WDMApp: High-Fidelity Whole Device Model of Magnetically Confined Fusion Plasma

A. Bhattacharjee (PI), C.-S. Chang (Co-PI), J. Dominski, S. Ethier, R. Hagar, G. Hammett, S. Ku, R. Kube, A. Mollen, B. Sturdevant, P. Trivedi, and G. Willkie, *Princeton Plasma Physics Laboratory (PPPL), Princeton University*
A. Siegel (Co-PI) and B. Allen*, Argonne National Laboratory*
F. Jenko, A. DiSiena, S. Jahnunen, and G. Merlo, *University of Texas at Austin*
S. Parker, Q. Cai, Y. Chen, and J. Cheng, *University of Colorado-Boulder*
J. Hittinger, T. Haut, L. LoDestro, L. Ricketson, and P. Tranquili, *Lawrence Livermore National Laboratory*
S. Klasky, E. D'Azevedo, A. Gainaru, S. Sreepathi, E. Suchyta, and M. Wolf, *Oak Ridge National Laboratory*
M. Parashar and P. Davis, *Rutgers University*
M. Shephard, and C. Davis, *Rensselaer Polytechnic Institute*
K. Germaschewski, *University of New Hampshire*
M. Adams, *LBNL*
A. Scheinberg, *Jubilee Development*
T. Egebo, Financial Control Officer, *PPPL*
L. Perkins, Project Manager, *PPPL*

EXASCALE COMPUTING PROJECT

www.ExascaleProject.org

Grand Challenge : *Integration* of the knowledge provided by plasma models to understand, predict, and control the performance of fusion experiments

"I think the…21st century will be the century of complexity. We have already discovered the basic laws that govern matter and understand all the normal situations. We don't know how the laws fit together, and what happens under extreme conditions….There is no limit to the complexity we can build using those basic laws."----Stephen Hawking

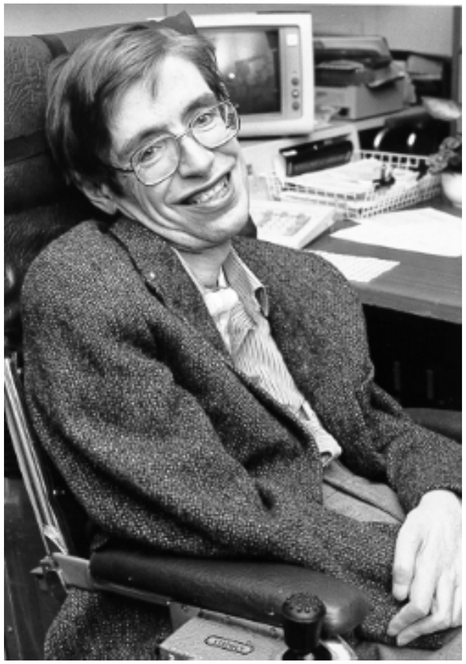

1942-2018

A bit of history: Need for extreme-scale computing resources

Exascale Computing Project (ECP)

- Mission Need: July 2015: President established **National Strategic Computing Initiative (NSCI)** to maximize the benefits of HPC for US economic competitiveness and scientific discovery
- Accelerating delivery of capable exascale computing systems that integrates hardware and software capability to deliver approximately 100 times the performance of current 10 petaflop systems across a range of applications representing government needs
  - To be deployed in 2021 (Aurora at ANL, Frontier at ORNL)

- **ECP was envisioned as a DOE 413.3b project with emphasis on software engineering that can be transformational for science**

ECP builds on the foundation of SciDAC Centers. Without SciDAC, with its blend of advanced computing and physics, strongly coupled with validation, ECP would not be possible.

FES Highest Priority (2015): Advancing Predictive Capability for Whole Device

"Developing an experimentally validated integrated predictive simulation capability that will reduce risk in the design and operation of next-step devices as well as enhance the knowledge gained from ITER."
2015 DOE Report to Congress

10-Year Problem Target of Fusion Exascale Computing Project (ECP) led by PPPL (*Acknowledgements*: S. Prager, M. Zarnstorff, and J. Mandrekas)
To carry out the simulation of non-inductively sustained burning plasma in a high-performance advanced regime, integrating the effects of energetic particles, extended MHD instabilities, and current drive.

Exascale Computing Program: Holistic Approach

Application Development	Software Technology	Hardware Technology	Exascale Systems
Science and mission applications	Scalable software stack	Hardware technology elements	Integrated exascale supercomputers

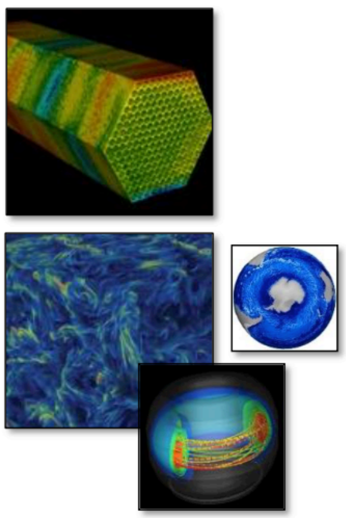

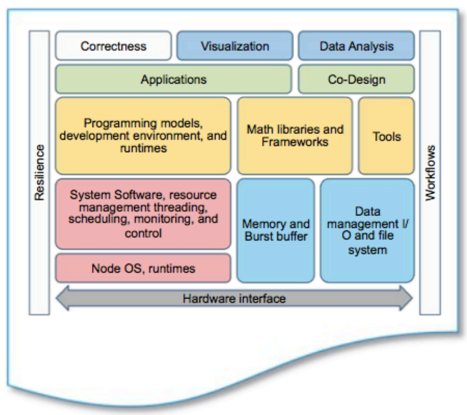

Exascale Applications Cover 6 DOE Strategic Pillars

National security

- Stockpile stewardship

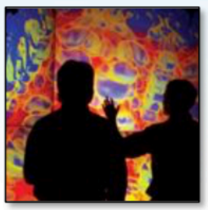

Energy security

- Turbine wind plant efficiency
- Design and commercialization of SMRs
- Nuclear fission and fusion reactor materials design
- Subsurface use for carbon capture, petro extraction, waste disposal
- High-efficiency, low-emission combustion engine and gas turbine design
- Carbon capture and sequestration scaleup
- Biofuel catalyst design

Economic security

- Additive manufacturing of qualifiable metal parts
- Urban planning
- Reliable and efficient planning of the power grid
- Seismic hazard risk assessment

Scientific discovery

- Cosmological probe of the standard model of particle physics
- Validate fundamental laws of nature
- Plasma wakefield accelerator design
- Light source-enabled analysis of protein and molecular structure and design
- Find, predict, and control materials and properties
- Predict and control stable ITER operational performance
- Demystify origin of chemical elements

Earth system

- Accurate regional impact assessments in Earth system models
- Stress-resistant crop analysis and catalytic conversion of biomass-derived alcohols
- Metagenomics for analysis of biogeochemical cycles, climate change, environmental remediation

Health care

- Accelerate and translate cancer research

WDMApp: $23.6M 2016-23

Vision: A High-Performant, First-Principles-Based Whole Device Model

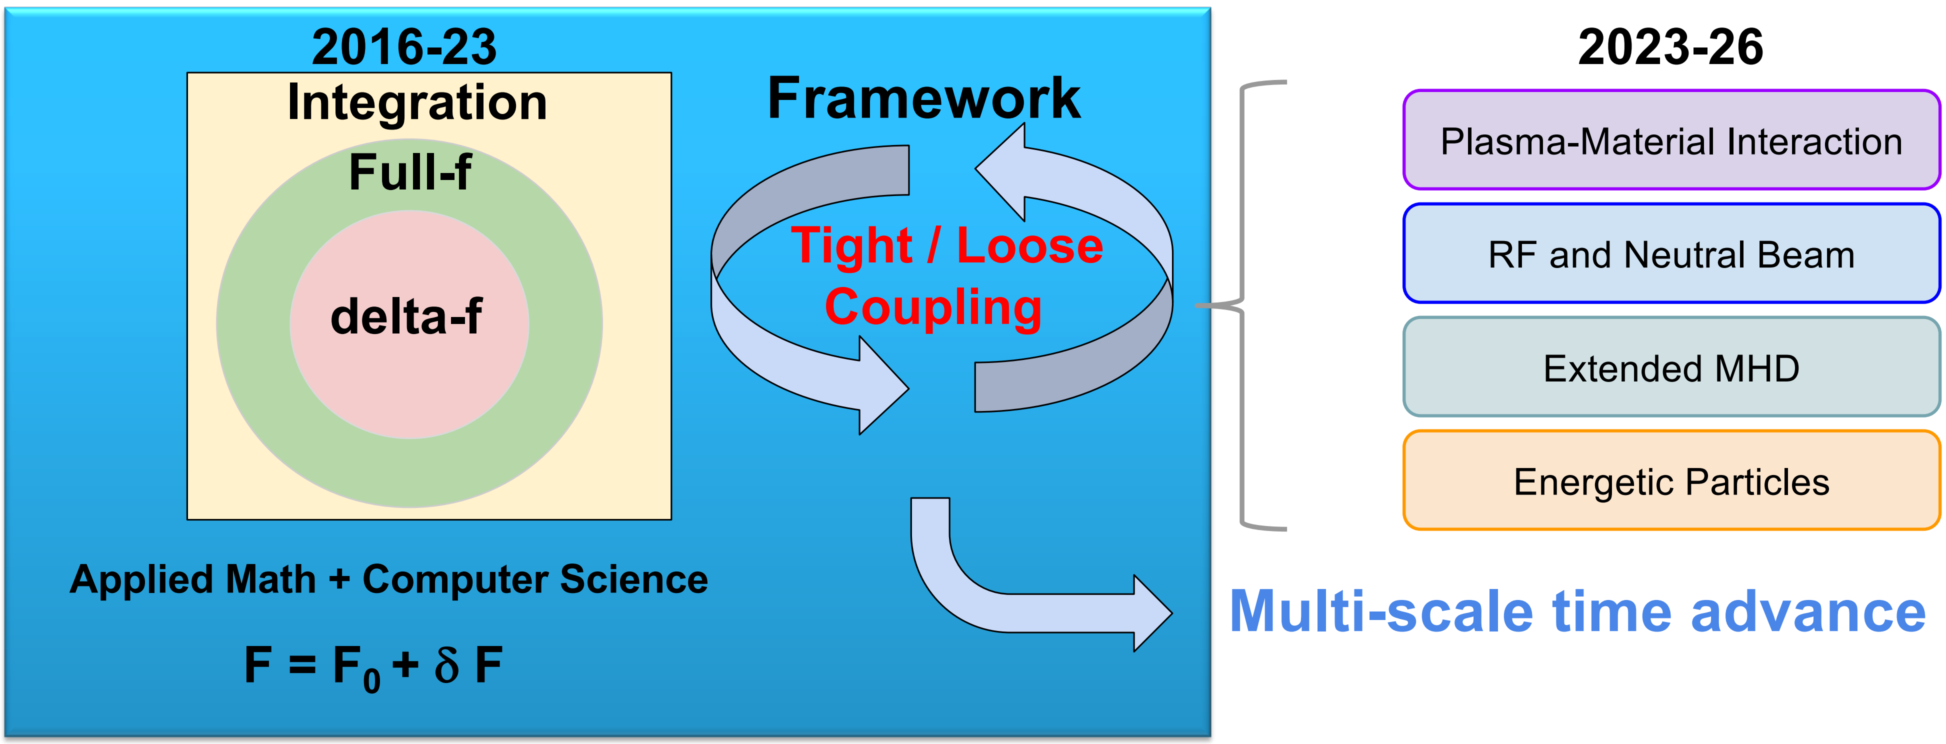

2019 NASEM Report: Committee on a Strategic Plan for U.S. Burning Plasma Research.

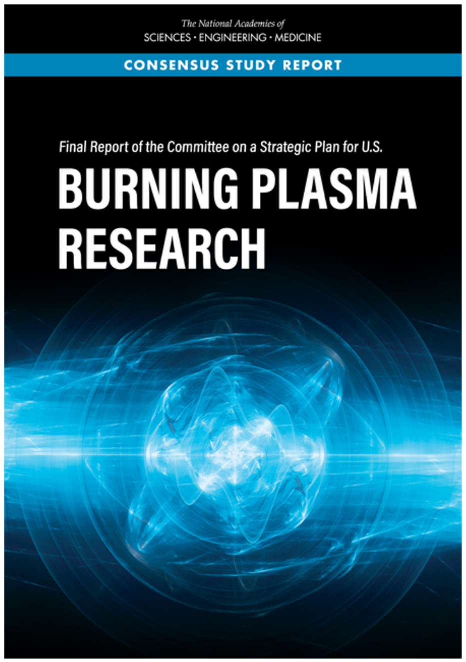

Key Recommendations:

First, the United States should remain an ITER partner as the most cost-effective way to gain experience with a burning plasma at the scale of a power plant.

Second, the United States should start a national program of accompanying research and technology leading to the construction of a compact pilot plant that produces electricity from fusion at the lowest possible capital cost.

The second recommendation motivated the 2021 NASEM Report on "Bringing Fusion to the Power Grid"

Sources

- **Accelerating Fusion Through Integrated Whole Device Modeling**

A. Bhattacharjee, P. Bonoli, A. Boozer, J. Callen, J. Candy, C.-S. Chang, A. Friedman, L. LoDestro, G. Staebler, F. Waelbroeck (*Chair, Theory Coordinating Committee*)

Endorsees: B. Breizman, P. Catto, R. Fitzpatrick, N. Gorelenkov, G. Hammett, D. Hatch, C. Hegna, J. King, P. Morrison, M. Peuschel, A. Ram, A. Reiman, C. Sovinec, D. Spong, X. Tang, R. Waltz, A. White, J. Wright

- **2015 Report of the Workshop on Integrated Simulations for Magnetic Fusion Energy Sciences (Bonoli-Curfman Report)**

  P. Bonoli, *Chair*, and L. Curfman-McInnes, *Co-Chair*

- **2016 Report on the Fusion Energy Sciences Exascale Requirements Review (Chang-Greenwald Report)**

  C.-S. Chang, *Chair*, and M. Greenwald, *Co-Chair*

Questions from the NAS Panel (2018)

- What is the current status of U.S. research that supports advanced scientific computing and modeling for fusion energy?
- What are your views of strategic initiatives that would foster the international development of fusion energy, including participation in the ITER project?
- How advancements in both theoretical understanding and computing infrastructure might broadly strengthen fusion research in the U. S. and promote leadership in the field?

WDM hierarchy: High-fidelity to reduced models

- WDM requires a *fidelity hierarchy* of computational models.
- Pyramid structure

1. Apex of the structure are high-fidelity kinetic models that require leadership-class computing resources, used to calibrate reduced models that form successive layers of the hierarchy.
2. Fast, reduced models are used for large dataset validation of physics basis.
3. Computing infrastructure needed depends on the level of the hierarchy---- from exascale computing (and beyond) at the highest fidelity to capacity computing for reduced models.
4. Necessary to develop a large database at all levels of hierarchy for prediction and uncertainty quantification. Machine learning can be very useful
5. Strong effort in analytical theory needed at all levels.

Advancing Predictive Capability for Whole Device Model is a leading priority for FES community

From the *"Final Report of the Committee on a Strategic Plan for US Burning Plasma Science Research (2019)" by the US National Academy of Sciences* :

"The creation of the DOE Exascale Computing Project in fiscal year 2017 has provided the impetus for the realization of a high-fidelity whole device model of fusion plasma applicable to a high performance advanced tokamak regime, integrating the effects of turbulent transport, plasma collisions, and neutral particles, energetic particles, plasma-material interactions, as well as heating and current drive…Whole device modeling holds the promise of being a powerful predictive tool for current and future fusion devices that will access hitherto unrealized plasma regimes, and has the potential to produce scientific discoveries of new and emergent behavior of burning plasmas that have been so far studied piecemeal…. The project will be developed in a computational ecosystem that brings together plasma physicists with applied mathematicians, computer scientists, and other application scientists using a diverse range of software technologies and several co-design centers."

FES Community Consensus on ECP WDMApp

Evolution of Kinetic Modeling Capability with Increasing Computing Power

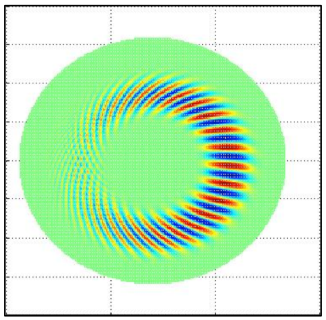

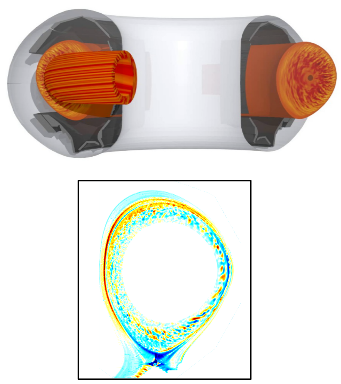

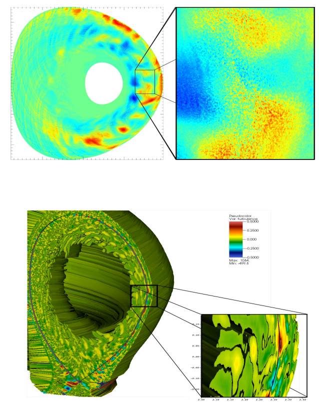

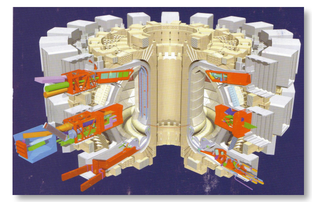

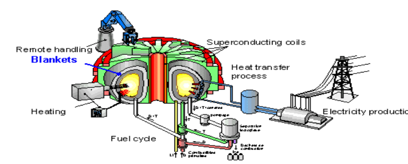

Gigaflops

5-D electrostatic ion physics in simplified circular cylindrical geometry

Teraflops

Core: 5D ion-scale electromagnetic physics in torus

Edge: ion+neutral electrostatic physics in torus

Petaflops

Core: 5D i+e scale electrostatic physics

Edge: ions + electrons + neutrals, electrostatic physics

Exaflops

Core-edge coupled 5D electromagnetic study of whole-device ITER, including turbulence, transport, large-scale instability, plasma-material interaction, rf heating, and energetic particles

Beyond

6D whole device modeling of all the relevant fusion reactor science with burn control

SciDAC Centers

A partnership between FES and ASCR divisions in the DOE Office of Science. PPPL/Princeton University leads three, and is a co-participant in four others:

- High-Fidelity Boundary Plasma Simulation (C.-S. Chang, *PI*)
- Center for Tokamak Transient Simulations (S. Jardin, *PI*)
- Simulation Center for Runaway Electrons Avoidance and Mitigation (D. Brennan (Princeton University, *PI*; A. Bhattacharjee, *Co-PI*)
- Tokamak Disruption Simulation (X. Tang (LANL), *PI*; W. Wang, *Co-PI*)
- Integrated Simulation of Energetic Particles (Z. Lin (UCI), PI; N. Gorelenkov, *Co-PI*)
- Advanced Tokamak Modeling Environment (J. Candy (GA), *PI*; J. Sachdev, *Co-PI*)
- Multi-Scale Gyrokinetics (D. Hatch (UT Austin), PI; G. Hammett, *Co-PI*)

Coupling the core and edge: first in fusion history

The core evolves more slowly than the edge

Core Turbulence from GENE

Edge Turbulence from XGC

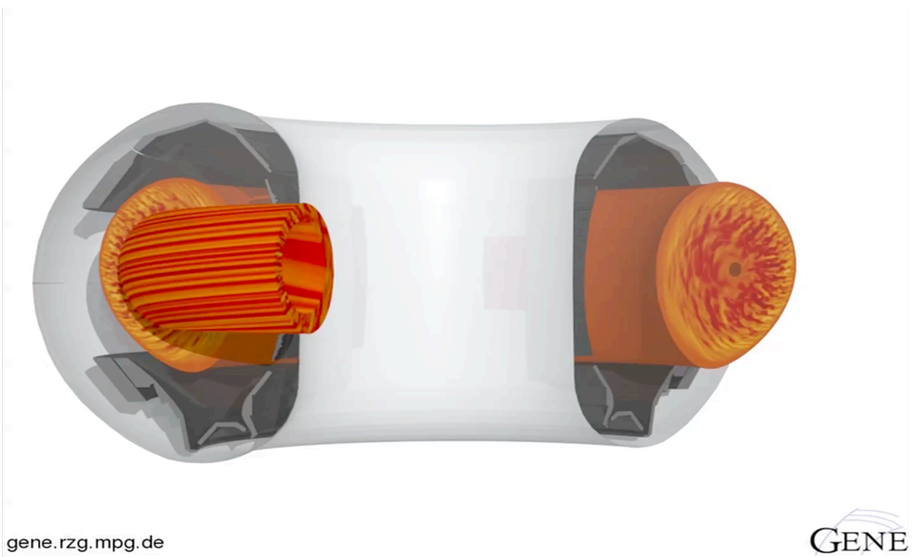

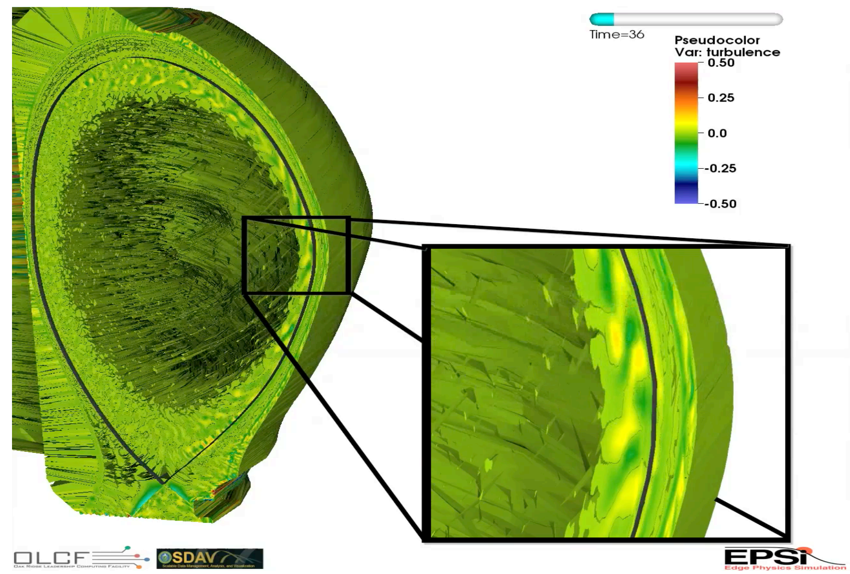

Principal WDMApp Goals

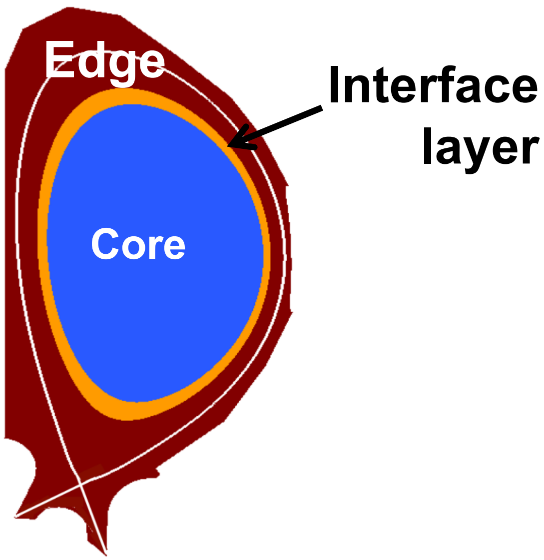

- Demonstration and assessment of WDM **gyrokinetic physics** on experimental transport time-scale in a challenge problem for pedestal formation
- **Figure of Merit (FOM) of >50** for coupled code on exascale platforms, accomplished through algorithmic advancement, performance engineering and hardware improvement (*to be discussed in detail in answer to Q3*)
- Completion of **extensible integration framework** EFFIS 2.0 (End-to-End Framework for Fusion Integrated Simulations 2.0) and demonstration on exascale platform

WDMApp Challenge Problem

High-fidelity simulation of a whole-device burning plasma (specifically, ITER with full plasma current) operating in “high-mode” (H-mode), and prediction of the plasma pressure “pedestal” shape (height and width)

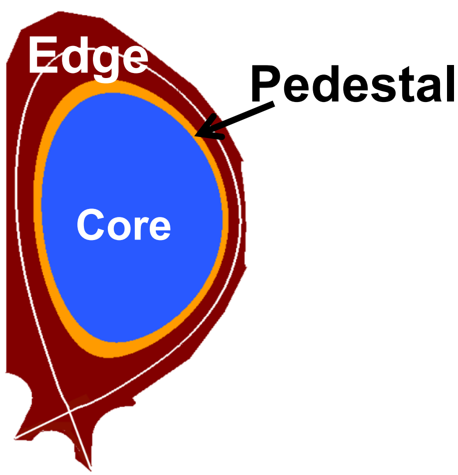

Pedestal determines the plasma pressure, hence fusion yield, in the burning core

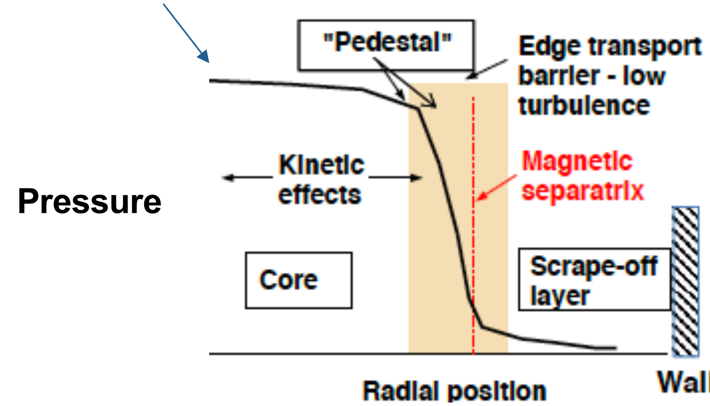

WDMApp Core-Edge Coupled Simulation shows seamless turbulence coupling

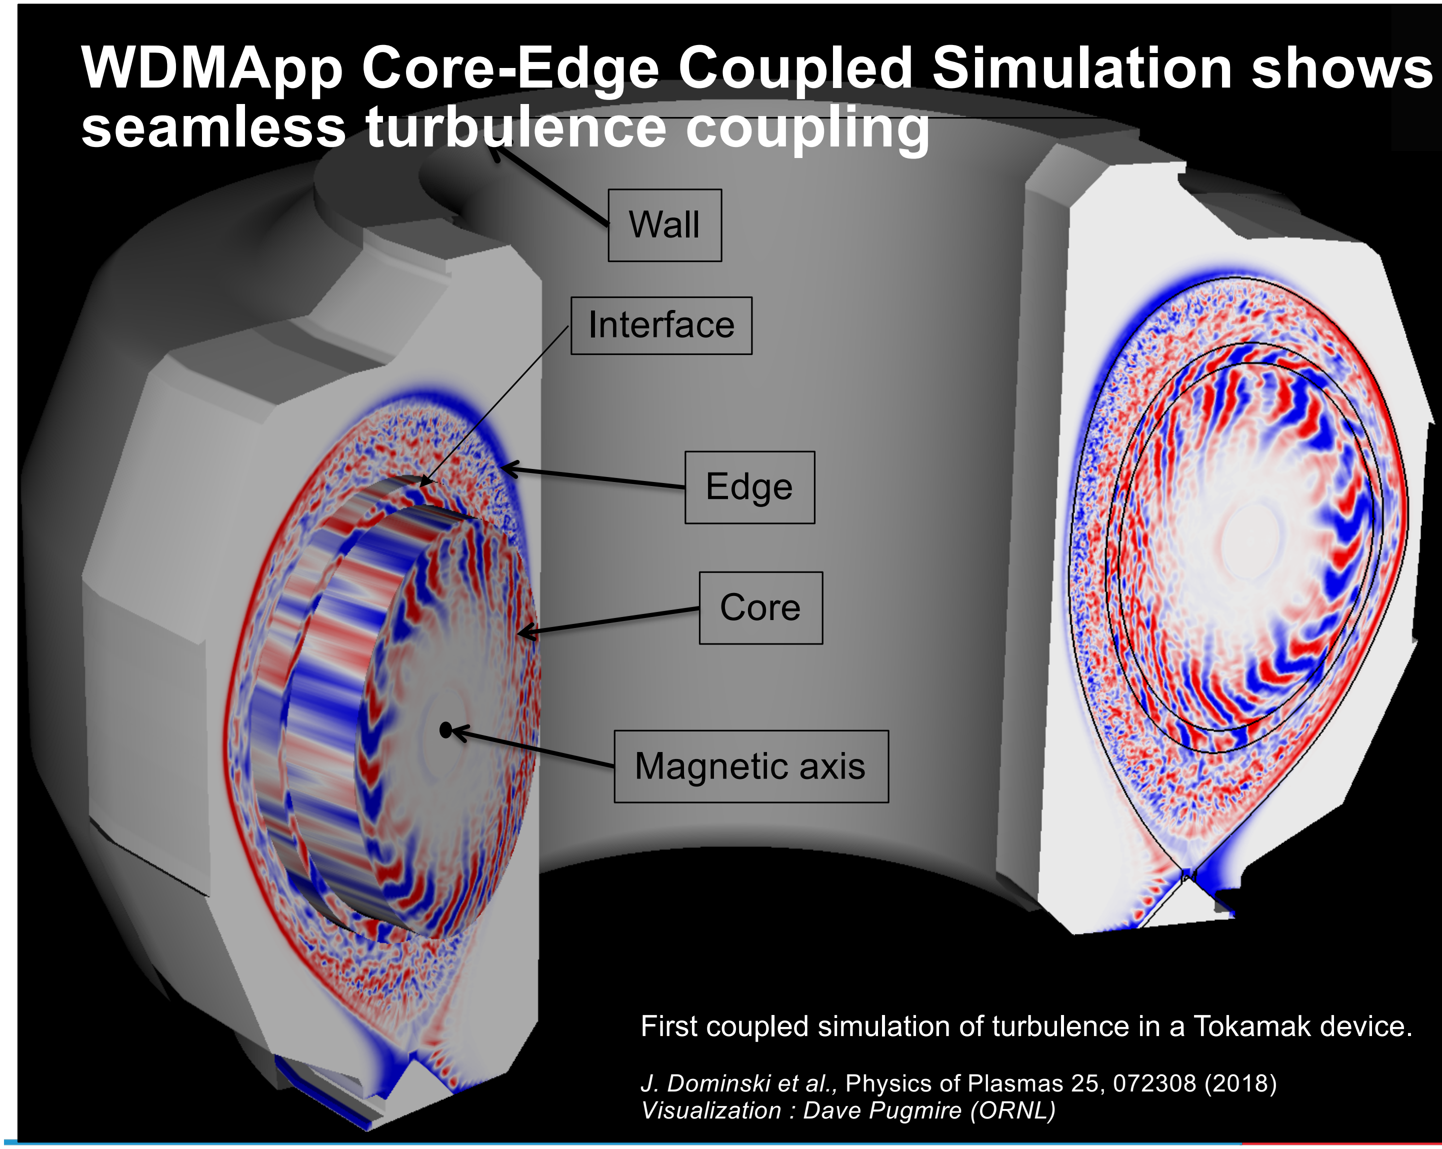

First coupled simulation of turbulence in a Tokamak device.

J. Dominski et al., Physics of Plasmas 25, 072308 (2018)
Visualization : Dave Pugmire (ORNL)

Developing core-edge coupling of technology

1. We first use XGC-XGC coupling to develop the technology
2. Apply the technology to GENE/GEM and XGC coupling

XGC is the leading gyrokinetic code for simulating edge region, including a separatrix.

GENE and GEM are leading gyrokinetic codes for simulating core region.

Cross-verification between GENE and XGC:

Non-linear ITG instability (J. Dominski, S. Ku, G. Merlo, C.S. Chang, F. Jenko, S. Parker)

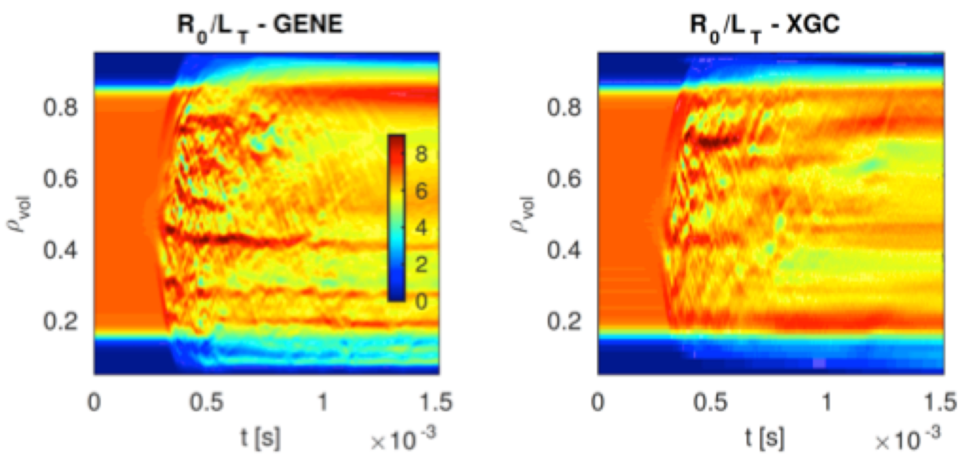

Time-radius dynamics of the logarithmic gradient R_o/L_T

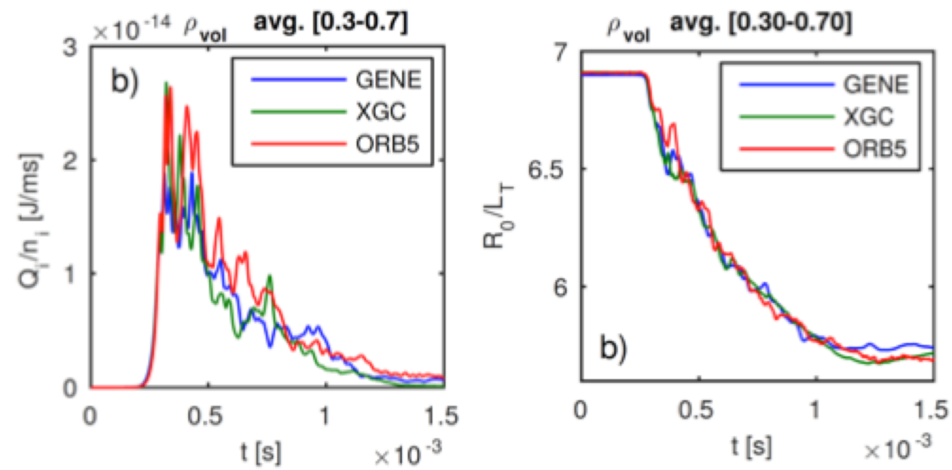

Excellent agreement in the time-evolution of the global ion heat flux and temperature gradient. Radial average is taken over the widest region (0.3-0.7) after removing the simulation-boundary area.

Core-edge coupling

Coupling of XGC-core and XGC-edge

Implemented by J. Dominski, S.H. Ku, and C.S. Chang.

- True kinetic coupling between executables
- The coupled simulation is statistically equivalent to the reference simulation
- Study how to replace the XGC core simulation with a GENE simulation

Difficulty was in avoiding turbulence suppression from nonlinear de-phasing between two codes.

Failed example

Ion heat flux, Qi

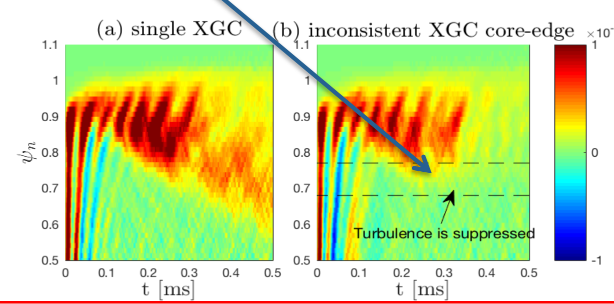

XGC – GEM coupling

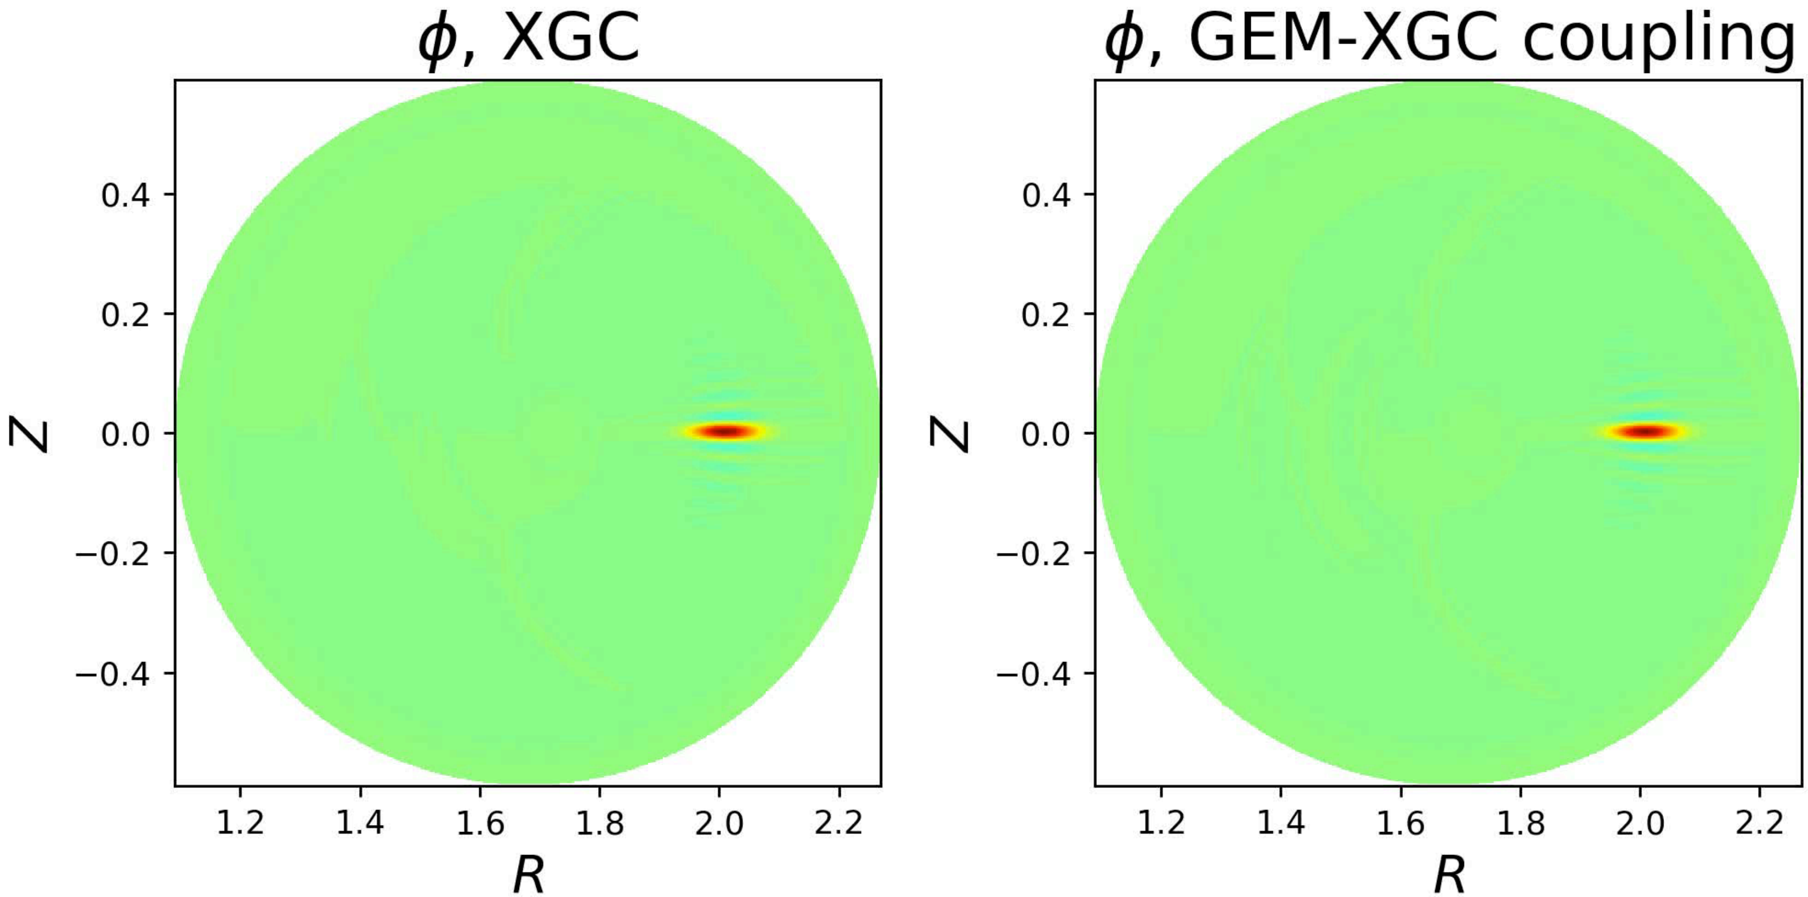

WDMApp coupling results

- Linear results (n=24): frequency almost the same, growth rate differs ~1%, between coupled code and XGC reference
- Nonlinear results (n=3,6,9,12…) show ~4% difference for the saturation level of heat flux, between coupled code and XGC reference
- Coupled code adds little cost when using parallelized grid-quantity mapping (algorithm and performance enhancement). For example, 24.62s for XGC only, 25.23s for coupling with parallelized mapping.

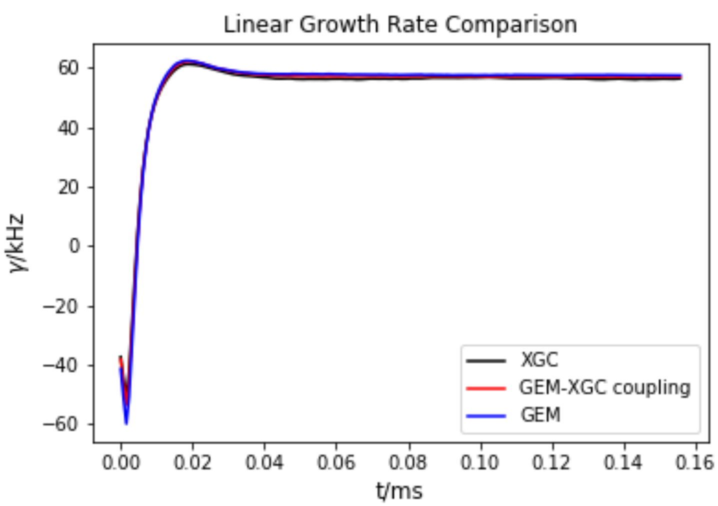

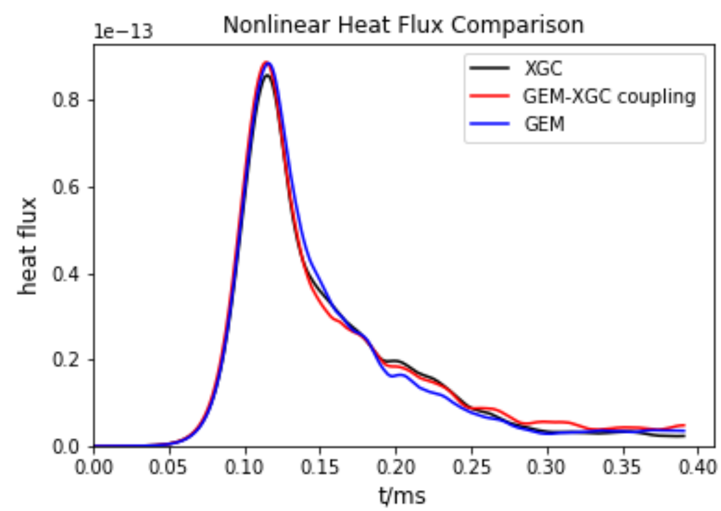

What is EFFIS 2.0?

Workflow Composition | Communication | Monitoring/ Provenance | Toolkit | Output

- EFFIS 2.0 is a workflow coordinator for the WDM App
  - A collection of services to compose, launch, monitor, and communicate between coupled applications
  - Automates "easy" deployment on DOE systems
  - Facilitates "easy" integration to analysis and visualization tools, components, frameworks, etc.
  - Unique features: in-situ memory-based data movement, placement options (e.g. same node), wide-area network, automated visualization

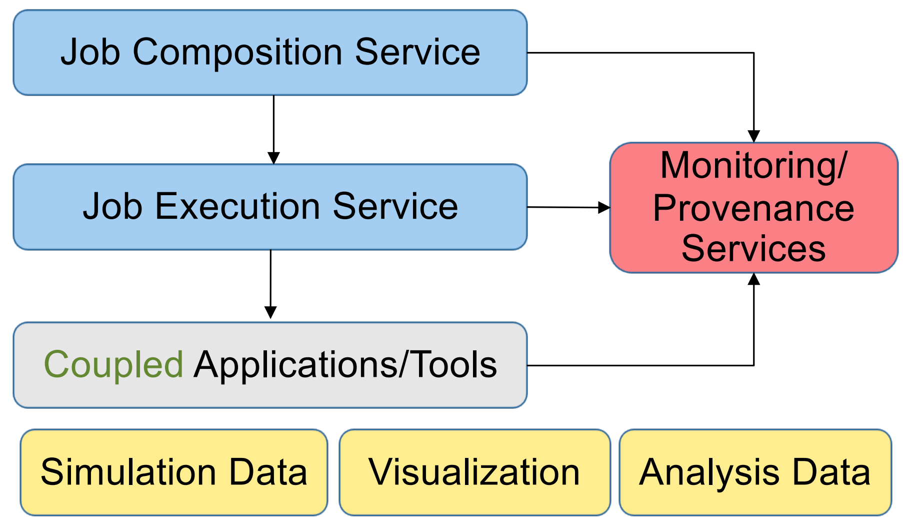

Tiered Science Challenge Goals

- **Basic challenge goal: for consistent measure of FOM**
  - Predict pedestal shape (height and width) in full-current ITER with electrostatic turbulence
  - Gyrokinetic ions and drift kinetic electrons
  - Execute simulation until pedestal reaches steady state
  - Pedestal shape largely determines energy-production efficiency of burning plasma in ITER
- **Advanced challenge goal: for more complete prediction with less uncertainty**
  - Include electromagnetic turbulence
  - Include tungsten impurity particles (ITER has tungsten wall)
  - Predict pedestal shape, and measure the divertor heat-load density

– Expected bottleneck: performance of preconditioner for electromagnetic turbulence solver
– Mitigation strategy: Use 3D PETSc solver and machine learning to speed up preconditioner

GENE-XGC Coupling

Optimization of GENE, XGC and GEM for large-scale Summit

ECP WBS	ADSE12-16 WDMApp
PI	Amitava Bhattacharjee, PPPL
Members	PPPL, ORNL, ANL, UNH, UT-Austin, U.Colorado-Boulder

Scope and objectives

- Port and optimize WDMApp codes on Summit computer in preparation for exascale systems
- Leverage the ECP Co-Design and Software Technologies projects for portability and performance
- Scale WDMApp codes GENE, XGC, and GEM to 20% of Summit

Impact

- Achieving high performance and scalability on a multi-GPU system is a critical requirement towards running WDMApp on Frontier and Aurora

Project accomplishments

- Successful porting to GPU of all three codes used in the WDM application: XGC, GENE, and GEM
- Use of CoPA-developed “Cabana” library in XGC, leading to high portability without loss of performance or scalability
- All 3 codes successfully ran on 1,024+ nodes on SUMMIT

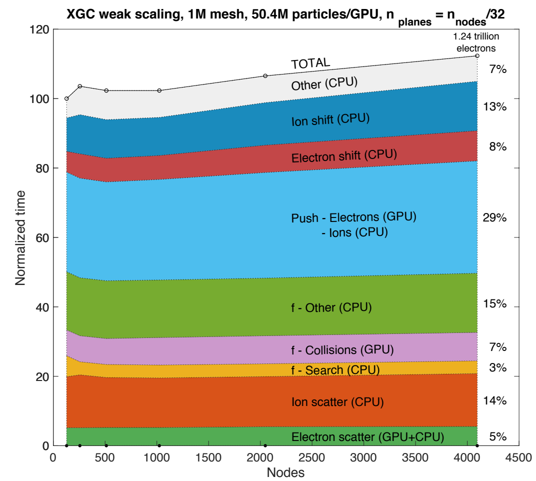

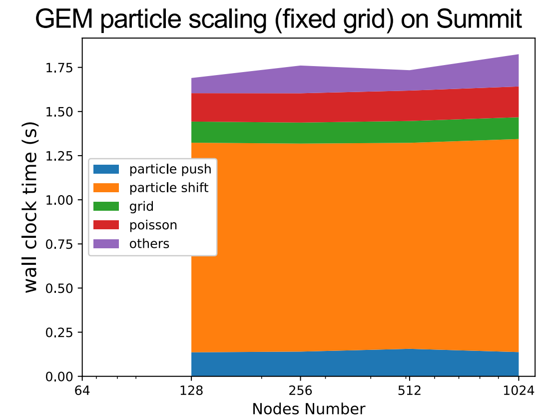

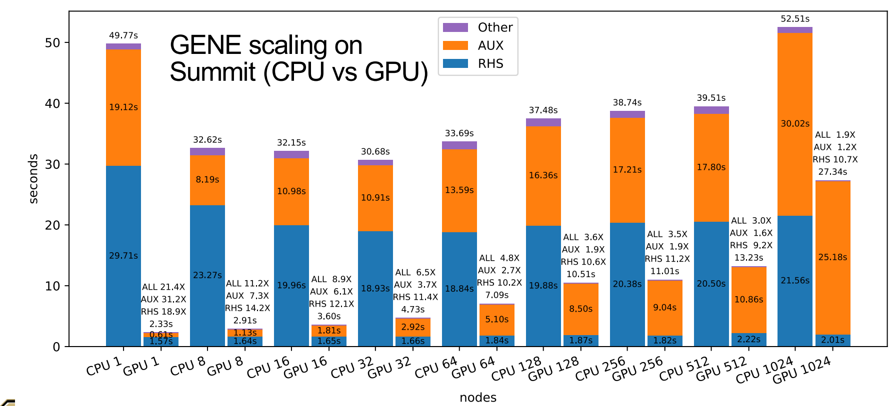

ST & CD Project Collaborations with WDMApp

- **ST**
  - ADIOS-ECP:ADIOS Using ADIOS for I/O, code coupling, and in situ analysis/visualization
    - Indirect collaboration for analysis: VTK-M team, CODAR (FTK)
  - ExaMPI: MPI is used in all codes (no support from ExaMPI)
  - SOLLVE:OpenMP: OpenMP is being used in all codes (no support from SOLLVE)
  - SPACK: Used for package management in EFFIS (no support from SPACK)
  - PETSc: Used in WDMApp, for the field and collision solver: (M. Adams, not supported by WDMApp yet)
- **CD**
  - MGARD (CODAR): investigating the use for the unstructured XGC1 datasets (M. Ainsworth)
  - Cabana (COPA) : Used in XGC1 to provide platform portable performance for GPU optimization (S. Slattery)
  - Savanna (CODAR): Used for the job submission and placement for EFFIS (K. Mehta)
  - TAU (CODAR): Used to obtain performance information for the coupled runs (K. Huck)

Multiple-Time-Scale Coupling for advanced challenge problem

WDMApp turbulence capability ~1ms time-scale

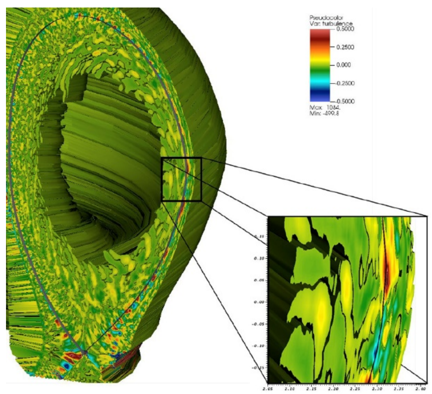

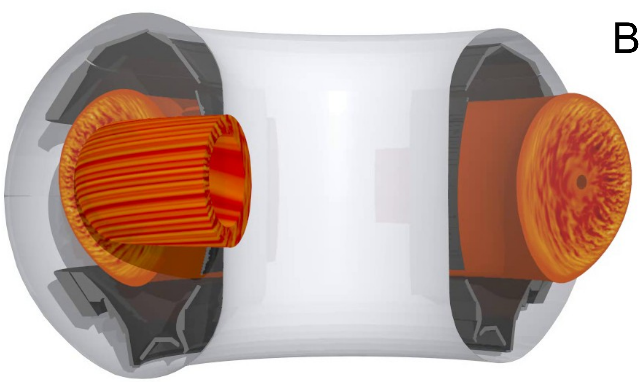

Multi-scale Coupling

Brute force requires resources far beyond exascale

Algorithmic innovation required

Full Experiment ~1s time-scale

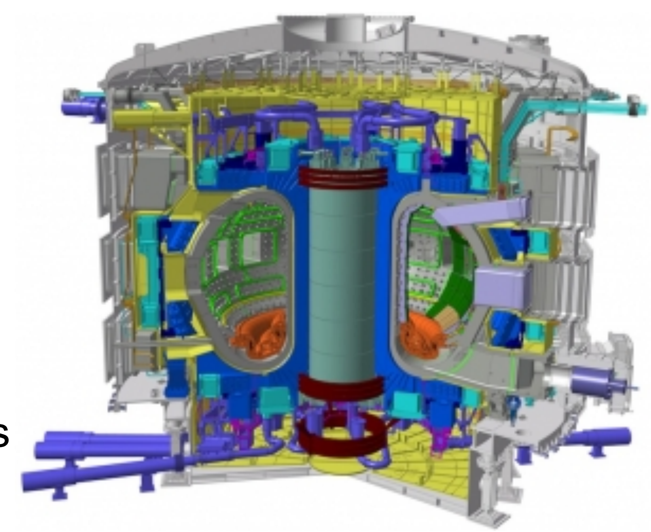

Challenge of Multiple-Time-Scale Coupling (continued)

Advanced challenge problem requires multiple approaches for risk mitigation

Tango-GENE

- GENE coupled to 1-D transport
- Recent innovation accelerates convergence near marginal stability

Relative errror vs. iteration

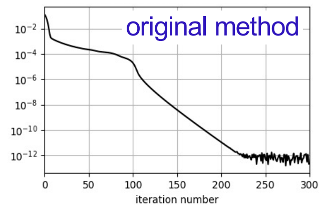

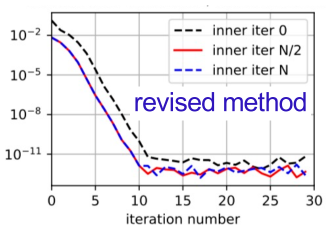

Equation-free

- Kinetic coupled to fluid moments
- Novel “lifting” operator preserves kinetic effects
- Currently implemented in XGCa
- 4.5x speedup in test cases

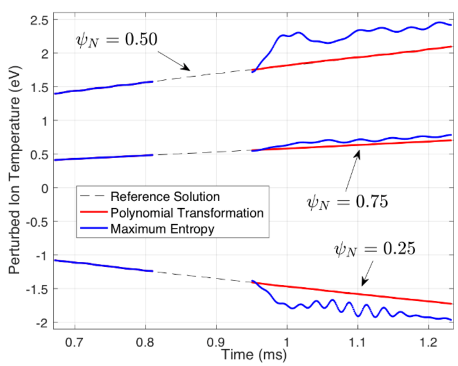

Parallel-in-Time

- Builds on MGRIT method
- New wrinkle: model reduction
- Synergistic with other approaches
- Additional parallelism improves scalability
- General framework compatible with wide variety of physics phenomena

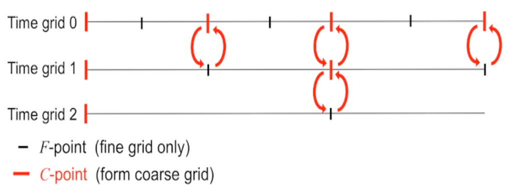

Image courtesy of XBraid

Conclusions

- WDMApp is a leading priority of the fusion community, and will deliver a computational tool of unprecedented power and versatility.
- We have focused here on two primary goals: (1) Coupling of core gyrokinetic code (GENE and GEM) and edge gyrokinetic code (XGC), and performance of the coupled code with FOM > 50 (2) Development of a user-friendly extensible framework EFFIS 2.0 for code-coupling in WDMApp.
- The science is extremely interesting, and compute power will help realize Hawking's vision for science in the 21st century.