High performance computing & data sciences for reactor systems

Workshop on Computational Nuclear Science and Engineering

Jean Ragusa,

jean.ragusa@tamu.edu

Nuclear Engineering / Institute for Scientific Computation / Center for Large Scale Scientific Simulations / Center for Exascale Radiation Transport

Texas A&M University

TEXAS A&M UNIVERSITY
Department of Nuclear Engineering

July 16, 2021 IAEA (online)

Outline

- Brief (picturesque) bio
- High-performance computing (HPC)
  - Some history
  - Some well-recognized software used in nuclear engineering
  - A few application examples:
    - Thermal-hydraulics (from Argonne Nat'l Lab)
    - Neutron Transport (from Argonne Nat'l Lab)
    - Neutral-particle Transport (from Texas A&M U.)
    - Multiphysics simulations of molten salt reactor (from CNRS/Texas A&M U.)
- Data sciences with HPC
  - Motivations (multi-query problems)
  - Data-driven model-order reduction
  - Application to multiphysics simulation of molten salt reactor (from Texas A&M U.)
- Conclusions and Outlook

← That's me (on GitHub)

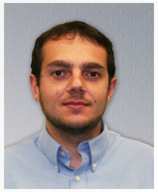

← That's me (in real life)

- Accelerator-driven production of tritium

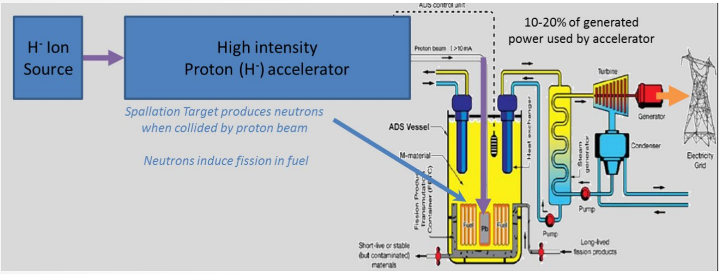

+ $^{7}LiF - BeF_2 - {}^{233}UF_4$

- Near-real-time PWR accident simulator for crisis management

Reactor physics and Applied Math Department

- 2004-present: Nuclear engineering, Texas A&M U.

Computational radiation transport, Multiphysics, and Predictive science https://multiphysics.engr.tamu.edu/

My research interests

- **Radiation transport**

  Predictive Science Academic Alliance Program (PSAAP)

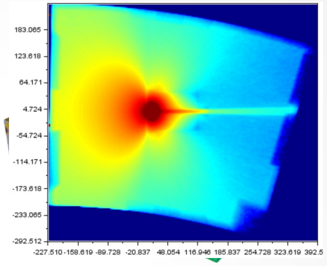

NNSA National Nuclear Security Administration

CERT
Center for Exascale Radiation Transport
http://class.tamu.edu/cert

  Stockpile stewardship Lawrence Livermore National Laboratory

- **Multiphysics** software development (RELAP-7, RattleSNake, Pronghorn)

MOOSE INL Idaho National Laboratory

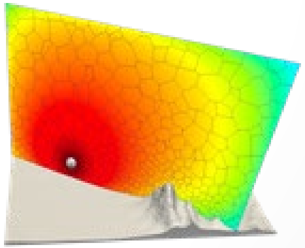

- **Data sciences and machine-learning**

  Nuclear radiation effects DTRA DEFENSE THREAT REDUCTION AGENCY *Detect. Deter. Defeat.*

  Multiphysics model reduction

NEUP Nuclear Energy University Program
U.S. Department of Energy

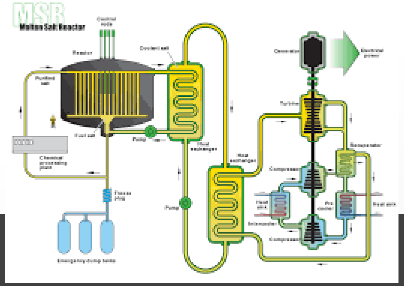

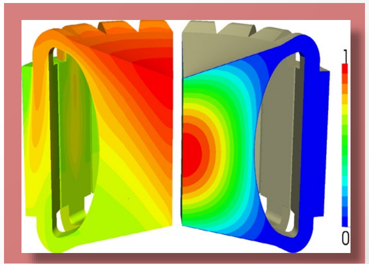

Outline

- Brief (picturesque) bio
- High-performance computing (HPC)
  - Some history
  - Some well-recognized software used in nuclear engineering
  - A few application examples
- Data sciences with HPC
  - Motivations (multi-query problems)
  - Data-driven model-order reduction
  - Application to multiphysics simulation of molten salt reactor
- Conclusions and Outlook

Evolution of processor speeds

Instruction-level parallelism:

- Architectural technique that allows the overlap of individual machine operations (add, mult, load, store, ...)

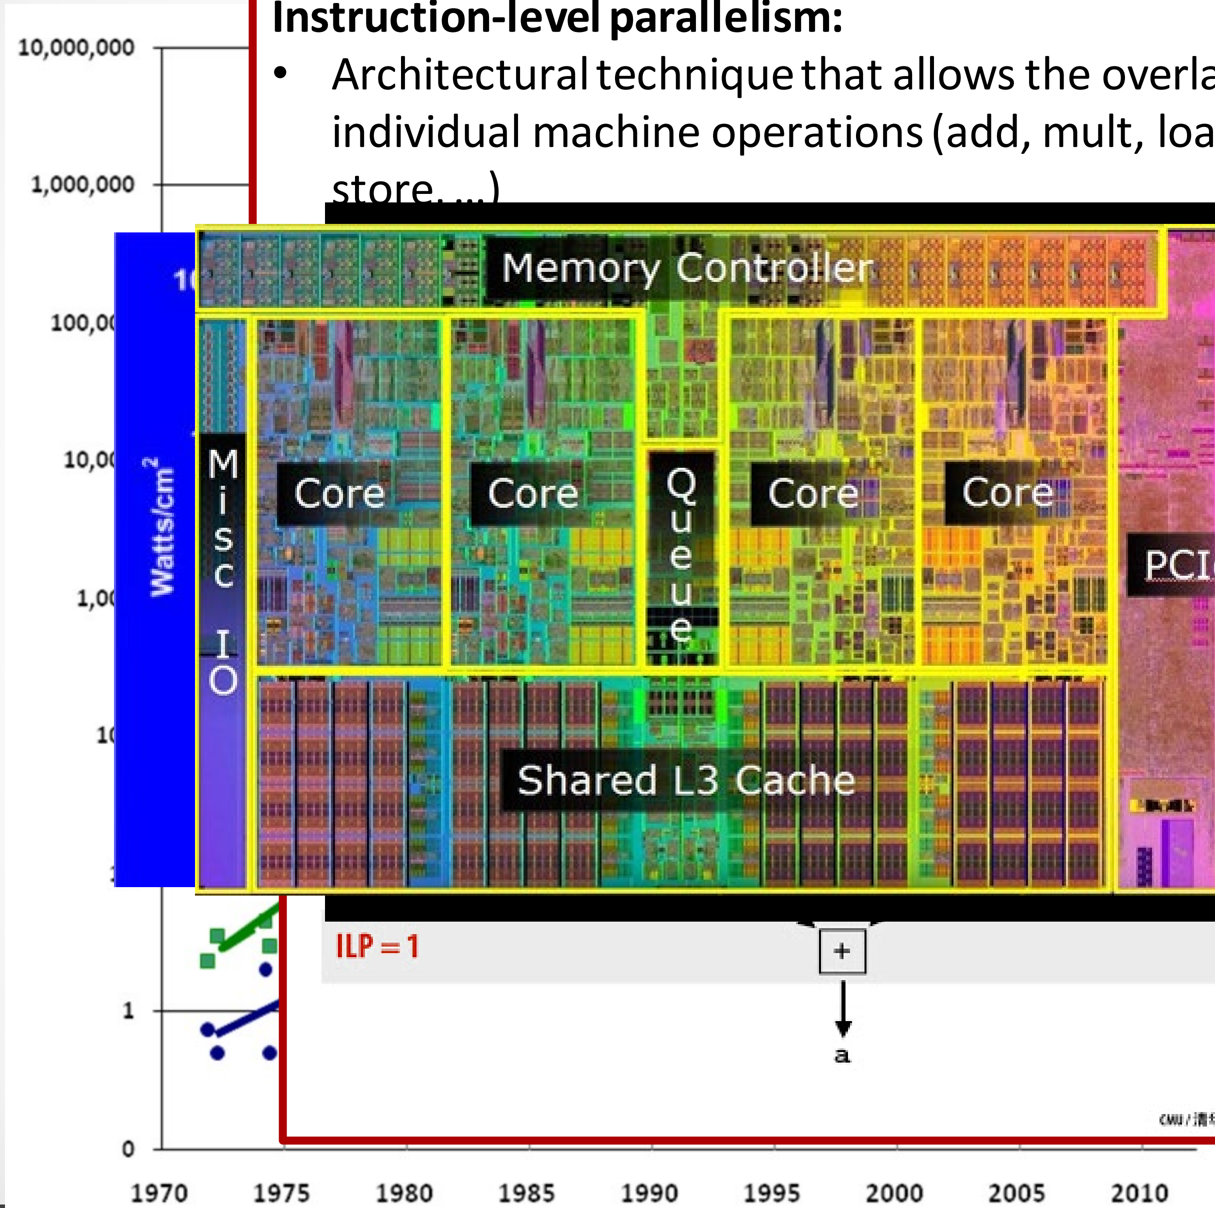

unt still rising
esigns);
re-interpreted"

d has flattened

nsumption
he issue)

structions
on-level
m or ILP)

Moore's law: ~straight line on semilog scale

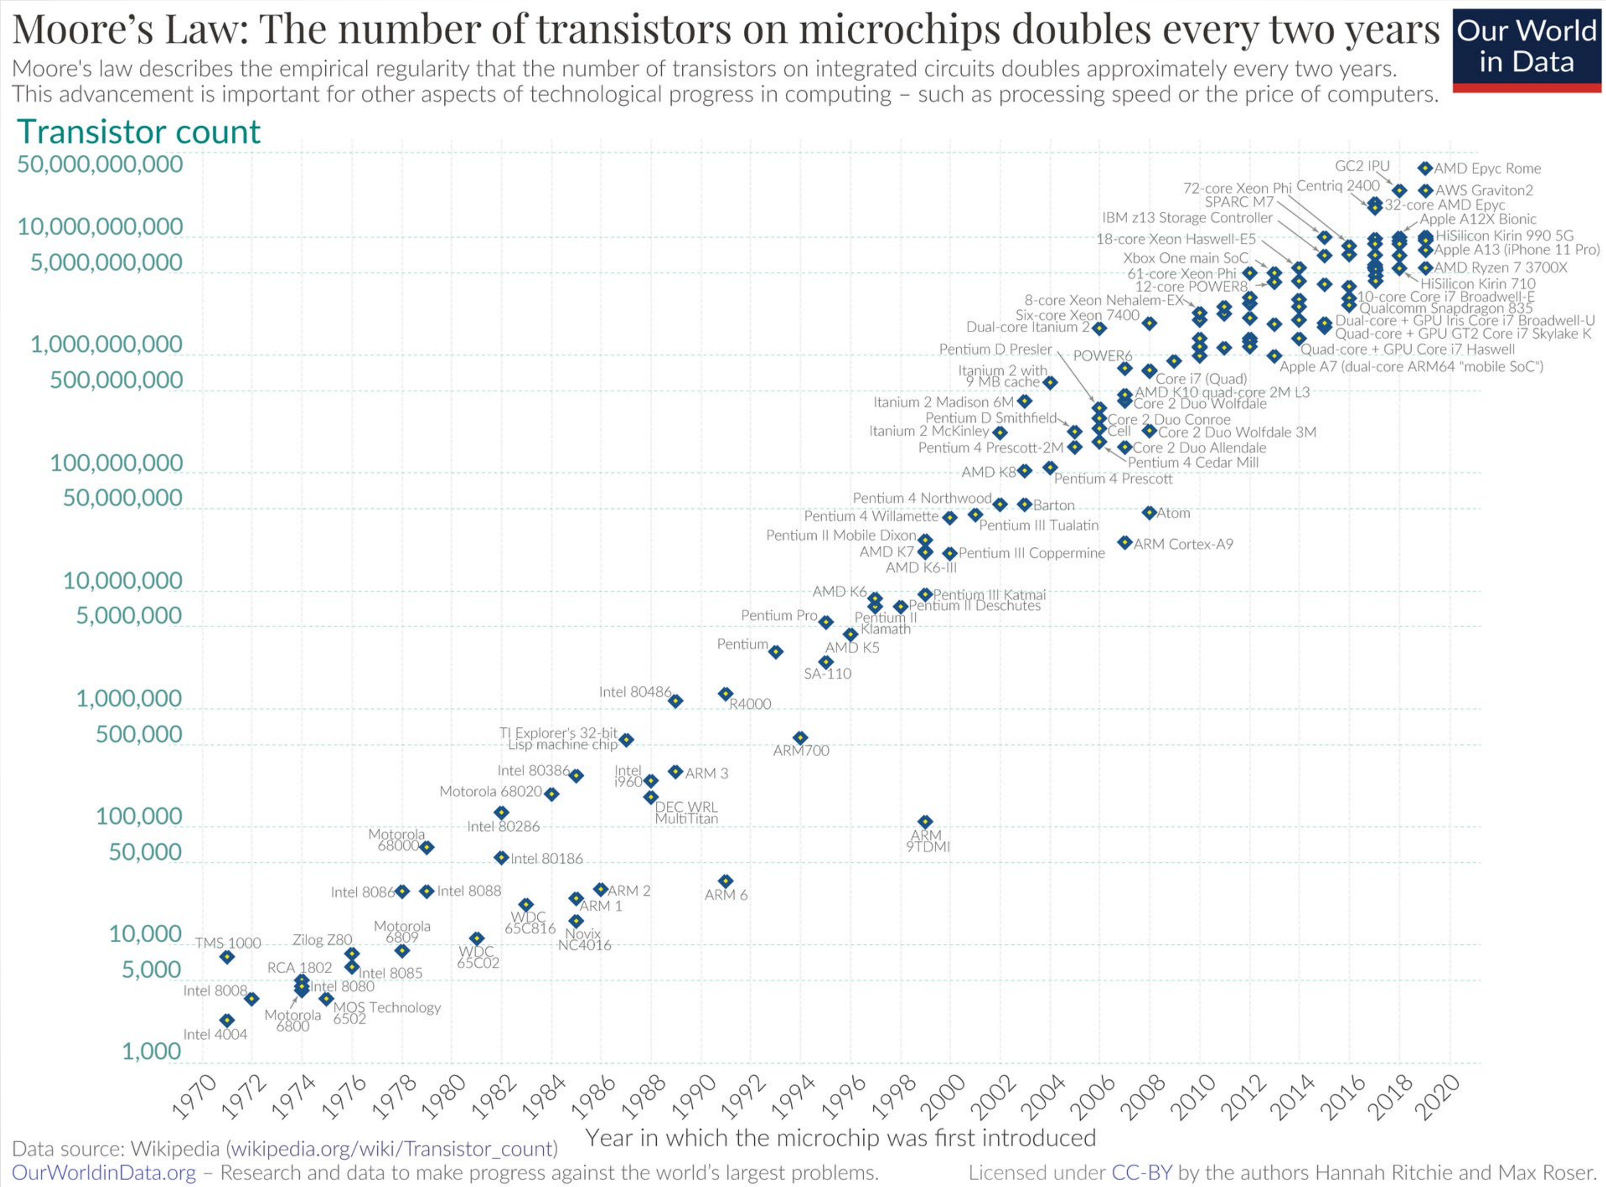

FLOP costs decreasing faster than RAM costs

- *Our ability to sense, collect, generate and calculate on data is growing faster than our ability to access, manage and even "store" that data*

Source: David Turek, IBM

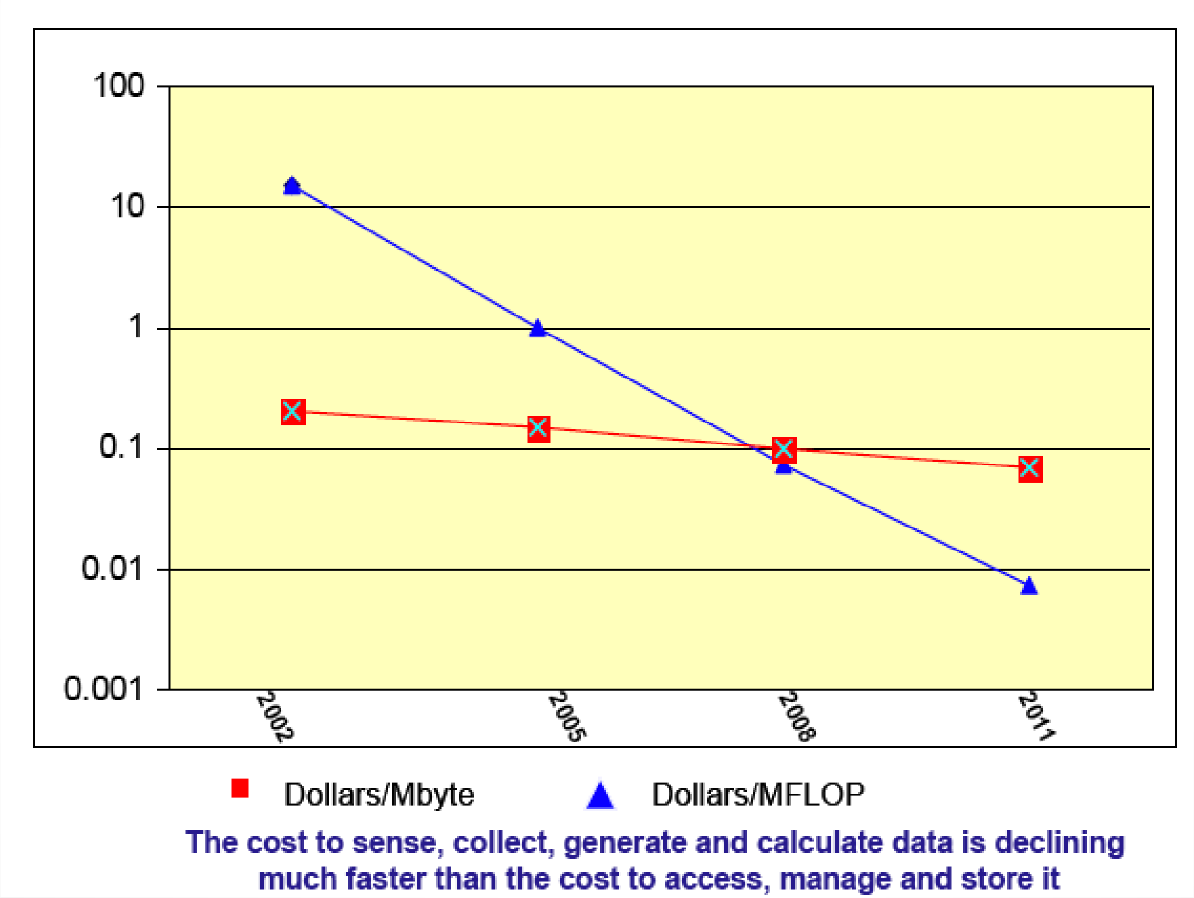

Perspective

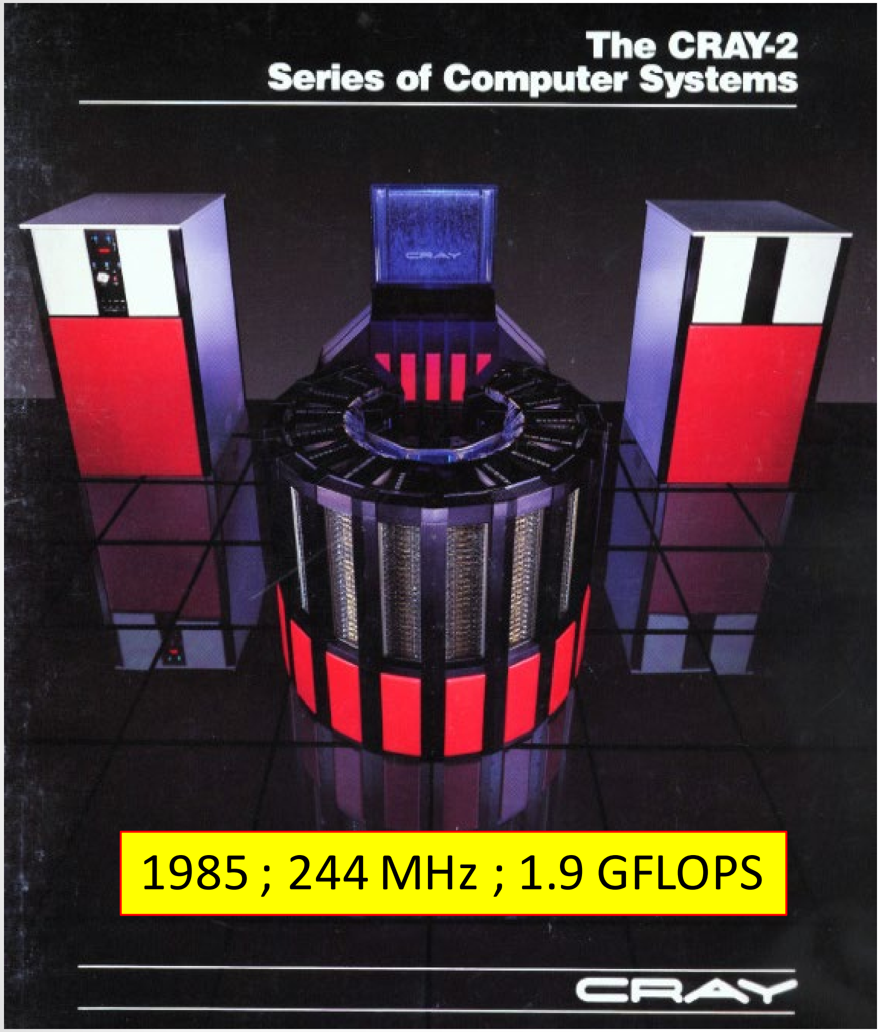

1985 ; 244 MHz ; 1.9 GFLOPS

2010; 800 MHz ; 1.6 GFLOPS

2015; 1,000 MHz ; 3 GFLOPS

Top 500 list

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	Supercomputer Fugaku - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848	442,010.0	537,212.0	29,899
	Eni S.p.A. Italy				
10	Frontera - Dell C6420, Xeon Platinum 8280 28C 2.7GHz, Mellanox InfiniBand HDR, Dell EMC Texas Advanced Computing Center/Univ. of Texas United States	448,448	23,516.4	38,745.9	

Rank	System	...ak (...op/s)	Power (kW)
1	Sunwa... Sunway... Nation...	...435.9	15,371
	... Joint C... Japan		
10	K comp... RIKEN Japan	280.4	12,660

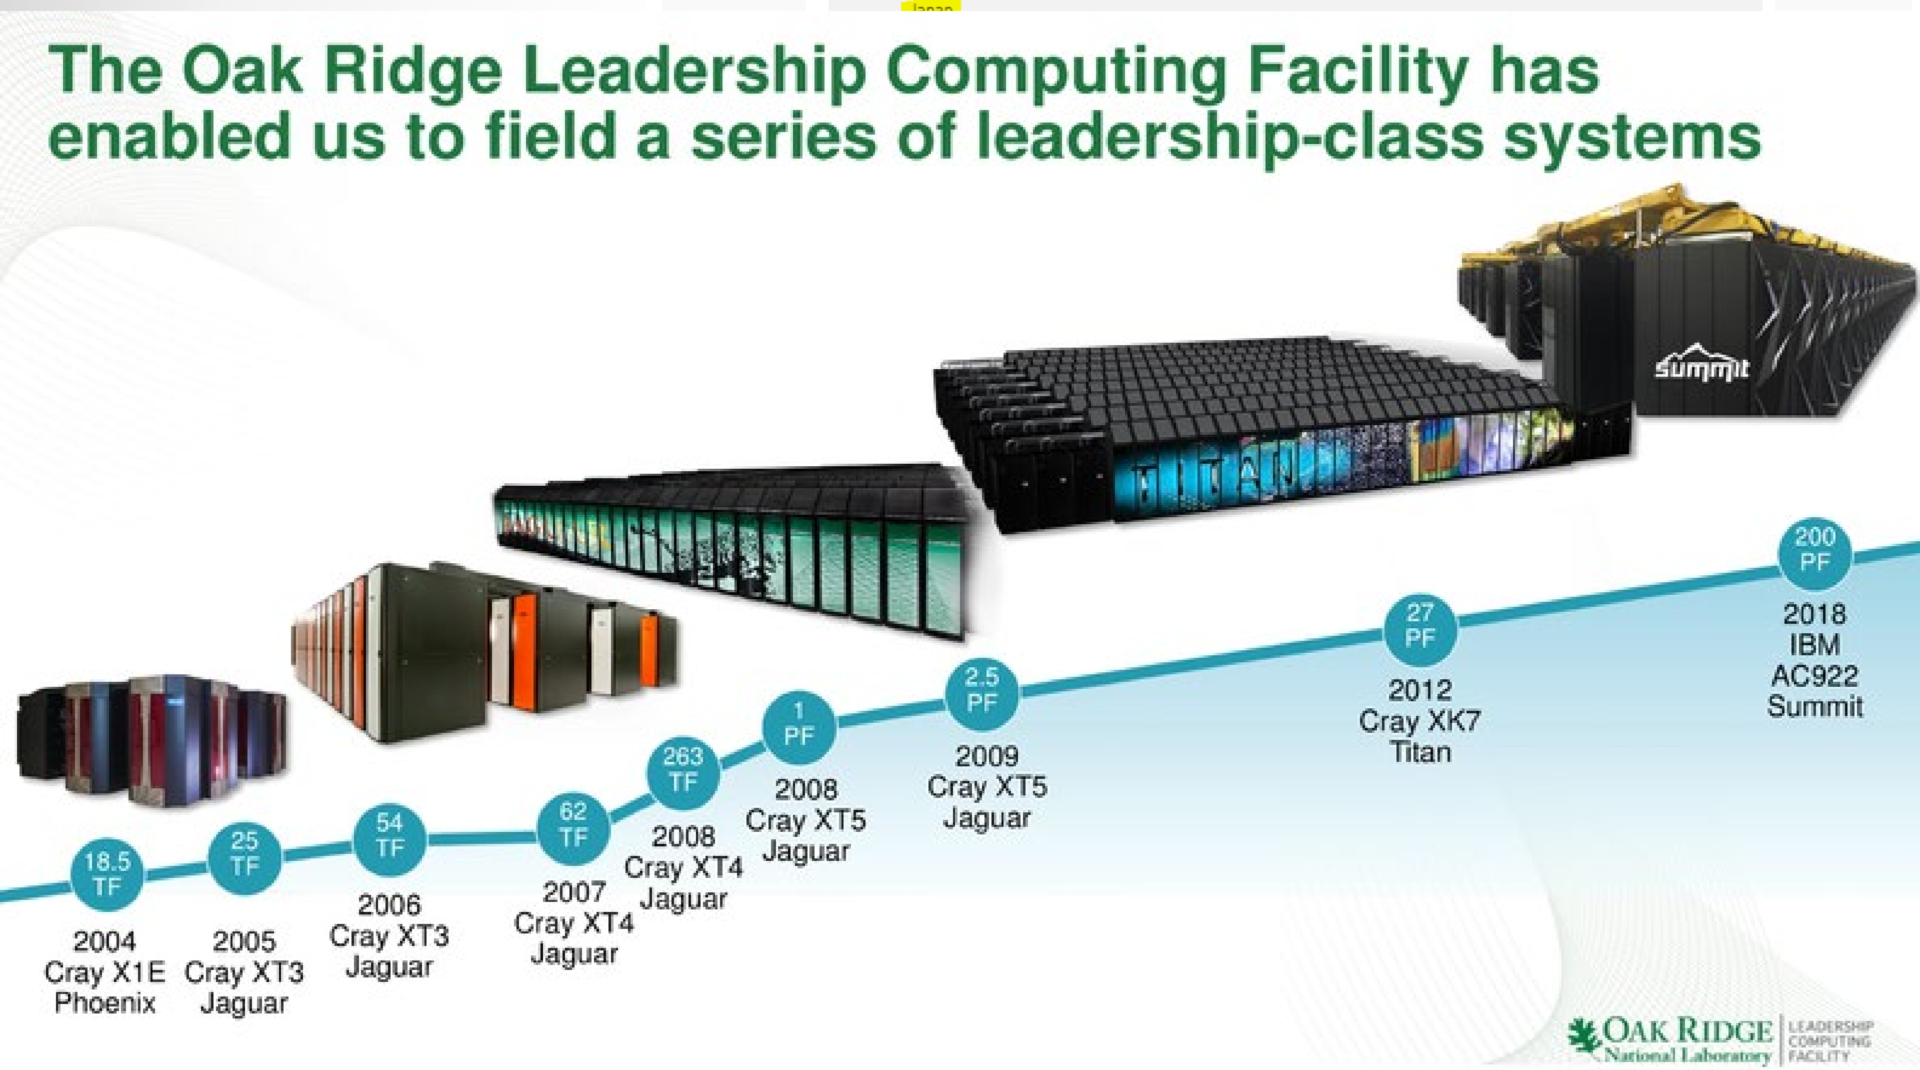

U.S. Presidential Information Technology Advisory Committee (PITAC)

- Computational science is a rapidly growing multidisciplinary field that uses advanced computing capabilities to understand and solve complex problems.
- Requires advances in hardware and software.

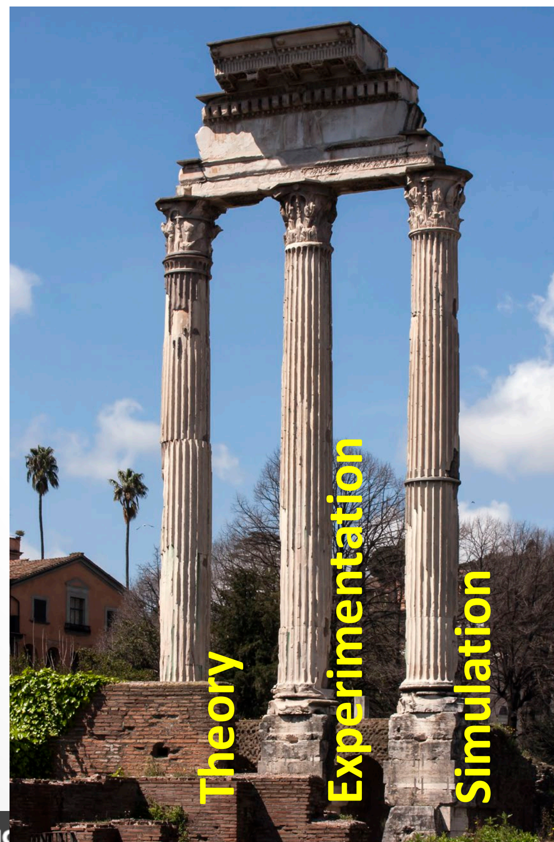

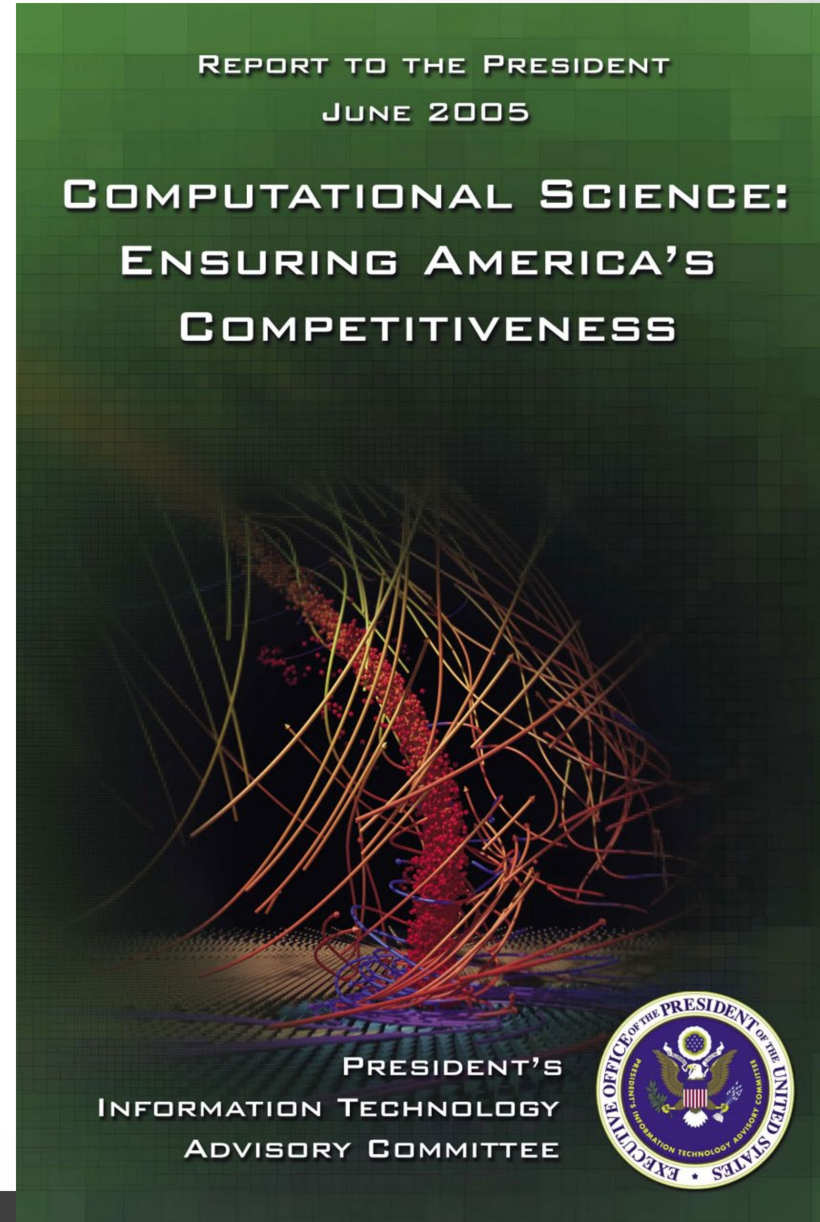

High Performance Scientific Computing

Introduction to HPSC

What is High Performance Scientific Computing?

- single process → **What is Scientific Computing**
  - Scaling Challenge →
    - application takes too long to run
    - Application doesn't fit in memory of single machine
    - Application doesn't fit & takes too long
  - Overview of computational landscape →
    - **Scientific Simulations** — Examples
    - **Data Analysis and Knowledge Extraction**
- Multi-process → **Multi-process**

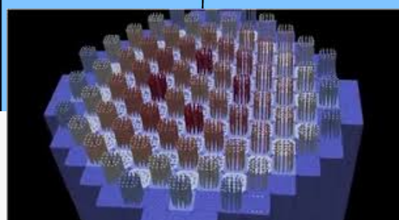

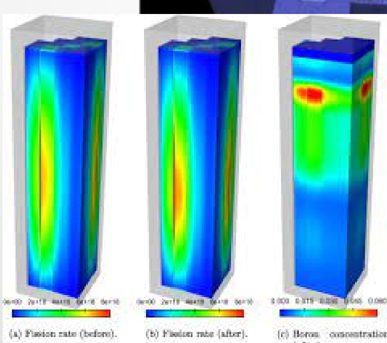

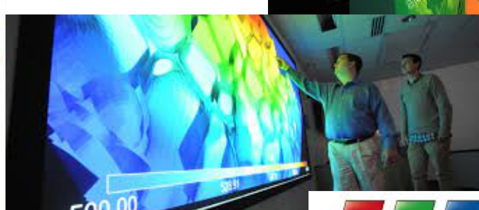

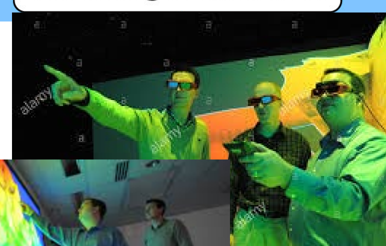

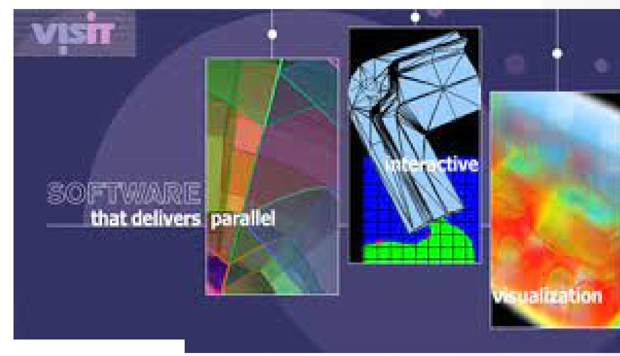

ParaView
Parallel Visualization Application

Partially adapted from the Lincoln Laboratory Supercomputing Center (MIT)

What is Scientifi Computin

Arch co

Parallel technologies: levels of parallelism

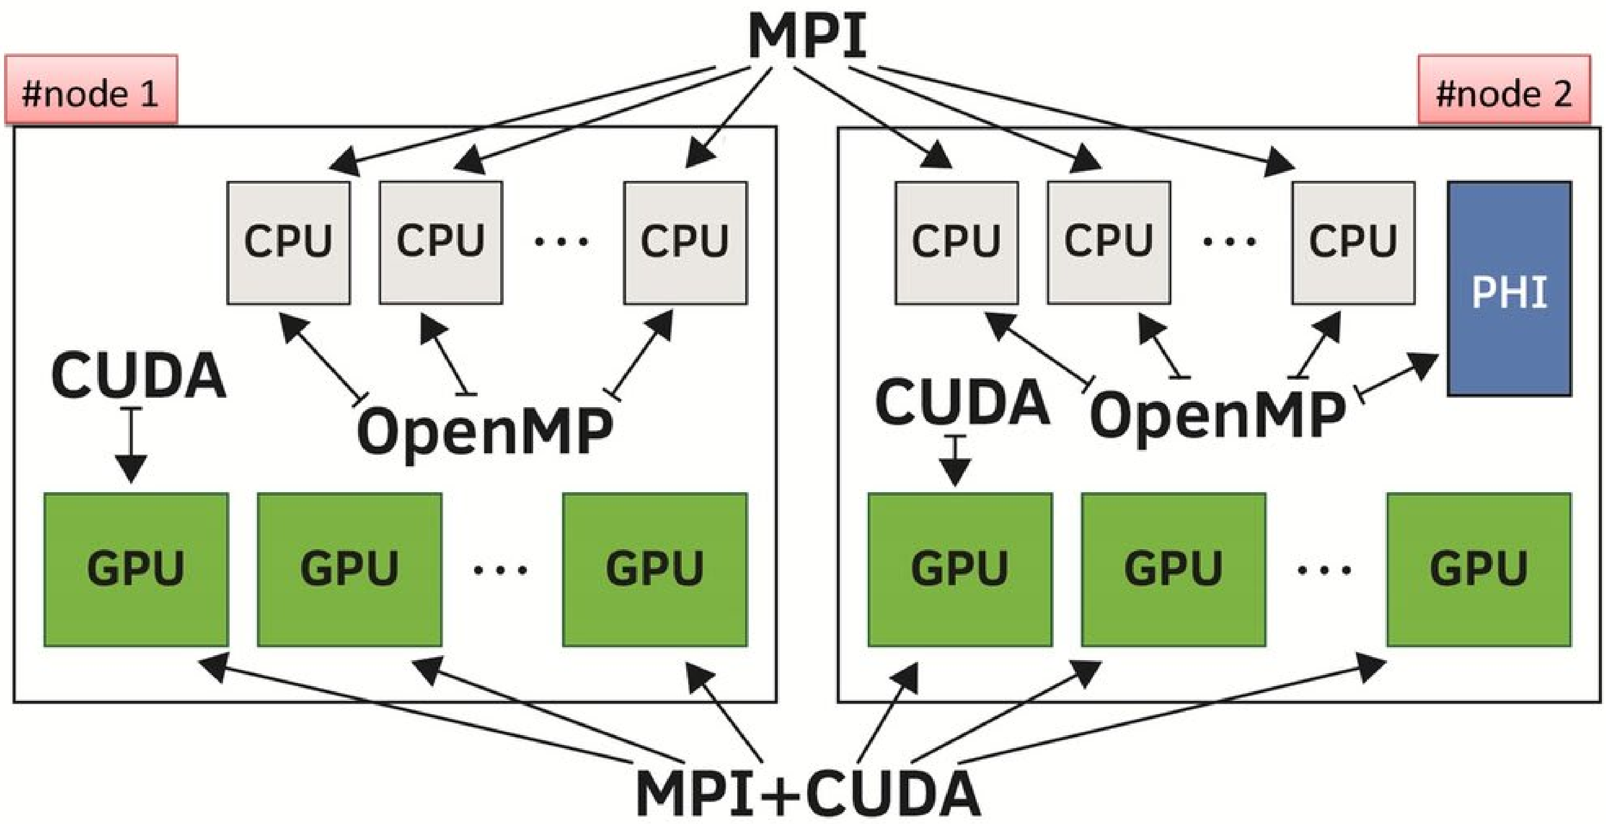

How to control hybrid hardware:
MPI – OpenMP – CUDA - OpenCL ...

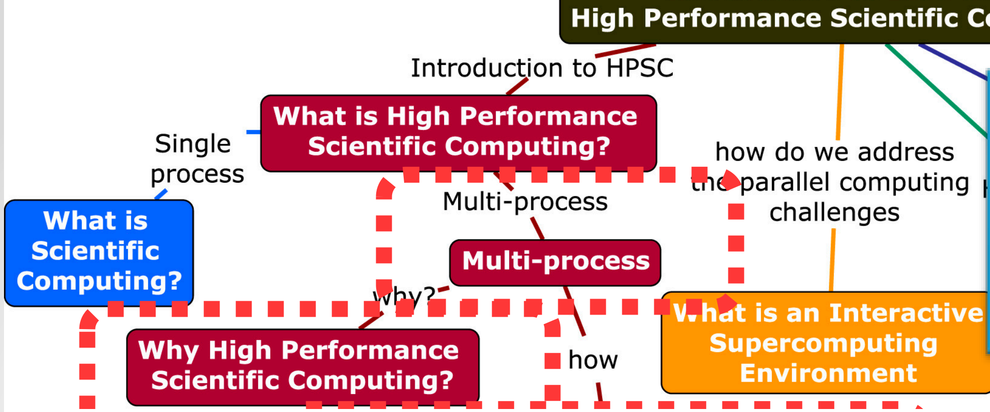
High Performance Scientific Computing
Introduction to HPSC
What is High Performance Scientific Computing?
Single process
What is Scientific Computing?
Multi-process
Multi-process
how do we address the parallel computing challenges
why?
Why High Performance Scientific Computing?
how
What is an Interactive Supercomputing Environment

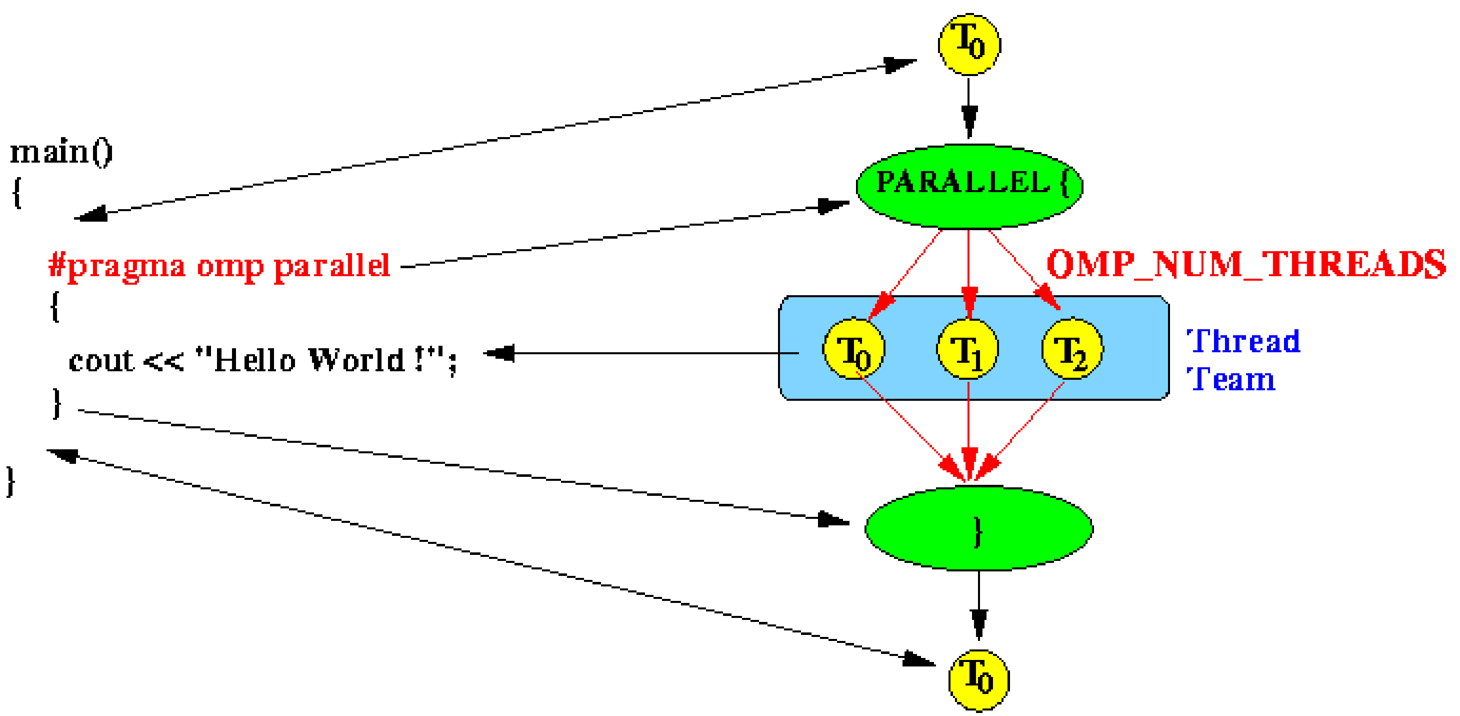
main()
{
#pragma omp parallel
{
cout << "Hello World !";
}
}
PARALLEL {
OMP_NUM_THREADS
Thread Team
}

Interacting with a supercomputer:

Job scheduling

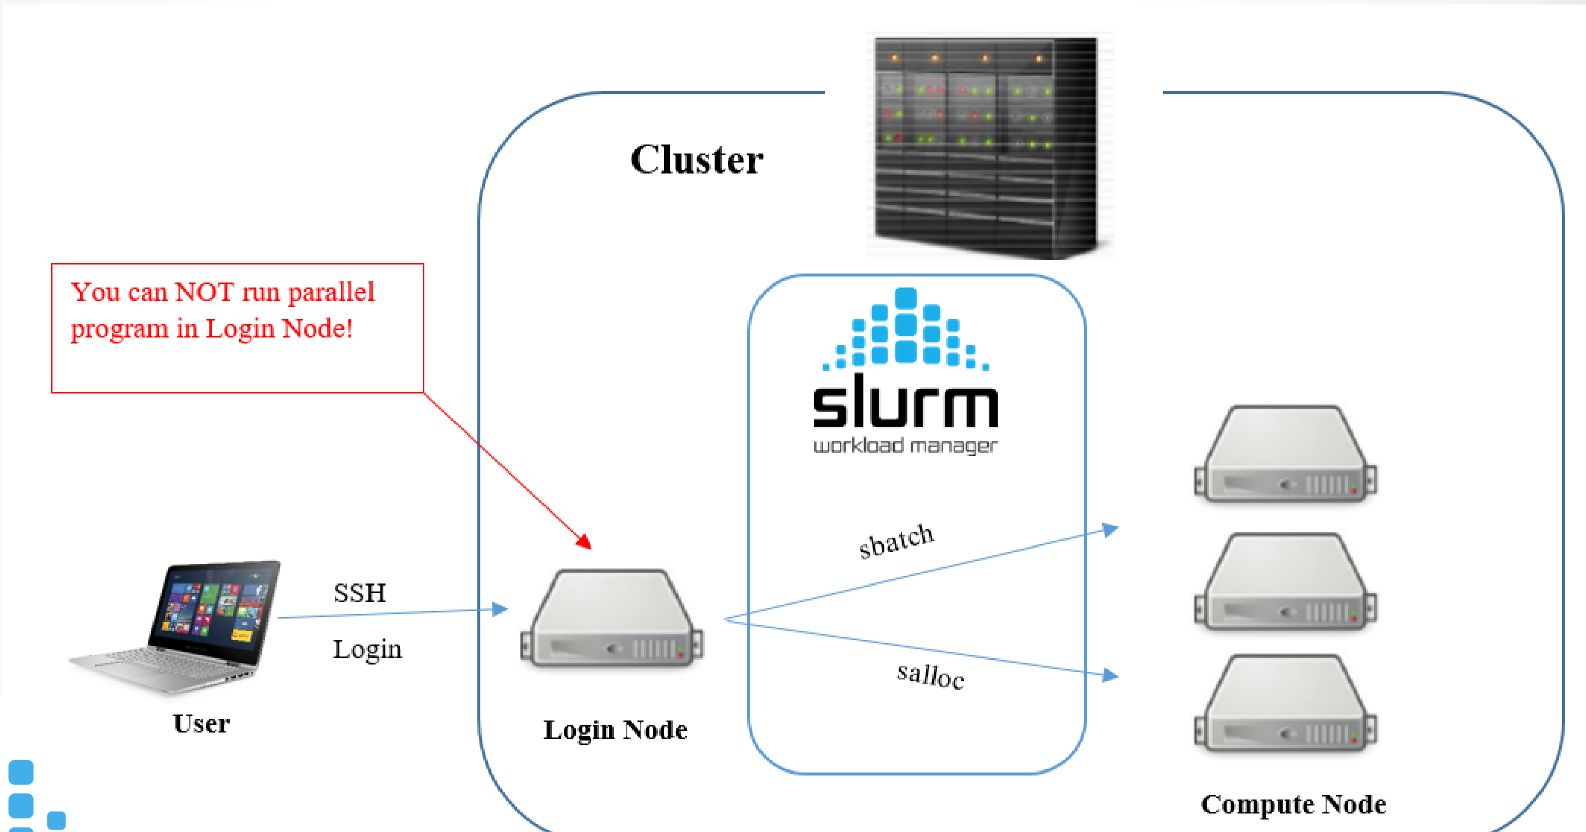

slurm
workload manager

HPC for Scientific computing (SC)

Verification
(Am I solving the equations correctly?)

Validation
(Am I solving the correct equations?)

Uncertainty Quantification
(What is the goal of my simulation? What are the QoI's?)

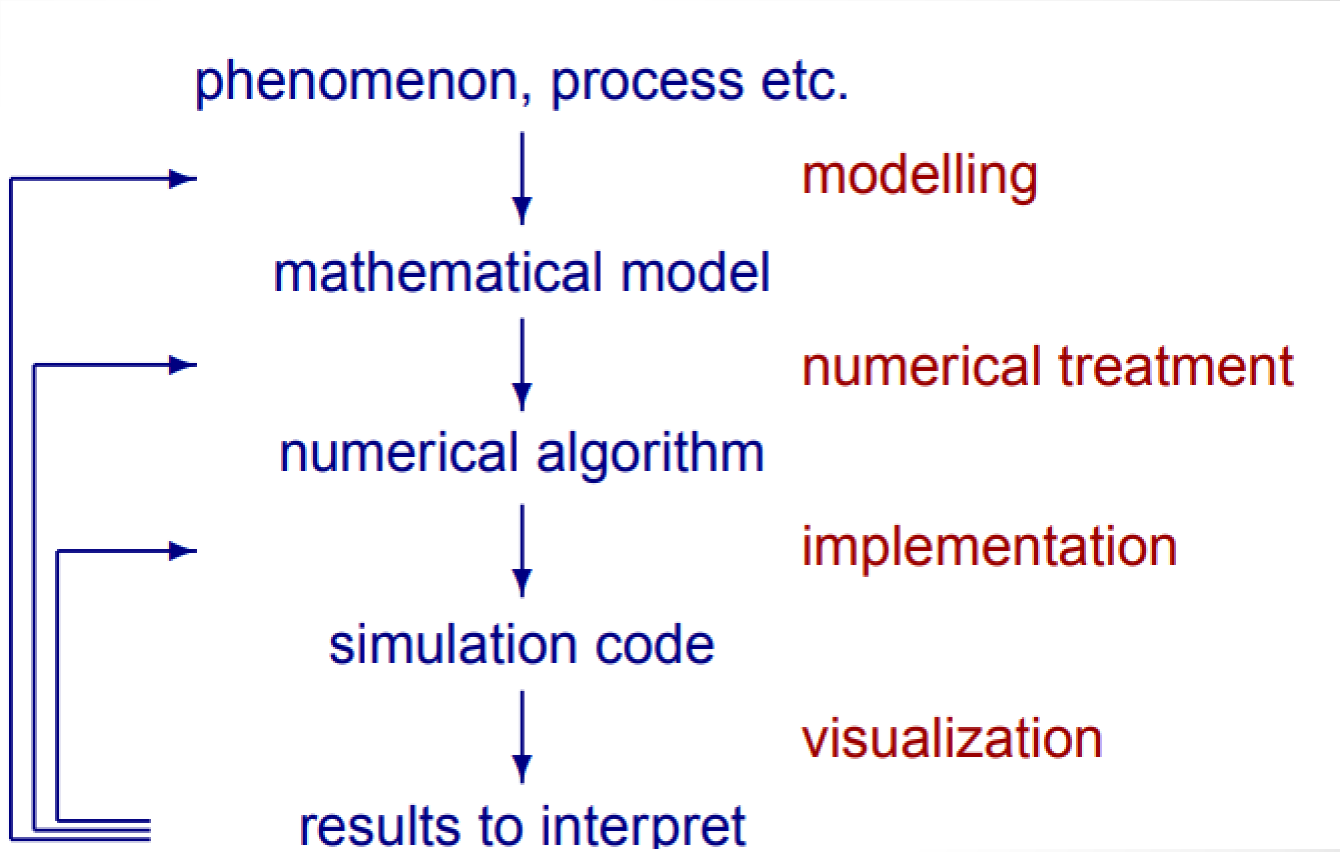

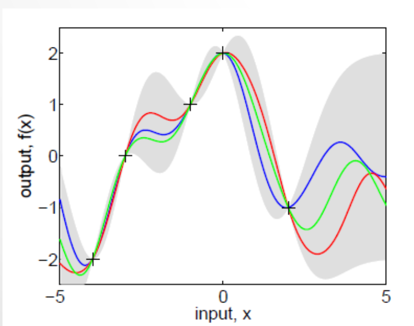

FastMath: Frameworks, Algorithms and Scalable Technologies for Mathematics

PETSc TAO

Portable, Extensible Toolkit for Scientific Computation

Toolkit for Advanced Optimization

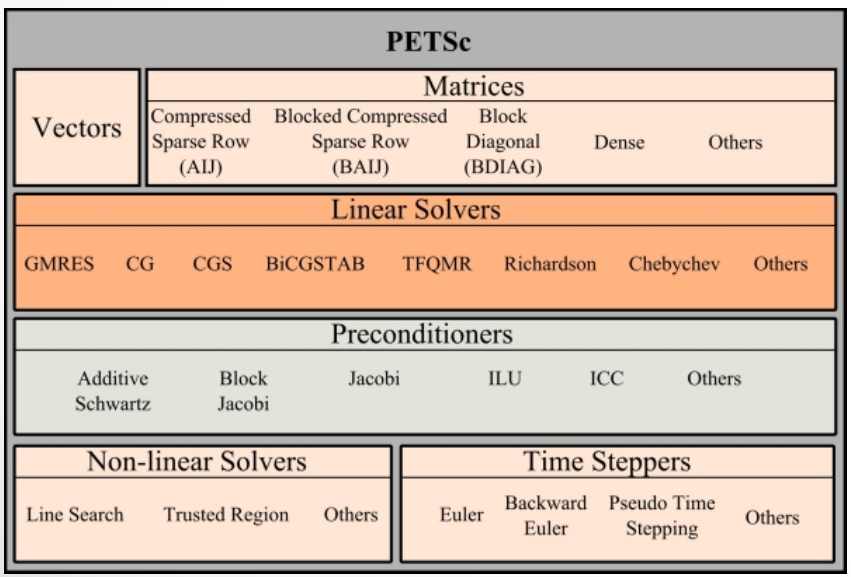

Example-1: Computational fluid dynamics

U.S. DEPARTMENT OF ENERGY

Nuclear Energy

User case: Using higher-resolution approaches to inform lower-resolution methods - 1

For complex geometries CFD-grade data is often not available.

- RANS approaches can benefit from comparison with DNS/LES

International collaboration (INERI) centered on wire-wrappers.

- Comparison between commercial codes and Nek5000
- Results are being used in the design of advanced reactors in Europe

[SFR]

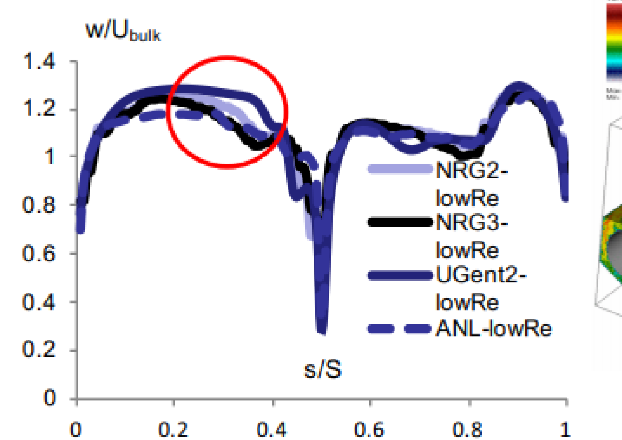

7 pin Assembly: streamwise velocity on diagonal

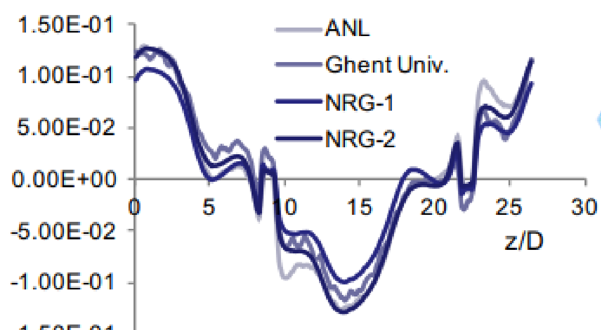

7 pin Assembly: Average Cross-flow

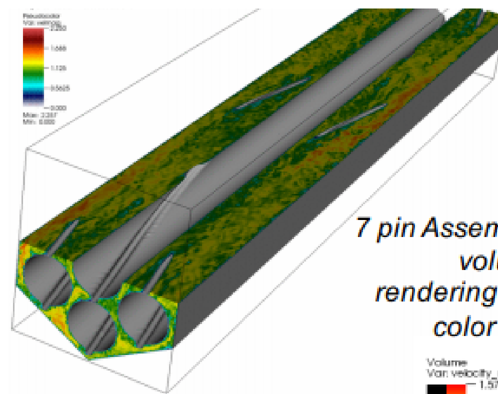

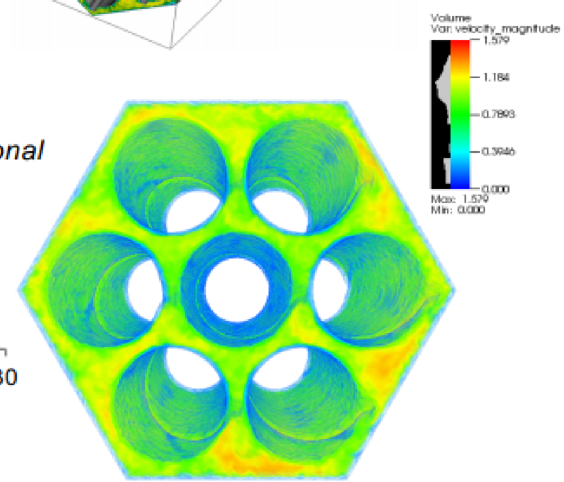

7 pin Assembly: volume rendering and color cuts

NEAMS
NUCLEAR ENERGY ADVANCED MODELING & SIMULATION PROGRAM

Example-2: MOOSE, a Multiphysics HPC platform

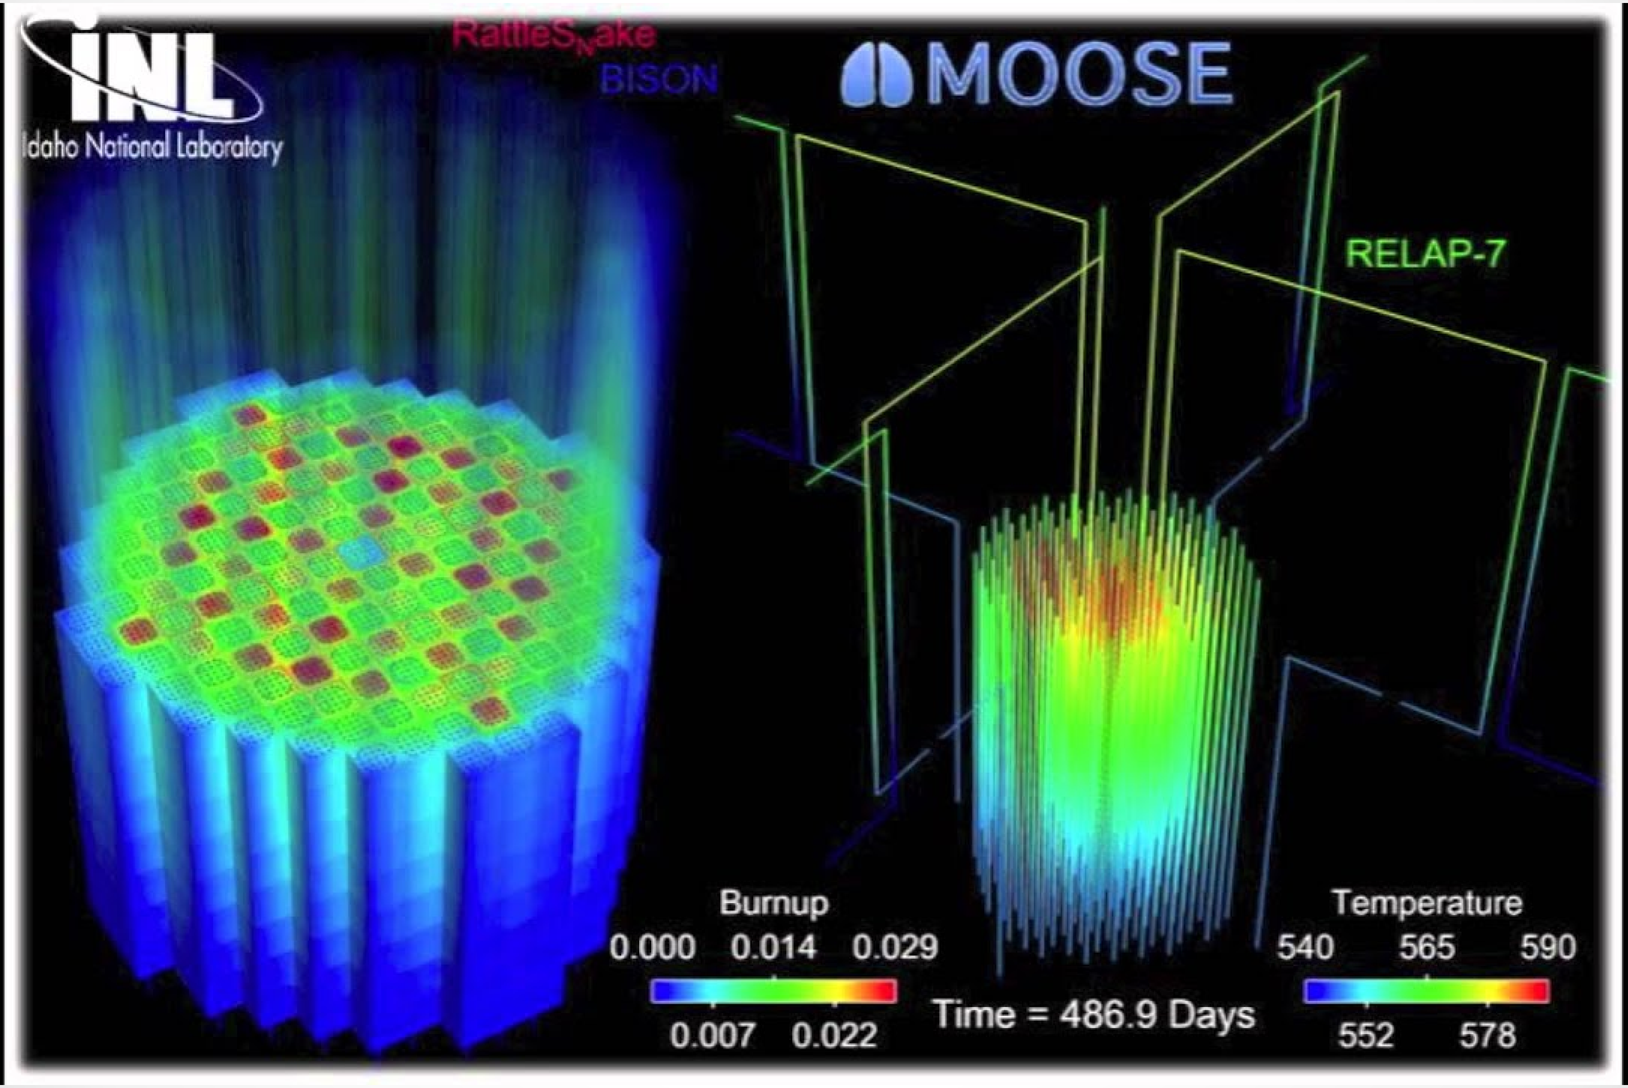

Example-3: Massively parallel radiation transport

- neutron, thermal radiation, gamma, electron
- steady-state, time-dependent, criticality, adjoint, etc.
- advanced solution techniques
- discretization in space/angle/energy
  - **Largest problem we have done: 20.8 Trillion unknowns**

TEXAS A&M UNIVERSITY
Department of
Nuclear Engineering

2018
NATIONAL DEBT OF UNITED STATES
$ 20,502,740,304,411

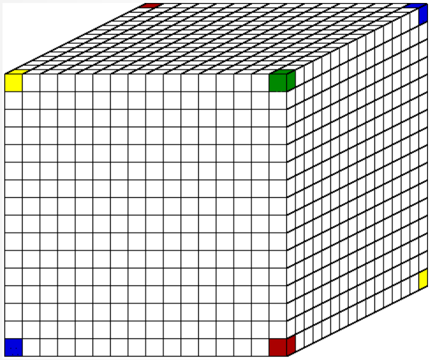

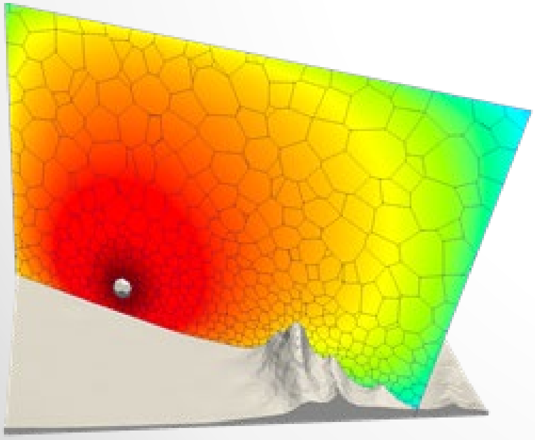

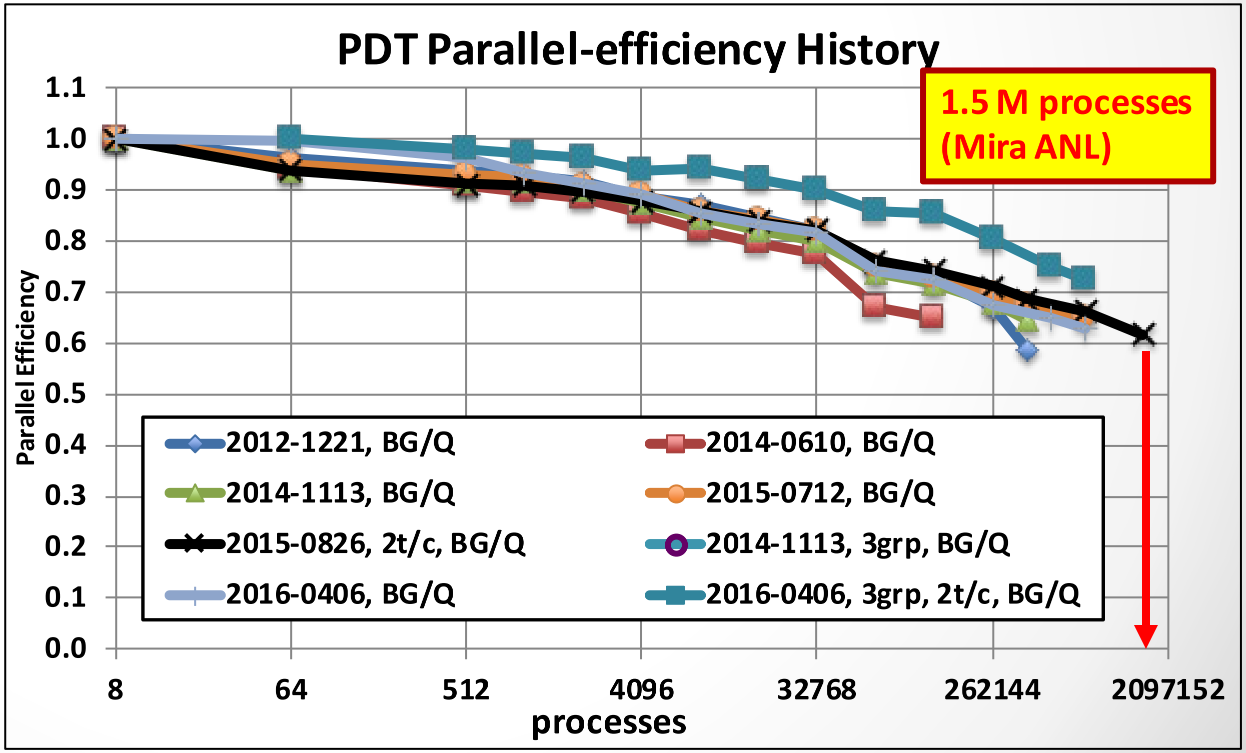

Example-4: Reactor containment

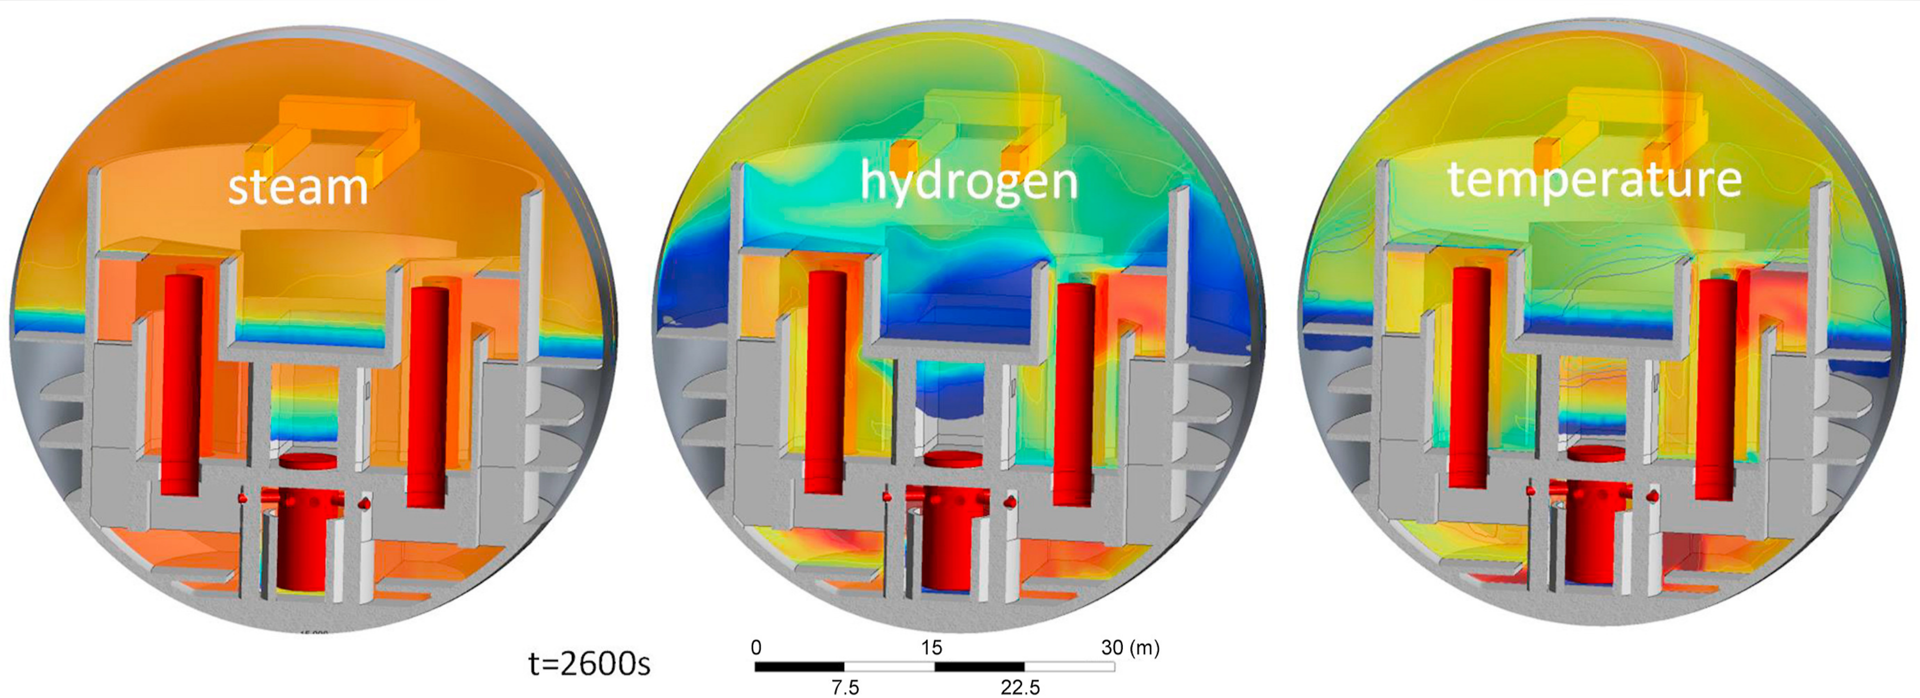

Gas distribution and pressurization inside the containment during an SB-LOCA (Julich, Germany, Kelm et al.).

Based on OpenFOAM for CFD

Example-5: Multiphysics of molten salt reactor

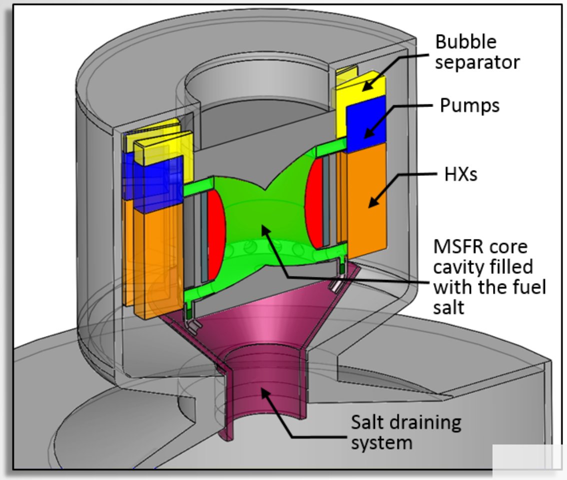

Conceptual Design of the MSFR core cavity

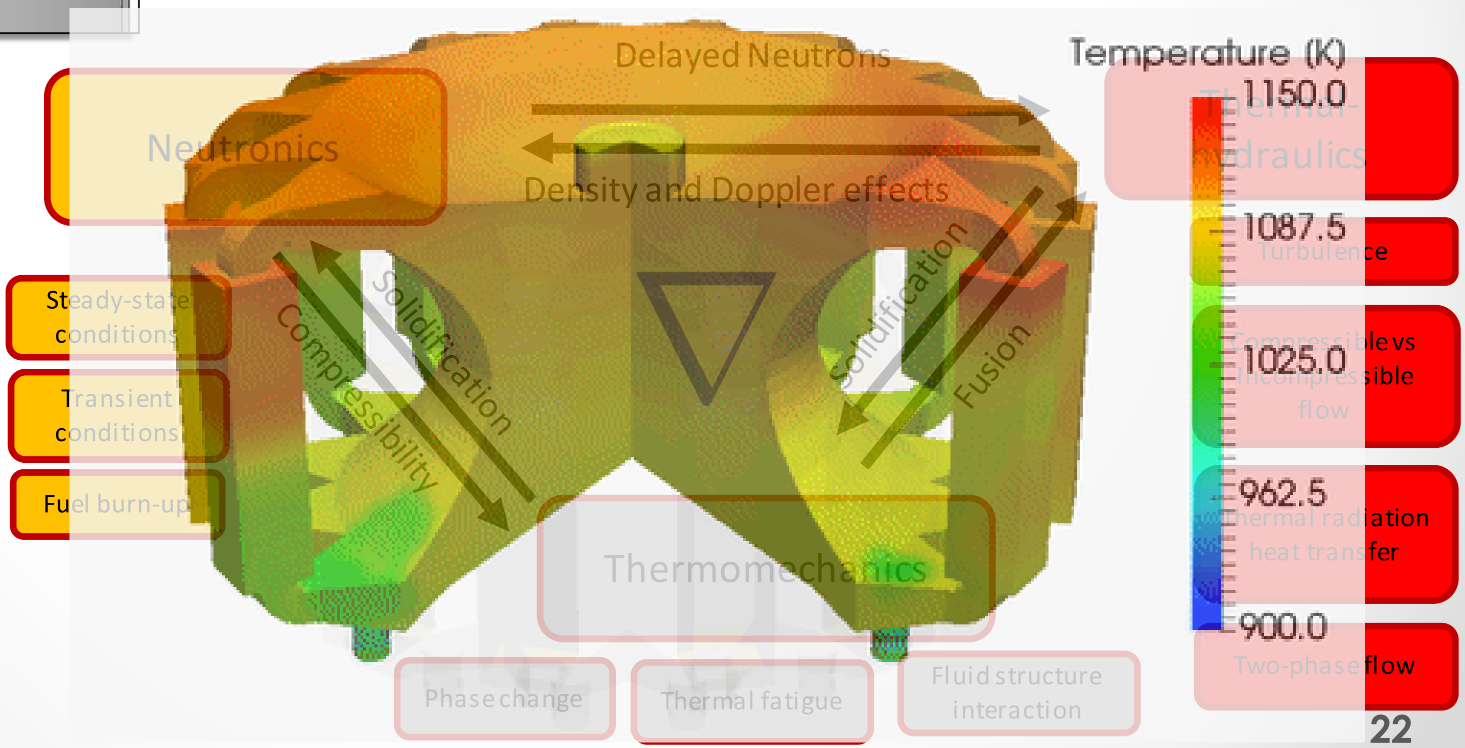

(CNRS, Tano)

Outline

- Brief (picturesque) bio
- High-performance computing (HPC)
  - Some history
  - Some well-recognized software used in nuclear engineering
  - A few application examples
- Data sciences with HPC
  - Motivations (multi-query problems)
  - Data-driven model-order reduction
  - Application to multiphysics simulation of molten salt reactor
- Conclusions and Outlook

What is SC/HPC for?

- Increase insight and understanding of physical phenomena
  - o DNS > LES > RANS > lumped-parameter TH
- Provide layers in a **<u>hierarchy</u>** of increasingly complex models
  - o Are we capturing the right physics?
  - o First-principle simulations to ascertain the range of applicability of (cheaper) low-order models
  - o New designs/configurations not handled by legacy codes
- Scaling bridging:
  - o Material science <-> continuum FEM
  - o NSF's Material Genome Initiative (MGI)
- Experimental design
  - o simulation-informed experiments
- Sometimes, too costly/dangerous experiments:
  - o Accidents (core meltdown), disasters, NW, …

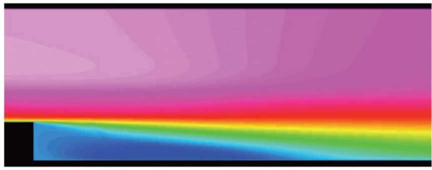

Source: Rémy Fransen, 3rd INCA colloquium, ONERA, Toulouse (2011)

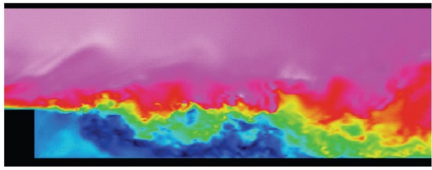

Source: Rémy Fransen, 3rd INCA colloquium, ONERA, Toulouse (2011)

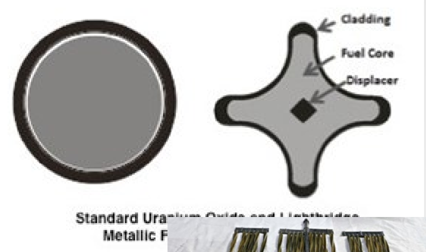

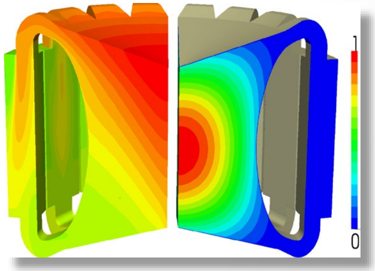

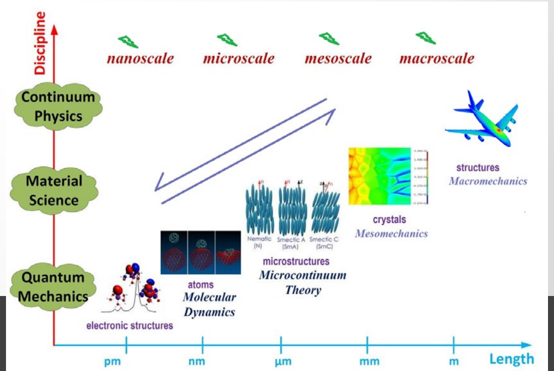

What is SC/HPC **not** for?

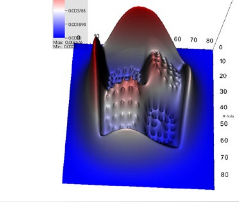

- Pretty pictures (the “viewgraph” norm = useless simulations).
- More demanding, complex simulations that do **NOT** increase your ability to **predict** the outcome (the quantities of interest)

$$\frac{\text{increased accuracy in prediction}}{\text{resources spent (\$)}}$$

- Kord Smith (ANS M&C 2003)
  - PWRs: 3D reaction rates 3% pin power, axially integrated power: 1%
  - +/-10 ppm critical boron at start-up = “close to perfect”
- Uncertainties in some physics (e.g., XS evaluations, fuel thermo-mechanics, TH, …) may overwhelm the accuracy of a solver for instance.

How about a much quicker turnaround for quality design calculations?

- High-dimensional (high-order/first-principles) models :
  - Generate a wealth of data
  - Require high-end HPC platforms
  - May need to be repeated for every change in the input parameter space.
- Multi-query HPC problems (repeated calculations with changes in the input) can become expensive
  - Design optimization
  - Uncertainty quantification
- **Data Sciences:**

Learn from HPC simulations to predict new cheap HPC-quality simulations

→ Think data assimilation ("image/video compression")

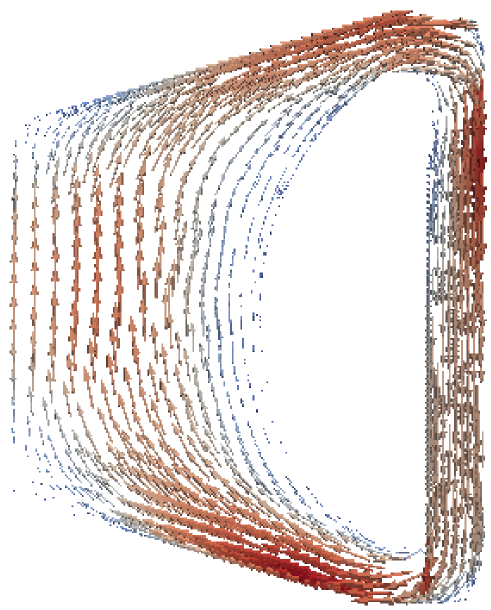

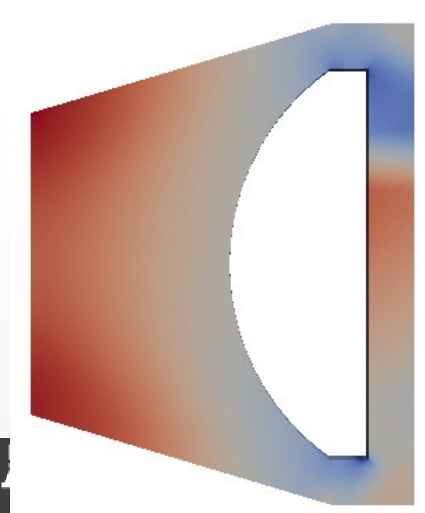

Many simulations with parametric variations

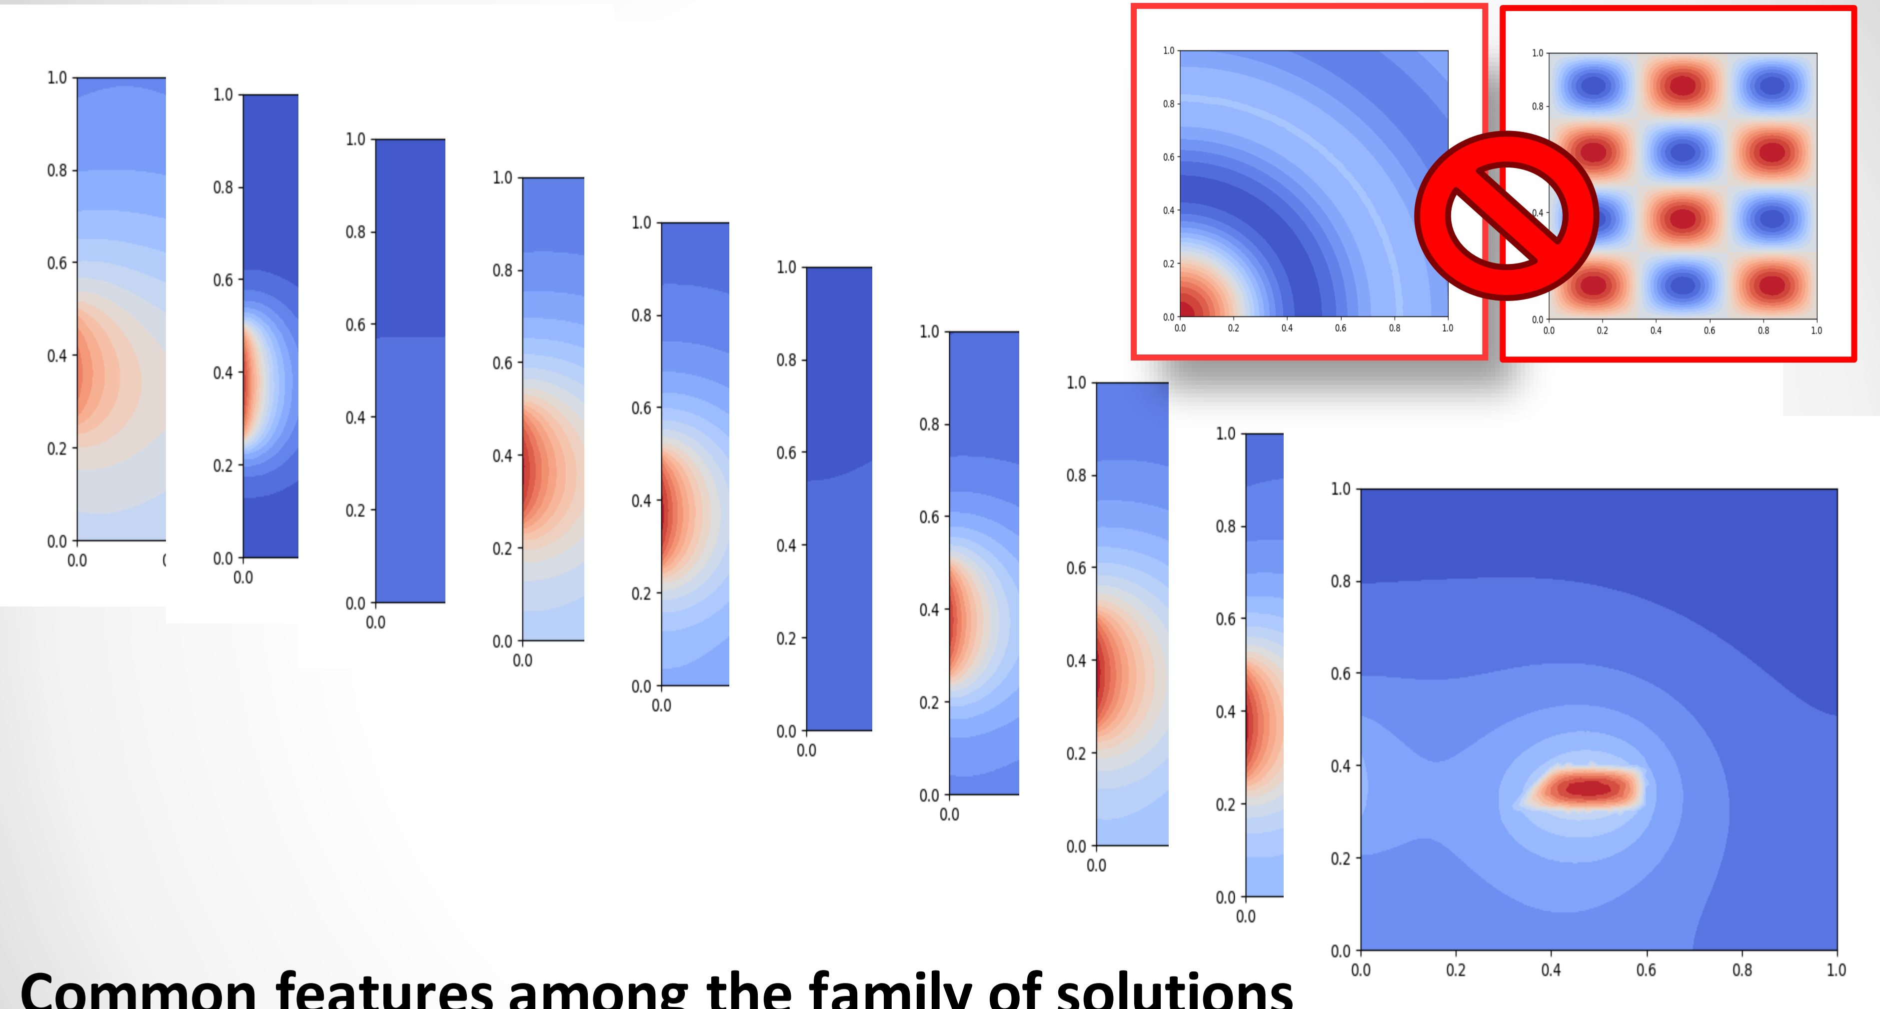

Common features among the family of solutions

Can we learn from that? → data-driven subspace discovery

Discovered subspace from learned data

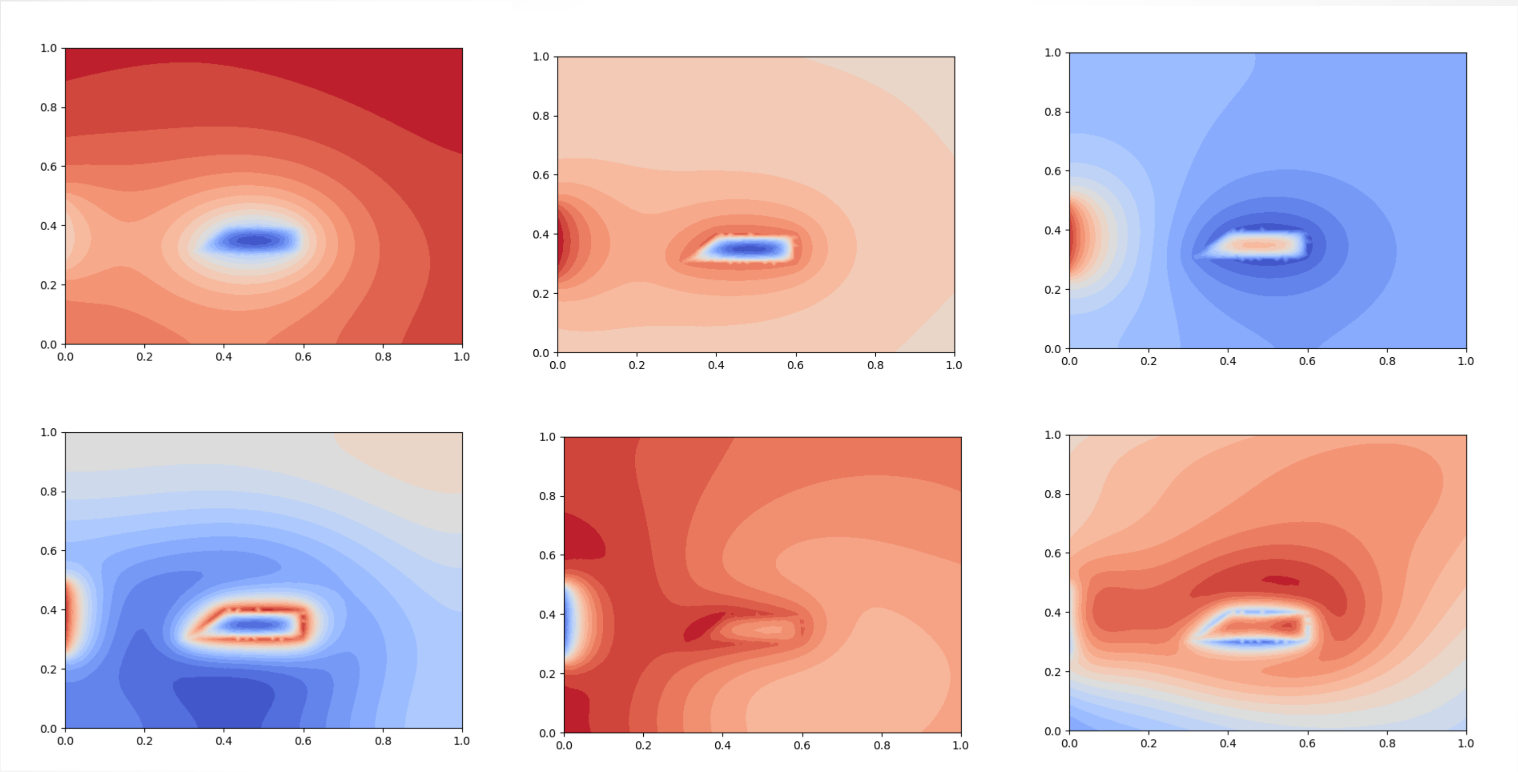

- Obtained via **Singular Value Decomposition** of the snapshots (learned data)
- Reduction comes from the low number of modes needed

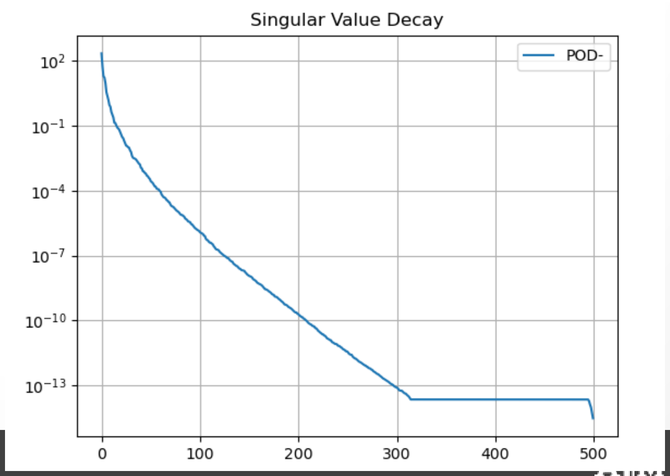

Model Order Reduction: to reduce the computational complexity

Definition

- Model order reduction (MOR) is a set of techniques aimed at reducing the computational complexity of mathematical models in numerical simulations.
- Description of reality (model) + problem input data (in) $\longrightarrow$ PDEs $\longrightarrow$ discretization $\longrightarrow$ large-scale model with a large number of unknowns (degrees of freedom, DoFs) N.

Full Order Model (FOM): Solve $\dot{x} = f(x(t), in(t))$ with $\boxed{x \in \mathbf{R}^N}$

- Model order reduction aims at lowering the computational complexity of such problems by reducing the # of DoFs ($r \ll N$)

Reduced Order Model (ROM): Solve $\dot{c} = f_r(c(t), in(t))$ $\boxed{c \in \mathbf{R}^r \quad \text{with } r \ll N}$

such that

$$\|x - Uc\| \leq C_r \|in\| \quad \text{with} \quad \lim_{r \to N} C_r = 0$$

U: reconstruction operator.

Key points:

- Full Order Model (FOM) $\boxed{x \in \mathbf{R}^N}$
- Reduced Order Model (ROM) $\boxed{c \in \mathbf{R}^r \quad \text{with } r \ll N}$
- Reconstruction: $\boxed{x \approx Uc}$ where U (size $N \times r$) is a data-driven discovered basis.

Need to determine the expansion coefficients c (as functions of the input parameters)

Data-driven sub-space discovery

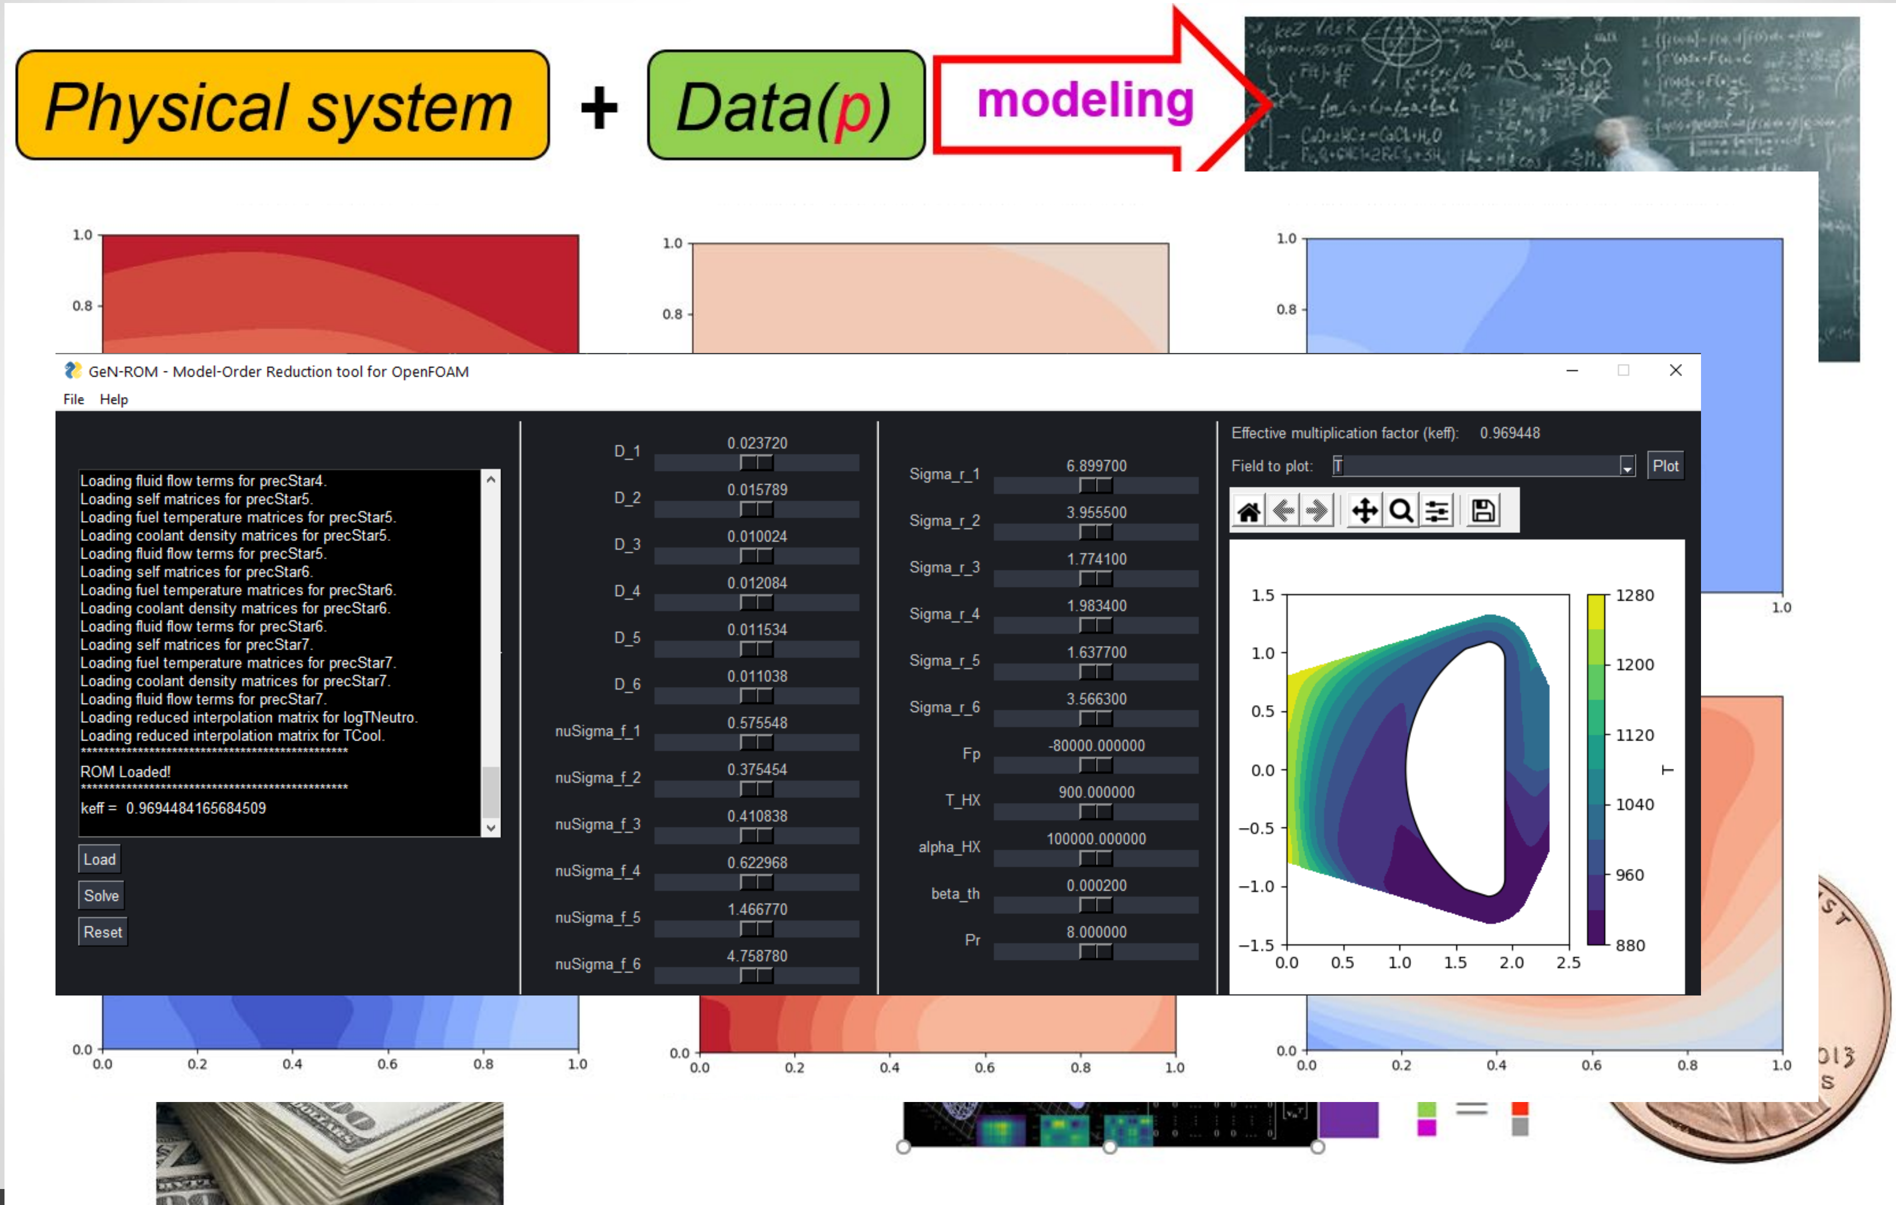

Model-order reduction for advanced reactors

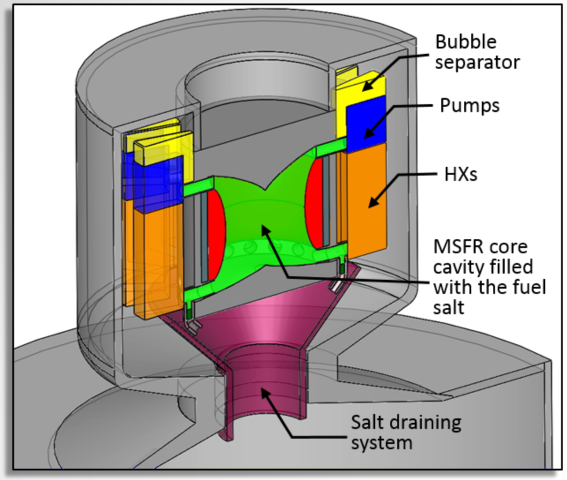

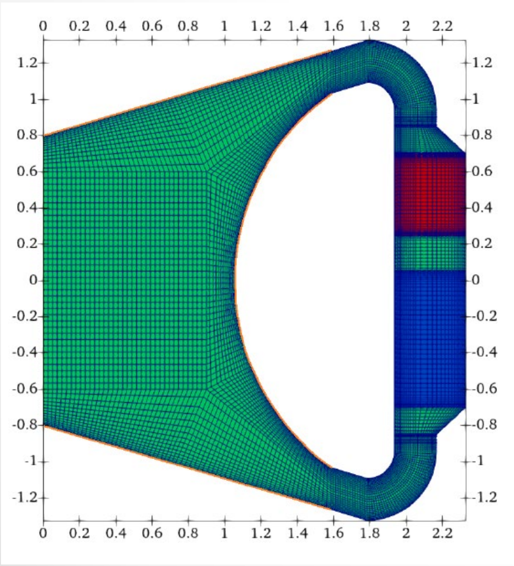

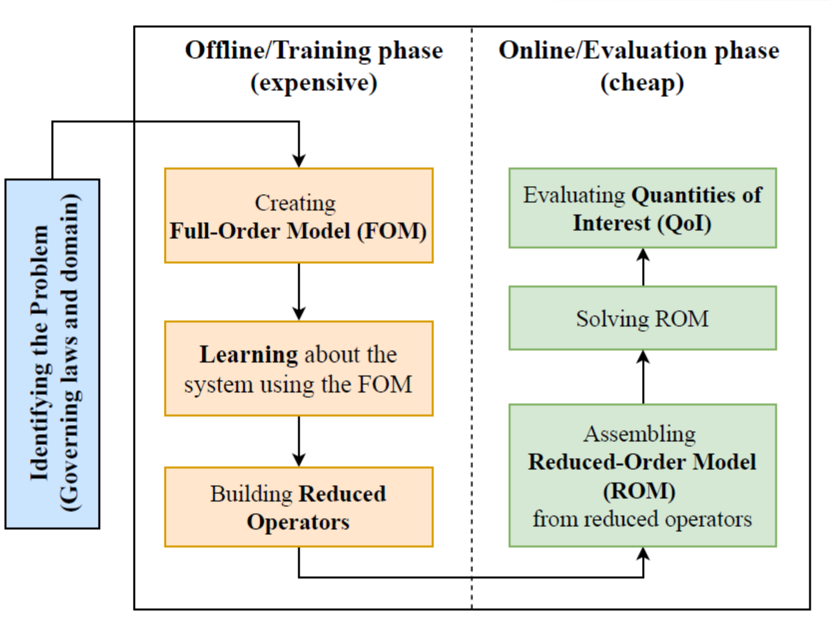

The reduced equation system

$$\rho \boldsymbol{M}\dot{\boldsymbol{c}}^{\boldsymbol{u}_D} + \rho \boldsymbol{c}^{\boldsymbol{u}_D,T}\underline{\underline{\boldsymbol{C}}}\boldsymbol{c}^{\boldsymbol{u}_D} - \boldsymbol{c}^{\eta,T}\underline{\underline{\boldsymbol{T}}}\boldsymbol{c}^{\boldsymbol{u}_D} - \eta\boldsymbol{D}\boldsymbol{c}^{\boldsymbol{u}_D} + \boldsymbol{P}\boldsymbol{c}^{p} + \Gamma(\boldsymbol{B}\boldsymbol{c}^{\boldsymbol{u}_D} - |\boldsymbol{u}_{D,in}|\boldsymbol{S}_r^{BD}) - \sum_{z=1}^{Z}\left(|\boldsymbol{F}_{p,z}|\boldsymbol{S}_{p,z} - \boldsymbol{S}_{fr,z}\boldsymbol{c}_z^{\boldsymbol{F}_{fr}}\right) - \rho\beta_e(\boldsymbol{A}\boldsymbol{c}^{T} - T_{ref}\boldsymbol{S}_T) = 0, \quad (16)$$

$$\rho\boldsymbol{G}\boldsymbol{c}^{\boldsymbol{u}_D} = 0 \quad (17)$$

Some of the reduced operators

$$\boldsymbol{M}_{i,j} = \langle \psi_i^{\boldsymbol{u}_D}, \psi_j^{\boldsymbol{u}_D} \rangle_{\mathcal{D}} \qquad \underline{\underline{\boldsymbol{C}}}_{i,j,k} = \left\langle \psi_j^{\boldsymbol{u}_D}, \frac{1}{\gamma}\nabla\cdot(\psi_i^{\boldsymbol{u}_D}\otimes\psi_k^{\boldsymbol{u}_D}) \right\rangle_{\mathcal{D}}$$

$$\boldsymbol{D}_{i,j} = \langle \psi_i^{\boldsymbol{u}_D}, \nabla\cdot[\nabla\psi_j^{\boldsymbol{u}_D} + (\nabla\psi_j^{\boldsymbol{u}_D})^T] \rangle_{\mathcal{D}} \qquad \boldsymbol{P}_{i,j} = \langle \psi_i^{\boldsymbol{u}_D}, \gamma\nabla\psi_j^{p} \rangle_{\mathcal{D}}$$

$$\boldsymbol{B}_{i,j} = \langle \psi_i^{\boldsymbol{u}_D}, \psi_j^{\boldsymbol{u}_D} \rangle_{\Gamma_{in}} \qquad \boldsymbol{S}_{p,z,i} = \left\langle \psi_i^{\boldsymbol{u}_D}, \gamma\frac{\delta_z(\boldsymbol{r})\boldsymbol{F}_{p,z}}{|\boldsymbol{F}_{p,z}|} \right\rangle_{\mathcal{D}}$$

$$\boldsymbol{S}_{r,i}^{BD} = \left\langle \psi_i^{\boldsymbol{u}_D}, \frac{\boldsymbol{u}_{in}}{|\boldsymbol{u}_{in}|} \right\rangle_{\Gamma_{in}} \qquad \boldsymbol{G}_{i,j} = \langle \psi_i^{p}, \nabla\cdot\psi_j^{\boldsymbol{u}_D} \rangle_{\mathcal{D}}$$

Reconstructed Flux (left) - Reconstruction Error (right)

Uncertain parameters (23 total):

- Diffusion coefficients, fission and removal cross sections ($\pm$ 10% around the nominal values)
- Pumping force, external coolant temperature, Heat transfer coefficient, Pr-number, thermal expansion coefficient

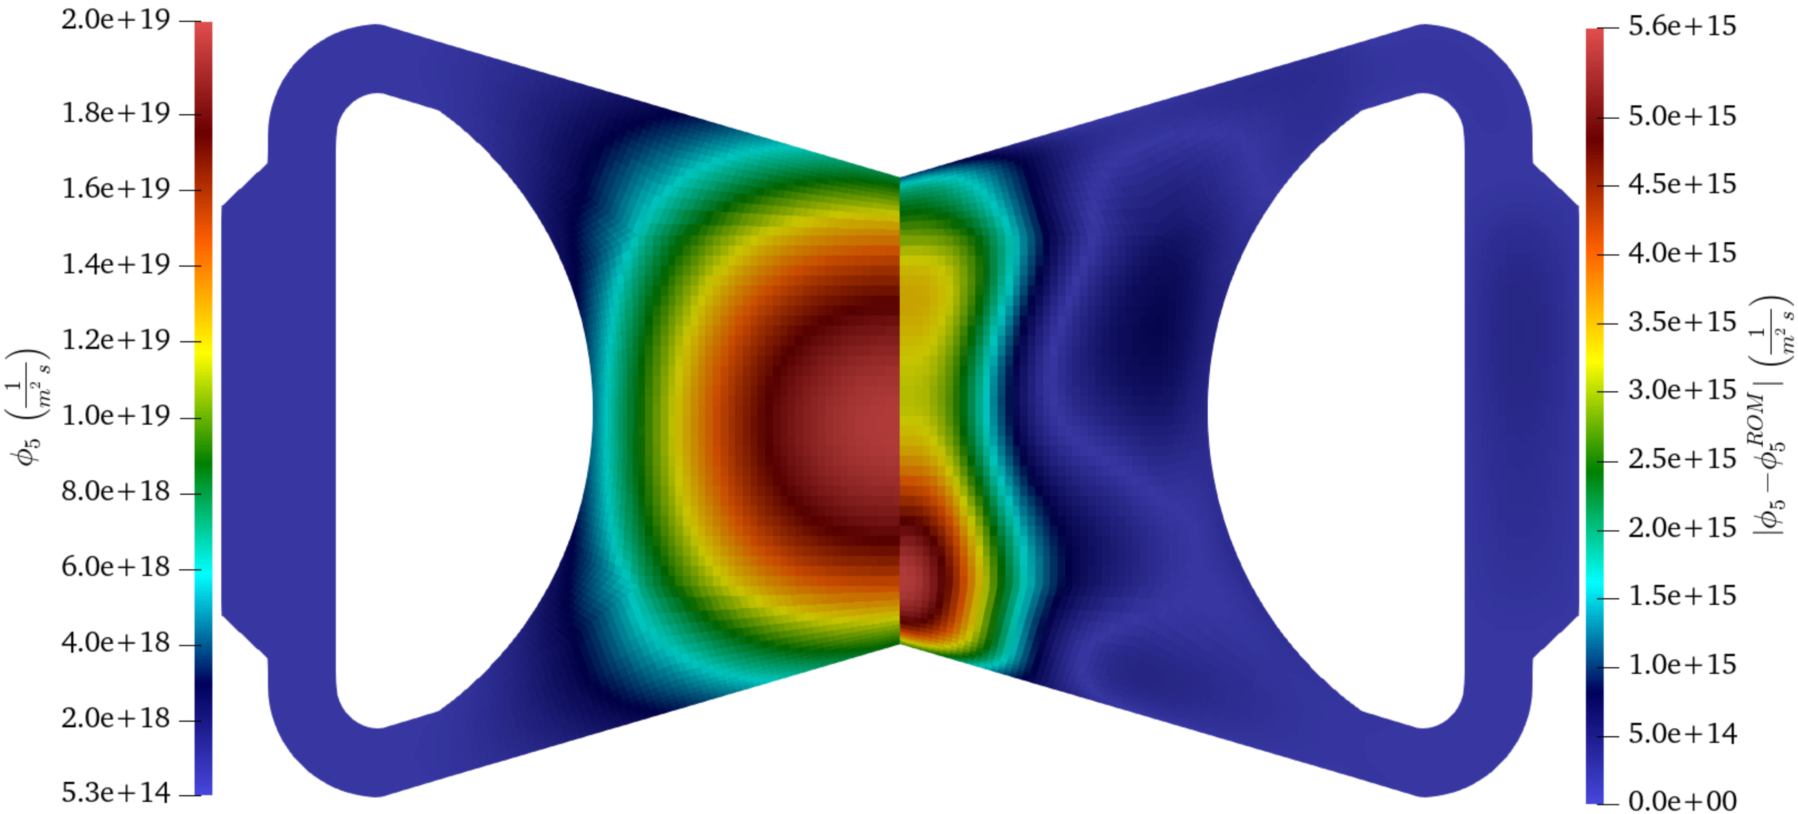

Reconstr. Temperature (left) – Reconstr. Error (right)

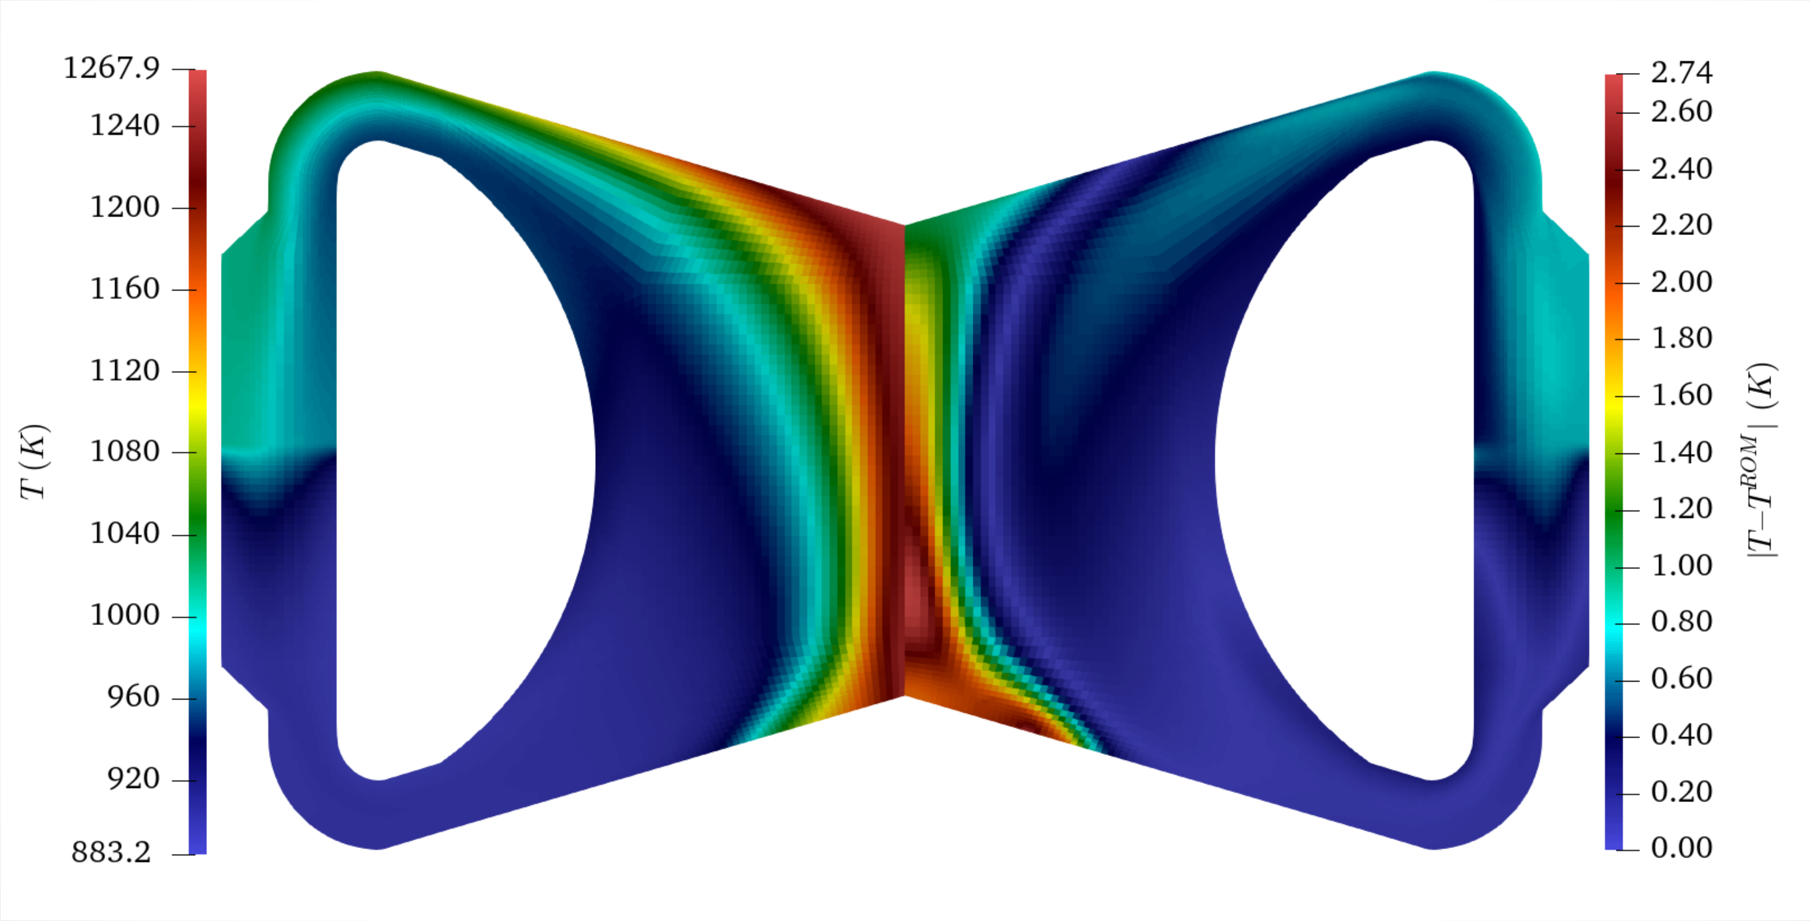

Reconst. Velocity (left) - Reconstruction Error (right)

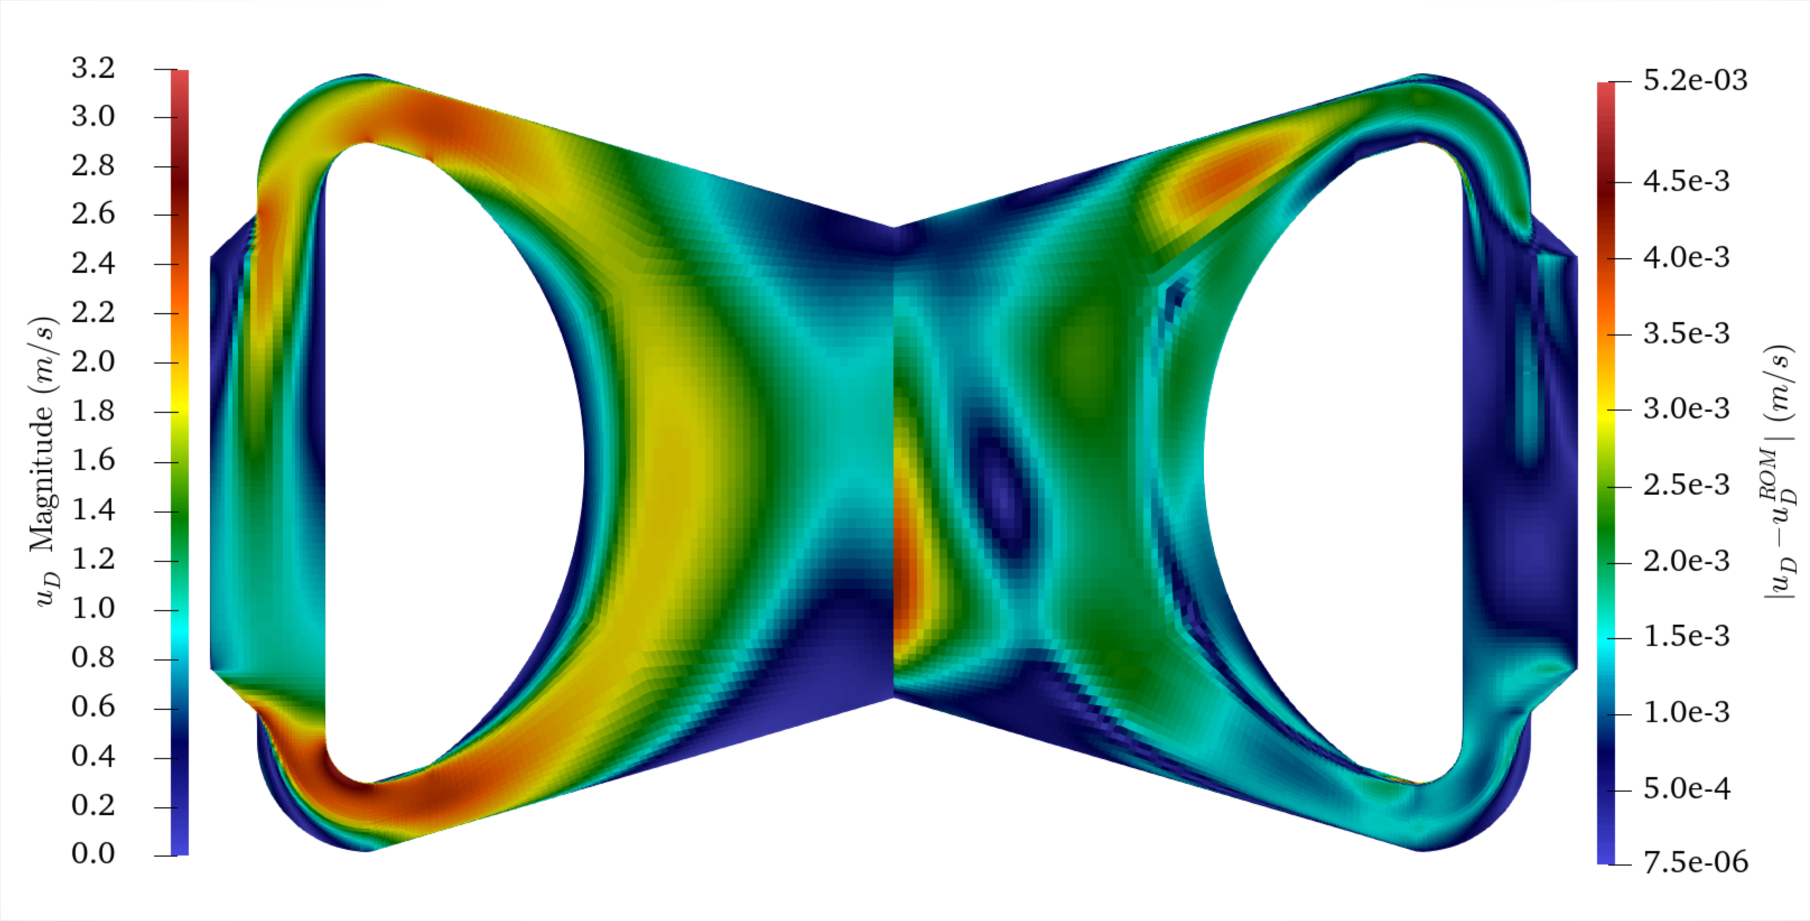

Model-order Reduction: huge speed ups

- Quantities of interest:
  - Effective multiplication factor (k_{eff})
  - Maximum temperature of the system (T_{max})
- Propagation of uncertainties: Monte Carlo approach with 50,000 samples
- Speedup in the UQ including training: approximately factor of 1,500

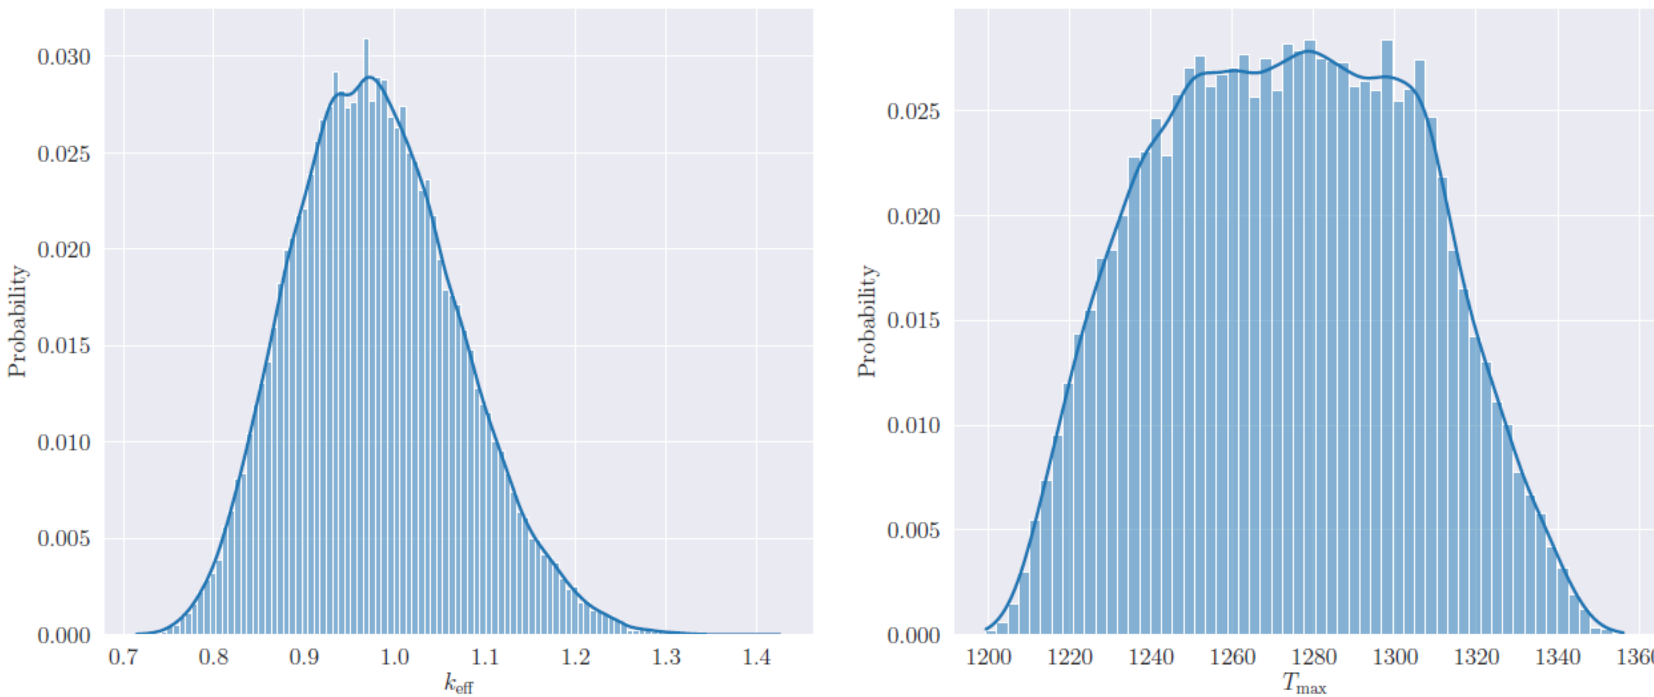

Graphical User Interface Demo

Conclusions and EOF

- Intro to HPC for scientific computing in nuclear engineering and sciences
- Focus SC/HPC efforts where gains are visible/tangible
- Emphasize Predictive Science (VVUQ) and whether the simulation efforts will have an impact?
- Seek certifiable reduced-order models for quick design cycle
    - Borrow from machine-learning and big-data science.