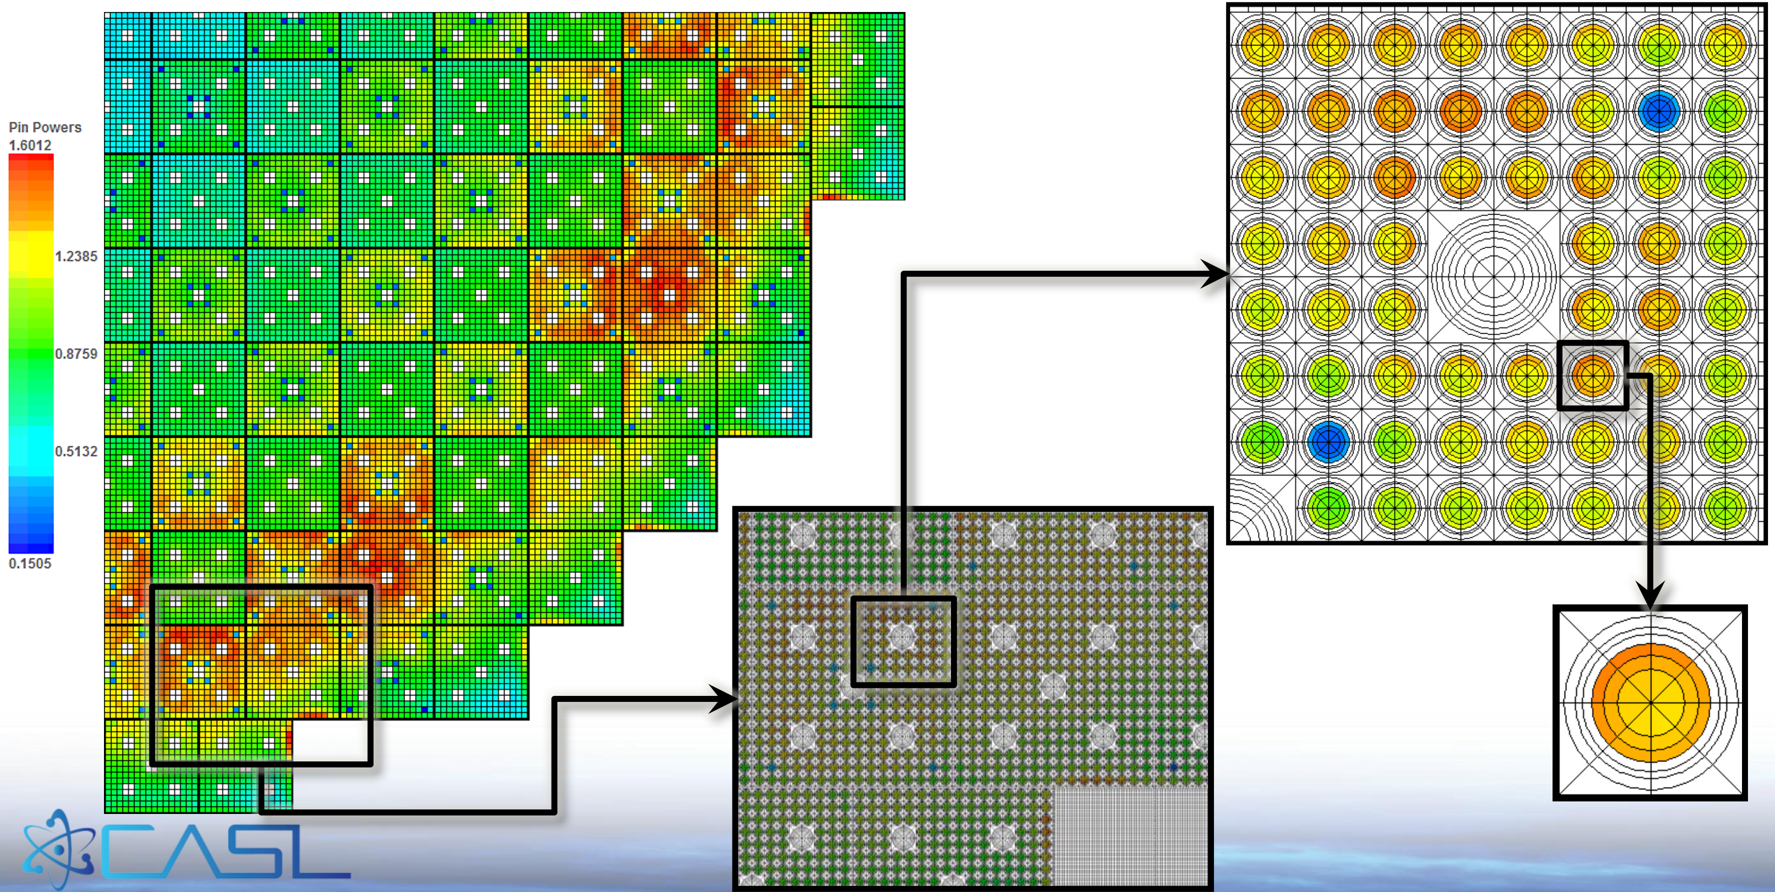

PWR-900 Whole-Core Reactor Problem

- 2 and 44-group, homogenized fuel pins
- 2×2 spatial discretization per fuel pin
- 17×17 fuel pins per assembly
- 289 assemblies
  - 157 fuel, 132 reflector
  - high, med, low enrichments
- Space-angle unknowns:
  - 233,858,800 cells
  - 128 angles (1 moment)
  - 1 spatial unknown per cell

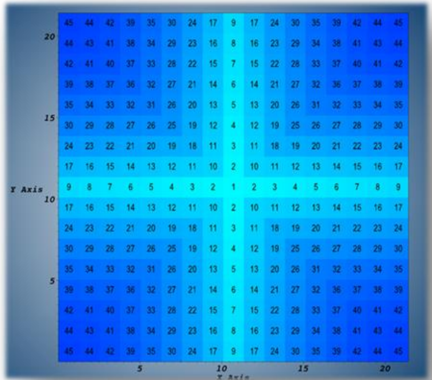

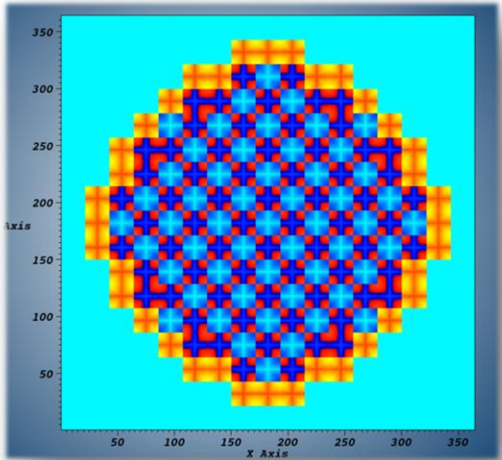

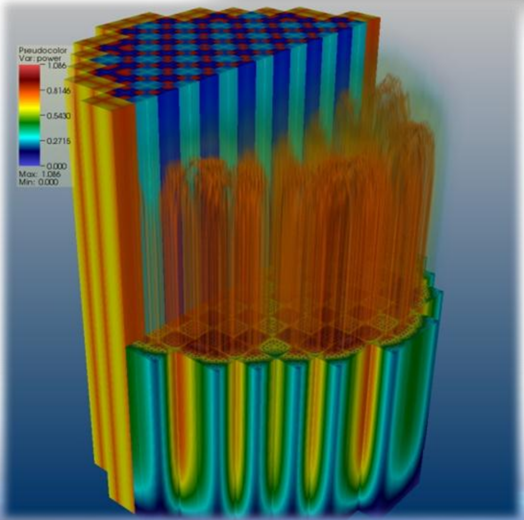

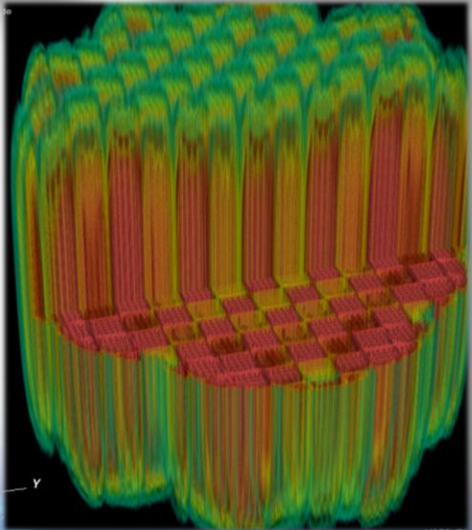

Performance at scaling on ORNL Titan (Cray XK6)

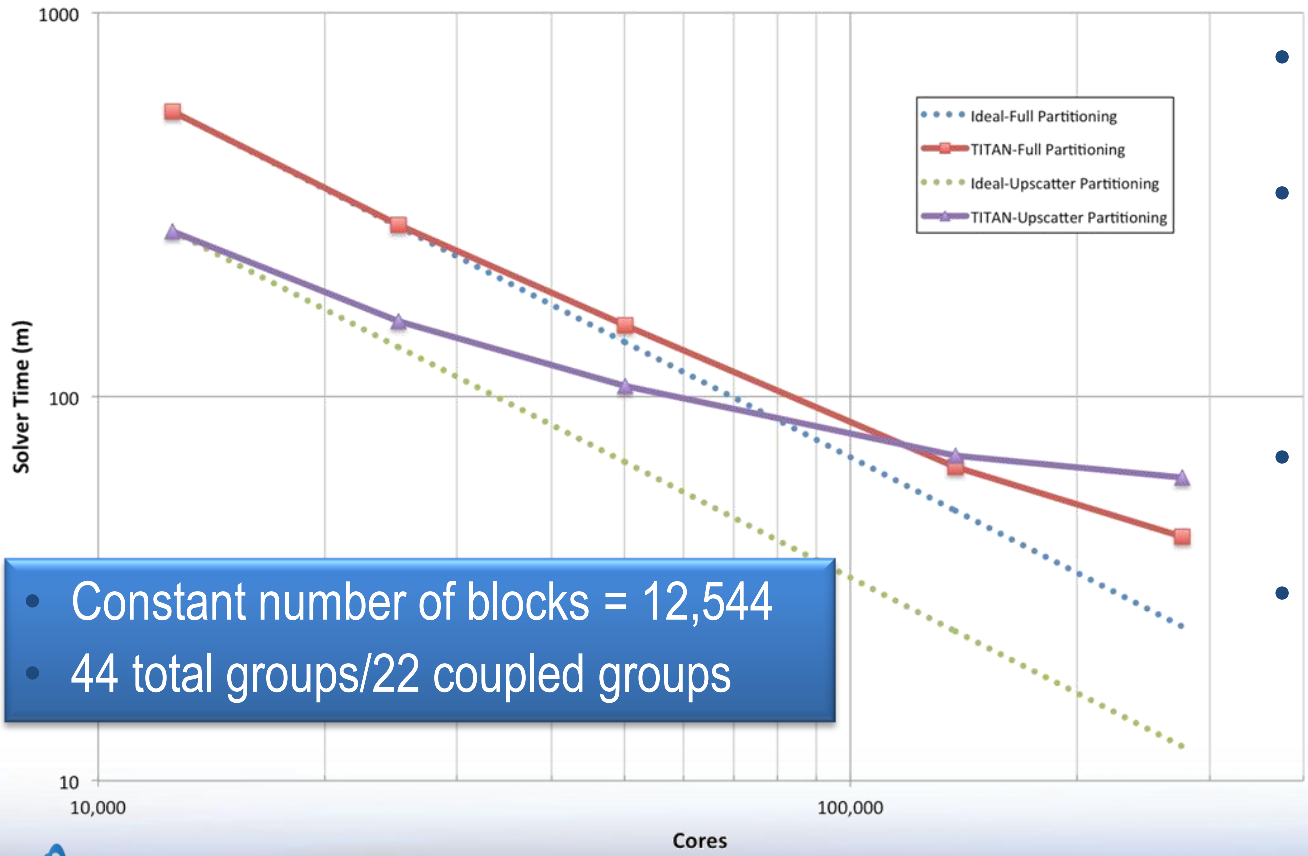

- Constant number of blocks = 12,544
- 44 total groups/22 coupled groups

- full partitioning scales well to 275K cores
- improved interconnect + reduce-scatter have dramatically reduced global reduction cost
- upscatter partitioning more efficient at lower set counts
- roll-over occurs between 4 and 11 sets (5 and 2 groups per set) where serial work in GS solver dominates

CASL

What does this mean?

Where we want to be…	Where we are…
• reproduce fidelity of 2D calculations using consistent 3D methods • produce all state-points for an 18-month depletion cycle in O(8 hours) • O(72) state points per cycle (1 week steps) • steady-state, coupled neutronics simulation with T-H feedback = $O(10^{19})$ unknowns	• assuming 2% peak, we can solve 1.7×10^{13} unknowns/hour (XT5) • we can solve a reduced 3D problem ($O(10^{15})$ unknowns) in 175 hours – assumes status quo on a 1 PF/s XT5 machine

So…

- to reach 2D fidelity at 3D we need to solve $\sim10^4$x more unknowns
- to run all state points in one day at this fidelity using existing code and methods would require ~140 EF/s

CASL

Is it hopeless?

- according to industry partners, a fully-consistent 3D calculation in 1 week would be acceptable
  - factor of 7 (20 EF/s)
- valuable insight possible without reproducing full 2D fidelity
  - factor of 150-200 (100 PF/s)
- utilize GPUs
  - if current projections hold, we can potentially get a factor of 3x-4x improvement by executing sweep kernels on the GPU
- further solver research (multigrid-in-energy) shows promise for reducing iteration counts as well

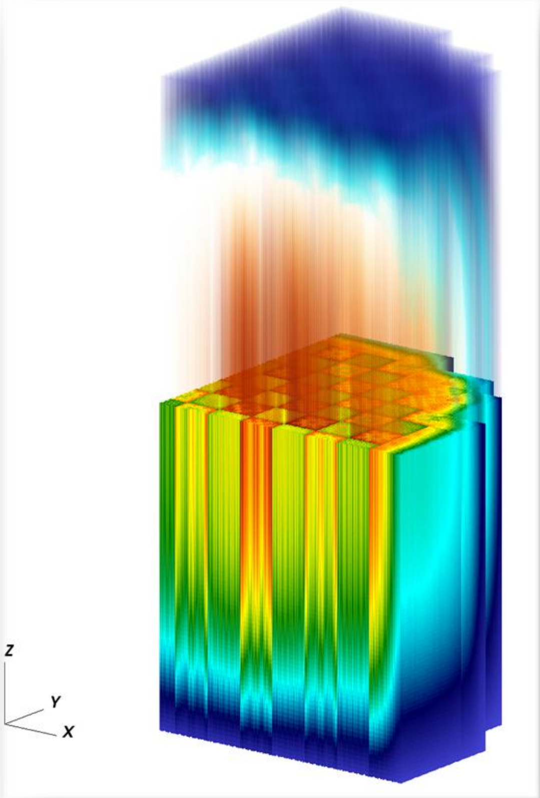

a 30-40 PF/s machine could allow fully-consistent, 3-D neutronics simulations

CASL

GPU Sweep Kernel

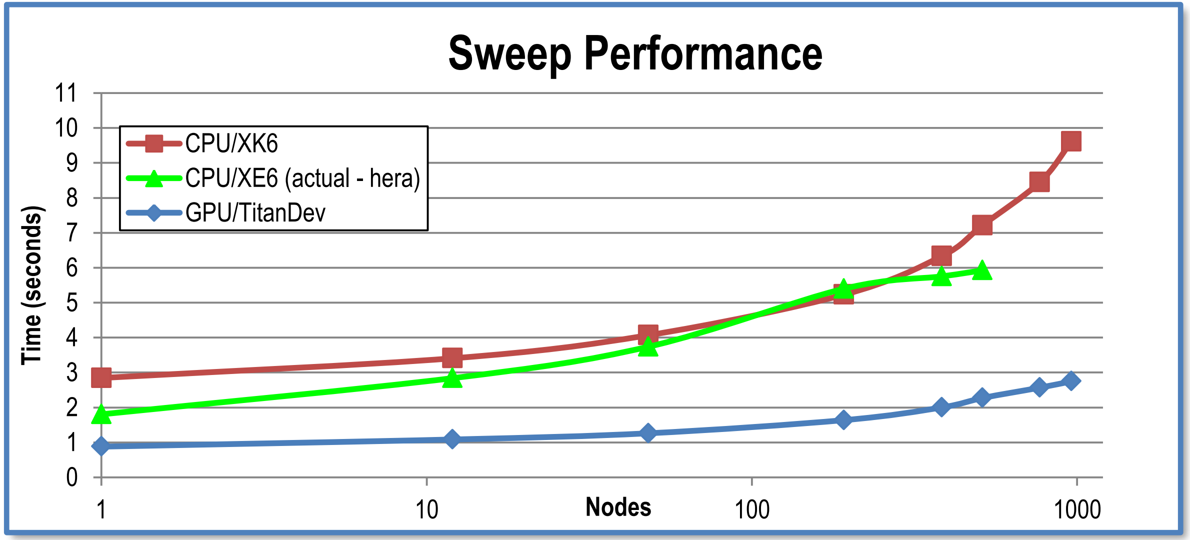

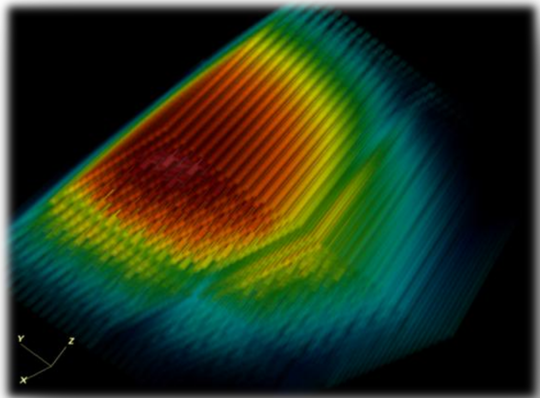

Performance Improvement factors		GPU
		XK6 Fermi
CPU	*XK6 / Interlagos*	**3.5**
	XE6 / dual Interlagos	**3.3**

- Krylov multigroup solvers allow space-angle sweeps to be performed over all groups concurrently
- ideal for exploiting thread-based concurrency on GPUs
- space-angle sweep for all groups on GPU

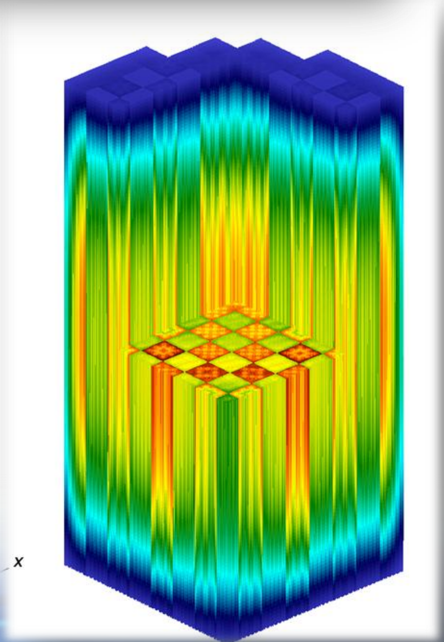

Future large-scale systems present challenges for applications

- Dramatic increases in node parallelism
  - 10 to 100× by 2015
  - 100 to 1000× by 2018
- Increase in system size contributes to lower mean time to interrupt (MTTI)
- Dealing with multiple additional levels of memory hierarchy
  - Algorithms and implementations that prioritize data movement over compute cycles
- Expressing this parallelism and data movement in applications
  - Programming models and tools are currently immature and in a state of flux

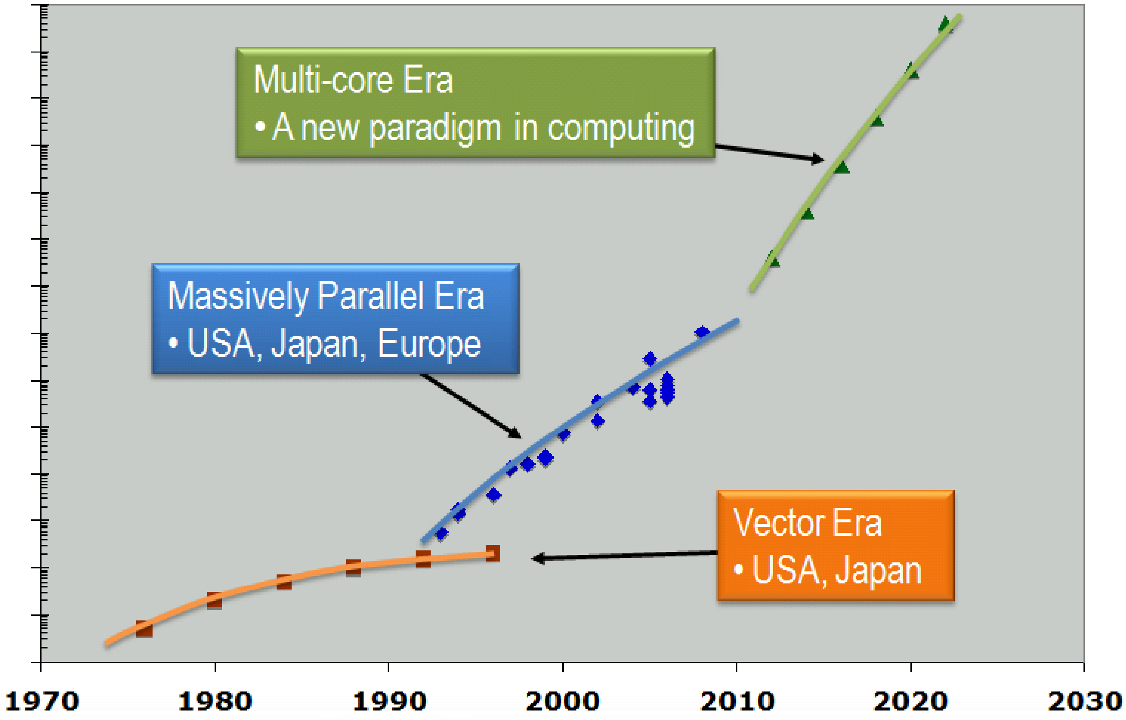

Exascale Initiative Steering Committee

Future ~~large-scale~~ desktop systems present challenges for applications

- Dramatic increases in node parallelism
  - 10 to 100× by 2015
  - 100 to 1000× by 2018
- Increase in system size contributes to
  TTI)
  levels of
  ons that
  r compu
  data
  - Programming models and tools are currently immature a

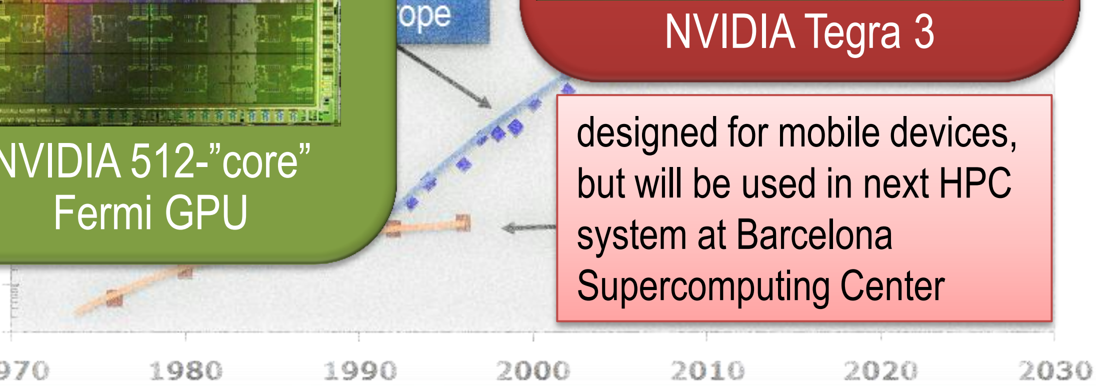

Intel 48-core experimental chip shipped in 2010

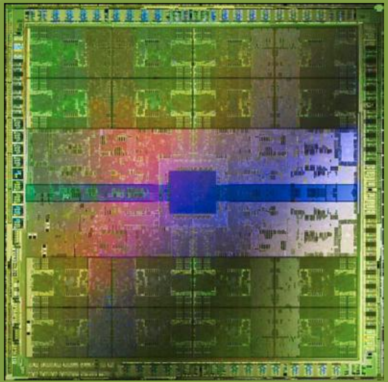

NVIDIA 512-"core" Fermi GPU

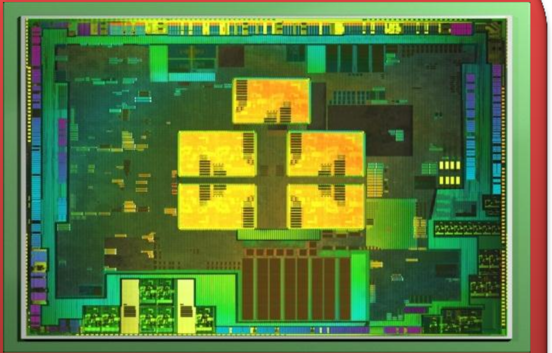

NVIDIA Tegra 3

designed for mobile devices, but will be used in next HPC system at Barcelona Supercomputing Center

ittee

Over the life of CASL, these challenges will become increasingly significant at the desktop level

CASL

Questions? http://www.casl.gov/ -or- info@casl.gov

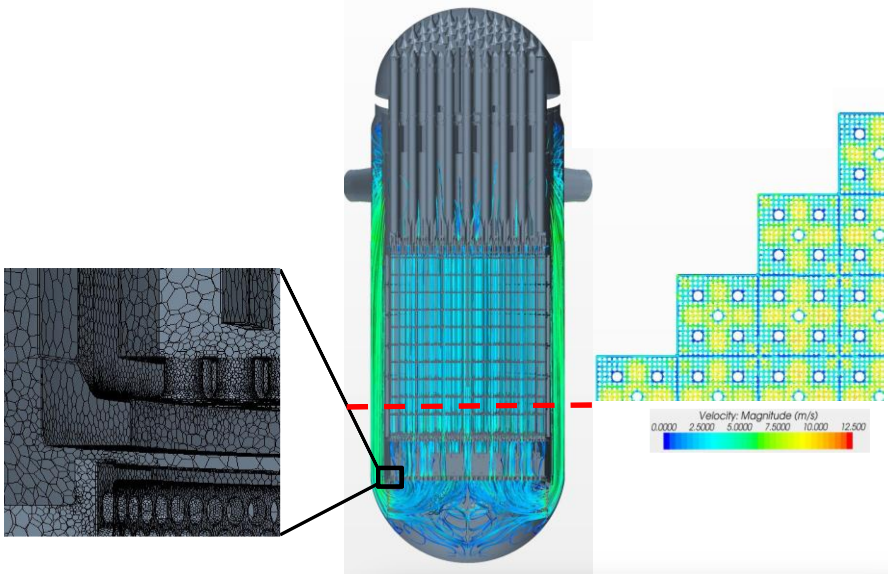

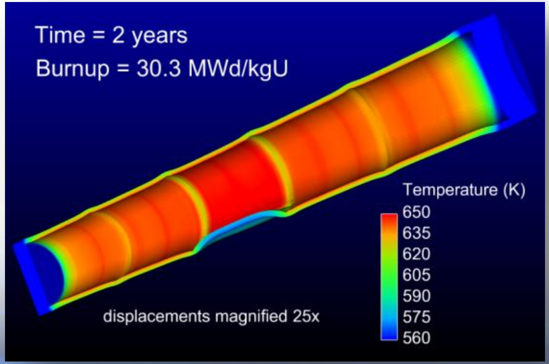

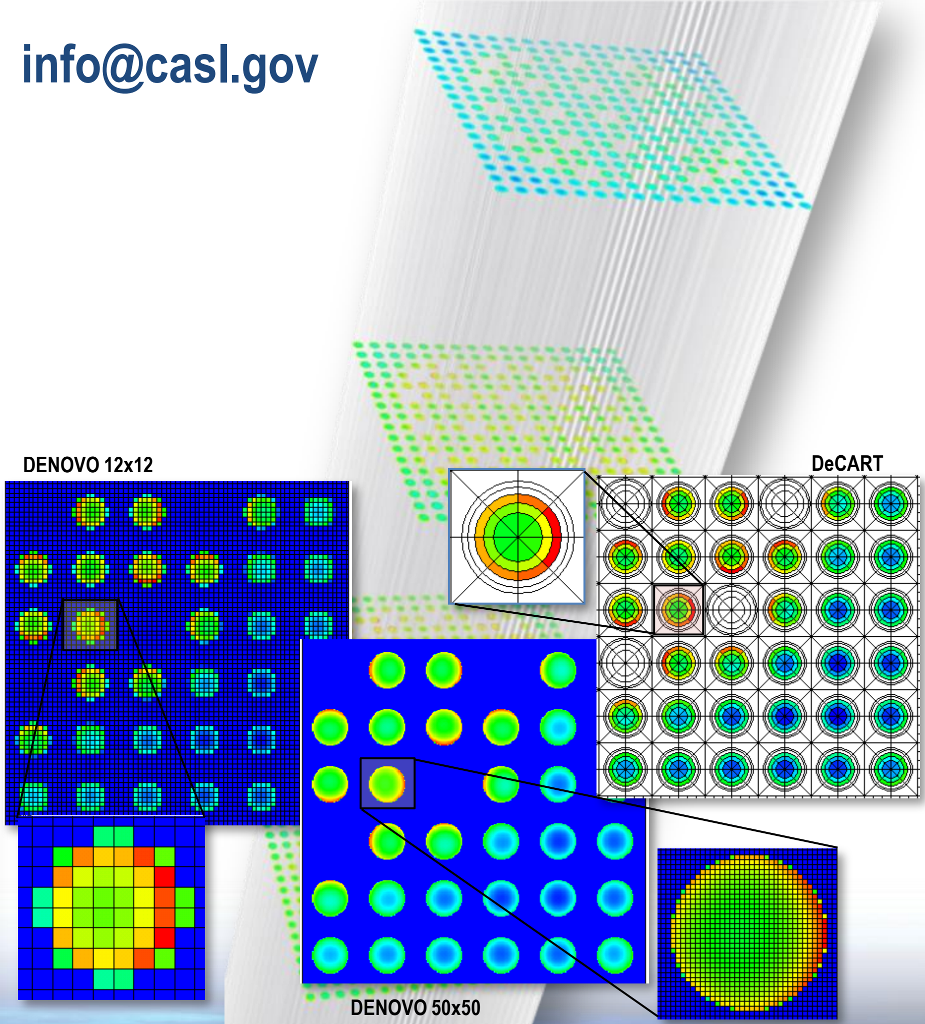

Supplemental

CASL Technical Focus Areas

- neutron behavior
- fluid flow and heat transfer

- V&V and calibration through data assimilation
- sensitivity analysis and uncertainty quantification

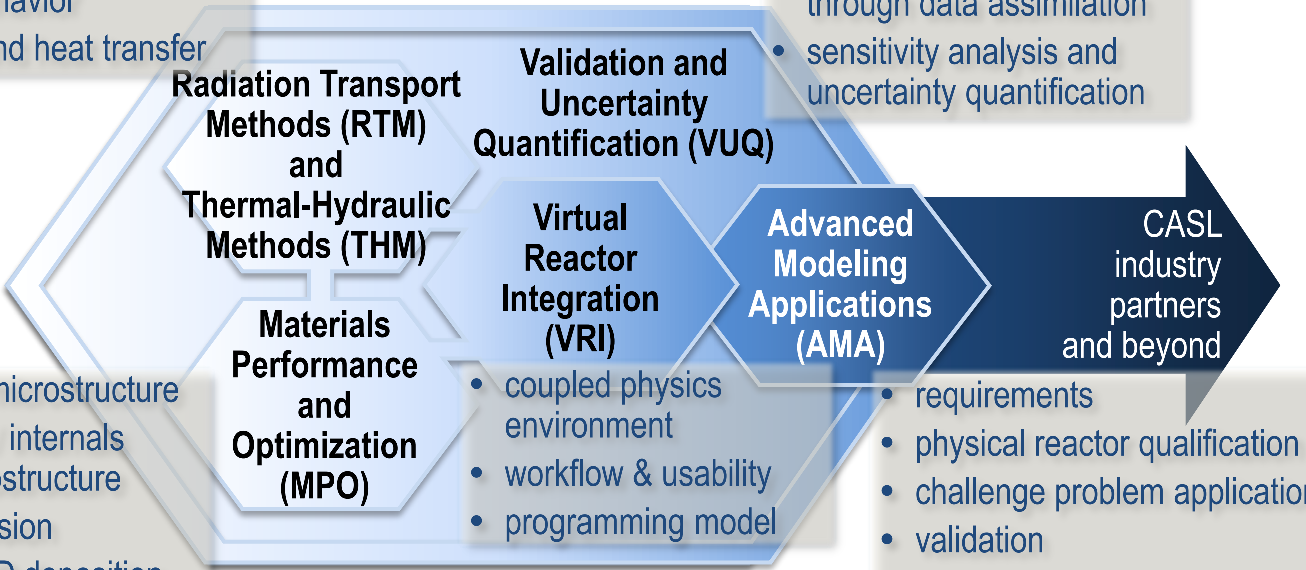

- fuel microstructure
- clad / internals microstructure
- corrosion
- CRUD deposition
- failure modes

- coupled physics environment
- workflow & usability
- programming model

- requirements
- physical reactor qualification
- challenge problem application
- validation
- NRC engagement

All Focus Areas span institutions (labs, universities, industry)

CASL

Virtual Environment for Reactor Applications (VERA)

A suite of tools for scalable simulation of nuclear reactor core behavior

- Flexible coupling of physics components
- Toolkit of components
  - Not a single executable
  - Both legacy and new capability
  - Both proprietary and distributable

- Attention to usability
- Rigorous software processes
- Fundamental focus on V&V and UQ

- Development guided by relevant challenge problems
- Broad applicability

- Scalable from high-end workstation to existing and future HPC platforms
  - Diversity of models, approximations, algorithms
  - Architecture-aware implementations

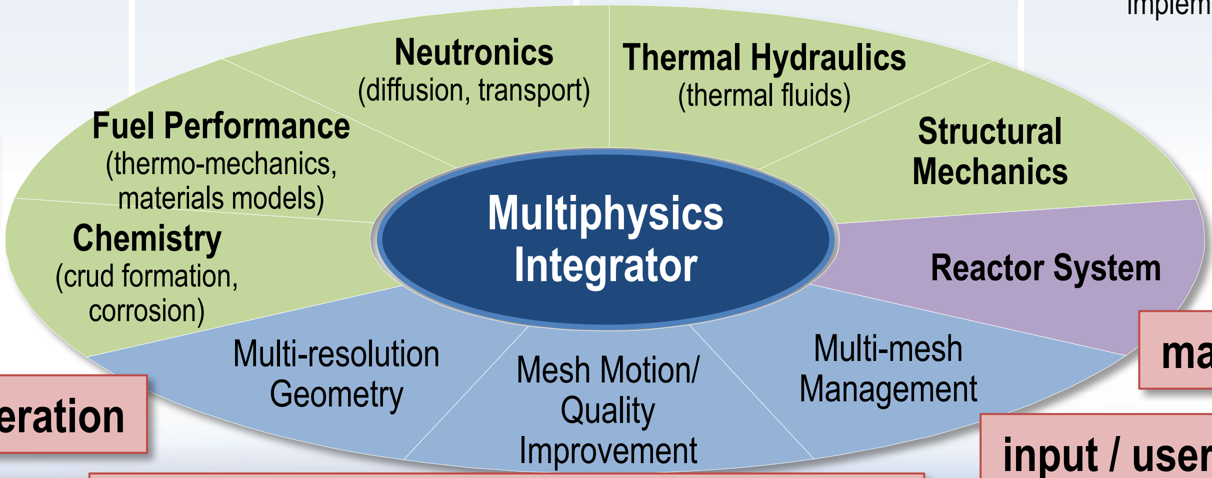

Missing…

geometry

mesh generation

material properties

input / user interface

workflow (analysis / design / optimization)

CASL has embraced Agile software development processes

- based on methodologies being used by partners
  - combine attributes of Scrum and Kanban methodologies
  - customized for CASL and refined as needed (iteratively)
- enabled diverse team to be productive very quickly

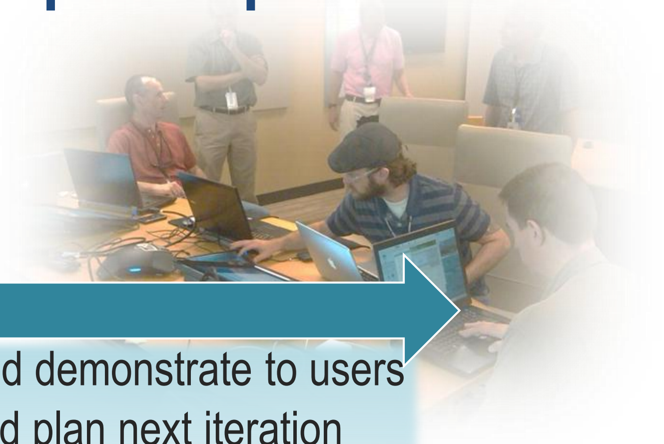

Start
- users prioritize goals
- team determines work assignments

Execute
- two 30-minute standup meetings each week

End
- deliver and demonstrate to users
- review and plan next iteration

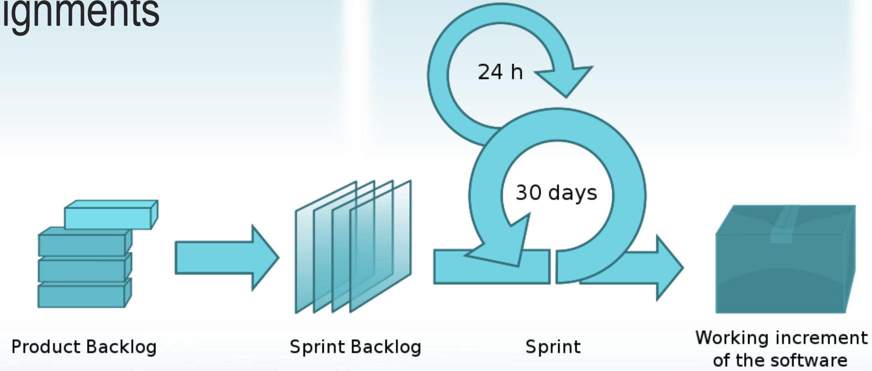

Scrum: http://en.wikipedia.org/wiki/Scrum_%28development%29

Desirable attributes
- emphasis on collaboration and adaptability
- constant communication / interaction
  - both within team and with user community
- accommodates changing requirements & unpredictability

Agility + Formality

CASL advanced CRUD modeling predictions

- Colored contours: boron concentration within crud layer
- Findings:
  - Crud thickness and boron vary with *T* variations on cladding surface
  - Crud and boron reduced by turbulence behind mixing vanes