G8 Exascale Software Applications: Fusion Energy

International Exascale Software Project (IESP) Workshop

San Francisco, CA

April, 2011

William Tang
Princeton University and Princeton Plasma Physics Laboratory
Princeton, NJ, USA

Context: G8 Exascale Software Applications Awards

Source; Irene Qualters, NSF Program Director for Cyberinfrastructure

---

- “Enabling Climate Simulation @ Extreme Scale” – US, Japan, France, Canada, Spain
- “Icosahedral-Grid Models for Exascale Earth System Simulations” – Japan, UK, France, Germany, Russia
- ➔ “Nuclear Fusion Simulations @ Exascale” – UK, US, Germany, Japan, France, Russia
- “Using Next-Generation Computers & Algorithms for Modelling Dynamics of Large Biomolecular Systems” -- Japan, UK, France, Germany, Russia
- “Modeling Earthquakes and Earth's Interior based upon Exascale Simulations of Seismic Wave Propagation” – US, Canada, France
- “ExArch: Climate Analytics on Distributed Exascale Data Archives” – UK, US, France, Germany, Canada, Italy

Fusion: an Attractive Energy Source

- Abundant fuel, available to all nations
  - Deuterium and lithium easily available for millions of years
- Environmental advantages
  - No carbon emissions, short-lived radioactivity
- Cannot "blow up or melt down," resistant to terrorist attack
  - Less than a minute's worth of fuel in the chamber
- Low risk of nuclear materials proliferation
  - No fissile materials required
- Compact relative to solar, wind and biomass
  - Modest land usage
- Not subject to daily, seasonal or regional weather variation; no requirement for local CO_2 sequestration
  - Not limited in its application by need for large-scale energy storage nor for long-distance energy transmission
- Fusion is complementary to other attractive energy sources

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

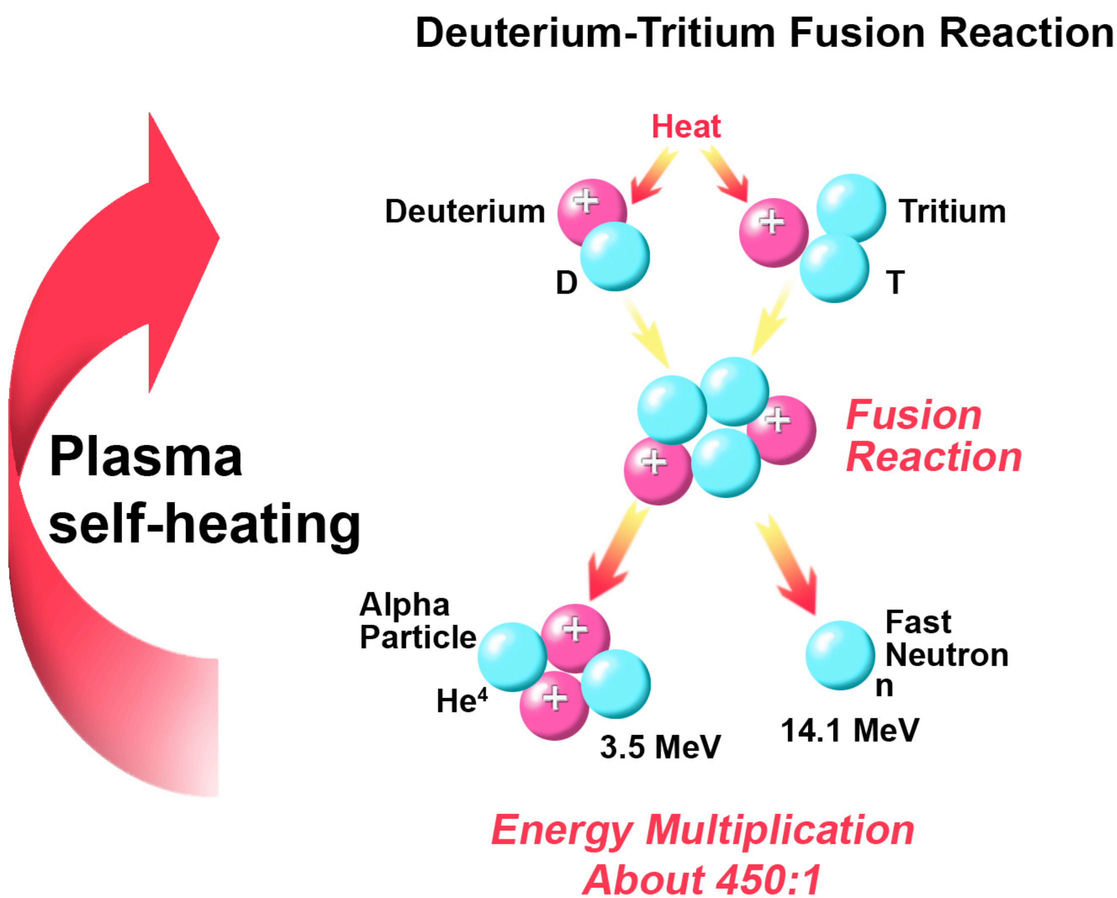

$$D^{+} + T^{+} \rightarrow {}^{4}He^{++}\ (3.5\ MeV) + n^{0}\ (14.1\ MeV)$$

Progress in Magnetic Fusion Research

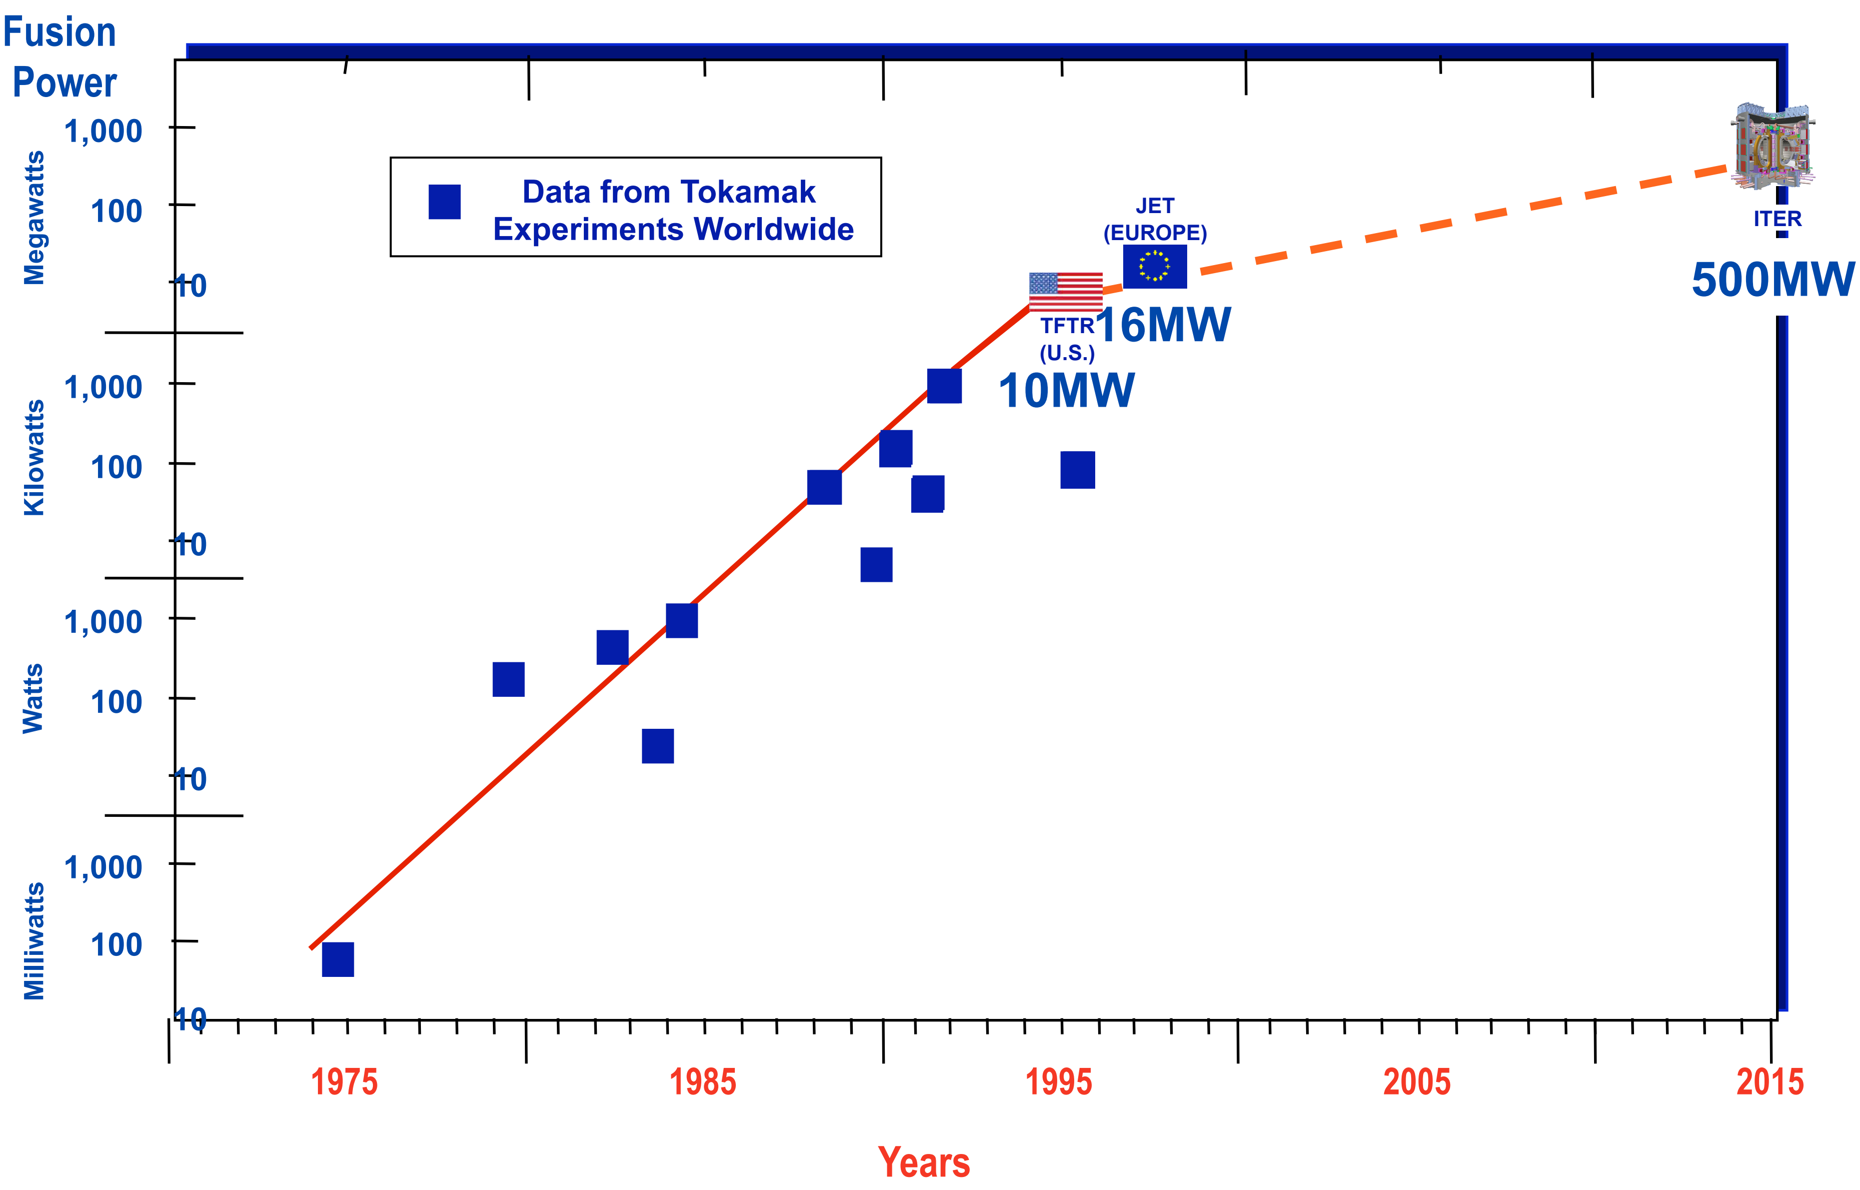

ITER Goal: *Demonstration of the Scientific and Technological Feasibility of Fusion Power*

- 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

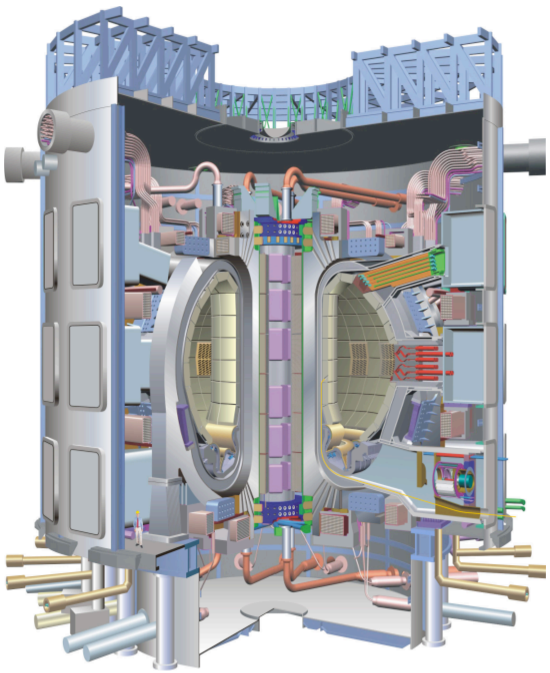

G8 Exascale Project: Fusion Energy Sciences

Some Criteria for Consideration

(1) **Demonstrated need for Exascale**

-- leading FES application codes currently utilize, e.g., Leadership Class Facilities (LCF's) at ORNL (Cray XT5) and ANL (IBM-BGP), demonstrating scalability of key physics with increased computing capability

(2) **Significant Scientific Impact:** *[identified: (i) "Grand Challenges in FES & Computing @ Extreme Scale (DoE Report, March 2010); and (ii) "Opportunities & Challenges of Exascale Computing – DoE Advanced Scientific Computing Advisory Committee Report (Fall, 2010)*

Priority Research Directions:

- *high physics fidelity predictive simulation capability for multi-physics, multi-scale FES dynamics*
- *ITER/burning plasmas physics simulation capability*

(3) **Productive Pathway (over 10 years) to Exploitation of Exascale**

-- demonstrated ability to carry out confinement simulations (e.g., turbulence-driven transport) of higher physics fidelity with access to increased computational capability

Exascale Application Domain: Fusion Energy Science

(Figure courtesy of D. Keyes)

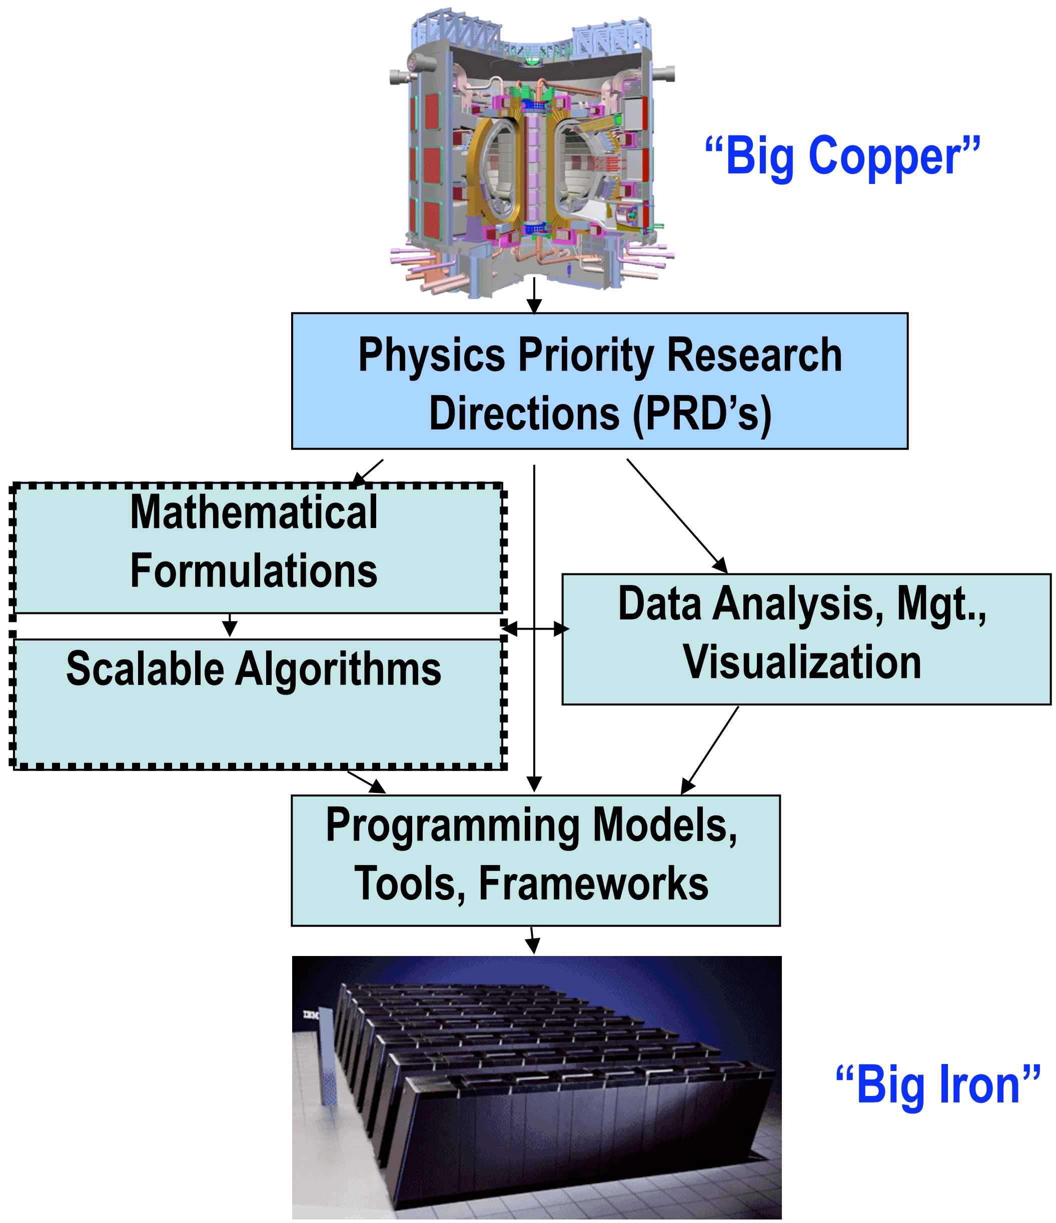

The petascale challenge of simulating magnetic confinement fusion experiments

- **The goal:** predictive simulation of fusion experiments at all time and spatial scales
- **The reality:** currently need at least 4 levels of codes to treat this extremely complex multiscale physics.

current experiment

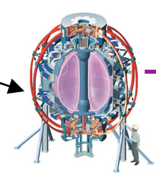

NSTX (@ PPPL)

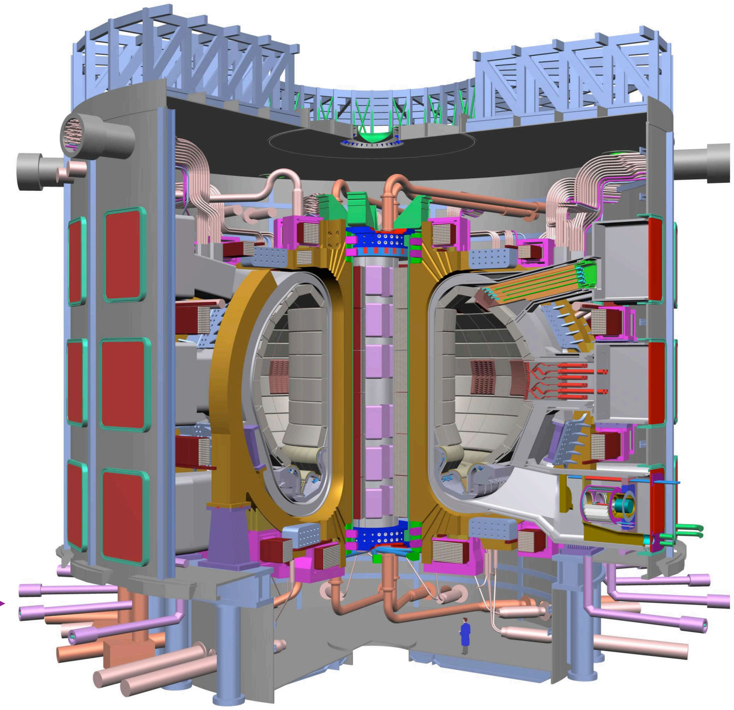

ITER

Elements of Fusion System Integrated Model

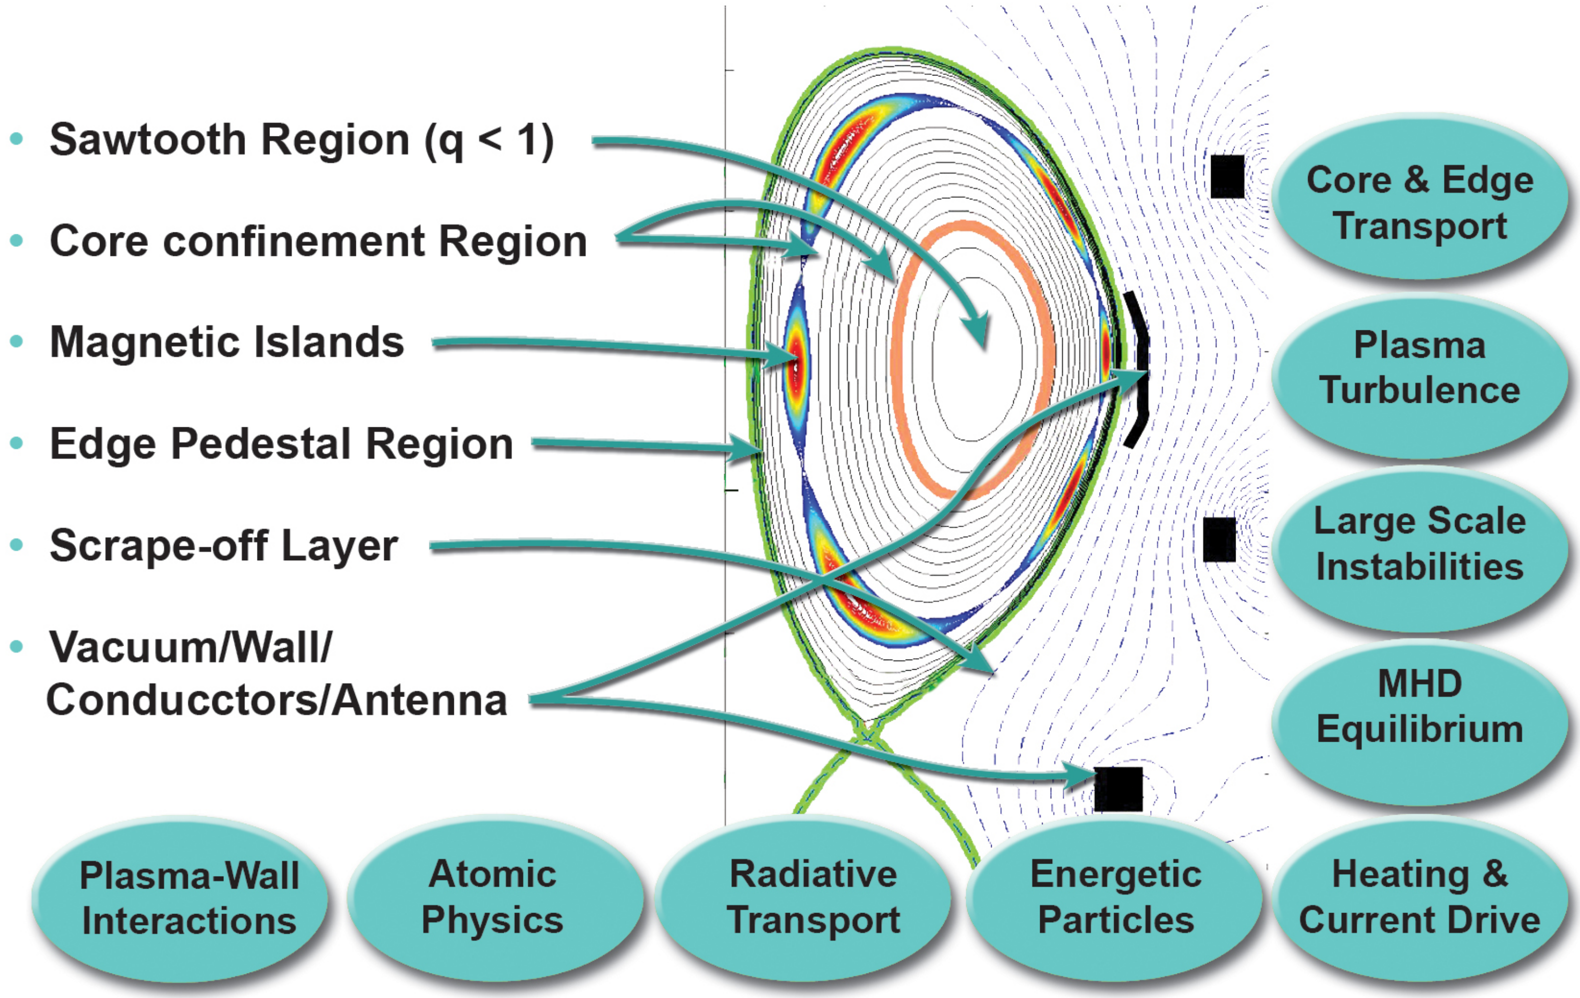

G8 Research Council's Initiative on Multilateral Research Funding

"NuFUSE" (Nuclear Fusion Simulations @ Exascale)

- UK -- Graeme Ackland (Univ. of Edinburgh) – Lead PI (Lead for Plasma Materials Physics) – EPCC, Edinburgh*
- USA -- William Tang (Princeton University/PPPL) – Co-PI (Lead for Plasma Physics) – ALCF, Argonne*
- France -- Xavier Garbet (CEA, Cadarache) – Plasma Physics
- Germany -- Detlev Reiter (Juelich) – Co-PI (Lead for Edge Physics) -- JSC, Juelich*

  & Frank Jenko (IPP, Garching) – Plasma Physics
- Japan -- Taisuke Boku (Tsukuba University) – Co-PI (Lead for Computation) – CCS, Tsukuba*
- Russia -- Boris Chetveruskin (Keldish Institute of Applied Math) – Edge Physics – IHEP, Moscow*

*Supercomputing Centres involved in the NuFUSE Project

NOTE: G8 funding process still in progress in some countries

➔ *1st "Face-to-Face" NuFUSE Project Meeting to be scheduled*

Importance of Turbulence in Fusion Plasmas

- Turbulence is believed to be primary mechanism for cross-field transport in magnetically confined plasmas:
  - ***Size and cost of a fusion reactor determined by particle and energy confinement time and fusion self-heating***
- Plasma turbulence is a complex nonlinear phenomenon:
  - Large time and spatial scale separations similar to fluid turbulence.
  - Self-consistent electromagnetic fields: many-body problem
  - Strong nonlinear wave-particle interactions: kinetic effects.
  - Importance of plasma spatial inhomogeneities, coupled with complex confining magnetic fields, as drivers for microinstabilities and the ensuing plasma turbulence.

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 Leadership Computing Facilities*
- *Excellent Scalability of Global PIC Codes enabled by strong ASCR-FES collaboration in SciDAC projects*
- *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)

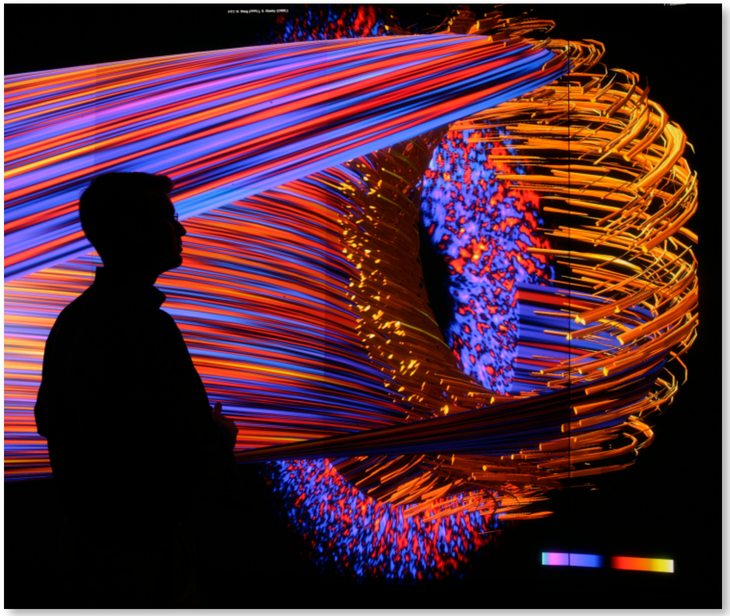

<u>Mission Importance:</u>

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

Scaling of 2D Domain Decomposition on IBM BG-P System

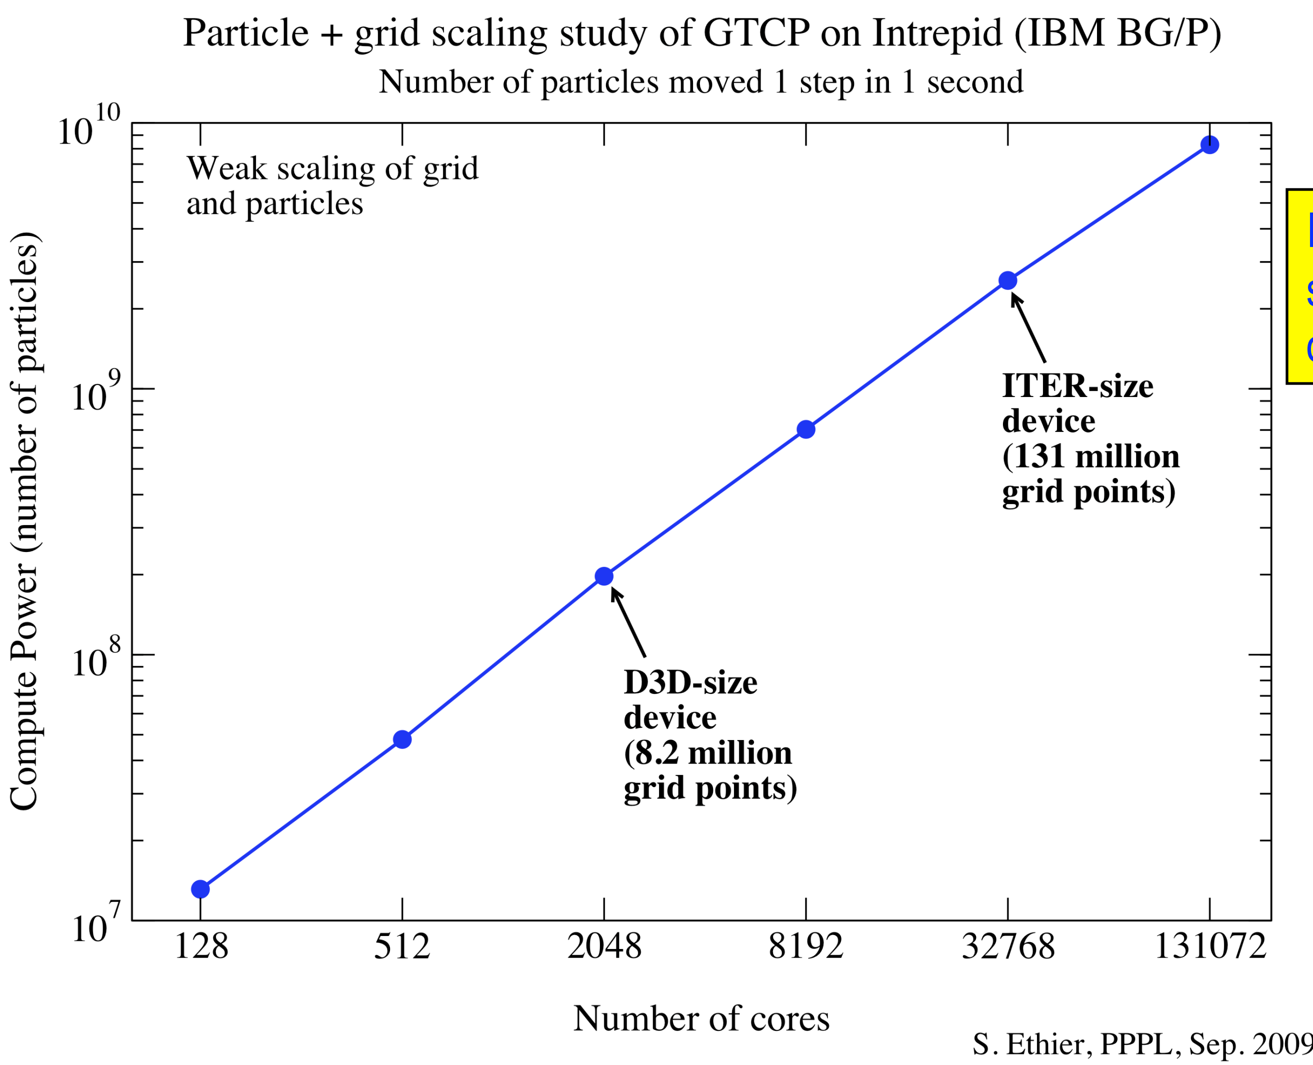

ITER-size simulation with only 512 MB/core!

Some PIC Challenges in Moving toward Exascale

Locality: *Need to improve data locality (e.g., by sorting particles according to their positions on grid)*

-- basic gather-scatter PIC algorithm expends too much time accessing random numbers in a large grid array

Latency: *Need to explore highly multi-threaded algorithms to address memory latency*

Flops vs. Memory: *Need to utilize Flops (cheap) to better utilize Memory (limited & expensive to access)*

-- redo some calculations instead of storing results that are reused somewhere else in the code (e.g. particle-grid positions).

Emerging Architectures: *Need to deploy innovative algorithms within modern codes that demonstrably deliver new science on GPU and GPU-CPU hybrid systems*

-- multi-threading within nodes, maximizing locality while minimizing communications, etc.

Large Future Simulations: *Need to likely work with >10 billion grid points and over 100 trillion particles!!*