The Future of Scientific Computing

Bruce Hendrickson
Associate Director, Computation

SIAM Annual Meeting
July 13, 2018

LLNL-PRES-738377
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Lawrence Livermore National Laboratory

Perhaps trying to predict the future is a bad idea

- *Prediction is very difficult, especially if it's about the future.*
  - Nils Bohr, Nobel laureate in Physics
- *Those who have knowledge, don't predict. Those who predict, don't have knowledge.*
  - Lao Tzu, 6th Century BC Chinese Poet
- And yet, we successfully predict ...
  - ... the weather days in advance
  - ... the climate decades in advance
  - ... eclipses centuries in advance
- Applied mathematicians know a lot about prediction.
  - We call it time integration!

What can we learn about prediction from applied mathematicians?

- Key insight: Identify conservation laws or invariants that constrain possible futures
- Methodology:
  - Observe the past to learn invariants
  - Then extrapolate them into the future
  - Sprinkle in any foreseeable disruptions
- Goal: A principled framework for thinking about the future
- *A good hockey player plays where the puck is. A great hockey player plays where the puck is going to be.*
  - Wayne Gretzky
- A great researcher plays where the *field* is going to be.

Invariants for scientific computing

1. Algorithms & models get more predictive, efficient & robust
   - Precise path is uncertain, but this is at the core of our discipline
   - Researchers in this room have played and will play a major role

- This is merely the scientific method. All scientific communities stand on the shoulders of the giants who preceded them.
- But not every community gets the added boost of Moore's Law
  - *Our giants are riding an up escalator!*

2. Computers get continually faster and bigger
   - Path can be disruptive, e.g. vector-to-parallel transition
   - Dennard scaling is over and Moore's Law is ending, but there is a clear path forward for at least the next decade

Consequences of continually faster computers

- Independent from unforeseeable research advances, we can continually do more and better simulations.
  - Paradoxically, rapid change may make prediction easier!
- *One decade's frontier research questions become the building blocks for the subsequent decade's more ambitious goals*
- Approach:
  - Propose a model
  - Validate model against data
  - Use validated model to predict

Prehistory of scientific computing

- First SIAM Conference on Computational Science and Engineering was held in 2000
- SIAM Activity Group on Computational Science & Engineering was created later that year
- But clearly the field is much older than this
  - Underlying PDEs formulated in the late 1800s and early 1900s
  - Early computers and numerical computations in 1940s and 1950s
  - Foundations of numerical methods and formulations established in 1950's to 1970's

1980s: A focus on linear algebra

- Research foci:
  - First edition of Golub/Van Loan (1983). Strong focus on matrix algorithms.
  - Linear solvers were very active area of research
    - Sparse direct methods were mature. Iterative methods were maturing. Multigrid emerging.
  - Significant developments in numerical discretizations
- Achievable simulation scope:
  - 1D and 2D problems, simple geometries, usually single physics
- Computing paradigm: vector supercomputers
- Tools for understanding simulations: Simple graphs & plots

1990s: the necessity of parallelism

- Research Foci:
  - Parallel algorithms
    - *We expected parallelism to generate lots of new algorithms, but parallelizing the old ones turned out to be hard enough.* Horst Simon (paraphrased)
  - Iterative and multigrid solvers – parallelizability & linear scaling essential
  - Efficient and accurate forward simulation
- Achievable simulation scope:
  - 2D and emerging 3D, more complex geometries, mostly single physics
  - Fairly highly resolved
- Computing paradigm: commodity cluster parallelism
- Tools for understanding simulations:
  - Rich 2D and 3D visualization frameworks emerging

2000s: The rise of uncertainty quantification

- Research Foci:
  - Forward simulations mature enough to be used as building blocks
  - Early "Outer loop" challenges: optimization, uncertainty quantification, etc.
  - Emerging focus on software engineering
- Achievable Simulation Scope:
  - Highly resolved, 3D with complex geometries, multi-physics
    - But closure models still essential for many phenomena
- Computing paradigm: massive, homogeneous parallelism
- Tools for understanding simulations:
  - Sophisticated spatial & temporal visualization environments

2010s: More complex landscapes

- Research Foci:
  - Complex "outer loop", e.g. model calibration, phase space analysis, inverse problems
  - Ensembles of simulations
  - Emerging roles for data science and machine learning
- Achievable Simulation Scope:
  - Optimal design, sensitivity & uncertainty analysis
  - Emerging predictive simulation
- Computing paradigm: heterogeneous parallelism
- Tools for understanding simulations:
  - Higher-level analysis tools, e.g. topological, or ensemble-focused

Exemplar 1: Topological abstractions of turbulent combustion simulations provide new insights

- Visualization of raw combustion data

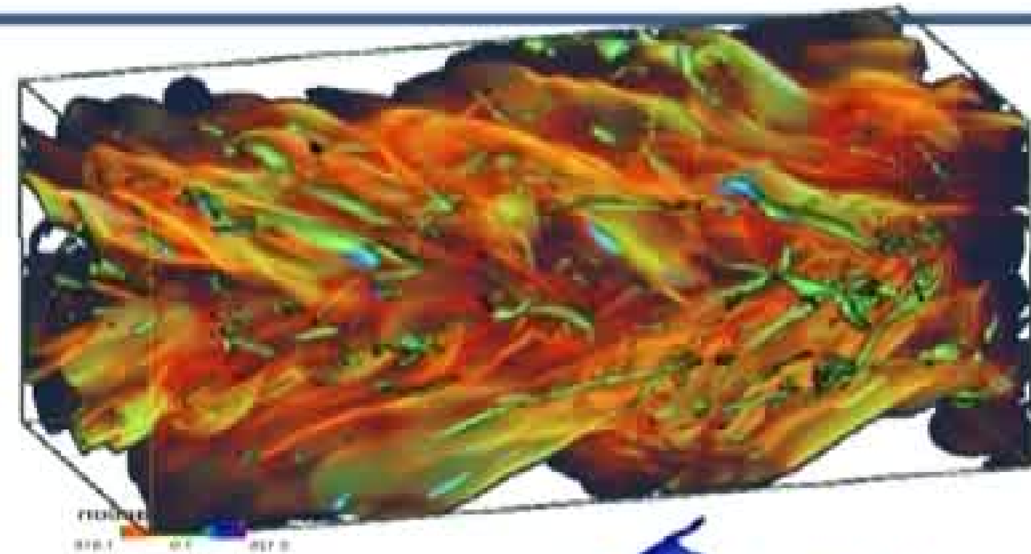

- Segmentation of extinction and re-ignition regions

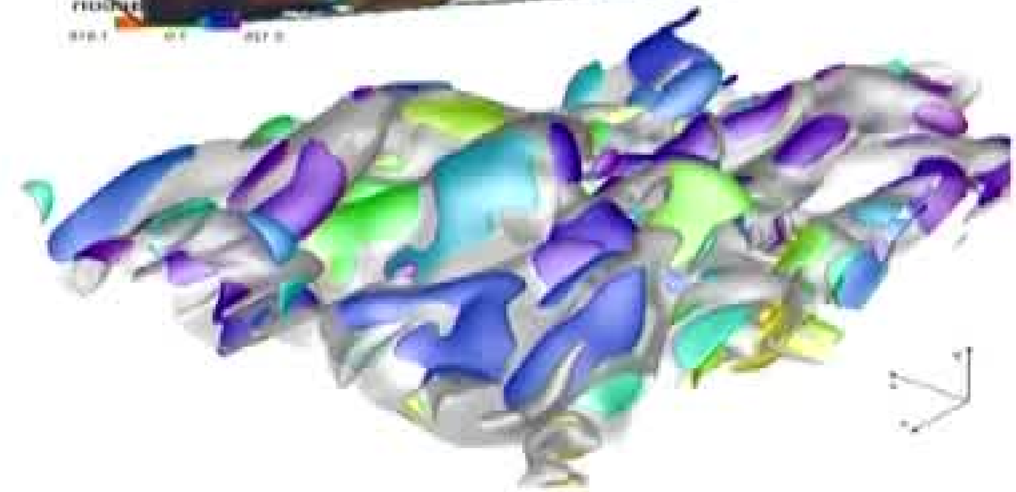

- Evolution of the features with temporal events: birth, split, merge, and death

Courtesy of Jackie Chen (SNL) & Timo Bremer (LLNL)

Exemplar 2: Tools for ensemble analysis enable deeper understanding

Sandia National Laboratories

Example HPC Ensemble Problem

- Tens-to-hundreds of thousands of HPC executions, each with
  - 10 input parameters,
  - 5 scalar outputs,
  - 8 outputs over time (time series),
  - 6 images, 5 movies (30s each)
- How do we efficiently perform parameter studies, sensitivity analysis, validation & verification on petabytes of data?

Slycat™ Provides

- Insights into previously unsuspected behaviors in simulation models (relationships, anomalies).
- Multiple models with varying perspectives on the data.
- On-demand remote exploration of terabytes of results without moving the data (reducing time/storage costs).
- Many-to-many correlations for sensitivity analysis.
- Ubiquitous web-based delivery for easy collaboration.

Courtesy of Pat Crossno (SNL)

Canonical Correlation Analysis (CCA)
- Correlations between two sets of variables
- Sensitivity analysis, anomaly detection

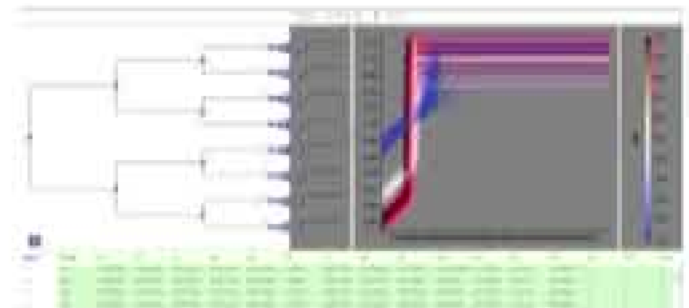

Time Series Clustering
- Time series similarity, shape filtering
- Map output variability to inputs, find outliers

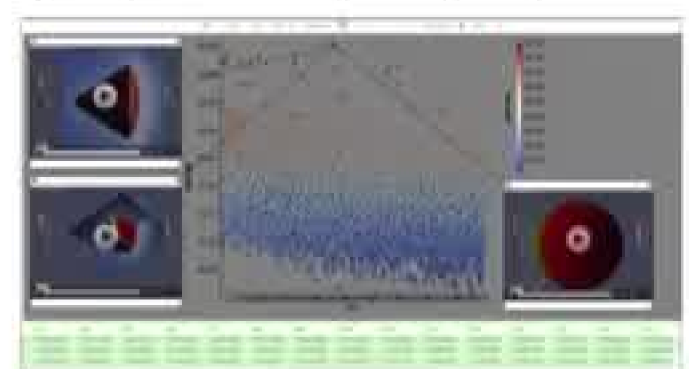

Parameter Space Exploration
- Visual exploration, filtering, image/video retrieval
- Parameter studies, multi-objective optimization

Advances outside our community can create disruption and opportunities

- Demise of Moore's Law will drive the need for new thinking on computer architectures
  - This will almost certainly impact algorithms and software engineering

Advances outside our community can create disruption and opportunities

- Demise of Moore's Law will drive the need for new thinking on computer architectures
  - This will almost certainly impact algorithms and software engineering
- Emergence of data science and machine learning tools create additional opportunities to make sense of our simulations
  - But this will require substantial changes to our workflows, and will require us to embrace a broader definition of scientific computing
- Exemplar 3: Machine learning from simulations provides new insight into NIF capsule design
  - *Following slides courtesy of Luc Peterson and Brian Spears (LLNL)*

To compress the capsule and fuel, indirect-drive ICF on NIF fires lasers into a hohlraum to form an x-ray drive

Hohlraum is filled with helium gas to permit cryogenic layering and to control interior wall expansion as lasers deposit energy

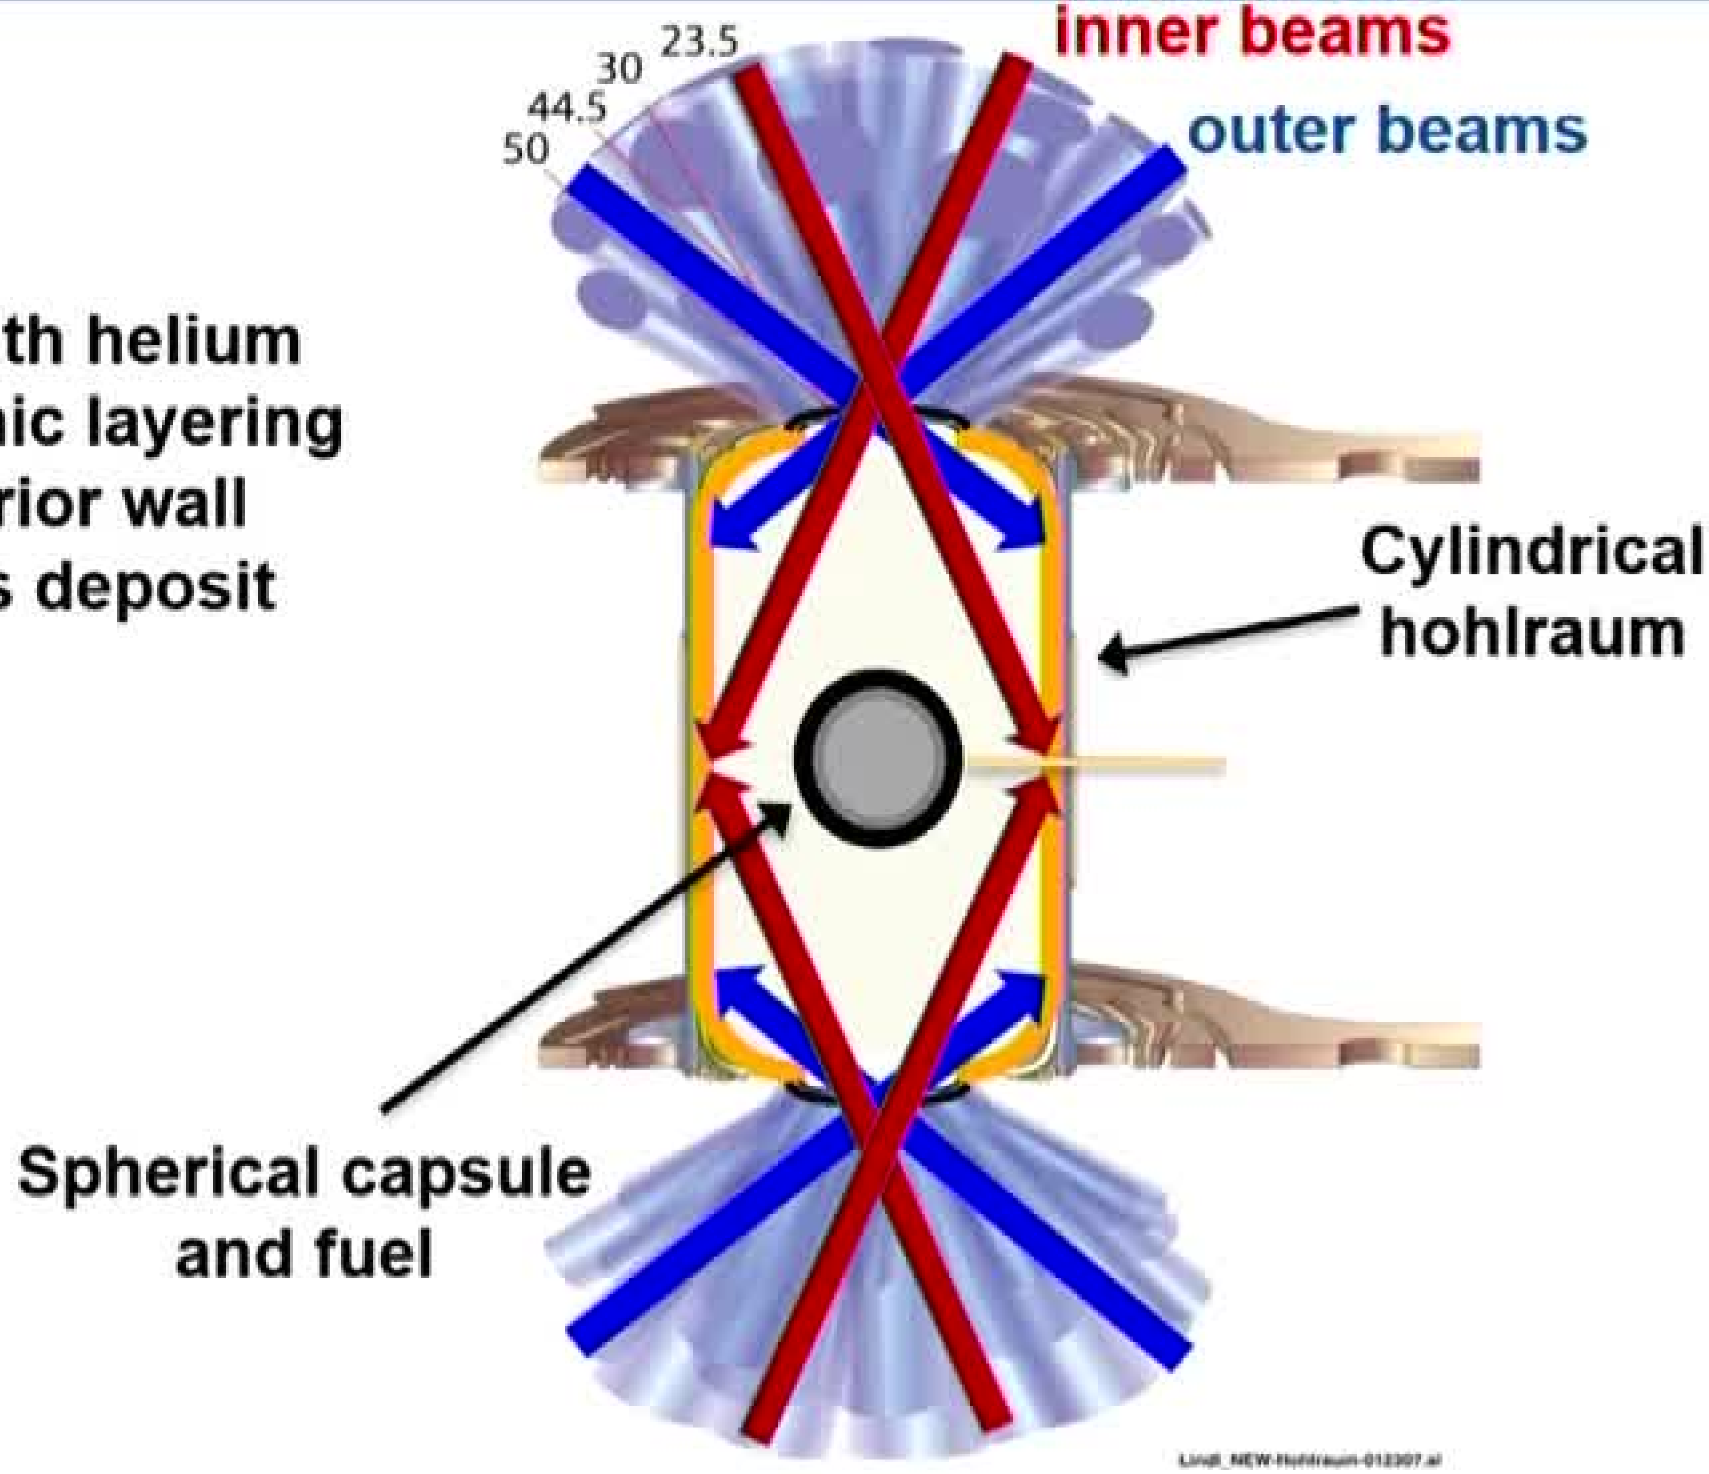

Slide: L. Berzak Hopkins

Round, symmetric implosion → critical challenge for ICF

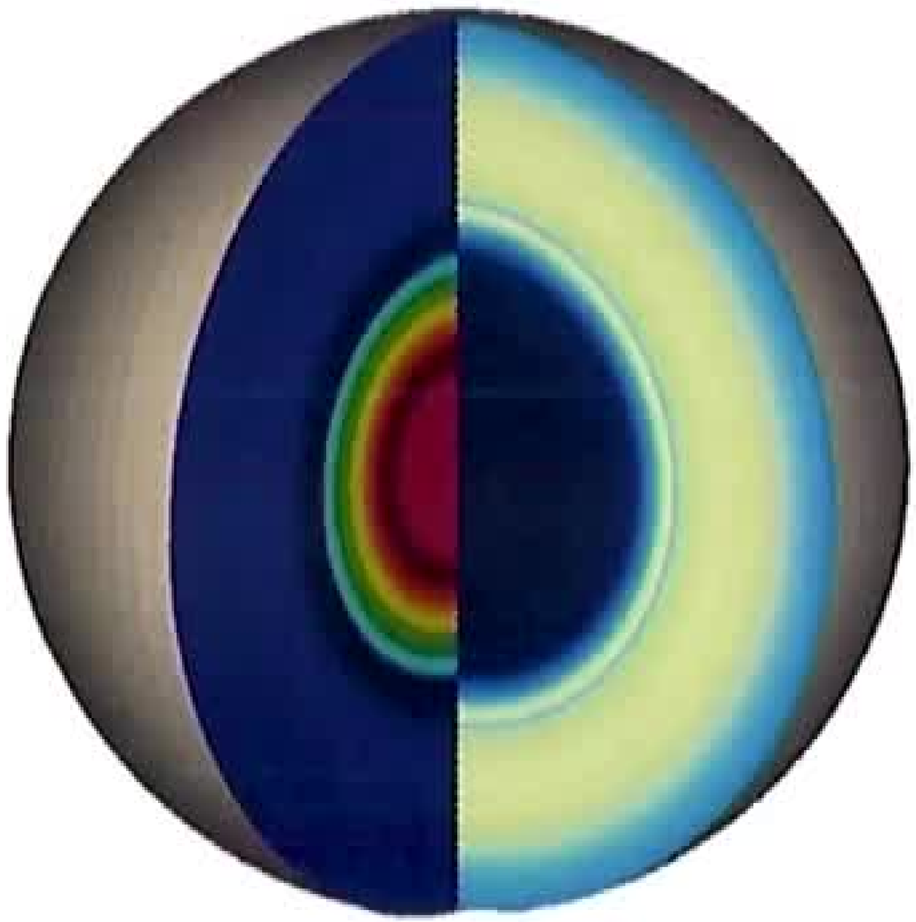

Highly efficient, highly symmetric simulated implosion

1D
500 zones
1 CPU
5 minutes runtime

Slide: L. Berzak Hopkins

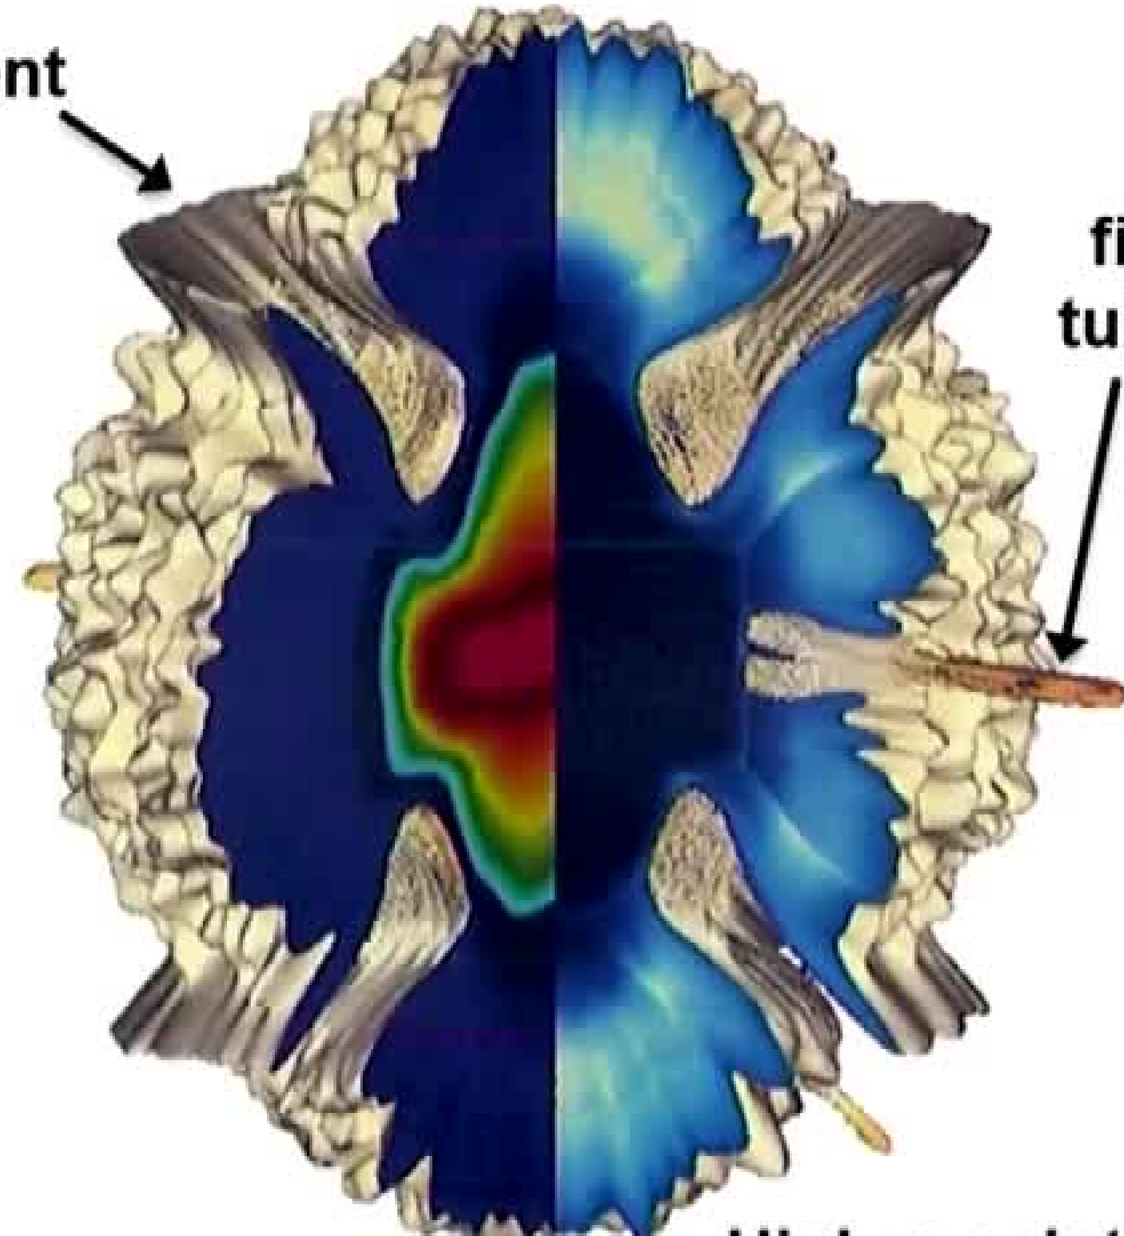

High-resolution postshot simulation of NIF experiment

3D full-res.
400,000,000 zones
6144 CPUs
1 month runtime

Idea: create and analyze a large database of ICF simulation data

- Varied 9 parameters
  - Time-varying drive asymmetry
    - Constant Legendre Modes P1, P4
    - Time-varying P2
  - Time-varying drive amplitude
  - Capsule gas fill density

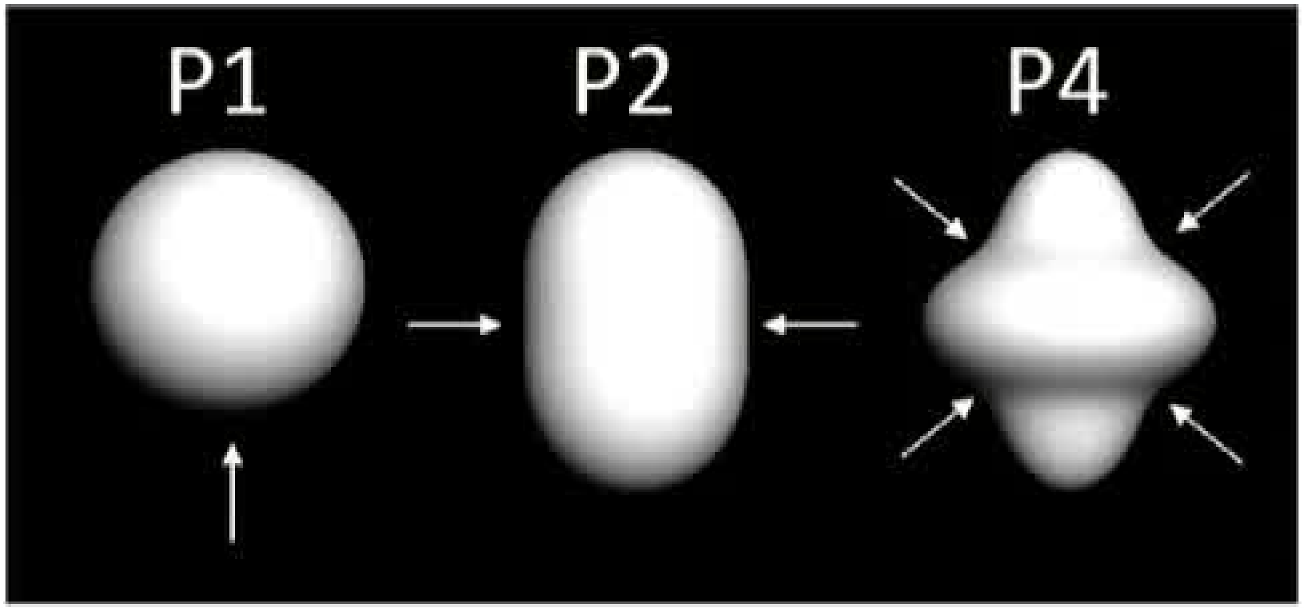

- Successfully completed 60k 2D simulations
- 39 Million CPU Hours on Trinity at LANL
- 5 PB Raw Data, 100 TB Processed & Zipped
- First chunk of data took ~2 months to get from LANL to LLNL

Optimization of surrogates found a robust, ovoid implosion

Key Features of optimal:

Time-Varying P2 Drive

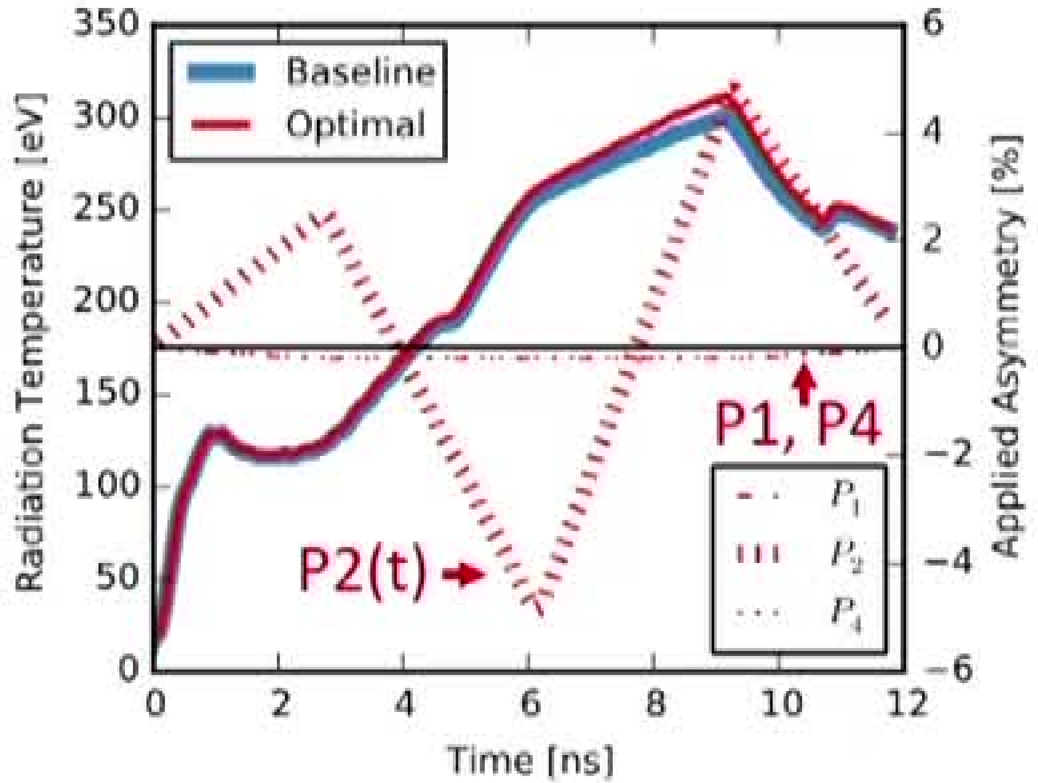

Oval Fuel Shape at Stagnation

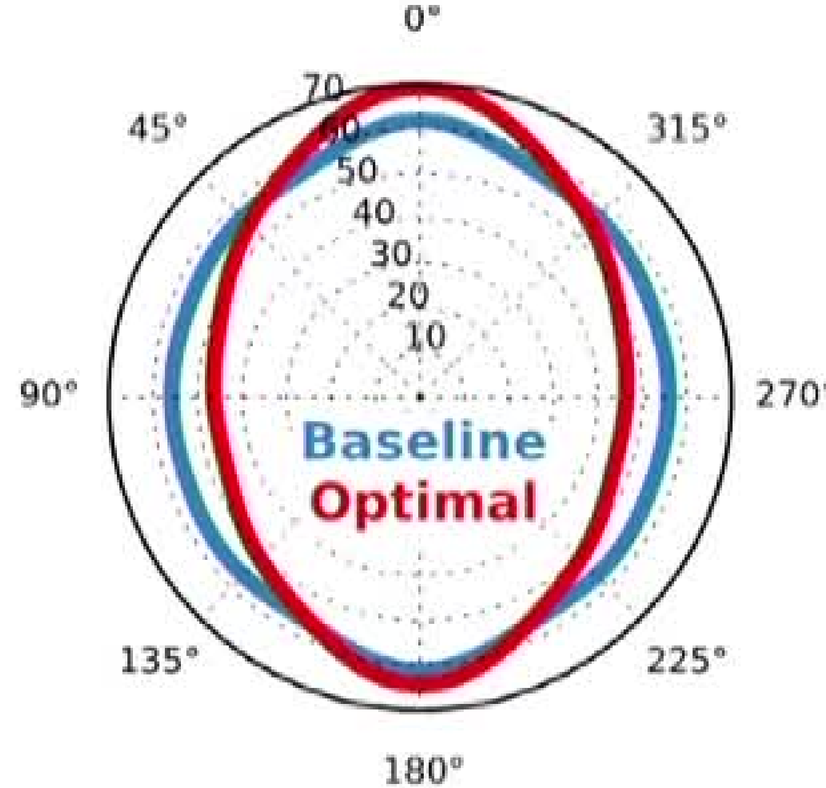

Surrogates: The ovoids are a family of robust implosions

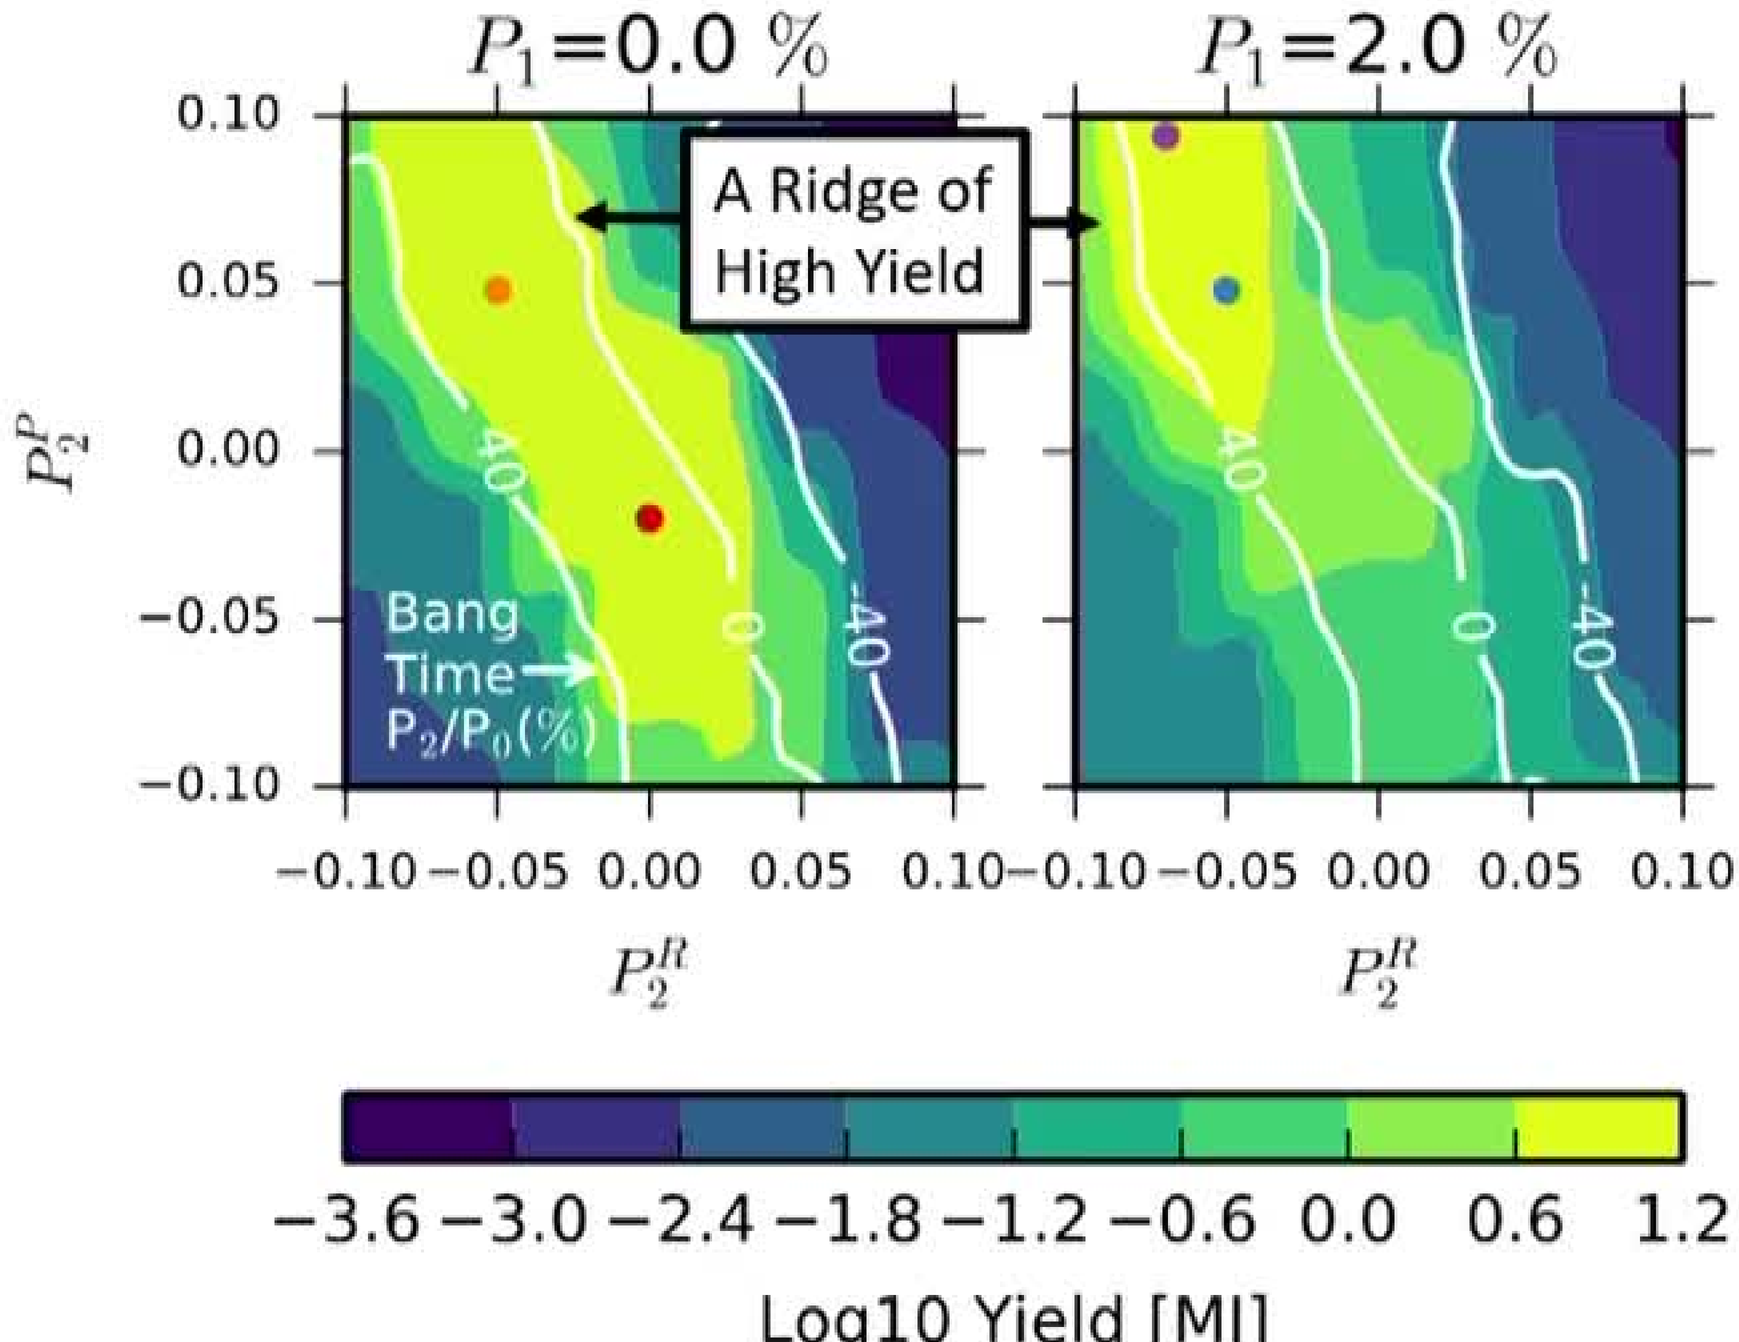

New HYDRA Simulations Confirm Surrogate's Prediction: Optimum is a high-yield ovoid (16.6 MJ)

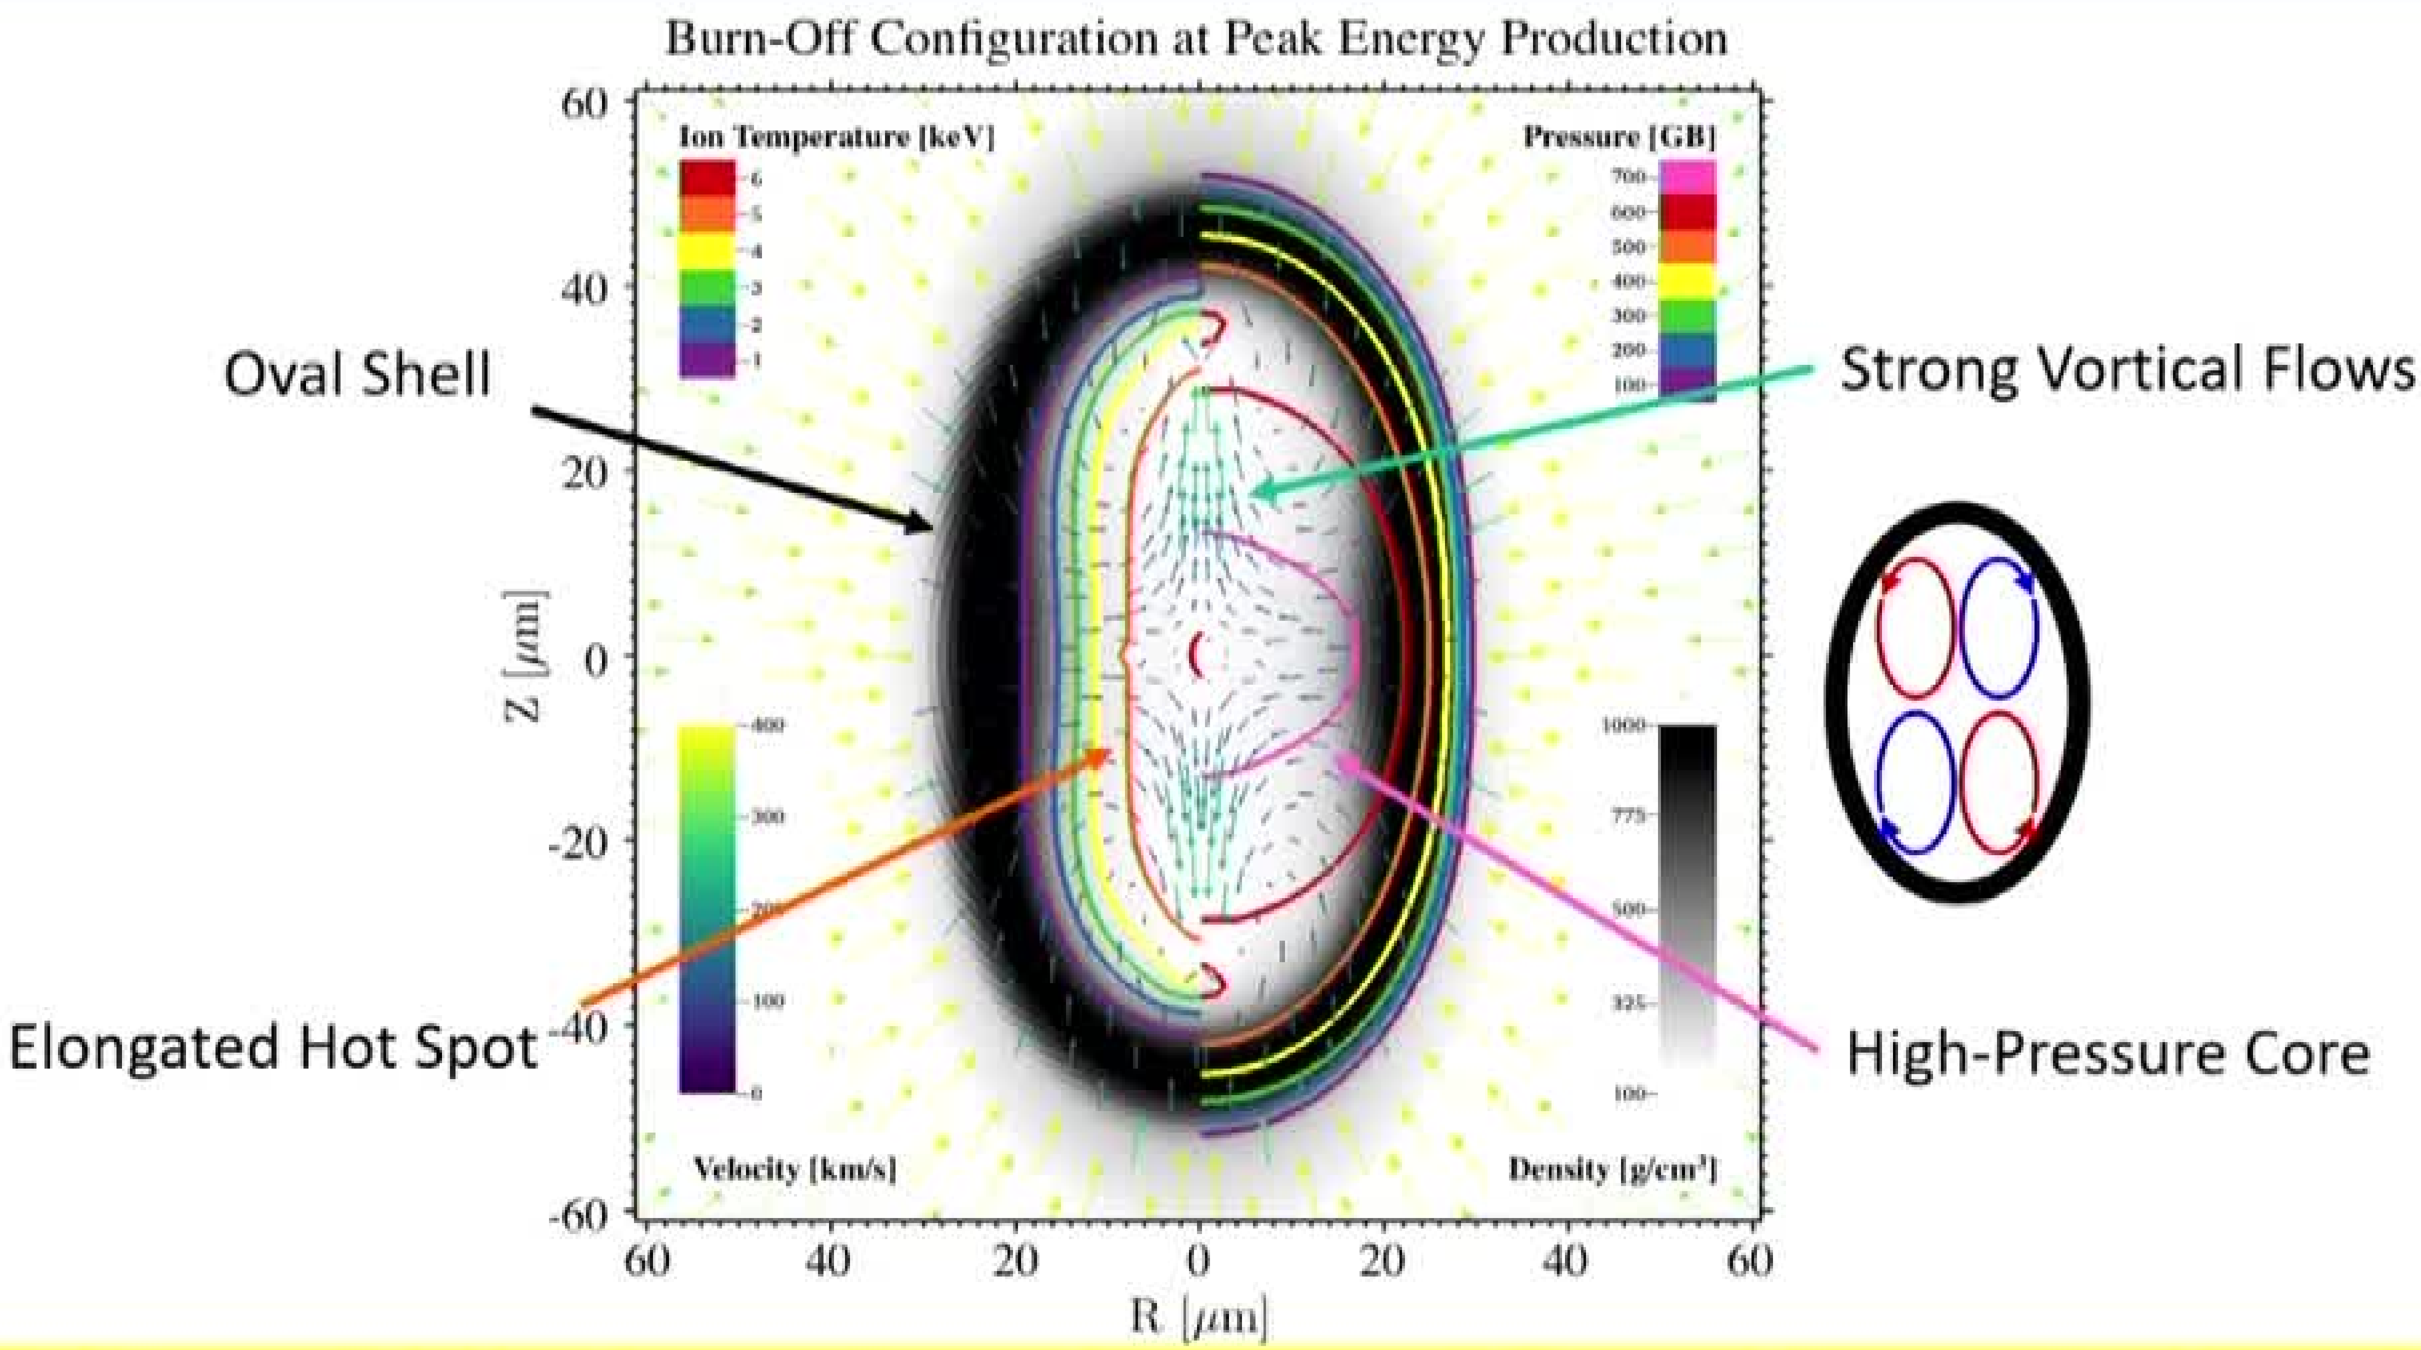

In the ovoid, fuel flows in on the equator, burns, leaves via the poles

In planetary atmospheres and magnetic fusion, these flows are 'zonal flows'

Flows in Jupiter's Core

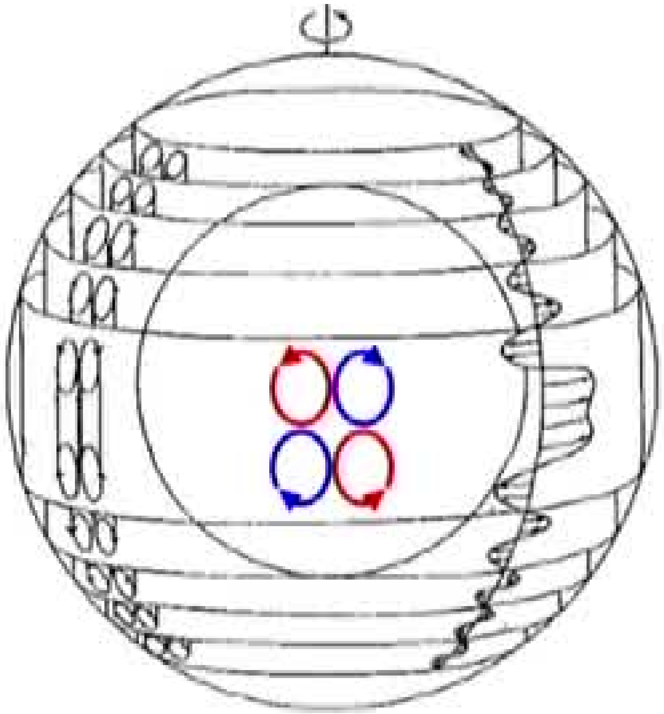

F. H. Busse, Chaos **4**, 123 (1994)

Tokamak Turbulence Suppression

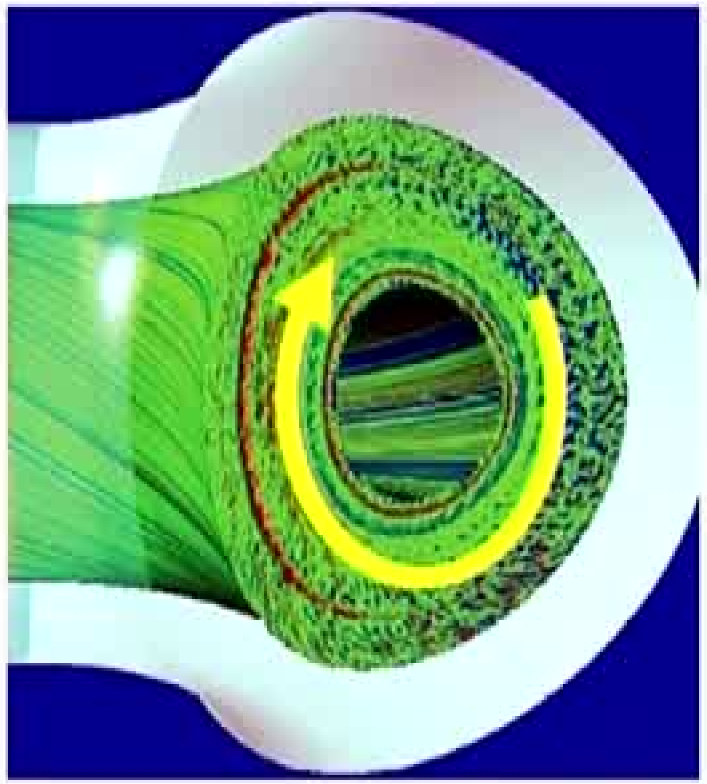

Candy et al GYRO, General Atomics

- Arise in inverse-cascade systems
- Energy can move from small to large scales
- Zonal flows can shear away smaller fluctuations, reduce radial convective transport

Zonal flows can suppress turbulent transport by feeding on and shearing away smaller scales

New science can be discovered by applying machine learning to a large ensemble of simulations

Prediction I: Today's research capabilities as building blocks for more ambitious goals

1. We will be able to generate ensembles of high-quality simulations
   - New abstractions and machine learning tools will be essential to making sense of this data
   - New workflows coupling different styles of computation
   - Numerous applications including
     - Better exploration of parameter space for margin analysis and design optimization – applications in material science, drug design, etc.
     - Automated design from functional specifications
     - Emerging applications involving stochastic or complex systems, e.g. biology

- Numerous research directions
  - PDE-constrained machine learning, uncertainty & robustness in machine learning, new UQ & optimization methods, data+simulation workflows, etc.

Prediction II: Same things faster

2. We will be able to do existing simulations much faster and more robustly
   - This will allow for simulations in real-, or faster than real-time
   - Add observation data to constrain simulation
   - Numerous applications including
     - Computational steering of physical experiments
       - Vary parameters for better data
       - Terminate early if going awry or data sufficient – allow for better facility utilization
     - Better control of dynamical systems, e.g. power-grid resilience

- Numerous research directions
  - New workflows, data-constrained simulations, new opportunities in experimental design, etc.

Prediction III: Outside advances create new opportunities – foreseeable disruptions

3. Machine learning will become a tool *within* our simulations
   - Use machine learning to improve our semi-empirical models, e.g.
     - Force fields in molecular dynamics
     - Parameterized material or constitutive models, etc.
   - Decision making for poly-algorithms
     - E.g. adaptive meshing, mesh untangling, nonlinear solver parameters, etc.

- Numerous research directions
  - Machine learning algorithms for scientific problems, data-driven modeling, etc.

Caveats

- Not everything is predictable
  - Conservation laws are necessary but not sufficient for prediction
  - The evolution of science has ill-conditioned and chaotic elements
- There remain opportunities for game-changing research ideas
  - E.g. processing in memory, randomized algorithms, quantum simulation…
- It's important to pay attention to other potential external disruptors
  - Post-Moore's Law computing paths & algorithmic implications
  - Ongoing revolution in machine learning
  - Others?