SciDAC PI Meeting, July 26-18, 2019

Partnership Center for High-fidelity Boundary Plasma Simulation (HBPS)

C.S. Chang[1],

Physics: **S. Parker**[2], Y. Chen[2], R.M. Churchill[1], D. Curreli[12], N. Ferraro[1], M. Greenwald[3], R. Hager[1], A. Hakim[1], G. Hammett[1], S. Ku[1], G. Wilkie[1]

Math: **L. Chacon**[4], M. Adams[5], C. Hauck[6], J. Hittinger/L. Ricketson[7], R. Moser[8], M. Shephard[9]

Computer Science: **S. Klasky**[6], J. Choi[6], W. Hoffman[10]

Performance Engineering: **E. Dazevedo**[6], P. Worley[11]

and **the HBPS Team**

[1]PPPL, [2]U. Colorado, [3]MIT, [4]LANL, [5]LBNL, [6]ORNL, [7]LLNL,

[8]U. Texas, [9]RPI, [10]Kitware, [11]PHWorley, [12]U. Illinois

ALCF

Outline$^{+,*}$

- **HBPS team leaders and the project OV**
- **The XGC gyrokinetic code**
  - **Performance enhancement**
- **Example scientific discoveries last year**
  - Turbulence interaction with external magnetic perturbation
  - Further investigation of ITER's divertor heat-flux width
- **Getting ready for WDM integration**
- **Collaboration with FASTMath and RAPIDS**
  - FASTMath: M. Shephard, J. Hittinger
  - RAPIDS: S. Klasky, W. Hoffman
- **Conclusion and discussion**

$^{+}$Funding provided by US DOE FES and ASCR via SciDAC-4
*Computational resources provided by INCITE (OLCF, ALCF) and NERSC

Leadership in HBPS

- **Advisory Committee**
- **Center for High-Fidelity Boundary Plasma Simulation**
  Lead PI: C.S. Chang (PPPL)

Executive Management

- **Physics** — **Parker** (Colorado)
  - **Kinetic** — Ku/Hammett (PPPL)
  - **MHD** — Ferraro (PPPL)
  - **Neutrals** — R. Kube/Stotler (PPPL)
  - **Validation** — Greenwald (MIT)
- **Applied Math** — **Chacon** (LANL)
  - **Algorithms** — Adams/Hauck (LBNL/ORNL)
  - **UQ** — Moser (Texas)
  - **Verification** — Hittinger/Ricketson (LLNL)
  - **Meshing** — Shephard (RPI)
- **Computer Science** — **Klasky** (ORNL)
  - **Performance** — **D'Azevedo (ORNL)**, Worley (PHWorley)
  - **Data Management** — Klasky/Choi (ORNL)
  - **Software Engineering** — Hoffman (Kitware)

Center for High-fidelity Boundary Plasma Simulation

(High-fidelity gyrokinetic simulation of the global BD plasma)

ASCR's Enabling Techology

XGC + DEGAS2* + M3D-C1+ + hPIC#

L-H transition
Pedestal shape
+ELM control
Divertor heat-flux width
Neutral particles and Impurities
#Sheath physics and integration with PMI

production component

GEM + GENE

E&M turbulence in omnigeneous B

cross-verification & instruction component

Continuum GK edge code **Gkeyll**

developmental component

BD plasma module for high-fidelity WDM on exascale computers

reduction

BD plasma module for Reduced WDMs

***DEGAS2 is coupled into XGC as a subroutine.**

The XGC Gyrokinetic Code

- **Particle-in-Cell, with added continuum technology**
- **In contact with material wall**
  - Far-from-equilibrium (non-Maxwellian)
  - Neutral recycling, impurity sputtering
- **Magnetic X-point and separatrix (one poloidal B winding→ infinite distance)**
  - Unconfined X-point orbit loss from pedestal
- **Multi-scale, multiphysics in space-time**
- **Unstructured triangular mesh**
- **PETSc (only ~2% of total compute time if done right)**
- **Large simulation-size (≳10k particles per grid-vertex) per time-step → ~ trillion particles for ITER**
- **Total-f XGC has been developed to study this kind of plasma**
- **XGC is not only a SciDAC code, but also an ECP code**
- **Summit-ESP, Aurora-ESP, NESAP, and INCITE**

XGC1

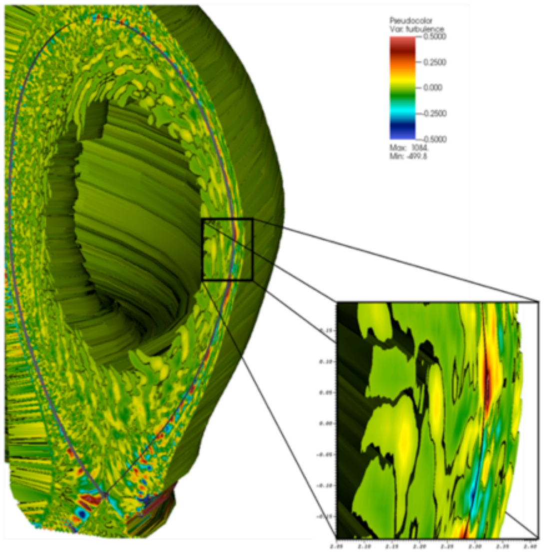

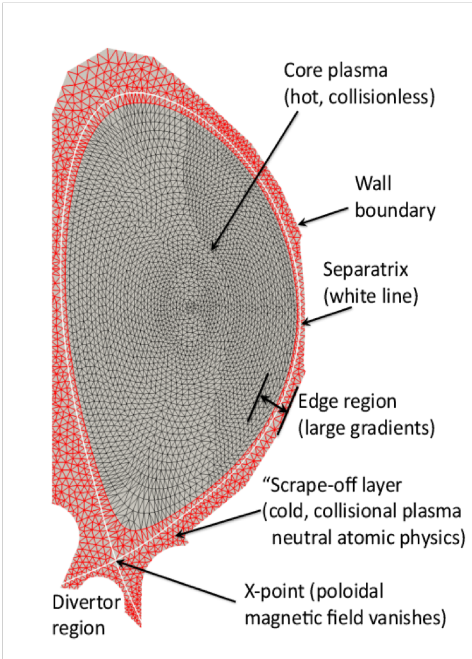

XGC perfomrs well on world #1 Summit

[E. D'Azevedo, A. Scheinberg, P. Worley, S. Ku, R. Hager, S. Slattery]

- XGC utilizes Kokkos for "future proofing," a programming model written in C++
  - Fortran interface has been developed for XGC via Cabana particle library
- Single version XGC can be performance portable not only to Summit and KNL, but also to different future architectures: Permutter, Aurora, Frontier

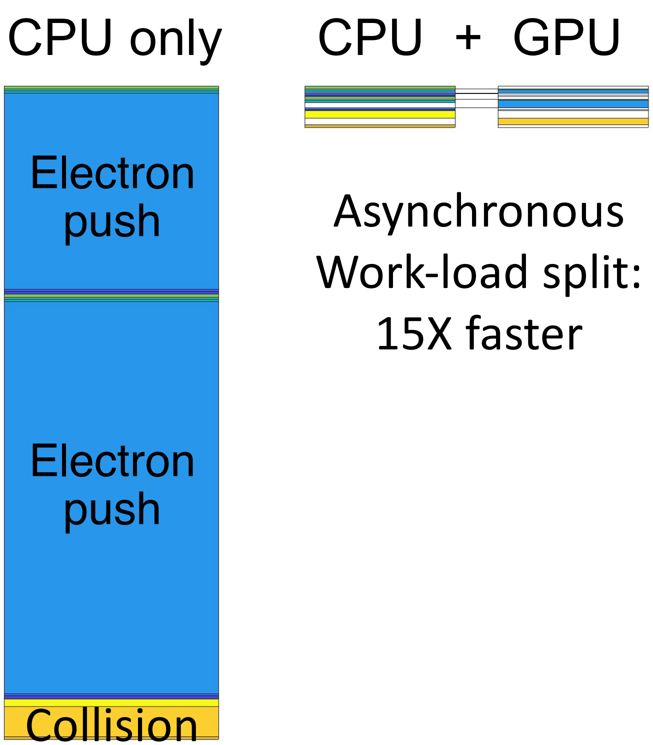

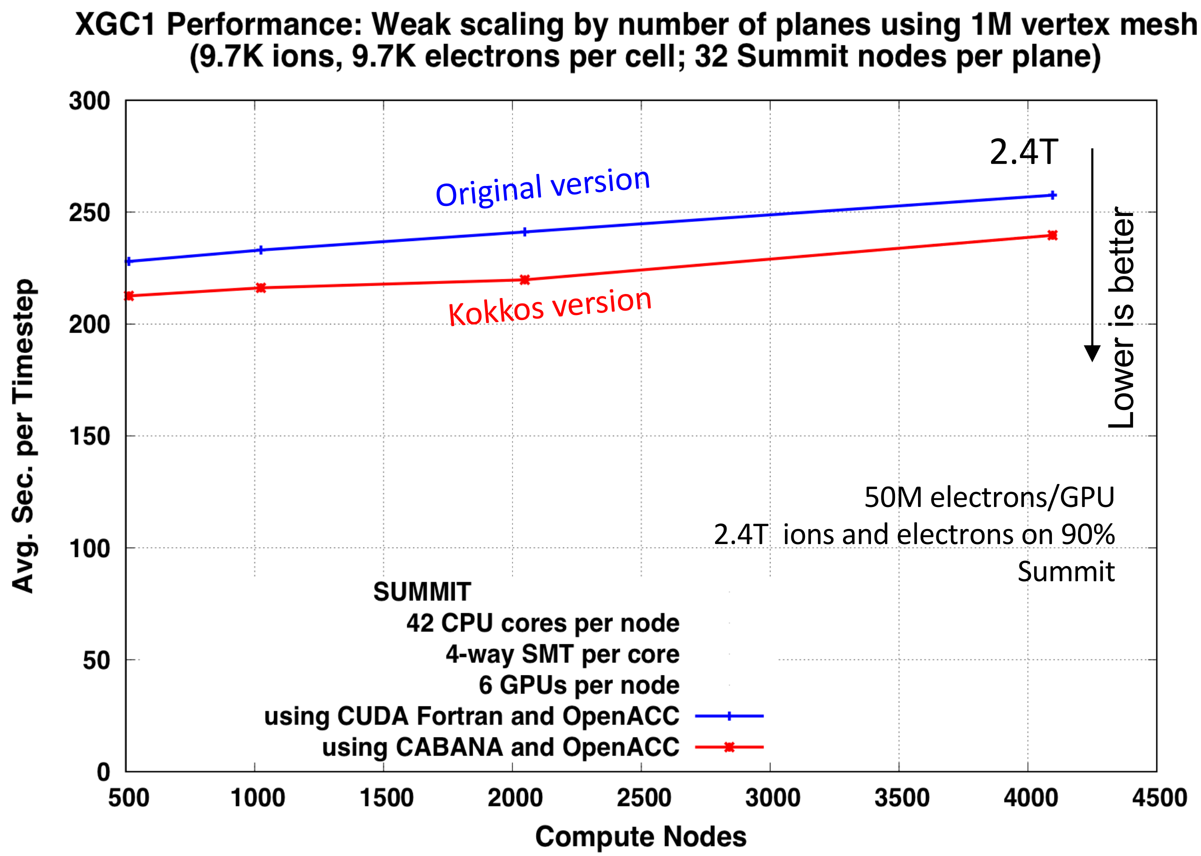

Example Scientific Discoveries last year enabled by SciDAC Institute collaborations

For the first time, the XGC code enabled total-f gyrokinetic simulation of RMP-neoclassical-turbulence-neutral interaction in diverted geometry

[R. Hager, 2019 APS-DPP Invited Talk]

Two big puzzles seen in the experimental RMP-driven edge transport:

1. Why is there the density pumpout?
2. How could an electron heat-barrier co-exist with density pumpout?

Puzzle 3: These two experimental observations are completely opposite to the existing "well-established" analytic model on stochastic transport [Rechester-Rosenbluth].

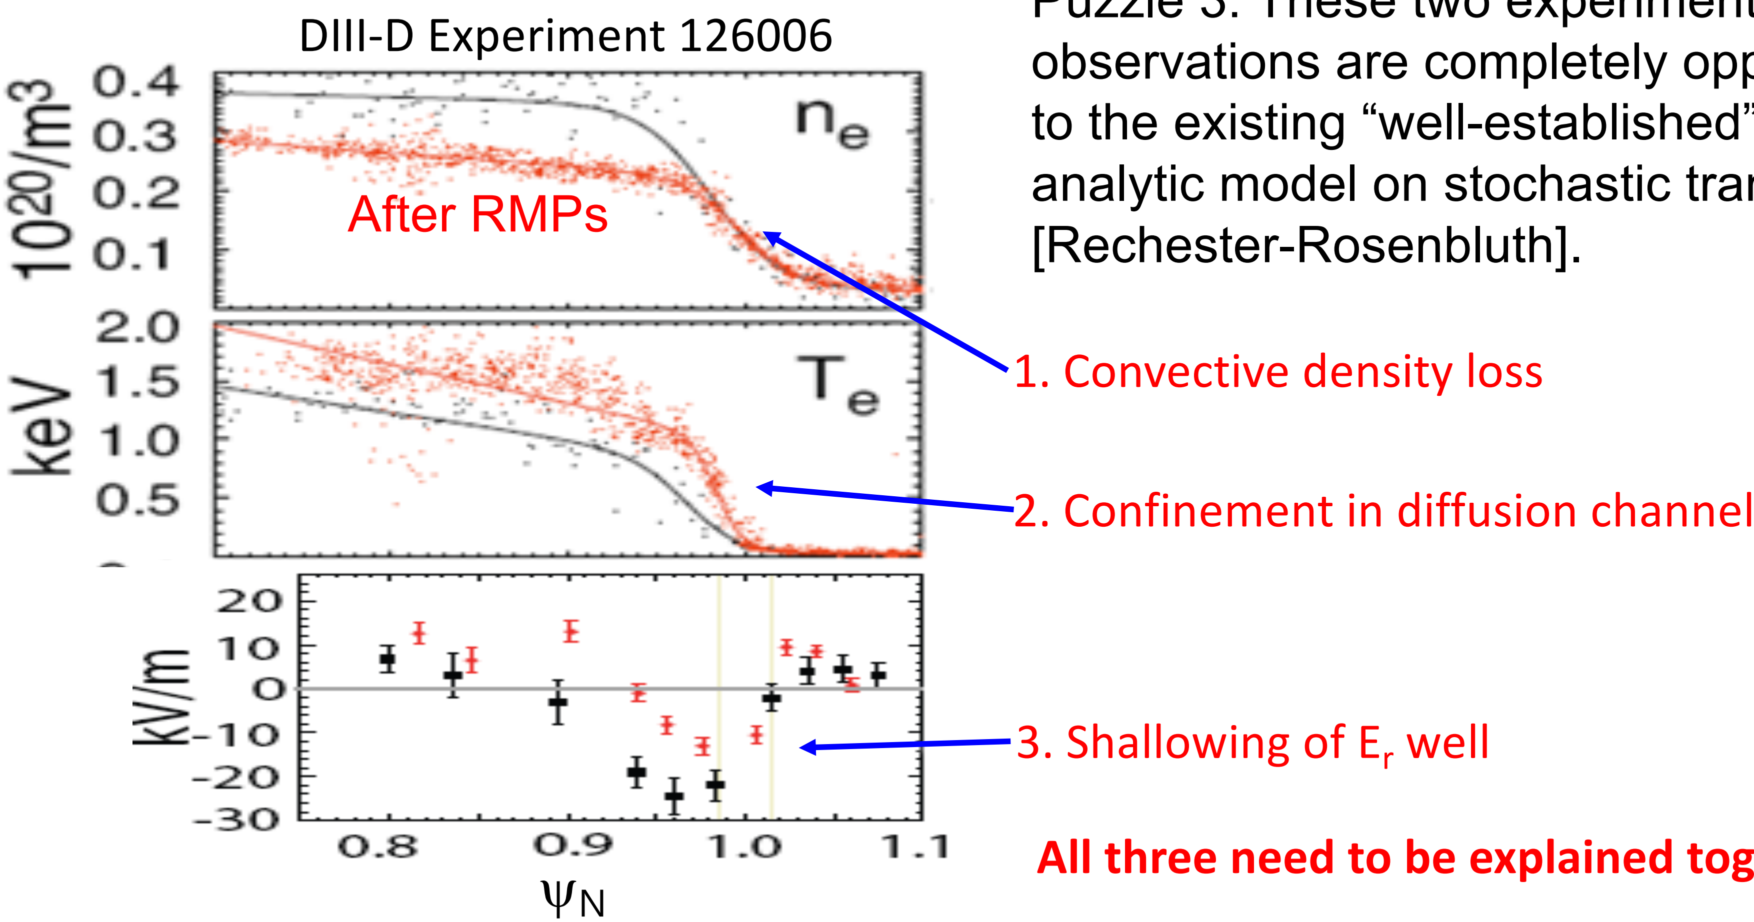

All three need to be explained together.

For the first time, XGC succedded in the total-f gyrokinetic simulation of RMP-neoclassical-turbulence-neutral interaction in diverted geometry

[R. Hager, 2019 APS-DPP Invited Talk]

- Finds enhanced <u>micro-turbulence</u> is responsible for the density pump-out in the M3D-C1 RMP field, not neoclassical as conjectured previously.
- Finds that the RMP-interacted turbulence preserves the electron heat barrier around $\Psi_N \approx 0.96$-0.98.
- Finds that the E_r-well shallows.
- Finds that the divertor heat-flux width widens by ~30% by RMPs, as seen in experiments.

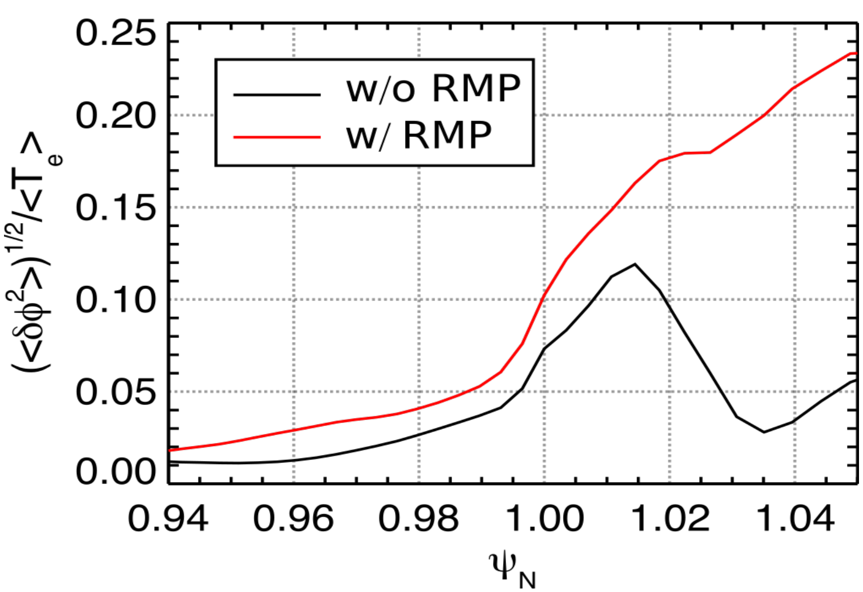

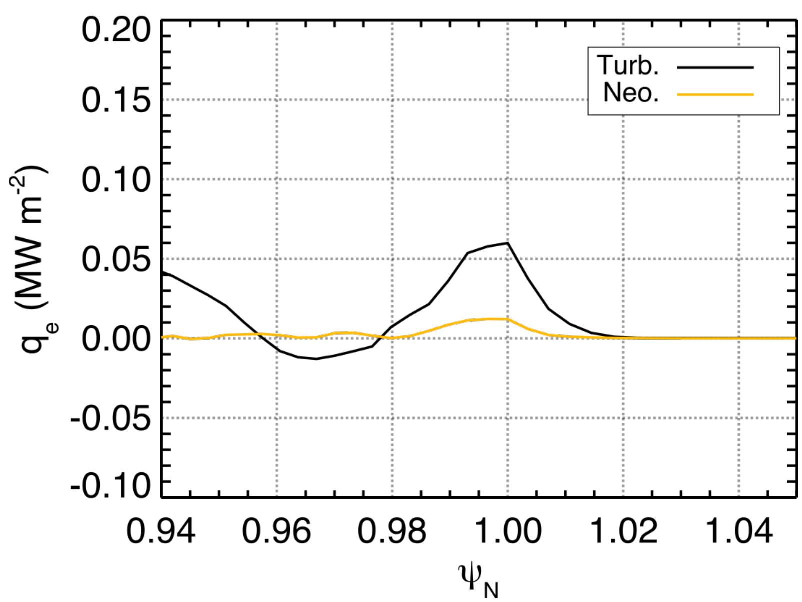

Answering the critical question on the (non)existence of pedestal in the target ITER plasma

- ITER's goal of achieving 10-fold energy production relies on the existence of high enough edge pedestal (H-mode), as ubiquitously obtained in today's tokamak experiments
- There is a concern that, due to the much greater physical size, the background ExB flow may be too weak to shear away the edge turbulence
  - Edge turbulence could become too strong to support the edge pedestal
- We find that the (electrostatic) turbulence self-organization with the background slope, in the presence of X-point orbit-loss, keeps ITER's edge pedestal still high
  - We will continue the big runs (computing time could be an issue).
  - Next question: How will the pedestal be modified by electromagnetic turbulence?

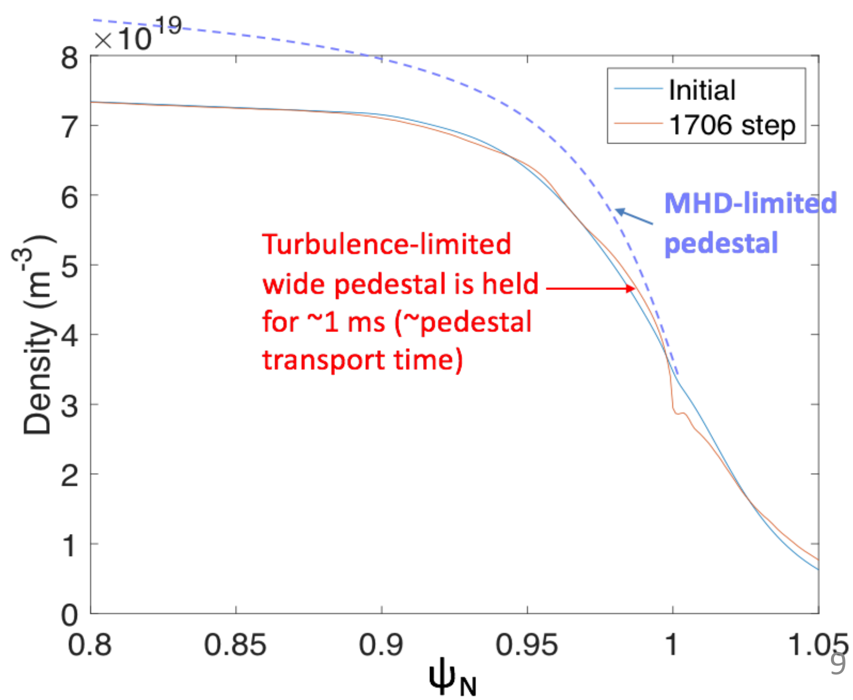

Further investigation of the ITER divertor heat-flux width

Shot	R (cm)	a (cm)	I_P (MA)	$B_{pol,OM}$ (T)
NSTX 132368	85	57	0.7	0.20
DIII-D 144977	169	55	1.0	0.30
DIII-D 144981	169	55	1.5	0.42
C-Mod 1100223026	66	24	0.5	0.50
C-Mod 1100223012	66	24	0.8	0.67
C-Mod 1100212023	66	24	0.9	0.81
C-Mod (high B)	69	20	1.4	1.11
JET 79692	303	91	4.5	0.89
ITER	620	199	15	1.21

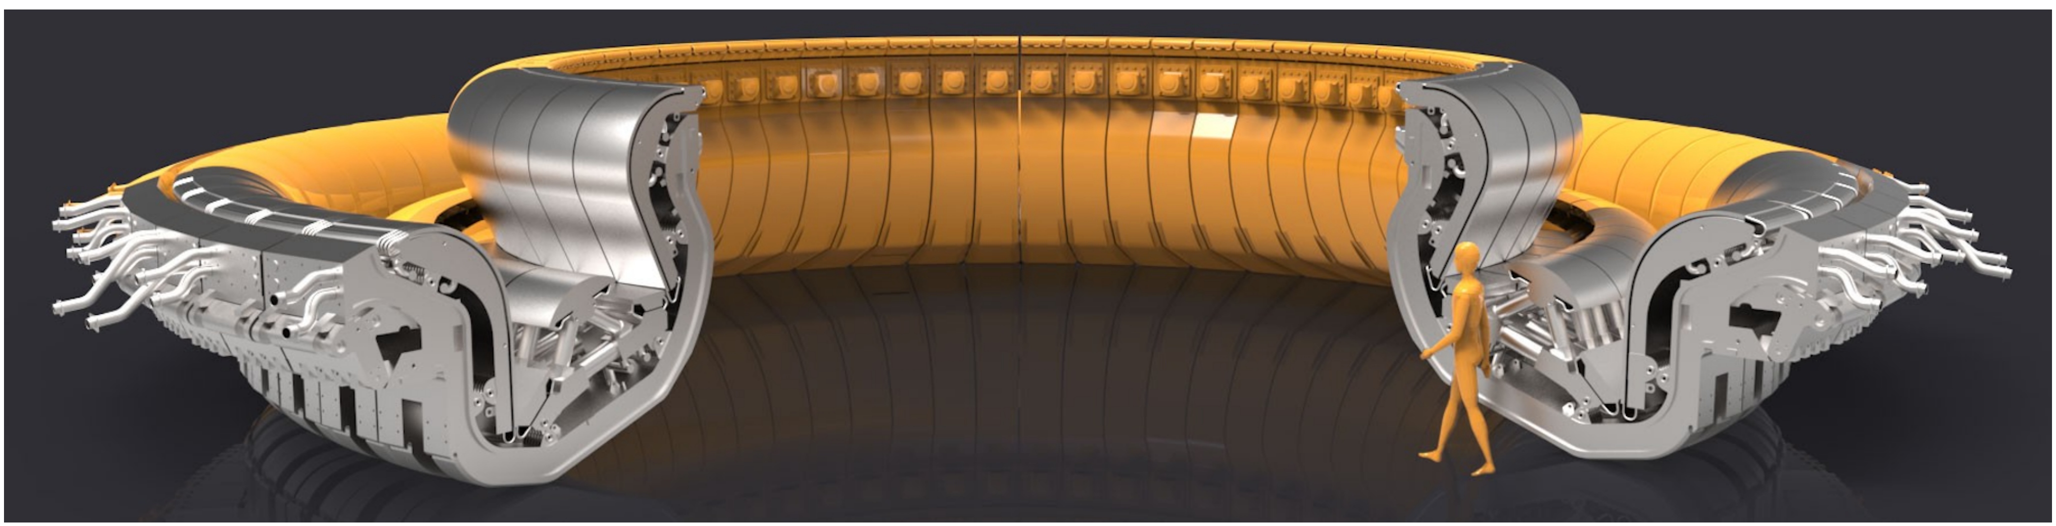

Further investigation of the ITER divertor heat-flux width

- The XGC-predicted divertor heat-flux width λ_q **has been well-validated**
- Why $\boldsymbol{\lambda_q}$**(XGC, ITER) >6X wider** than λ_q(Eich)?
  - To check if it is a pure size effct or B_{pol} effect, a lower B_{pol} ITER and a highest B_{pol} C-Mod case simulated → No is the answer.
- Community question: How come such a strong difference at similar B_{pol}?
- The suspect "**ρ_i/a effect**" remains: Texas [Kotschenreuther], PPPL [Chang]
  - If we correct the B_{pol} dependence to B_{pol} **(a/ρ_i)** dependence, there is a long distance between the high B_{pol} C-Mod to full current ITER.

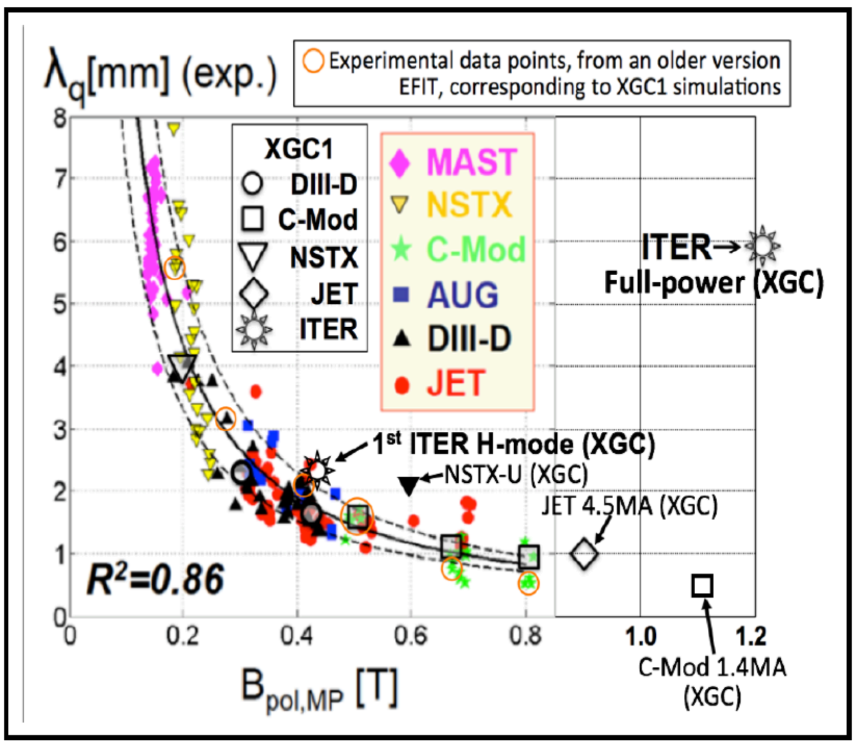

If we include the a/ρ_i effect, there is a long distance between C-Mod and ITER

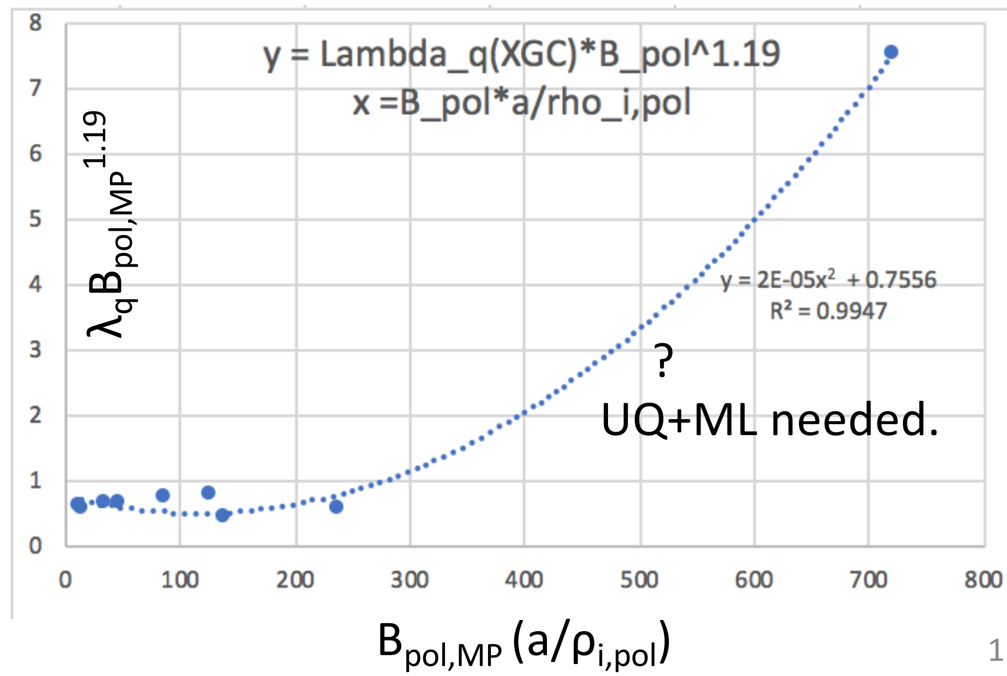

Evidence for different edge physics between higer and lower ρ_i/a values.

In all the higher ρ_i/a tokamaks, including low-current ITER, edge tubulence across the separatrix is blob type and the ExB shearing rate is high. In the high-current ITER, the turbulence is streamer type and the ExB shearing rate is low across the separatrix.

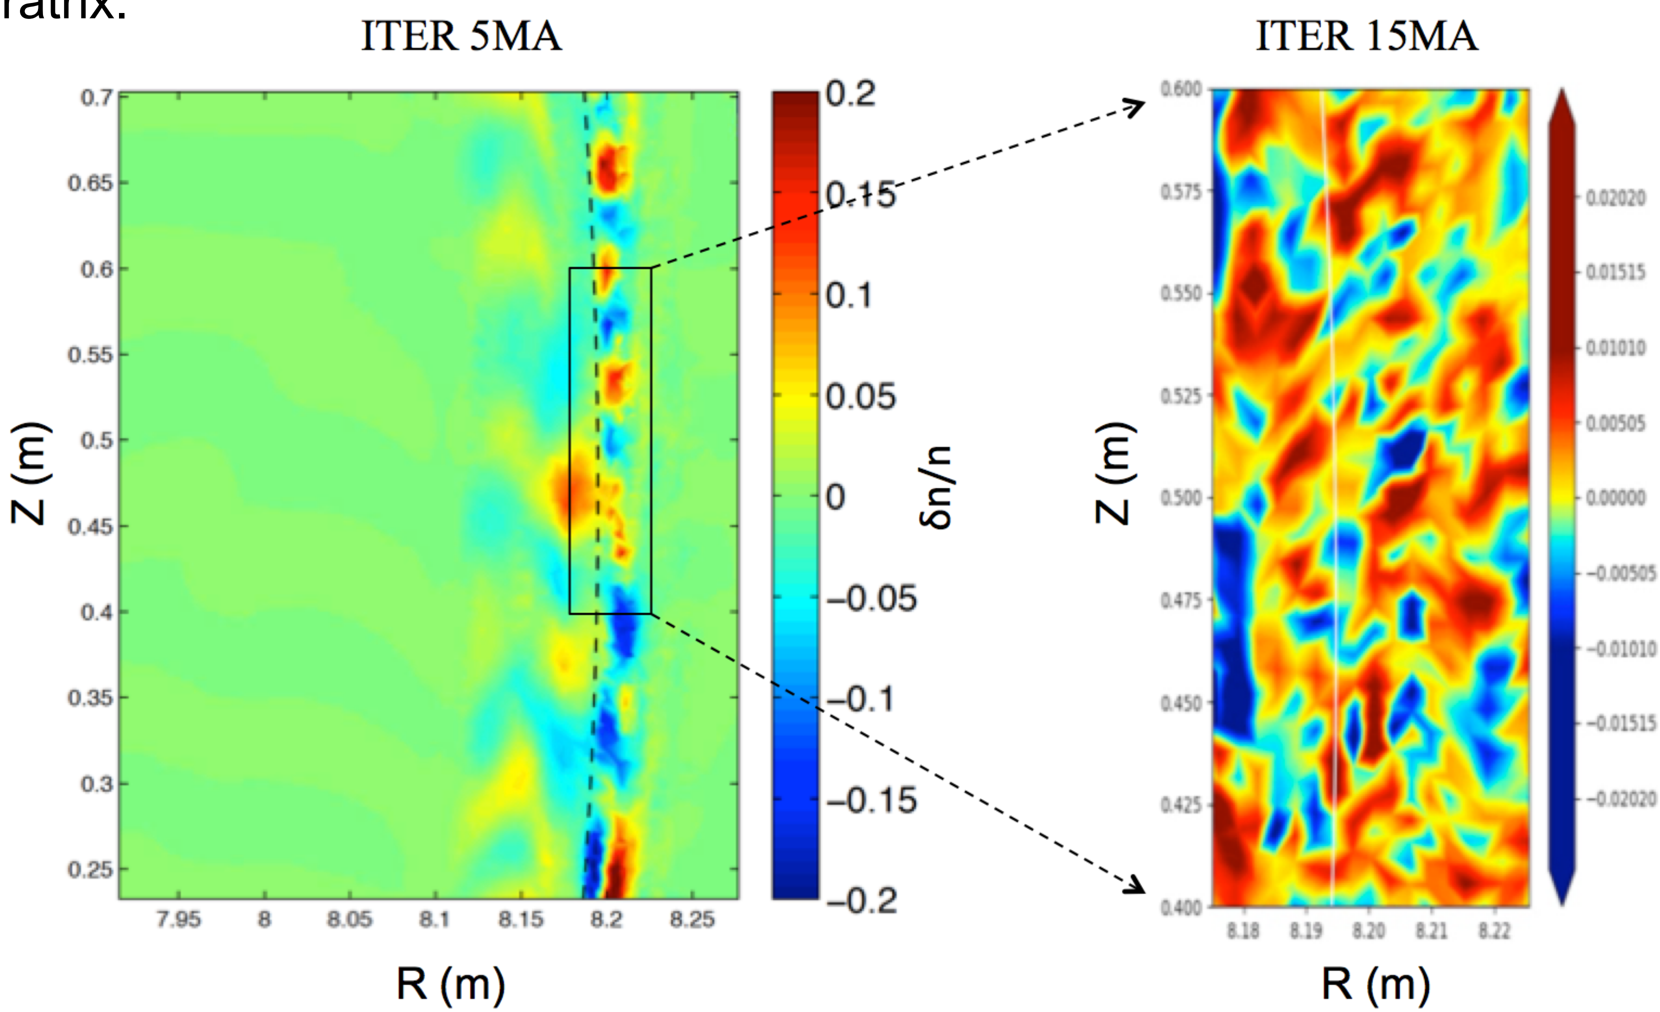

Unlike the blobby turbulence, the full-current ITER containes a strong non-adiabatic electron response across the magnetic separatrix,

as evidenced by a large phase difference between density and potential fluctuations (≳π/2) and a strong de-correlation between their amplitudes.

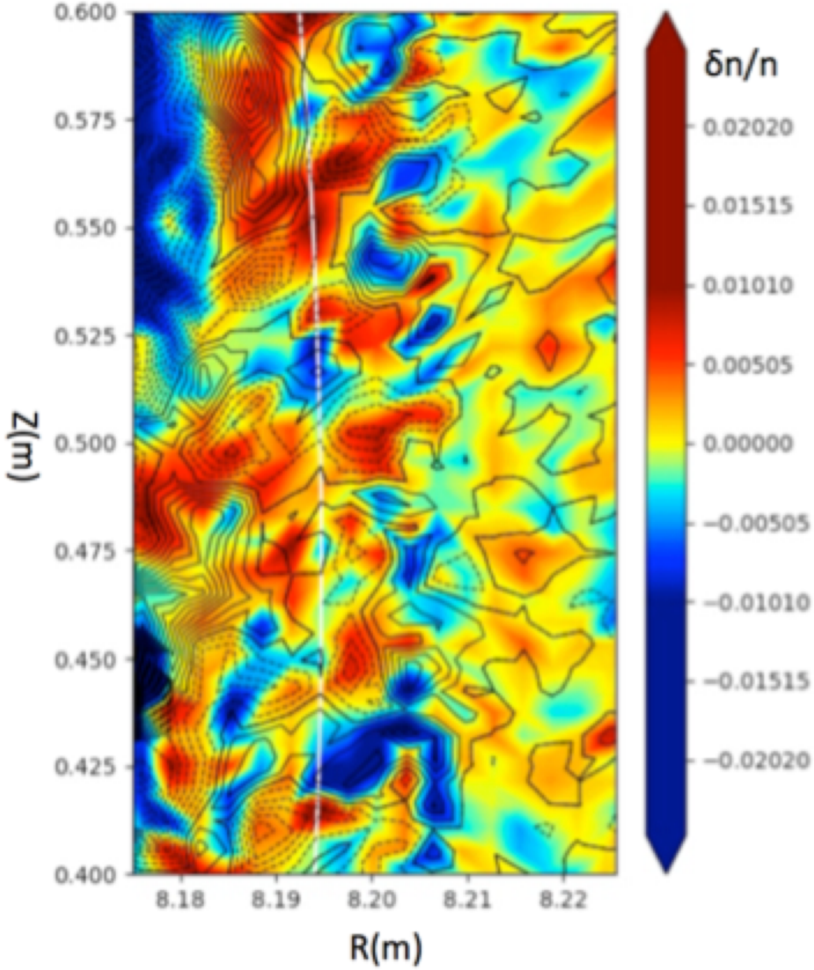

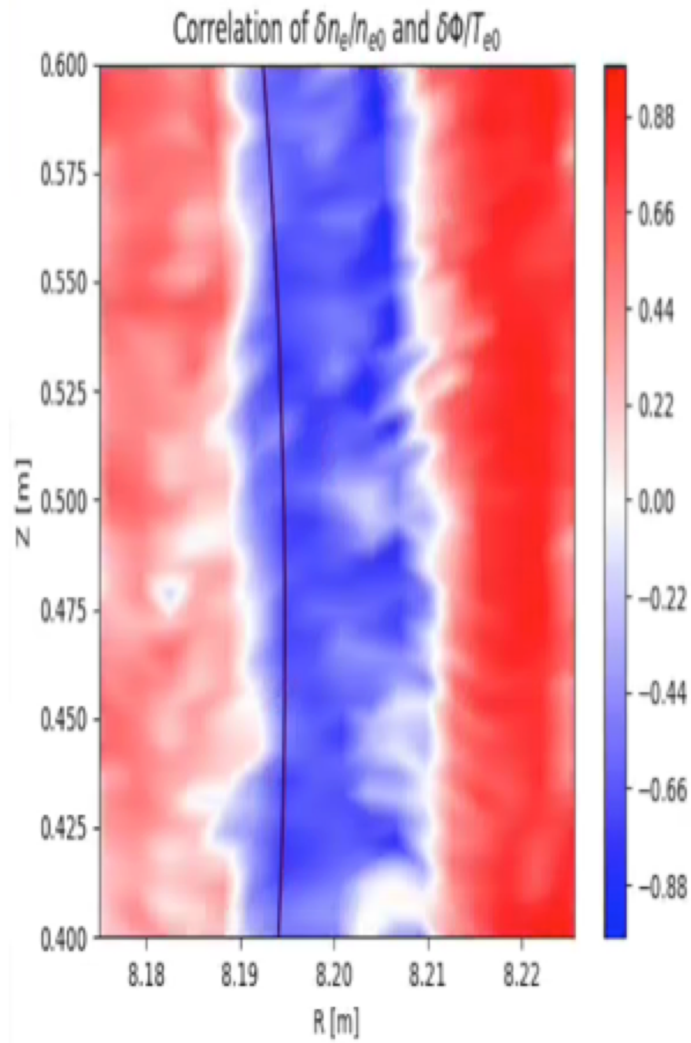

XGC-hPIC Coupling: enabling WDM integration with material codes (U. Illinois, PPPL, ORNL, RAPIDS, FASTMath)

- **Coupling via ADIOS-2/EFFIS**
- **hPIC (6D full-orbit PIC) allows to account for gyro-orbits near material surface, where a plasma sheath and a magnetic presheath are formed**
- **The 5D Gyro-Center distributions provided by the XGC code are converted into full 6D distributions and fed in as input to hPIC**
- **hPIC provides the Energy-Angle distributions of the ions across the sheath before they impact on the surface, a necessary input for material codes evaluating sputtering and reflection of the plasma particles at the surface (e.g. Fractal-TRIDYN, etc.)**
- **Next Steps: Porting to GPUs and Testing of coupled codes on HPC**

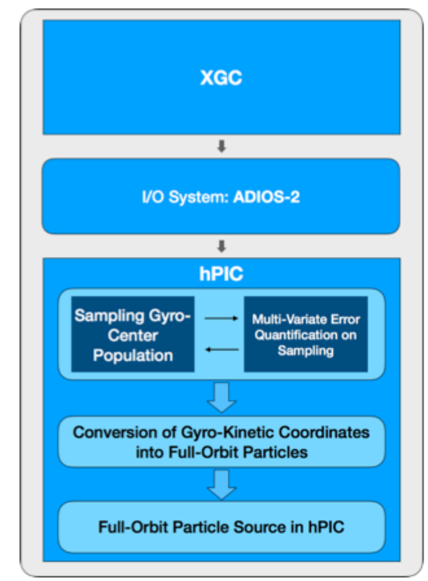

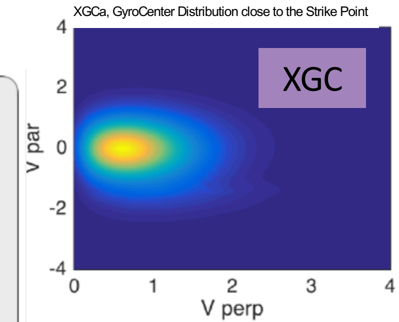

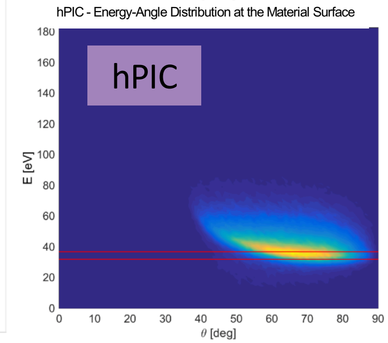

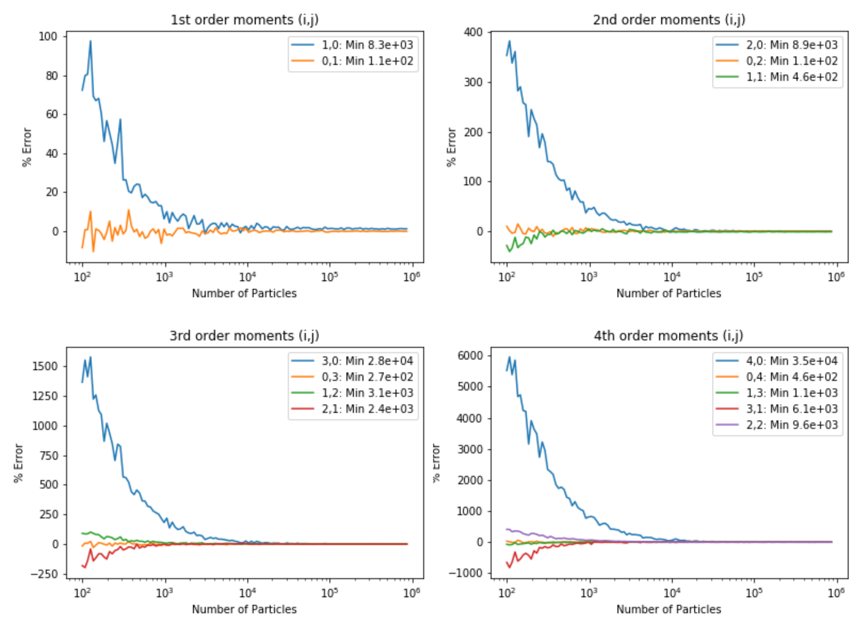

ILLINOIS

Collaboration with FASTMath

FASTMath(1): Meshing tool (M. Shephard)

- XGC uses unstructured triangular mesh due to separatrix and wall:
  - Mesh quality affects the simulation quality
  - Before the meshing by FASTMath, the meshing was performed manually. One case took a few days, but still overlooking details at many areas
  - M. Shephard developed an automated tool, by working with physicists

Last year, FASTMath enabled more accurate XGC solution by

- Improving mesh quality where flux curves interact with reactor wall
- Improving mesh gradation around the magnetic x-point
- Reordering mesh data for better memory access

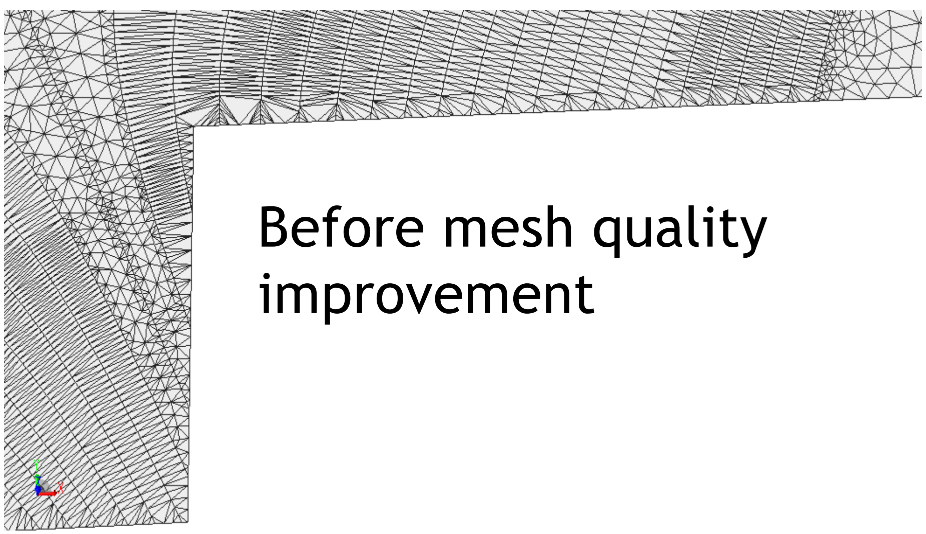

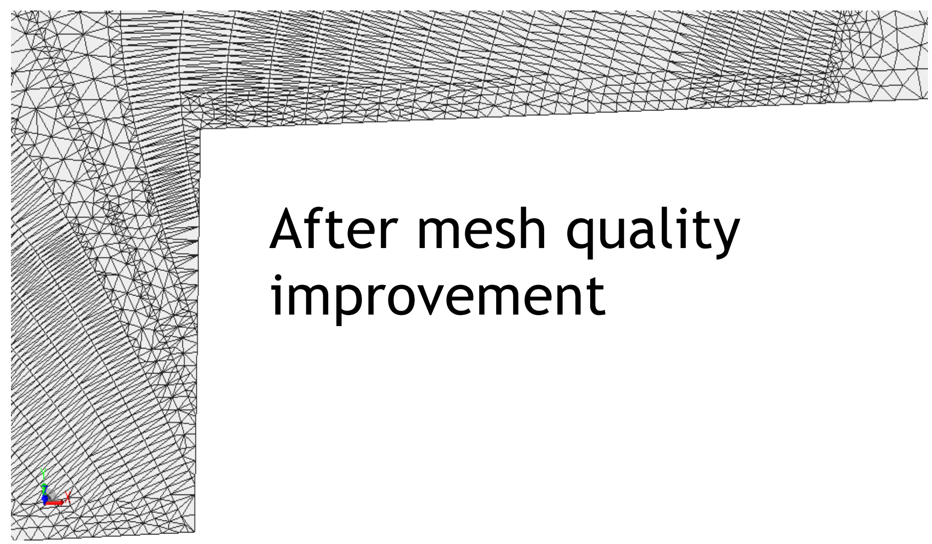

FASTMath (2): XGC-PUMIpic (M. Shephard)

PUMIpic: Parallel Unstructured Mesh PIC

- Mesh centric – no independent particle structure
- Particle migration and load balancing between pushes
- Focused on structures for execution on GPUs
  - Omega GPU-ready mesh-topology being integrated
- Since most of data structures are changed, XGC is rewritten in C++
- Test shows on-par performance using less memory

	no sorting	full sorting
ptcls (Ki)	time (s)	time (s)
128	2.298661	3.642041
256	2.895464	3.415048
512	3.79263	3.851178
1024	4.972283	4.090044
2048	7.089673	4.389198
4096	11.578984	4.799475

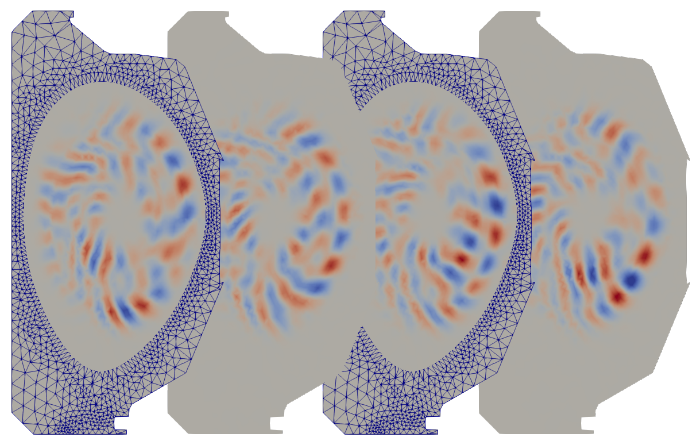

Ion turbulence from XGC-PUMIpic in tokamak geometry on a highly coarse-grained mesh

FASTMath(3): New 3D solver (M. Adams)

Chacon's fully implicit electromagnetic scheme implemented in XGC

- All the previous production runs were with electrostatic turbulence
- Summit is now powerful enough for electromagnetic simulation

- For electrostatic turbulence XGC used 2D solver + finite differencing in the third direction
- The EM scheme requires fully 3D solver
  - Implicit iteration requires efficient solver
  - PETSc/GMRES is used
- Significant effort invested in reducing iteration #
- Machine learning is being explored to optimize the preconditioner [J. Hittinger, R. Archbald]

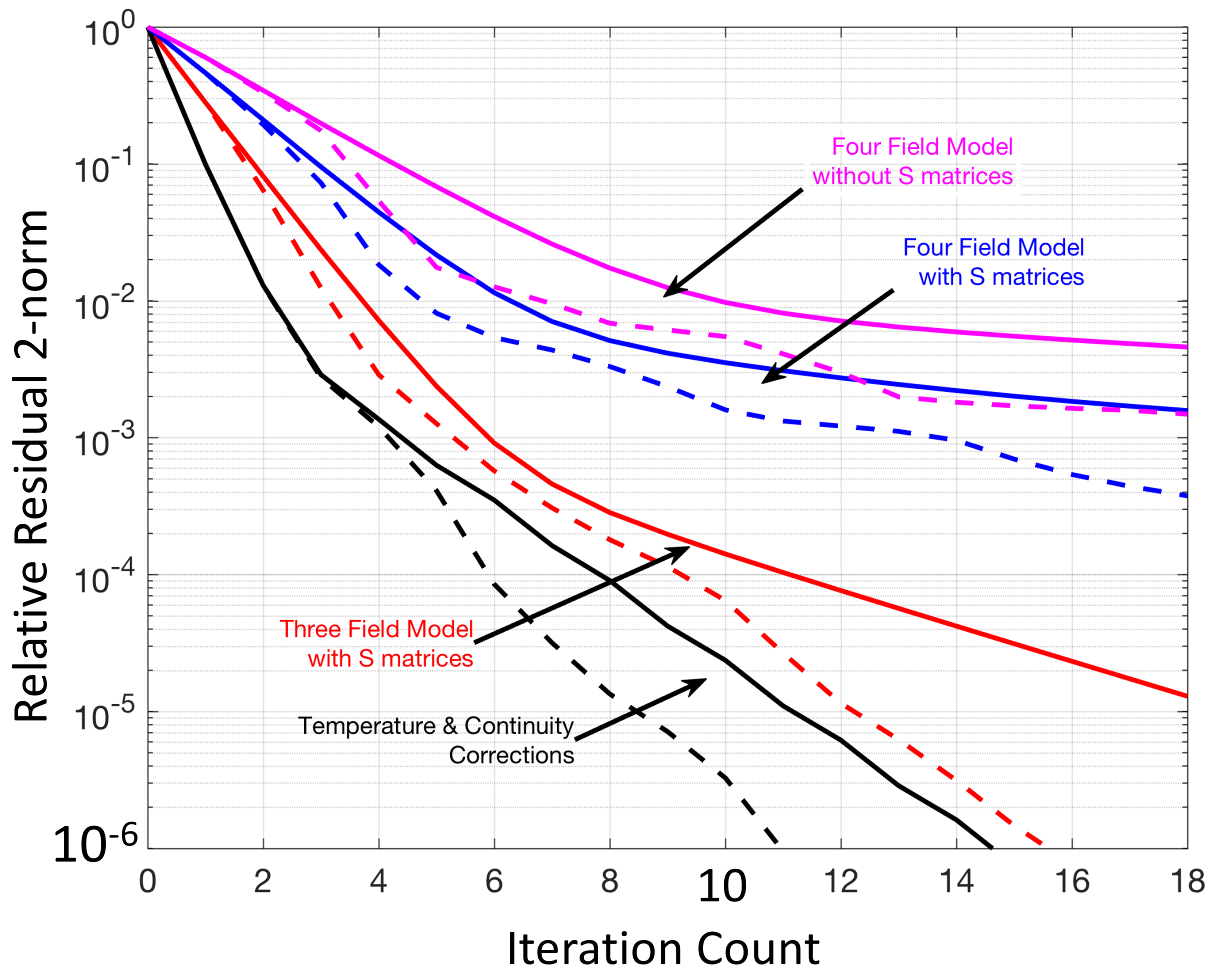

Iteration count is low enough for production on Summit [B. Studervant, L. Chacon]

Other on-going collaborations with FASTMath

- Machine learning for preconditioner: ***J. Hittinger and R. Archbald***
- Nonlinear Fokker-Planck collision operator that is more efficient and WDM-ready than the N^2 operation currently used in XGC: ***Mark Adams***
- Method of Manufactured Solution for PIC solution to Vlasov-Poisson equation: ***J. Hittinger, L. Ricketson, P.J. Tranquilli***
  - Expanded and corrected previous work by Riva et al. [Phys. Plasmas 2017]
  - Demonstrated implementation with 2D Vlasov-Poisson

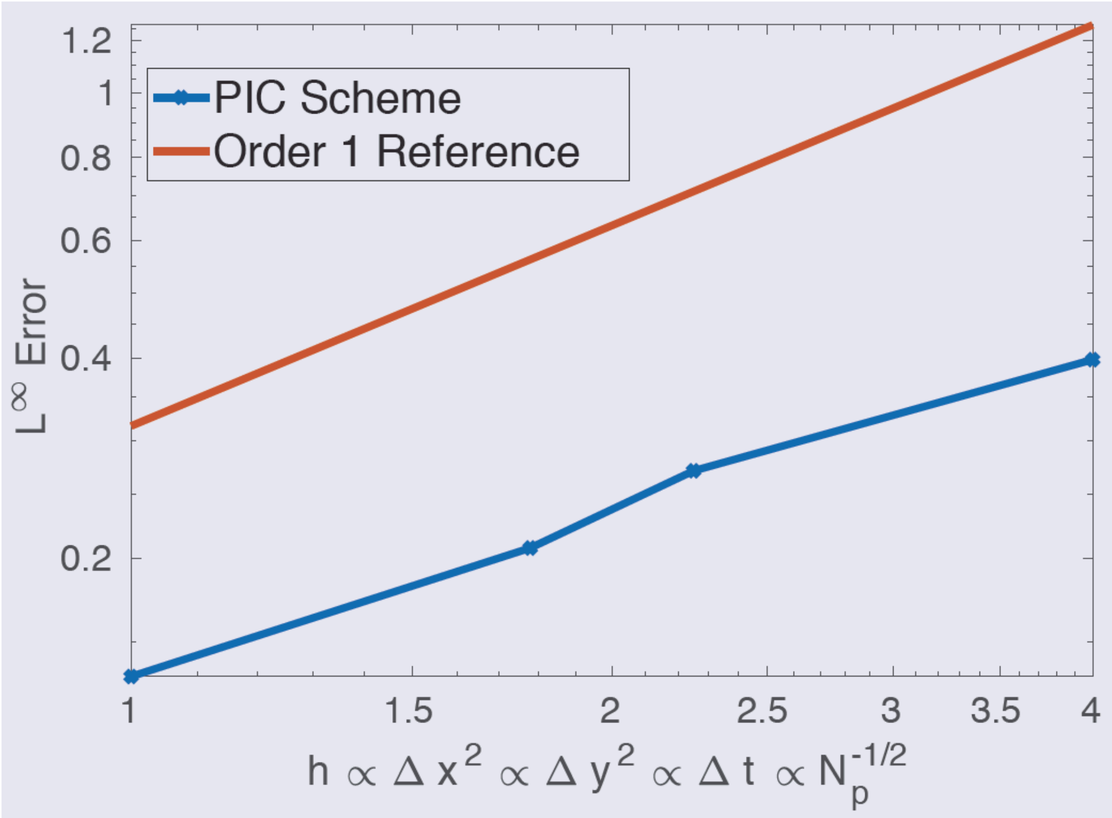

Collaboration with RAPIDS

RAPIDS(1): Big data I/O (S. Klasky)

I/O was a big concern for XGC. ADIOS changed it to "No Concern."

- ITER simulation produces ~20TB restart data per hour using ½ Summit.
  - Takes < 1 minute for writing on GPFS
  - Takes ~ 1 sec for writing on NVMe
- ADIOS is fast enough to write out 50PB a day for XGC, but way too big for GPFS → In situ, streaming data analysis is in progress.

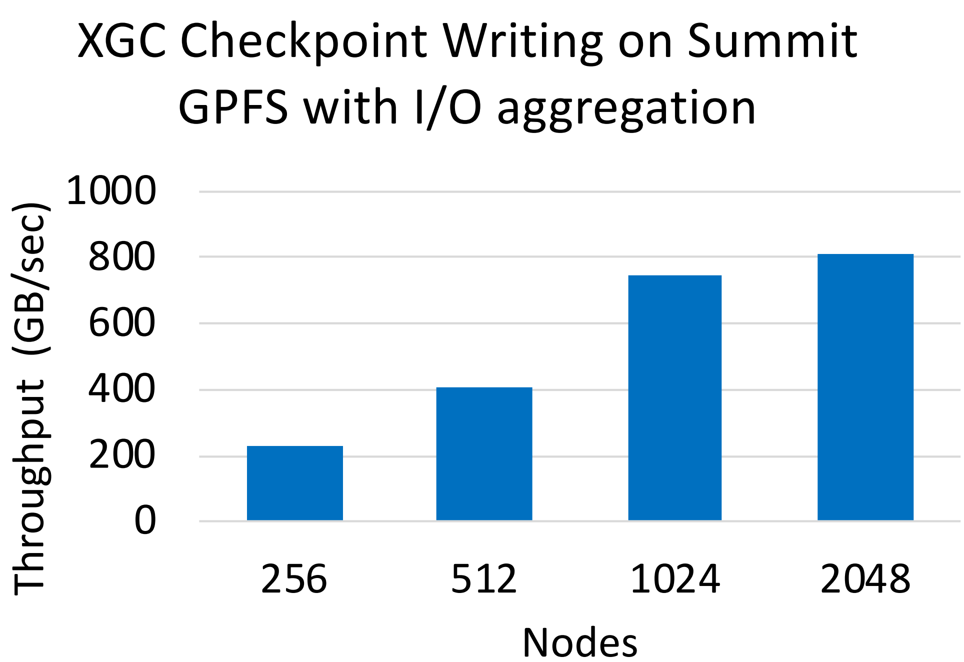

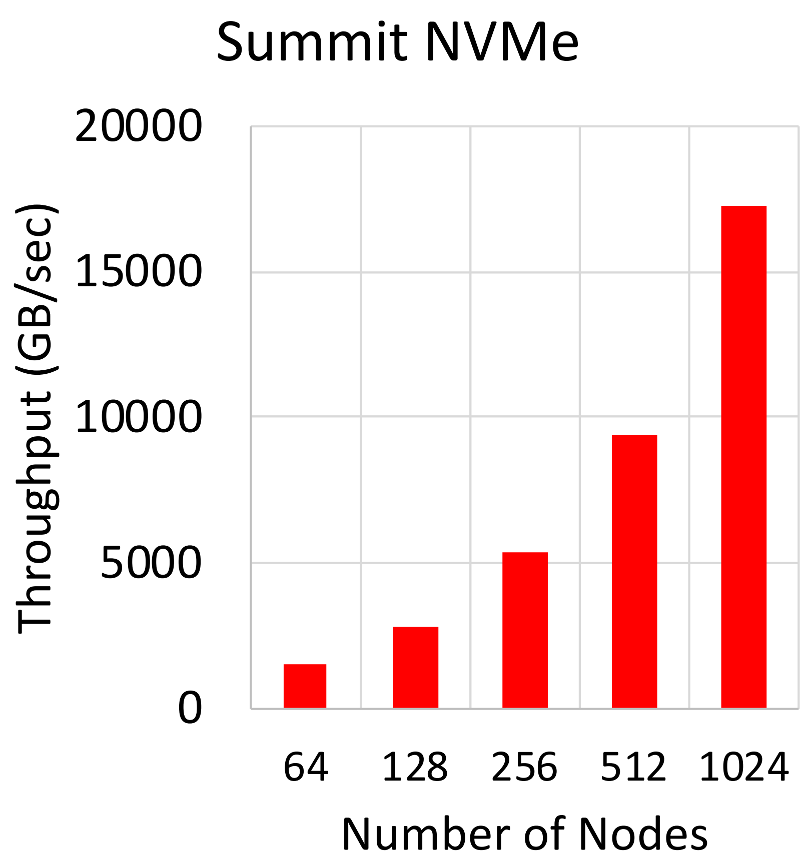

RAPIDS(2): Code coupling (S. Klasky)

- XGC is coupling with
  - the MHD code M3Dc1 for
    - magnetic equilibrium reconstruction as pedestal buildup
    - edge localized mode crash of the pedestal
  - the plasma-wall interaction code hPIC for
    - Plasma influx to the surface
    - Particle outfux from the surface
    - As described in the physics slides
- ADIOS enables code couplings
  - Coupling manager
  - In-situ staging
  - In-situ visualization
  - Coupling performance monitoring
  - Heterogeneous memory/process utilization

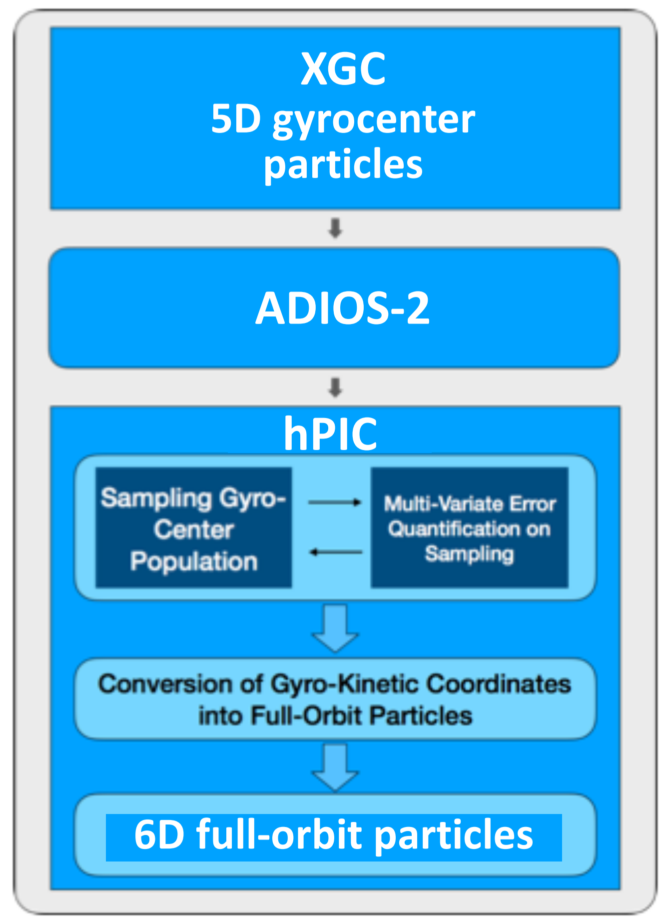

RAPIDS(3): ADIOS Coupling with analysis codes (S. Klasky)

XGC asynchronously offloads the analysis computations to designated processors via ADIOS

ADIOS enables:
- In-memory analysis
- Inter-network analysis
- WAN-coupled analysis

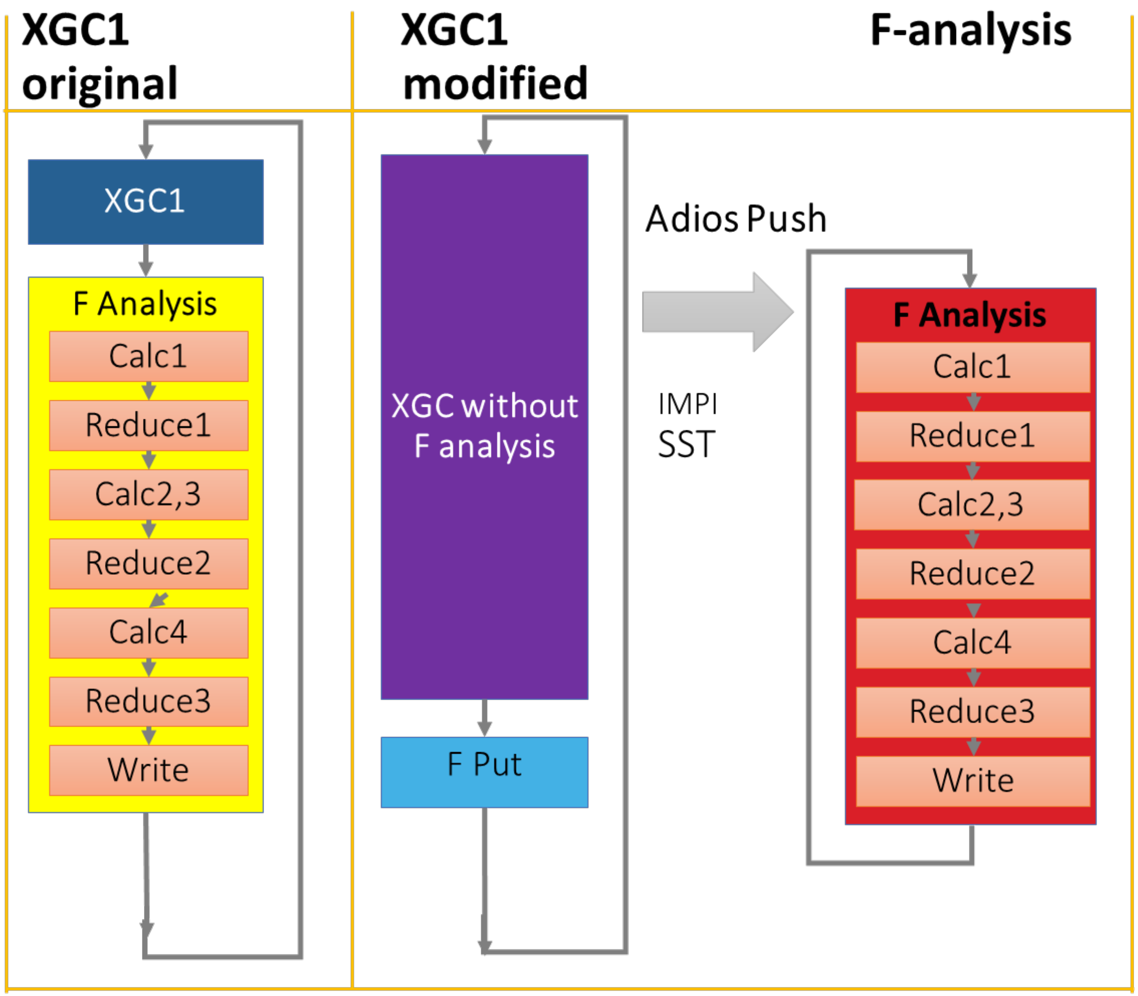

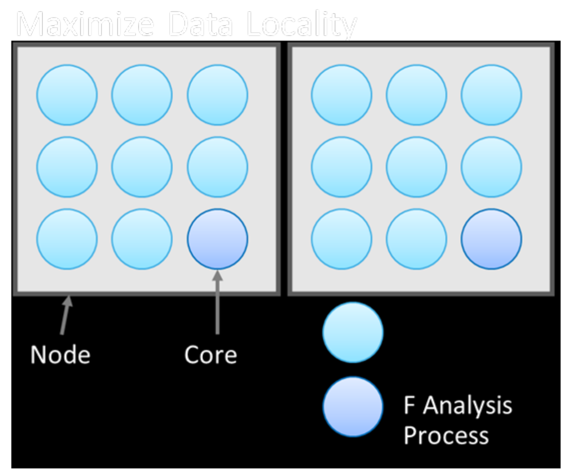

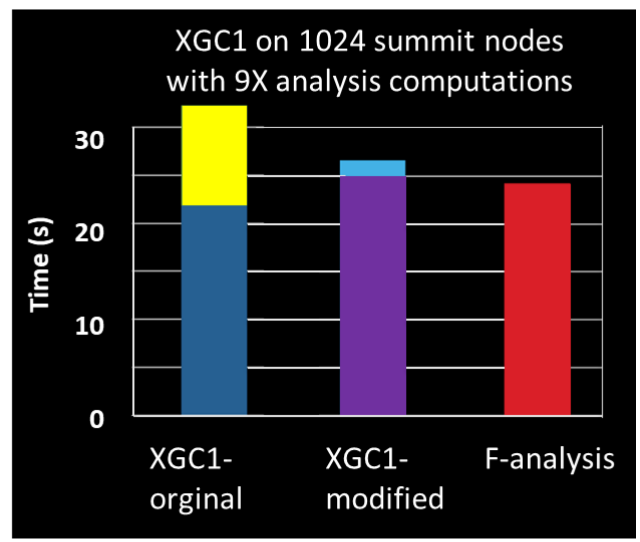

Example Data Analysis Code of XGC1 Turbulence

I. Keramidas[1], J. Myra[2], S. Parker[1]

Univ. of Colorado[1], Lodestar Corp.[2]

- One XGC1-run requires 10^{17} flop-hrs
- Enormous scientific interest in the data
- Powerful data science tools are widely available

Detailed analysis of neoclassical and turbulent particle fluxes

Turbulent particle flux at the last closed flux surface is shown

[I. Keramidas, et al., *Phys. Plasmas* **25** 072306 (2018)]

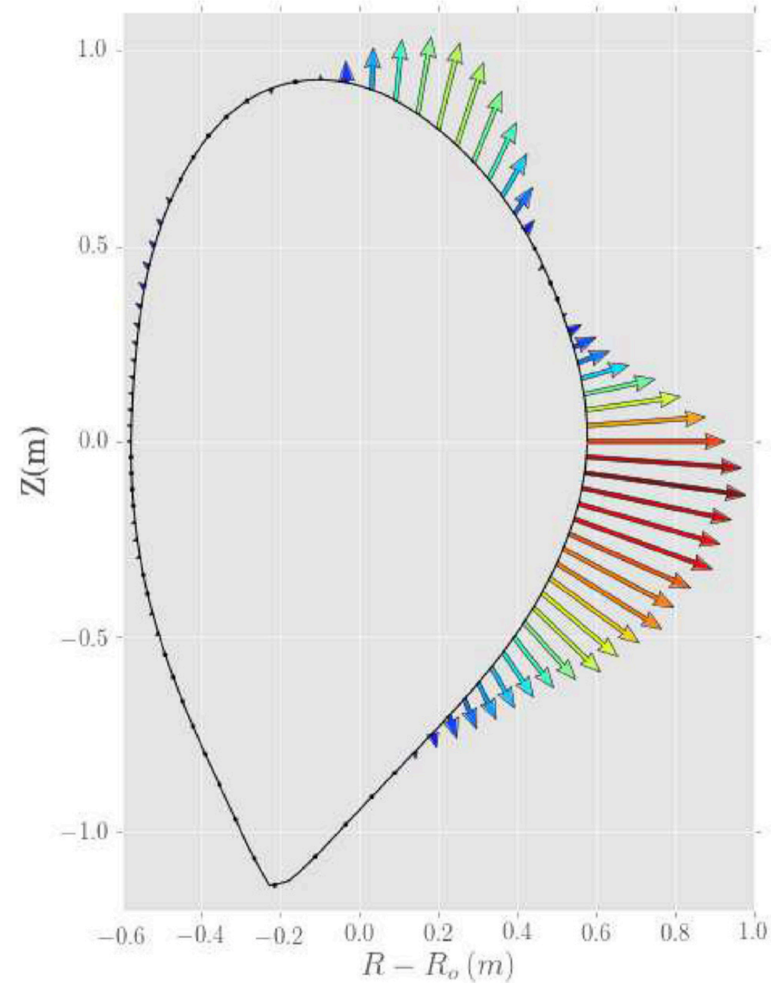

FIG. 5. Poloidal pattern of turbulent radial $E \times B$ flux [associated dataset available at https://doi.org/10.5281/zenodo.1230040] (Ref. 41).

Other collaborations with RAPIDS

- Software process (Hoffman/Galbreath, Kitware)
  - Continuous Integration (CI) testing system
  - Incorporate a modern CMake build
  - Git workflow incorporated with CI
  - Integrate CDash into github

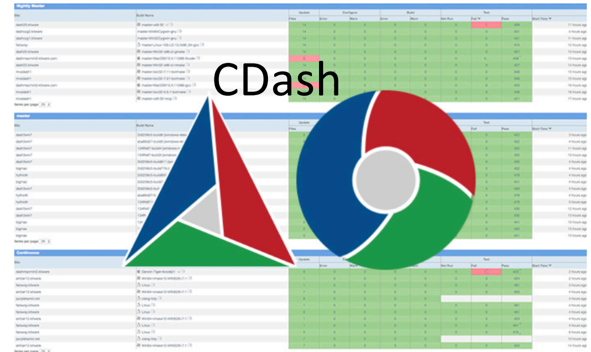

- EFFIS (Klasky, ORNL): End-to-end Framework for Fusion Integrated Simulation
  - Initial development in SciDAC-2 CPES (PI: Chang)
  - Integrated platform of services to compose, launch, monitor, analyze, and control coupled applications on all leadership HPCs
  - High Performant I/O for coupled codes
  - Process placement (inter or intra node)
  - Online dashboard functionality

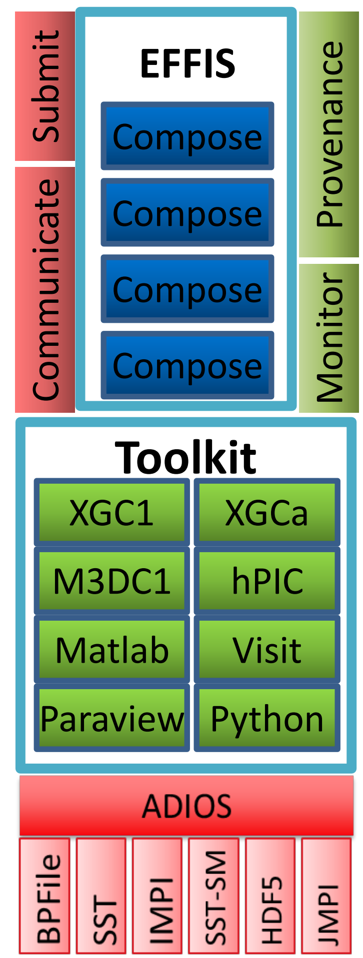

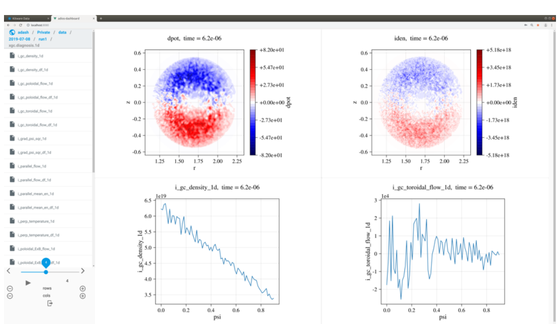

Example EFFIS specification file

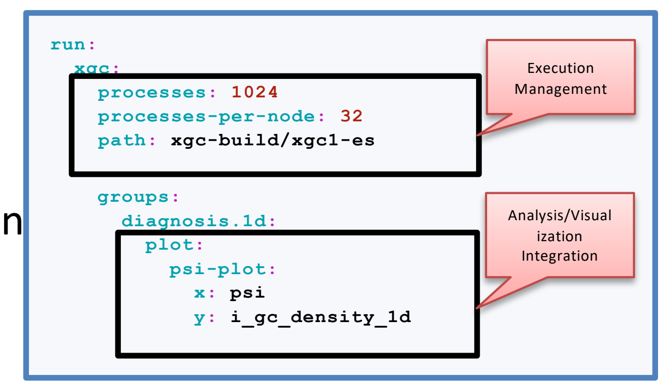

Summary

- HBPS is making scientific discoveries that would not have been possible without the SciDAC framework and the US Leadership Class Computers

  Since last SciDAC Conference
  - Many Invited/Plenary Talks at major scientific conferences, including Smokey Mt 2018, IAEA-FEC 2018, APS-DPP 2019, International Data Driven Plasma Sciences 2019, KSTAR 2019, PACS 2019, + many more
  - 7 SciDAC-funded physics papers + many ASCR papers, (co-authored)
- Strong collaboration with FASTMath and RAPIDS in many subjects
- Other ASCR collaborations outside of SciDAC Institutes also strong
  - **R. Moser (**Texas): UQ, Multifidelity Monte Carlo (MFMC) methods
  - **L. Chacon** (LANL): math leadership
  - **V. Carey** (U. Colorado Denver): particle re-sampling
  - **C. Hauck** (ORNL): super time-stepping
- Preparation for WDM coupling with most Fusion SciDACs
- AI/ML is becoming an important part of the project
  - In-situ data analytics
  - Validation
  - Simulation acceleration
  - Anchored ML for prediction, by adding high-fidelity simulation data