Scientific and Computational Challenges in Fusion Energy Sciences

W. M. TANG

Princeton University, Plasma Physics Laboratory

Princeton, NJ USA

Physics Department

University of California @ Irvine (UCI)

Irvine, CA

May 19, 2011

Computational Science, Exascale Computing & Leadership in Science & Technology

- **Scientific & technological challenges:**
  - *New fuels and reactors (fission & fusion)*
  - *Stewardship without nuclear tests*
  - *Carbon sequestration alternatives*
  - *Regional climate impacts*
- **Broader application of exascale computing can significantly accelerate progress in both applied and fundamental science**
  - *Renewable energy and energy storage*
  - *Prediction and control of materials in extreme environments*
  - *Understanding dark energy and dark matter*
  - *Clean and efficient combustion in advanced engines*

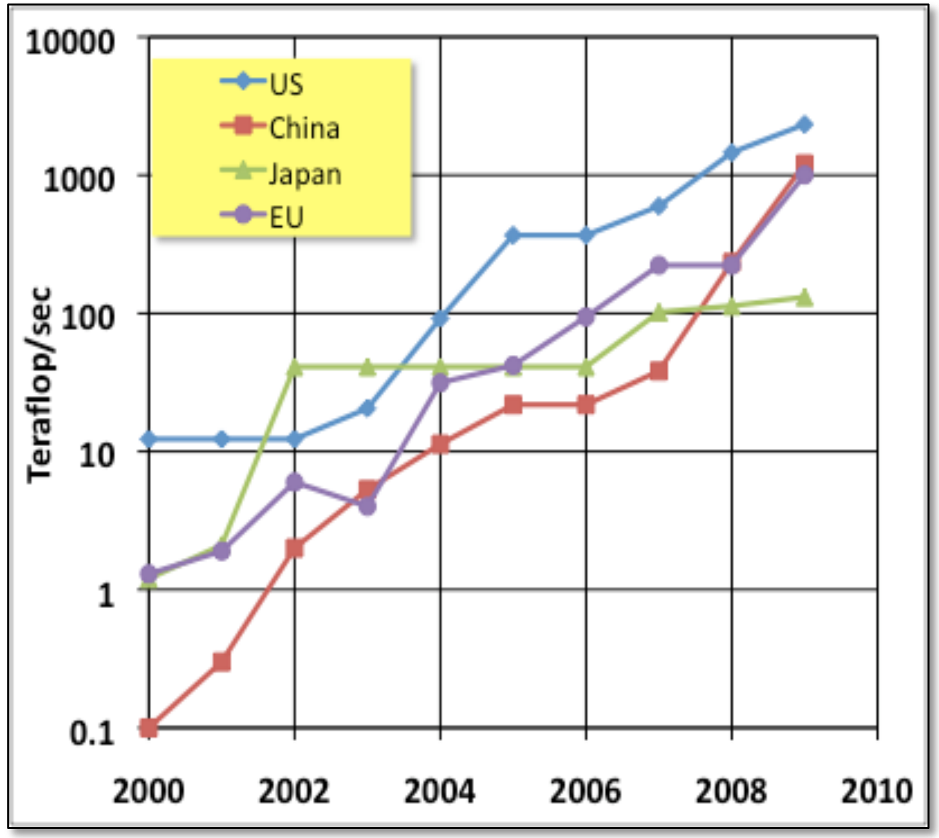

International Competition in HPC

Chart shows most capable system for each year in TOP500

Pre-eminence in 21st Century science & technology requires leadership in computational science and high performance computing => exascale applications

Advanced Computing can Transform Many Domain Applications Areas (including FES)

- ***Practical Considerations:*** [achieving "buy-in" from general scientific community]
  - Need to distinguish between *"voracious"* (*more of same - just bigger & faster*) *vs.* *"transformational"* (*achievement of major new levels of scientific understanding*)
  - Need to improve *significantly* on *experimental validation* together with *verification & uncertainty quantification* to enhance realistic predictive capability

- ***Associated Extreme Scale Computing Challenges:***
- *Hardware complexity*: Heterogenous multicore *(e.g., cpu+gpu -- LANL, ORNL, …),* power management, memory, 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)*

  (MPI + threads; CUDA; rewriting code focused on data movement over arithmetic; …..)

  ***People: *Major challenge to attract, train, & assimilate the next generation of simulation/modeling-oriented CS, Applied Math and applications-oriented computational scientists and engineers.**

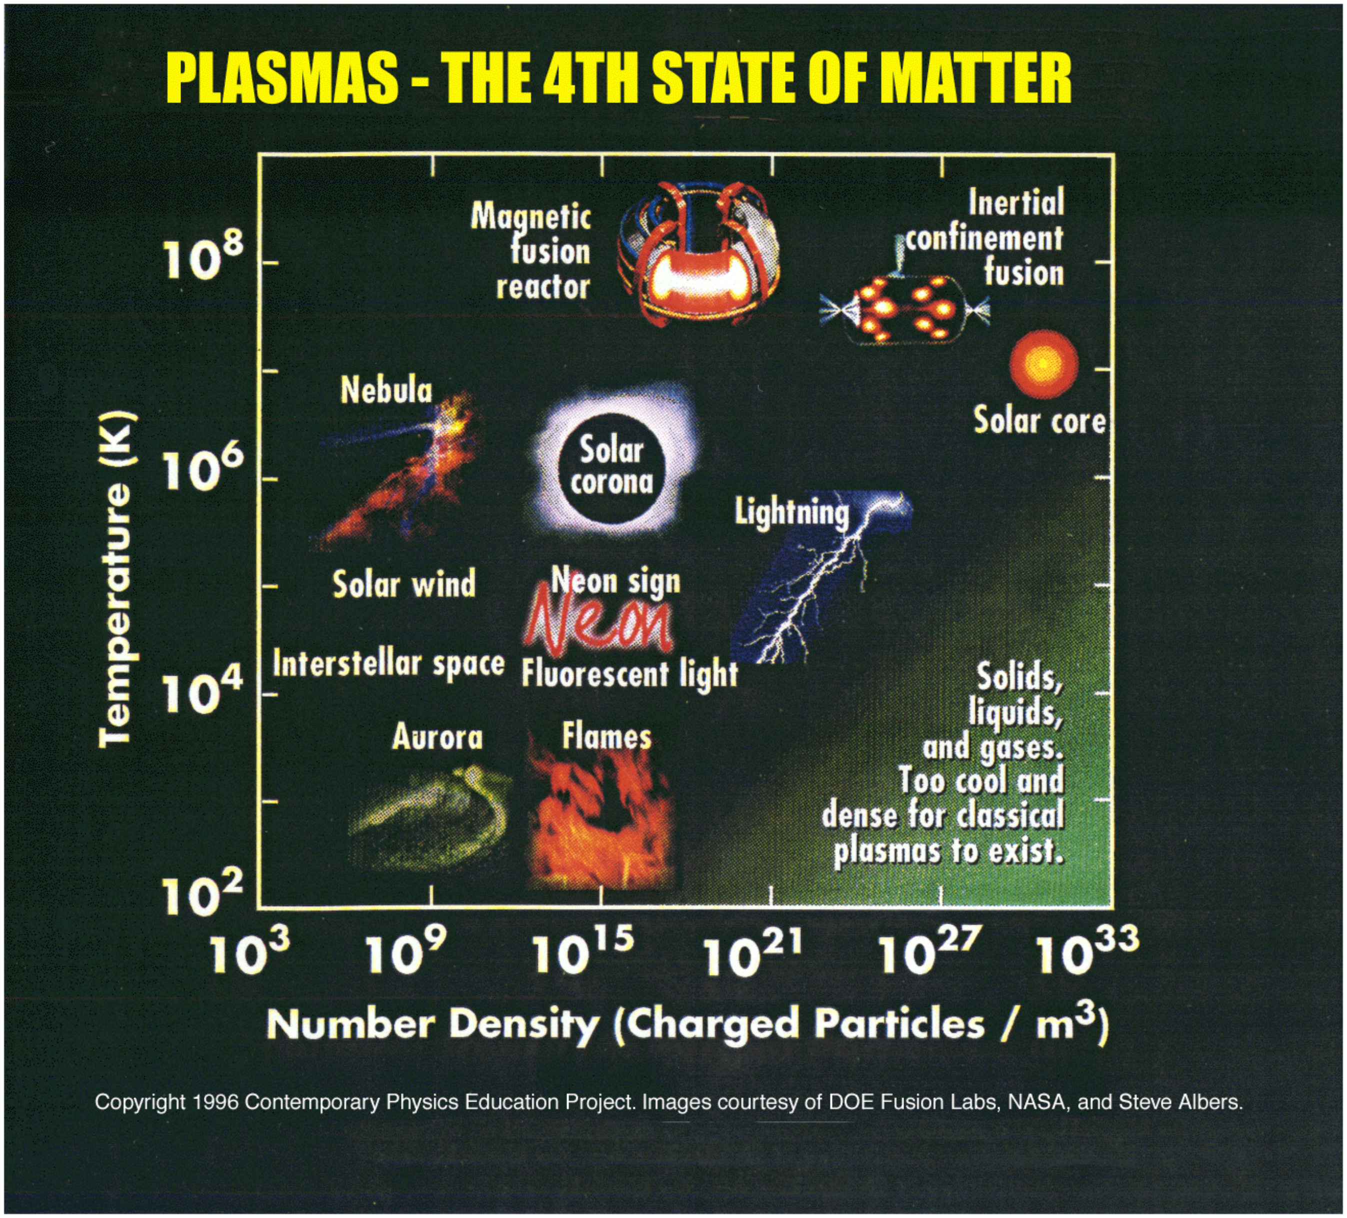
PLASMAS - THE 4TH STATE OF MATTER
Magnetic fusion reactor
Inertial confinement fusion
Nebula
Solar corona
Solar core
Lightning
Solar wind
Neon sign
Neon
Interstellar space
Fluorescent light
Aurora
Flames
Solids, liquids, and gases. Too cool and dense for classical plasmas to exist.
Temperature (K)
10^8
10^6
10^4
10^2
10^3
10^9
10^{15}
10^{21}
10^{27}
10^{33}
Number Density (Charged Particles / m^3)
Copyright 1996 Contemporary Physics Education Project. Images courtesy of DOE Fusion Labs, NASA, and Steve Albers.

ADVANCED COMPUTING IS AN INCREASINGLY POWERFUL TOOL FOR SCIENTIFIC DISCOVERY

- Advanced computation in tandem with theory and experiment is powerful tool for scientific understanding and innovation in research
- Plasma Science is effectively utilizing the exciting advances in Scientific Computing
  - References:

    (1) "Advances and Challenges in Computational Plasma Science"
    Plasma Physics 47 (February, 2005)

    (2) "Scientific and Computational Challenges of the Fusion Simulation Project "
    Journal of Physics: Conference Series, SciDAC Conference (June, 2008)

    (3) DoE 2010: "Grand Challenges in Fusion Energy Sciences and Computing at the Extreme Scale" http://www.er.doe.gov/ascr/ProgramDocuments/Docs/FusionReport.pdf

- Primary goal is to accelerate progress toward reliable/validated predictions of complex properties of high temperature plasmas
  - Predictive models superior to those based on empirical scaling trends demand acquisition of improved scientific insights and understanding

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

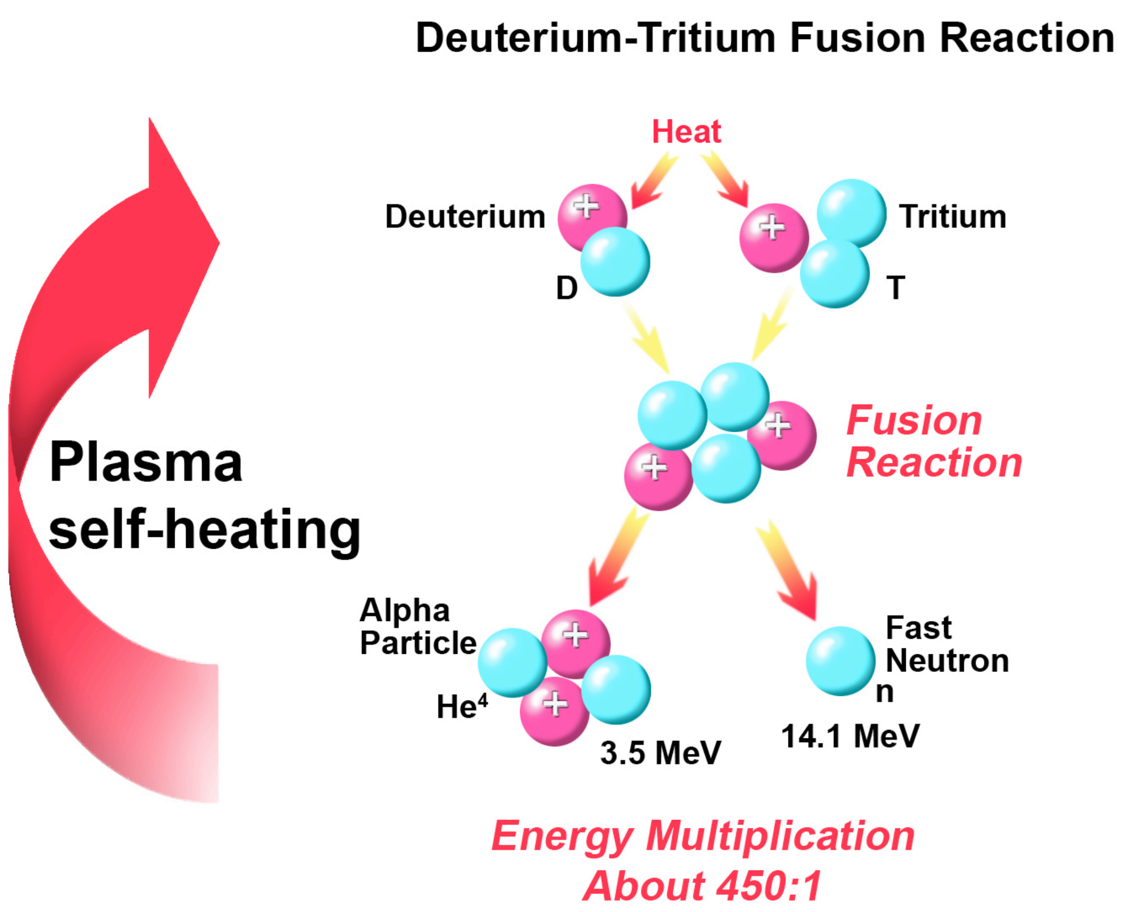

$$D^{+} + T^{+} \rightarrow {}^{4}He^{++}\ (3.5\ MeV) + n^{0}\ (14.1\ MeV)$$

Approaches to Fusion Confinement

Magnetic Confinement

Dilute: 10^{15} particles/cubic cm, long-lived: continuous

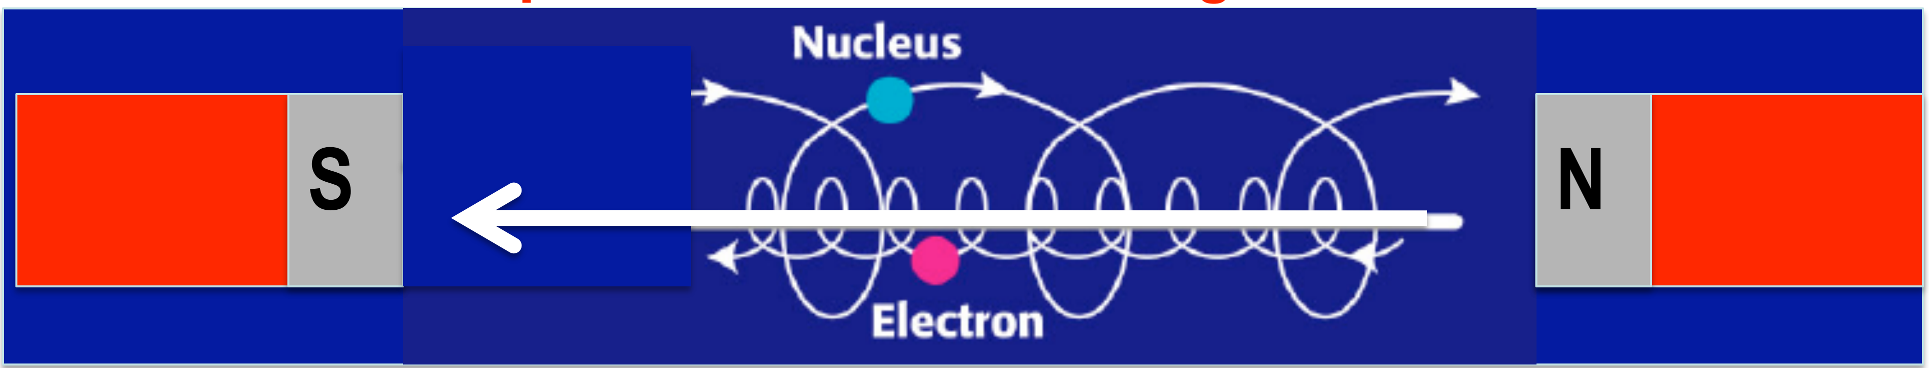

Inertial Confinement

Extremely dense: $10^{25}/cm^3$,

Short-lived ≈ one billionth of a sec

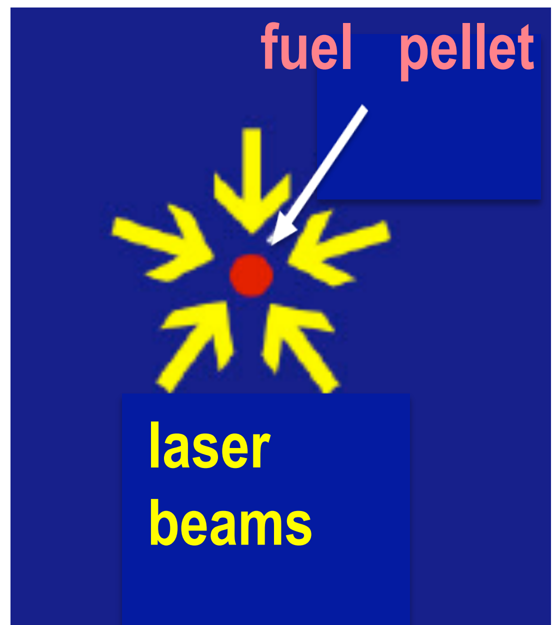

Inertial Fusion Energy (IFE)

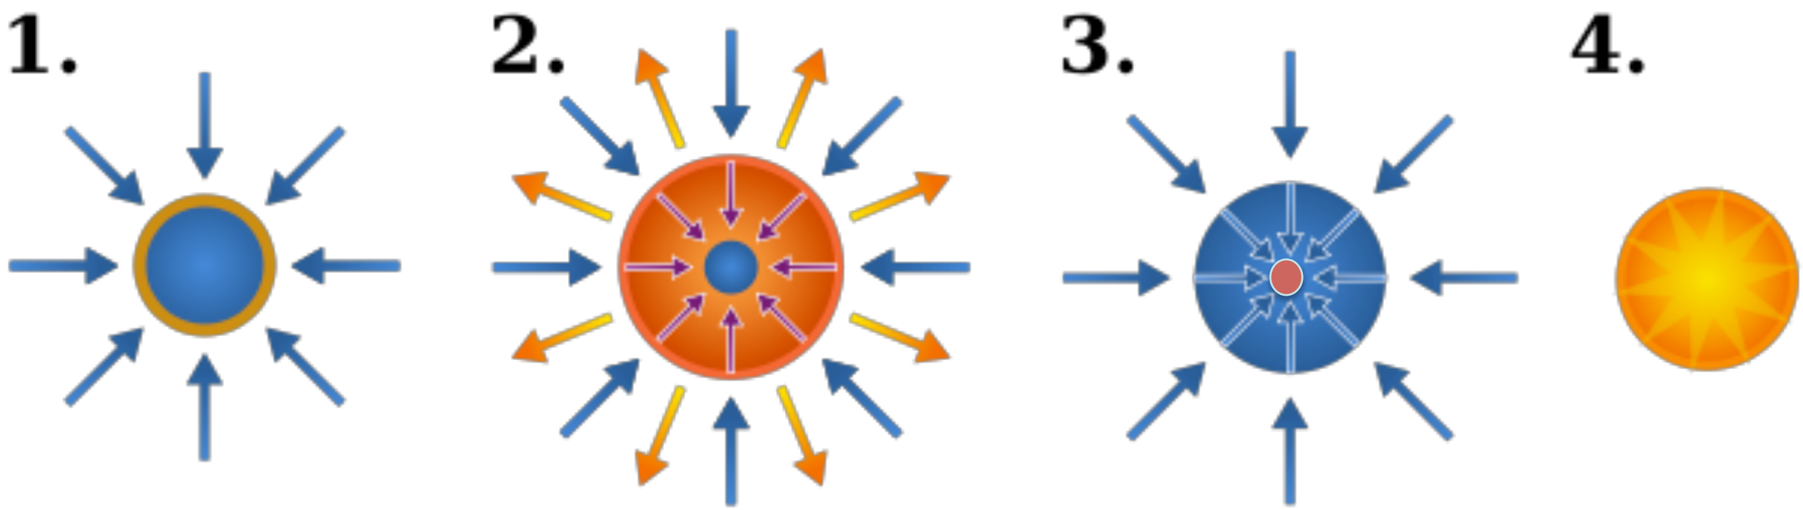

lasers heat surface

surface blows off, fuel compresses by rocket effect

fuel heats, fusion begins

fusion burn

Progress in Magnetic Fusion Energy (MFE) Research

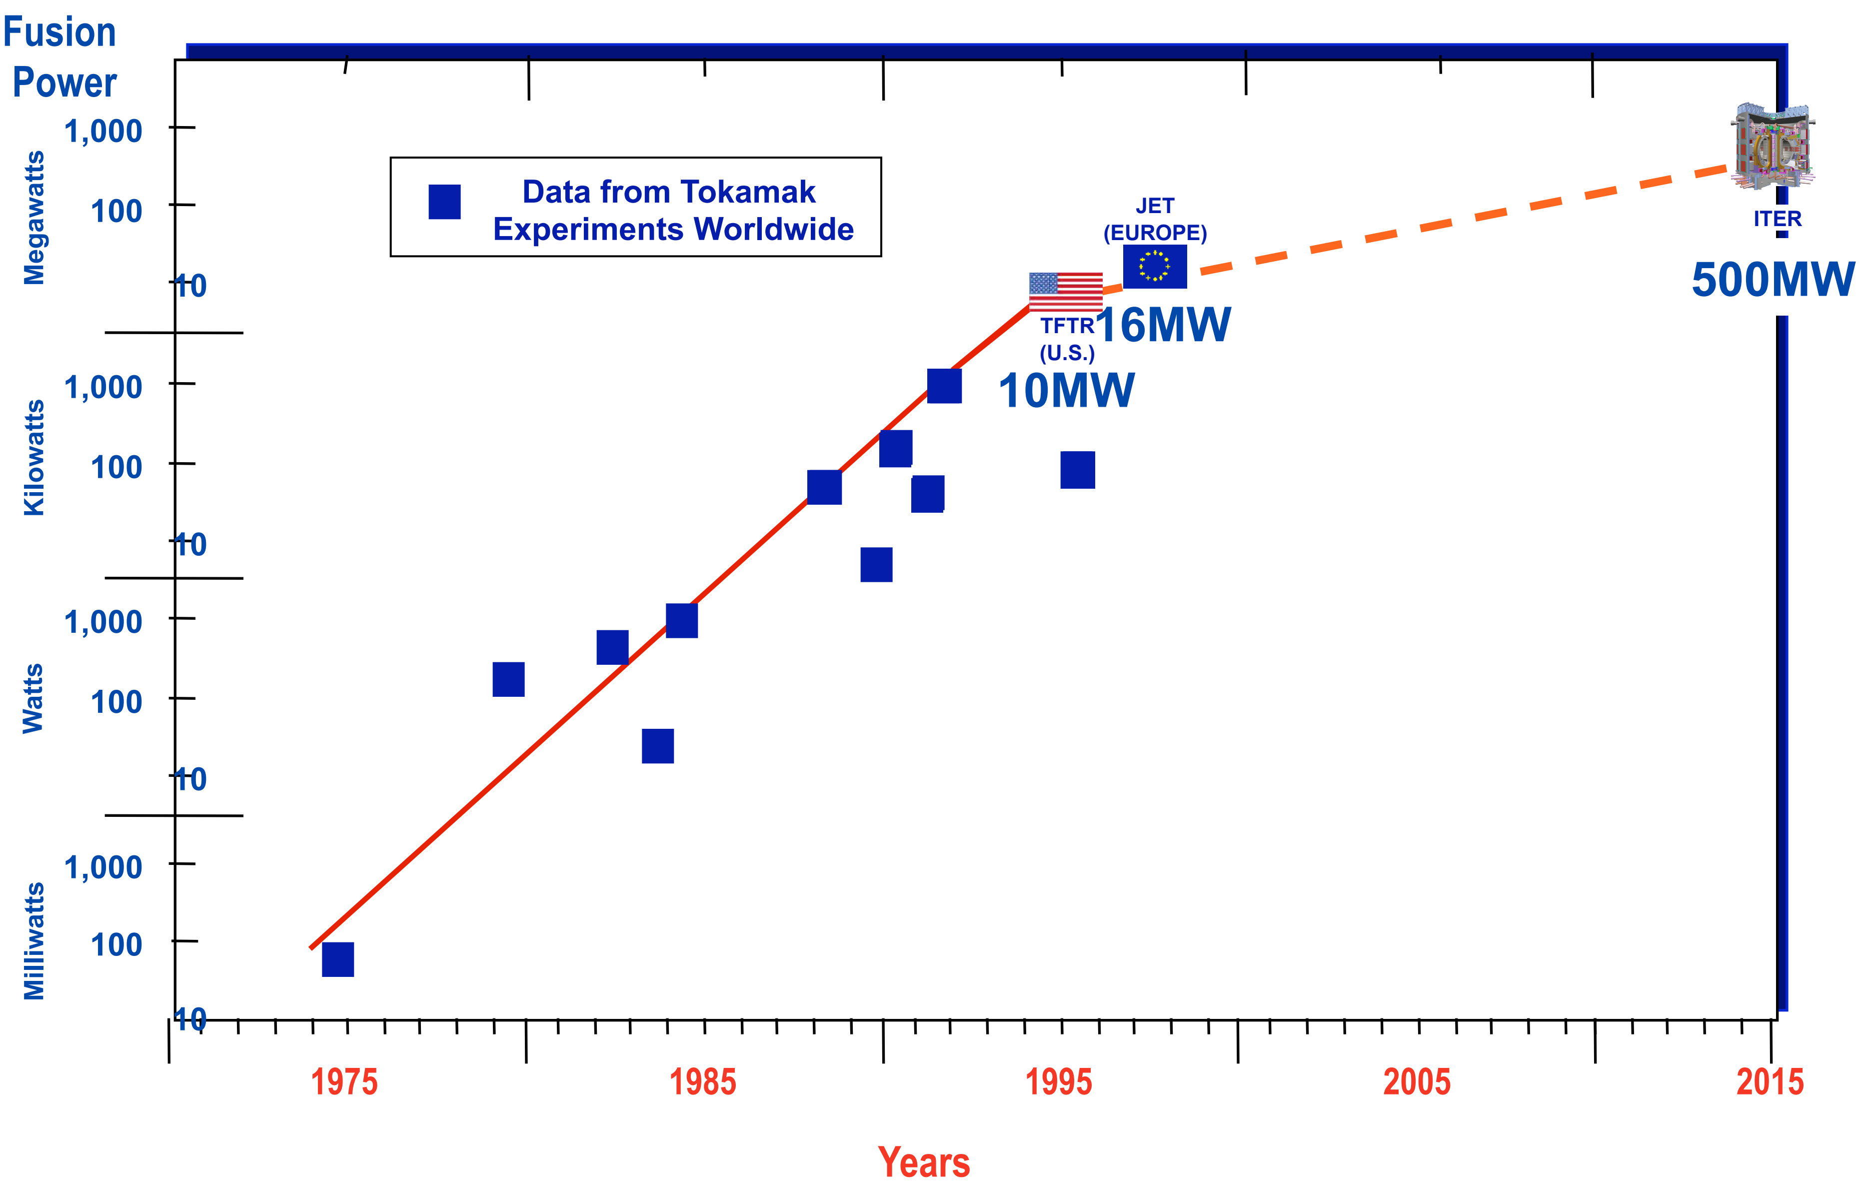

ITER Goal: *Demonstration of the Scientific and Technological Feasibility of Fusion Power*

- ITER *is a dramatic next-step for Magnetic Fusion Energy (MFE)*:

  -- 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

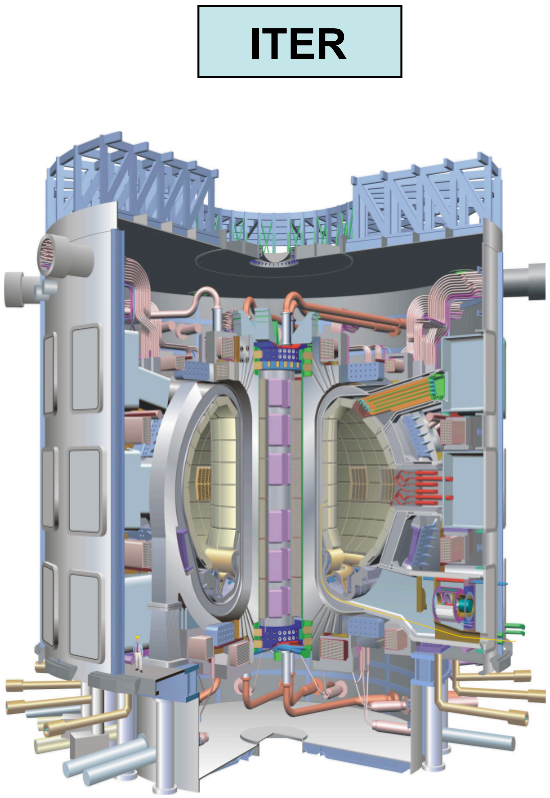

Magnetically confined plasmas in a tokamak are complex and demand integrated analysis

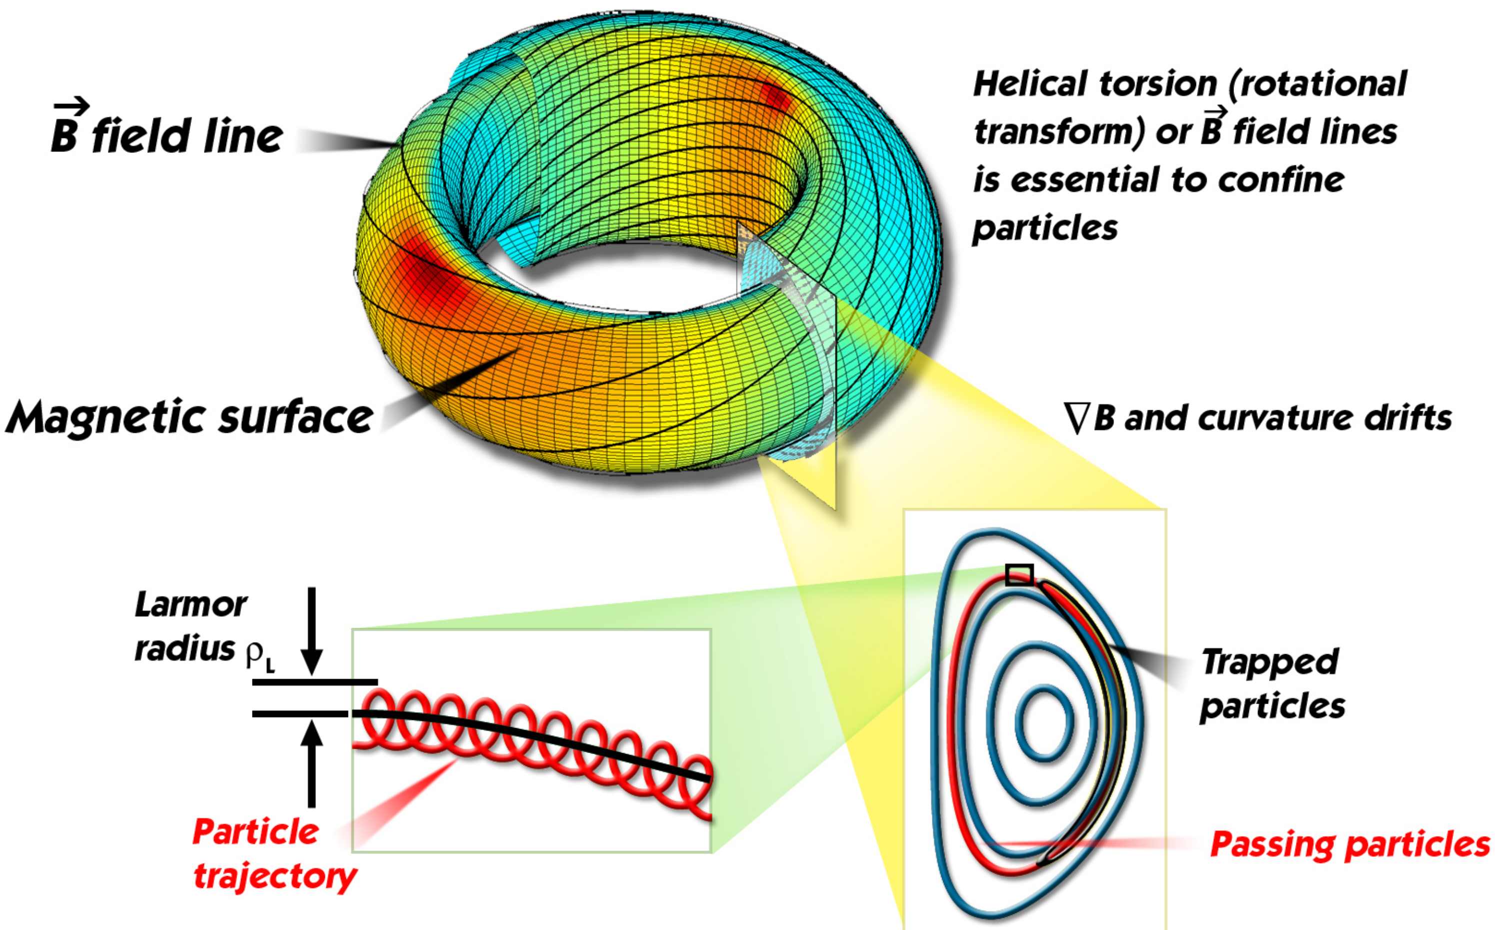

Integrated predictive models must span huge range of spatial & temporal scales
-- major challenges to theory and simulation

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

- **Needed to effectively harvest insights from ITER and to plan for DEMO**

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

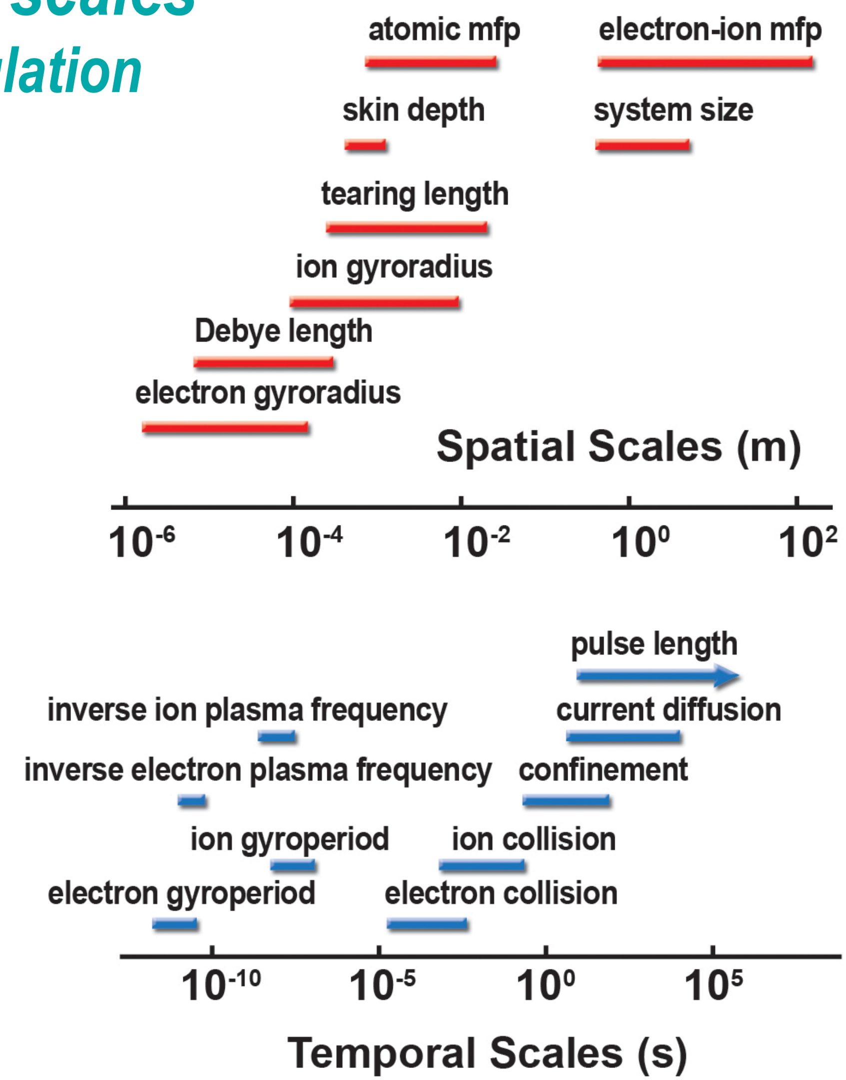

Advanced Scientific Codes --- *"a measure of the state of understanding of natural and engineered systems"* (T. Dunning, 1st SciDAC Director)

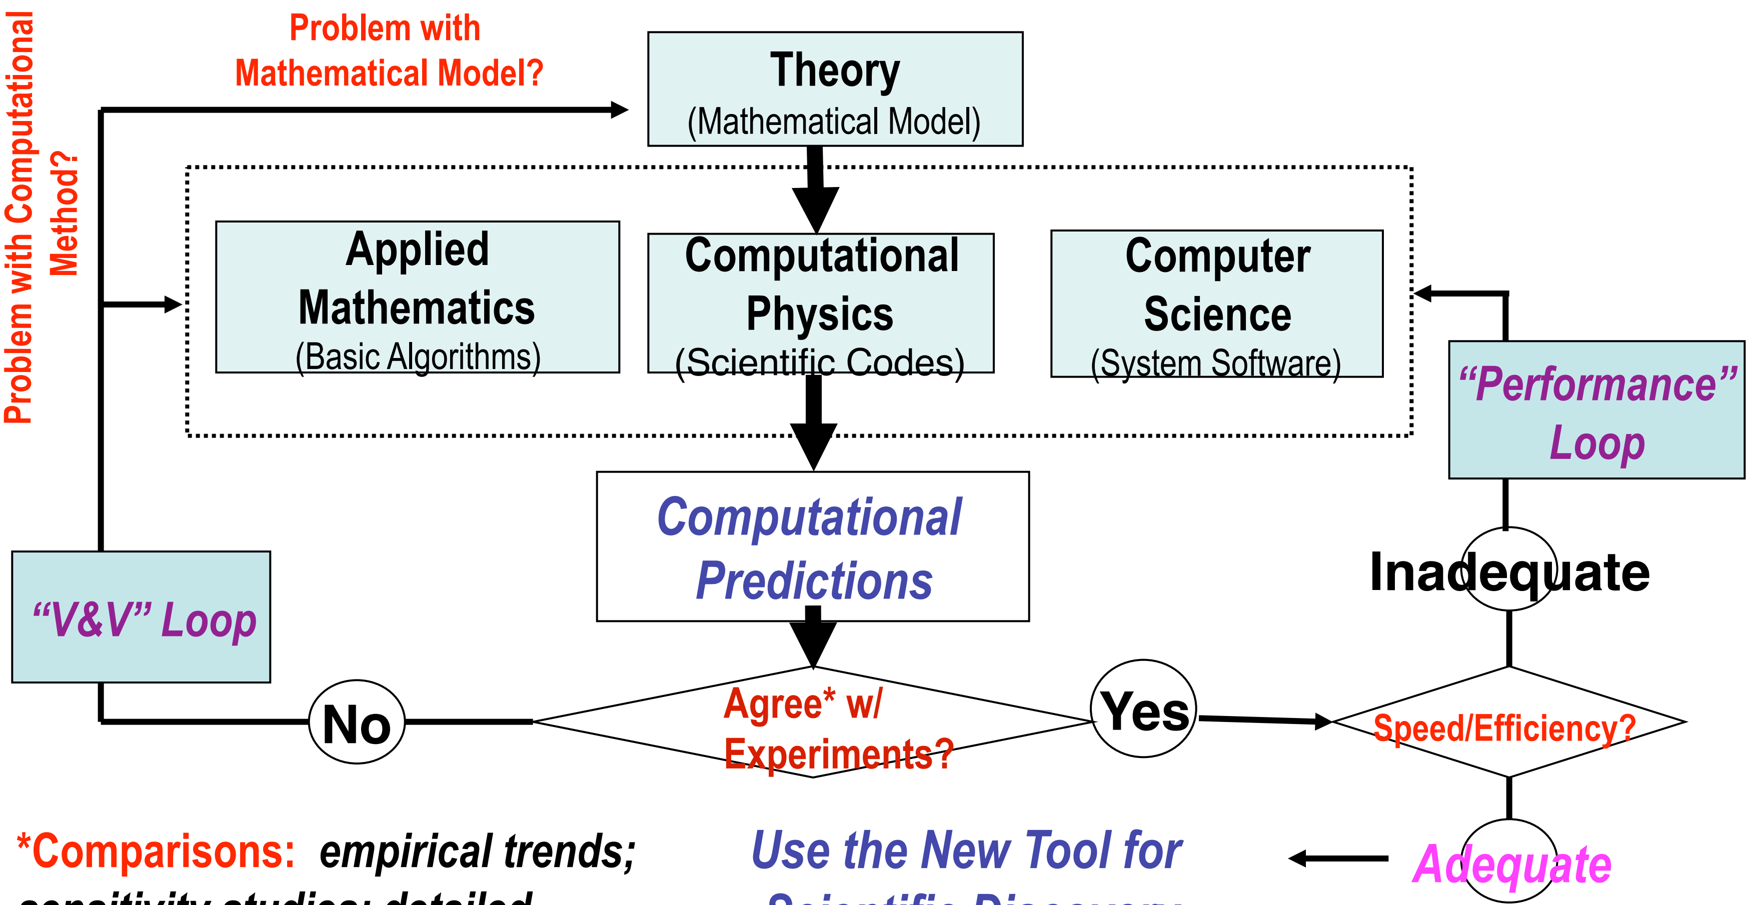

*Comparisons: *empirical trends; sensitivity studies; detailed structure (spectra, correlation functions, …)*

(Repeat cycle as new phenomena encountered)

Verification, Validation, & Uncertainty Quantification Challenges

- **Establishing the physics fidelity of modern plasma science simulation tools demands proper Verification & Validation (V&V) and Uncertainty Quantification (UQ) -- *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.

--- also requires*:* (1) comparisons with theoretical predictions; and (2) cross-code benchmarking (codes based on different mathematical formulations/algorithms but targeting the same generic physics)

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

--- also requires: (1) deployment of modern diagnostics; and (2) application of "synthetic diagnostics" in advanced simulations

• **Uncertainty Quantification is the quantitative characterization & reduction of uncertainty in applications related to variability of input data/model parameters & uncertainties due to unknown processes or mechanisms (e.g., sensitivity analysis)**

Elements of an MFE Integrated Model

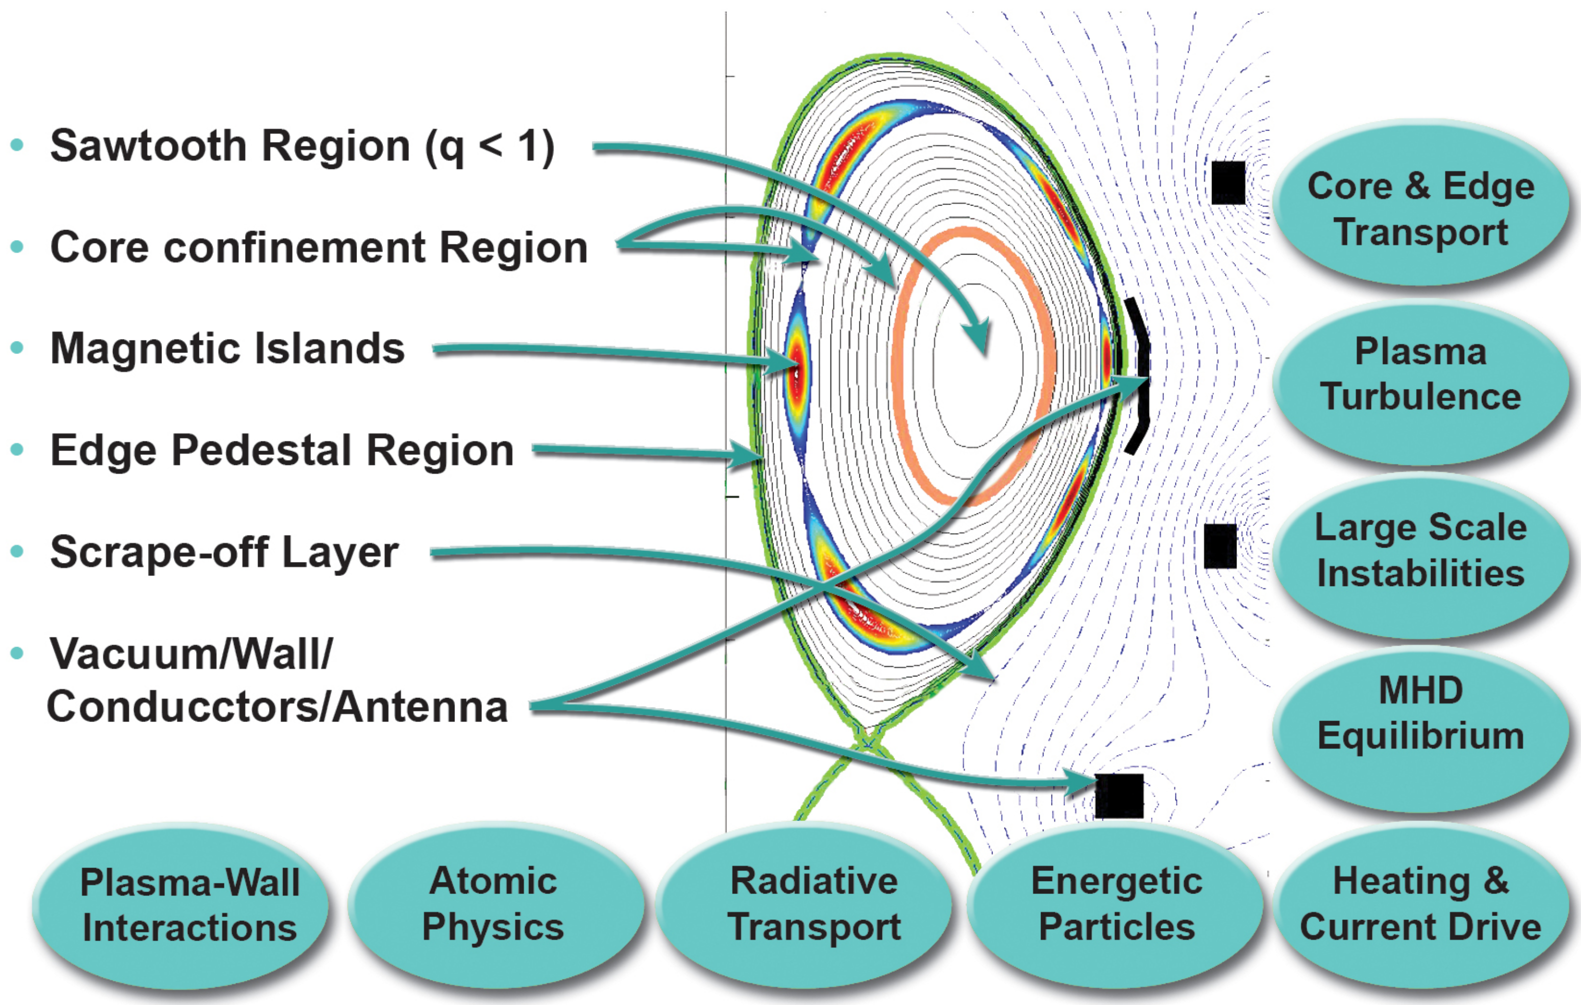

MFE Science Drivers

- ***Disruptions***: **Large-scale macroscopic events leading to rapid termination of plasma discharges**

Goal: Avoid or mitigate because ITER can sustain only a limited number of full-current disruptions with ***severe heat loads, JxB forces, run-away electron generation***

- ***Pedestals***: **Steep spatial gradients whose formation leads to transient heat loads in plasma periphery (divertor region)**

Goal: Predict (i) onset & growth because pedestal height observed to control confinement; (ii) frequency and size of edge localized mode (ELM) crashes to ***mitigate erosion of divertor and plasma-facing components***

- ***Core Profiles***: **Plasma profiles in core confinement region of MFE plasmas.**

Goal*: (i) determine operational limits (e.g., sustainable plasma pressure) and ***optimize plasma performance*** (e.g., fusion yield); and (ii) provide confidence in extrapolating core confinement predictions to future devices

**Requires prediction of temperature, density, current, and rotation profiles in plasma core, including the internal transport barrier region*

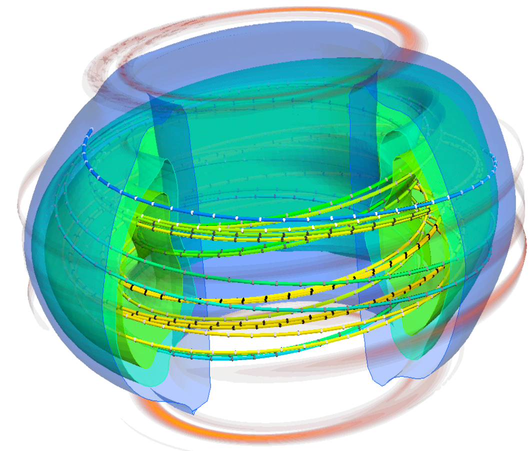

Plasma Disruptions in DIII-D

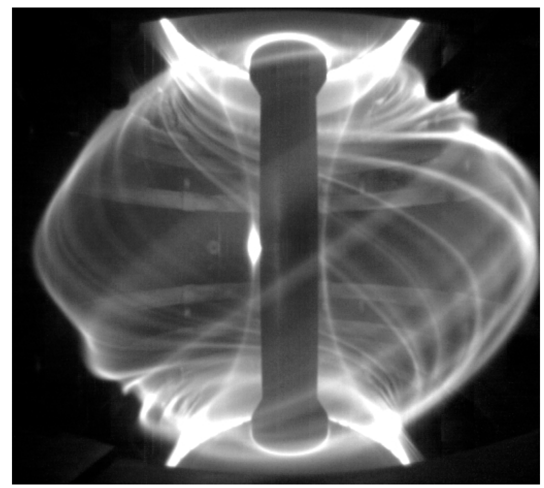

ELMs in MAST

MFE Science Drivers

- ***Wave-Particle Interactions***: **Dynamics between energetic particles and EM waves impacting efficiency of auxiliary heating and the fast-particle confinement of fusion products & supra-thermal particles from RF and energetic beam heating**

Goal: ensure steady-state (long-pulse) performance in burning plasmas such as ITER

-- ***Burning plasma regime is fundamentally new with stronger self-coupling and weaker external control***

- ***Plasma Boundary***: **Region at plasma periphery where complex plasma-materials interactions occur**

Goal: effective modelling of heat loads impacting: (i) divertor design and operational strategies; (ii) erosion of divertor and plasma facing components; and (iii) ***tritium retention and removal***

- ***Whole Device Model***: **Experimentally-validated integrated predictive code to simulate entire MFE system, including dynamics of other 5 Science Drivers**

Goal: provide reliable ***scenario modelling*** for existing and planned (ITER) experimental systems & enable effective design of future devices (DEMO)

Plasma Boundary Layer
Plasma-wall interactions

Magnetically Confined Burning Plasmas:
Unique opportunities for scientific discoveries

- BP/ITER physics elements raise mission-critical questions
  - Unprecedented size
  - Self-heating
  - Large energetic particle population
  - Multiple instabilities with unknown consequences for fast ion confinement

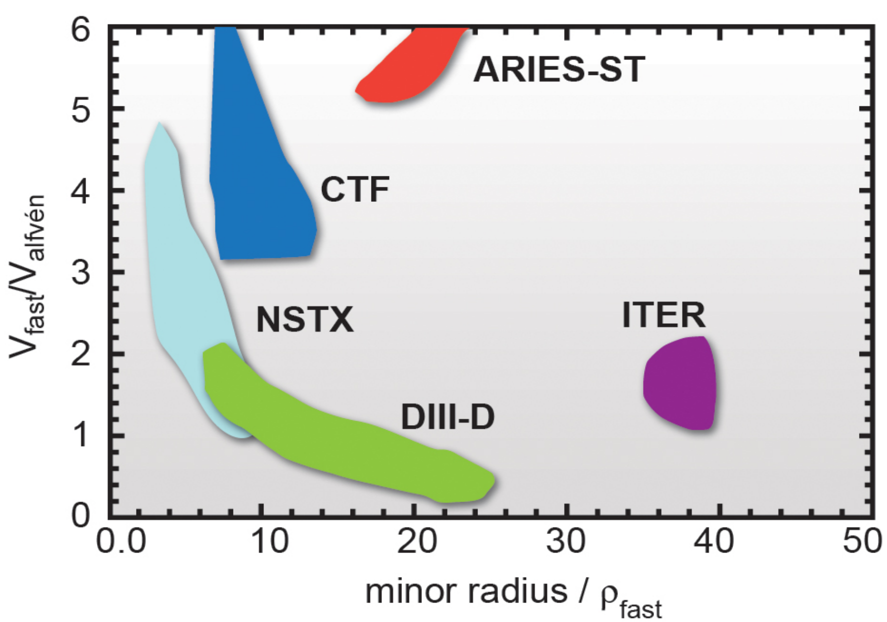

Predicting fast ion confinement: Critical for sustaining a burning plasma

•What is nonlinear interaction between energetic particles and "sea of Alfvén modes?"

•How is transport affected by presence of multiple instabilities?

•How can predictive numerical tools be properly validated?

•What scale of computational resources will be needed to answer BP/ITER mission-relevant questions?

Microturbulence in Fusion Plasmas

- Primary mechanism for cross-field transport in magnetically confined plasmas
  - Size and cost of a fusion reactor determined by balance between particle and energy confinement and fusion self-heating rates
- Challenge: complex multi-scale nonlinear problem
  - Large time and spatial scale separations similar to fluid turbulence (CFD)
  - Self-consistent accounting for electromagnetic fields: many-body problem
  - Strong nonlinear wave-particle interactions: kinetic dynamics
  - Microinstabilities driving turbulence require realistic representation of spatial inhomogeneities together with complex confining EM fields

PROBLEM DESCRIPTION: Particle-in-cell Simulation of Plasma Microturbulence

- **Key Issue**: confinement of high temperature plasmas by magnetic fields in 3D geometry

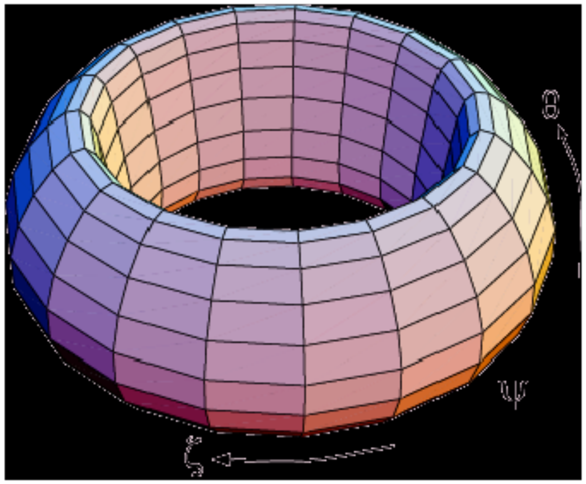

- Pressure gradients drives instabilities producing loss of confinement due to turbulent transport

- Plasma turbulence is nonlinear, chaotic, 5-D problem

- Particle-in-cell simulation

→distribution function - integrate along characteristics with particles advanced in parallel

→interaction - self-consistent EM fields

Particle Simulation of the Boltzmann-Maxwell System

- **The Boltzmann equation (*Nonlinear PDE in Lagrangian coordinates*):**

$$\frac{dF}{dt} = \frac{\partial F}{\partial t} + \mathbf{v}\cdot\frac{\partial F}{\partial \mathbf{x}} + \left(\mathbf{E} + \frac{1}{c}\mathbf{v}\times\mathbf{B}\right)\cdot\frac{\partial F}{\partial \mathbf{v}} = C(F).$$

- **"Particle Pushing" (*Linear ODE's*)**

$$\frac{d\mathbf{x}_j}{dt} = \mathbf{v}_j, \qquad \frac{d\mathbf{v}_j}{dt} = \frac{q}{m}\left(\mathbf{E} + \frac{1}{c}\mathbf{v}_j\times\mathbf{B}\right)_{\mathbf{x}_j}.$$

- **Klimontovich-Dupree representation**,

$$F = \sum_{j=1}^{N}\delta(\mathbf{x}-\mathbf{x}_j)\delta(\mathbf{v}-\mathbf{v}_j),$$

- **Poisson's Equation:** [***Linear PDE in Eulerian coordinates (lab frame)***]

$$\nabla^2\phi = -4\pi\sum_{\alpha} q_{\alpha}\sum_{j=1}^{N}\delta(\mathbf{x}-\mathbf{x}_{\alpha j})$$

- **Ampere's Law and Faraday's Law** [***Linear PDE's in Eulerian coordinates (lab frame)***]

Particle-in-Cell Simulations

• Early attempts [*Buneman (1959); Dawson (1962)]*

• Finite-Size Particles and Particle-in-Cell Simulation [*Dawson et al. (1968) and Birdsall et al. (1968)*]

- *Coulomb potential modified by Debye shielding*

- *Short-range forces within Debye sphere ignored*

- ***Point particles replaced by finite sized particles -- uniformly charged spheres of Debye-length radius***

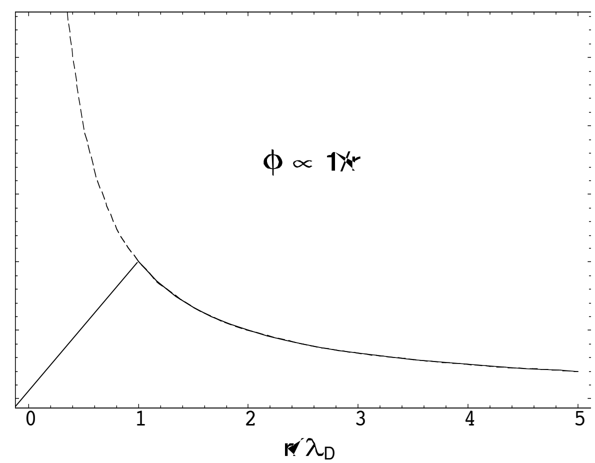

• Number of calculations for N particles

- ***N^2 for direct interactions and N for PIC***

• Collisions are treated as sub-grid phenomena via Monte-Carlo methods [*Shanny, Dawson & Greene (1976)*]

Gyrokinetic Particle Simulation

- *[W. Lee, PF ('83); JCP ('87)]*

- Gyrophase-averaged Vlasov-Maxwell equations for low frequency microinstabilities.

- Spiral motion of a charged particle is modified as a rotating charged ring subject to guiding center electric and magnetic drift motion as well as parallel acceleration -- ***speeds up computations by 3 to 6 orders of magnitude in time steps and 2 to 3 orders in spatial resolution***

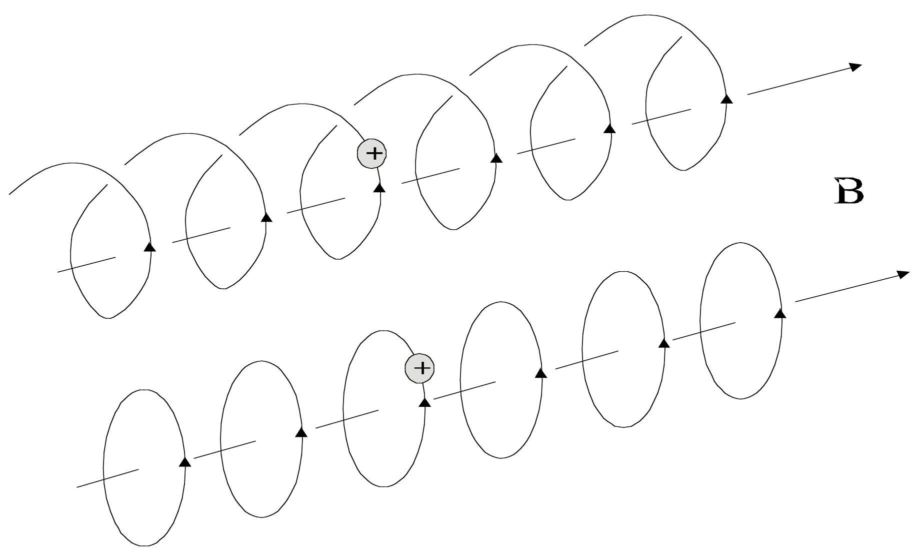

Simulation of Turbulence in Future Ignition-Scale Experiments Require Leadership-Class Computers

- Microturbulence Simulations for range including:
  - a/ρ_i = 400 (JET, largest present lab experiment) through
  - a/ρ_i = 1000 (ITER, ignition experiment)
- Results enabled by use of MPP platforms (e.g., from multi-TF runs @ NERSC)
- These PIC simulations: 1 billion particles, 125M spatial grid points; 7000 time steps
- Such larger-scale simulations indicate transition to more favorable scaling of plasma confinement

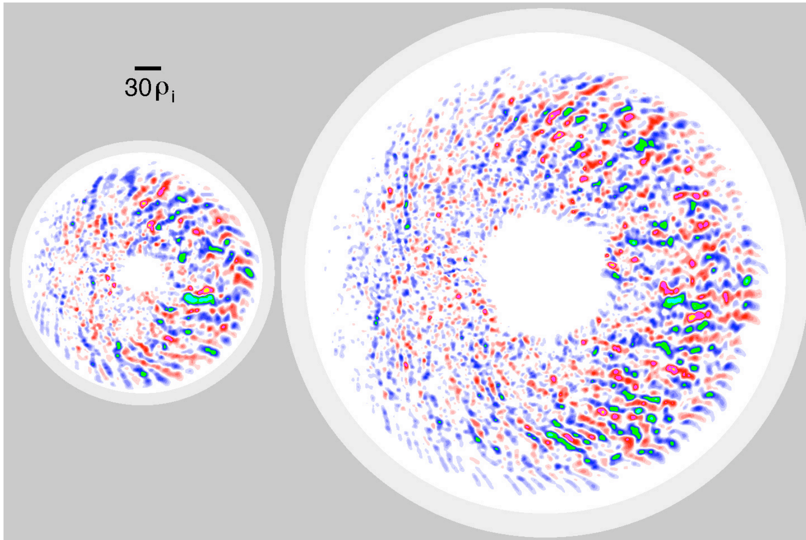

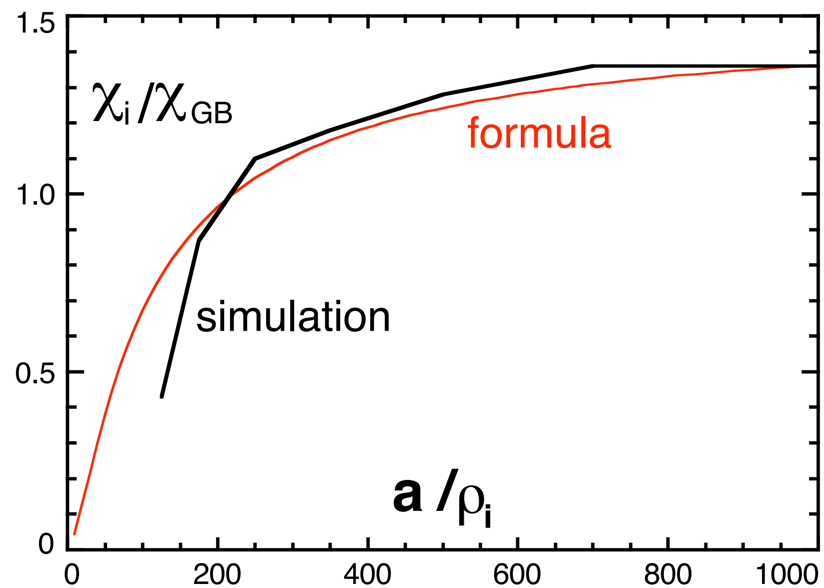

3D Particle Simulation of Plasma Turbulence: Massively Parallel Computation

Turbulent Transport Reduction by Zonal Flows

Princeton Plasma Physics Laboratory
Princeton University

PPPL
PRINCETON PLASMA PHYSICS LABORATORY

High-Resolution Simulations of Loss Mechanisms caused by Microturbulence

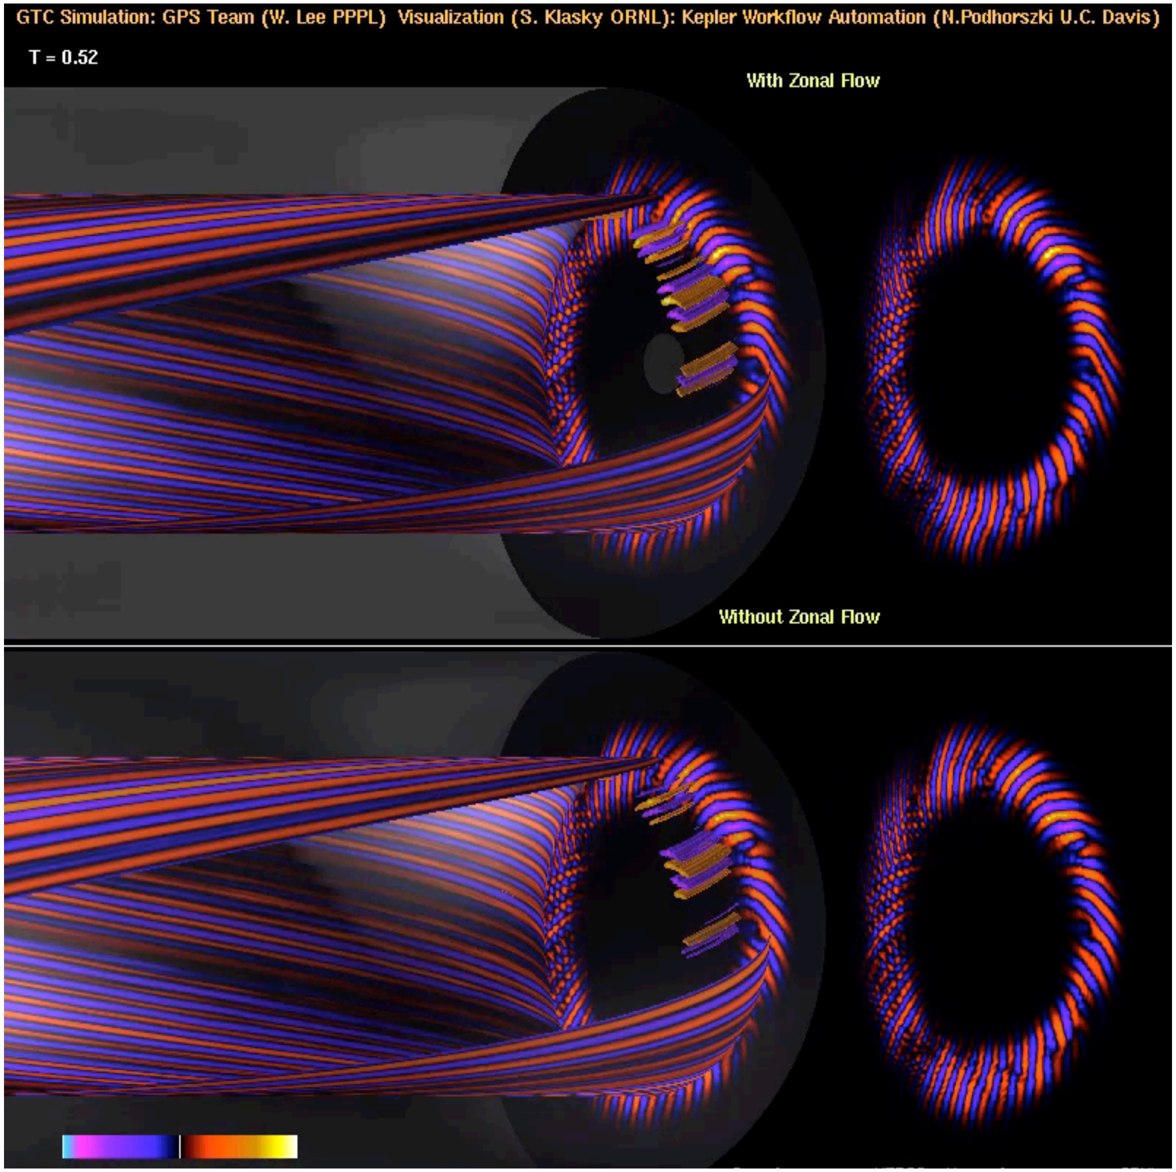

- High-resolution *realistic shaped-cross section toroidal plasma simulations* on leadership class computers
  - [SciDAC GPS Center & ORNL's Jewel Milestone project (W. Wang, et al.)]

- Efficiently generated via "Workflow Automation" -- *automation of data movement, data reduction, data analysis, and data visualization* [SciDAC SDM Center's Kepler workflow project (S. Klasky, et al.)]

Recent LCF-enabled simulations provide new insights into nature of plasma turbulence

Teraflops-to- petaflops computing power have accelerated progress in understanding heat losses caused by plasma turbulence

Multi-scale simulations accounting for fully global 3D geometric complexity of problem (*spanning micro and meso scales*) have been carried out on DOE-SC LCF's

Excellent Scalability of Global PIC Codes *enabled by strong ASCR-FES collaborations* in SciDAC

Exascale-level production runs are needed to enable running codes with even higher physics fidelity and 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)*

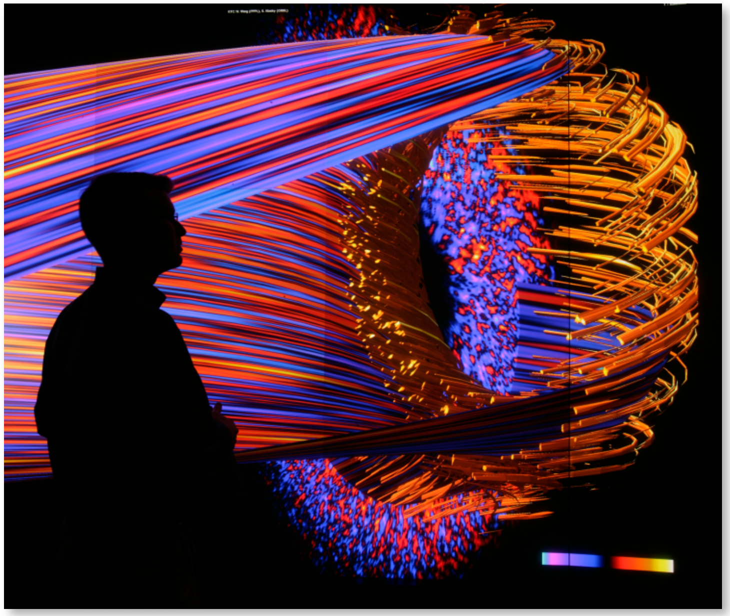

Mission Importance:

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

Scaling GTC-P on IBM BG-P at ALCF

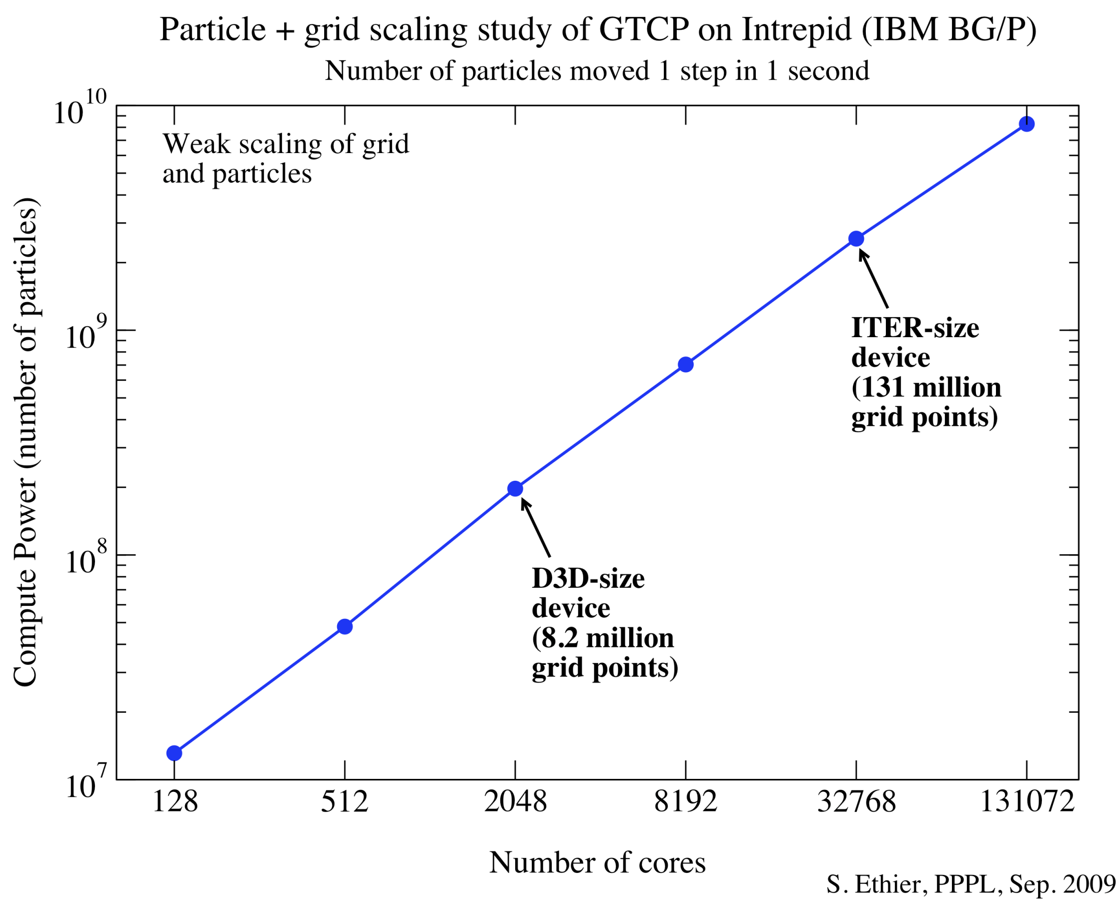

Excellent scalability demonstrated – promising basis for performance on new (multi-petaflop) IBM BG-Q

U.S. Leadership Computing Facility (LCF) Resource Capability

e.g., @ Oak Ridge National Laboratory LCF – "OLCF"

Increased over 750-fold in last 5 years & fusion science applications have been among the largest and most effective consumers

Hardware scaled from single-core through dual-core to quad-core and dual-socket , 12-core SMP nodes	Scaling applications and system software is the biggest challenge

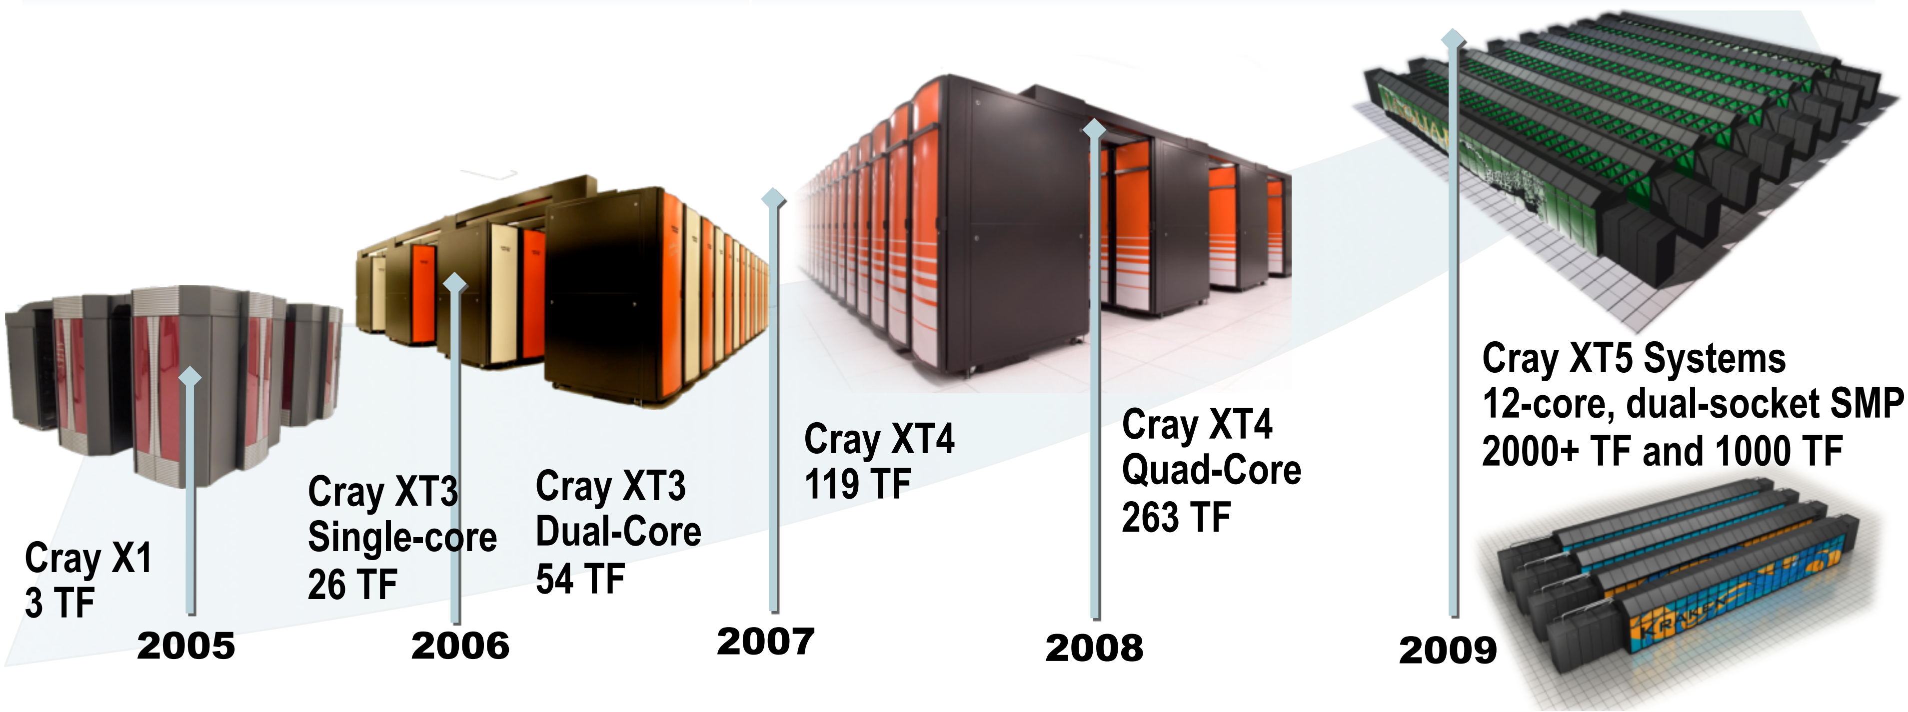

Domain Applications (e.g., FES) Must be Prepared to Exploit *Local Concurrency* to Take Advantage of Coming Hybrid Supercomputing Systems

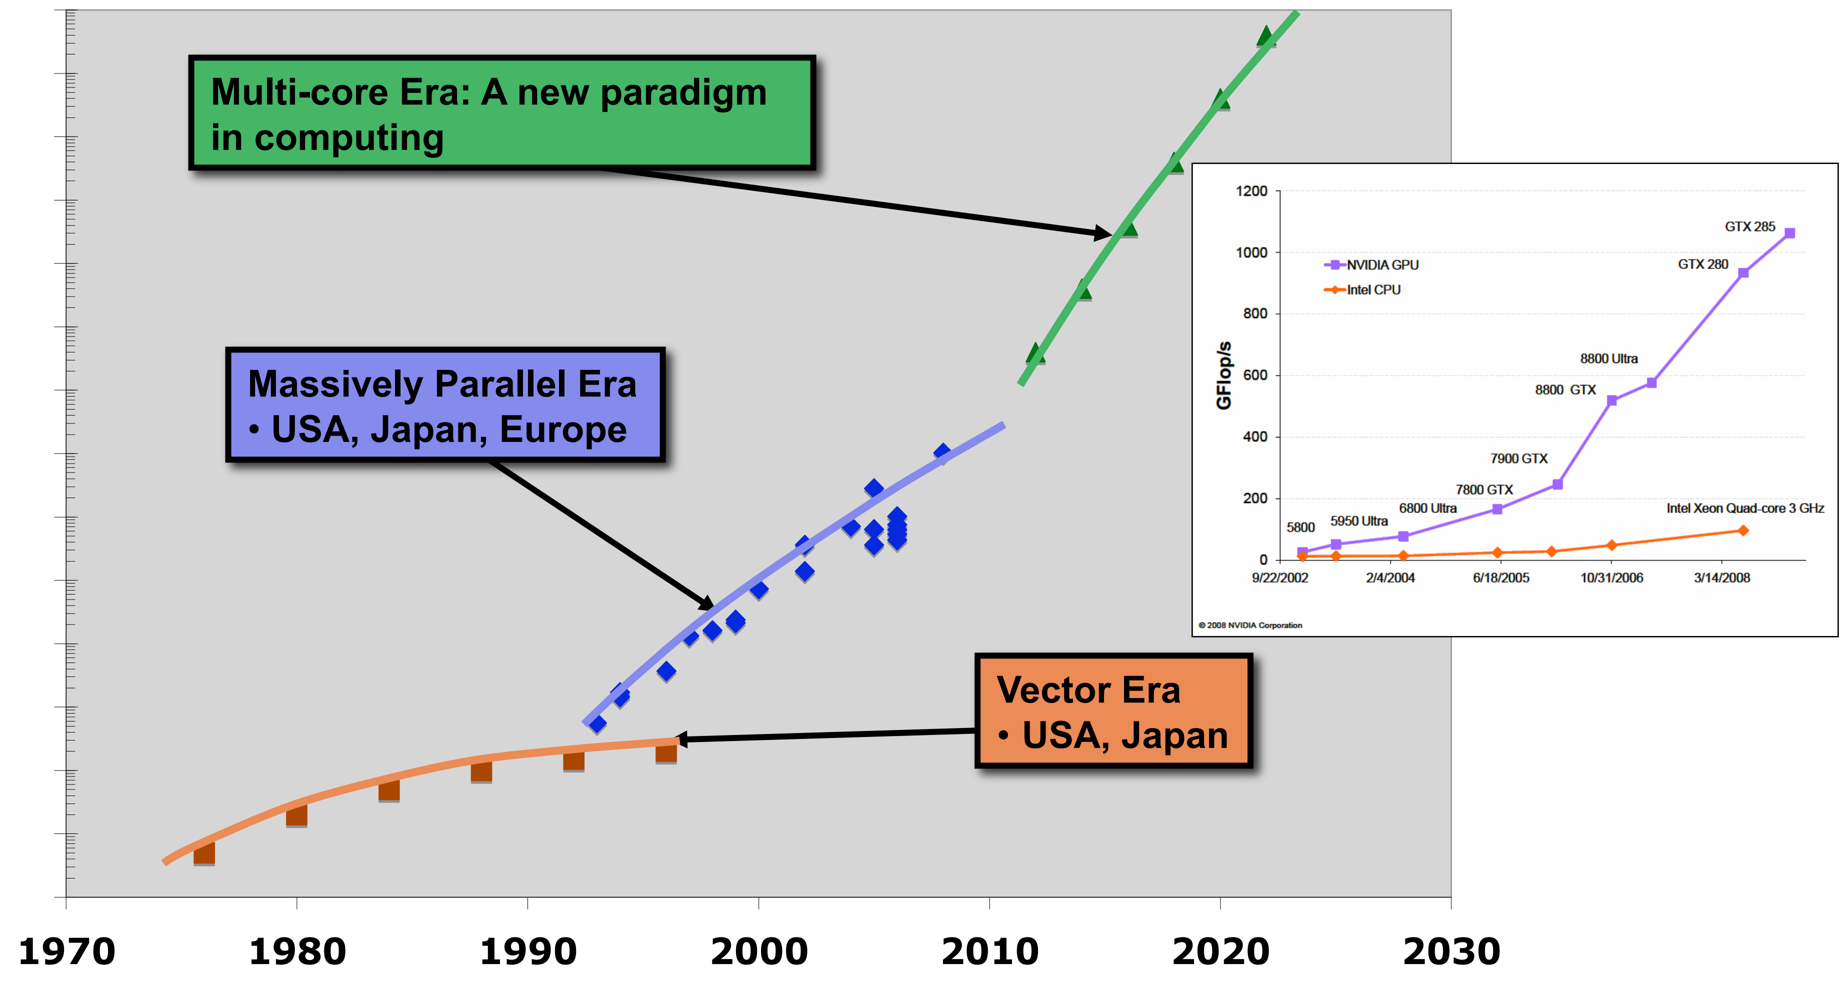

Future HPC Interests very likely include preparing for Hybrid Architectures

General Purpose GPUs, Floating Point Accelerators, etc.

- Large GPU-based systems springing up everywhere
  - *NSF Track 2D in negotiation with Georgia Tech/ORNL*
  - *Japan: “Tsubame” at Tokyo Institute of Technology*
  - *“Orbit” ORNL 100 TF NVIDIA testbed*
  - *Oil and gas industry deploying large clusters*
- Features for computing on GPUs
  - *Added high-performance 64-bit arithmetic*
  - *Critical for a large system*
  - *Larger memories*
  - *Dual copy engines for simultaneous execution and copy*
  - *Development of CUDA and recently announced work with PGI on Fortran CUDA*
- Large and growing pool of people who know how to program accelerators and who will develop tools
  - *Every laptop has a processor and GPU*
    - *Macintosh, PC, Linux ports of CUDA available*
  - *Most computer science programs now teach GPU programming*

300+ accelerated applications listed on NVIDIA’s web site: http://www.nvidia.com/object/cuda_home.html

FSP -- A Strategic *Opportunity* to Accelerate Scientific Progress in FES

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

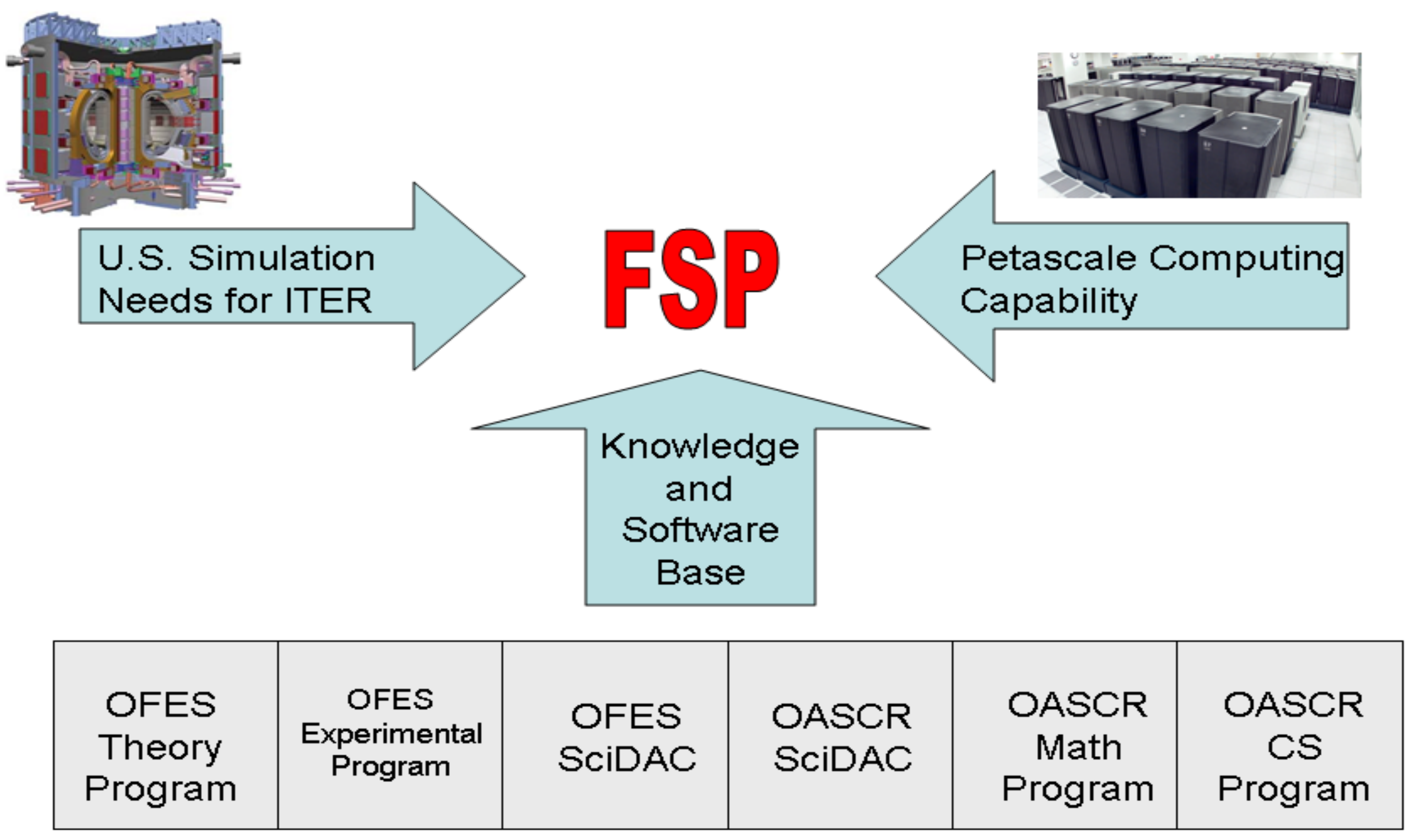

FSP MISSION & VISION

national FSP web-site [*http://www.pppl.gov/fsp/*]

VISION: ***The Fusion Simulation Program (FSP) will enable scientific discovery of important new plasma phenomena with associated understanding that emerges only upon integration. It will provide a predictive integrated simulation capability for magnetically-confined fusion plasmas that are properly validated against experiments in regimes relevant for producing practical fusion energy.***

MISSION: ***The Fusion Simulation Program (FSP) will develop advanced HPC-enabled tools to help accelerate understanding of magnetized toroidal plasmas via efficient integration of multiple, coupled physical processes. This task will engage theory, experiment, and advanced HPC resources to deliver unprecedented capability for harvesting information from experiments and designing new devices with improved performance.***

Moving to the Exascale

- **DoE (SC/NNSA) held series of workshops in 2009-2010 (including FES) to assess the *opportunities and challenges of exascale computing for the advancement of science, technology, and Office of Science missions.***
- **ASCR strategy to address the challenges and deliver on such opportunities involves working with:**
  - **-- domain applications areas such as *FES* to scale applications to each of the new computer systems**
  - **-- LCF's at ORNL & ANL in providing series of increasingly powerful computer systems**

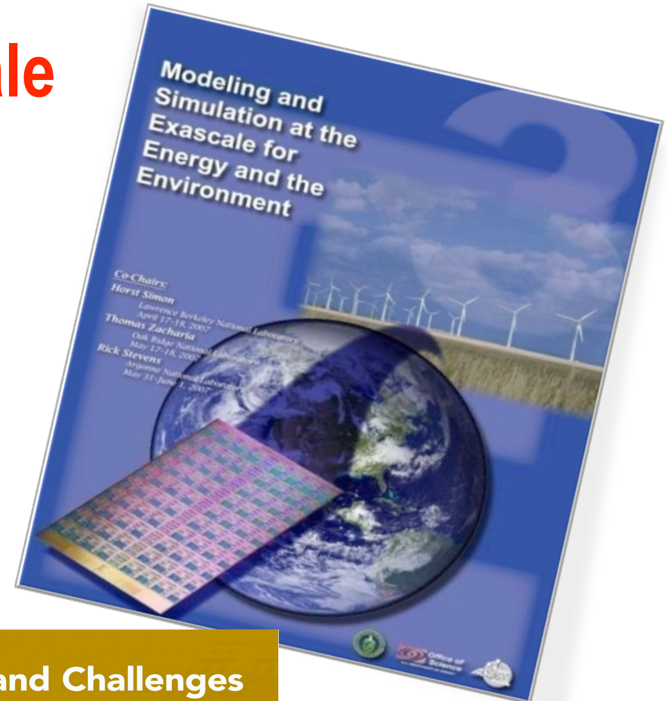

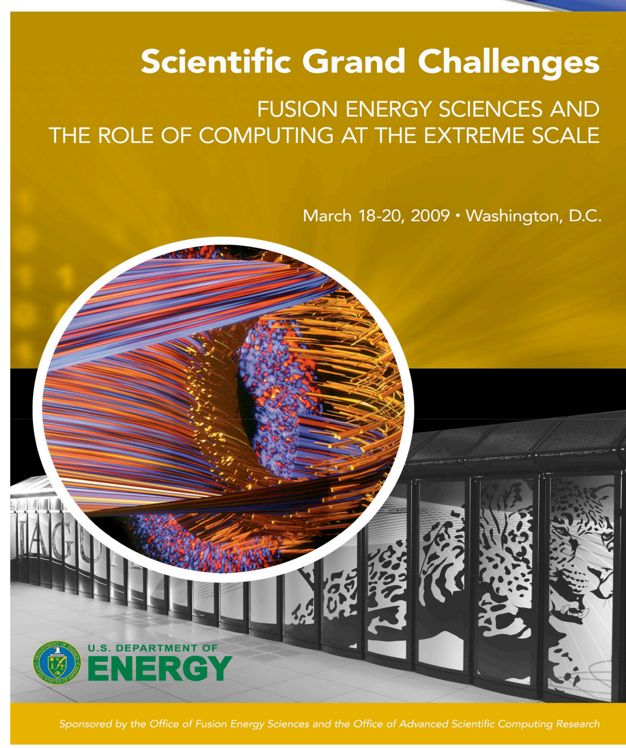


as featured in the recent movie:
"Wall Street - Money Never Sleeps"
FUSION
SIMULATION
PROJECT


“Keep buildin’ that baby star, Doc!”
Shia Labeouf playing “Jake Moore”

- **U. S. Energy Undersecretary** *Steven Koonin:*

3 November 2009 – American Physical Society Meeting, Atlanta,Georgia

"Validated predictive simulation capability is key to advancing fusion science towards energy … the Fusion Simulation Program is a start along this path."

- **U.S. Energy Secretary** *Steven Chu:*

27 September 2010 – "All Hands Meeting" at PPPL, Princeton, NJ

"The world's energy challenge requires a strong continued commitment to plasma and fusion science."

"Progress in fusion has to be grounded in validated predictive understanding: the DoE is clearly interested in your planning and progress for a strong Fusion Simulation Program (FSP)."

➔ **U. S. Energy Undersecretary** *Steven Koonin:*

18 March 2011 -- *House Sub-Committee on Energy & Water Testimony, DC*

"With respect to US investment in ITER … need to make sure our program is best positioned to take advantage of that … so we'll put a lot of money into diagnostics, simulation, and the human capability … Executing simulation is in part about the hardware but it's also about the software and the expertise to meld experiments and observations together with the codes .. and there the U.S. is second to none."

Future Challenges and Opportunities

(1) Energy Goal in FES application domain is to increase availability of clean abundant energy by first moving to a *burning plasma experiment* -- the multi-billion dollar *ITER* facility located in France & involving the collaboration of 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

(2) *HPC Goal* is to harness increasing HPC power ("Moore's law) to ensure timely progress on the scientific grand challenges in FES as described in DoE-SC report (2010) on *"Scientific Grand Challenges: Fusion Energy Sciences and Computing at the Extreme Scale."*

(3) *Interdisciplinary Computational Sciences Goal* is to leverage advances/"lessons learned" from successful U.S. DoE national cross-disciplinary programs such as SciDAC .

Mission of FSP (Fusion Simulation Program):

Accelerate progress in delivering reliable integrated predictive capabilities -- benefiting significantly from access to HPC resources – from petascale to exascale & beyond -- together with a vigorous verification, validation, & uncertainty quantification program