Salishan Conference on High-Speed Computing

Computational Challenges at the Petascale and Beyond for Fusion Energy Sciences

W. M. TANG
Princeton University, Plasma Physics Laboratory,
Princeton, New Jersey

Salishan Lodge
Gleneden Beach, Oregon
28 April 2009

Fusion Energy: *Burning plasmas are self-heated and self-organized systems*

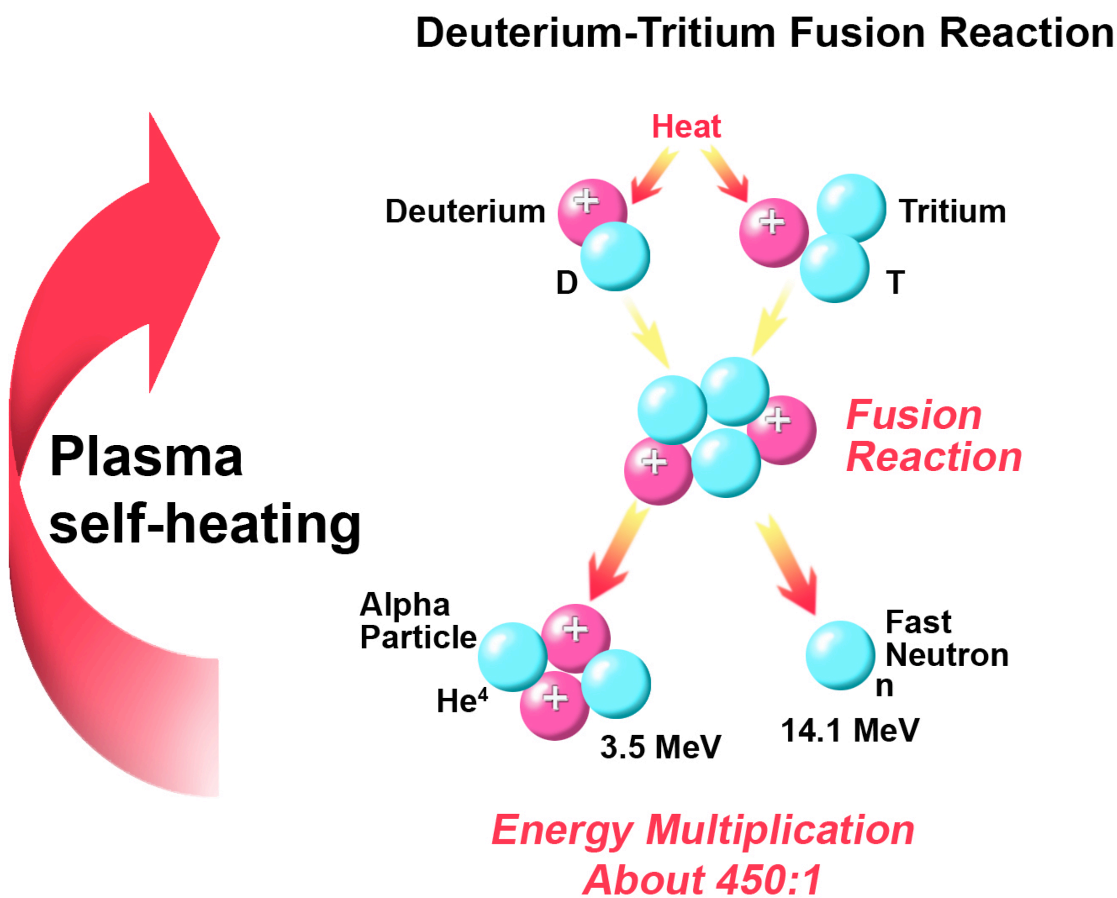

$$D^{+} + T^{+} \rightarrow {}^{4}He^{++}\ (3.5\ MeV) + n^{0}\ (14.1\ MeV)$$

DOE Office of Science Facilities Plan

"*ITER* is an international collaboration to build the ***first fusion science experiment capable of producing a self-sustaining fusion reaction, called a 'burning plasma.'***

It is the ***next essential and critical step on the path toward demonstrating the scientific and technological feasibility of fusion energy***."

Progress in Magnetic Fusion Research

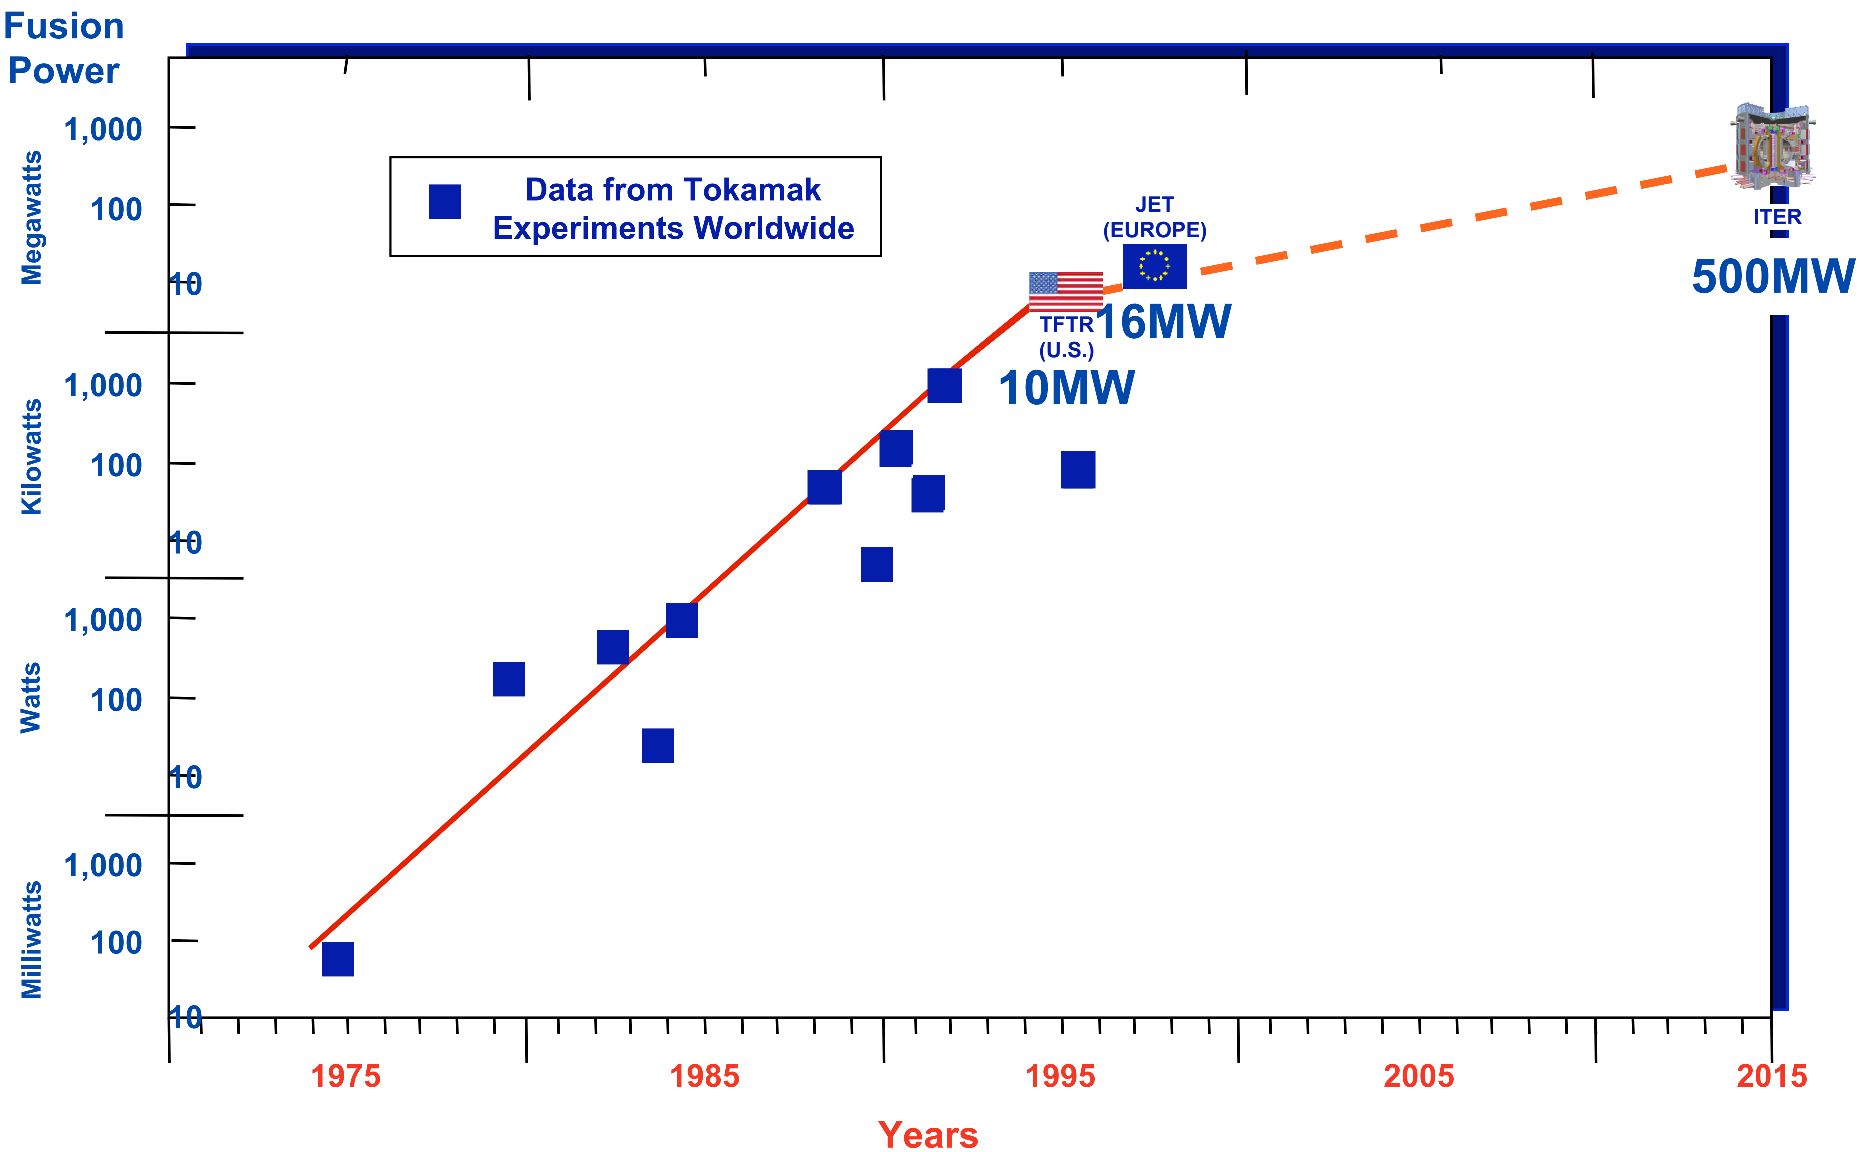

ITER Goal: *Demonstration of the Scientific and Technological Feasibility of Fusion Power*

- ITER *is a dramatic next-step for Fusion*:

-- Today: 10 MW(th) for 1 second with gain ~1
-- ITER: 500 MW(th) for >400 seconds with gain >10

- **Many of the technologies used in ITER will be the same as those required in a power plant *but additional R&D will be needed***

-- "DEMO": 2500 MW(th) continuous with gain >25, in a device of similar size and field as ITER
* Higher power density
* Efficient continuous operation

- **Strong R&D programs are required to support ITER and leverage its results.**

-- Experiments, theory, computation, and technology that support, supplement and benefit from ITER

ITER

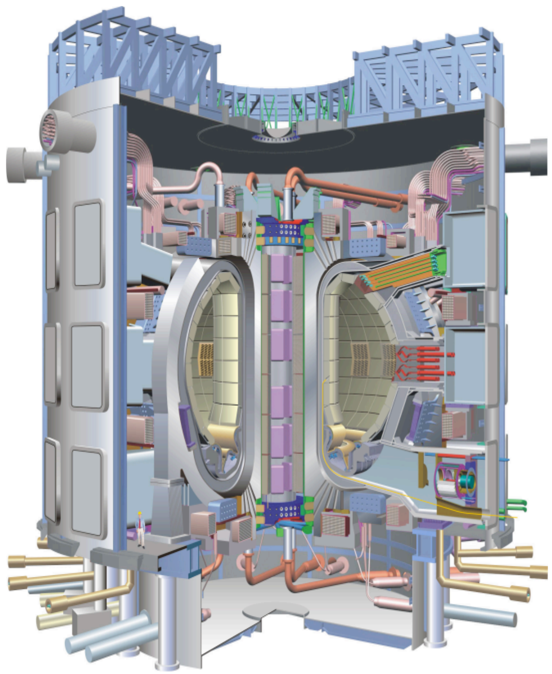

Magnetically confined plasmas in a tokamak are complex and demand integrated analysis

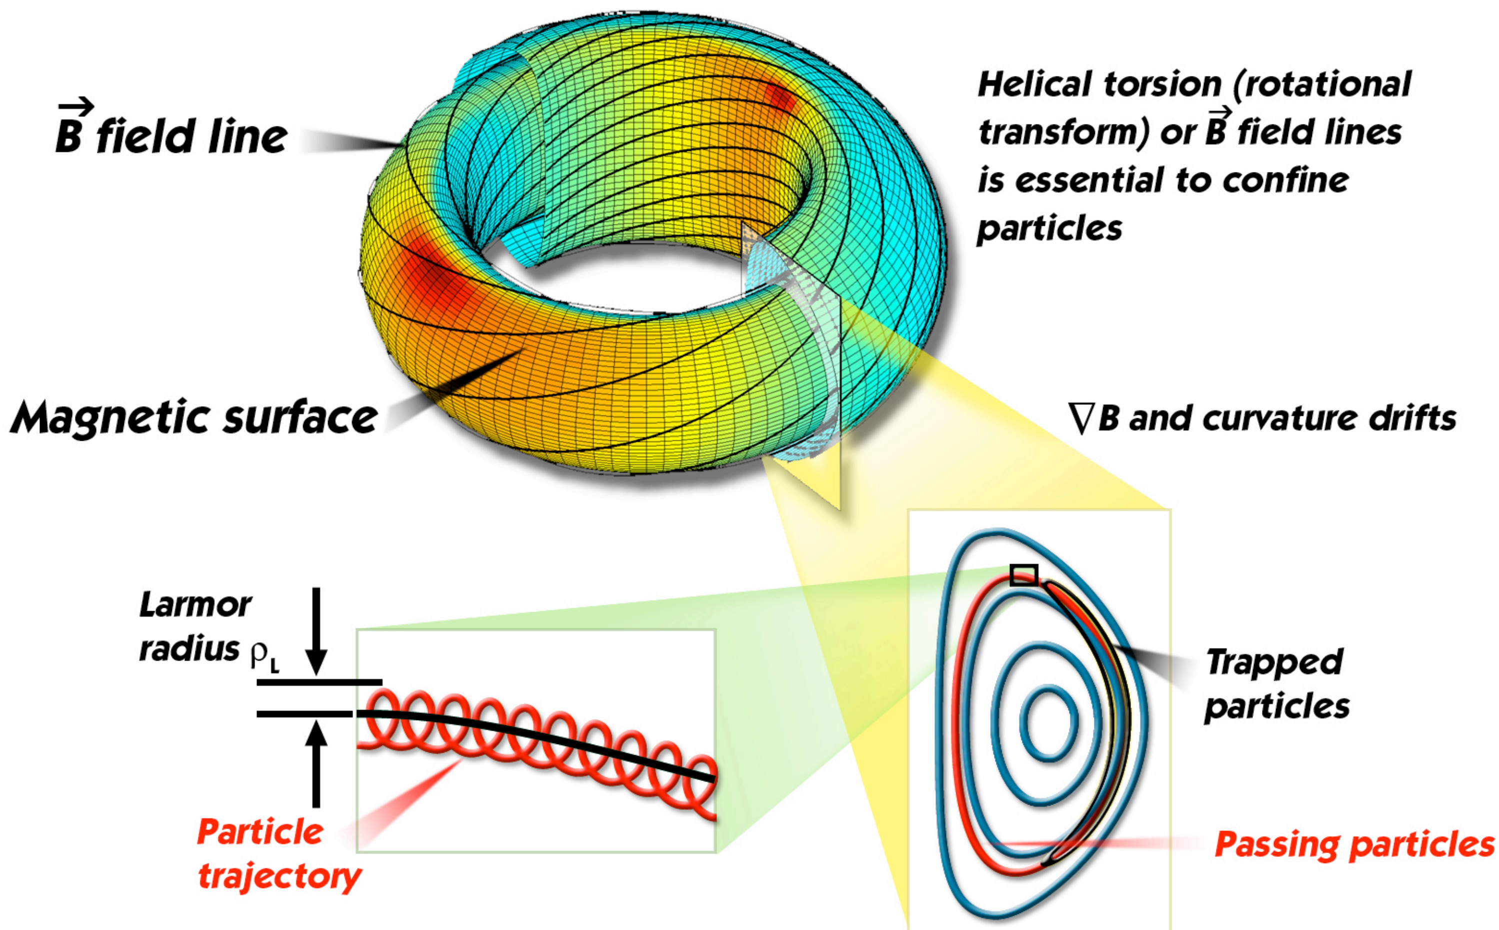

The huge range of spatial and temporal scales presents major challenges to theory and simulation

- **Overlap in scales often means strong (simplified) ordering is not possible**

- **Effective simulations at the petascale** *(10^{15} floating point operations per second)* **and beyond are required to address grand challenges in plasma science**

  e.g., understanding burning plasmas in magnetically-confined fusion systems

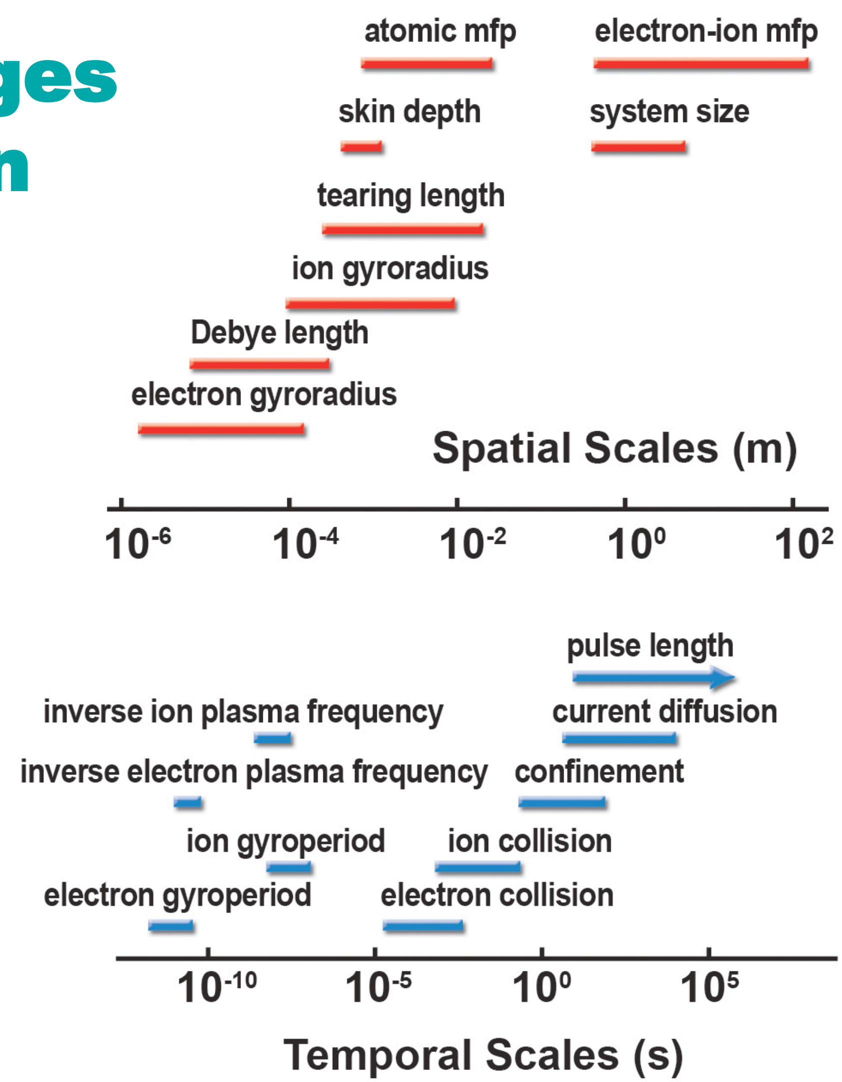

Nature of Physics & Computational Challenges

- Critical *physics* issues for fusion come from *"gaps analysis"* of the *most needed predictive capabilities* from *advanced scientific codes* that traditional *theory or experiment, by themselves, cannot readily deliver*

- Critical *computational* issues come from *"gaps analysis"* of *capabilities missing from current state-of-art tools to effectively utilize advanced computing facilities* for dealing with critical scientific issues
  - Coding/algorithmic challenges in face of *<u>increased computer architecture complexity (multi/many core)</u>*
    - *Exascale likely to arrive before fully-operational ITER*
  - Aligning computer science objectives (both hardware and software aspects) plus applied math objectives (algorithms, etc.) with the critical physics issues for fusion

Exciting Opportunities for Computational FES

- Need for reliable predictive simulation capability for *BP/ITER* *(especially in the US)*
- Powerful ("Leadership Class") Computational Facilities worldwide moving rapidly toward petascale & beyond
- Interdisciplinary *collaborative experience*, knowledge, & software assembled over 8 years under SciDAC plus OFES and OASCR base research programs in the US

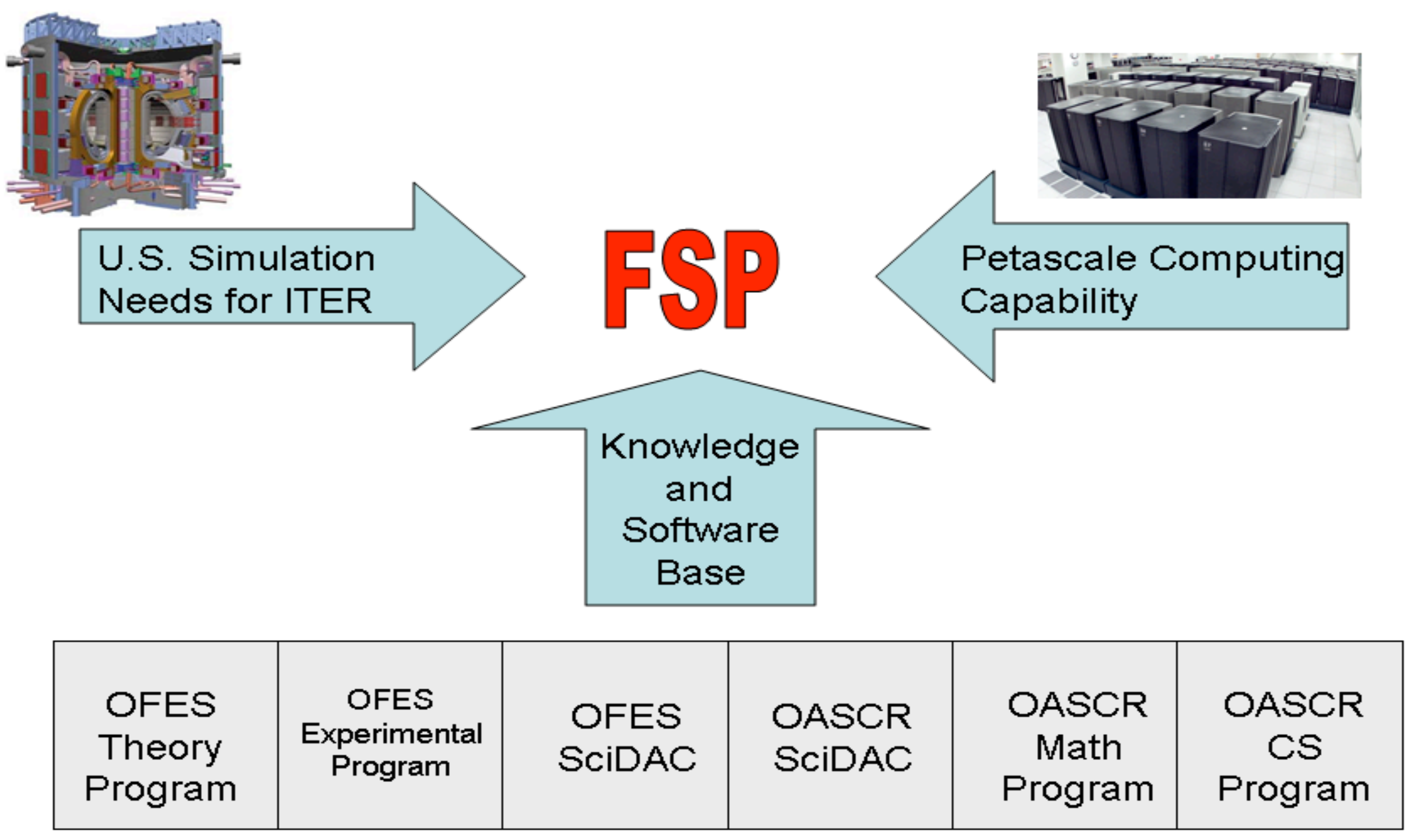

Advanced Scientific Codes --- *"a measure of the state of understanding of natural and engineered systems"* (T. Dunning)

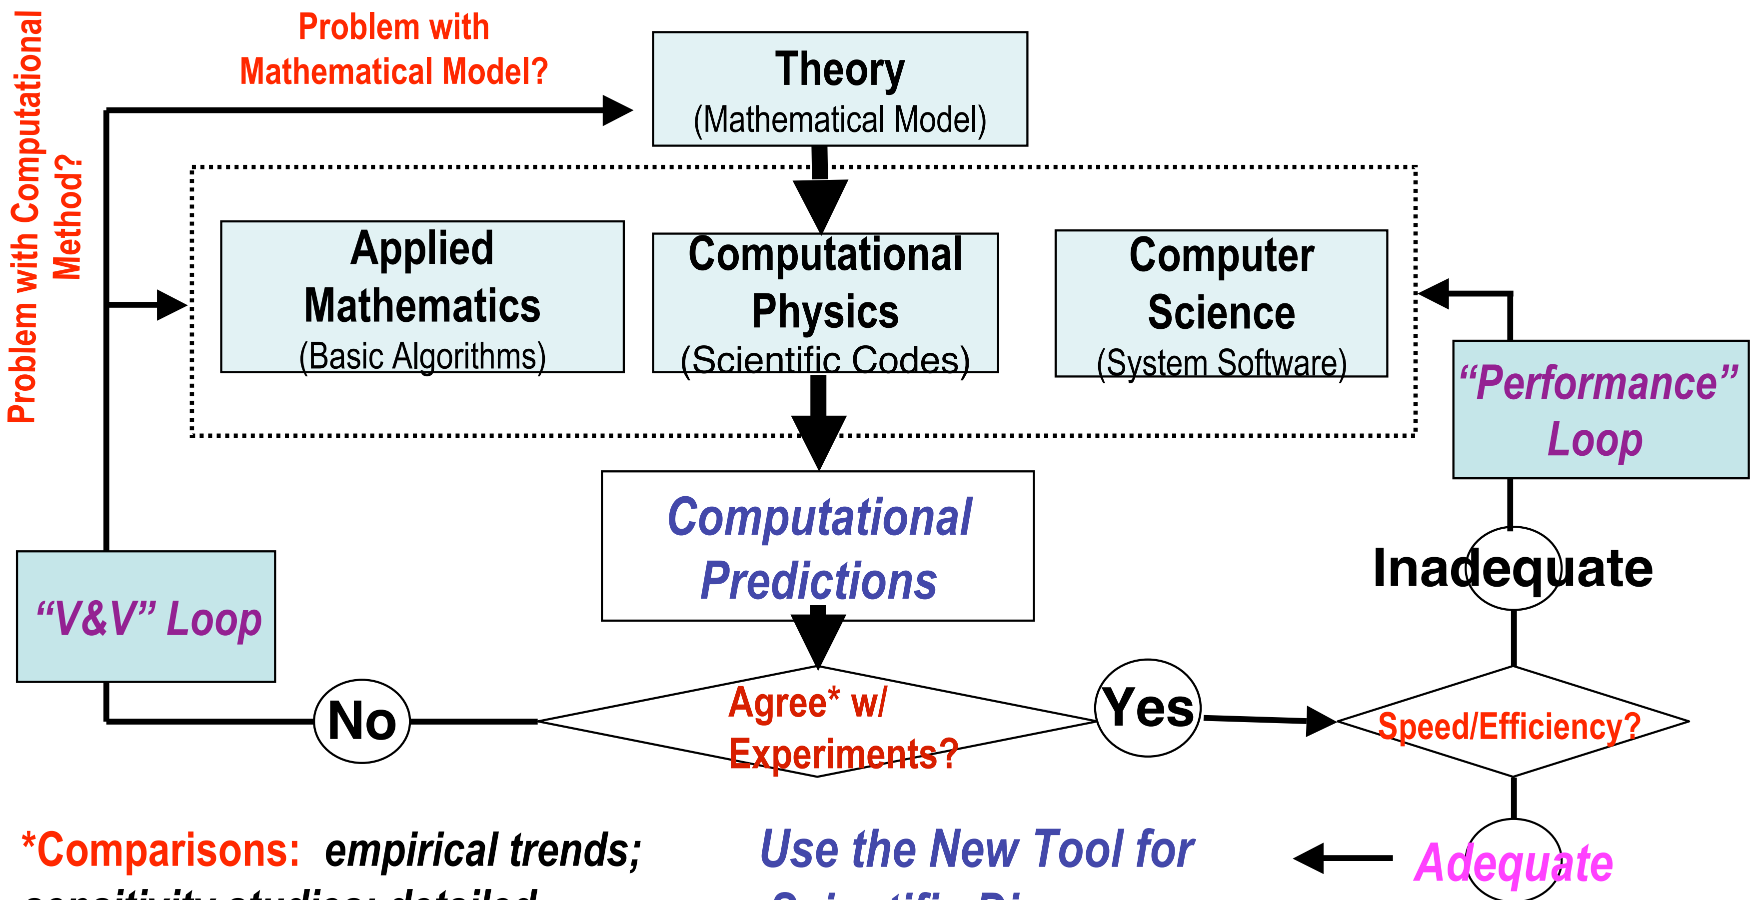

*Comparisons: *empirical trends; sensitivity studies; detailed structure (spectra, correlation functions, …)*

Use the New Tool for Scientific Discovery

(Repeat cycle as new phenomena encountered)

Physics Integration Challenges in Fusion Energy Sciences

Requirement: **HPC predictive software that embodies theoretical and experimental understanding of confined thermonuclear plasmas (extending to the ITER scale) in a realistic integrated modeling capability**

- **Goal**: *develop reliable simulation capability to predict behavior of plasma discharges in toroidal magnetic fusion devices on all relevant time and space scales in context of self-consistent calculations*
  - Verified for fidelity vs theoretical models
  - Validated vs experimental data

- **Approach**: *assessment & development of suite of codes and models that constitute two near-separate disciplines* --
  - Largest-scale (direct numerical simulation/DNS) codes addressing multi-scale physics of mostly individual phenomena in realistic 3D geometry *(largely enabled by leadership-class resources)*
  - Integrated models with much smaller-scale lower dimensionality and some empirical elements for experiment interpretation and design *(largely enabled by mid-range computing resources)*

- **Challenge**: *effectively use beyond-petascale to exascale multi-core supercomputers with associated algorithmic advances to accelerate progress in understanding complex plasma phenomena*

Elements of an Integrated Model

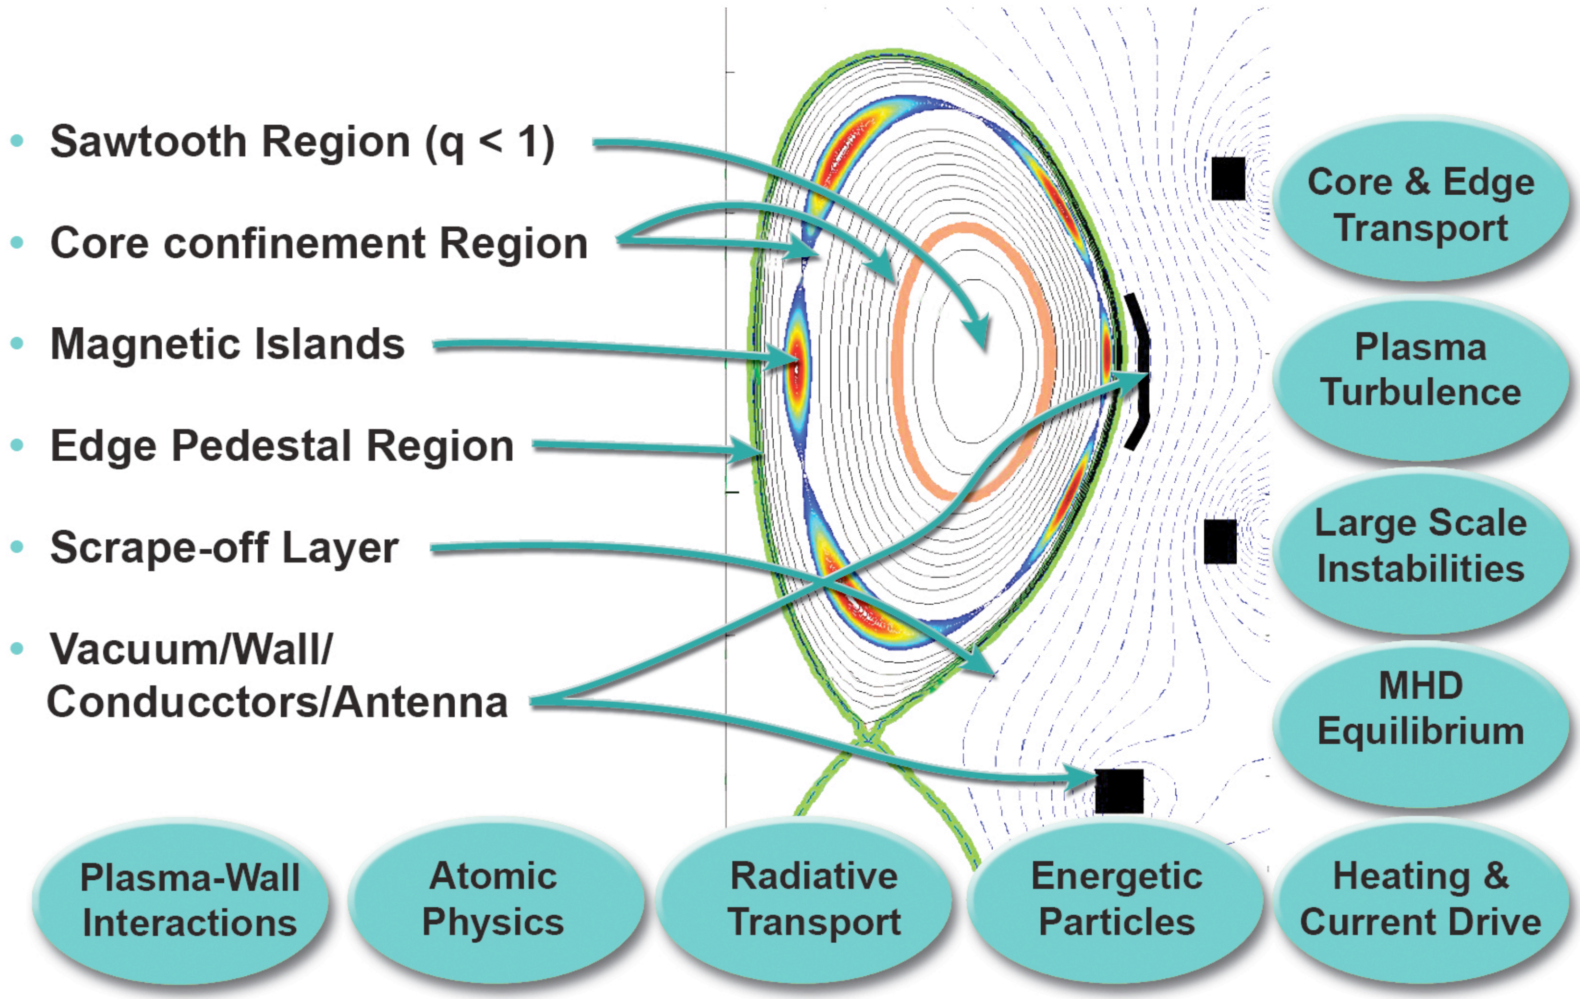

Integrated Real-time Plasma Control System

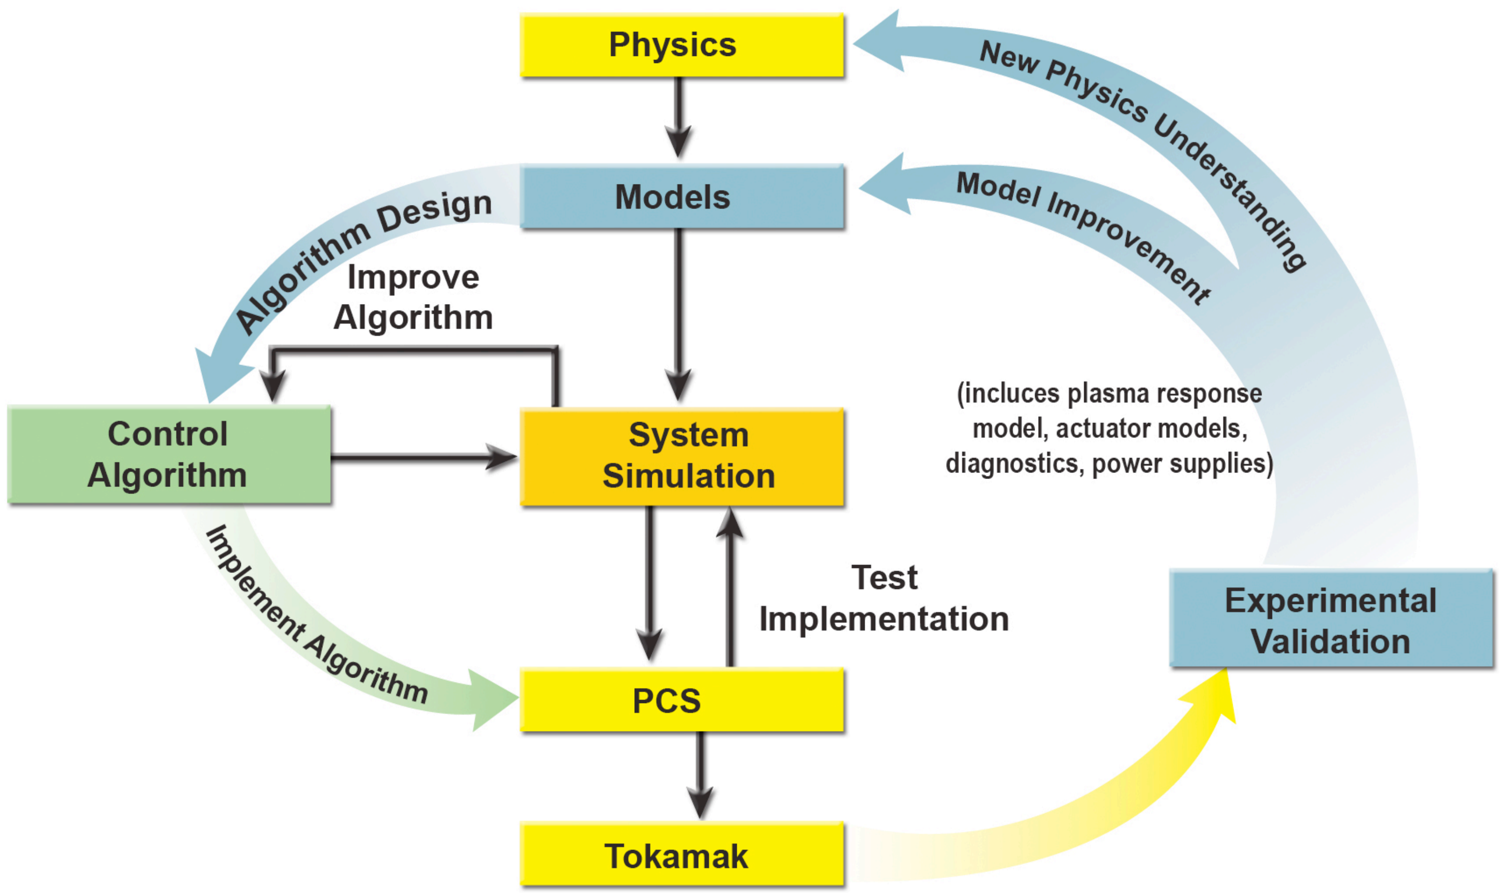

Insights from first-principles approaches should be folded into this type of reduced control-level model

Key Scientific Challenges for Burning Plasmas

- ***Disruptions***: Large-scale macroscopic events leading to rapid termination of plasma discharges
  - Avoid or mitigate because ITER can sustain only a limited number of full-current disruptions
- Pedestals: Formation of steep spatial gradients leading to transient heat loads in plasma periphery (divertor region)
  - Predict onset and growth because pedestal height is observed to control confinement
  - Predict frequency and size of ***edge localized mode (ELM) crashes*** to mitigate erosion of divertor and plasma-facing components
- Tritium migration/retention and impurity transport
- Performance optimization and scenario modeling
- Plasma feedback control
  - ***Burning plasma regime is fundamentally new with stronger self-coupling and weaker external control***

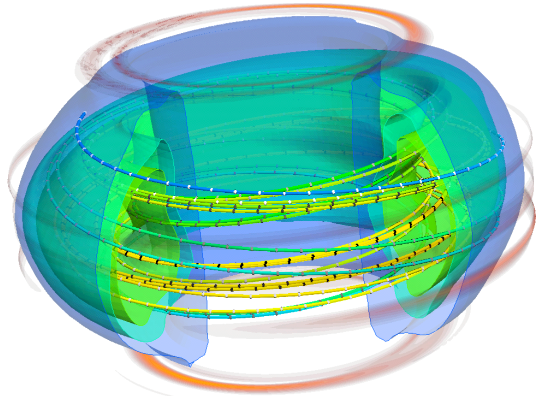

Plasma disruption in DIII-D

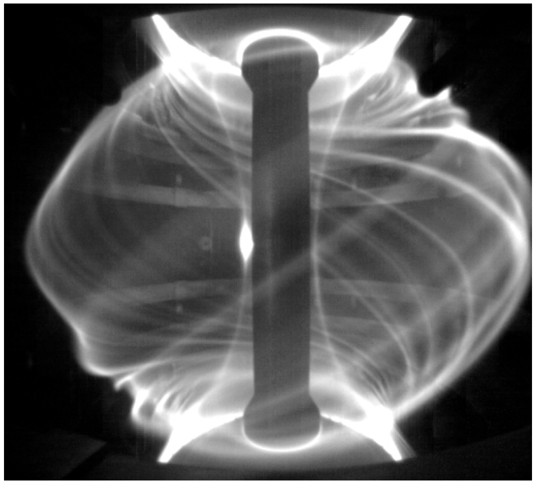

ELMs in MAST

Plasma Feedback: ELM Control/ Mitigation

Amplitude of uncontrolled ELM heat pulse in ITER expected to be order of magnitude above tolerable level for divertor plasma facing components

Magnetic Control (DIII-D tokamak)

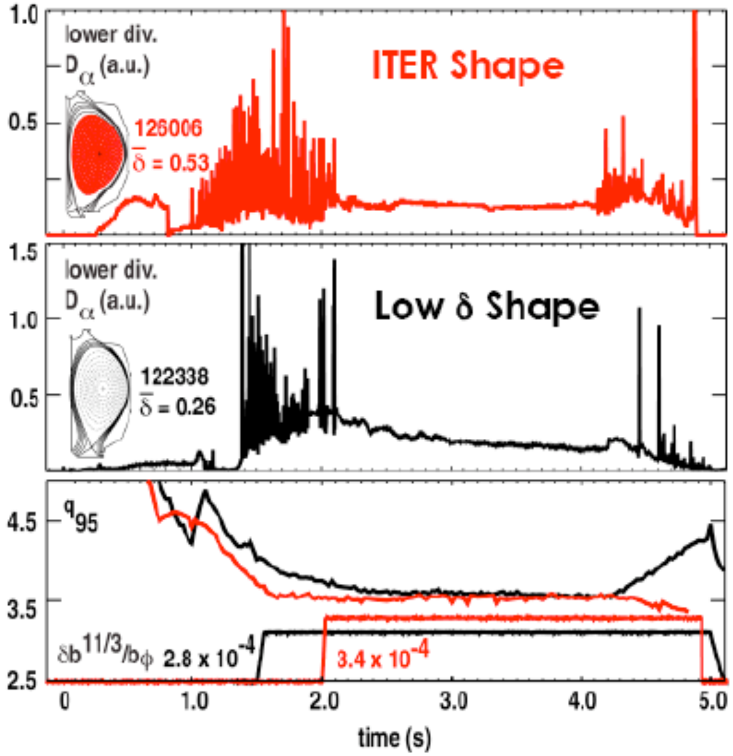

Pellet Pacemaking

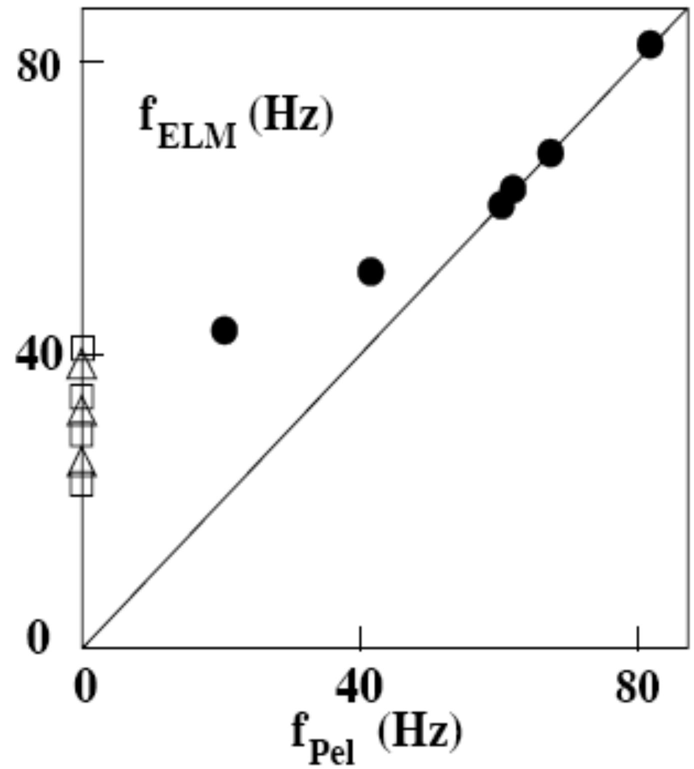

- **Two principal approaches currently under development:**

  -- edge ergodization by Resonant Magnetic-Field Perturbation (RMP) coils

  -- pellet pacemaking

Modelling of Pellet Fuelling: Using AMR *(Adaptive Mesh Refinement)*

Formidable multi-scale/multi-physics problem: resolving both small-scale pellet physics and large-scale MHD ELMs

---

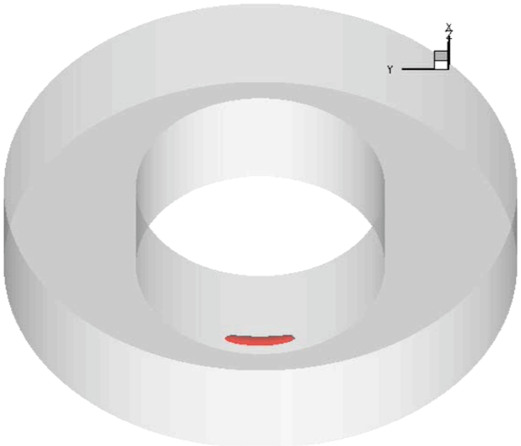

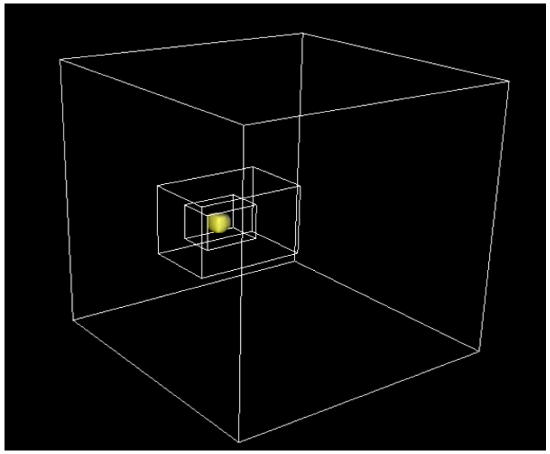

Adaptive Mesh Refinement

Inside Pellet Launch

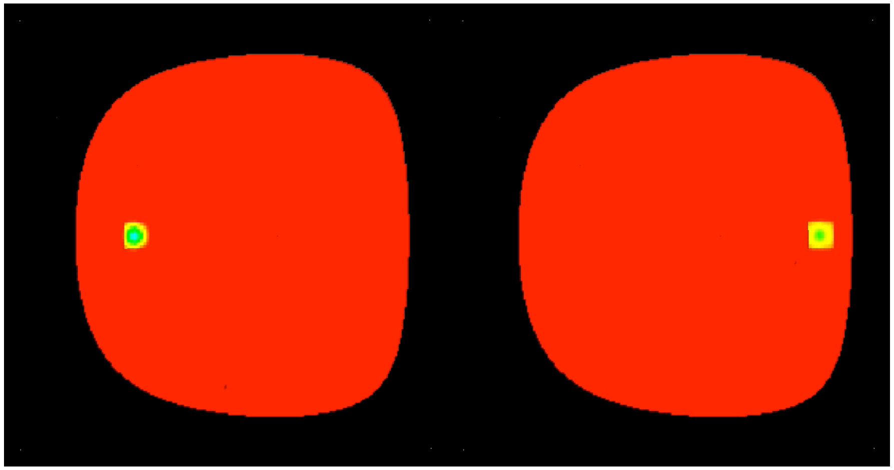

Outside Pellet Launch

Recent LCF-enabled simulations provide new insights into plasma turbulence

Applications of sustained multi-teraflops-to-petaflops computing power have accelerated progress in understanding heat losses caused by plasma turbulence

Simulations accounting for fully global 3D geometric complexity of problem have been carried out with unprecedented resolution on DOE-SC Leadership Computing Facilities

Exascale-level production runs are needed to ensure even higher physics fidelity and to support more comprehensive & realistic integrated dynamics

e.g. -- Current petascale-level production runs on ORNL's Jaguar LCF require 24M CPU hours (100,000 cores × 240 hours)

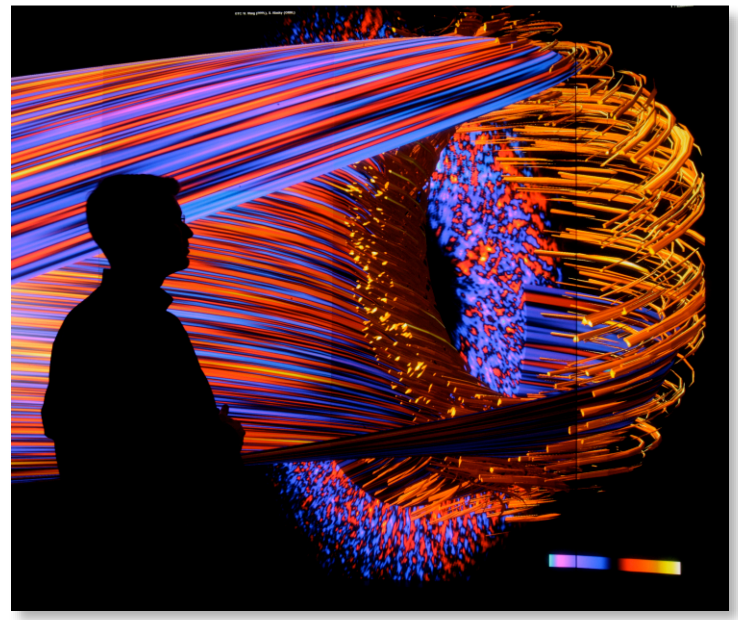

Mission Importance:

Fusion reactor size and cost are determined by balance between loss processes and self-heating rates

Scaling study of the GTS particle-in-cell (PIC) global fusion turbulence code on quad-core Jaguar

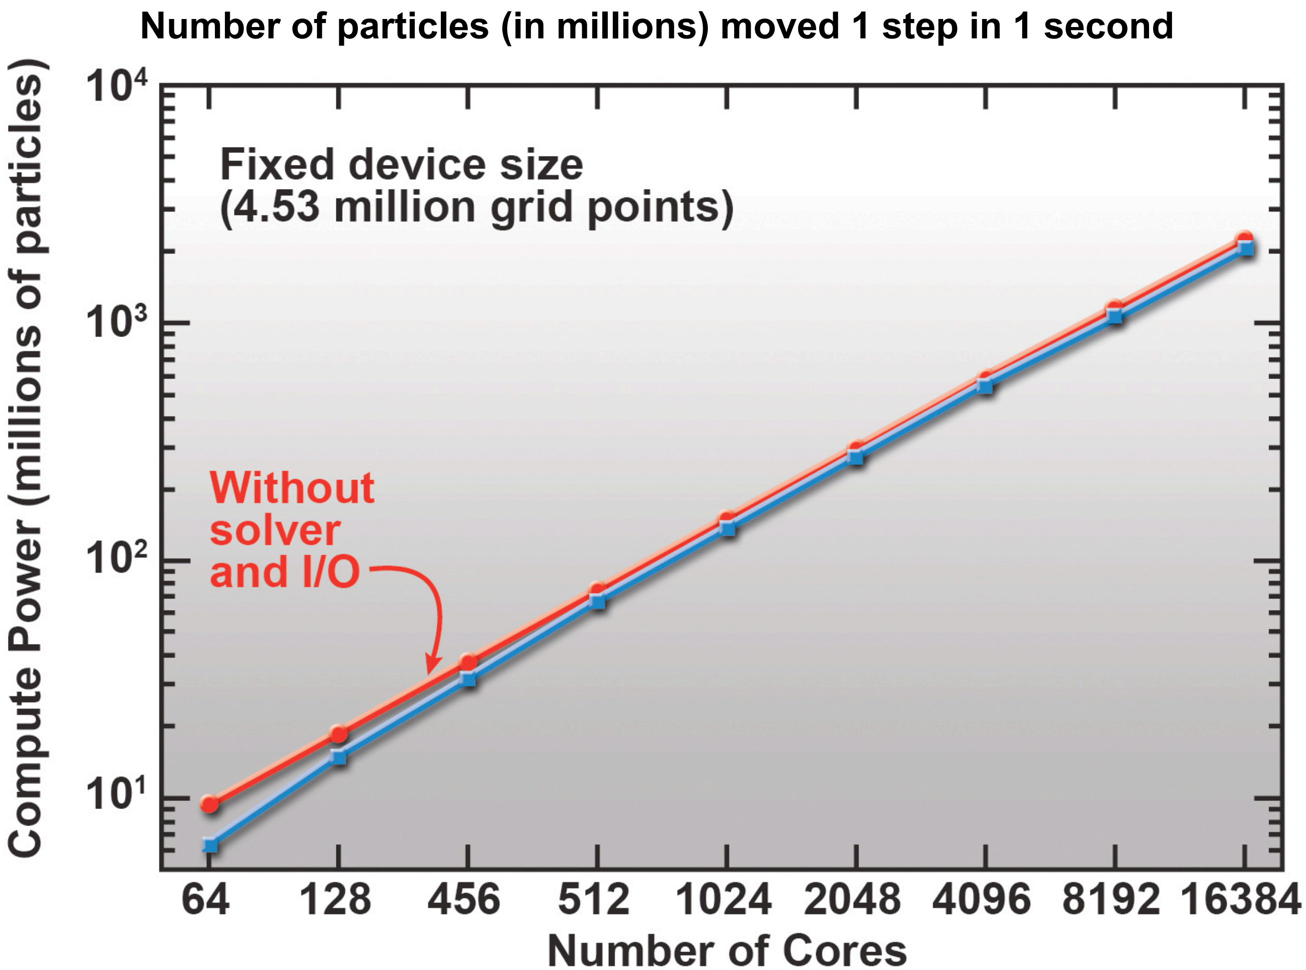

S. Ethier, PPPL, May 2008

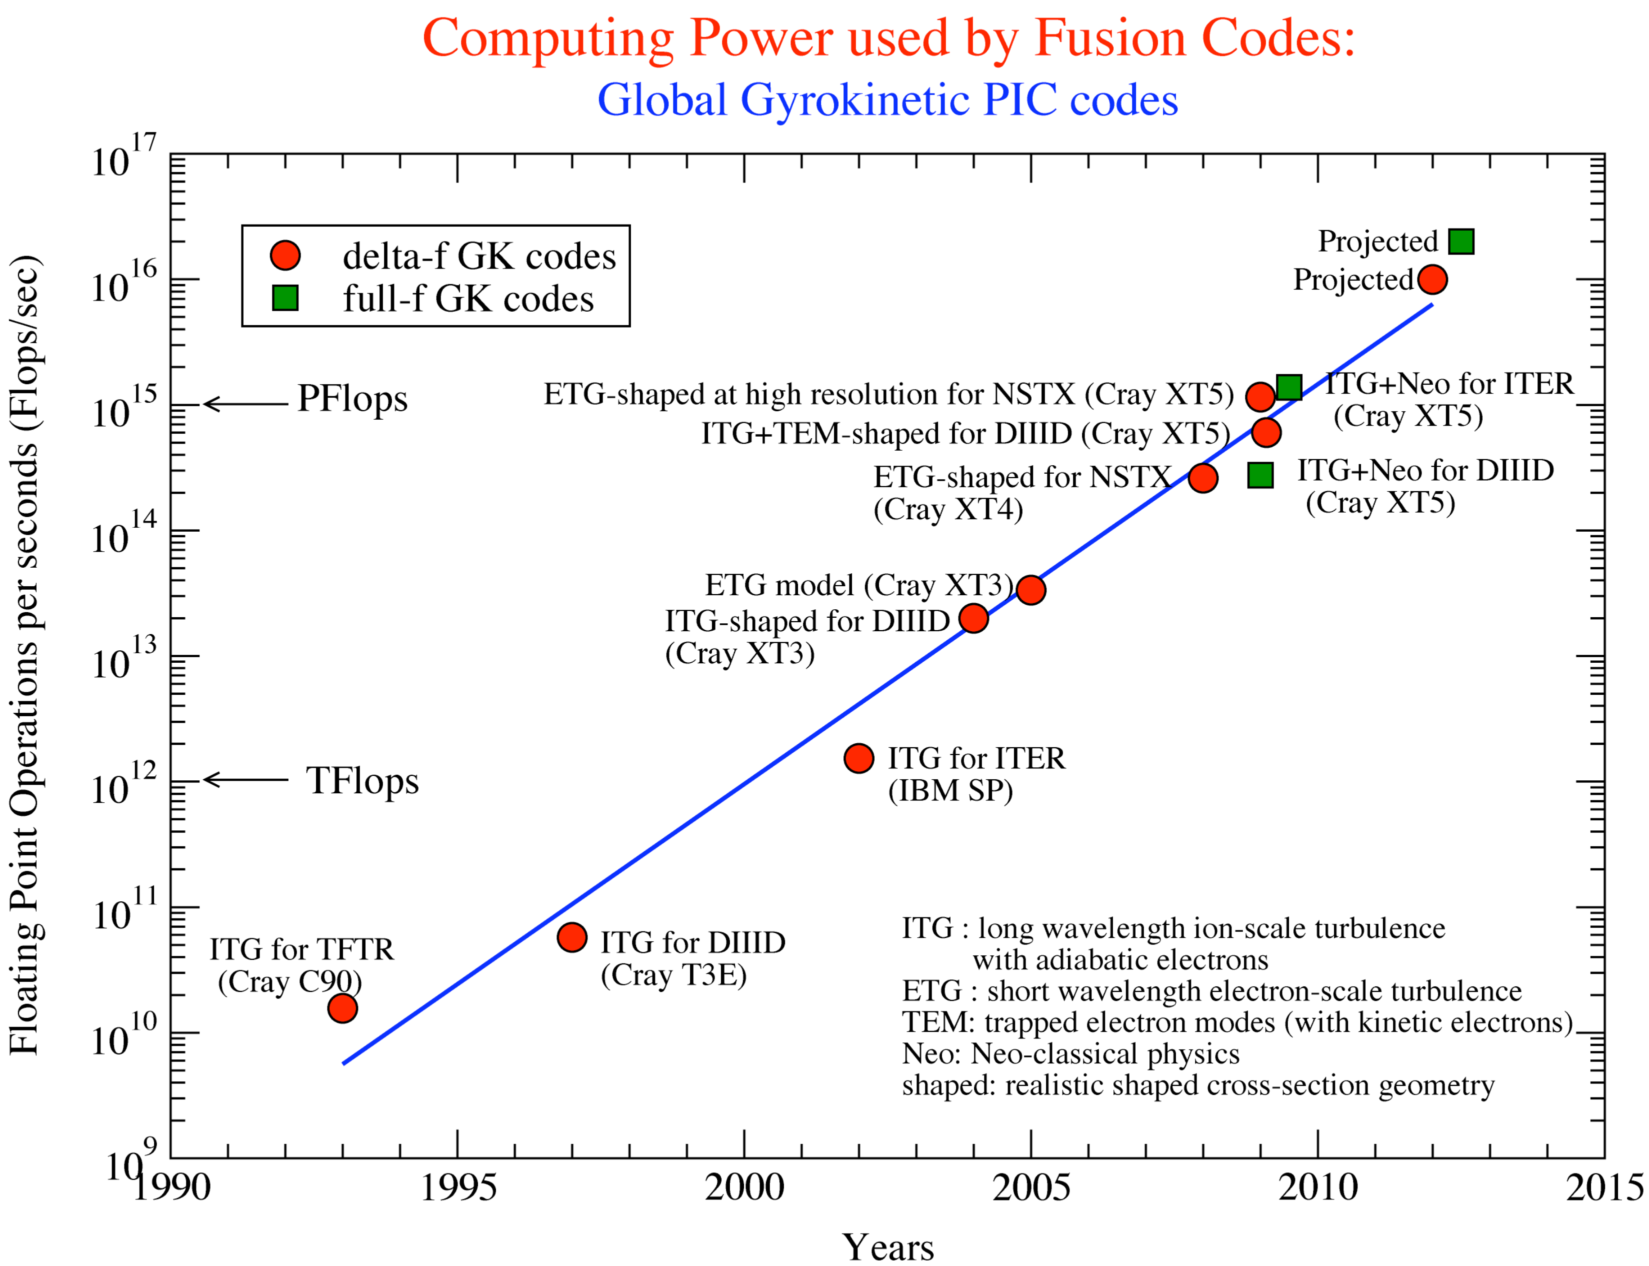
Computing Power used by Fusion Codes:
Global Gyrokinetic PIC codes
delta-f GK codes
full-f GK codes
Floating Point Operations per seconds (Flops/sec)
Years
PFlops
TFlops
Projected
Projected
ETG-shaped at high resolution for NSTX (Cray XT5)
ITG+Neo for ITER (Cray XT5)
ITG+TEM-shaped for DIIID (Cray XT5)
ETG-shaped for NSTX (Cray XT4)
ITG+Neo for DIIID (Cray XT5)
ETG model (Cray XT3)
ITG-shaped for DIIID (Cray XT3)
ITG for ITER (IBM SP)
ITG for TFTR (Cray C90)
ITG for DIIID (Cray T3E)
ITG : long wavelength ion-scale turbulence with adiabatic electrons
ETG : short wavelength electron-scale turbulence
TEM: trapped electron modes (with kinetic electrons)
Neo: Neo-classical physics
shaped: realistic shaped cross-section geometry

Microturbulence in Fusion Plasmas: *Size & Cost of reactor from balance between confinement & fusion self-heating rates*

- *"Scientific Discovery"* - Transition to favorable scaling of confinement observed for large plasmas of future
- *Data Streaming Technology* enabled moving terabytes of data from NERSC to PPPL

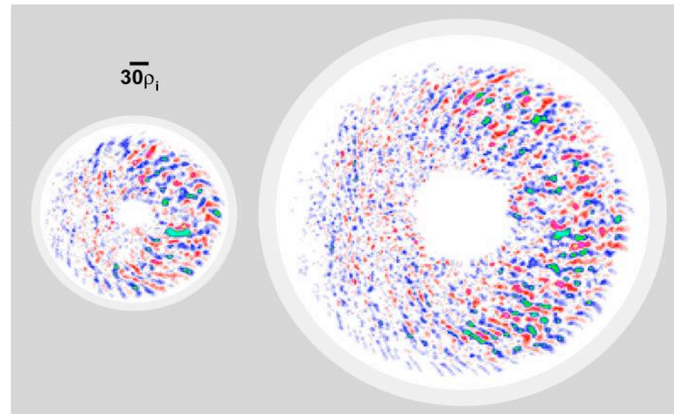

Good news for ITER!

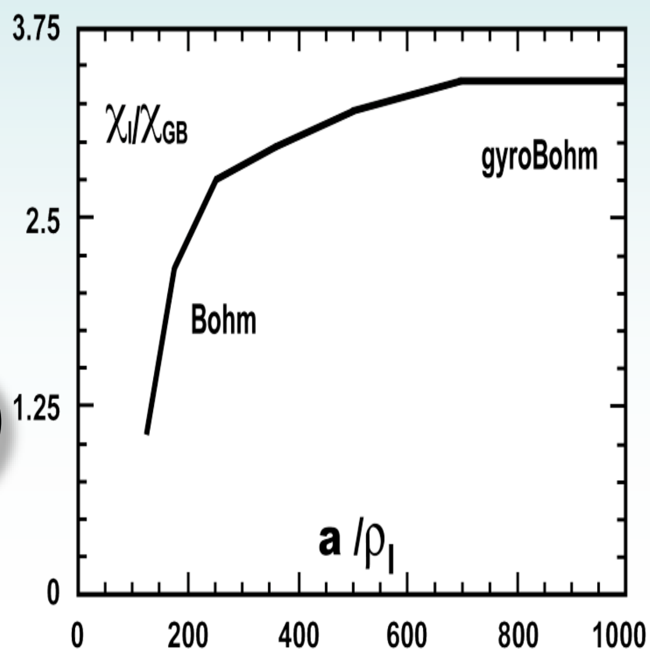

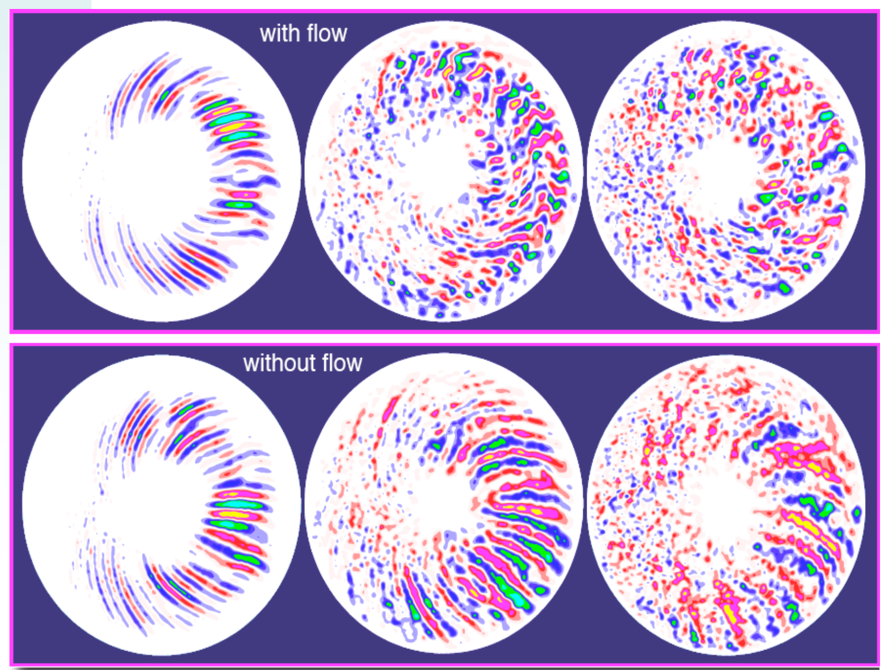

Recent High-Resolution Simulations

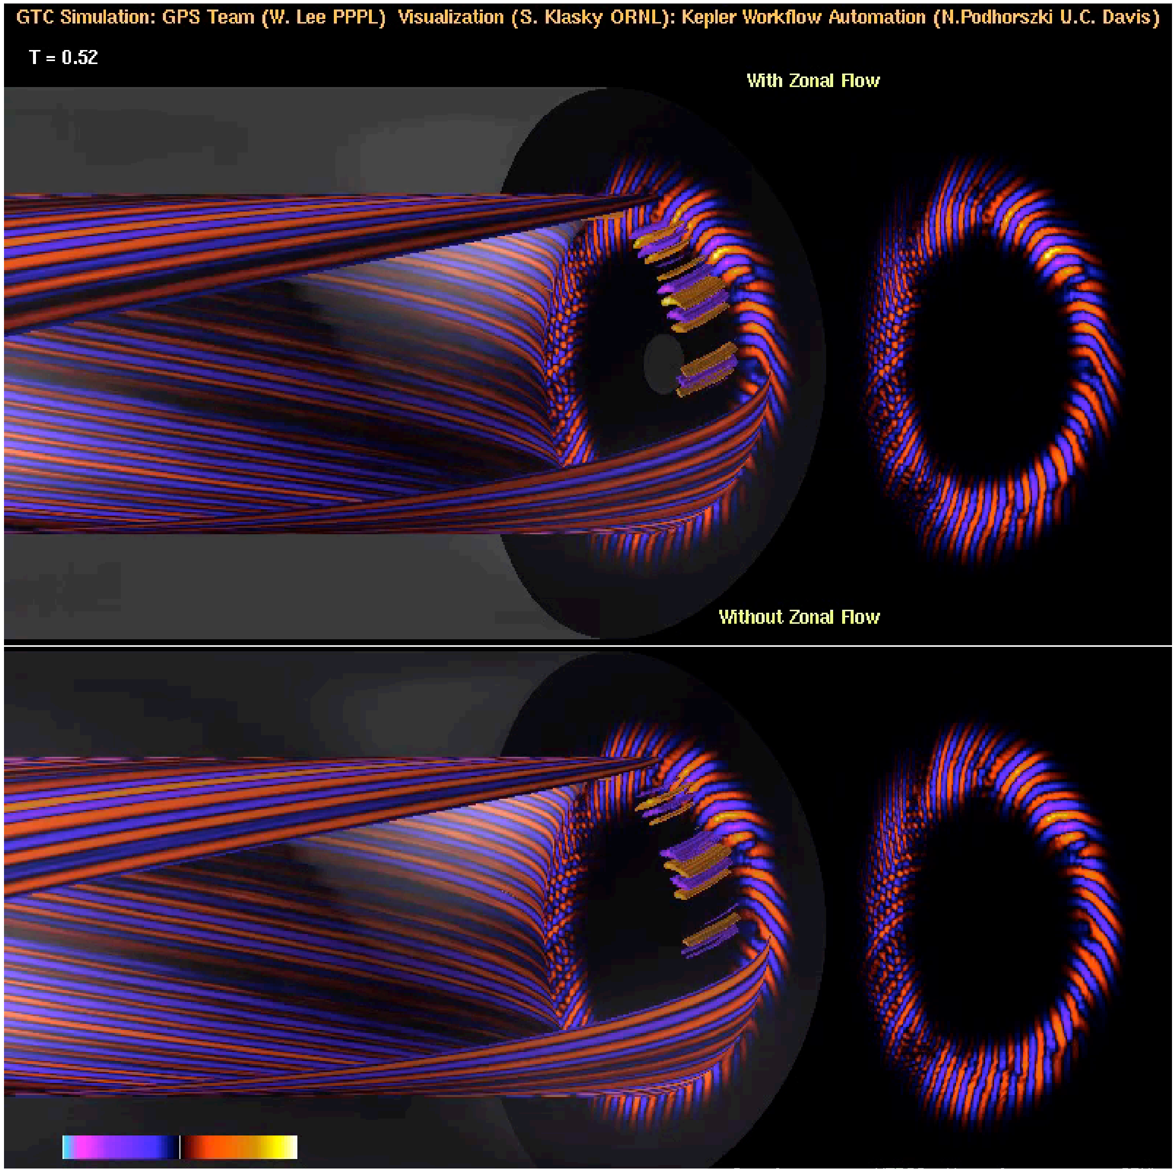

- High-resolution visualization from ***realistic shaped-cross section toroidal plasma simulations*** on leadership class computers

- Efficiently generated via "Workflow Automation" -- ***automation of data movement, data reduction, data analysis, and data visualization***

Verification & Validation Challenges

- **Establishing the physics fidelity of modern plasma science simulation tools demands proper Verification & Validation (V&V) -- *Reliable codes demand solid theoretical foundations and careful experimental validation***

  • Verification assesses degree to which a code (*both in the advanced direct numerical simulation (DNS) and reduced models categories*) correctly implements the chosen physical model

  --- more than "essentially a mathematical problem"

  e.g., accuracy of numerical approximations, mesh/space and temporal discretization, statistical sampling errors, etc.

--- Special emphasis should be placed on code verification via:

(1) comparisons with theoretical predictions

e.g. -- threshold/onset conditions for instabilities; weakly nonlinear evolution; nonlinear saturation estimates; etc.

(2) cross-code benchmarking (codes based on different mathematical formulations/algorithms but targeting the same generic physics)

e.g. -- finite difference, finite elements, spectral methods, implicit schemes, etc.and/or models such as Particle-in-Cell, Vlasov/Continuum, Hybrid PIC-Fluid, etc.

Verification & Validation Challenges

• Validation assesses degree to which a code (within its domain of applicability) "describes the real world," e.g.

Schematic: *Combined Efforts from Theory/Modeling/Experiment for Realistic Predictive Transport Capability in Plasma Core*

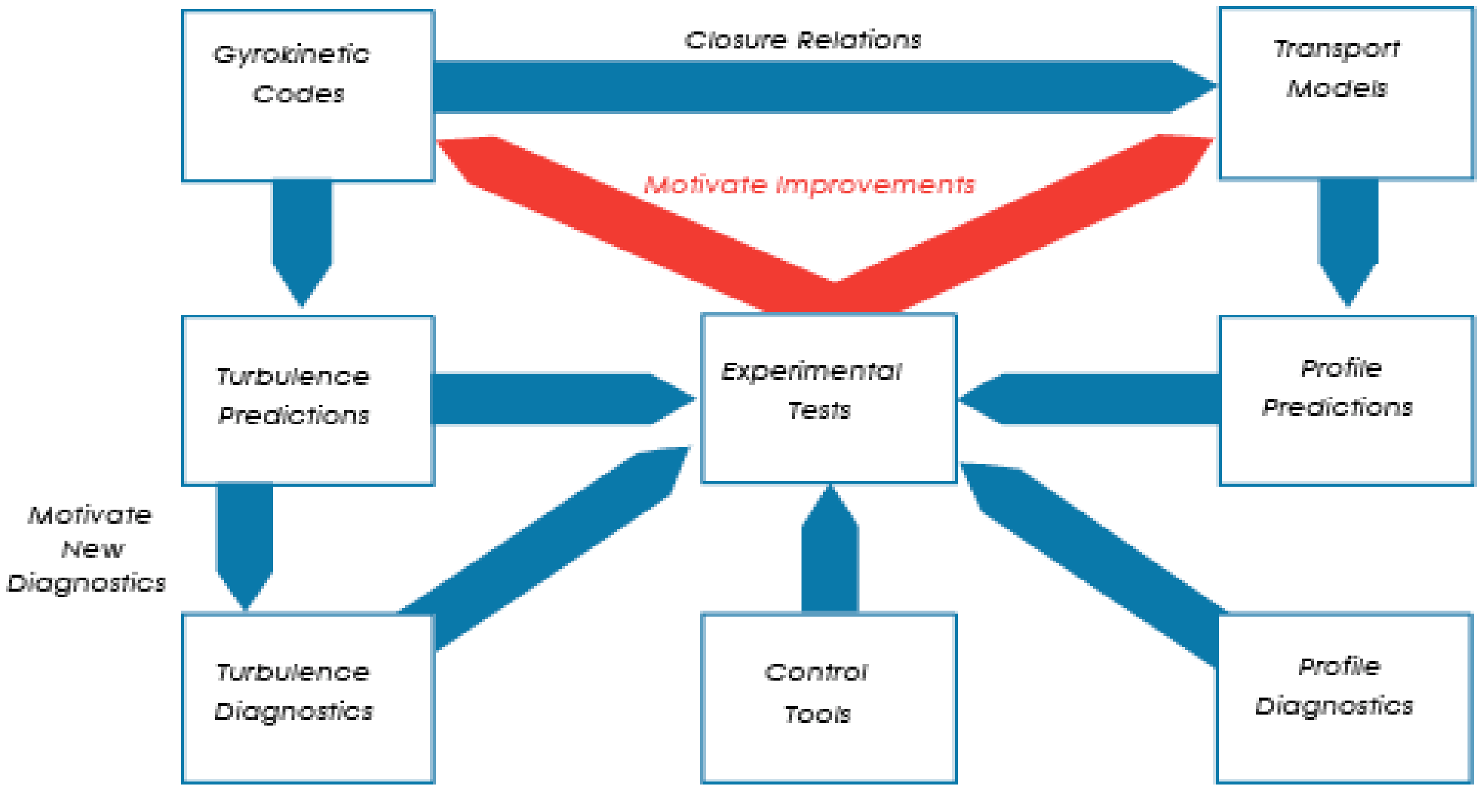

• *V & V in FES/Plasma Science can benefit from "lessons learned" from other prominent applications domains featuring large scale simulations -- e.g., climate modeling, combustion, ASCI, etc.*

Looking toward the Future: “Exascale”

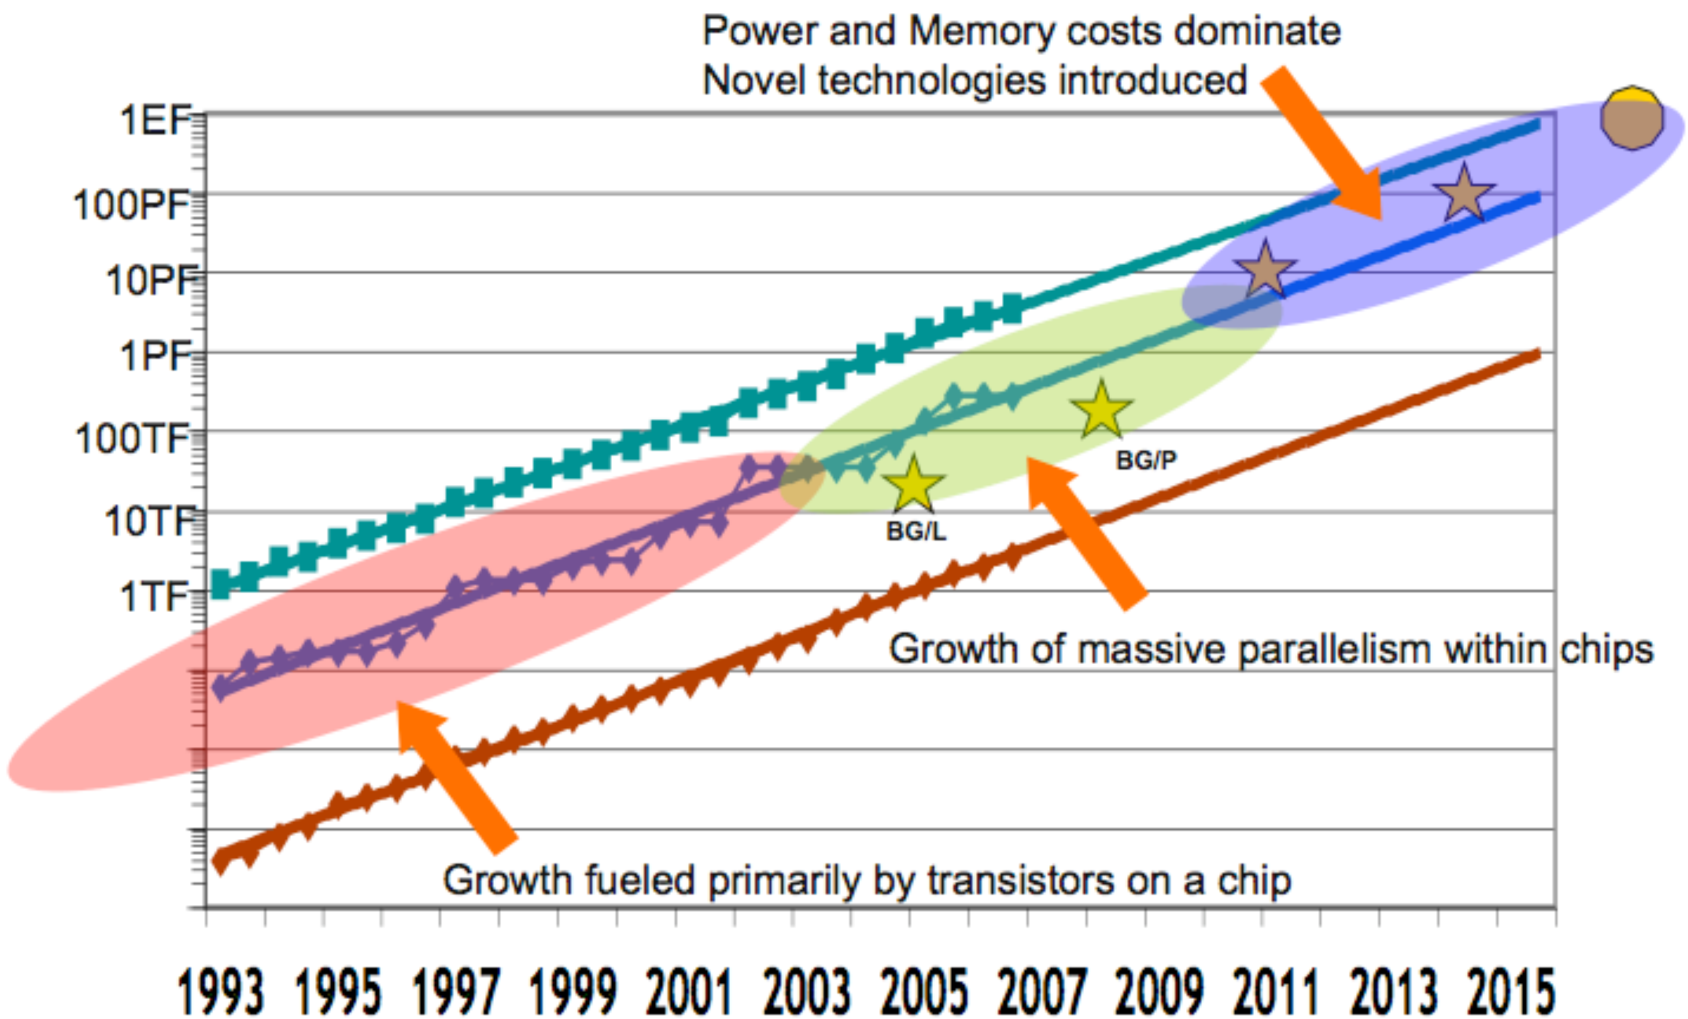

Slide courtesy Argonne Leadership Computing Facility (ALCF)/Argonne National Laboratory

Associated Major Challenges

- ***Some are obvious …***
  - Hardware complexity: Heterogenous multicore, power management, error control, communications, storage, …
  - Software challenges: Operating systems, I/O and file systems, and coding/algorithmic needs in the face of ***increased computer architecture complexity*** … ***"parallelism doubles every two years" (as a new form of Moore's Law)***
- ***Some are less so …***
  - Achieving greater "buy-in" from broader scientific community:
    - Distinguish between ***"voracious"*** (*more of same - just bigger & faster*) vs. ***"transformational"*** (*achievement of major new levels of scientific understanding*)
    - Improve ***significantly*** on ***experimental validation*** and ***theoretical verification*** to enhance realistic predictive capability
  - **People:** ***Training the next generation of simulation/modeling-oriented CS, Applied Math and applications-oriented computational scientists and engineers ….***

Concluding Comments

- The next major milestone in MFE research is a *burning plasma experiment* -- leading to *ITER* -- a multi-billion dollar international collaboration centered in France & involving 7 governments representing over half of world's population
- ITER targets 500 MW for 400 seconds with gain > 10 to demonstrate *technical feasibility of fusion energy* & *DEMO* (*demonstration power plant*) will target 2500 MW with gain of 25
- Clear need for using *advanced computation to harvest knowledge from ITER and for designing DEMO*
- Future Integrated Modeling Tools will target realistic simulations of fusion and energy systems with unprecedented physics fidelity

    -- involves delivering shorter-term opportunistic HPC software tools (built largely from modestly improved existing tools); &

    -- parallel longer-term development emphasizing new, more rigorous, more engineered performance capabilities

- *In general, progress in delivering reliable* predictive capabilities *in Fusion Energy Science will benefit significantly from access to HPC resources -- from terascale to petascale & beyond -- together with a vigorous* verification & validation program