Toward Predictive Modeling of Nuclear Reactor Performance: Application Development Experiences, Challenges, and Plans in CASL

The Consortium for Advanced Simulation of Light Water Reactors
A DOE Energy Innovation Hub

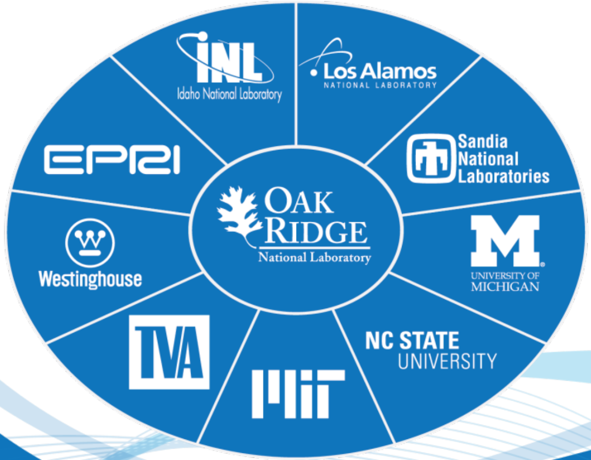

Douglas B. Kothe
Oak Ridge National Laboratory
Director, CASL

U.S. DEPARTMENT OF
ENERGY
Nuclear Energy

CASL
Consortium for Advanced Simulation of LWRs

Development Platforms

Cray Systems at LANL in 1980s

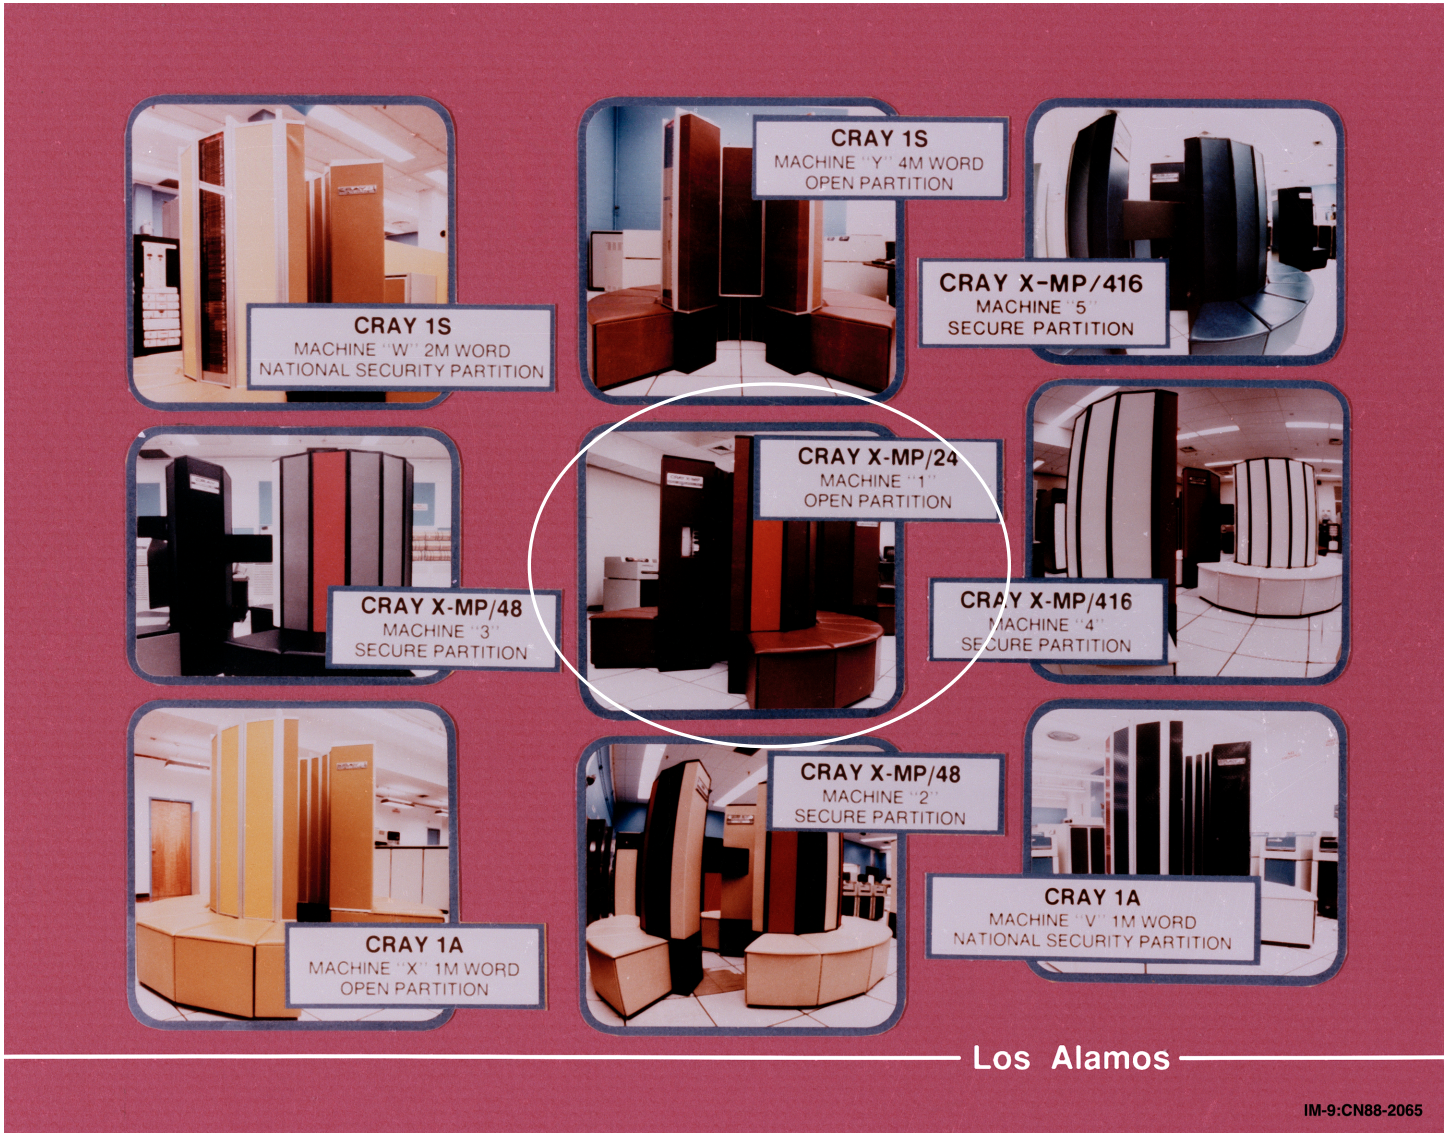

Memories

- DDT debugger
- Cray Assembly Language (CAL)
- Fortran Command Language (FCL)
- tiny
- Common File System (CFS)
- Bank points

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

The PAGOSA code

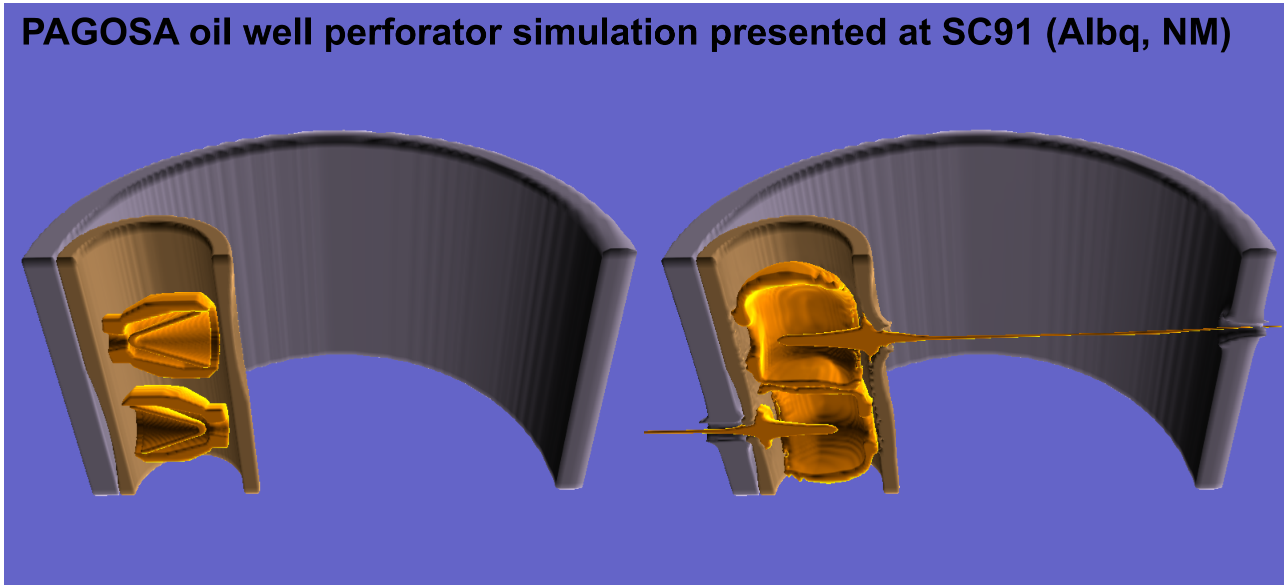

PAGOSA oil well perforator simulation presented at SC91 (Albq, NM)

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
  - ○ After showing SNL what was under the hood, we backed out of this program and deferred to SNL CTH code (boy was that ever a mistake)

Development Platforms
Connection Machines at LANL in Early 1990s

CM-2, CM-200 (1989)

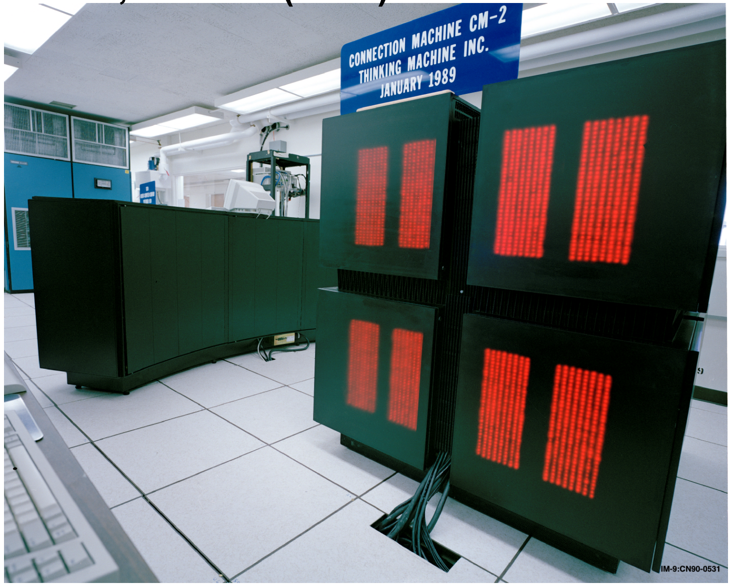

CM-5 (1992)

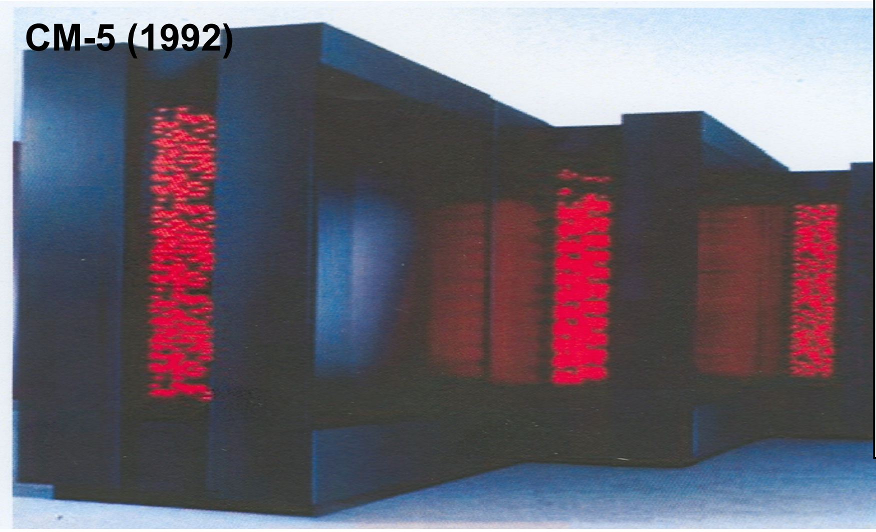

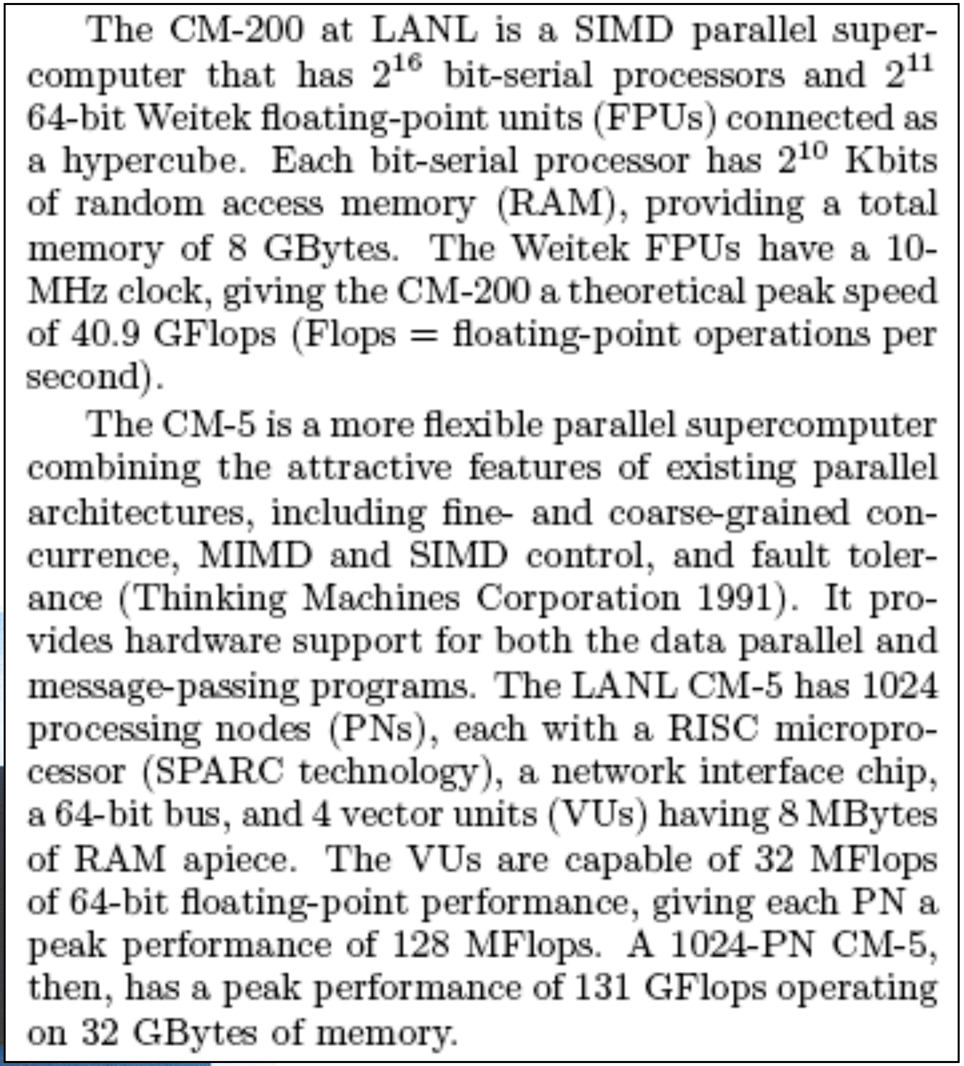

The CM-200 at LANL is a SIMD parallel supercomputer that has 2^{16} bit-serial processors and 2^{11} 64-bit Weitek floating-point units (FPUs) connected as a hypercube. Each bit-serial processor has 2^{10} Kbits of random access memory (RAM), providing a total memory of 8 GBytes. The Weitek FPUs have 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

Application Example
Gravity Casting Mold Fill

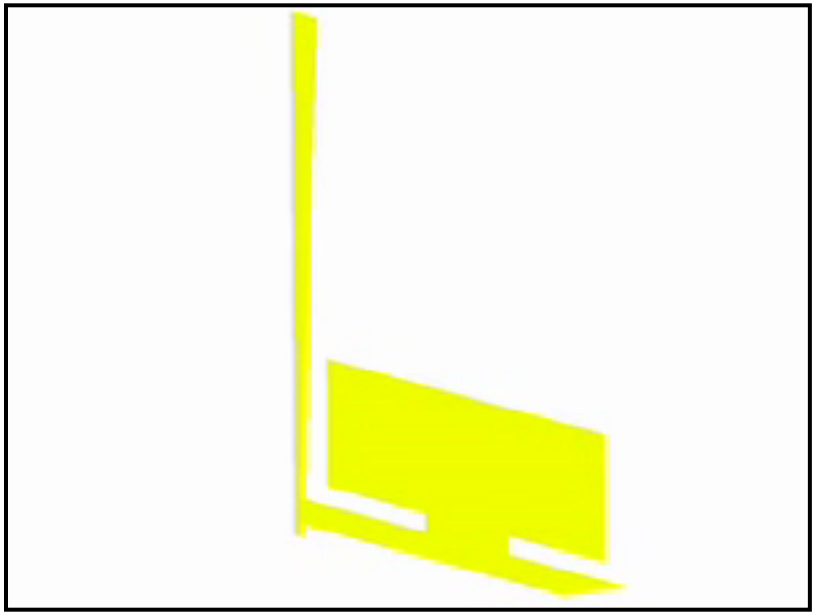

- International casting mold-fill benchmark put together by manufacturing industries and R&D agencies
  - Does the offshore hydro community do this often?
- Molten aluminum into air
  - Turbulence free surface flow of Navier-Stokes fluid
- Blind test tried by *many* codes – and many failed
- We "cheated"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

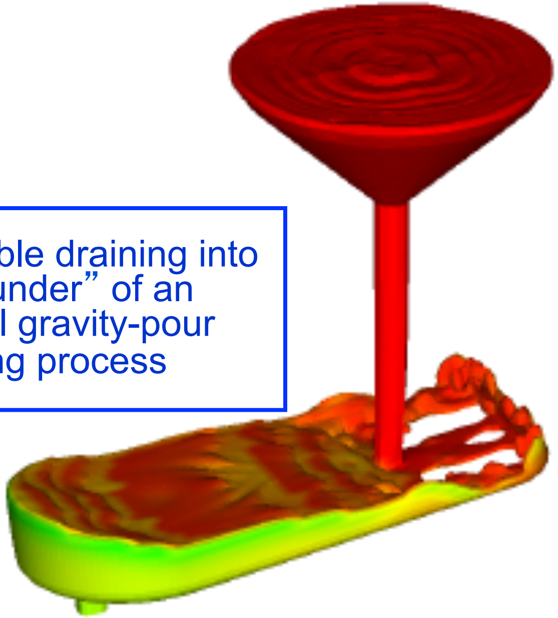

Crucible draining into a "launder" of an actual gravity-pour casting process

Los Alamos Cylindrical Shell Test Part

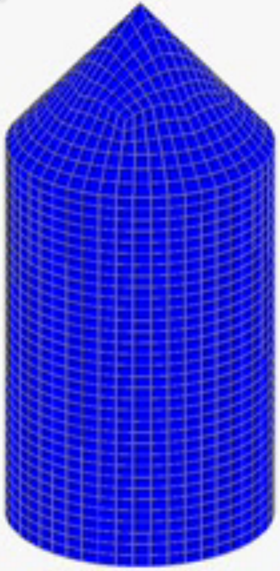

Application Requirements @ the PF & Beyond

- We have recently surveyed, analyzed, and documented the science drivers and application requirements envisioned for leadership systems out to the 2020 timeframe
- These studies help to
  - Provide a roadmap for the ORNL Leadership Computing Facility
  - Uncover application needs and requirements
  - Focus our efforts on those disruptive technologies and research areas in need of our and the HPC community's attention

COMPUTATIONAL SCIENCE REQUIREMENTS FOR LEADERSHIP COMPUTING

JULY 2007

NATIONAL CENTER FOR COMPUTA
LEADERSHIP COMPUTING FACILITY • OAK RIDG

OAK RIDGE
NATIONAL LABORATORY
MANAGED BY UT-BATTELLE
FOR THE DEPARTMENT OF ENERGY

ORNL/TM-2007/232

Science Prospects and Benefits
with Exascale Computing

OAK RIDGE
NATIONAL LABORATORY
MANAGED BY UT-BATTELLE
FOR THE DEPARTMENT OF ENERGY

ORNL/TM-2007/238

Scientific Application Requirements
for Leadership Computing
at the Exascale

December 2007

Prepared by
Computing Requirements Team
National Center for Computational Sciences

UT-BATTELLE
ORNL-27 (4-00)

PF Application Findings
(with some opinion)

- A rigorous & evolving apps reqms process pays dividends
  - Needs to be quantitative: apps cannot “lie” with performance analysis
- Algorithm development is evolutionary
  - Can we break this mold?
  - Ex: Explore new parallel dimensions (time, energy)
- Hybrid/multi-level programming models virtually nonexistent
- No algorithm “sweet spots” (one size fits all)
  - But algorithm footprints share characteristics
- V&V and SQA not in good standing
  - Ramifications on compute systems as well as apps results generated
- No one is really clamoring for new languages
- MPI until the water gets too hot (frog analogy)
- Apps lifetimes are >3-5x machine lifetimes
  - Refactoring a way of life
- Fault tolerance via defensive checkpointing de facto standard
  - Won’t this eventually bite us? Artificially drives I/O demands
- Weak or strong scale or both (no winner)
- Data analytics paradigm must change
- The middleware layer is surprisingly stable and agnostic across apps (and should expand!)

Scalable Applications
How We Can Accelerate Development & Readiness

- Automated diagnostics
  - Drivers: performance analysis, application verification, S/W debugging, H/W-fault detection and correction, failure prediction and avoidance, system tuning, and requirements analysis
- Hardware latency*
  - Won’t see improvement nearly as much as flop rate, parallelism, B/W in coming years
  - Can S/W strategies mitigate high H/W latencies?
- Hierarchical algorithms*
  - Applications will require algorithms aware of the system hierarchy (compute/memory)
  - In addition to hybrid data parallelism, and file-based checkpointing, algorithms may need to include dynamic decisions between recomputing and storing, fine-scale task-data hybrid parallelism, and in-memory checkpointing
- Parallel programming models*
  - Improved programming models needed to allow developer to identify an arbitrary number of levels of parallelism and map them onto hardware hierarchies at runtime
  - Models continue to be coupled into larger models, driving the need for arbitrary hierarchies of task and data parallelism

Scalable Applications
How We Can Accelerate Development & Readiness

- Solver technology and innovative solution techniques*
  - Global communication operations across $10^{6\text{-}8}$ processors will be prohibitively expensive, solvers will have to eliminate global communication where feasible and mitigate its effects where it cannot be avoided. Research on more effective local preconditioners will become a very high priority
  - If increases in memory B/W continue to lag the number of cores added to each socket, further research needed into ways to effectively trade flops for memory loads/stores
- Accelerated time integration*
  - Are we ignoring the time dimension along which to exploit parallelism? (Ex: climate)
- Model coupling*
  - Coupled models require effective methods to implement, verify, and validate the couplings, which can occur across wide spatial and temporal scales. The coupling requirements drive the need for robust methods for downscaling, upscaling, and coupled nonlinear solving
  - Evaluation of the accuracy and importance of couplings drives the need for methods for validation, uncertainty analysis, and sensitivity analysis of these complex models
- Maintaining current libraries
  - Reliance of current HPC applications on libraries will grow
  - Libraries must perform as HPC systems grow in parallelism and complexity

Applications at the Exascale
A Speculative Look

- We are in danger of failing because of a software crisis unless concerted investments are undertaken to close the H/W-S/W gap
  - H/W has gotten way ahead of the S/W (same ole – same ole?)
- Structured grids and dense linear algebra continue to dominate, but …
  - Increase projected for Monte Carlo algorithms, unstructured grids, sparse linear algebra, and particle methods (relative decrease in FFTs)
  - Increasing importance for AMR, implicit nonlinear systems, data assimilation, agent-based method, parameter continuation, optimization
- Priority of computing system attributes
  - Increase: interconnect bandwidth, memory bandwidth, mean time to interrupt, memory latency, and interconnect latency
    - Reflect desire to increase computational efficiency to use peak flops
  - Decrease: disk latency, archival storage capacity, disk bandwidth, wide area network bandwidth, and local storage capacity
    - Reflect expectation that computational efficiency will not increase
  - Per-core requirements relatively static, while aggregate requirements will grow with the system

Applications at the Exascale
A Speculative Look

- System software must possess more stability, reliability, and fault tolerance during application execution
  - New fault tolerance paradigms must be developed and integrated into applications
    - Job management and efficient scheduling of those resources will be a major obstacle faced by computing centers
- Systems must be much better “science producers”
  - Strong software engineering practices must be applied to systems to ensure good end-to-end productivity
  - Data analytics must empower scientists to ask “what-if” questions, providing S/W & H/W infrastructure capable of answering these questions in a timely fashion (Google desktop)
  - Strong data management will become an absolute at the exascale
- Just like H/W requires disruptive technologies for acceleration of natural evolutionary paths, so too will algorithm, software, and physical model development efforts need disruptive technologies (invest now!)

What Will an EF System Look Like?

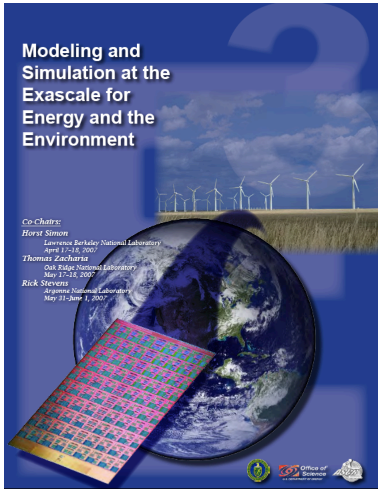

Modeling and Simulation at the Exascale for Energy and the Environment

Co-Chairs:
Horst Simon
Lawrence Berkeley National Laboratory
April 17-18, 2007
Thomas Zacharia
Oak Ridge National Laboratory
May 17-18, 2007
Rick Stevens
Argonne National Laboratory
May 31-June 1, 2007

- All projections are daunting
  - Based on projections of existing technology both with and without “disruptive technologies”
  - Assumed to arrive in 2020-22 timeframe
- Example 1
  - 115K nodes @ 10 TF per node, 50-100 PB, optical interconnect, 150-200 GB/s injection B/W per node, 50 MW
- Examples 2-4 (DOE “Townhall” report*)

Example system	Ops/ cycle	Freq [GHz]	Cores/ socket	Peak/ socket [TF/s]	Sockets	Total cores	Peak/ system [EF/s]	Power [MW]
A	4	3.0	64	0.768	1300k	85M	1.0	130
B	8	16.0	128	16.0	120k	15M	2.0	60 - 80
C	8	1.5	512	6.1	200k	100M	1.8	20 - 40

*www.er.doe.gov/ASCR/ProgramDocuments/TownHall.pdf

Current requirements

Application	**Structured**	**Unstructured**	**FFT**	**Dense**	**Sparse**	**Particles**	**Monte Carlo**
Molecular		X	X	X		X	
Nanoscience	X			X		X	X
Climate	X		X		X	X	
Environment	X	X			X		
Combustion	X						
Fusion	X		X	X	X	X	X
Nuc. energy		X		X	X		
Astrophysics	X	X		X	X	X	
Nuc. physics				X			
Accelerator		X			X		
QCD	X						X
#X	7	5	3	6	6	5	3

Current hardware requirements

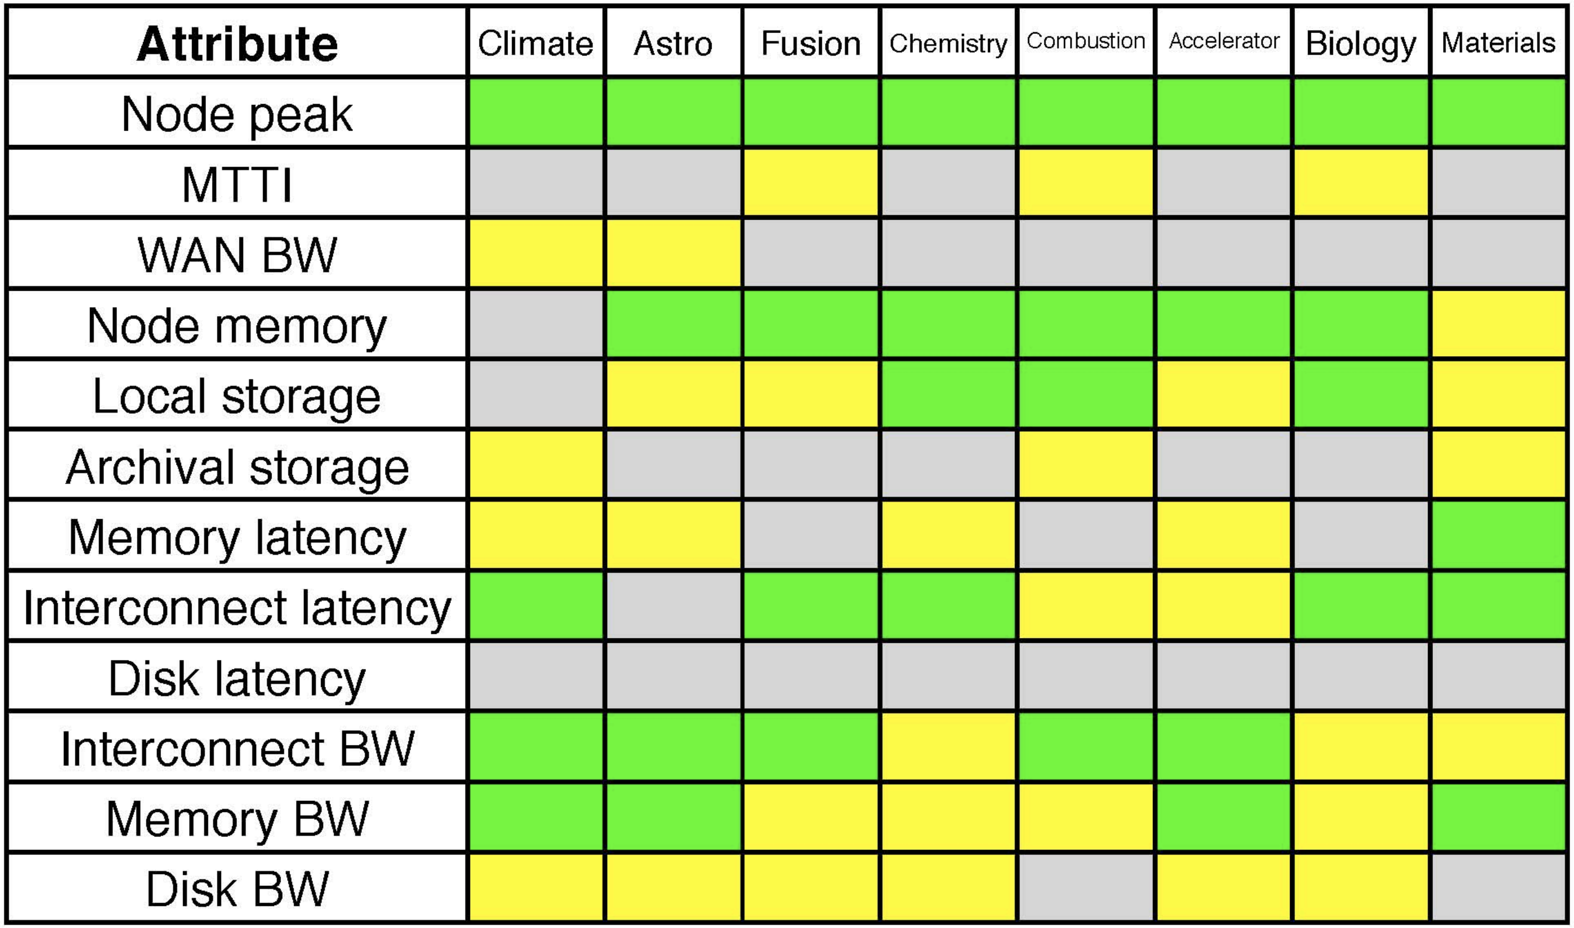

Attribute	Climate	Astro	Fusion	Chemistry	Combustion	Accelerator	Biology	Materials
Node peak	🟩	🟩	🟩	🟩	🟩	🟩	🟩	🟩
MTTI	⬜	⬜	🟨	⬜	🟨	⬜	🟨	⬜
WAN BW	🟨	🟨	⬜	⬜	⬜	⬜	⬜	⬜
Node memory	⬜	🟩	🟩	🟩	🟩	🟩	🟩	🟨
Local storage	⬜	🟨	🟨	🟩	🟩	🟨	🟩	🟨
Archival storage	🟨	⬜	⬜	⬜	🟨	⬜	⬜	🟨
Memory latency	🟨	🟨	⬜	🟨	⬜	🟨	⬜	🟩
Interconnect latency	🟩	⬜	🟩	🟩	🟨	🟨	🟩	🟩
Disk latency	⬜	⬜	⬜	⬜	⬜	⬜	⬜	⬜
Interconnect BW	🟩	🟩	🟩	🟨	🟩	🟩	🟨	🟨
Memory BW	🟩	🟩	🟨	🟨	🟨	🟩	🟨	🟩
Disk BW	🟨	🟨	🟨	🟨	⬜	🟨	🟨	⬜

Nuclear Energy Overview

Source: Nuclear Energy Institute (NEI)

- World nuclear power generating capacity
  - 436 plants (U.S. - 100 plants in 31 states)
  - U.S. electricity generation (2012): nuclear is 0.77 out of 4.05 TWh
  - 72 nuclear plants under construction in 15 countries (5 in U.S.!)
- Electricity from nuclear: 19.0% in U.S. (12.3% worldwide)
- U.S. electricity demand projected to grow 25% by 2030
  - 2007: 3.99 TWh
  - 2030: 4.97 TWh
- Nuclear accounts for 64% of emission-free electricity in U.S.

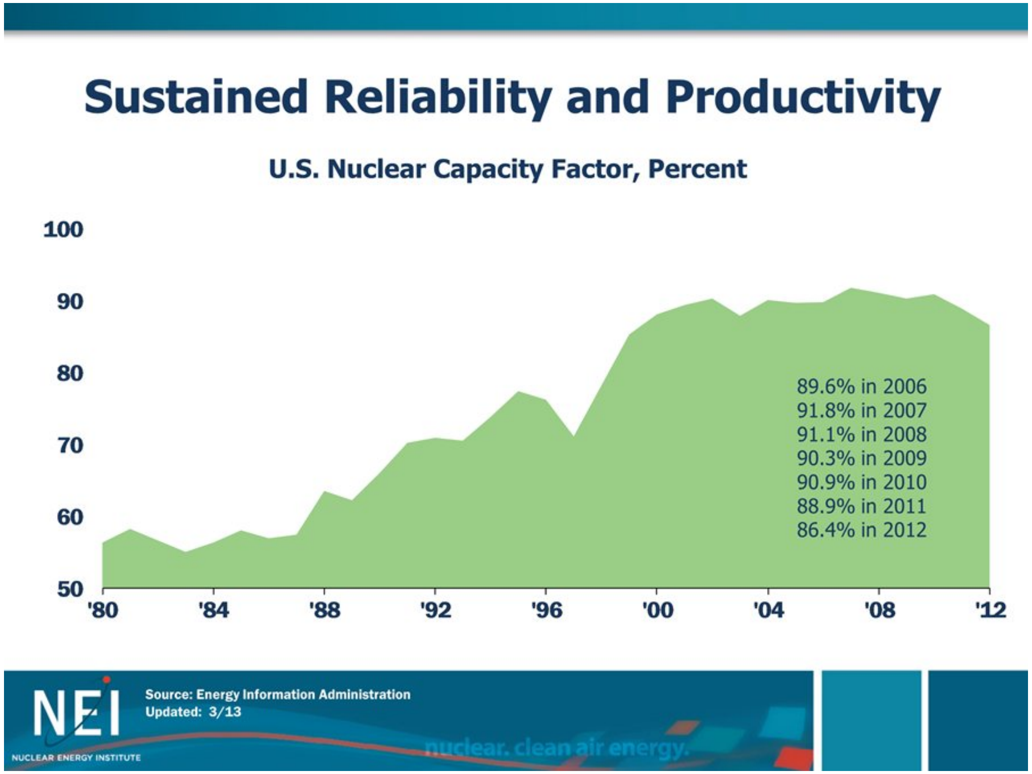

Sustained Reliability and Productivity

U.S. Nuclear Capacity Factor, Percent

89.6% in 2006
91.8% in 2007
91.1% in 2008
90.3% in 2009
90.9% in 2010
88.9% in 2011
86.4% in 2012

Source: Energy Information Administration
Updated: 3/13

Nuclear Energy: Clean and Reliable

Greenhouse Gas-Free Electricity Production 2012

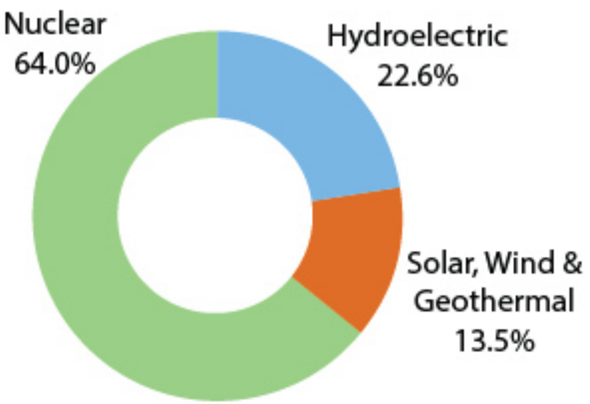

Source: U.S. Energy Information Administration

Average Operating Efficiency by Source of Electricity, 2012

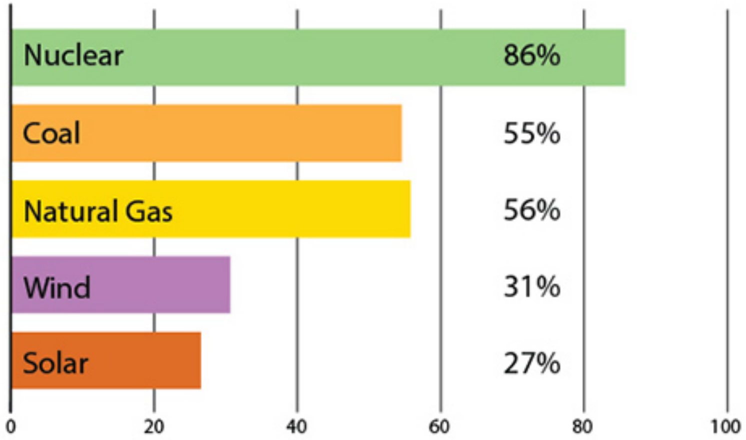

Source: Ventyx / U.S. Energy Information Administration

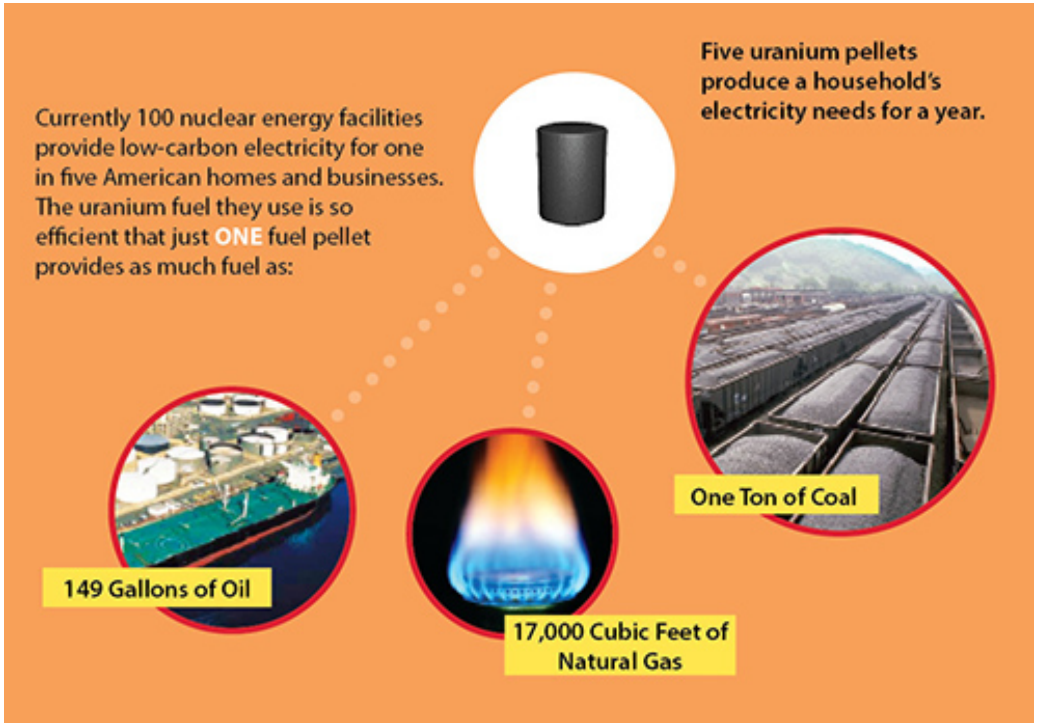

Currently 100 nuclear energy facilities provide low-carbon electricity for one in five American homes and businesses. The uranium fuel they use is so efficient that just ONE fuel pellet provides as much fuel as:

Five uranium pellets produce a household's electricity needs for a year.

More info at nei.org (Nuclear Energy Institute)

Anatomy of a Nuclear Reactor

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

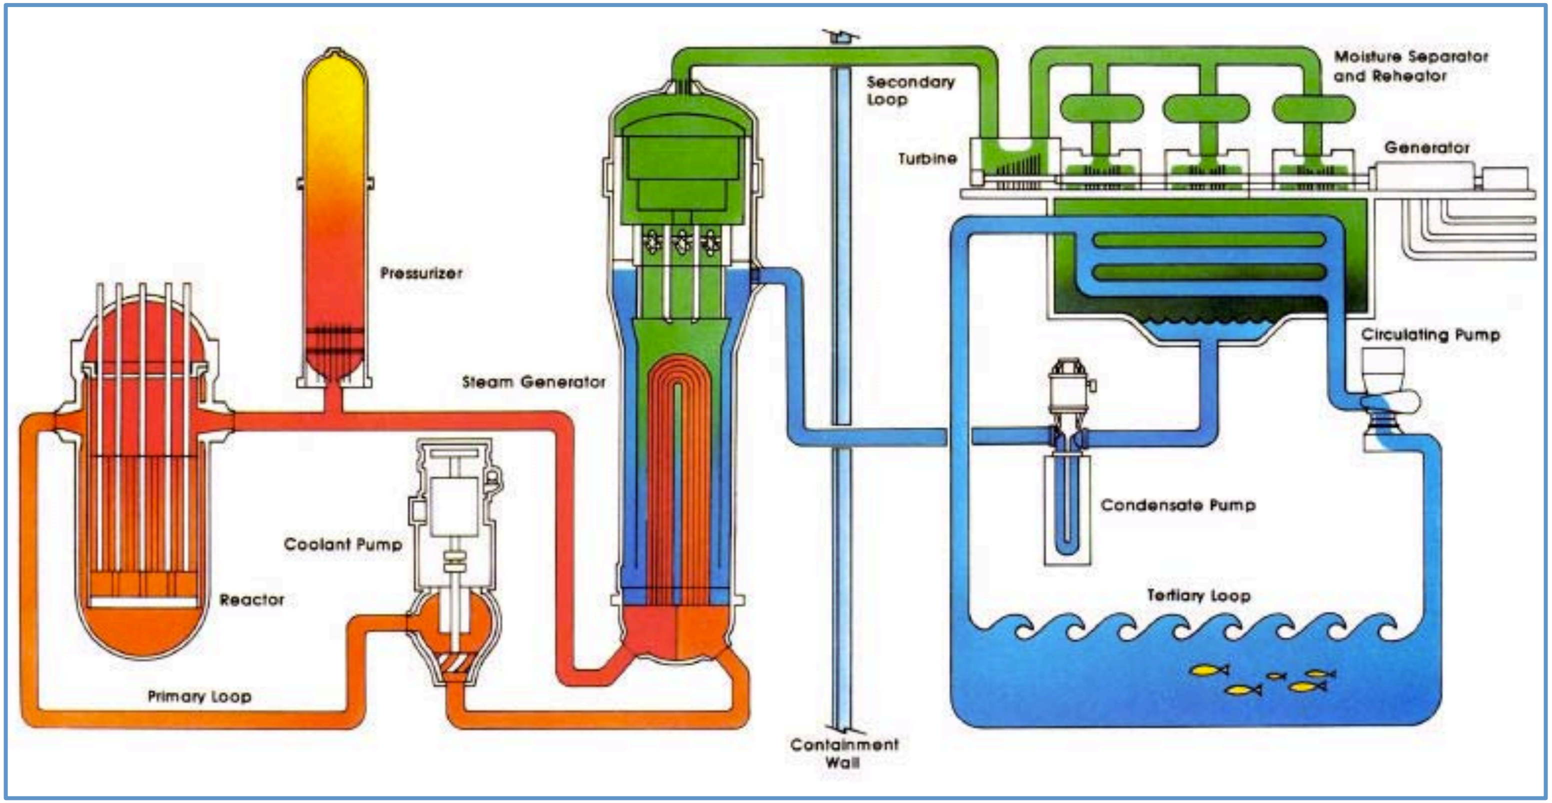

Power: ~1170 MWe (~3400 MWth)

Core: 11.1’ diameter x 12’ high, 193 fuel assemblies, 107.7 tons of UO_2

Coolant: pressurized water (2250 psia), T_{in} ~ 545°F, T_{out} ~ 610°F, 134M lb/h (4 pumps)

Pressure Vessel: 14.4’ diameter x 41.3’ high x 0.72’ thick alloy steel

Containment Building: 115’ diameter x 156’ high steel / concrete

Anatomy of a Nuclear Reactor

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

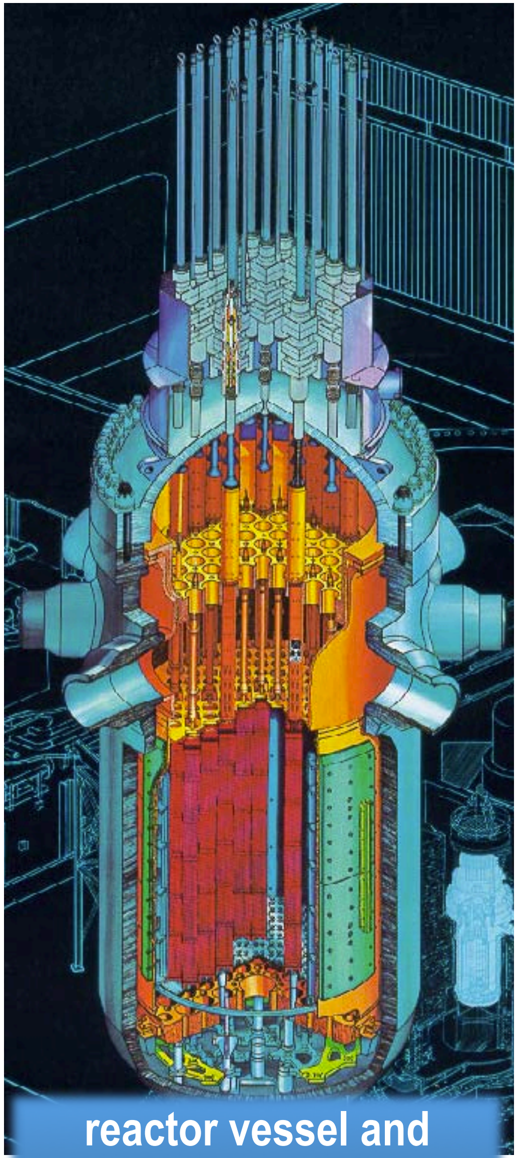

reactor vessel and internals

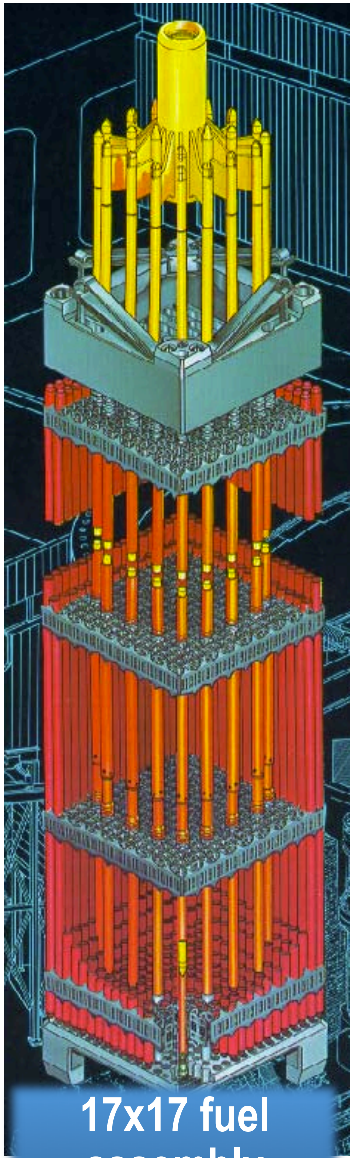

17x17 fuel assembly

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's Charter

Mission is to provide leading-edge modeling and simulation capabilities to improve the performance of currently operating light water reactors

Vision	*Predict, with confidence, the performance and assured safety of nuclear reactors, through comprehensive, science-based M&S technology deployed and applied broadly by the U.S. nuclear energy industry*
Goals	• Develop and effectively apply modern virtual reactor technology • Provide more understanding of safety margins while addressing operational and design challenges • Engage the nuclear energy community through M&S • Deploy new partnership and collaboration paradigms
Strategies	• Virtual Environment for Reactor Applications (VERA) • Industry Challenge Problems • Technology Delivery • Targeted, Enabling R&D • Education and Training • Collaboration and Ideation 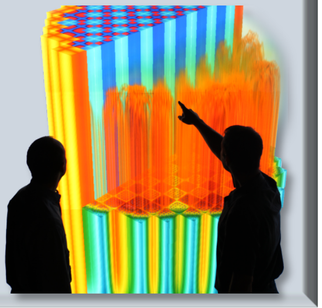

Scope

- ☐ Address, through new insights afforded by advanced M&S technology, key nuclear energy industry challenges
  - ✓ furthering power uprates
  - ✓ higher fuel burnup
  - ✓ lifetime extension

  while providing higher confidence in enhanced nuclear safety
- ☐ Focus on performance of pressurized water reactor core, vessel, and in-vessel components to provide greatest impact within 5 years

CASL Components

US team with a remarkable set of assets – Address tough industry challenges that matter – Urgent and compelling plan

Collaborate creatively – Target and foster innovation - Deliver industry solutions with predictive simulation

CASL Background

- **What is CASL doing?**
  - *Create* an advanced coupled multi-physics "virtual reactor" technology by adapting existing and developing new modeling and simulation (M&S) tools
  - *Effectively apply* the virtual reactor technology to provide more understanding of safety margins while addressing selected operational and design challenges of operational light water reactors
- **Why?**
  - Improve the performance and energy output of existing nuclear reactors by focusing on important industry defined challenge problems
  - M&S technology has long been a mainstay in the nuclear industry (vendors, owner/operators), helping to inform consequential operational and safety decisions codes daily. Current nuclear industry M&S technology, though continuously improved, has failed to capitalize on the benefits that more precise predictive capability and fundamental understanding offer
- **Why do this in the Hub R&D business model?**
  - Solution requires clear deliverables & products promoted by Hub R&D approach ("fierce sense of urgency")
  - Public-private partnership essential for adaptation, application, and "useful and usable" deployment of advanced M&S technologies under development at DOE national labs and universities to nuclear enterprise
- **What is working?**
  - Several elements have proven effective: partnerships, industry pull, technology deployment, clear deliverables and plans, effective and agile project management, 5-year time horizon, S&T guidance/review

Strong Dependency on Modeling and Simulation

Need to assure nuclear safety but limited by inability to perform full-scale experimental mockups due to cost, safety & feasibility [1% power derating translates to \$(5-10)M annual loss of revenue for 1 GWe unit]

Need to minimize economic uncertainty associated with new product introduction (e.g. fuel) by employing precise predictions [1% error in core reactivity has \$4M annual fuel cycle cost impact for 1 GWe unit]

The CASL Team

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

Contributing Partners

ASCOMP GmbH
CD-adapco
City College of New York
Florida State University
Imperial College London
Rensselaer Polytechnic Institute
Texas A&M University
Pennsylvania State University
University of Florida
University of Tennessee – Knoxville
University of Wisconsin
University of Notre Dame
Anatech Corporation
Core Physics Inc.
Pacific Northwest National Laboratory
G S Nuclear Consulting, LLC
University of Texas at Austin
University of Texas at Dallas

First Introduced by Secretary in the President's FY2010 Budget

A Different Approach

- "Multi-disciplinary, highly collaborative teams ideally working under one roof to solve priority technology challenges" – *Steven Chu*
- "Create a research atmosphere with a <u>fierce sense of urgency</u> to deliver solutions." – *Kristina Johnson*
- Characteristics
  - Leadership – Outstanding, independent, scientific leadership
  - Management – "Light" federal touch
  - Focus – Deliver technologies that can change the U.S. "energy game"

CASL Challenge Problems

Key safety-relevant reactor phenomena that limit performance

Departure from Nucleate Boiling

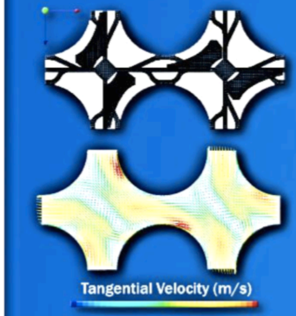

Cladding Integrity

- During LOCA
- During reactivity insertion accidents
- Use of advanced materials to improve cladding performance

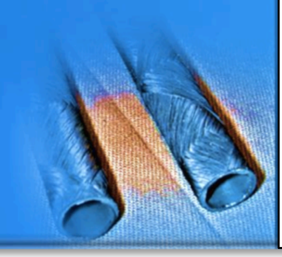

Safety Related Challenge Problems

CASL Challenge Problems

- Are relevant industry problems whose solutions remain elusive
- Are amenable to insight afforded by advanced M&S
- Help to direct RD&D activities on CASL M&S technology
- Help to establish clear performance metrics

Crud

- Deposition
- Axial offset anomaly
- Hot spots

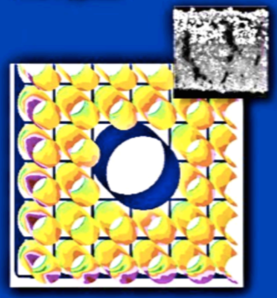

Grid-to-Rod Fretting

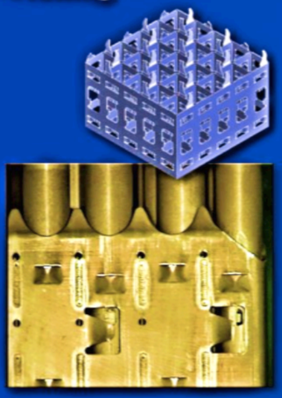

Pellet-Clad Interaction

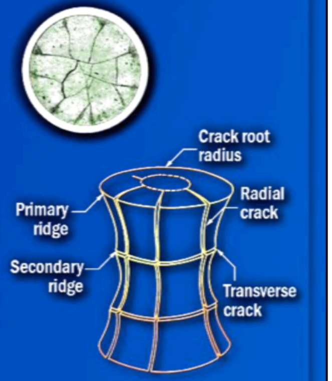

Operational Challenge Problems

Nuclear Energy Drivers and Payoffs for M&S technology

- **Extend licenses of existing fleet (to 60 years and beyond)**
  - Understand material degradation to reduce inspection & replacements
- **Up-rate power of existing fleet (strive for another 5-10 GWe)**
  - Address power-limiting operational & design basis accident scenarios
- **Inform flexible nuclear power plant operations**
  - Load follow maneuvering & coolant chemistry to enhance reliability
- **Design and deploy accident tolerant fuel (integrity of cladding)**
  - Concept refinement, test planning, assessment of safety margins
- **Margin quantification, recovery, tradeoff**
  - Plant parameters, fuel hardware, reload flexibility, regulatory changes
- **Resolve advanced reactor design & regulatory challenges**
  - Support Gen III+ reactors under construction (AP1000), refine SMR designs
- **Fuel cycle cost savings**
  - More economical core loadings and fuel designs
- **Used fuel disposition**
  - Inform spent fuel pools, interim storage, and repository decisions

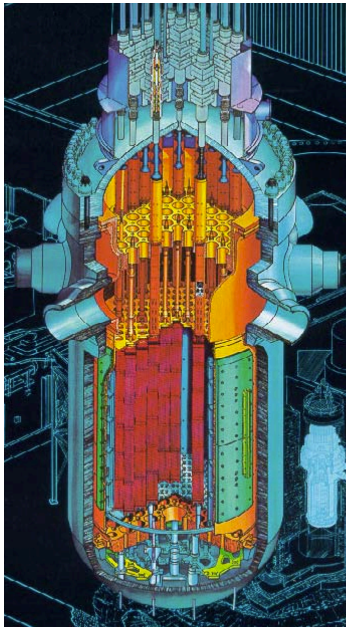

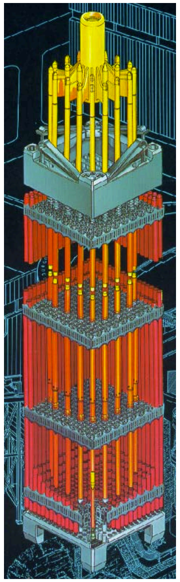

CASL Consortium for Advanced Simulation of LWRs

Power Uprates

Source: Heather Feldman (EPRI)

- **Measurement uncertainty recapture (MUR)**
  - Ex: Feed water flow rate (<2%)
- **Stretch Power Up-rate (SPU)**
  - Ex: Instrument set points (2% to 7%)
- **Extended Power Up-rate (EPU)**
  - Ex: Design changes (7% to 20%)
- **Ultra Power Up-rate (UPU)**
  - Ex: Extensive fuel and BOP changes (> 20%)
  - None have been performed
- **Equivalent to ~6 large nuclear power plants (6,440 Mwe) added to the grid thru uprates**
  - 143 power up-rates approved since 1977
- **About 6,000 MWe remains available for EPU**
  - 17 applications currently under review (9 MURs, 8 EPUs)
  - 15 new applications are expected in the next 5 years (8 MURs, 7 EPUs)

Westinghouse Experience

MUR COMPLETED

- 25 Americas PWRs
- 5 European PWRs
- 2 Asian PWRs

SPU COMPLETED

- 10 Americas PWRs
- 0 European PWRs
- 4 Asian PWRs

EPU COMPLETED

- 7 Americas PWRs
- 5 European PWRs
- 0 Asian PWRs
- Completed/planned: 5 of 6 2-loops, 4 of 13 3-loops, 0 of 30 4-loops, 5 of 14 CE Design

Where is NRC’s Focus on Up-rates?

- Japan Event Follow-up
- GSI -191 Post-LOCA Debris Effects
- Containment Accident Pressure (CAP)
- Thermal Conductivity Degradation (TCD)
- Boron Precipitation/Long Term Cooling
- Gas Accumulation
- LOCA Analysis
- Spent Fuel Pool Issues
- Digital I&C
- Alternate Source Term
- Steam Dryers
- Steam Generator Issues
- Single Failure Concerns
- High Energy Line Break
- Licensing Amendment Issues
- Licensing Conditions and Commitments

More at: http://www.nrc.gov/reactors/operating/licensing/power-uprates.html

Margin Management
Source: Sumit Ray (Westinghouse)

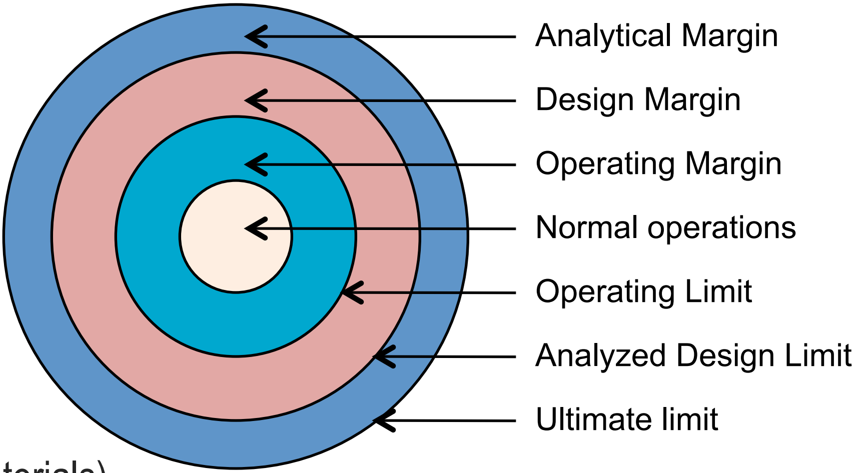

- Requires a strategic approach
  - How much is needed? How to allocate?
  - How can margin be transferred from one bucket to another?
- Key considerations
  - Plant operating parameters & assumptions (plant optimization & flexibility, load follow)
  - Fuel hardware (advanced product features & materials)
  - Design software and methodology (advanced technologies)
  - Core monitoring, In-core fuel management
  - Margins for the unknown or uncertain
  - Reload flexibility
  - Regulatory changes

Margin trade-offs and evaluation of risks require involvement of many stakeholders within the Utility (Fuels and Plant Operations) and suppliers (BOP, NSSS, T/G, etc.)

- Margins can be “recovered”
  - Change in design or operation or testing, reduced safety factor
  - Reduced calculational conservatism (possibly employing advanced analytic tools)
  - Changes to design characteristics of a limiting variable
  - Decrease in the margin of one parameter to increase the margin in another
  - Modification of system or component

One of the strategic targets for the CASL VERA toolkit is to provide enhanced insights in the area of critical reactor margins

An Example Nuclear Industry M&S Workflow

Crud Induced Power Shift Risk Evaluation

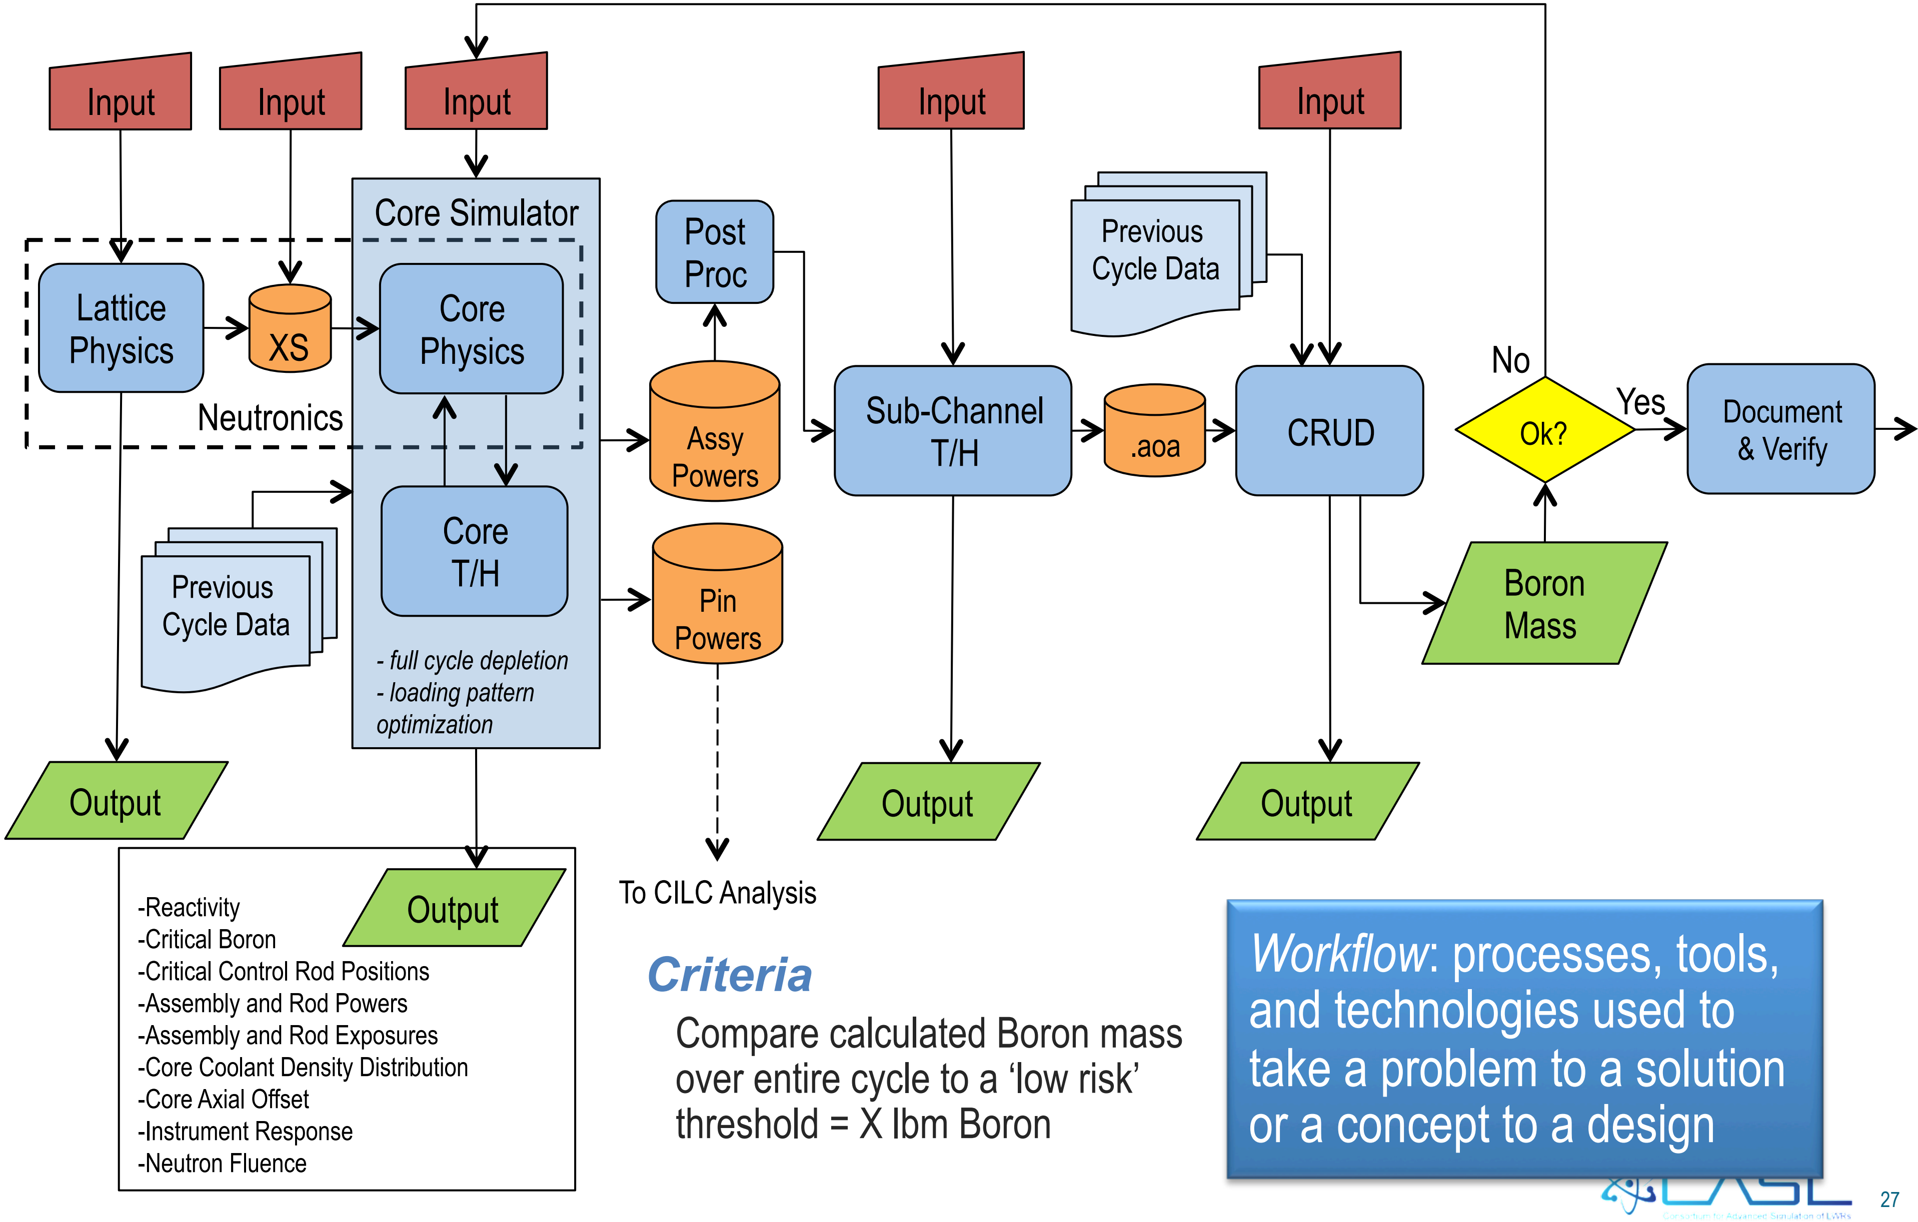

-Reactivity
-Critical Boron
-Critical Control Rod Positions
-Assembly and Rod Powers
-Assembly and Rod Exposures
-Core Coolant Density Distribution
-Core Axial Offset
-Instrument Response
-Neutron Fluence

Criteria

Compare calculated Boron mass over entire cycle to a 'low risk' threshold = X lbm Boron

Workflow: processes, tools, and technologies used to take a problem to a solution or a concept to a design

CASL Targets the Multi-Scale Challenge of Predictively Simulating a Reactor Core

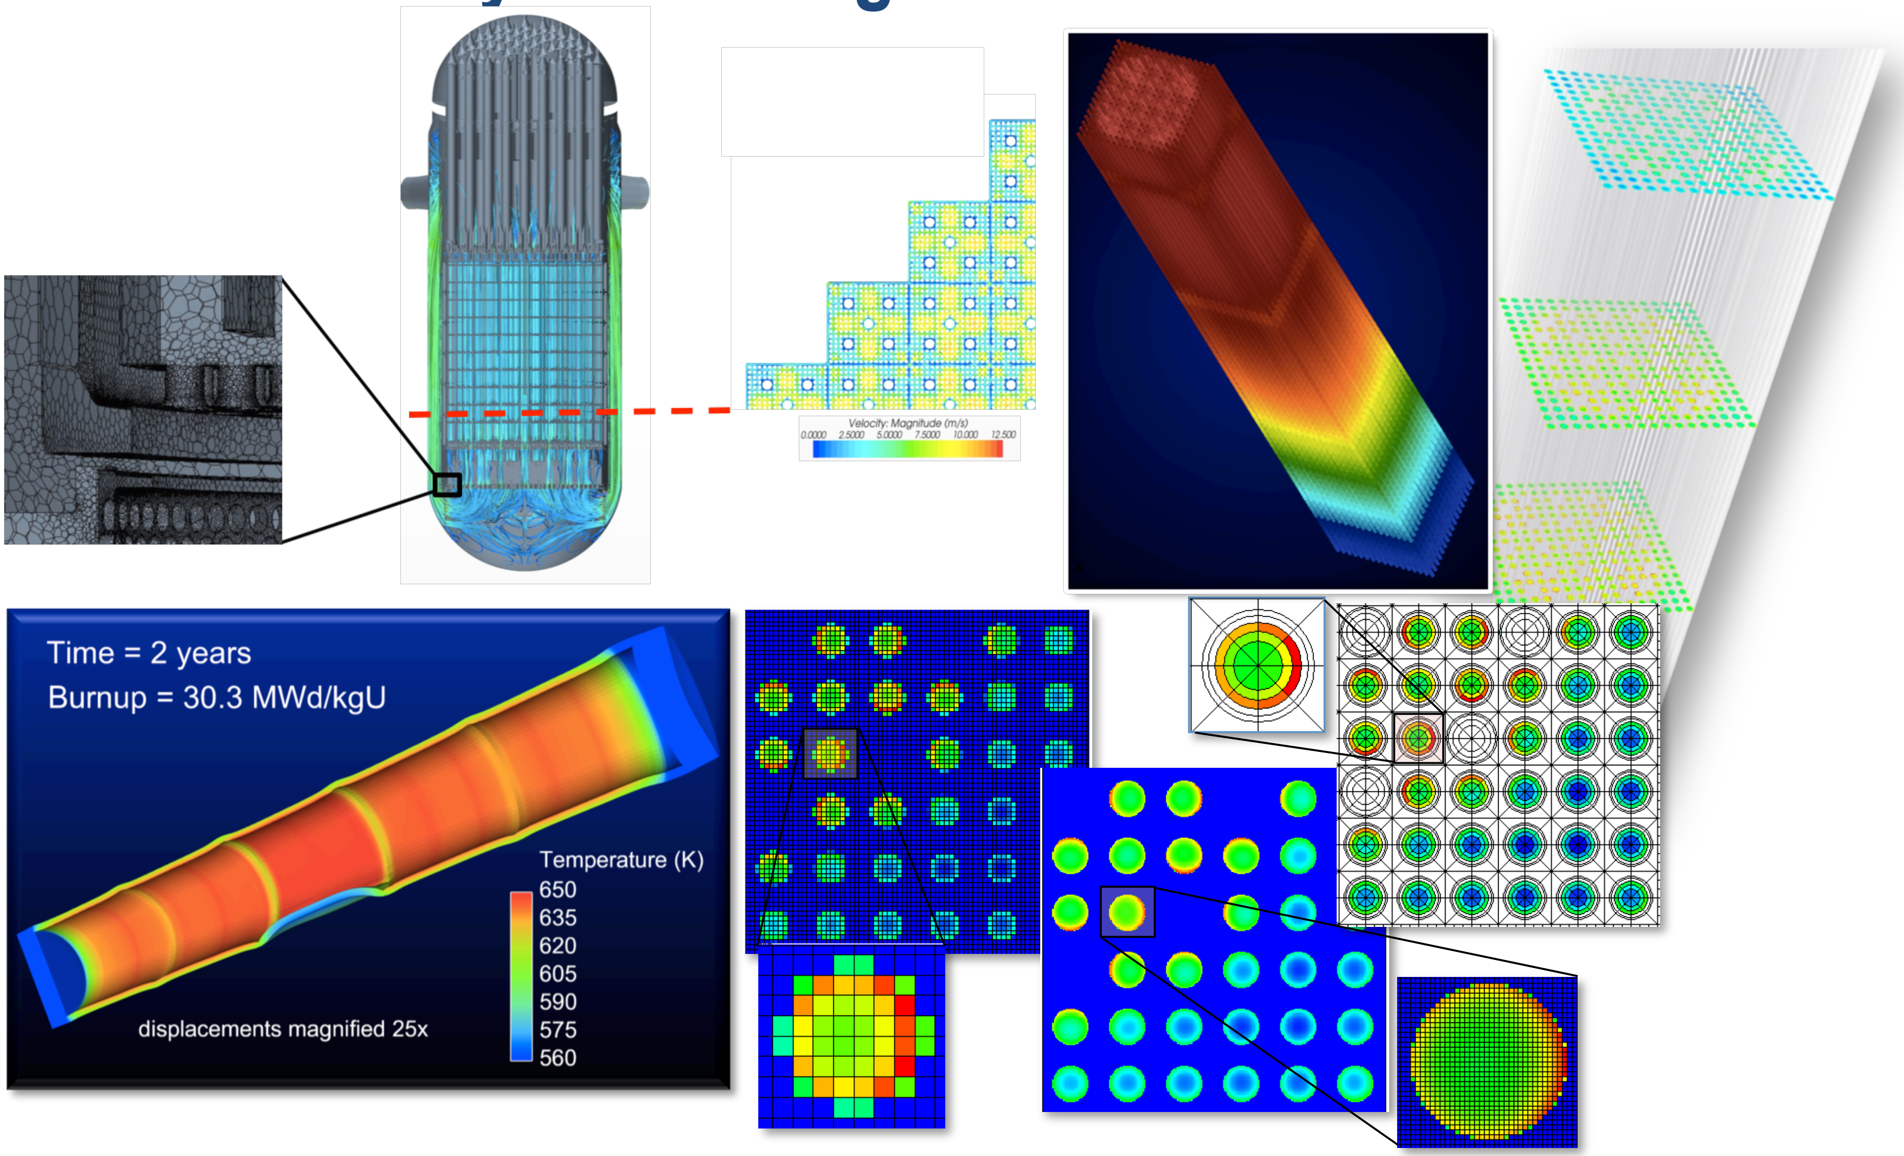

From full core to fuel assembly to fuel subassembly to fuel pin/pellet

Nuclear Applications Must Support a Wide Range of Spatial and Temporal Scales

- Nuclear fuel behavior and performance
  - Spatial scale: fuel pellet to fuel pin to fuel sub-assembly (3x3 pins)
    - From dislocations/voids/cracks (< 1 μm) to grains (<100 μm) to clad (<1 mm) to pellet (<5 cm) to pins (<4 m)
- Single-phase thermal hydraulics
  - Spatial scale: fuel sub-assembly (3x3 pins) to fuel assembly (17x17 pins)
    - From mixing vanes (<1 mm) to boundary layers (<1 cm) to turbulent structures (<10 cm) to assemblies (5 m)
- Multi-phase thermal hydraulics
  - Spatial scale: fuel assembly (17x17 pins) to full core (193 assemblies or >51K pins)
    - Same as single phase except now add bubbles (<1 mm to 1 cm) and full core (<10 m)
- Neutron transport
  - Spatial scale: fuel pellet to fuel pin to fuel assembly to full core; also 2D lattice
    - From burnable absorber layers (<1 mm) to pellet (< 1 cm) to lattice (<1 m) to full core (< 10 m)
- Coolant chemistry and CRUD deposition/buildup
  - Spatial scale: fuel pellet to fuel pin to fuel subassembly(?)
    - From oxide/hydride layers (<10 μm) to CRUD layers (< 0.1 mm) to pellets (<5 cm) to pins (<4 m)

Operational time scales: hours to days to years to decades
Safety time scales: sec to min to hours to days

Nuclear Power Reactors

Thermal hydraulics

Terascale	Petascale	Exascale
Lumped parameter models for full reactor core ▪ Calibrated subchannel models capture large scale axial flow effects - estimate transverse flow	**_High-fidelity telescoping_ of localized regions** ▪ Large-Eddy Simulations locally: sub-assembly, steam generator, upper plenum	**_High-fidelity full core_ simulations** ▪ Large-eddy simulations over entire reactor core
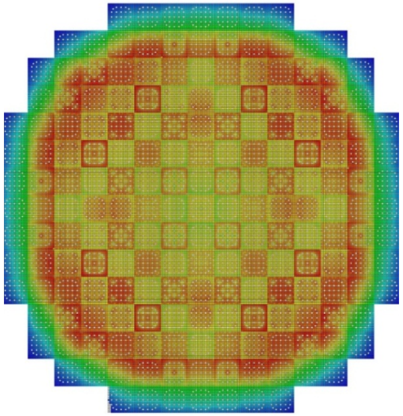	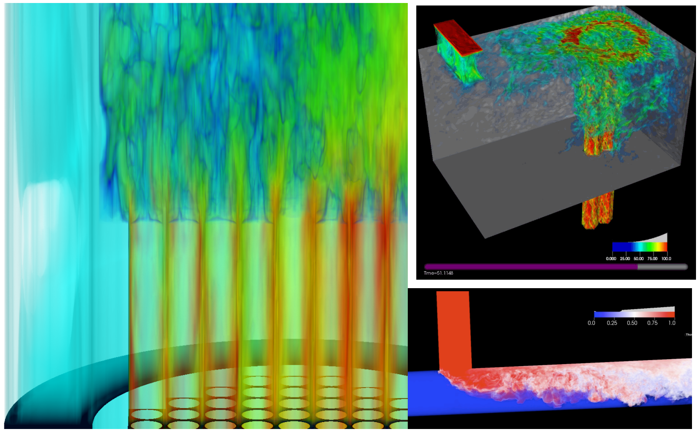	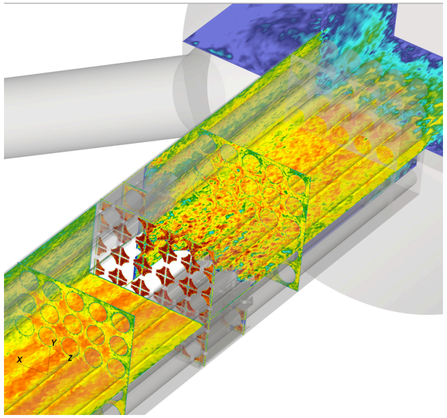
▪ Experiments required to calibrate simplified physics methods and closure models ▪ Thermal, multiphase, and boiling effects modeled – not resolved	• Reveal complicated heat transfer physics + mechanism for instabilities difficult to attain with experiments	• Experimental "data sets" of whole core effects. • Enable virtual prototyping.

Nuclear Power Reactors

Fuel performance

Terascale	Petascale	Exascale
Limited assessment of fuel pellet cladding interaction	***Capture 3D failure modes for fuel and cladding***	***Virtual test reactor for advanced fuels***
▪ Fuel type and conditions pinned to experimental database	▪ Additional experimental data on 3D effects required	▪ Requires separate effects experiments for validation
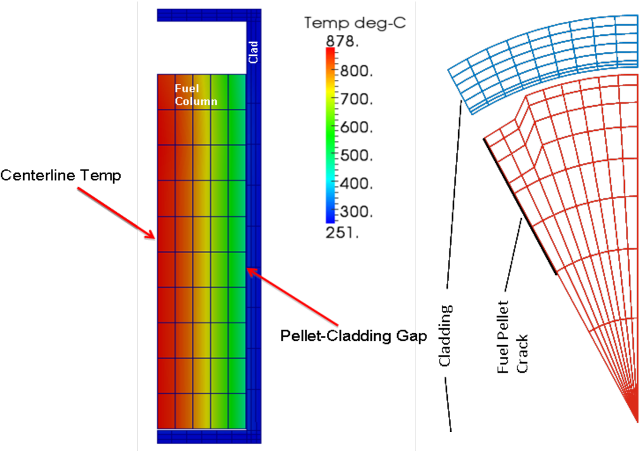	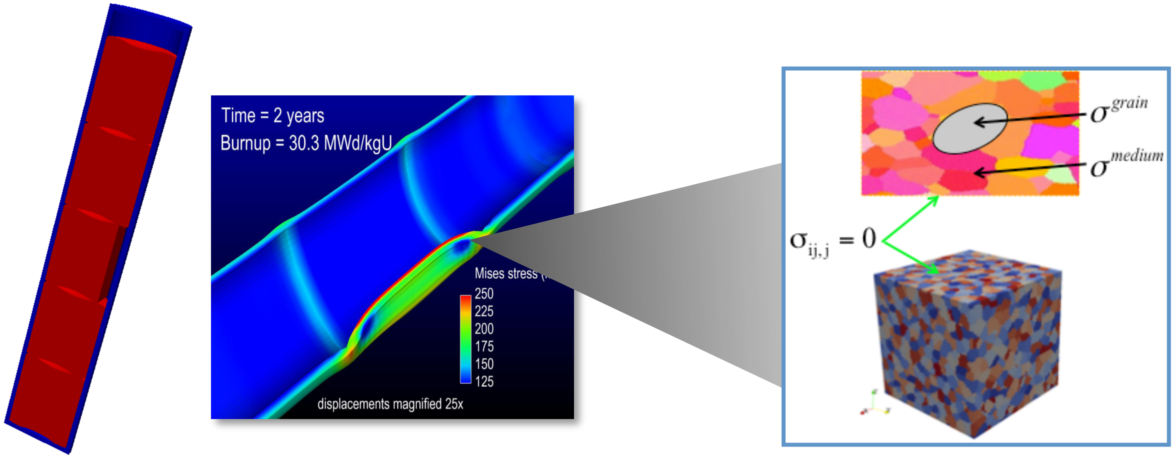	
▪ 2D finite element analysis of fuel thermo-mechanics ▪ Empirical material property and response models calibrated in 2D to experimental data	▪ 3D finite element analysis ▪ Reformulation of material models calibrated in 2D ▪ Selected upscaling of microscale phenomena (crack propagation, fission product release)	▪ Physics-based fuel performance ▪ Replacement of empirical material property and response models ▪ Expand applicability outside of test database

Nuclear Power Reactors
Coupled-Physics Core Simulations

Terascale

***Engineering** 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

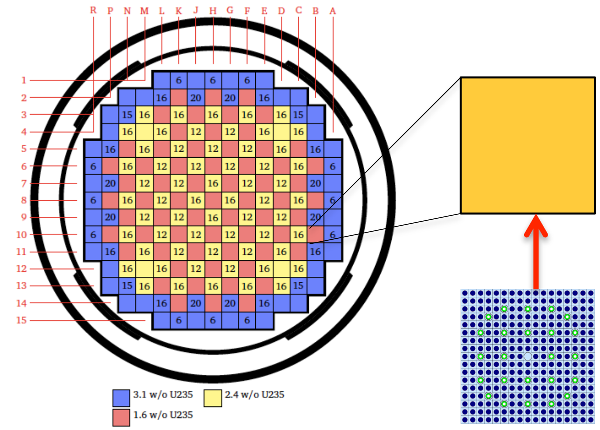

*Homogenized Fuel **Assemblies***

- 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

Petascale

***High-Fidelity** Core Analysis*

- Criticality and safety set-points
- Core pin power predictions
- Cycle isotopic fuel depletions
- Localized sub-channel feedback
- Assembly or full core structural models

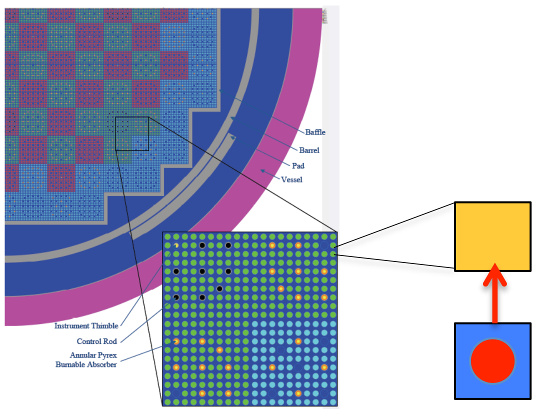

*Homogenized Fuel **Pin-Cells***

- 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

Exascale

***Extreme-Fidelity** Analysis*

- 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

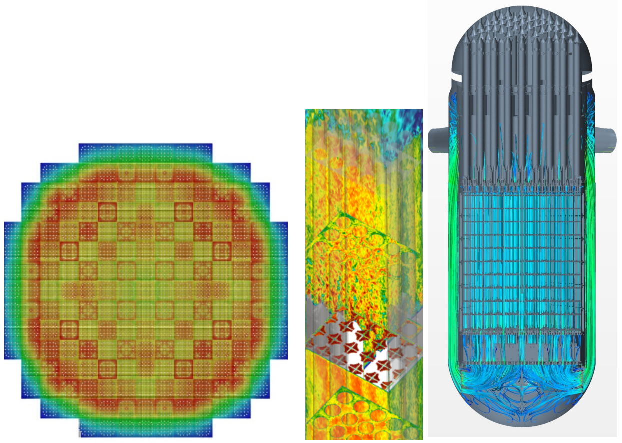

***Explicit** Fuel/Clad/Fluids &Vessel*

- 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

Solutions Realizeable at the Petascale

A Step Change in Technology

Predictive full core neutron state points for operating reactors

- Consistent method for 3D pin-resolved deterministic and Monte Carlo transport

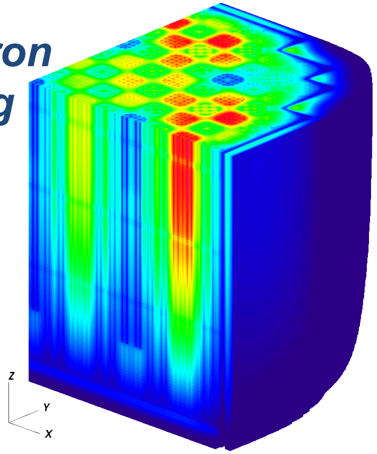

Calibrated CFD of full core turbulent multi-phase flows

- With boiling and upscaled closure models, mechanistic DNB assessment

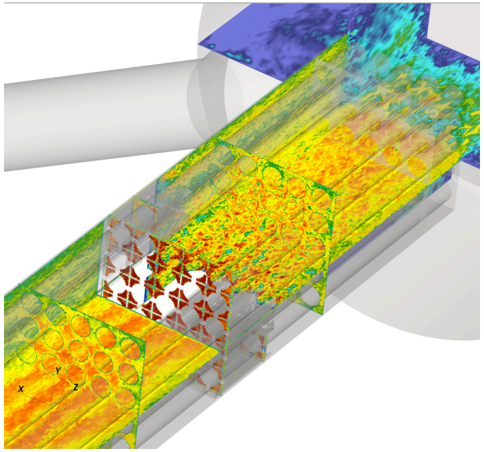

Reliable 3D assessment of nuclear fuel performance

- Inform assessments of operational risks & identify solutions (PCI,CRUD)
- Functional capability for fuel response in reactor transients (RIA, LOCA)

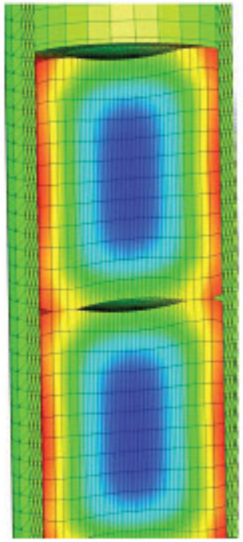

Identify vulnerabilities to operational & safety performance-limiting reactor phenomena

- Tightly-coupled multi-physics assessment of system performance
- Understand best-estimate system response and associated uncertainties upset events

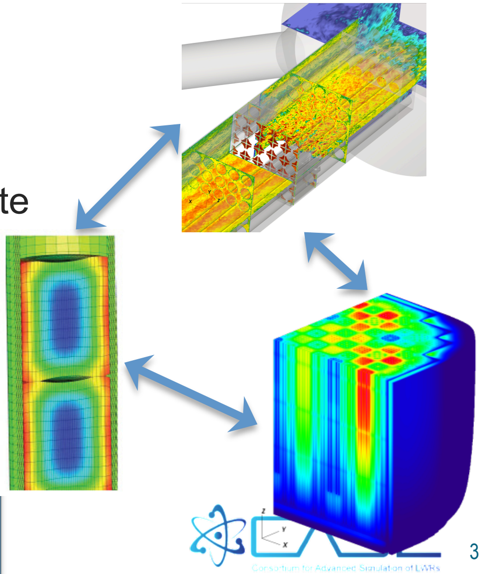

Leapfrogs calibrated industry core simulators that use lumped homogenization & correlation based closures

Solutions Expected at the Exascale

Reactor Core Physics

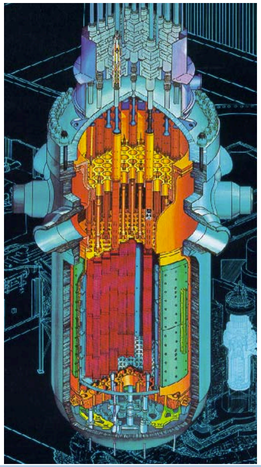

- **Predictive load-follow simulation for feedback to operational reactors for plant maneuvering and upset event recovery**
  - Core-wide multi-physics: radiation transport, fluids, fuels, chemistry, material, local wear and contacts, full core structural response
- **Expand to entire power plant**

Thermal Hydraulics

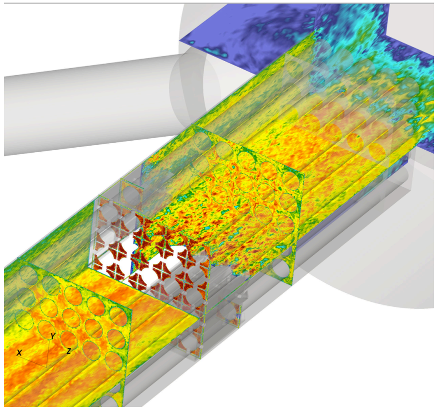

- **Enhanced heat transfer from fuel to coolant within quantified safety margins**
  - Core-wide multi-scale fluids performance with two-way DNS-to-LES CFD coupling

Nuclear Fuels

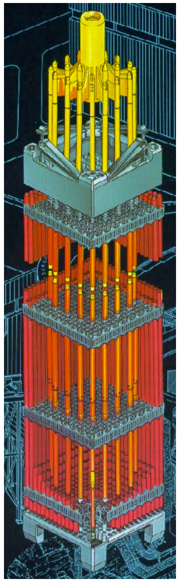

- **Investigation and suppression of barriers to higher fuel burnup**
  - Micro-scale informed fuel behavior and transient transport
- **Reliable, targeted fuel designs outside of principal test base**
  - Multi-scale fuel performance (active two-way coupling micro to meso to macro)
- **Fuel failure prediction**
  - Move from empirical failure thresholds to mechanistic models of actual failure

V&V, Data Assimilation & UQ

- **Execution of quality virtual experiments to fill validation data gaps**
  - Computational (pseudo experimental) data generated by full resolution core simulator used for data assimilation & UQ of design tools
- **Optimum experimental design & usage**
  - Integrated data assimilation, UQ & mathematical optimization to design experiments & process resulting data to recover margins to safety & operational limits

Underpinned by enabling science-based engineering models

Creating a Virtual Reactor

Enable assessment of fuel design, operation, and safety criteria

Deliver improved predictive simulation of PWR core, internals, and vessel	• Couple Virtual Reactor (VR) to evolving out-of-vessel simulation capability • Maintain applicability to other nuclear power plant (NPP) types
Execute work in 6 technical focus areas	• Equip VR with necessary physical models and multiphysics integrators • Build VR with a comprehensive, usable, and extensible software system • Validate and assess the VR models with self-consistent quantified uncertainties

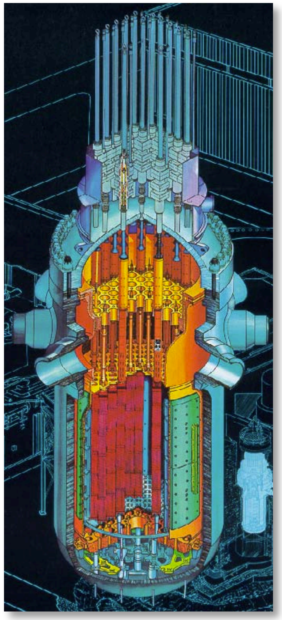

MPO
Materials performance and optimization

RTM
Radiation transport methods

THM
Thermal hydraulics methods

VUQ
Validation and uncertainty quantification

PHI
Physics integration

AMA
Advanced modeling applications

Integrated and interdependent projects span the range from basic science to application

CASL Innovations

Advanced Modeling Applications

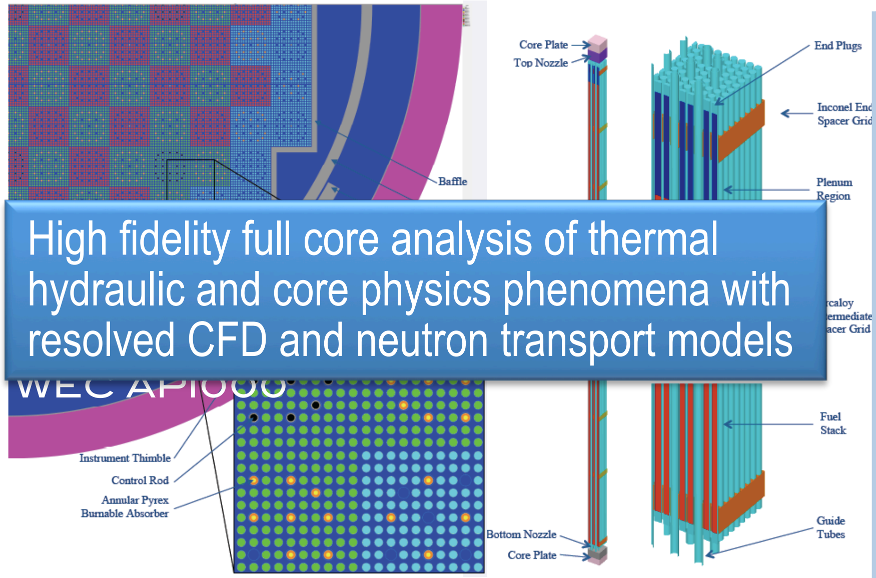

High fidelity full core analysis of thermal hydraulic and core physics phenomena with resolved CFD and neutron transport models

Physics Integration

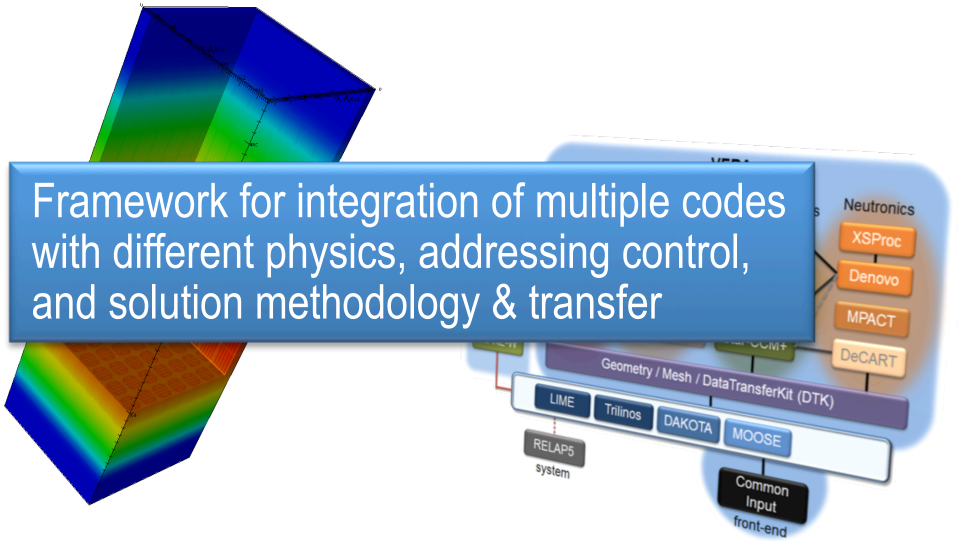

Framework for integration of multiple codes with different physics, addressing control, and solution methodology & transfer

Thermal Hydraulic Methods

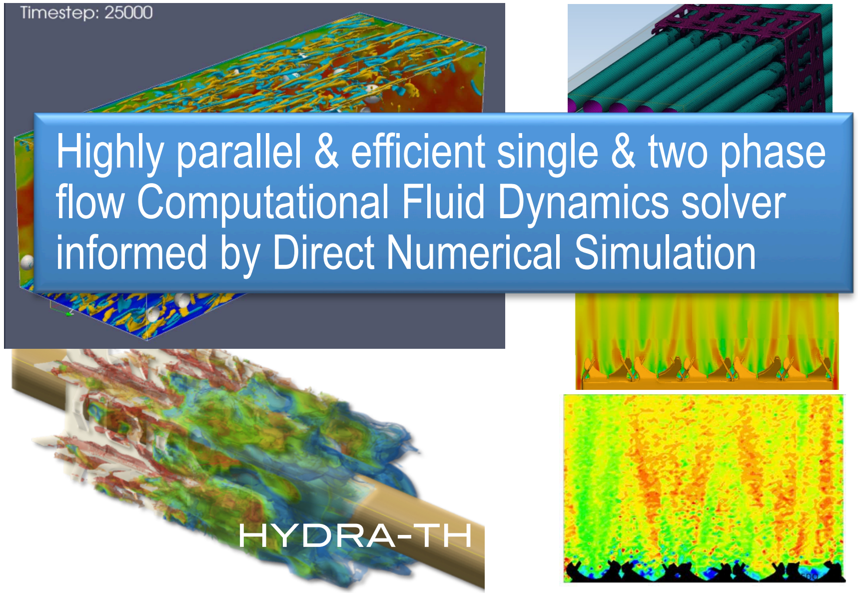

Highly parallel & efficient single & two phase flow Computational Fluid Dynamics solver informed by Direct Numerical Simulation

Radiation Transport Methods

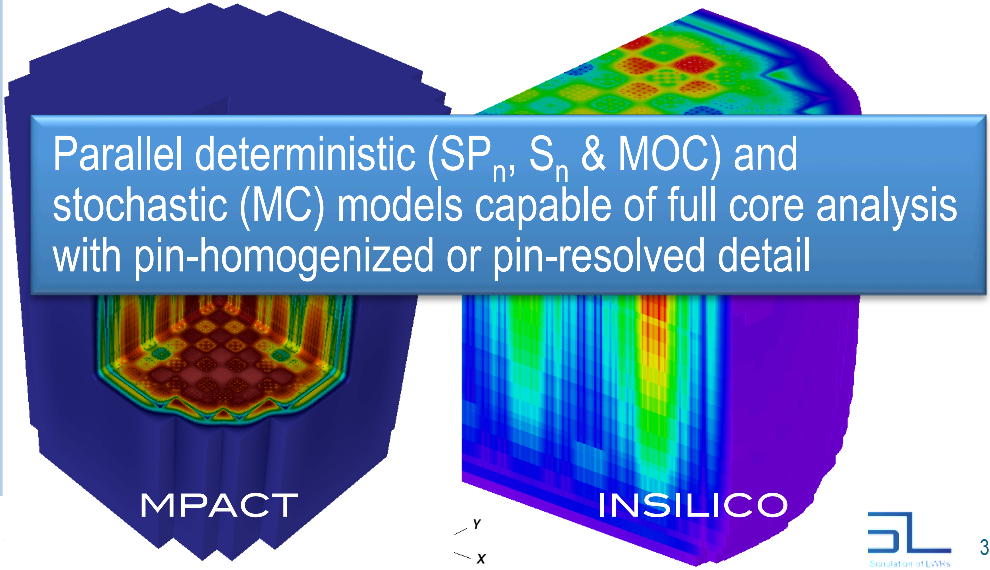

Parallel deterministic (SP_n, S_n & MOC) and stochastic (MC) models capable of full core analysis with pin-homogenized or pin-resolved detail

CASL Innovations

Validation & Uncertainty Quantification

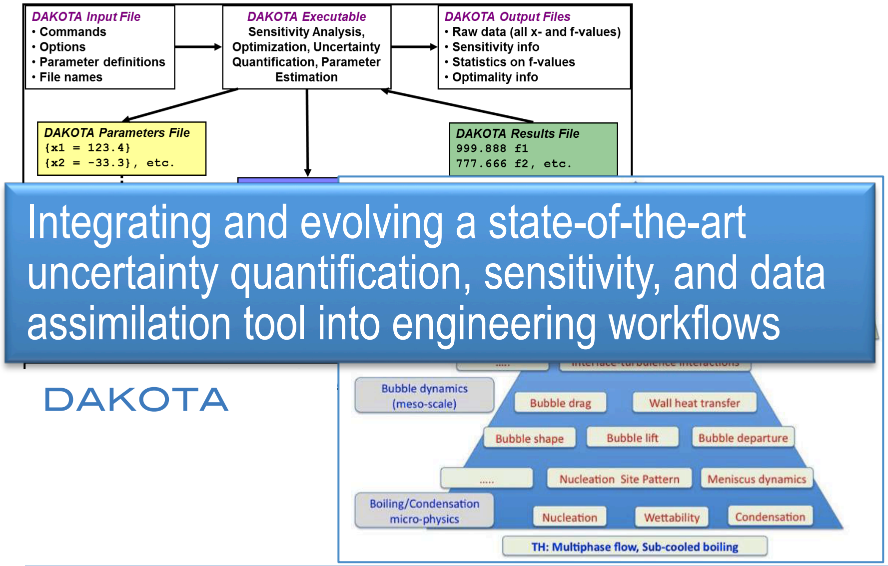

Integrating and evolving a state-of-the-art uncertainty quantification, sensitivity, and data assimilation tool into engineering workflows

VOCC

Bringing together local ("physical") and geographically distributed ("virtual") contributors in a meaningful and productive way

Materials Performance and Optimization

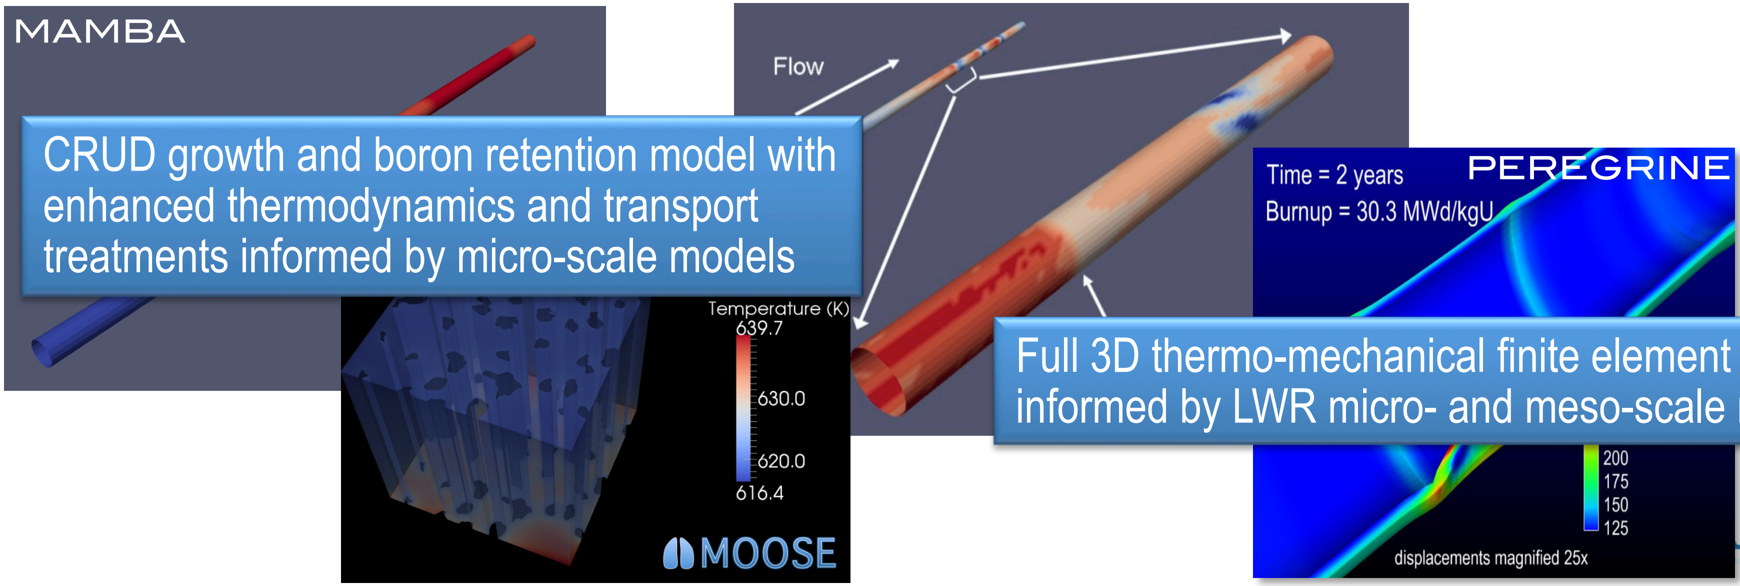

CRUD growth and boron retention model with enhanced thermodynamics and transport treatments informed by micro-scale models

Full 3D thermo-mechanical finite element model informed by LWR micro- and meso-scale models

VERA 2.1 snapshot (06/2012)

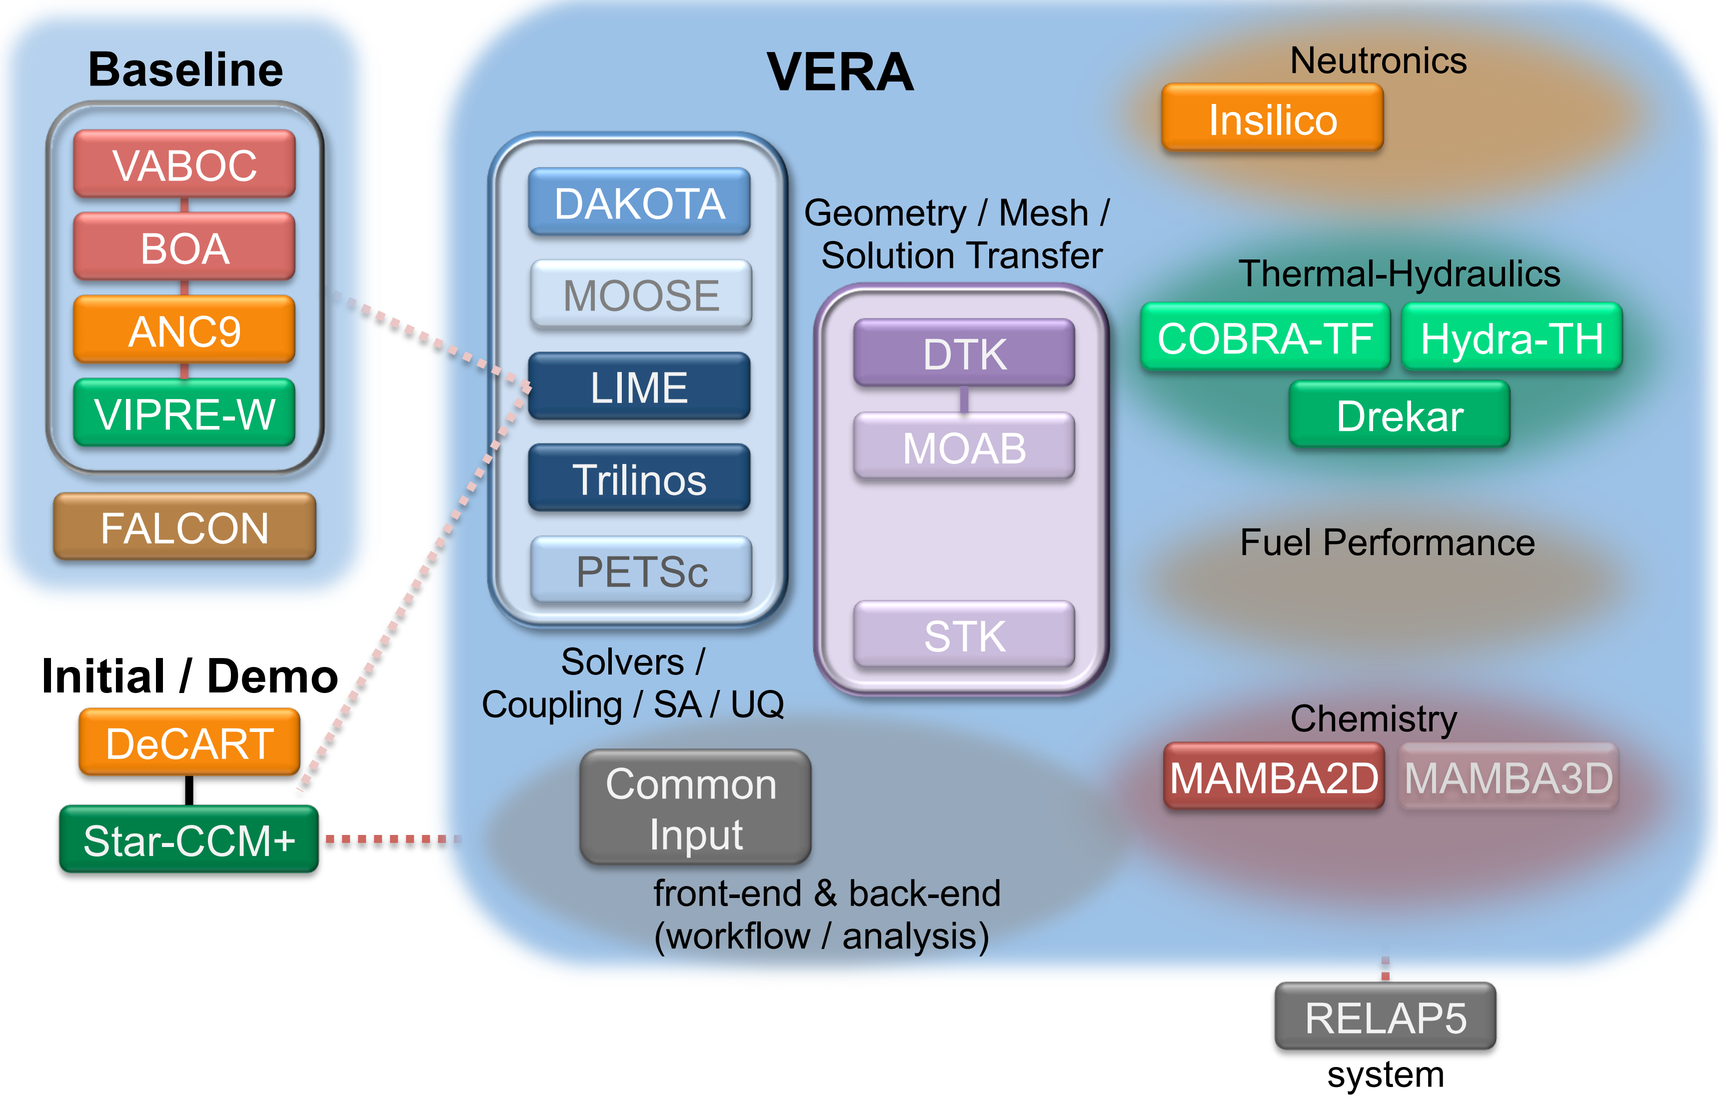

Baseline

VABOC
BOA
ANC9
VIPRE-W

FALCON

Initial / Demo

DeCART
Star-CCM+

VERA

DAKOTA
MOOSE
LIME
Trilinos
PETSc

Solvers / Coupling / SA / UQ

Geometry / Mesh / Solution Transfer

DTK
MOAB
STK

Neutronics

Insilico

Thermal-Hydraulics

COBRA-TF
Hydra-TH
Drekar

Fuel Performance

Chemistry

MAMBA2D
MAMBA3D

Common Input

front-end & back-end (workflow / analysis)

RELAP5

system

VERA 3.1 snapshot (07/2013)

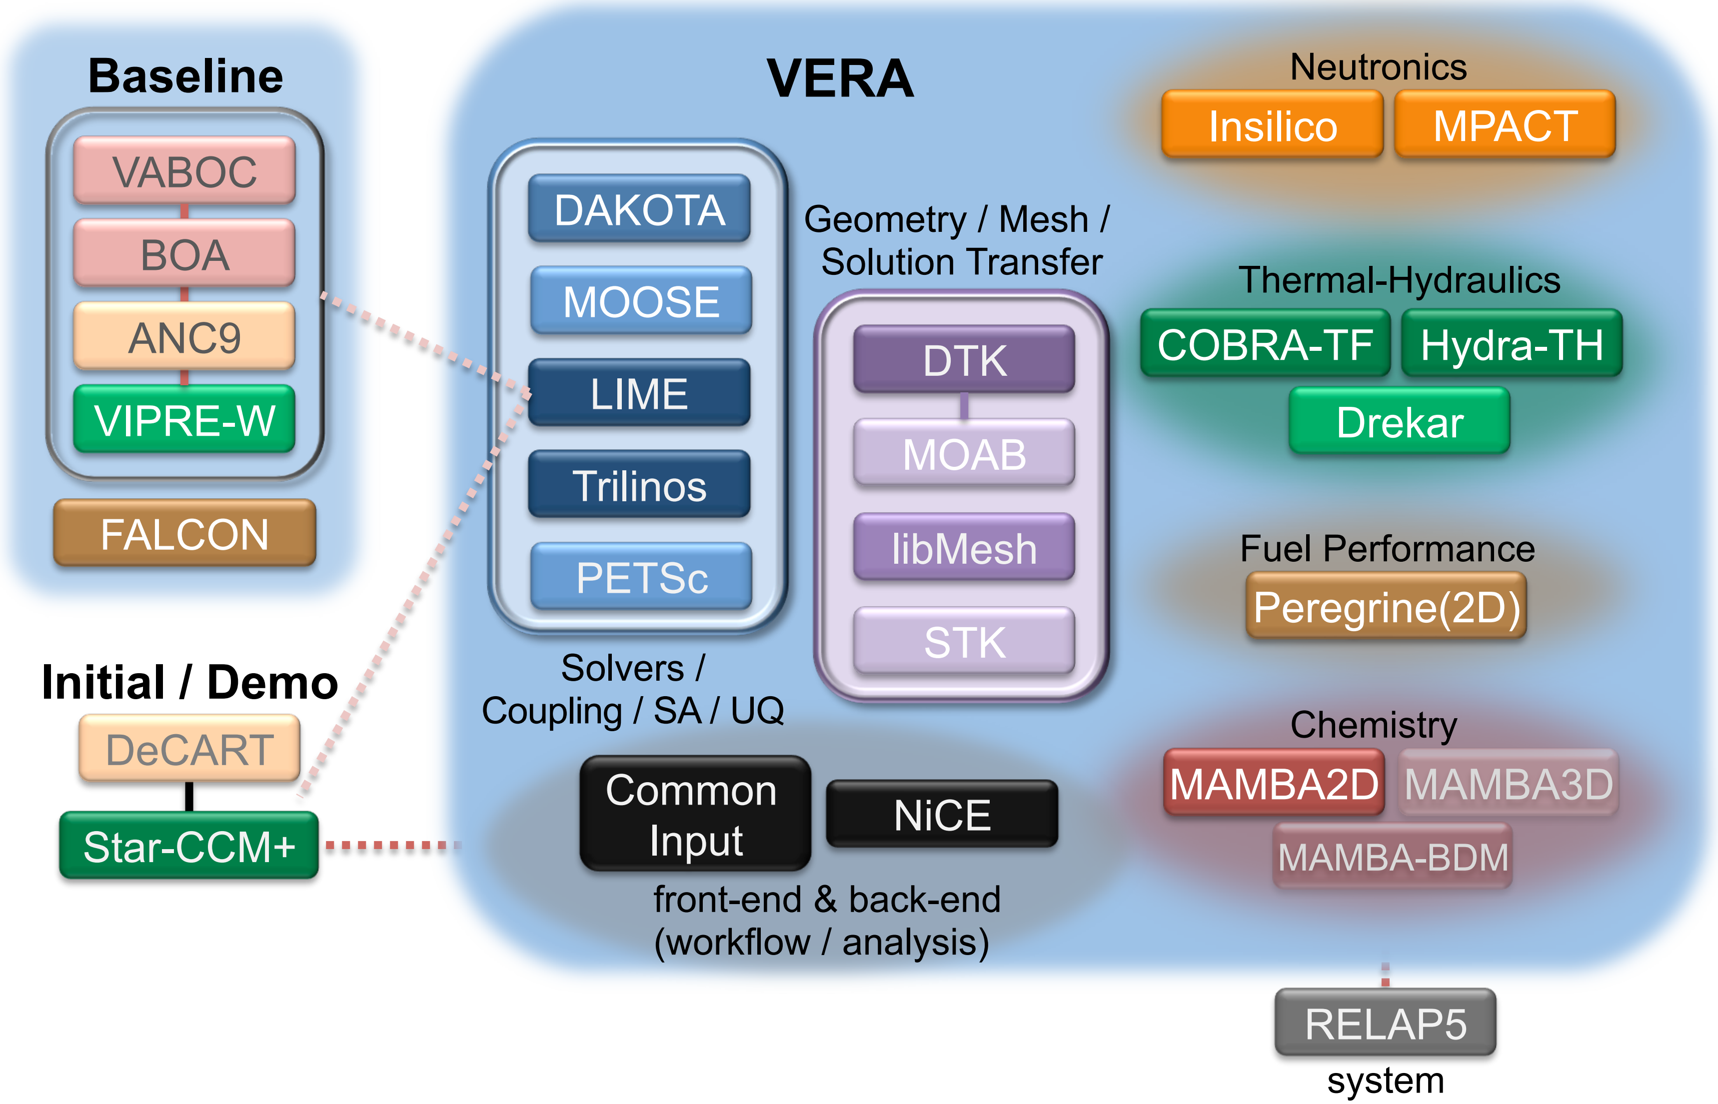

VERA-CS is a subset of VERA capabilities

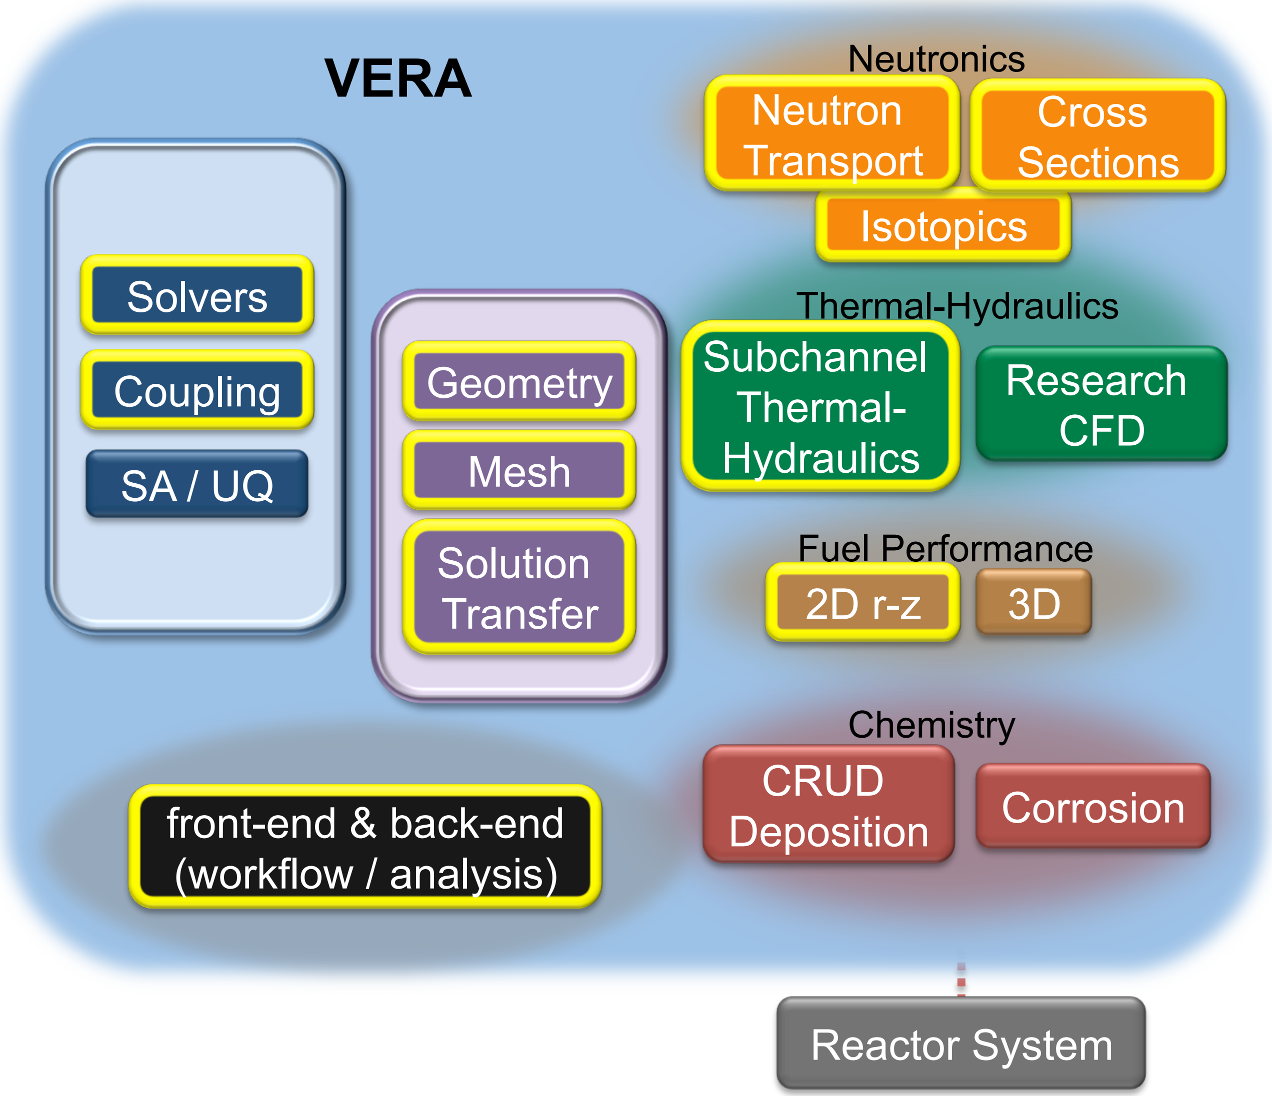

VERA

Solvers

Coupling

SA / UQ

Geometry

Mesh

Solution Transfer

front-end & back-end (workflow / analysis)

Neutronics

Neutron Transport

Cross Sections

Isotopics

Thermal-Hydraulics

Subchannel Thermal-Hydraulics

Research CFD

Fuel Performance

2D r-z

3D

Chemistry

CRUD Deposition

Corrosion

Reactor System

VERA: Virtual Environment for Reactor Applications

CASL's evolving virtual reactor for in-vessel LWR phenomena

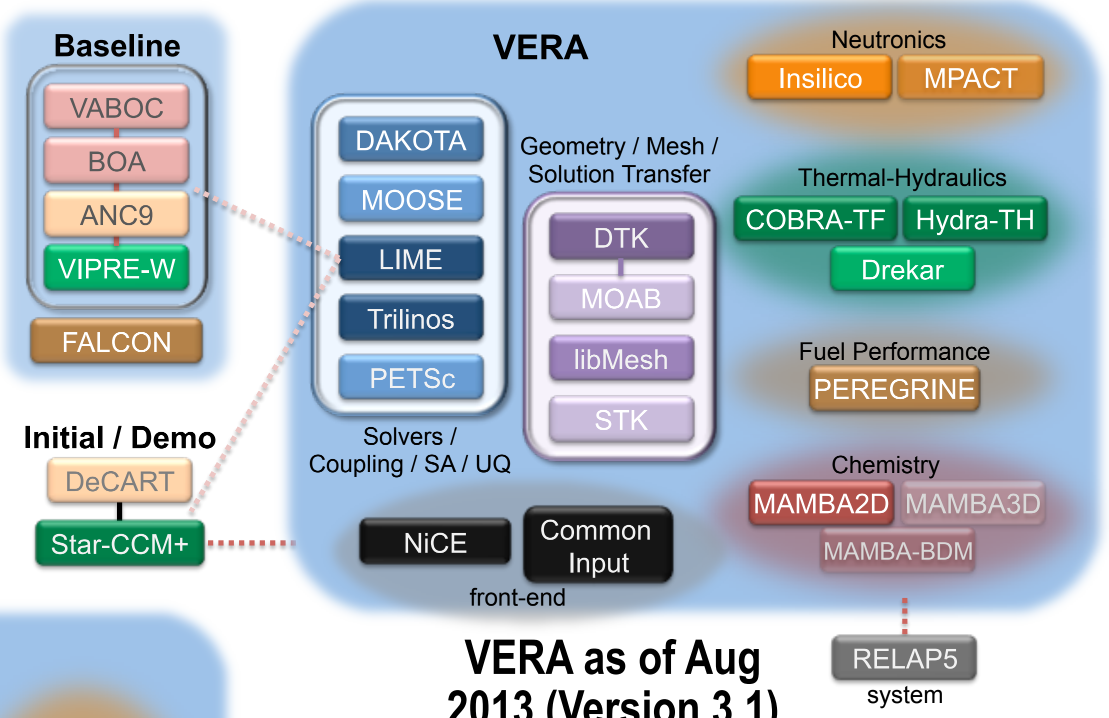

Baseline

VABOC
BOA
ANC9
VIPRE-W
FALCON

Initial / Demo

DeCART
Star-CCM+

VERA

DAKOTA
MOOSE
LIME
Trilinos
PETSc

Solvers / Coupling / SA / UQ

Geometry / Mesh / Solution Transfer

DTK
MOAB
libMesh
STK

Neutronics

Insilico
MPACT

Thermal-Hydraulics

COBRA-TF
Hydra-TH
Drekar

Fuel Performance

PEREGRINE

Chemistry

MAMBA2D
MAMBA3D
MAMBA-BDM

NiCE
Common Input

front-end

RELAP5
system

VERA as of Aug 2013 (Version 3.1)

Required functional capabilities

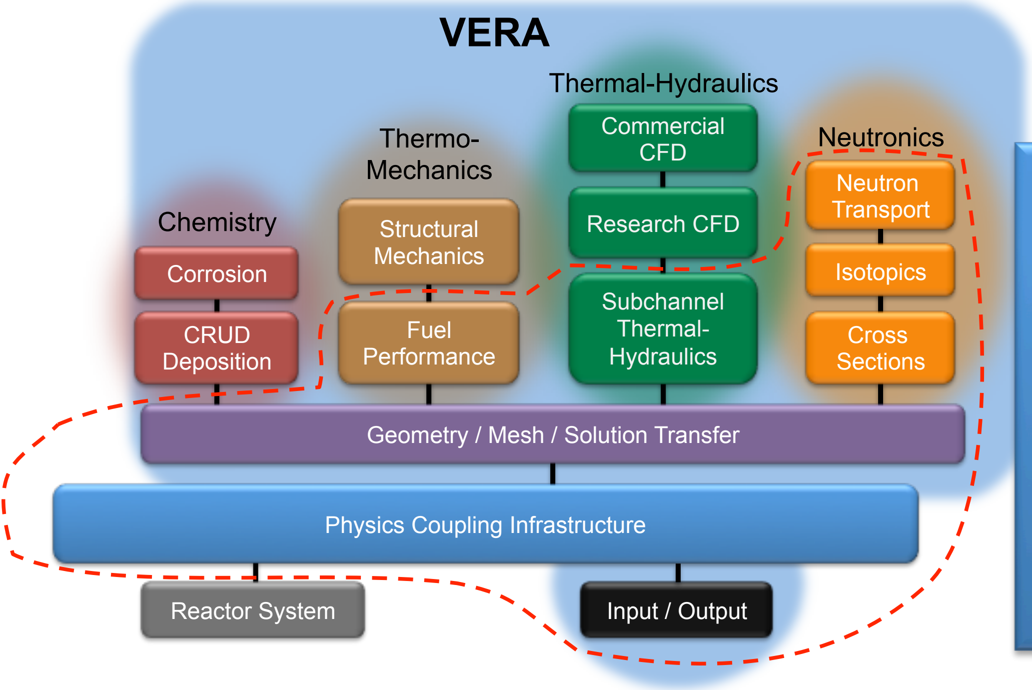

VERA

Chemistry

Corrosion
CRUD Deposition

Thermo-Mechanics

Structural Mechanics
Fuel Performance

Thermal-Hydraulics

Commercial CFD
Research CFD
Subchannel Thermal-Hydraulics

Neutronics

Neutron Transport
Isotopics
Cross Sections

Geometry / Mesh / Solution Transfer

Physics Coupling Infrastructure

Reactor System

Input / Output

CASL has 3 M&S technology products

1. VERA-CS as the fast running core simulator, which has value both standalone and for providing power histories, etc for more detailed codes
2. Engineering suite of standalone codes with ability to couple 2 or more within VERA or in other environments
3. Leadership suite of high fidelity codes used to drive improvements in 1 and 2

CASL Innovations

CASL vs. Industry Core Simulators

Physics Model	Industry Practice	CASL (VERA-CS)
Neutron Transport	3-D diffusion (core) 2 energy groups (core) 2-D transport on single assy	3-D transport 23+ energy groups
Power Distribution	nodal average with pin-power reconstruction methods	explicit pin-by-pin
Thermal-Hydraulics	1-D assembly-averaged	subchannel (w/crossflow)
Fuel Temperatures	nodal average	pin-by-pin 2-D or 3-D
Xenon/Samarium	nodal average w/correction	pin-by-pin
Depletion	infinite-medium cross sections quadratic burnup correction history corrections spectral corrections reconstructed pin exposures	pin-by-pin with actual core conditions
Reflector Models	1-D cross section models	actual 3-D geometry
Target Platforms	workstation (single-core)	1,000 – 300,000 cores

CASL current and planned capabilities will leapfrog calibrated industry core simulators that use lumped homogenization and correlation-based closures

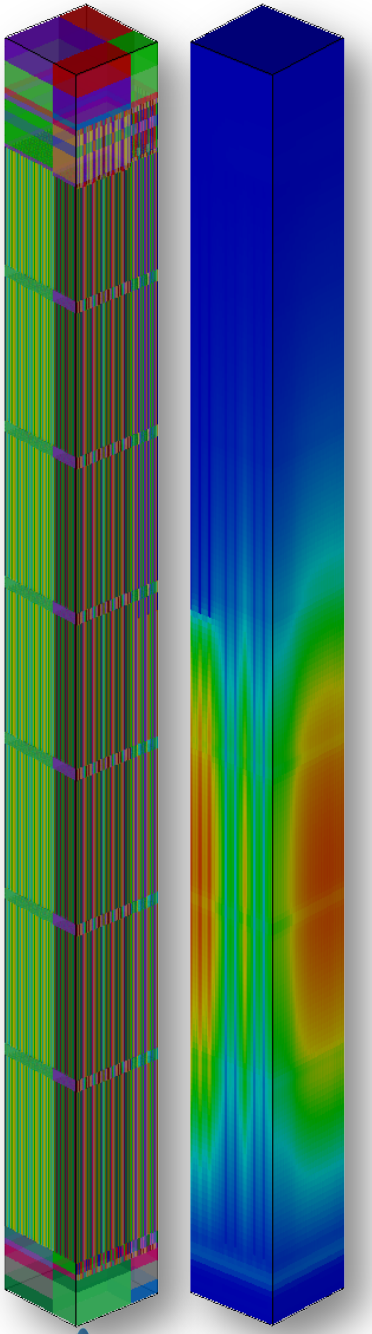

CASL

Consortium for Advanced Simulation of LWRs

VERA can appear complex...

essentially a collection of capabilities sharing infrastructure

Baseline
ANC + VIPRE-W
ANC + VIPRE-W + BOA
ANC + VIPRE-W + BOA + VABOC

Initial and/or Demonstrations
DeCART
DeCART + Star-CCM+
DeCART + Star-CCM+ + MAMBA3D

Leveraged / External
Star-CCM+
SIERRA
RELAP5
RELAP7

VERA (**VERA-CS in bold**)	Status
Insilico (XSProc,Denovo,SP_N,Shift,etc)	**deployed**
Drekar	integrated
Hydra-TH	integrated
COBRA-TF (CTF)	**integrated**
Insilico + Drekar	demo
Insilico + CTF	**integrated**
Insilico + CTF + MAMBA2D	under dev.
Insilico + CTF + Peregrine(2D)	**testing**
MPACT	**testing**
MPACT + CTF	**under dev.**
MPACT + CTF + MAMBA2D	future
MPACT + CTF + Peregrine(2D)	**future**
MPACT + Hydra-TH	future
etc...	

number of repositories	18
number of TriBITS VERA pkgs	184
repository source code statistics	43k files, 8M source
number of VERA-specific tests	504
number of nightly builds	Baseline: 3, VERA: 2
compilers/platforms tested nightly	Baseline: GCC 4.5.1 + Intel 11.1.064 VERA: GCC 4.6.1

TriBITS facilitates mgmt. of software complexity. Third-party libraries (TPLs) are a particular challenge, esp. as we couple more components.

VERA input is comfortable for current industry users and extensible.

- ability to create, archive, compare, and modify input similar to current industry workflows
- attributes of real reactors
  - assemblies, poisons, control rods, non-fuel structures, baffle, power, flow, depletion, boron search, detectors, etc.
- eliminate inconsistencies between physics components through use of a common geometry description
- will evolve as needed
- **currently using VERA input**
  - **Insilico (S_N, SP_N, Monte Carlo)**
  - **COBRA-TF**
  - **MPACT**
  - **Peregrine**

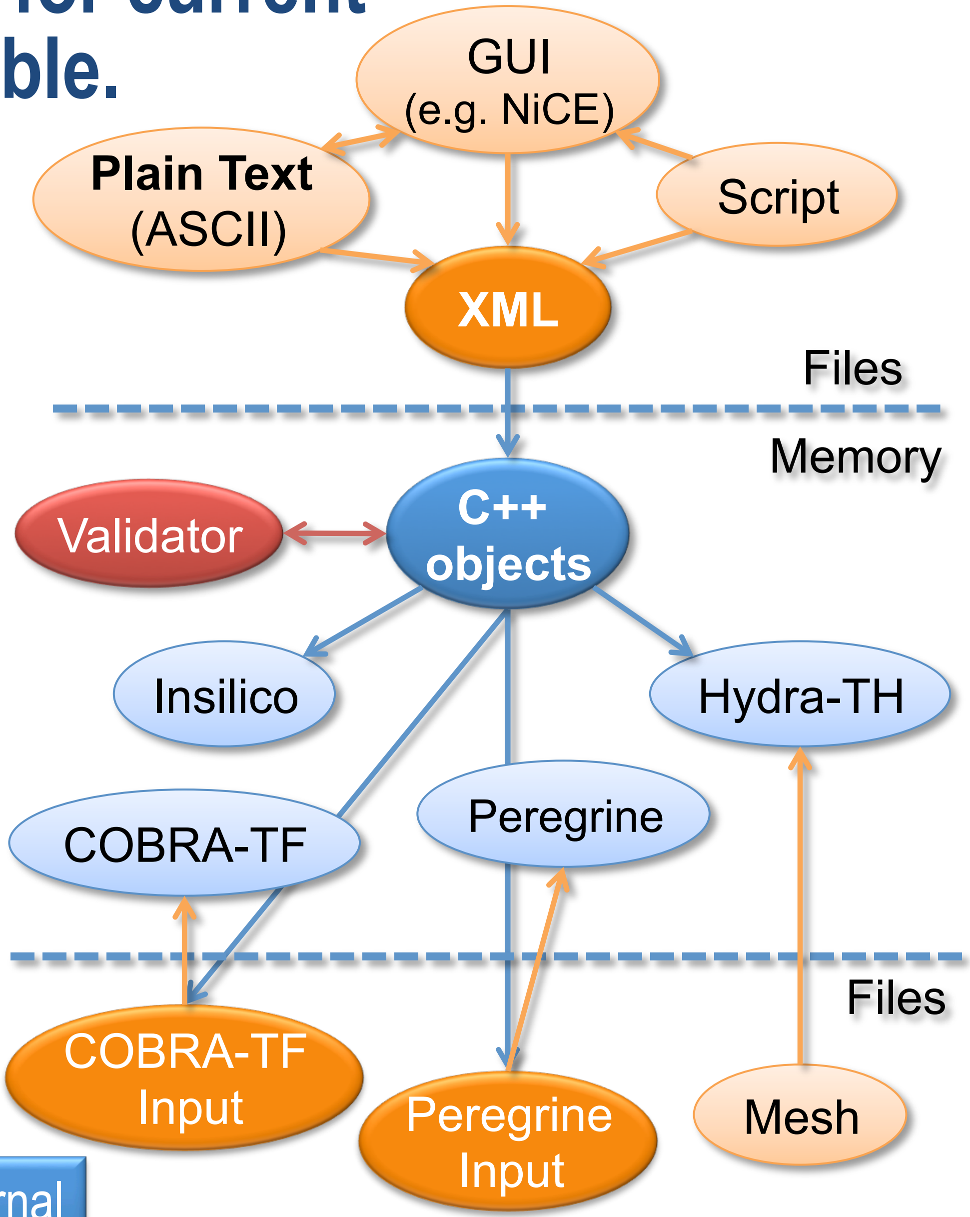

Discussing with VUQ how best to expose internal component model parameters for SA/UQ.

AMA Progression Problem 3: Hot Zero Power Assy.

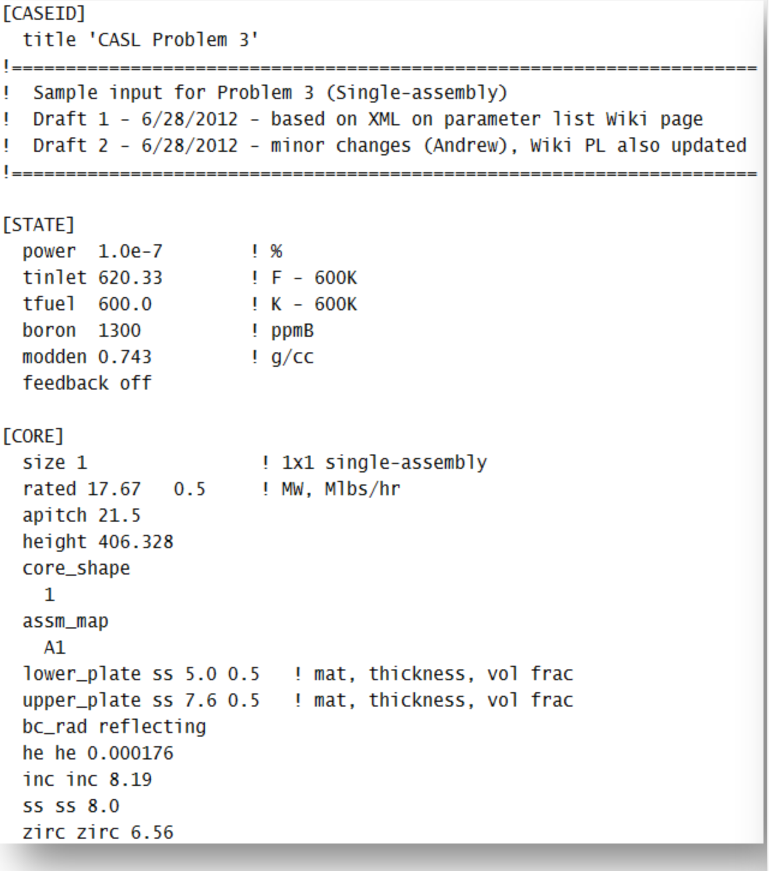

```
[CASEID]
  title 'CASL Problem 3'
!=======================================================================
!  Sample input for Problem 3 (Single-assembly)
!  Draft 1 - 6/28/2012 - based on XML on parameter list Wiki page
!  Draft 2 - 6/28/2012 - minor changes (Andrew), Wiki PL also updated
!=======================================================================

[STATE]
  power  1.0e-7        ! %
  tinlet 620.33        ! F - 600K
  tfuel  600.0         ! K - 600K
  boron  1300          ! ppmB
  modden 0.743         ! g/cc
  feedback off

[CORE]
  size 1                ! 1x1 single-assembly
  rated 17.67   0.5     ! MW, Mlbs/hr
  apitch 21.5
  height 406.328
  core_shape
    1
  assm_map
    A1
  lower_plate ss 5.0 0.5   ! mat, thickness, vol frac
  upper_plate ss 7.6 0.5   ! mat, thickness, vol frac
  bc_rad reflecting
  he he 0.000176
  inc inc 8.19
  ss ss 8.0
  zirc zirc 6.56
```

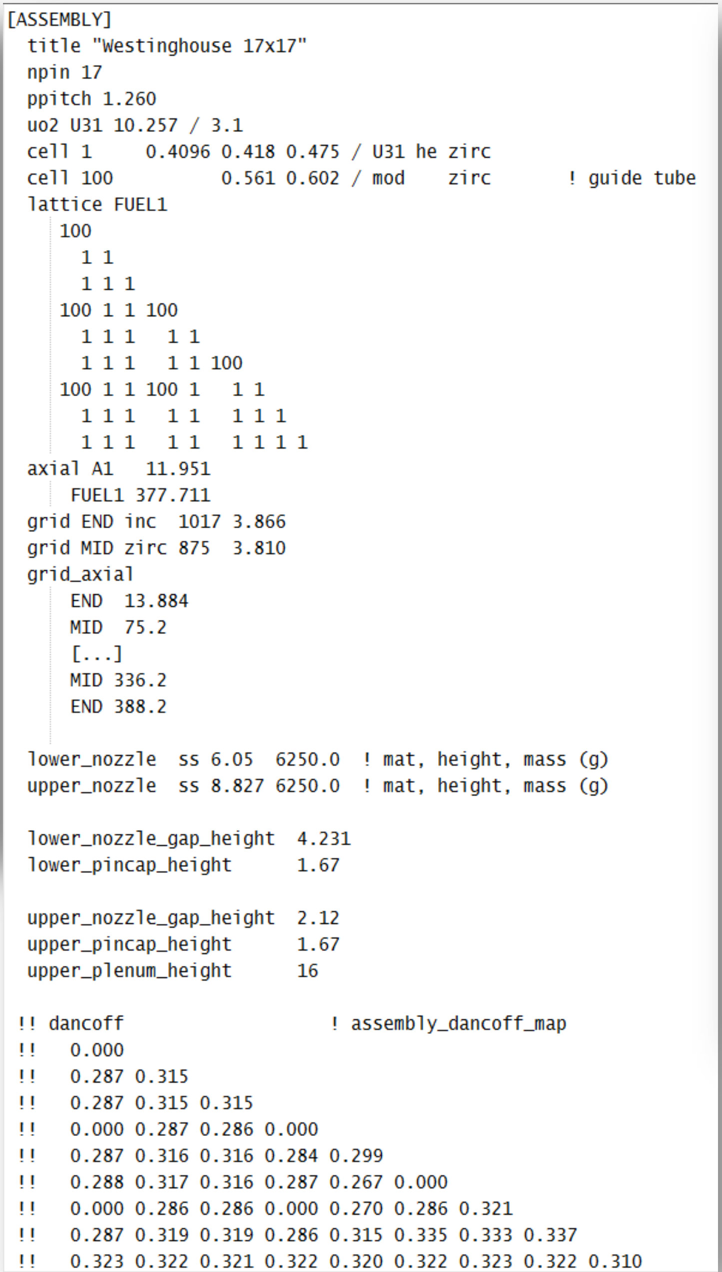

```
[ASSEMBLY]
  title "Westinghouse 17x17"
  npin 17
  ppitch 1.260
  uo2 U31 10.257 / 3.1
  cell 1     0.4096 0.418 0.475 / U31 he zirc
  cell 100          0.561 0.602 / mod    zirc       ! guide tube
  lattice FUEL1
     100
       1 1
       1 1 1
     100 1 1 100
       1 1 1   1 1
       1 1 1   1 1 100
     100 1 1 100 1   1 1
       1 1 1   1 1   1 1 1
       1 1 1   1 1   1 1 1 1
  axial A1    11.951
      FUEL1 377.711
  grid END inc  1017 3.866
  grid MID zirc 875  3.810
  grid_axial
      END  13.884
      MID  75.2
      [...]
      MID 336.2
      END 388.2

  lower_nozzle  ss 6.05  6250.0  ! mat, height, mass (g)
  upper_nozzle  ss 8.827 6250.0  ! mat, height, mass (g)

  lower_nozzle_gap_height  4.231
  lower_pincap_height      1.67

  upper_nozzle_gap_height  2.12
  upper_pincap_height      1.67
  upper_plenum_height      16

!! dancoff                    ! assembly_dancoff_map
!!    0.000
!!    0.287 0.315
!!    0.287 0.315 0.315
!!    0.000 0.287 0.286 0.000
!!    0.287 0.316 0.316 0.284 0.299
!!    0.288 0.317 0.316 0.287 0.267 0.000
!!    0.000 0.286 0.286 0.000 0.270 0.286 0.321
!!    0.287 0.319 0.319 0.286 0.315 0.335 0.333 0.337
!!    0.323 0.322 0.321 0.322 0.320 0.322 0.323 0.322 0.310
```

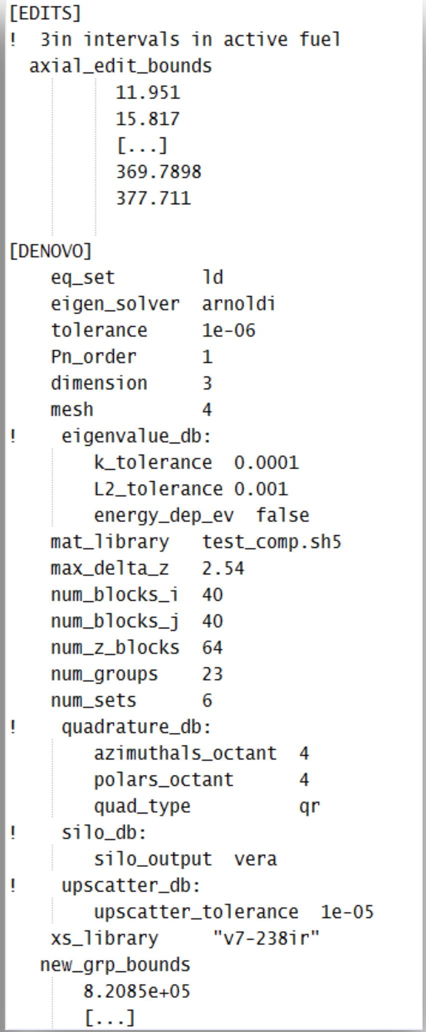

```
[EDITS]
!  3in intervals in active fuel
  axial_edit_bounds
          11.951
          15.817
          [...]
          369.7898
          377.711

[DENOVO]
    eq_set        ld
    eigen_solver  arnoldi
    tolerance     1e-06
    Pn_order      1
    dimension     3
    mesh          4
!    eigenvalue_db:
       k_tolerance  0.0001
       L2_tolerance 0.001
       energy_dep_ev  false
    mat_library   test_comp.sh5
    max_delta_z   2.54
    num_blocks_i  40
    num_blocks_j  40
    num_z_blocks  64
    num_groups    23
    num_sets      6
!    quadrature_db:
       azimuthals_octant  4
       polars_octant      4
       quad_type          qr
!    silo_db:
       silo_output  vera
!    upscatter_db:
       upscatter_tolerance  1e-05
    xs_library     "v7-238ir"
   new_grp_bounds
      8.2085e+05
      [...]
```

VERA Common Output

- fine-mesh results written to **SILO** files for visualization in tools such as VisIt / ParaView
- pin-by-pin distributions (from multiple codes) written to a common **HDF5** format that can be post-processed to create user edits
  - 2D/3D pin distributions
  - 2D/3D assembly distributions
  - peaking factors
  - Compare distributions (e.g. Keno vs. VERA)
- recognition that industrial users need **both** visualization and “real numbers”

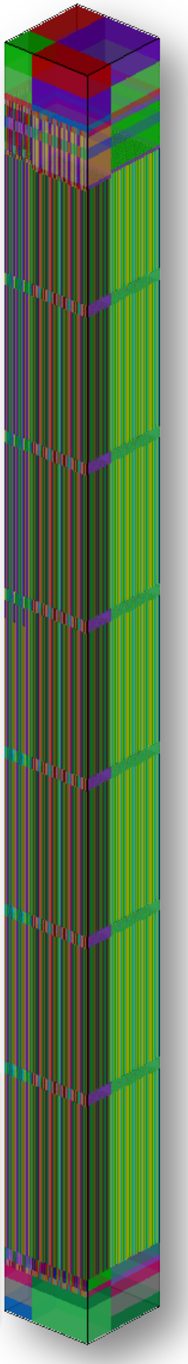

Insilico fission rate for full assembly, generated SILO file and VisIt

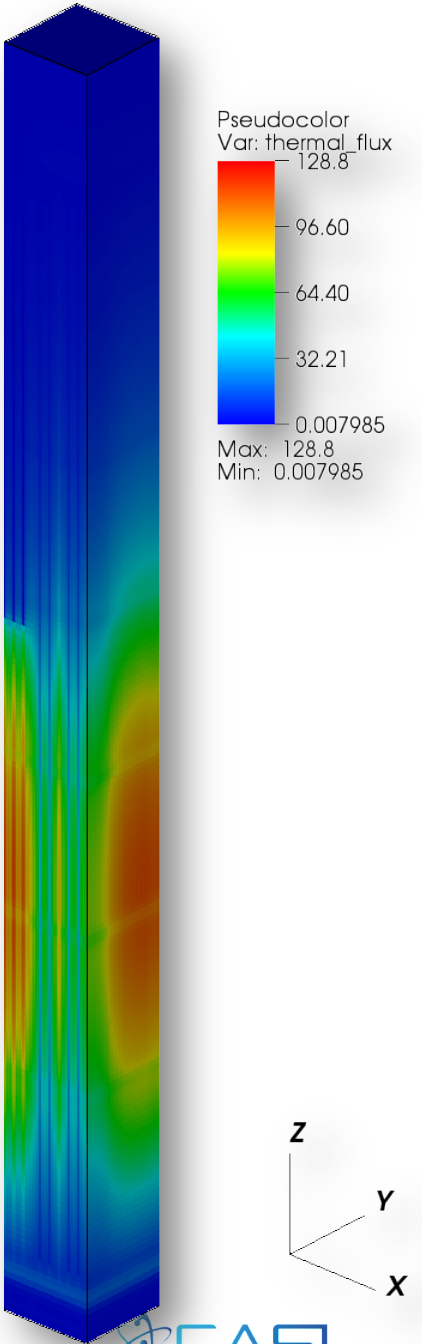

Coupling Challenges and Solutions

- **Challenges**
  - Many codes assume they are the “master”
  - Conflicting dependencies and build systems
  - Existing codes that have a life of their own outside CASL
  - Multiple languages (primarily Fortran and C++)
  - Disparate input and output formats and conventions
  - Different meshes and discretizations
- **Solutions**
  - Common build system that extends widely-used standards
  - Philosophy of continuous integration (catch and fix issues and conflicts as early as possible)
  - Standardize input / output (and restart)
  - Develop infrastructure components as necessary (e.g. DTK for solution transfer)

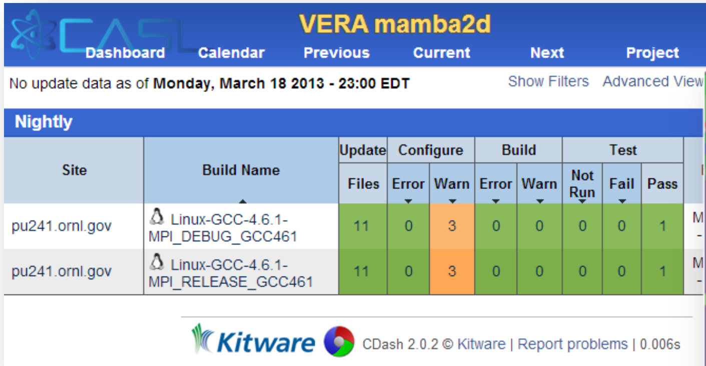

No update data as of **Monday, March 18 2013 - 23:00 EDT**

Nightly

Site	Build Name	Update	Configure		Build		Test		
		Files	Error	Warn	Error	Warn	Not Run	Fail	Pass
pu241.ornl.gov	Linux-GCC-4.6.1-MPI_DEBUG_GCC461	11	0	3	0	0	0	0	1
pu241.ornl.gov	Linux-GCC-4.6.1-MPI_RELEASE_GCC461	11	0	3	0	0	0	0	1

Kitware CDash 2.0.2 © Kitware | Report problems | 0.006s

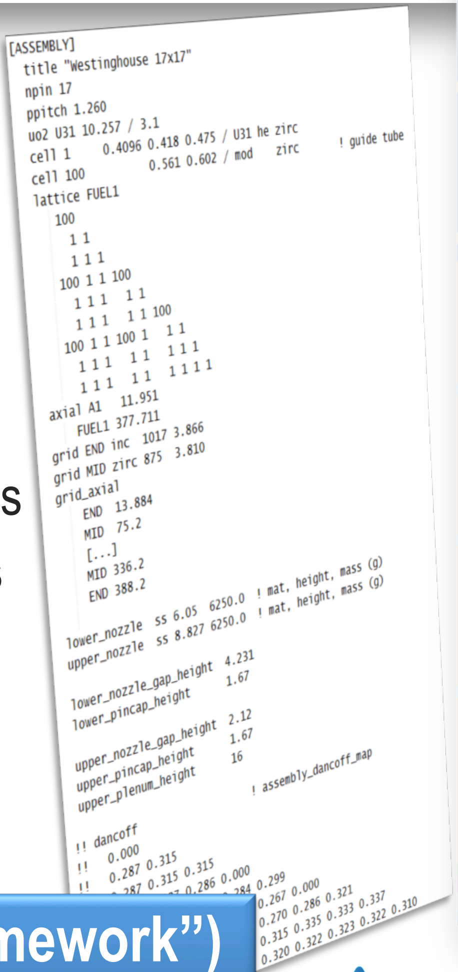

```
[ASSEMBLY]
  title "Westinghouse 17x17"
  npin 17
  ppitch 1.260
  uo2 U31 10.257 / 3.1
  cell 1     0.4096 0.418 0.475 / U31 he zirc
  cell 100          0.561 0.602 / mod   zirc      ! guide tube
  lattice FUEL1
      100
        1 1
        1 1 1
      100 1 1 100
        1 1 1   1 1
        1 1 1   1 1 100
      100 1 1 100 1   1 1
        1 1 1   1 1   1 1 1
        1 1 1   1 1   1 1 1 1
  axial A1   11.951
      FUEL1 377.711
  grid END inc  1017 3.866
  grid MID zirc 875  3.810
  grid_axial
      END  13.884
      MID  75.2
      [...]
      MID 336.2
      END 388.2

  lower_nozzle  ss 6.05  6250.0  ! mat, height, mass (g)
  upper_nozzle  ss 8.827 6250.0  ! mat, height, mass (g)

  lower_nozzle_gap_height  4.231
  lower_pincap_height      1.67

  upper_nozzle_gap_height  2.12
  upper_pincap_height      1.67
  upper_plenum_height      16
                                 ! assembly_dancoff_map

!! dancoff
!!    0.000
!!    0.287 0.315
!!    0.287 0.315 0.315
                 0.286 0.000
                       0.284 0.299
                             0.267 0.000
                             0.270 0.286 0.321
                             0.315 0.335 0.333 0.337
                             0.320 0.322 0.323 0.322 0.310
```

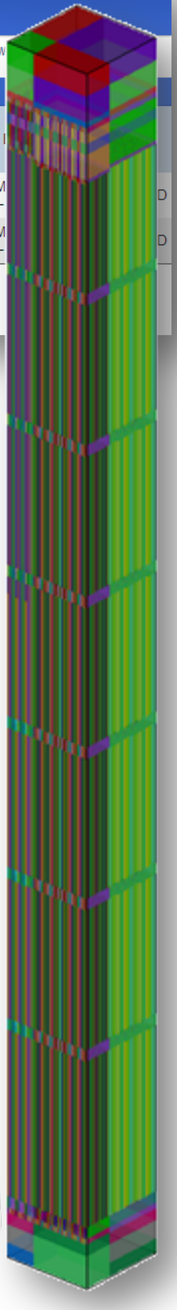

The “mechanics” of nonlinear iteration (the “framework”) is not the most challenging aspect of coupling.

Core Simulator Progression Problems Drive VERA Development

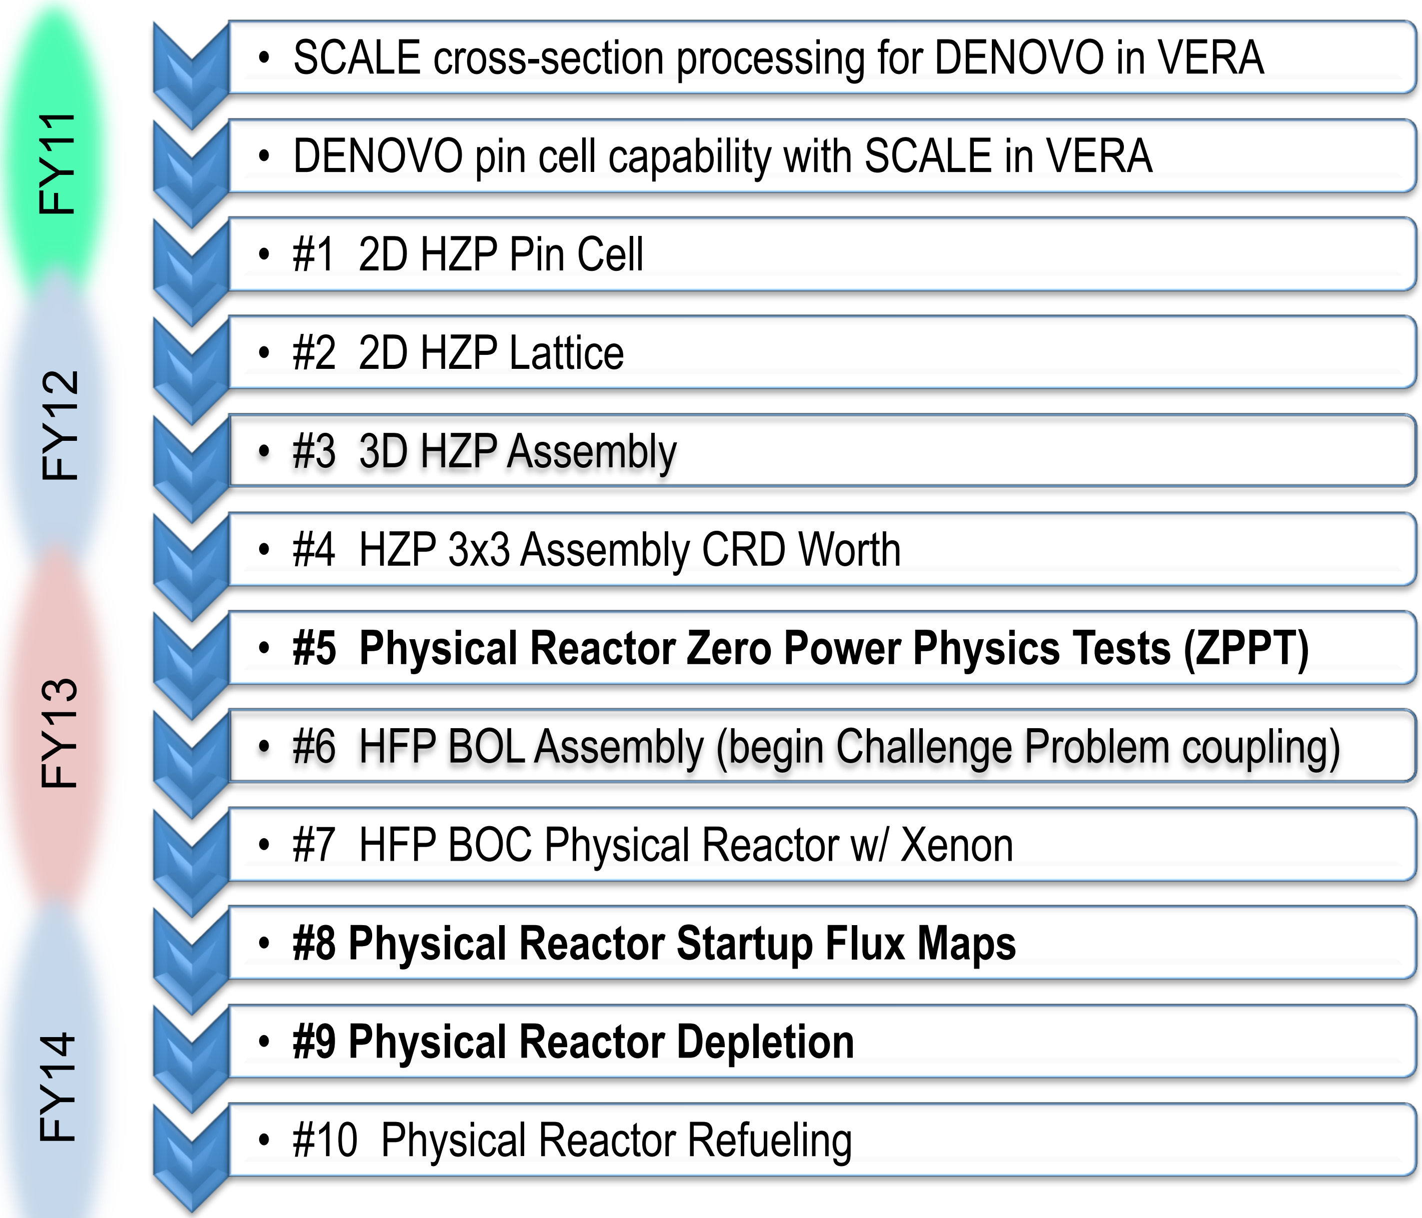

- SCALE cross-section processing for DENOVO in VERA
- DENOVO pin cell capability with SCALE in VERA
- #1 2D HZP Pin Cell
- #2 2D HZP Lattice
- #3 3D HZP Assembly
- #4 HZP 3x3 Assembly CRD Worth
- **#5 Physical Reactor Zero Power Physics Tests (ZPPT)**
- #6 HFP BOL Assembly (begin Challenge Problem coupling)
- #7 HFP BOC Physical Reactor w/ Xenon
- **#8 Physical Reactor Startup Flux Maps**
- **#9 Physical Reactor Depletion**
- #10 Physical Reactor Refueling

* **Bold text signifies ability to compare to measured plant data**

Problem 6 – PWR Single Assembly

17x17 Westinghouse Fuel Assembly

Watts Bar Unit 1 Cycle 1

Hot Full Power (HFP)

Beginning of Life (BOL) – no depletion

- Fuel Pins
- Plenum
- End Plugs
- Cladding
- Guide Tubes
- Spacer Grids
- Nozzles
- No Control Rods

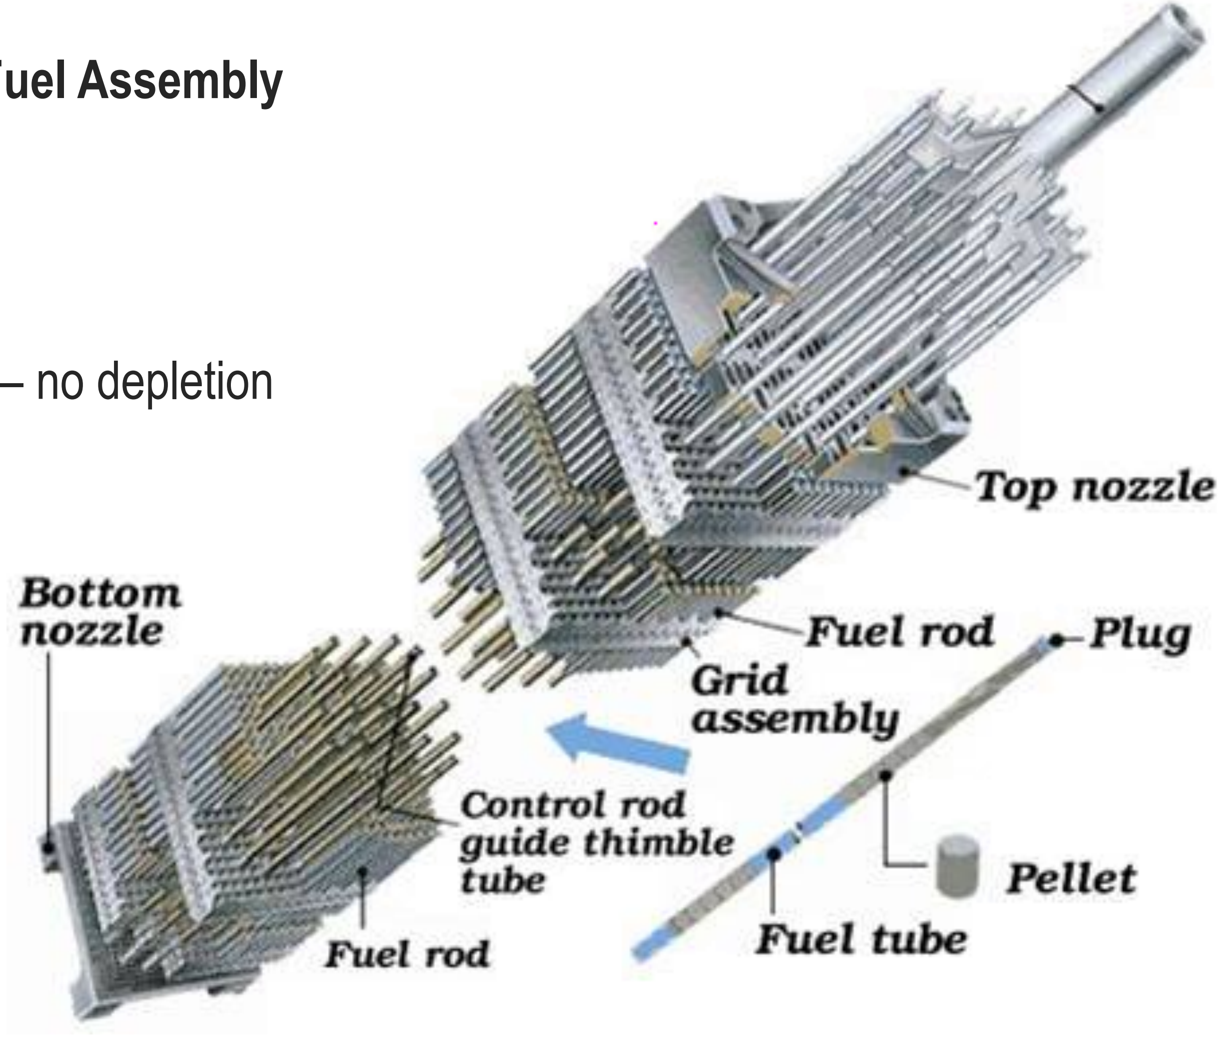

Prob. #6

Coupling of neutronics and thermal-hydraulics components for hot full-power beginning of life assy.

- AMA progression problem 6
  - neutronics (cross sections + neutron transport)
  - thermal-hydraulics (fluid flow and fuel/clad temperatures)

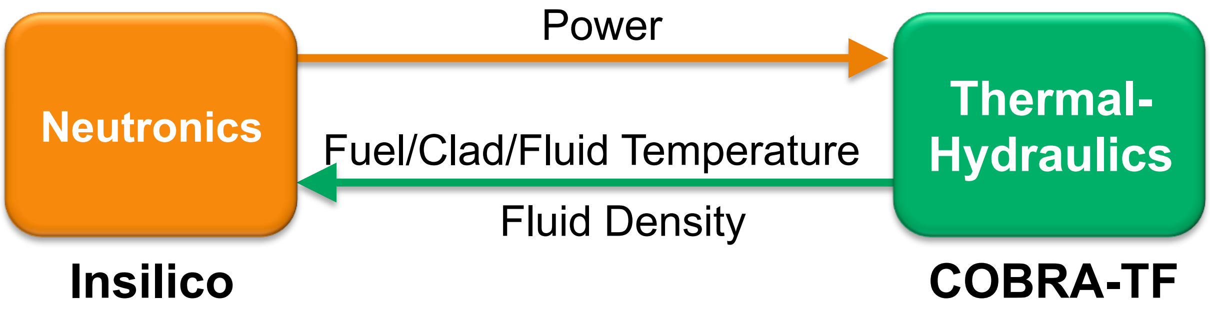

Power

Neutronics

Fuel/Clad/Fluid Temperature

Thermal-Hydraulics

Fluid Density

Insilico

COBRA-TF

- Coupling becomes more complicated with more codes, but we’ve done it.
- Challenges are related more to data transfer than “framework”.

Data Transfer Kit (DTK) will be discussed later.

Common Input – VERAin

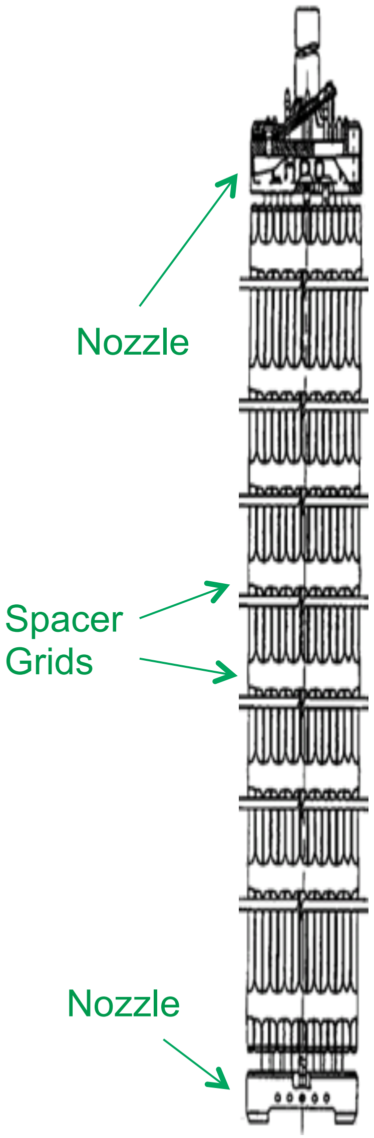

```
[ASSEMBLY]
  title "Westinghouse 17x17"
  npin 17
  ppitch 1.260

  fuel U31 10.257 95.0 / 3.1

  cell 1    0.4096 0.418 0.475 / U31 he zirc
  cell 10          0.561 0.602 / mod    zirc       ! guide tube
  cell 20          0.561 0.602 / mod    zirc       ! instrument tube
  cell 7           0.418 0.475 / mod    mod        ! empty location
  cell 8           0.418 0.475 /     he zirc       ! plenum
  cell 9                 0.475 /        zirc       ! pincap

  lattice FUEL1
     20
      1 1
      1 1 1
     10 1 1 10
      1 1 1  1 1
      1 1 1  1 1 10
     10 1 1 10 1  1 1
      1 1 1  1 1  1 1 1
      1 1 1  1 1  1 1 1 1

 lattice PLEN1
     20
      8 8
      8 8 8
     10 8 8 10
      8 8 8  8 8
      8 8 8  8 8 10
     10 8 8 10 8  8 8
      8 8 8  8 8  8 8 8
      8 8 8  8 8  8 8 8 8
```

```
axial A1    6.050
    LGAP1  10.281
    PCAP1  11.951
    FUEL1 377.711
    PLEN1 393.711
    PCAP1 395.381
    LGAP1 397.501

grid END inc  1017 3.866  ! grid mass, height (cm)
grid MID zirc 875  3.810  ! grid mass, height (cm)

grid_axial
    END  13.884
    MID  75.2
    MID 127.4
    MID 179.6
    MID 231.8
    MID 284.0
    MID 336.2
    END 388.2

lower_nozzle  ss 6.05  6250.0  ! mat, height, mass
upper_nozzle  ss 8.827 6250.0  ! mat, height, mass
```

Information shown originates from FSAR document
(ref 4 from L1:CASL.P7.01)

Neutronics – Insilico

- Part of the Exnihilo environment
- Transport solver is S_N (pin-resolved) or SP_N (pin-homogenized)
- Built in cross section processing with XSProc
- S_N uses the KBA implementation which solves the problem on a Cartesian grid and can scale efficiently to over 100,000 processors

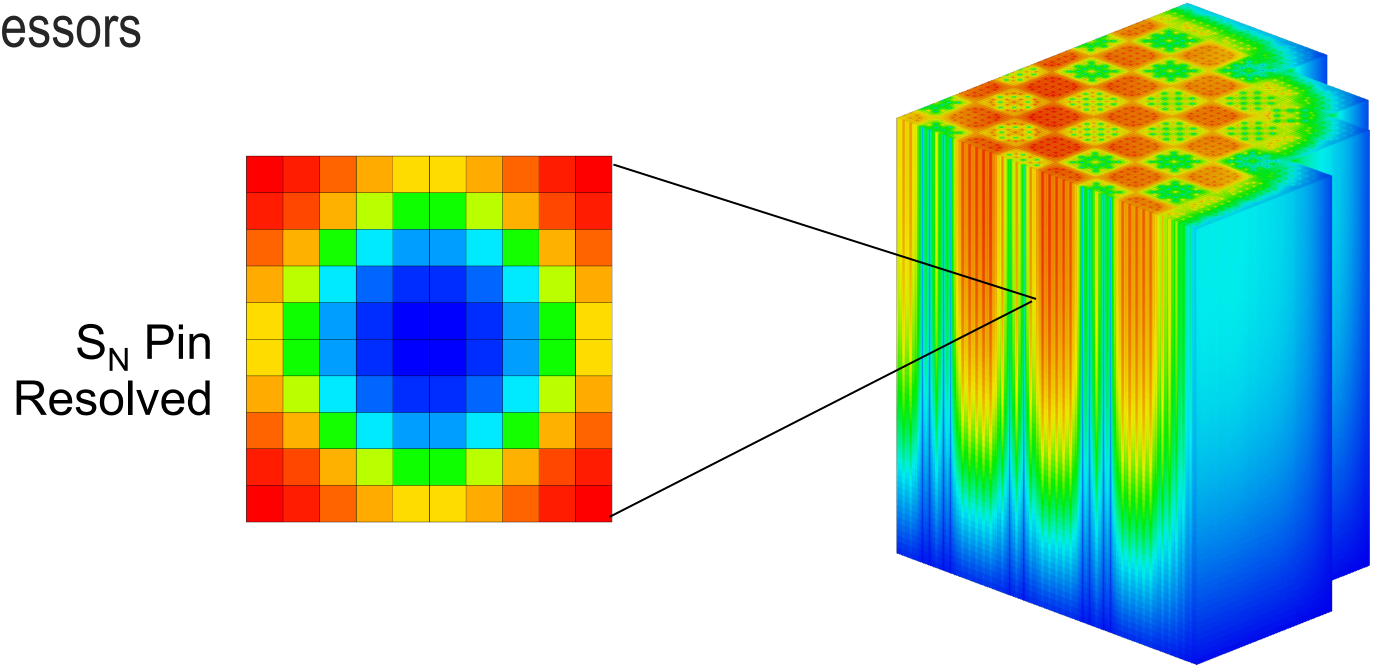

Efficient Scaling to Large 3D Problems

Thermal Hydraulics – COBRA-TF

- COBRA-TF (CTF) subchannel code from Penn. State Univ.
- Two-fluid, three-field representation of the two-phase flow
  - Continuous vapor (mass, momentum and energy)
  - Continuous liquid (mass, momentum and energy)
  - Entrained liquid drops (mass and momentum)
  - Non-condensable gas mixture (mass)
- Spacer grid models
- Pin conduction model
- Built-in material properties
- Since bringing in for CASL / VERA:
  - dramatically reduced memory usage
  - dramatically increased performance
  - dramatically expanded test coverage
  - implementing parallel version for further reduction in run-times

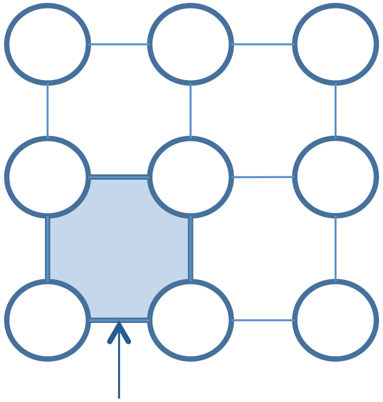

Subchannel area x 49 axial levels

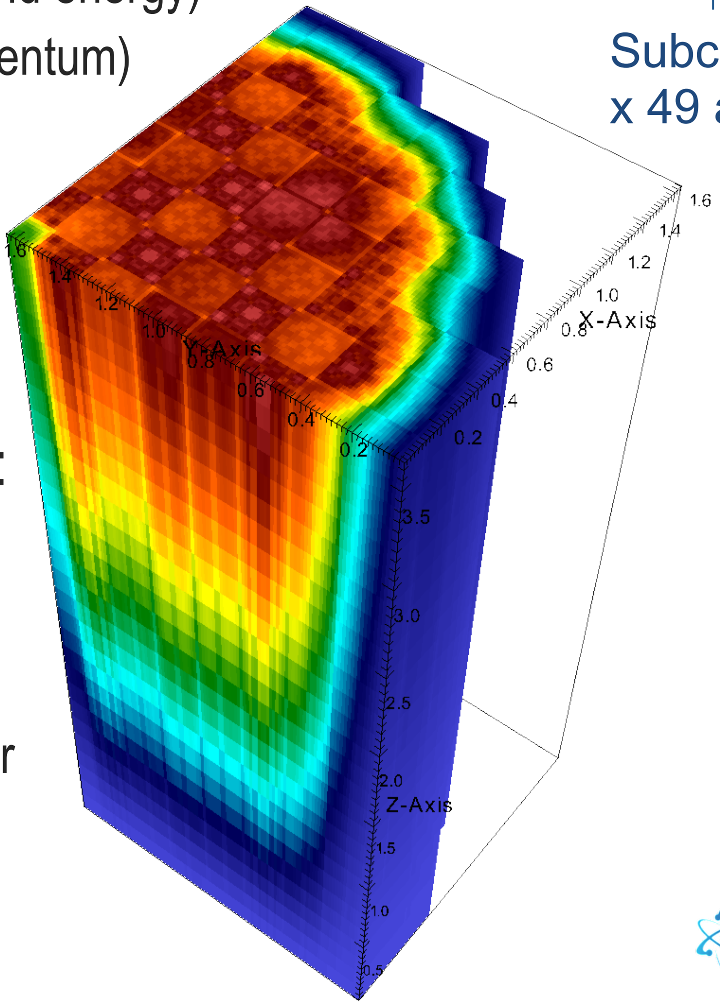

Thermal-Hydraulics (T-H)

- Subchannel run-times
  - Few minutes for single assembly
  - Scalable to full core (>62,000 subchannels) using one assembly per compute core
- CFD can provide higher fidelity for smaller assembly-sized problems
  - HYDRA and STAR-CCM
  - Currently impractical for full-core depletion calculations

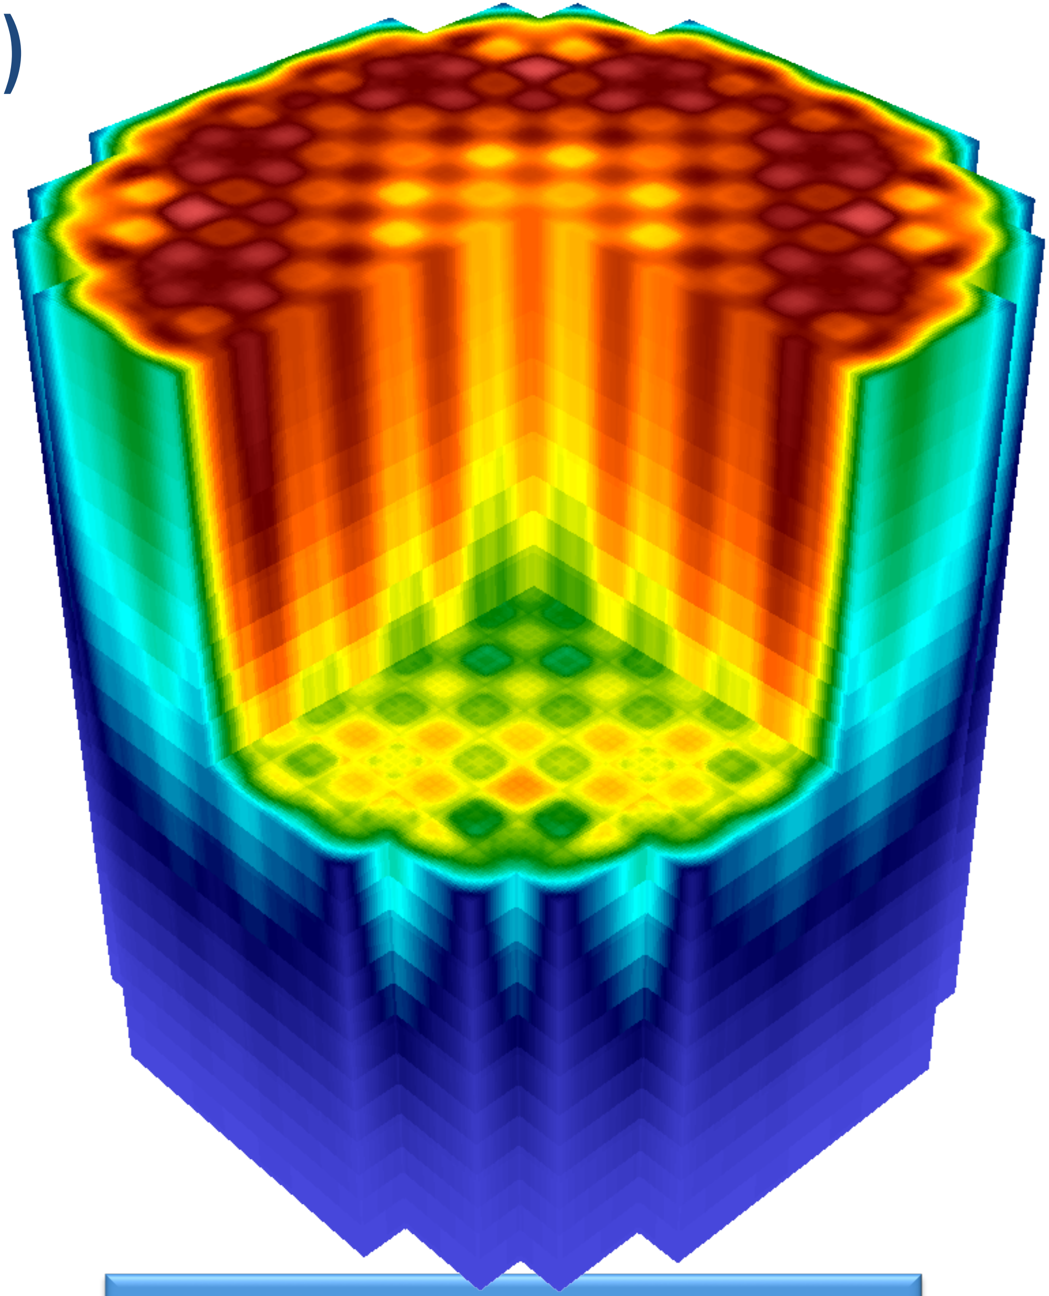

CTF results for coolant enthalpy in qtr-core Watts Bar model

Coupled Results – Boron Concentration

Prob. #6

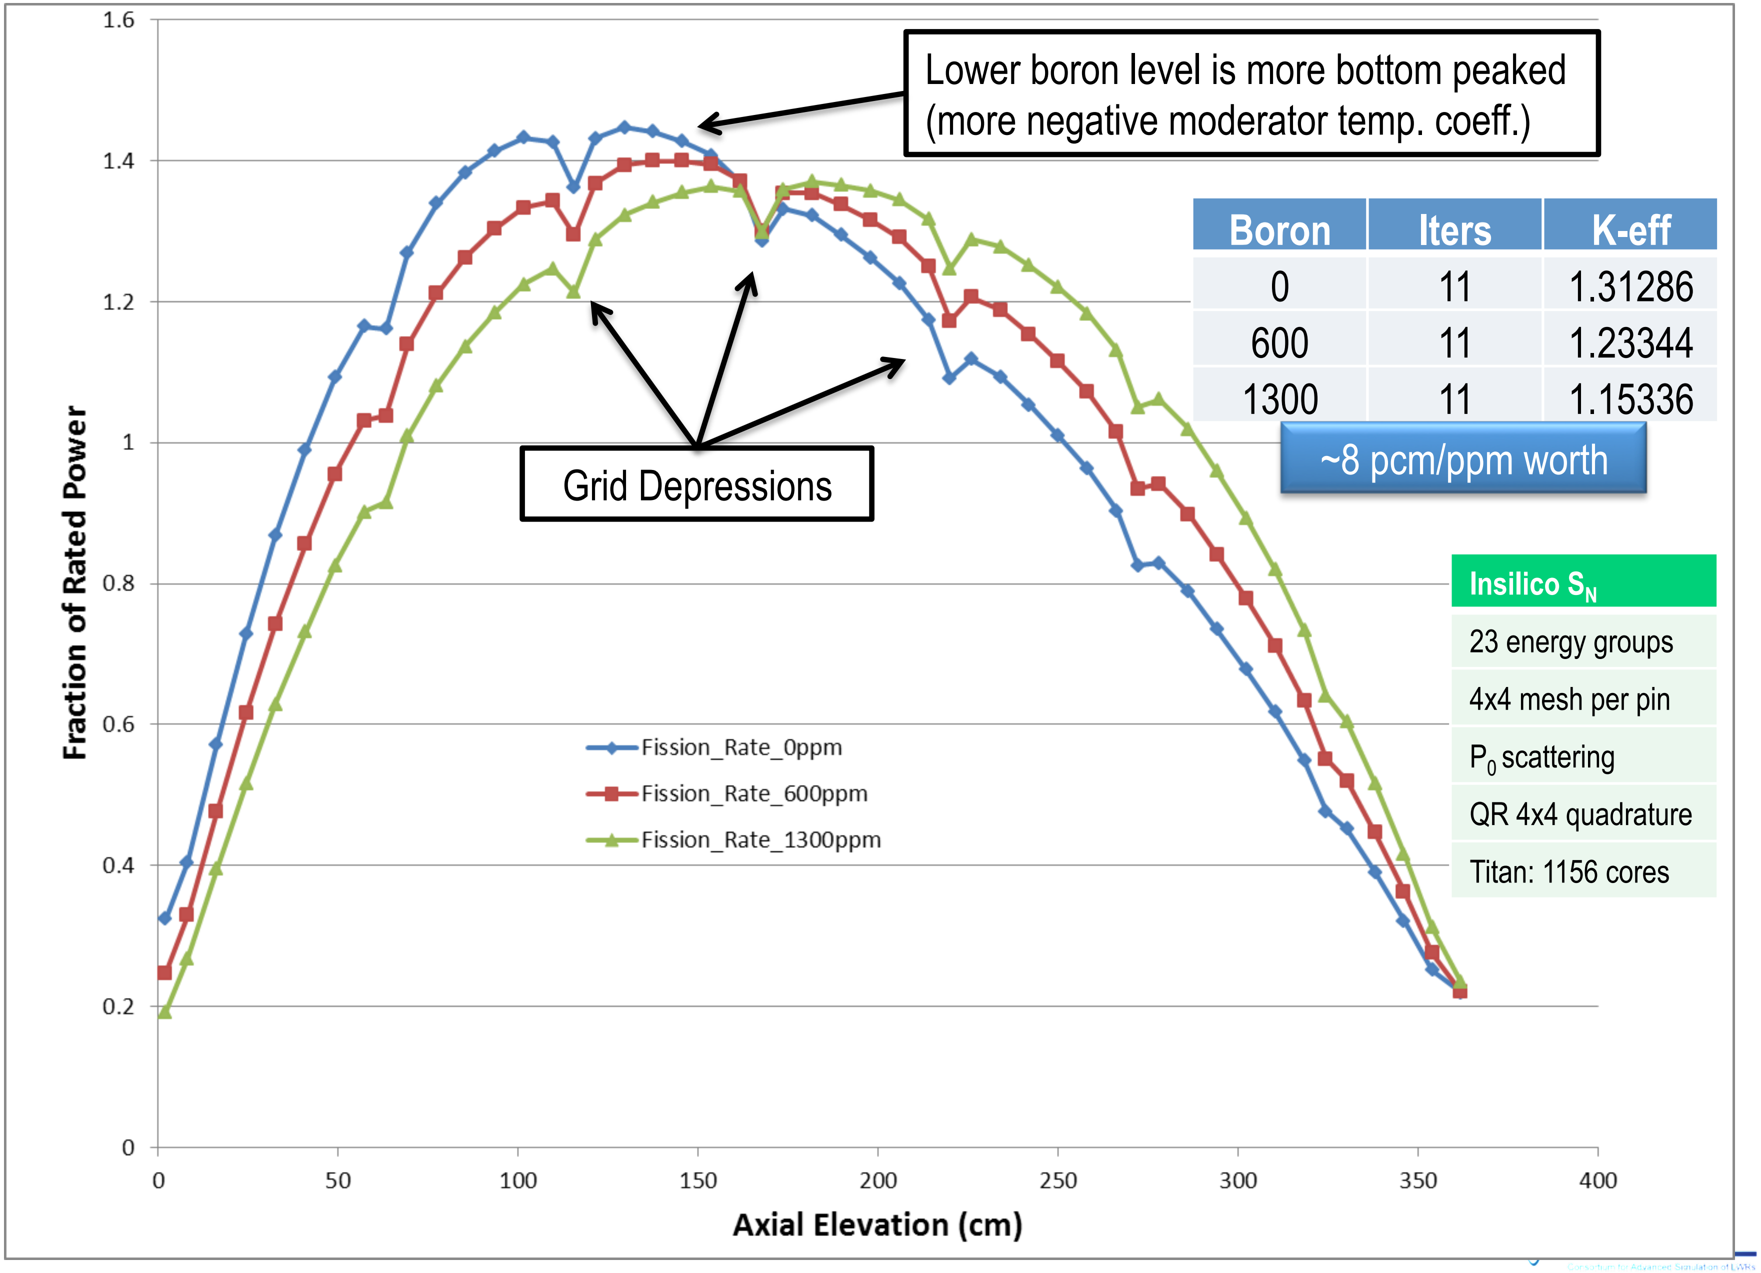

Lower boron level is more bottom peaked (more negative moderator temp. coeff.)

Grid Depressions

Boron	Iters	K-eff
0	11	1.31286
600	11	1.23344
1300	11	1.15336

~8 pcm/ppm worth

Insilico S_N

- 23 energy groups
- 4x4 mesh per pin
- P_0 scattering
- QR 4x4 quadrature
- Titan: 1156 cores

Prob. #6

VERA Simulation of Hot Full Power Assembly
(neutronics with fluid / moderator temp. feedback)

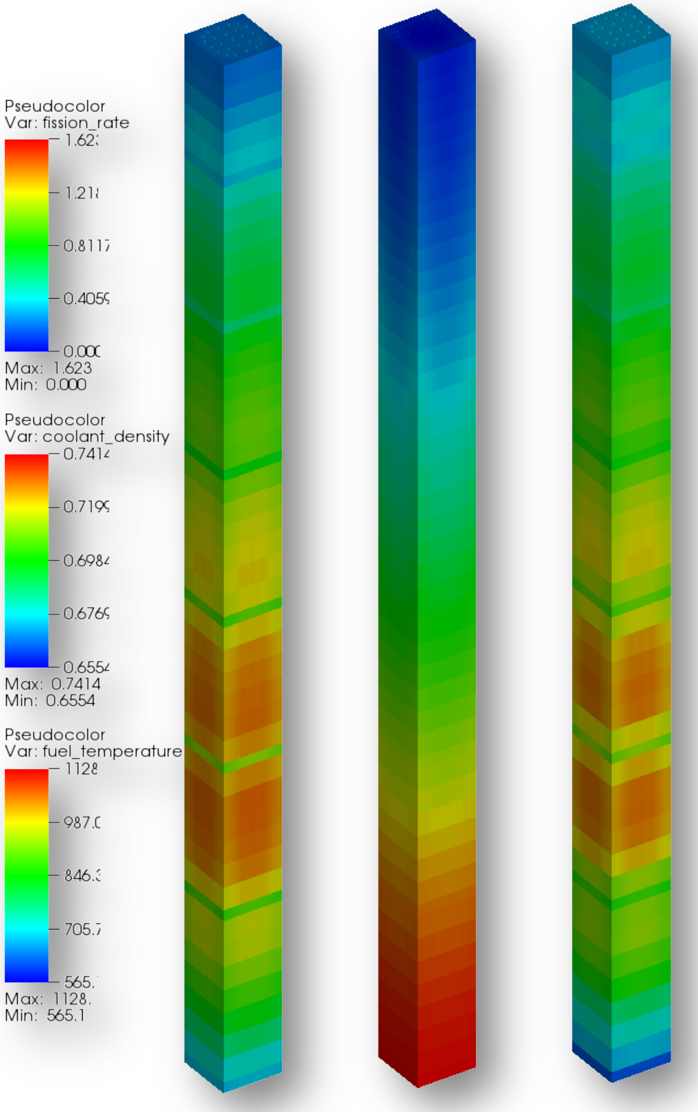

0 ppmB

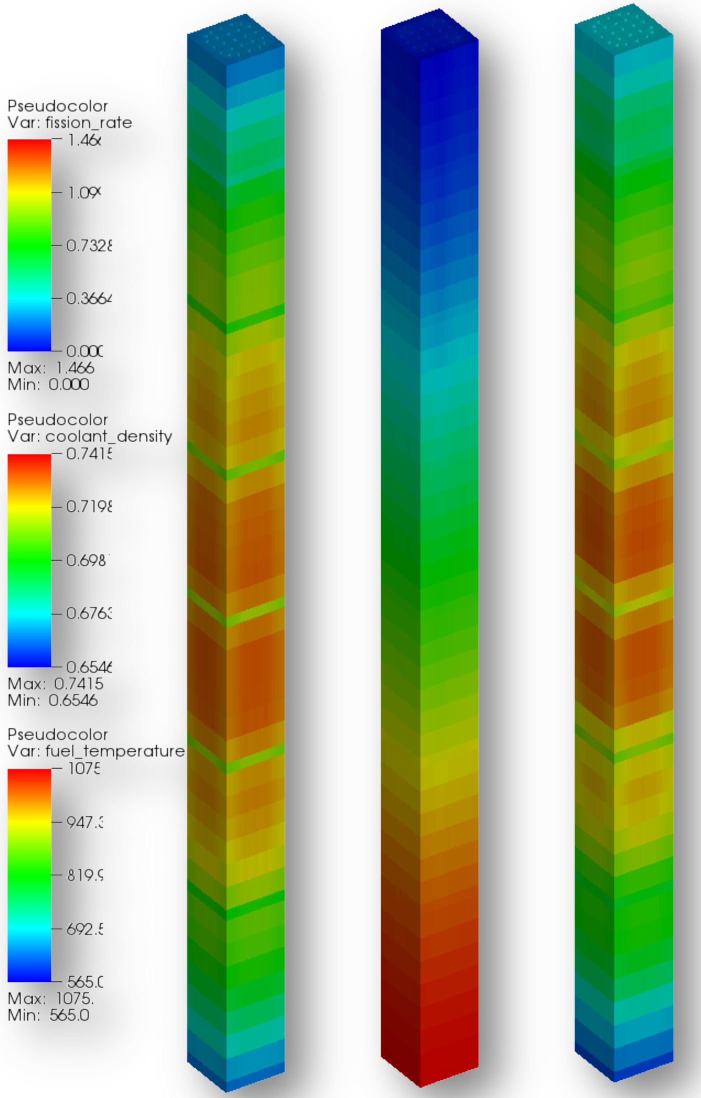

1300 ppmB

Coupled Results – Power Level

Prob. #6

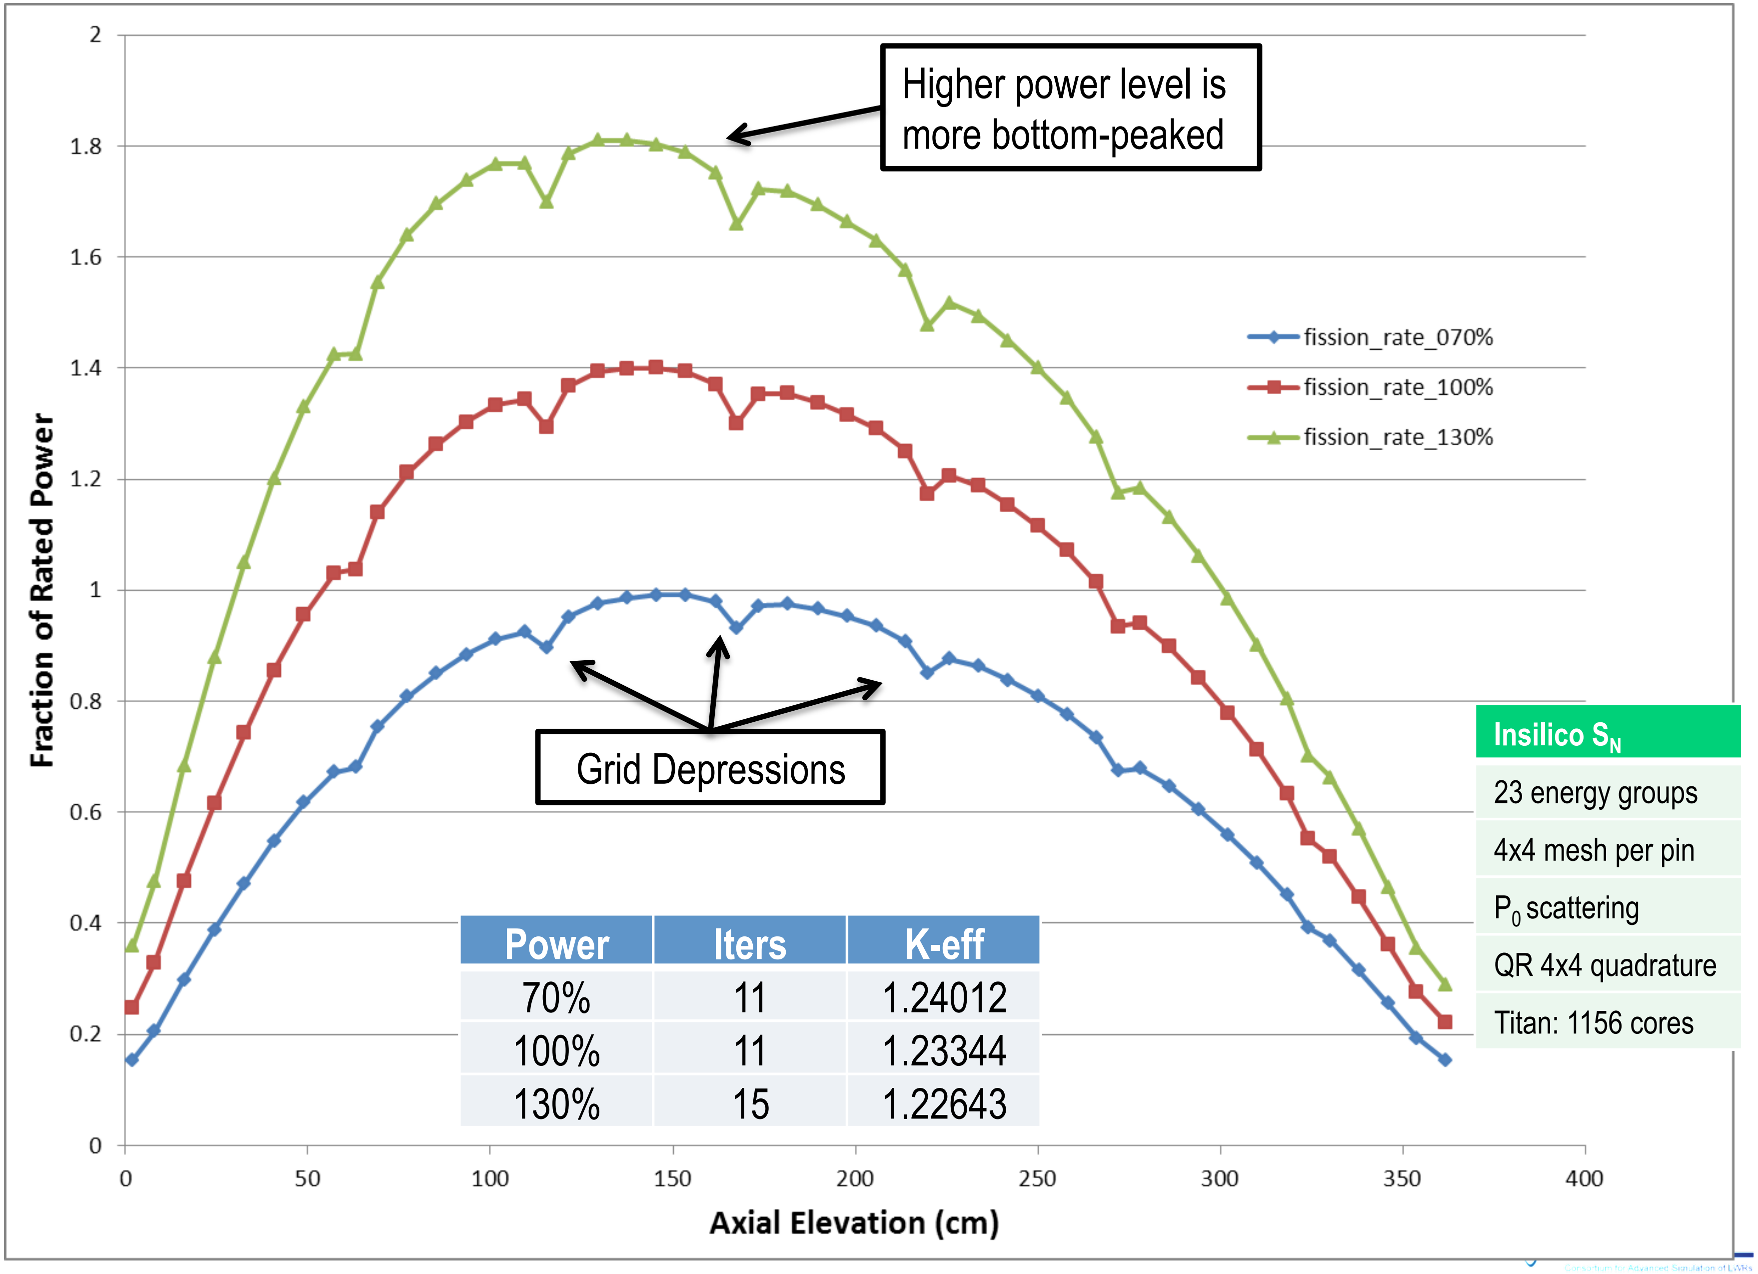

Higher power level is more bottom-peaked

Grid Depressions

Power	Iters	K-eff
70%	11	1.24012
100%	11	1.23344
130%	15	1.22643

Insilico S$_N$

23 energy groups

4x4 mesh per pin

P$_0$ scattering

QR 4x4 quadrature

Titan: 1156 cores

Prob. #6

VERA Multiphysics Simulation of PWR Fuel Assembly
- Base for adding Challenge Problems Physics

- Coupled multiphysics model of WEC PWR fuel assembly
  - Neutron transport to calculate power distribution (Insilico)
  - Thermal-Hydraulics in coolant (COBRA-TF)
  - Heat conduction in fuel rods (COBRA-TF)
  - Neutron cross sections as function of temperature and density (XSProc)
- Next step is scaling up to a 1/4 reactor core simulation in support of a DOE reportable milestone.

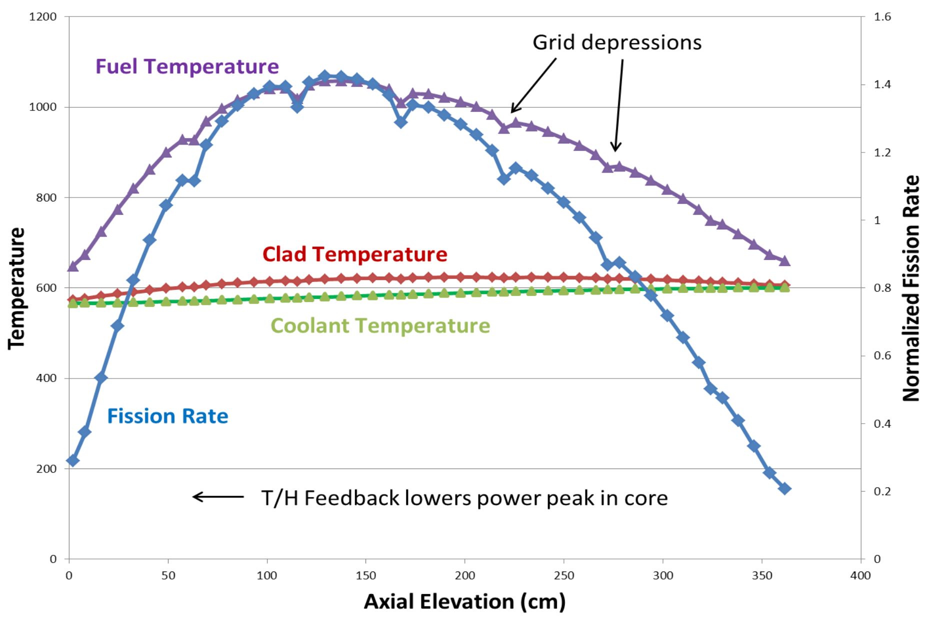

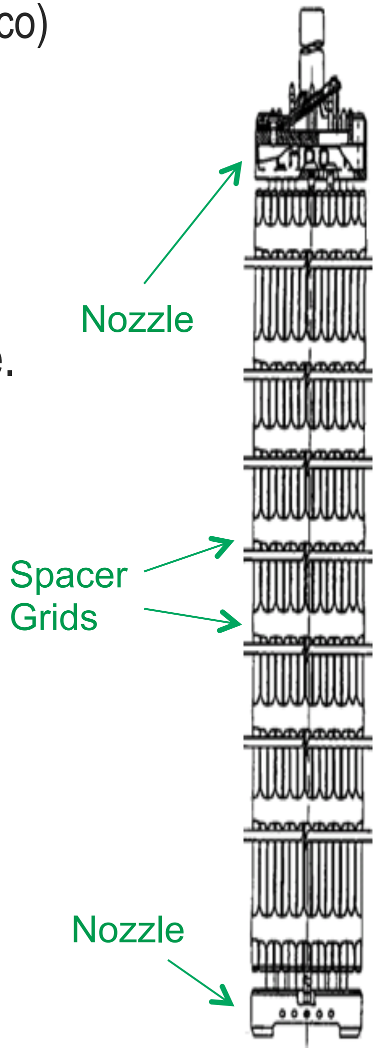

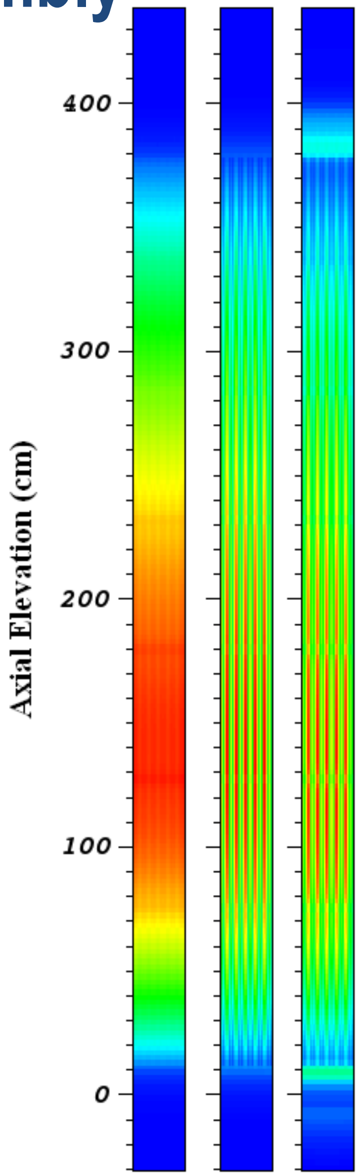

Fast, Epithermal, and Thermal Flux Profiles

CTF / Insilico / Peregrine Coupled Driver

Roger Pawlowski (SNL)

- Named **Tiamat**, multi-headed dragon (Babylonian mythos)
  - easy to add new “heads” provided the apps meet integration criteria
- All applications are run in their own MPI process space
  - can overlap if desired
  - reduce collisions and improved algorithms performance
- Data Transfers are handled through DTK with MPI sub-communicators

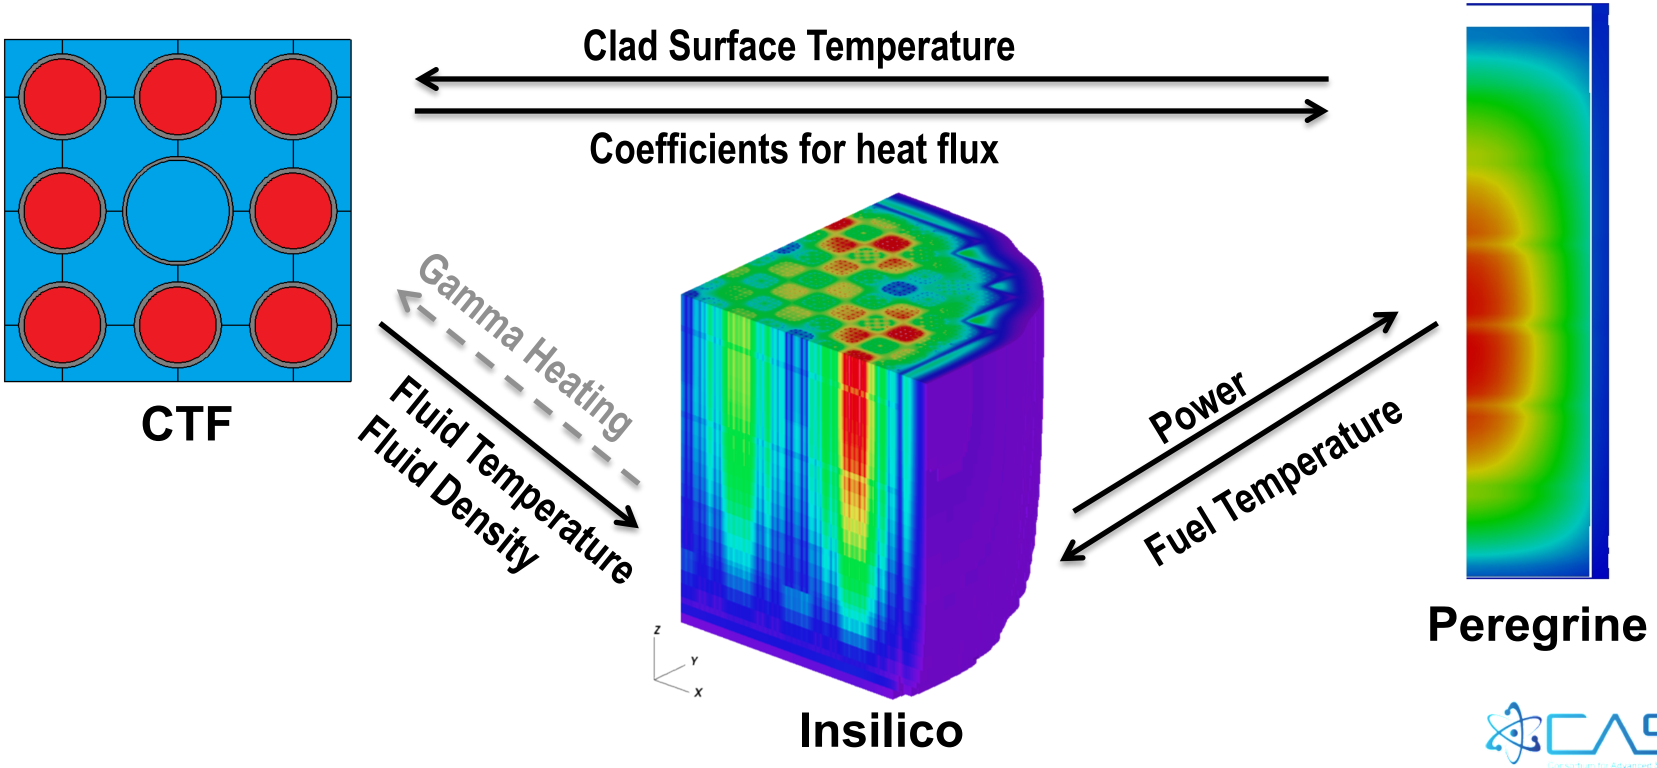

Clad Surface Temperature

Coefficients for heat flux

Gamma Heating

Fluid Temperature

Fluid Density

Power

Fuel Temperature

CTF

Insilico

Peregrine

Roger Pawlowski (SNL)

VERA: Virtual Environment for Reactor Applications

(components currently comprising Tiamat highlighted)

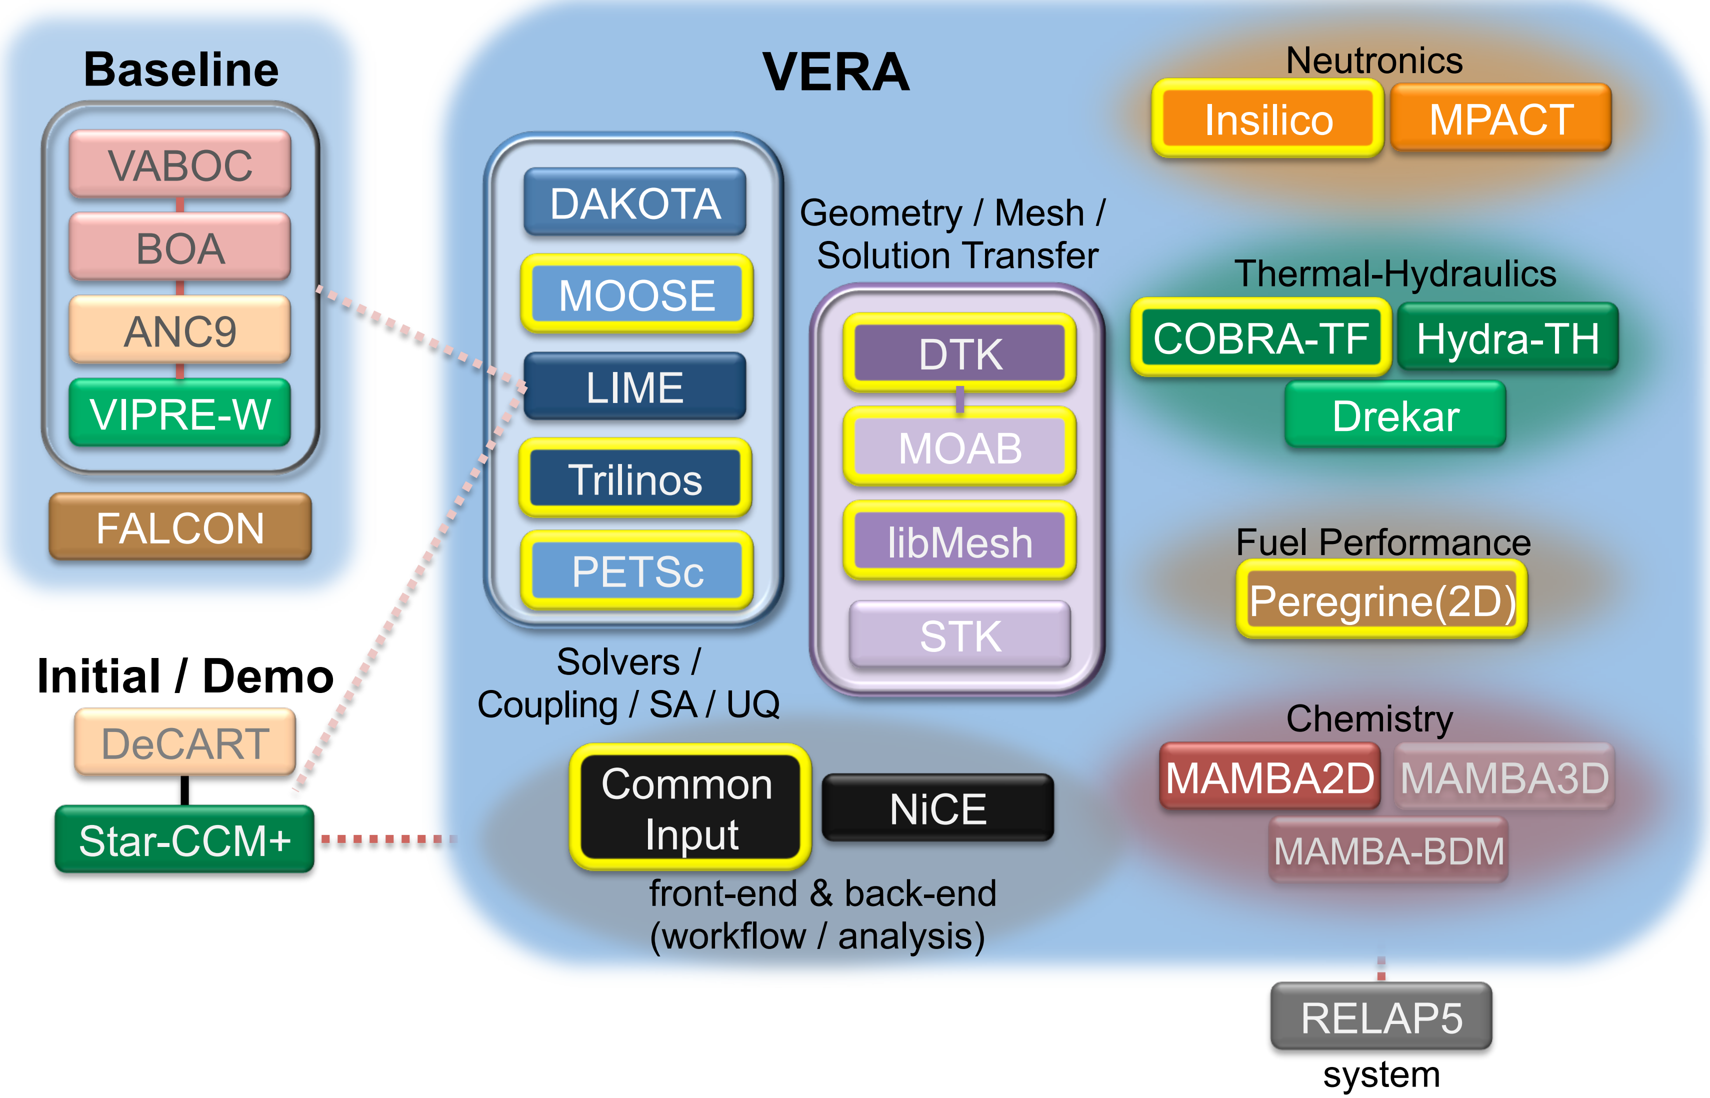

Baseline

VABOC

BOA

ANC9

VIPRE-W

FALCON

Initial / Demo

DeCART

Star-CCM+

VERA

DAKOTA

MOOSE

LIME

Trilinos

PETSc

Solvers / Coupling / SA / UQ

Geometry / Mesh / Solution Transfer

DTK

MOAB

libMesh

STK

Common Input

NiCE

front-end & back-end (workflow / analysis)

Neutronics

Insilico

MPACT

Thermal-Hydraulics

COBRA-TF

Hydra-TH

Drekar

Fuel Performance

Peregrine(2D)

Chemistry

MAMBA2D

MAMBA3D

MAMBA-BDM

RELAP5

system

MOOSE/Peregrine Integration Comments

Roger Pawlowski (SNL)

- We connect to MOOSE (INL) via the DTK MultiApp interfaces
  - MOOSE controls sub-commmunicator layout for individual pins
- CASL_MOOSE: Peregrine built under the VERA build system (TriBITS)
  - Can run any stand alone Peregrine input file. Verifies correct integration.
- "TiamatApp": A new MOOSE Application
  - To properly unit test we have implemented a new MOOSE App that pulls in Peregrine kernels and then adds extra unit testing MOOSE kernels
    - Runs all Peregrine input files, but extends for coupling
  - New kernels:
    - Multiphase surface flux from CTF
    - Heat source from Insilico
    - All incoming quantities are stored as MOOSE aux variables

Modifying a code for integration can be challenging, but not necessarily in obvious ways.

- Exposing input parameters and output responses is one of the most challenging aspects
  - Parallel distribution, data structures, units and coordinates, …
- Codes are no longer top of the software food chain (main()):
  - No global variables, using namespace declarations in headers
  - “Solve” can be called multiple times
  - Must be able to reset if any physics fails a “step”
  - Can not control/manipulate the parse of input
  - Can not redirect output streams, must allow ostreams to be set
  - Can not assume MPI_COMM_WORLD anywhere in your code (must accept an MPI communicator)
  - In-source builds are dangerous
- Memory management strategies are critical (RCPs)
- Robust error handling

Data Transfer Kit (DTK) (Slattery, Wilson, Pawlowski)

Stu Slattery (ORNL)
Roger Pawlowski (SNL)

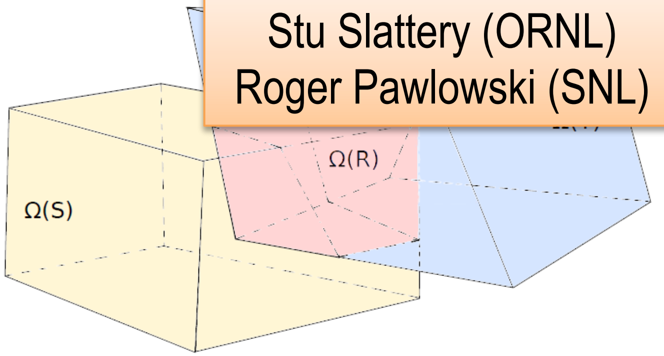

$$\mathbf{M} : \mathbb{R}^D \rightarrow \mathbb{R}^D, \forall r \in \Omega_R$$

$$\mathbf{G}(t) \leftarrow \mathbf{M}(\mathbf{F}(s))$$

- Collection of geometry-based data mapping algorithms for shared domain problems
  - Rendezvous Algorithm, Initially developed by the Sandia SIERRA team in mid-2000's for parallel mesh-based data transfer
- Data maps allow for efficient movement of data in parallel
  - e.g. between meshes of a different parallel decomposition
- Ideally maps are generated in desirable time complexity (logarithmic)
- Does not provide general interface for all physics codes to couple to all other physics codes
- Does not provide discretization services (e.g. basis functions)
- Open-source BSD 3-clause license - https://github.com/CNERG/DataTransferKit

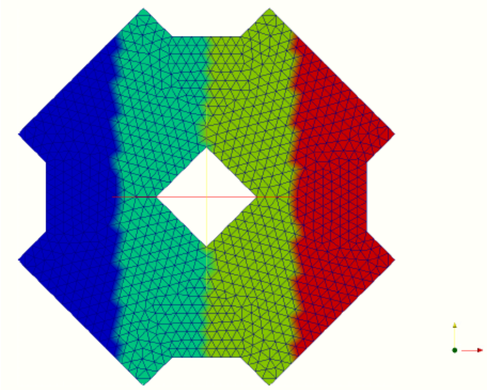

W_source

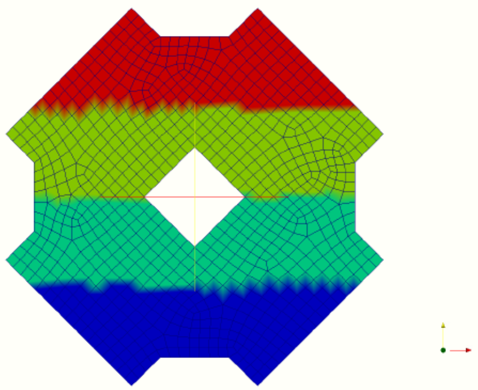

W_target

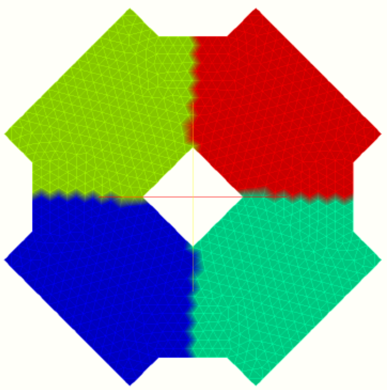

W_rendezvous

Stu Slattery (ORNL)
Roger Pawlowski (SNL)

DTK Implements Mappings for Required Transfers (Rendezvous used by all Mappings)

Shared Domain Map

Mesh →Point

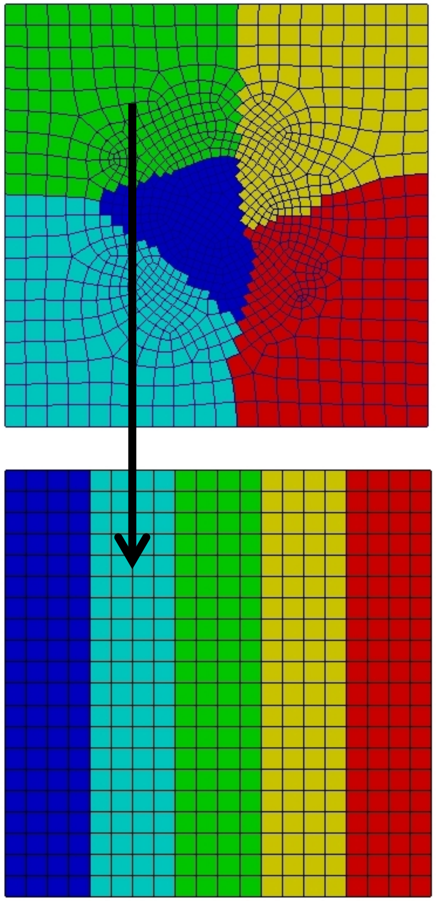

Integral Assembly Map

Mesh →Geometry

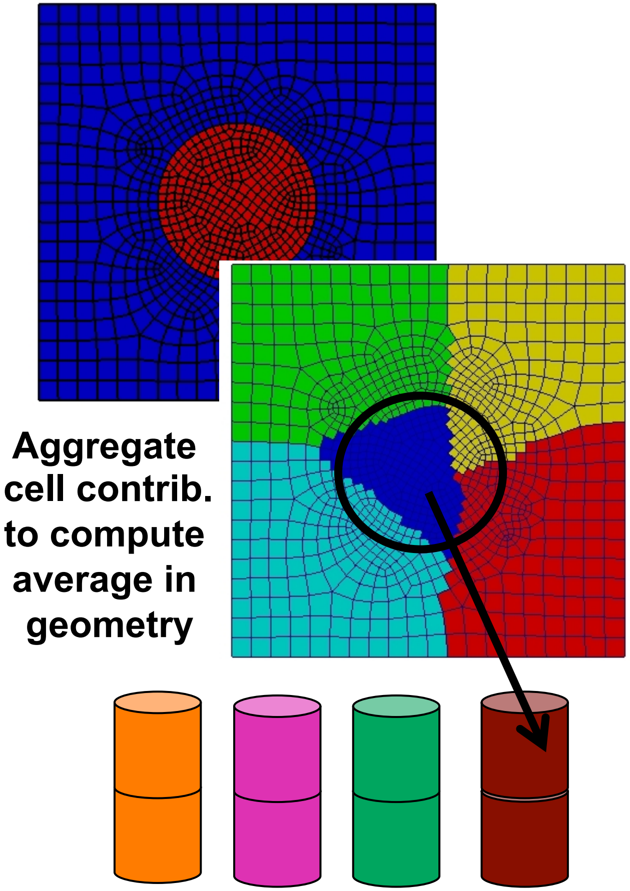

Aggregate cell contrib. to compute average in geometry

Shared Volume Map

Geometry → Point

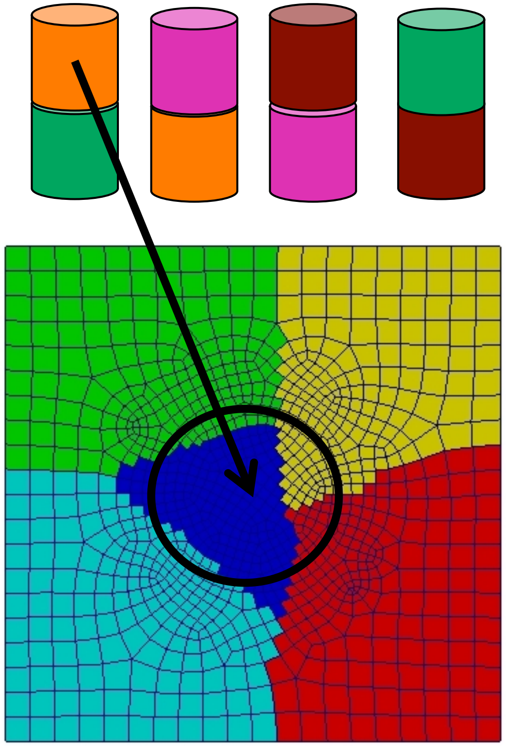

Colors represent different MPI processes

DTK has been released as open-source, and has been integrated into MOOSE (INL) for solution transfer.

Stu Slattery (ORNL)
Roger Pawlowski (SNL)

Data Transfer Kit (DTK)

Weak Scaling Study (16 to 16K cores)

- **Worst case scenario**
  - all-to-all with 10K random points per core
  - Applications will have significantly better data locality
- Scaling study on Titan (Cray XK7)
- Largest test problems so far over 10^9 elements and 10^5 cores

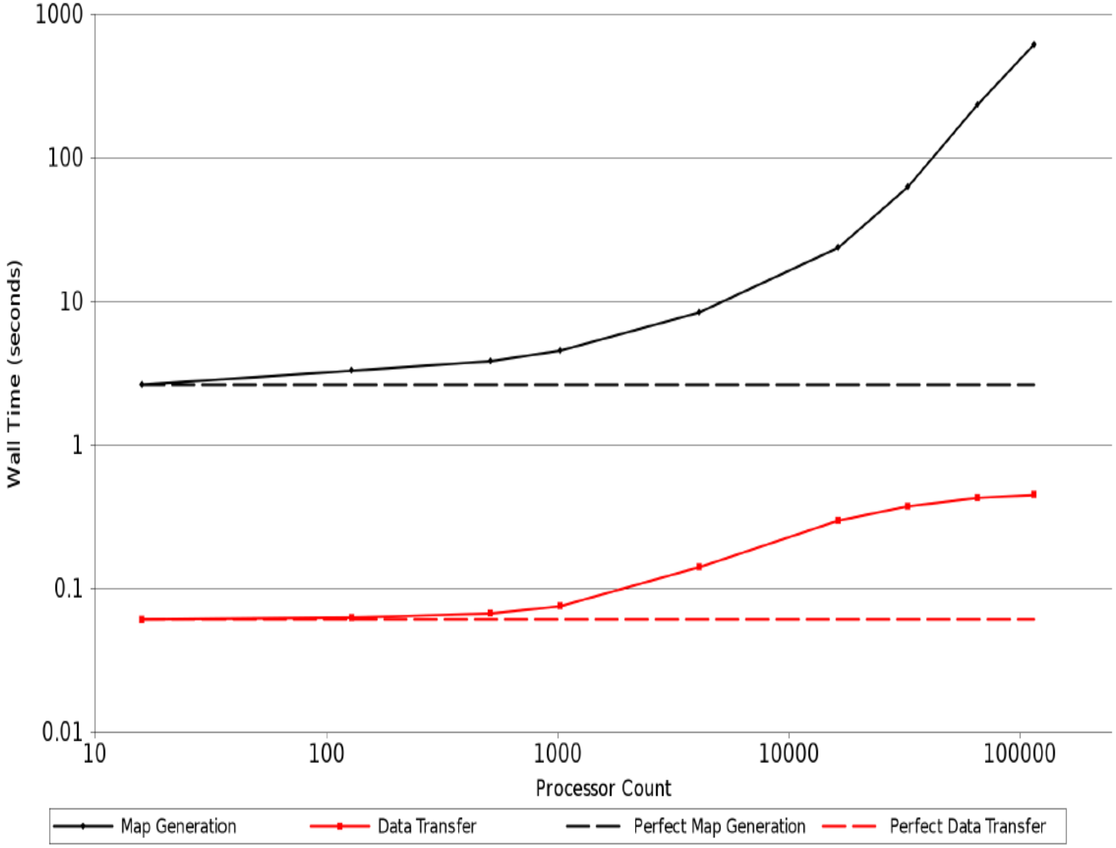

Excellent performance to 116K cores!

A Multiphysics Distributor
(Four levels of MPI Communicators)

Roger Pawlowski (SNL)

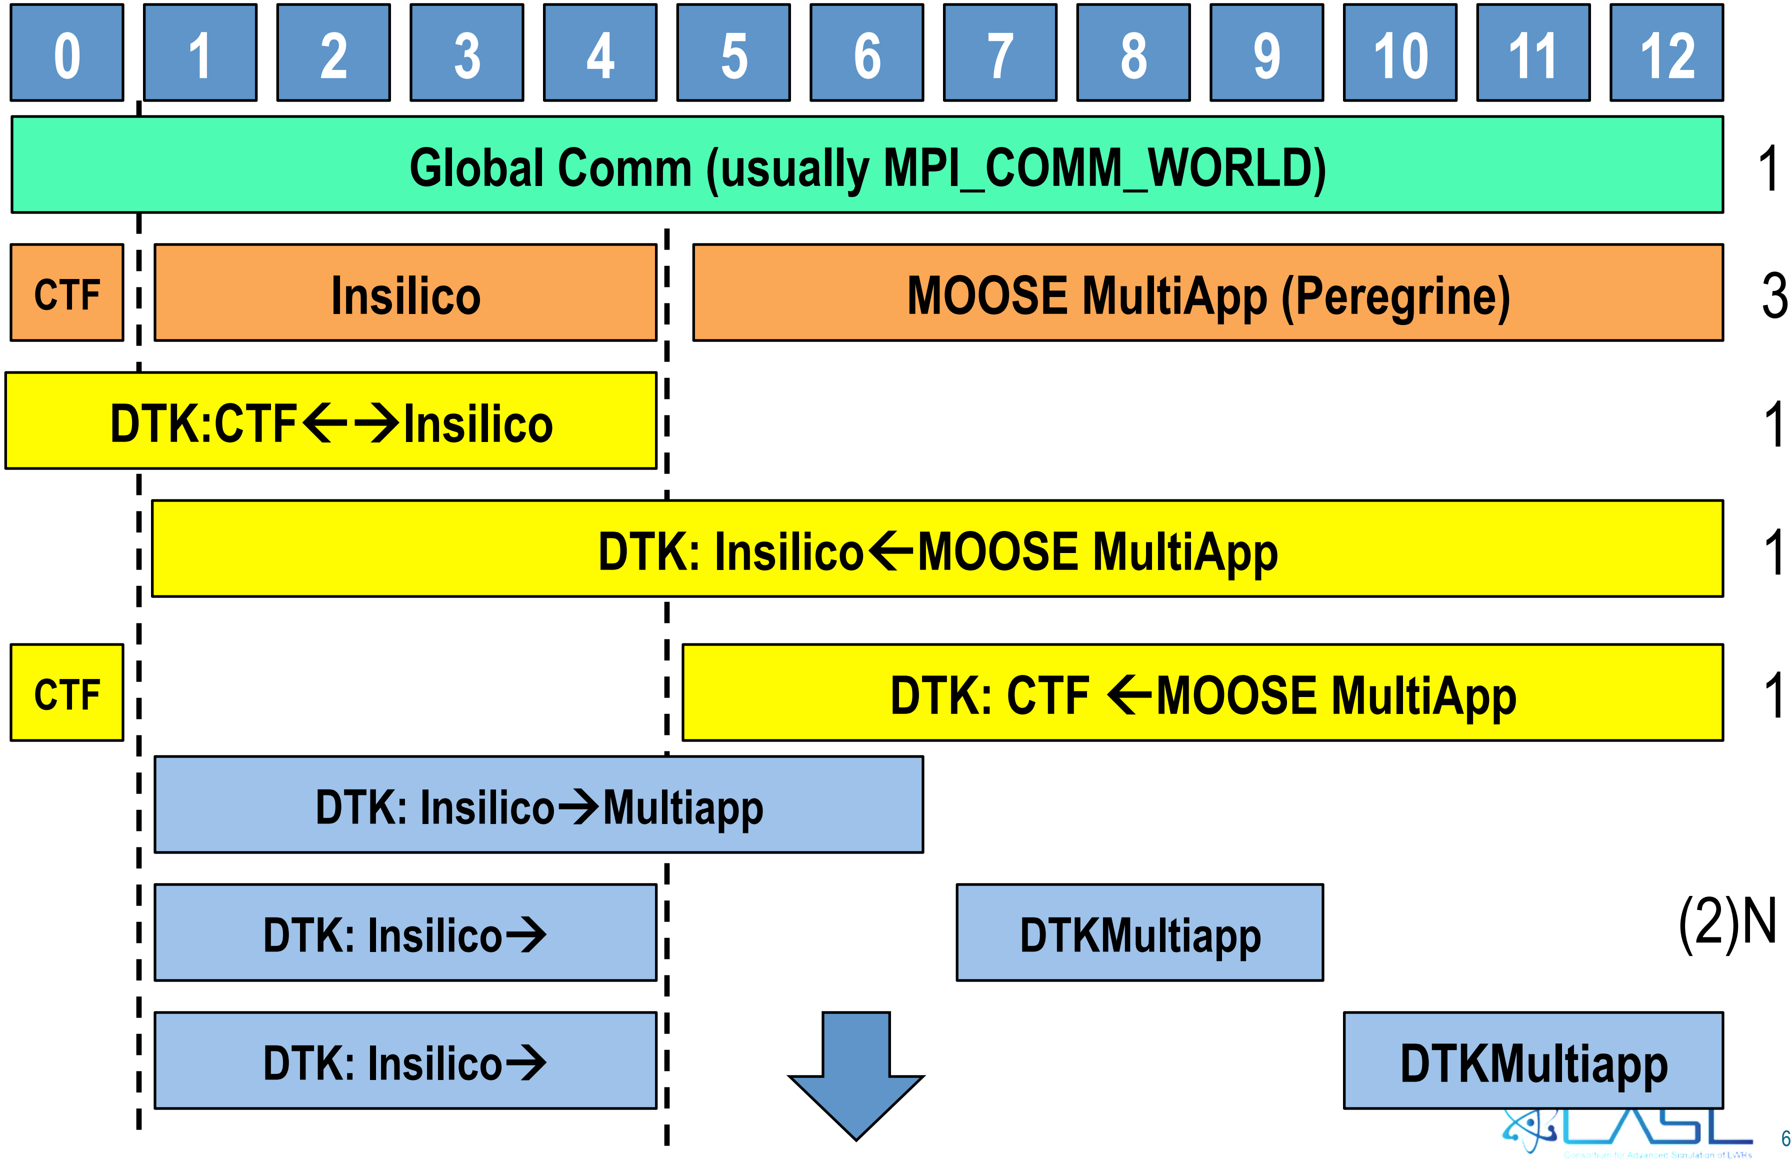

0 1 2 3 4 5 6 7 8 9 10 11 12

Global Comm (usually MPI_COMM_WORLD) 1

CTF | Insilico | MOOSE MultiApp (Peregrine) 3

DTK:CTF←→Insilico 1

DTK: Insilico←MOOSE MultiApp 1

CTF | DTK: CTF ←MOOSE MultiApp 1

DTK: Insilico→Multiapp

DTK: Insilico→ | DTKMultiapp (2)N

DTK: Insilico→ | DTKMultiapp

Coupled results for 17x17 WEC Assembly (AMA progression problem 6)

Roger Pawlowski (SNL)

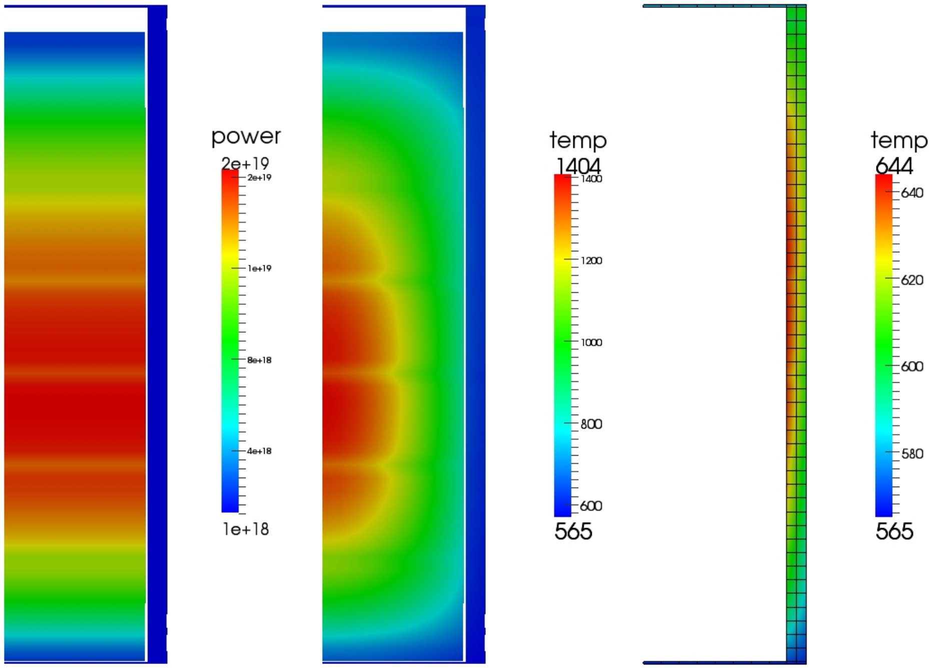

Fission rate (from Insilico) and temperature in Peregrine for a selected rod. The plot on the right is scaled to show clad temps.

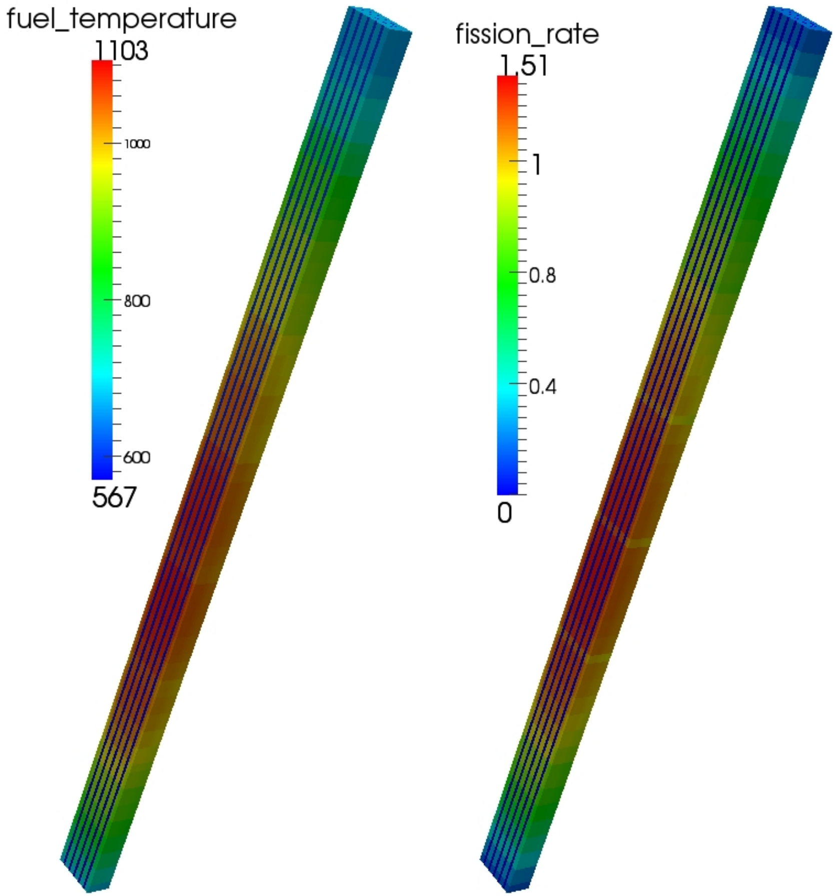

Insilico averaged fuel temp. and fission rate

Coupled results for 17x17 WEC Assembly
(AMA progression problem 6)

Roger Pawlowski (SNL)

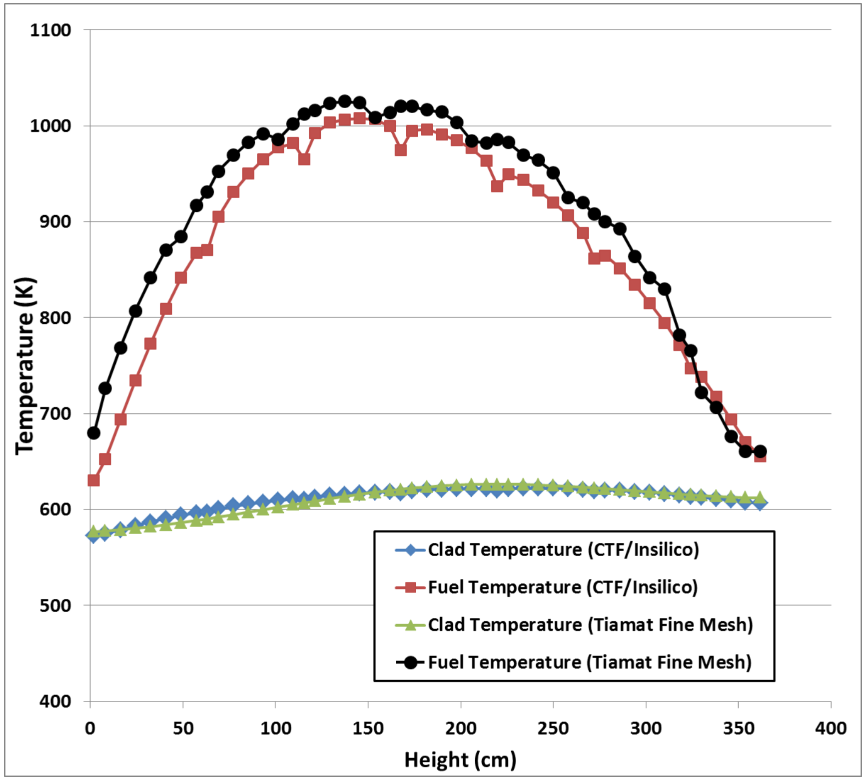

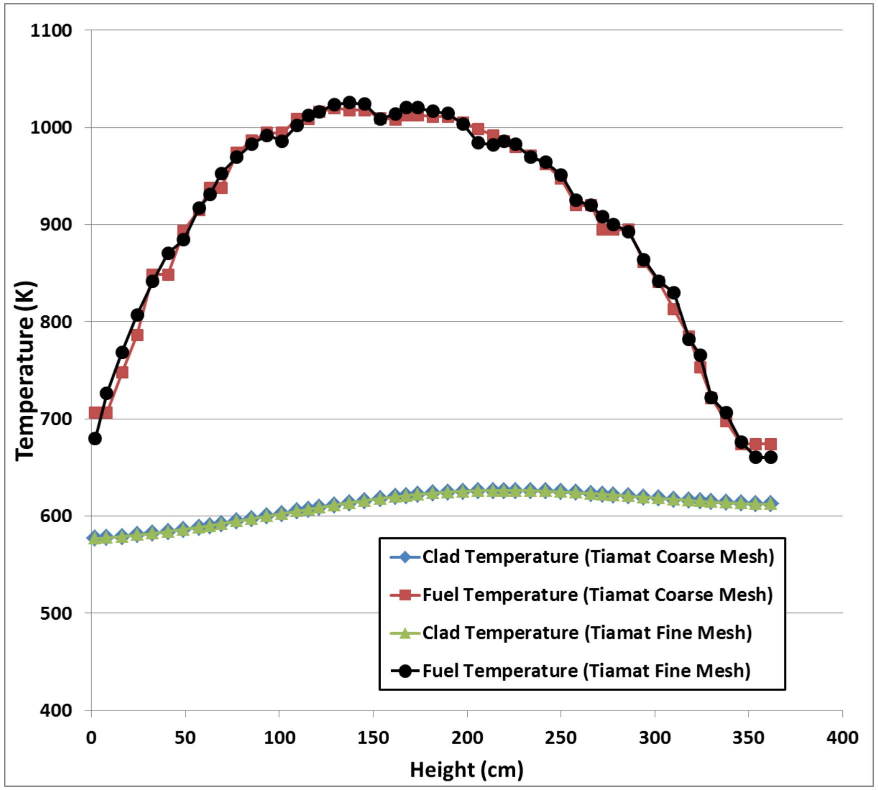

Tiamat vs. CTF+Insilico coupled capability

- need to investigate consistently higher Tiamat / Peregrine temperatures and apparent offset
  - substantial differences between CTF and Peregrine rod models
  - boundary condition treatments

Coarse mesh vs. “fine” mesh

- more smearing with coarse mesh (as expected)
- further refinement studies needed

Tiamat represents a significant new VERA capability.

Roger Pawlowski (SNL)

- Successfully integrated MOOSE/Peregrine into VERA
- Successfully coupled MOOSE/Peregrine with COBRA-TF and Insilico
- Developed new multiphysics driver ("LIME2")
  - evolution of LIME first formally described in PoR-3 milestone L3:VRI.PSS.P3.01, "LIME 2.0 Design Report"
- Developed new data transfer mechanisms for specific applications

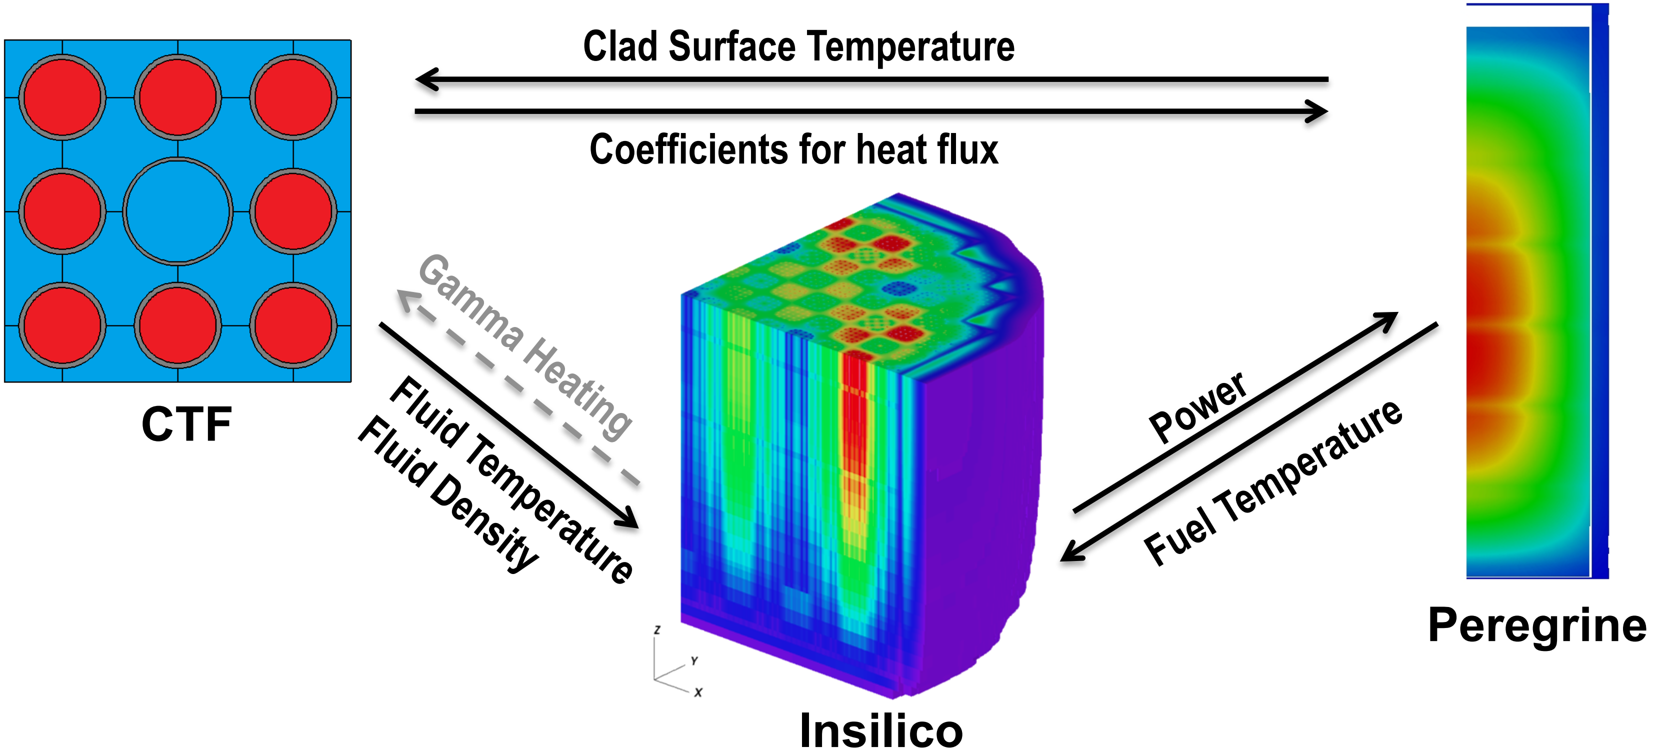

Demonstration of a New High Fidelity Multi-Physics Simulation Model of PWR Reactor

Purpose

- First large-scale coupled multi-physics model of operating PWR reactor using Components of CASL's Virtual Environment for Reactor Applications (VERA)
- Features resolved are based on the dimensions and state conditions of Watts Bar Unit 1 Cycle 1: geometry for fuel, burnable absorbers, spacer grids, nozzles, and core baffle

Execution

- Common input used to drive all physics codes
- Multigroup neutron cross sections calculated as function of temperature and density (SCALE/XSPROC)
- SP_N neutron transport used to calculate power distribution (DENOVO)
- Subchannel thermal-hydraulics in coolant (COBRA-TF)
- Rod-by-Rod heat conduction in fuel rods (COBRA-TF)
- Simulation ran in 17.5 hours on Titan using 18,769 cores – over 1M unique material (fuel/coolant/internals) regions resolved

Next Steps

- Add fuel depletion and core shuffling
- Compare results to plant measured data

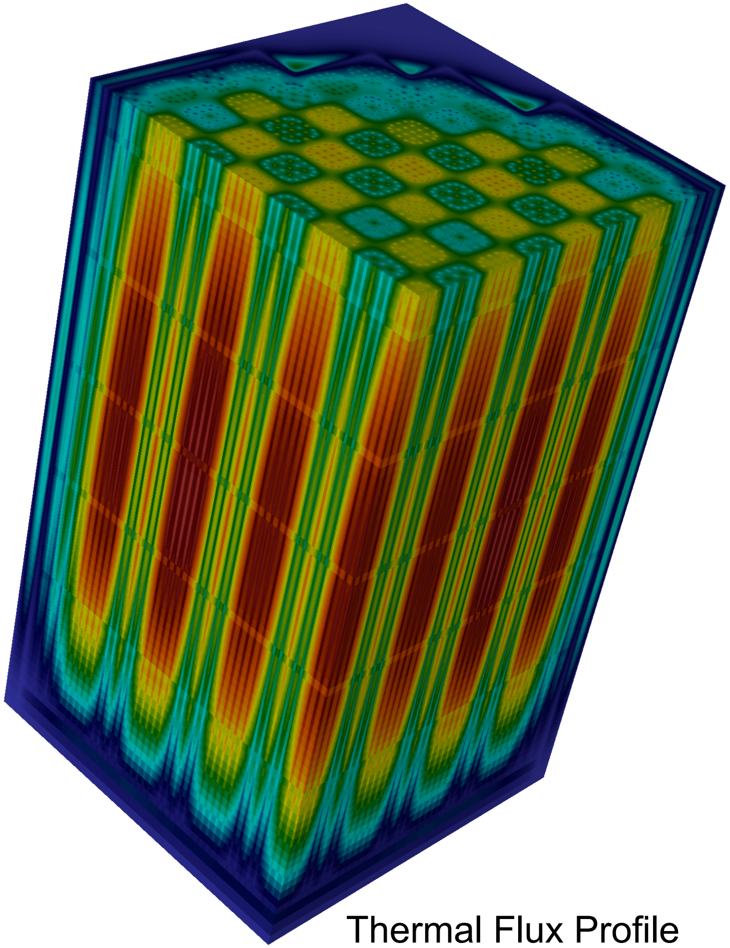

Thermal Flux Profile in Reactor Core

"Demonstration of Neutronics Coupled to Thermal-Hydraulics for a Full-Core Problem using VERA," CASL-U-2013-0196-000, Dec 2013.

Evaluate New VERA Continuous-Energy Monte Carlo Capability (Shift) – Quarter-Core Zero Power Physics Test

Goals

- Compare fidelity and performance of Shift against Keno, SP_N, and S_N (Denovo)
- Generate high-fidelity neutronics solution for code comparison of solutions for predicting reactor startup and physics testing

Contributors

Tom Evans
Fausto Franceschini
Andrew Godfrey
Steve Hamilton
Wayne Joubert
John Turner

Execution

- Proposal submitted to OLCF as part of Titan Early Science program
- Awarded 60 million core-hours on Titan (worth >$2M)
- AP1000 model created and results generated for reactor criticality, rod worth, and reactivity coefficients
- Identical VERA Input models used for Shift, SP_N, and S_N
  - dramatically simpler than KENO-VI input model

Results

- Some of the largest Monte Carlo calculations ever performed (1 trillion particles) have been completed
  - runs use 230,000 cores of Titan or more
- Excellent agreement with KENO-VI
- Extremely fine-mesh S_N calculations, which leverage Titan's GPU accelerators, are under way

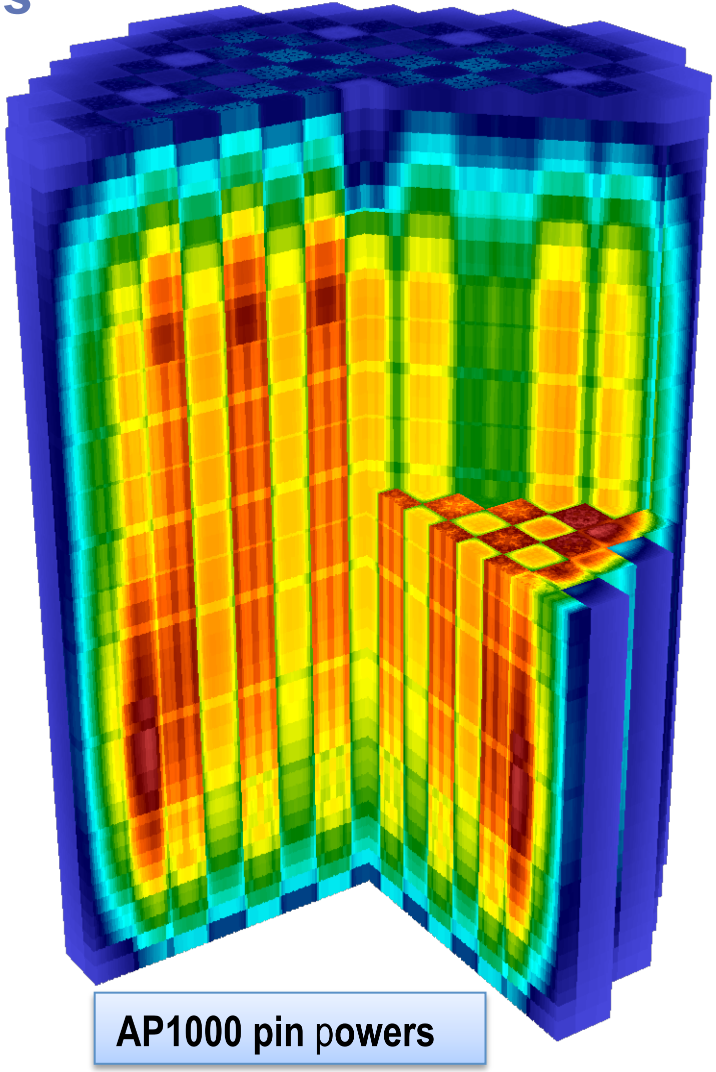

AP1000 pin powers

AP1000 core models generated with VERA-CS & KENO

AP1000 First Core

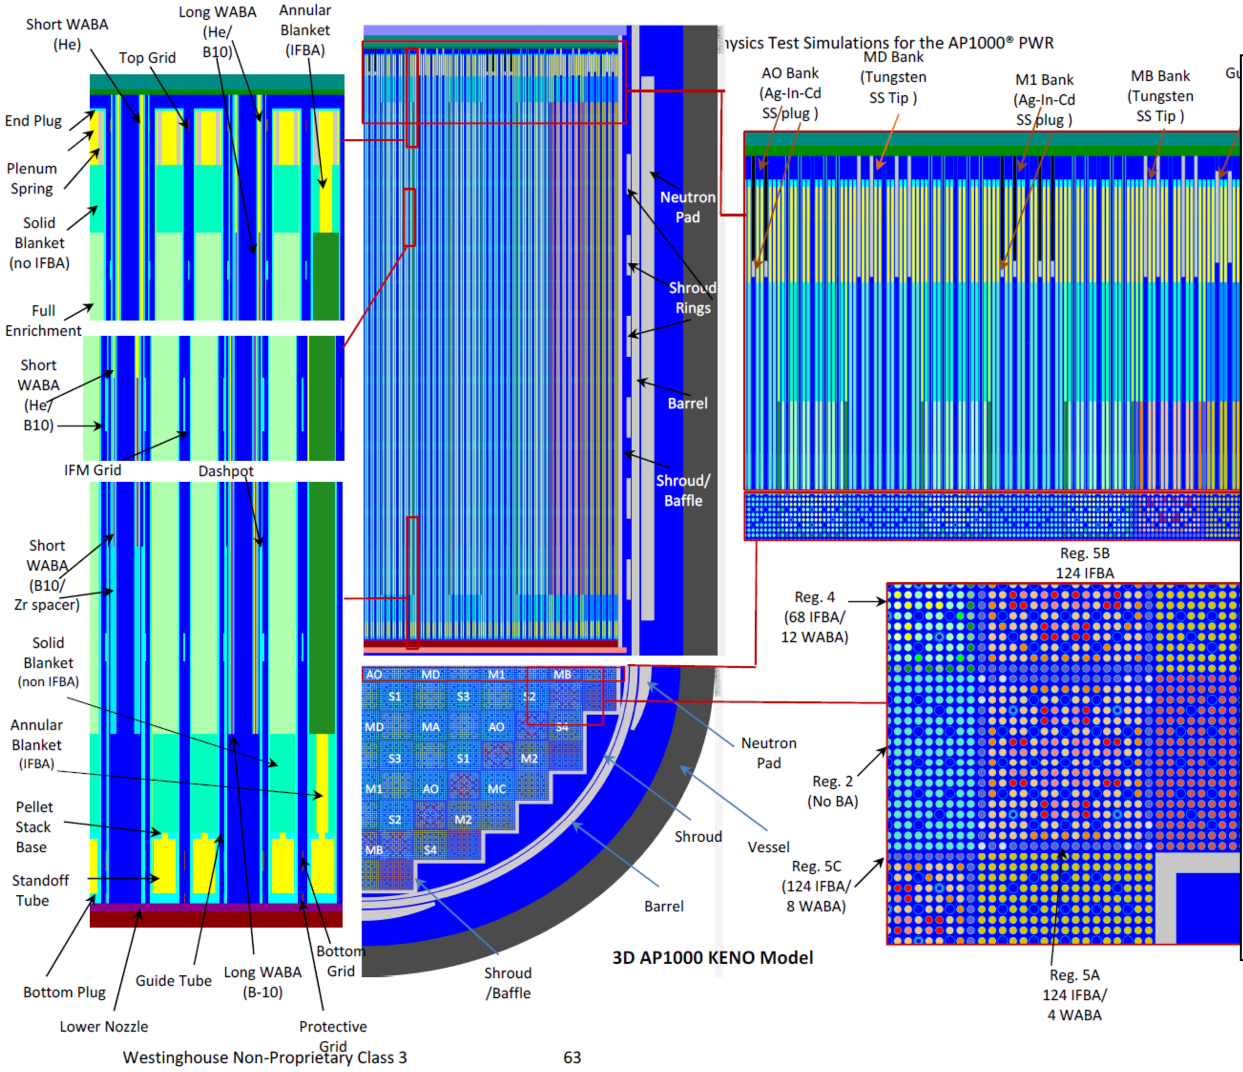

3D AP1000 KENO Model

VERA Input

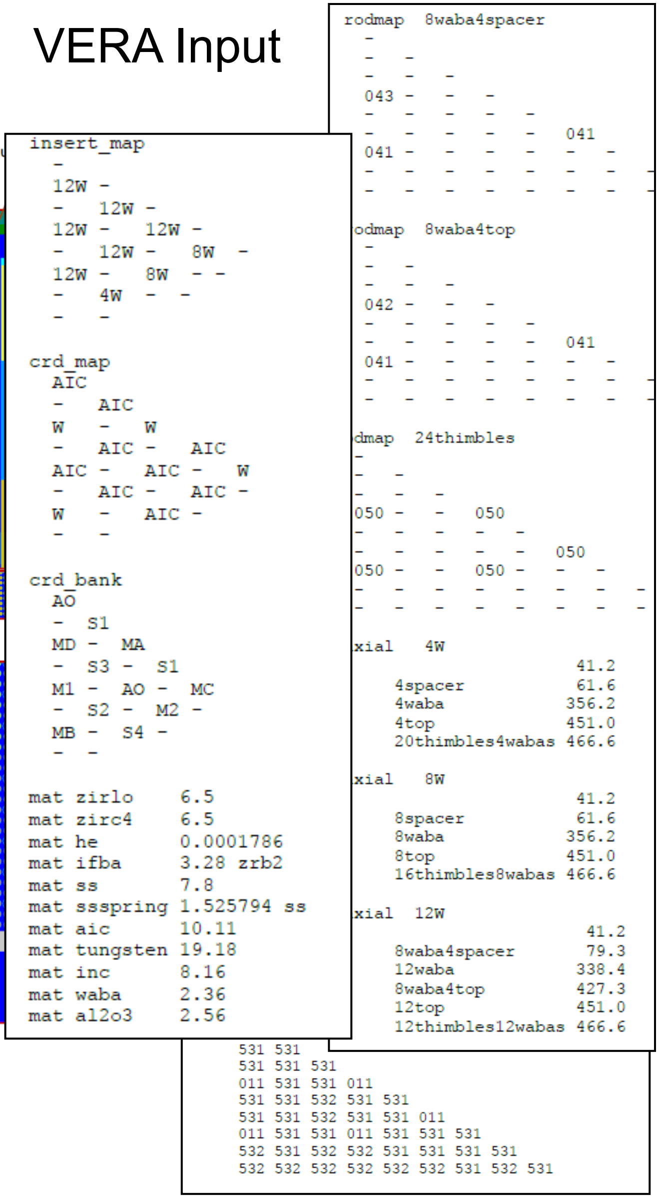

```
insert_map
  -
  12W -
  -   12W -
  12W -   12W -
  -   12W -   8W  -
  12W -   8W  - -
  -   4W  -  -
  -   -

crd_map
  AIC
  -   AIC
  W   -   W
  -   AIC -   AIC
  AIC -   AIC -   W
  -   AIC -   AIC -
  W   -   AIC -
  -   -

crd_bank
  AO
  -  S1
  MD -  MA
  -  S3 -  S1
  M1 -  AO -  MC
  -  S2 -  M2 -
  MB -  S4 -
  -  -

mat zirlo    6.5
mat zirc4    6.5
mat he       0.0001786
mat ifba     3.28 zrb2
mat ss       7.8
mat ssspring 1.525794 ss
mat aic      10.11
mat tungsten 19.18
mat inc      8.16
mat waba     2.36
mat al2o3    2.56
```

```
rodmap  8waba4spacer
  -
  -   -
  -   -   -
  043 -   -   -
  -   -   -   -   -
  -   -   -   -   -   041
  041 -   -   -   -   -   -
  -   -   -   -   -   -   -   -
  -   -   -   -   -   -   -   -

odmap  8waba4top
  -
  -   -
  -   -   -
  042 -   -   -
  -   -   -   -   -
  -   -   -   -   -   041
  041 -   -   -   -   -   -
  -   -   -   -   -   -   -   -
  -   -   -   -   -   -   -   -

dmap  24thimbles
-
-   -
-   -   -
050 -   -   050
-   -   -   -   -
-   -   -   -   -   050
050 -   -   050 -   -   -
-   -   -   -   -   -   -   -
-   -   -   -   -   -   -   -

xial   4W
                          41.2
    4spacer               61.6
    4waba                356.2
    4top                 451.0
    20thimbles4wabas     466.6

xial   8W
                          41.2
    8spacer               61.6
    8waba                356.2
    8top                 451.0
    16thimbles8wabas     466.6

xial  12W
                          41.2
    8waba4spacer          79.3
    12waba               338.4
    8waba4top            427.3
    12top                451.0
    12thimbles12wabas    466.6
```

```
531 531
531 531 531
011 531 531 011
531 531 532 531 531
531 531 532 531 531 011
011 531 531 011 531 531 531
532 531 532 532 531 531 531 531
532 532 532 532 532 532 531 532 531
```

<1,000 lines of input for VERA-CS vs. Million lines for KENO