HBPS
High-fidelity
Boundary
Plasma Simulation

SciDAC PI Meeting, July 23-24, 2018

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

C.S. Chang[1],

S. Parker[2], N. Ferraro[1], M. Greenwald[3], G. Hammett[1], S. Ku[1],
L. Chacon[4], M. Adams[5], C. Hauck[6], J. Hittinger/L. Ricketson[7],
R. Moser[8], M. Shephard[9]
S. Klasky[6], J. Choi[6], E. Dazevedo[6], W. Hoffman[10], 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

ALCF

OLCF Oak Ridge Leadership Computing Facility

NERSC

Outline+,*

- **Leadership in HBPS and the project OV**
- **The XGC gyrokinetic code**
- **Example scientific discoveries**
  - **L-H transition**
  - **Divertor Heat-Flux Width**
- **Enabling Techonolgies**
  - **Performance optimization**
  - **Data management**
  - **Applied mathematics**
- **Integration with other Fusion SciDAC codes for WDM**
- **Summary**

+Funding provided by US DOE
*Computational resources provided by OLCF, ALCF and NERSC

Leadership in HBPS

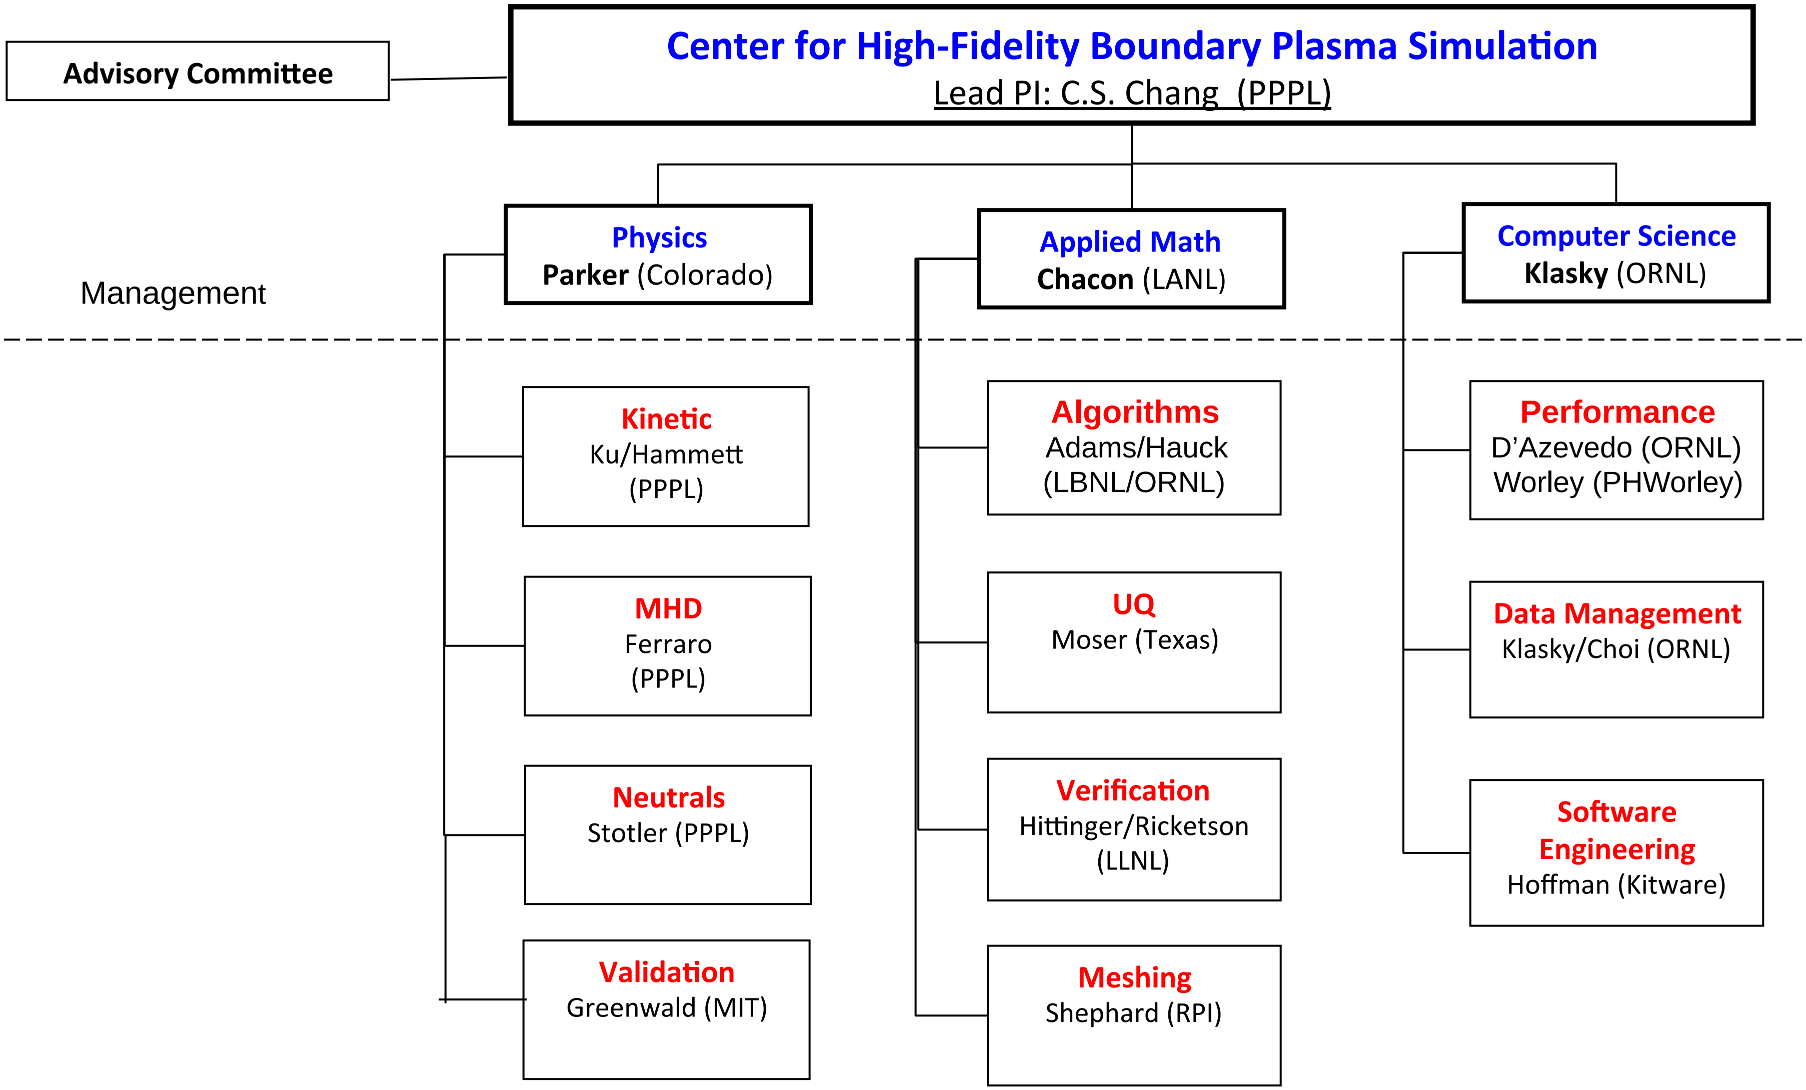

Center for High-fidelity Boundary Plasma Simulation

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

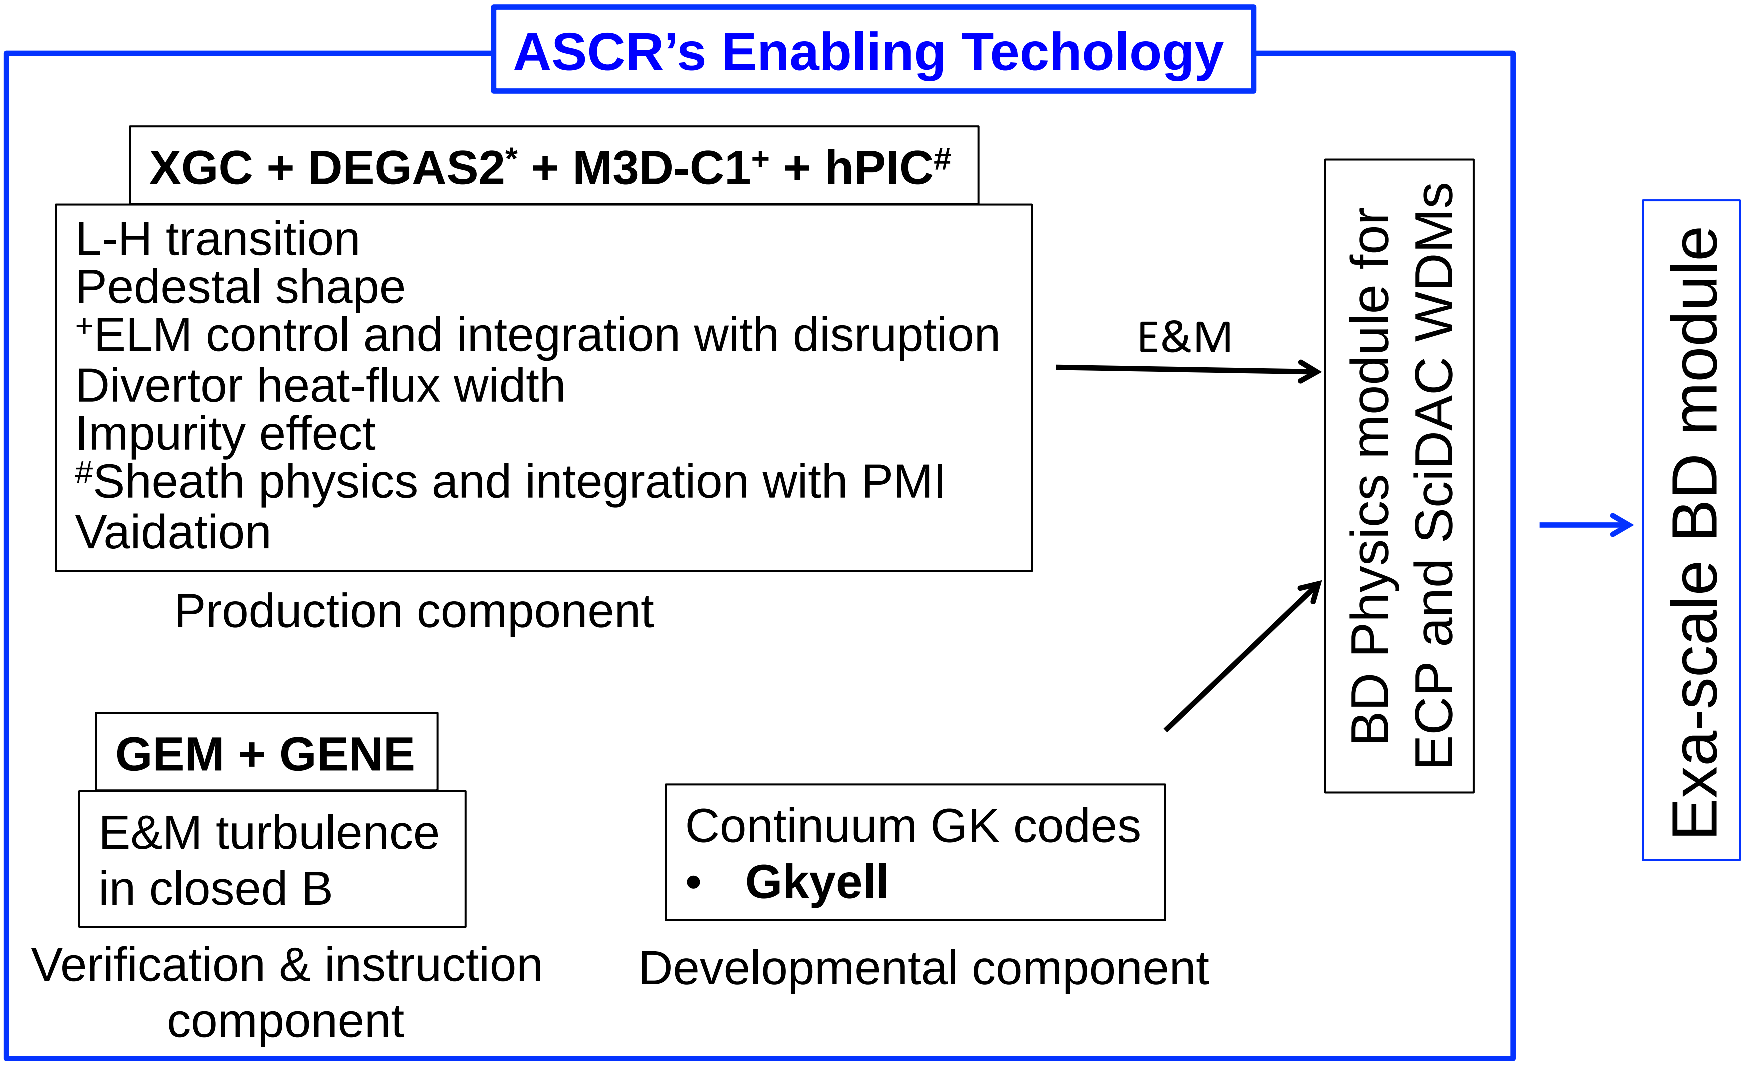

***DEGAS2 is coupled into XGC as a subroutine.**

Outline+,*

- Leadership in HBPS and the project OV
- **The XGC gyrokinetic code**
- Example scientific discoveries
  - L-H transition
  - Divertor Heat-Flux Width
- Enabling Techonolgies
  - Performance optimization
  - Data management
  - Applied mathematics
- Integration with other Fusion SciDAC codes for WDM
- Summary

The XGC Gyrokinetic Code

- **Particle-in-Cell, with added continuum technology**
- **In contact with material wall**
  - Far-from-equilibrium (non-Maxwellian)
  - Neutral particles
- **Magnetic X-point and separatrix (q→∞)**
  - X-point orbit loss from pedestal
- **Multi-scale, multiphysics in space-time space**
- **Unstructured triangular mesh**
- **PETSc (only ~2% of total computing time)**
- **Large simulation-size (≳10k particles per grid-vertex) per time-step**
- **Total-f XGC has been developed to study this kind of complicated plasma**
- **Most of the production runs are on ~90% Titan, ~50% Theta, and ~50% Cori.**
- **XGC is not only a SciDAC code, but also in all three existing exa- or pre-exa programs (CAAR, Aurora-ESP, NESAP), ECP, and INCITE**

XGC1

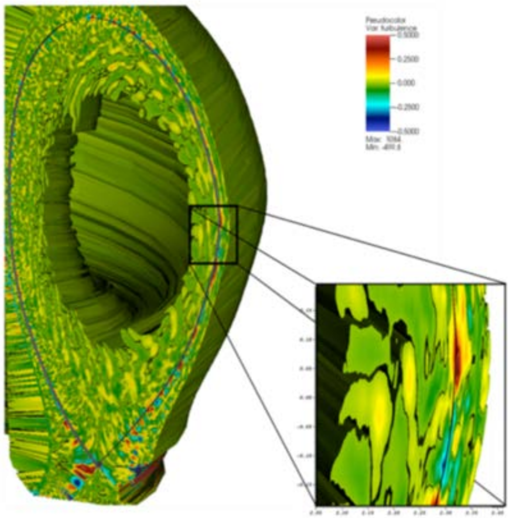

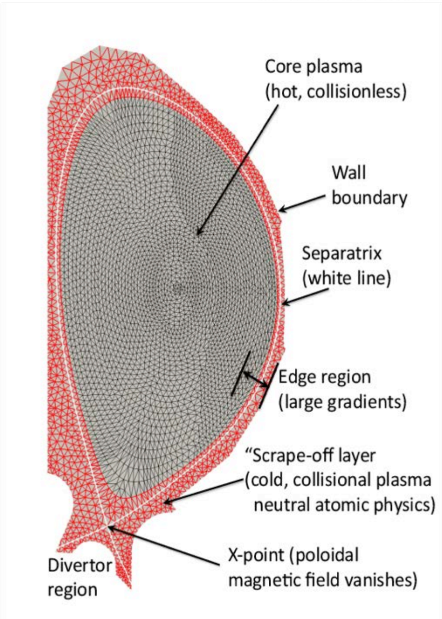

In the core plasma, *f* evolves slowly

For this argument, let's use the drift kinetic equation for simplicity

$$\partial f/\partial t+(\mathbf{v}_{||}+\mathbf{v}_d)\cdot \nabla f + (e/m)E_{||}\, v_{||}\, \partial f/\partial w = C(f,f) + Sources/Sinks.$$

In near-thermal equilibrium, take the "transport ordering" (= diffusive ordering):

$$\partial f/\partial t=O(\delta^2),\ S=O(\delta^2),\ with\ \delta<<1$$

- Let $f=f_0+\delta f$, with $\delta f / f_o=O(\delta)$, $\delta <<1$, $v_d/v_{||} = O(\delta)$, $E_{||}/m = O(\delta\ or\ \delta^2)$

$O(\delta^0)$: $v_{||} \cdot \nabla f_0 = C(f_0,f_0) \rightarrow f_0=f_M$: H-theorem

$O(\delta^1)$: $\partial\delta f/\partial t + v_{||}\cdot \nabla\delta f + v_d \cdot \nabla f_0 + (e/m)\, E_{||}v_{||}\, \partial f_o/\partial w = C(\delta f)$

- ✧ Perturbative kinetic theories then yield transport coefficients $=O(\delta^2)$
- ✧ In this case, fluid transport equations ($f_o \rightarrow n,T$) can be used with analytic or delta-f kinetic closures

→ ***δf*-GK simulation is cheaper per physics time** (small computers)**, but equilbrates on a slow time scale $O(\delta^1\omega_{bi})^{-1}$ ~ *ms***: *Core GK simulation time scale*

A meaningful time evolution of f_0 can only be obtained in a long "transport-time" scale $O(\delta^2\omega_{bi})^{-1}$: Not yet reachable by GK simulation; Multiscale time integration is needed.

In edge, f equilibrates in zeroth-order time-scale

- Ion radial orbit excursion width ~ pedestal & scrape-off layer width; unconfined orbits with neutral recycling → Non-Maxwellian

All terms can be large: ~ either *$O(\omega_{bi})$ or $O(\nu_C)$*

- $\mathbf{v}_{||}\cdot \nabla f \sim \mathbf{v}_d \cdot \nabla f \sim C(f,f) \sim eE_{||}v_{||}/m\, \partial f/\partial w \sim O(\omega_{bi})$ ~ 0.05 ms in DIII-D
- *f* equilibrates very fast: $\partial f/\partial t + (\mathbf{v}_{||}+\mathbf{v}_d)\cdot \nabla f\,(e/m) + E_{||}v_{||}\partial f/\partial w = C(f,f)+S$

- Fast-evolving nonthermal kinetic system
  - **Fluid equations (with closure ordering) could give a long time scale.**

Edge turbulence around the separatrix saturates before the central core turbulence has even started to form

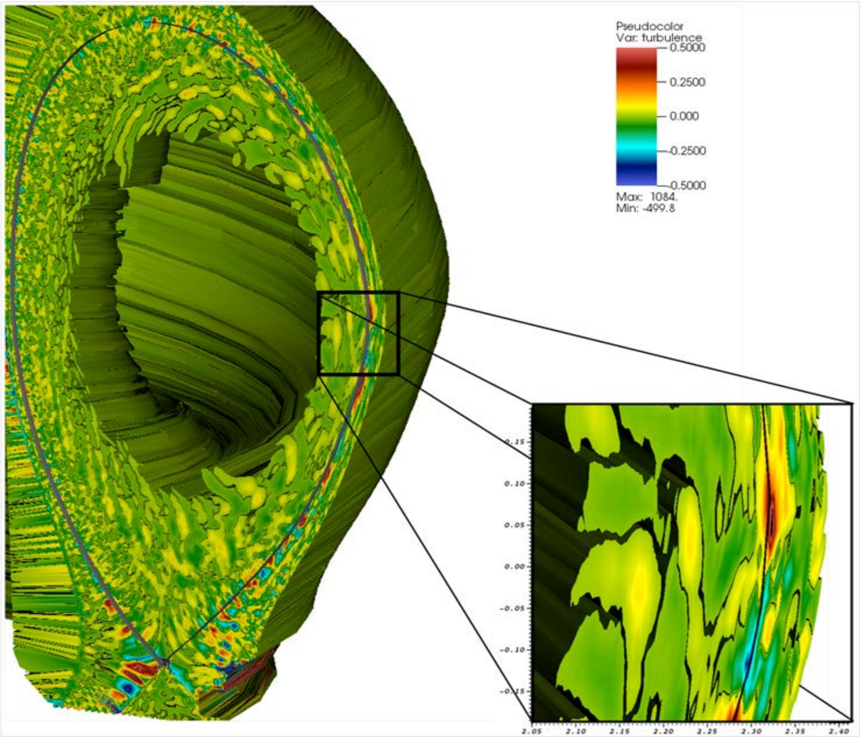

Ideal for extreme scale computing: big physics in short physics-time (small number of time steps)

Outline+,*

- Leadership in HBPS and the project OV
- The XGC gyrokinetic code
- **Example scientific discoveries**
  - **L-H transition**
  - Divertor Heat-Flux Width
- Enabling Techonolgies
  - Performance optimization
  - Data management
  - Applied mathematics
- Integration with other Fusion SciDAC codes for WDM
- Summary

Gyrokinetic L-H bifurcation study, using a low-beta C-Mod L-mode plasma in XGC1*,#

Plasma input condition

- C-Mod #1140613017 in L-mode, single-null (P_{LH}~1-1.5MW)
- β_e ≈0.01% < m_e/m_i in the bifurcation layer
- Ion magnetic-drift direction was flipped toward the divertor in the first study (favorable direction), then flipped back in the second study.

Simulations include the most important multiscale physics

- Neoclassical kinetic physics
- Nonlinear electrostatic turbulence
- ITG, TEM, Resistive ballooning, Kelvin-Helmholtz, other drift waves
- Neutral particle recycling with CX and ionization crossections
- Realistic diverted geometry

*Chang, Invited IAEA-FEC2016, PRL2018, and
#Ku, Invited APS-DPP2017 (PoP2018, cover), Invited IAEA-FEC2018

(EM correction to the present result is left for a near-future work.)

Input: an L-mode plasma from C-Mod (beta~0.01%)

Edge temperature increases from heat accumulation

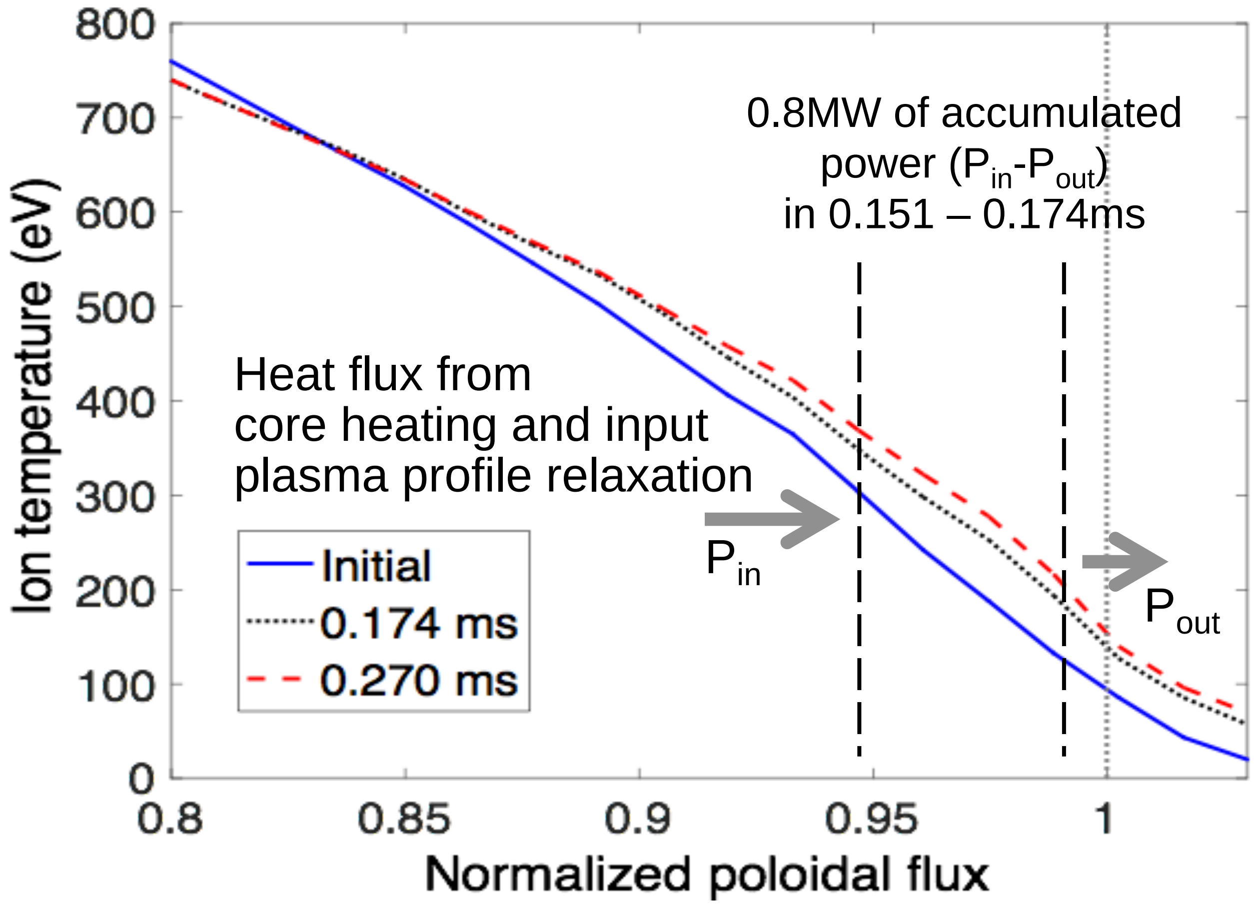

Gyrokinetic observation of the L-H bifurcation in a C-Mod model plasma

1. At t~0.175-0.21ms, lower frequency turbulence decays and higher frequency turbulence appears: through conservative **Reynolds work** via eddie tilting-absorption.
2. At t>0.21ms, suppression of all-frequency turbulence follows, with higher frequency part disappearing: through dissipative ExB shearing **w/o Reynolds** force.

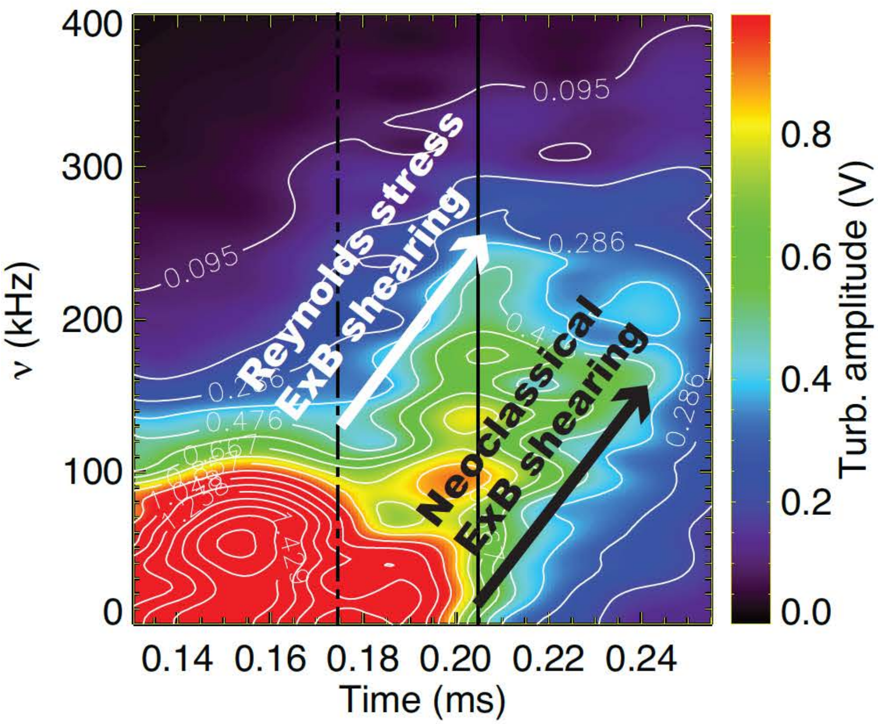

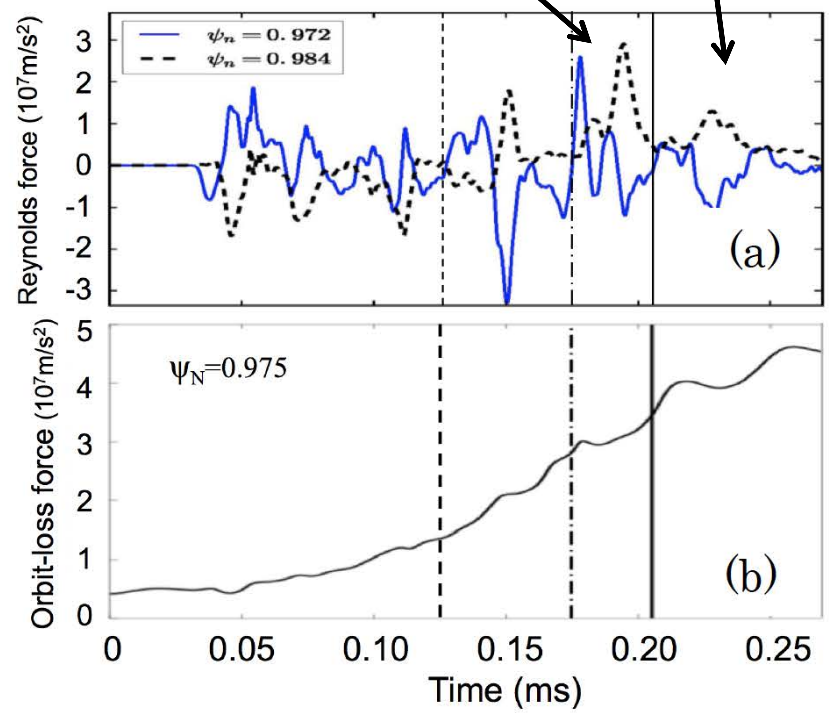

When the magnetic drift is changed back to the unfavorable direction, more interesting physics have been discovered.

1. GAM activity is stronger in the edge bifurcation layer ($0.96<\Psi_N<0.98$)*

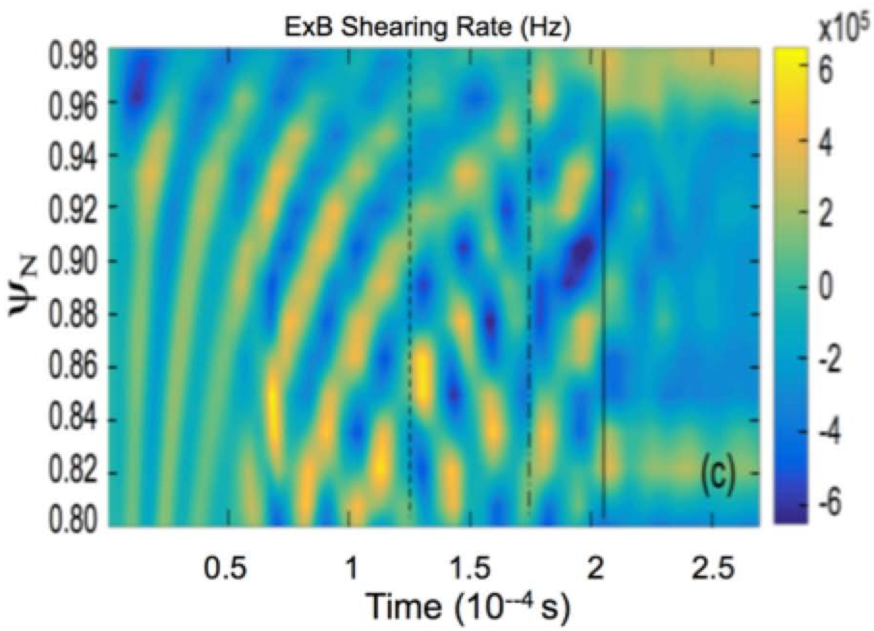

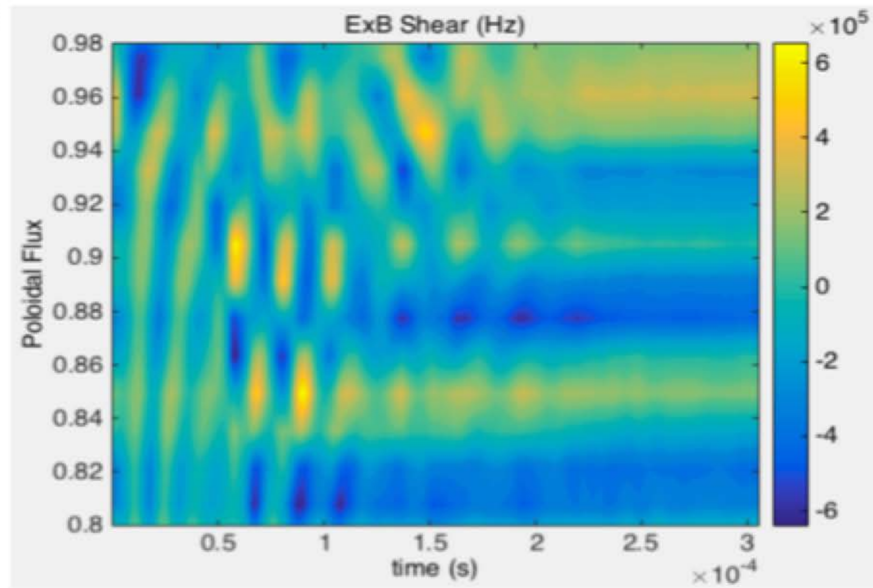

*Beginning to be observed experimentally [Czigler PRL2017]

2. Quasi-coherent modes appear even in the electrostatic simulation (but appears not to last long → EM needed?)

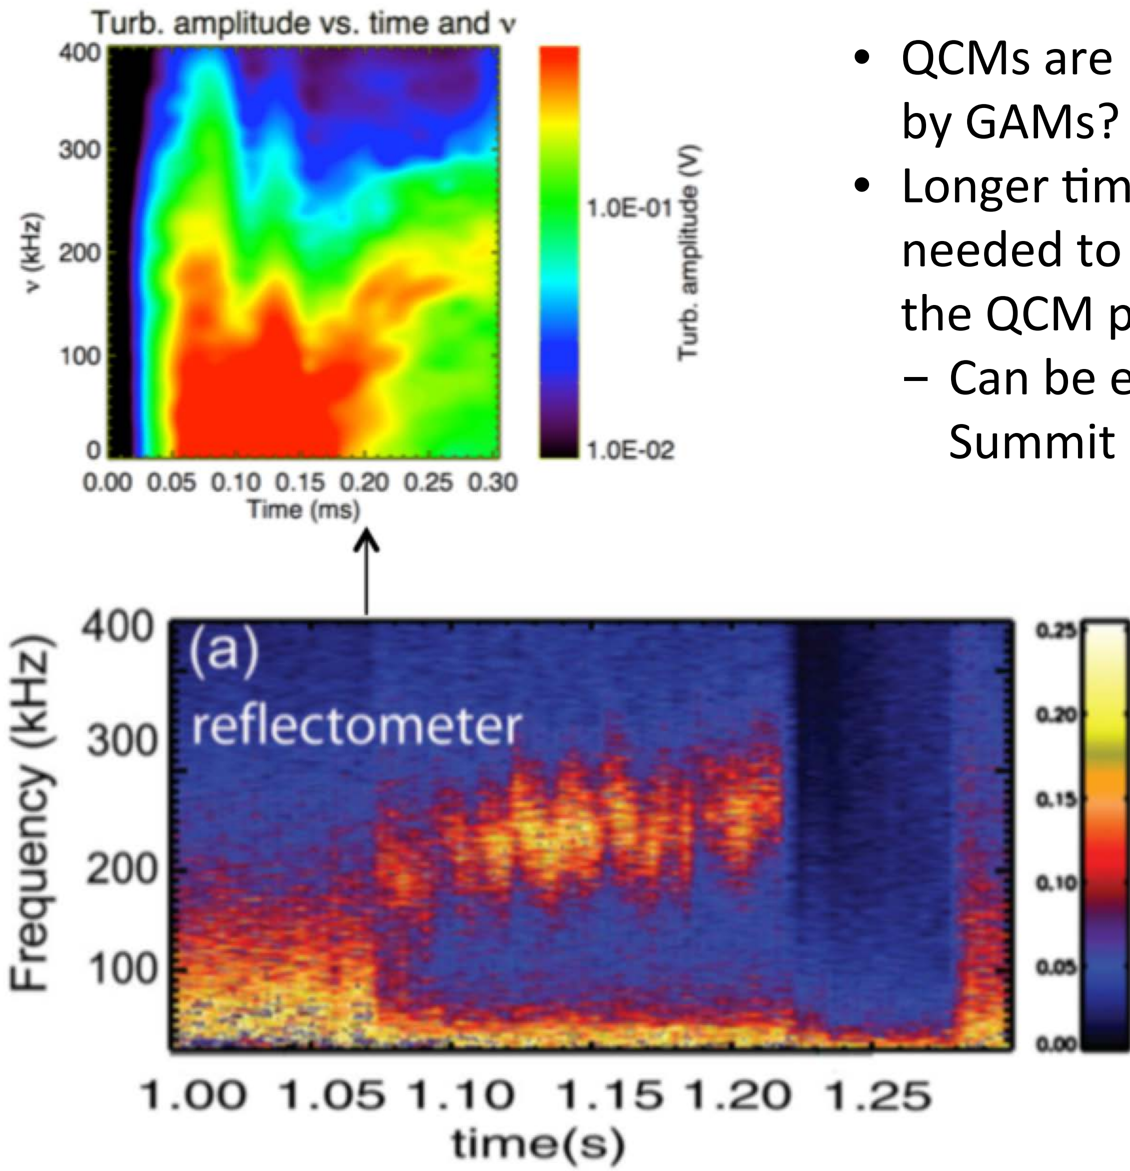

- QCMs are ringing modes induced by GAMs?
- Longer time simulation and EM needed to get to the bottom of the QCM physics:
  - Can be enabled by the 200PF Summit computer, soon.

Larger HPCs and more enabling technologies are needed for the first-principles-based prediction of the L-H bifurcation efficiency in ITER and fusion reactors

- Will the weak neoclassical effect due to the small ρ_i/a [1] in ITER hinder the second (dissipative) stage of the L-H bifurcation process?
  - Can the X-point orbit-loss effect help strengthen the mean ExB shearing and help the L-H transtion and the H-mode pedestal formation?
  - Is the planned external heating power strong enough to induce the needed L-H transition in ITER?
- Can the edge GAMs be used to control the L-H transition when needed?
- Can we utilize the I-mode in the future fusion reactor operation?
- How important is the EM effect in the L-H bifurcation dynamics?

☐ Longer physics-time simulation (for pedestal buildup) and/or higher flop-rate simulation (EM and ITER) are needed.

☐ Help needed from enabling technology: algorithm optimization, error reduction, performance enhancement, I/O improvement, on-memory data analysis and reduction, fault resilience, load balancing, machine learning, platform portability, UQ for extreme-scale simulation …

[1] Noticed in recent publications by Kotschenreuther-Hatch, and Chang

Outline$^{+,*}$

- Leadership in HBPS and the project OV
- The XGC gyrokinetic code
- **Example scientific discoveries**
  - L-H transition
  - **Divertor Heat-Flux Width**
- Enabling Techonolgies
  - Performance optimization
  - Data management
  - Applied mathematics
- Integration with other Fusion SciDAC codes for WDM
- Summary

GK simulation of Divertor heat-flux width:Validation on present devices and prediction for ITER

- Discharges are selected for wide distribution of $B_{pol,OM}$.
- Experimental eqdsk data are imported into XGC.

Shot	Time (ms)	B_T (T)	I_P (MA)	$B_{pol,OM}$ (T)
NSTX 132368	360	0.4	0.7	0.20
DIII-D 144977	3103	2.1	1.0	0.30
DIII-D 144981	3175	2.1	1.5	0.42
C-Mod 1100223026	1091	5.4	0.5	0.50
C-Mod 1100223012	1149	5.4	0.8	0.67
C-Mod 1100212023	1236	5.4	0.9	0.81
JET 79692		3.56	4.5	0.89

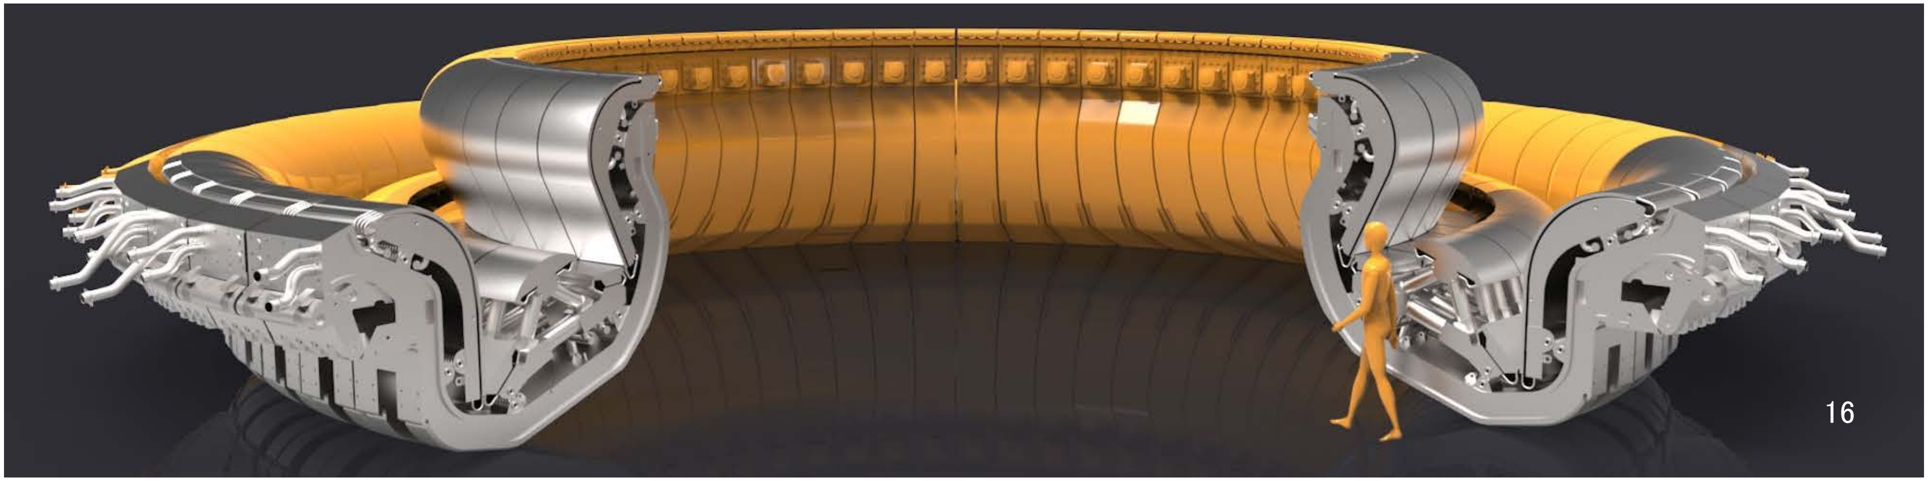

- The XGC-predicted divertor heat-flux width $\boldsymbol{\lambda_q}$ **has been well-validated** against various existing tokamak data.
- However, XGC predicts about **6X wider $\boldsymbol{\lambda_q}$ for ITER** than the regression value by Eich et al.: Why?
- Edge turbulence is **blob type** in the present tokamaks, but **streamer type** in the full-current ITER.
- To check if the enhanced λ_q in the full-current ITER is from the "absolute **size effect**" or from the "**$\boldsymbol{\rho_i/a}$ effect**," a reduced-current "first-phase" ITER has been simulated → λ_q agrees with the present tokamaks → **$\boldsymbol{\rho_i/a}$ effect**.

- The "absolute size effect" is related to the parallel physics and the neutral particle transport
- The "B_{pol} effect" is mostly from the perpendicular physics

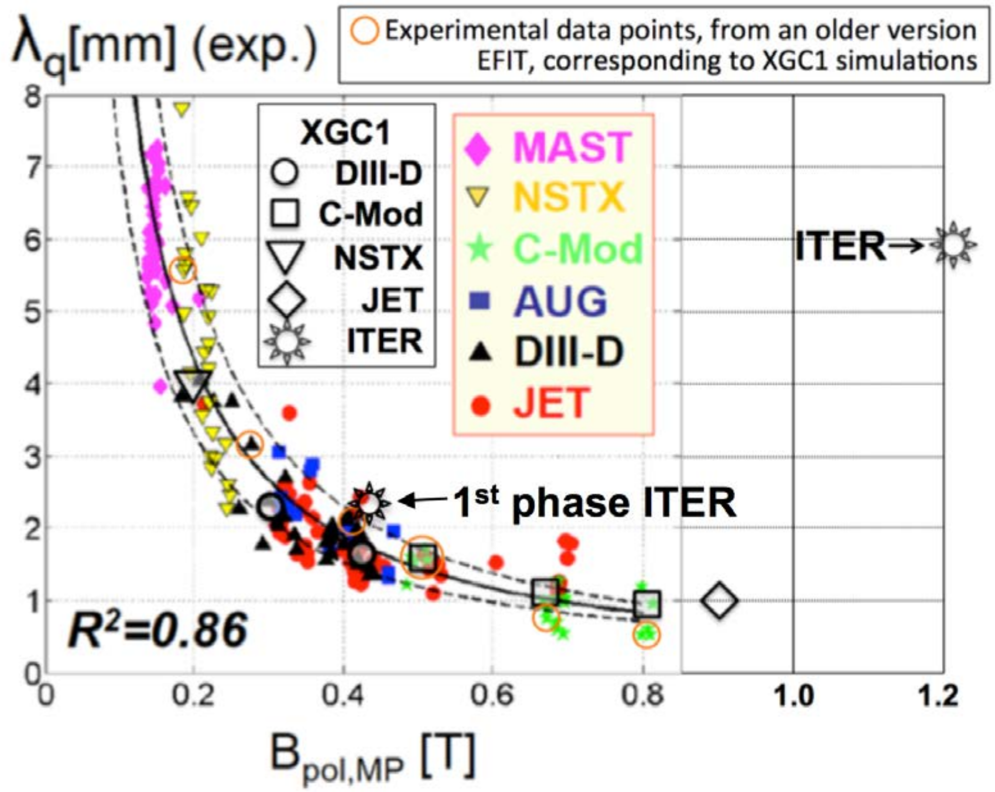

Need UQ/ML research to maximize divertor heat-flux width

- Sensitivity to B_{pol} has been well-known from experimental data
- Sensitivity to ρ_i/a has been discovered in our XGC simulation
- New: sensitivity to plasma shape and radiative cooling seen in XGC
- UQ/ML on large-scale simulation: Need a multi-fidelity method

NSTX-U plasma shows sensitivity of λ_q to plasma shaping and divertor cooling.

Maximize λ_q → huge impact on ITER operation and success

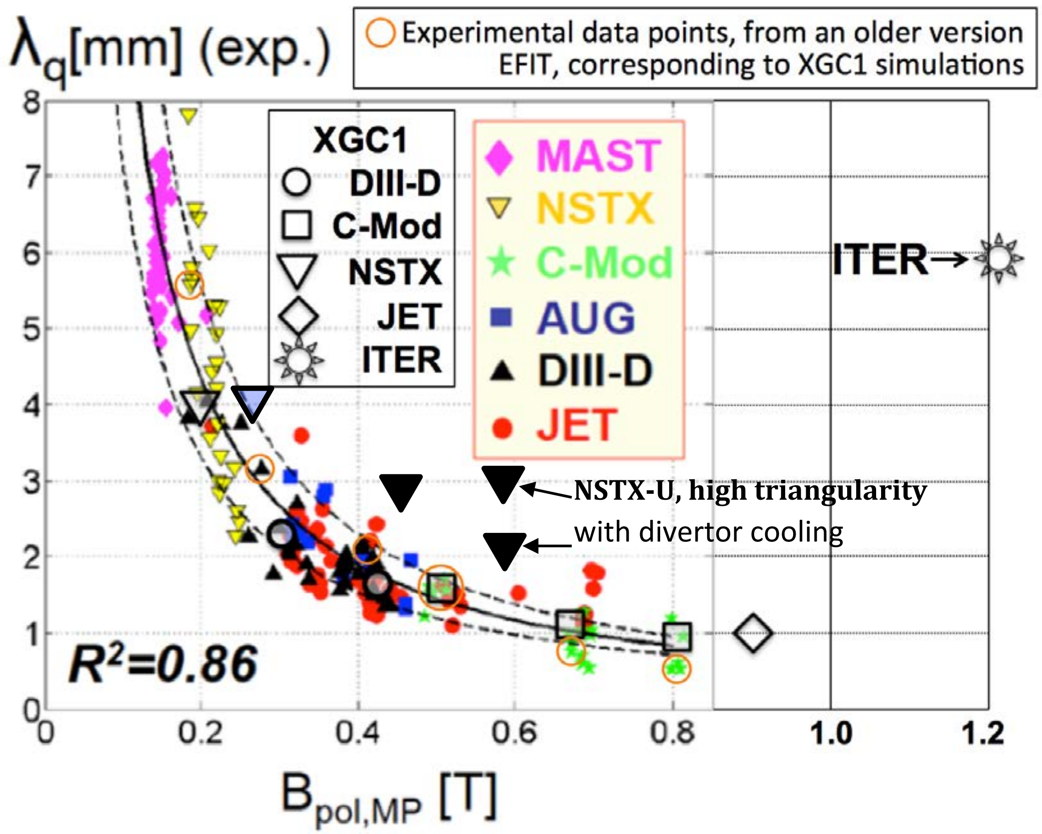

Evidence for an edge physics bifurcation between the higer and lower ρ_i/a values.

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.

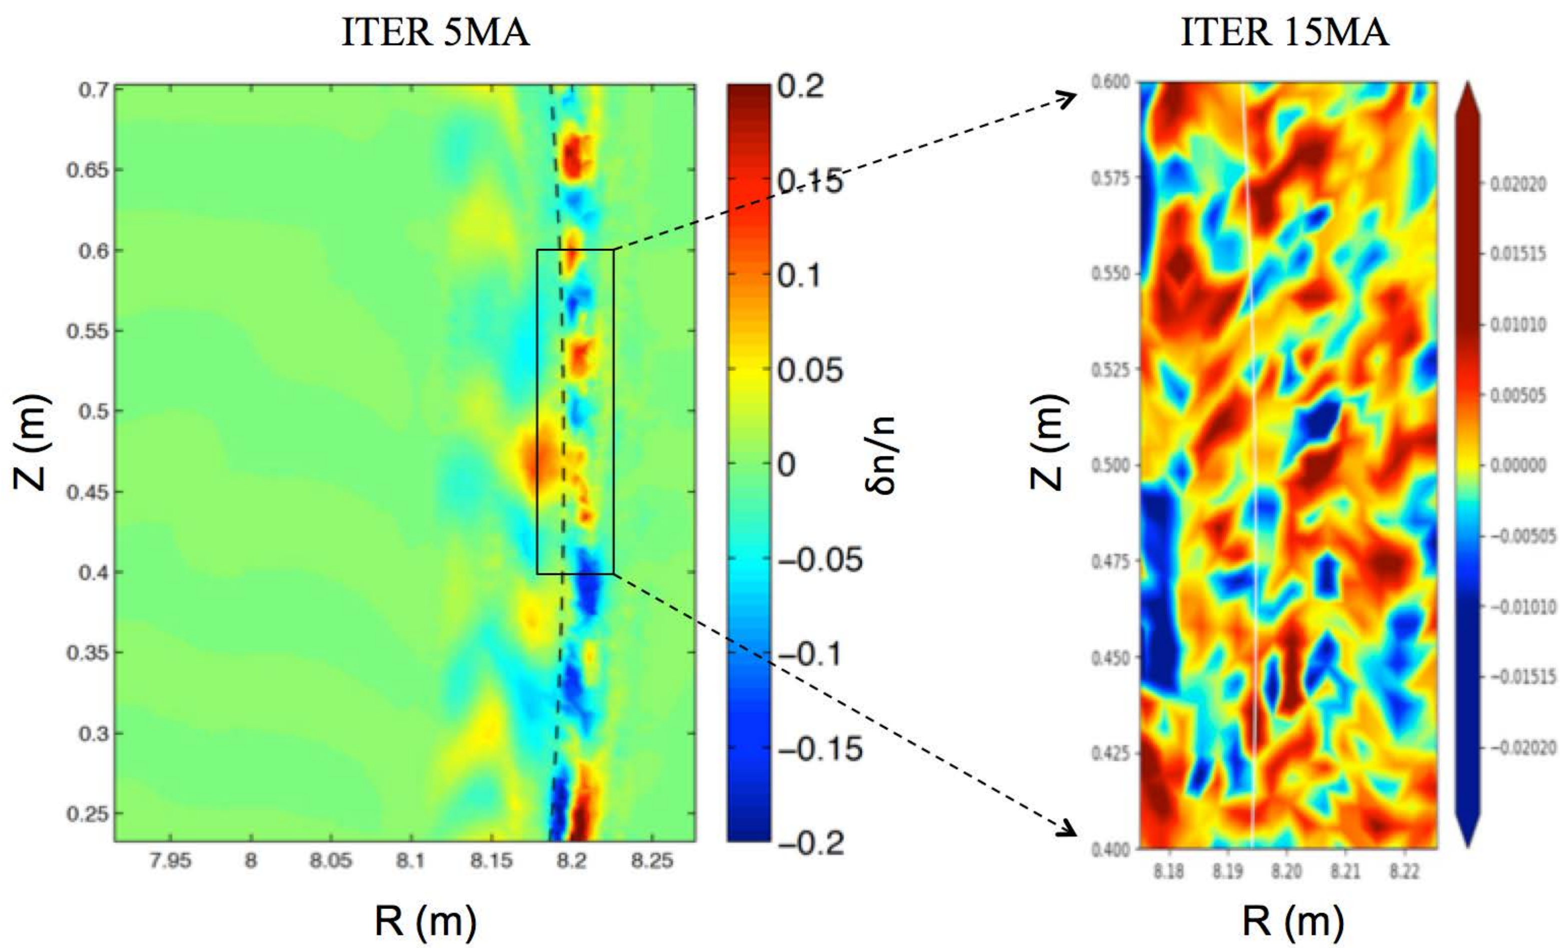

A careful study will be performed in the near future: needs large HPC time.

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.

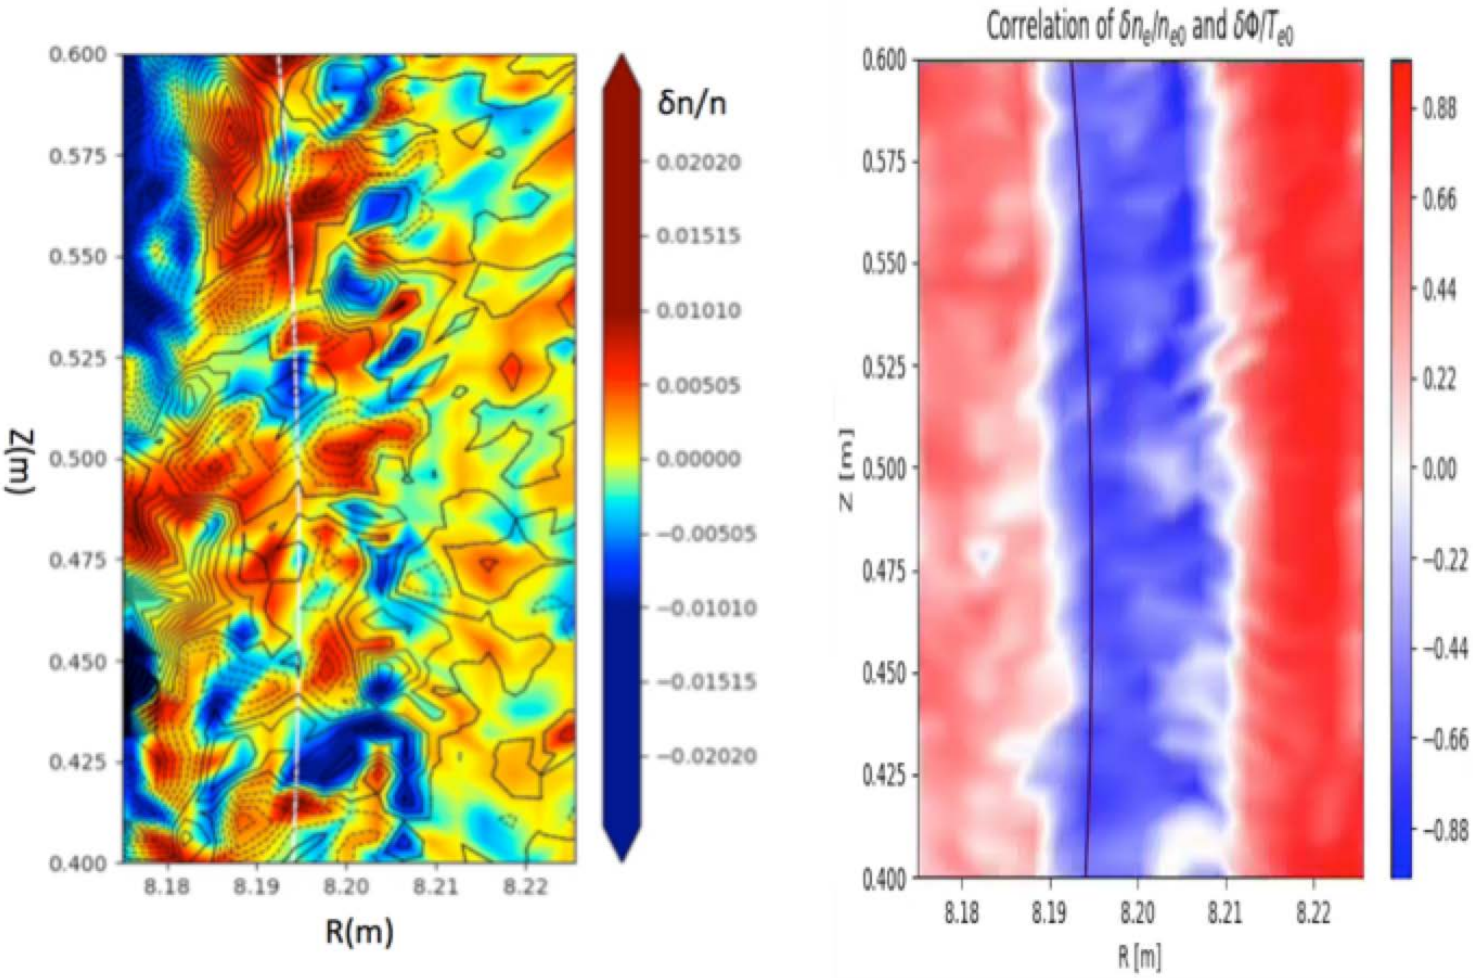

Outline+,*

- Leadership in HBPS and the project OV
- The XGC gyrokinetic code
- Example scientific discoveries
  - L-H transition
  - Divertor Heat-Flux Width
- **Enabling Techonolgies**
  - **Performance optimization**
  - **Data management**
  - **Applied mathematics**
- Integration with other Fusion SciDAC codes for WDM
- Summary

XGC Scales well on the new #1 Summit [Worley, D'Azevedo, …]

- XGC has been scaling well to the maximal Titan, Cori and Theta.
- CAAR project: XGC also scales well on the new world #1 Summit to the maximal available # nodes (2,048, near 50% capacity).
- Using a **present production** case (underusing the GPU capability) on 2048 Summit nodes, XGC shows 11.3x speedup on GPU+CPU from CPU only.

Multi-level parallelization:

MPI + OpenMP + CUDA + OpenACC

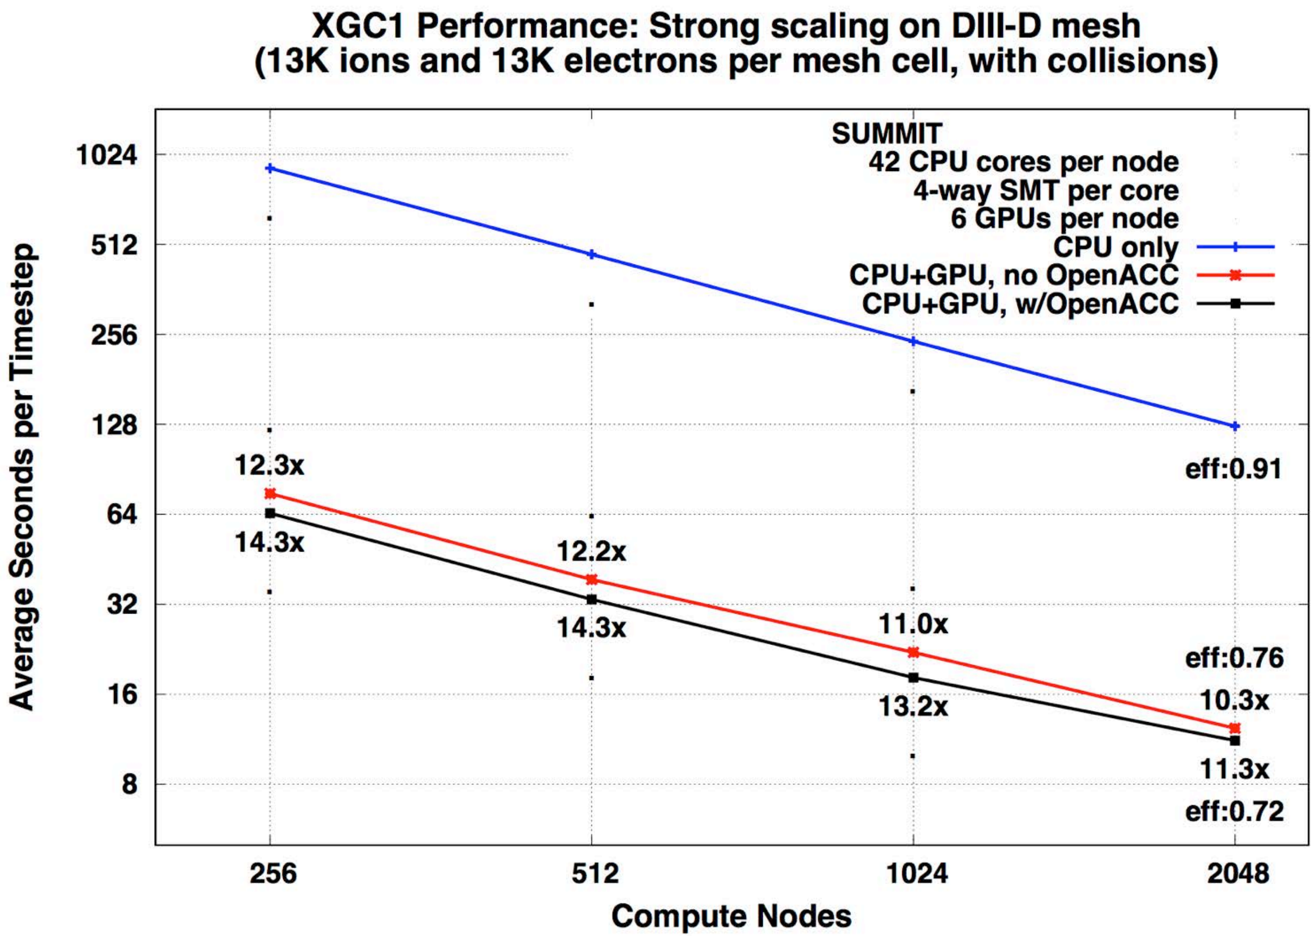

The production XGC is 3.8x faster than Titan on 2,048 Summit nodes, when matching #nodes to contain the same #GPUs.

(Theoretical ratio is ≈ 5)

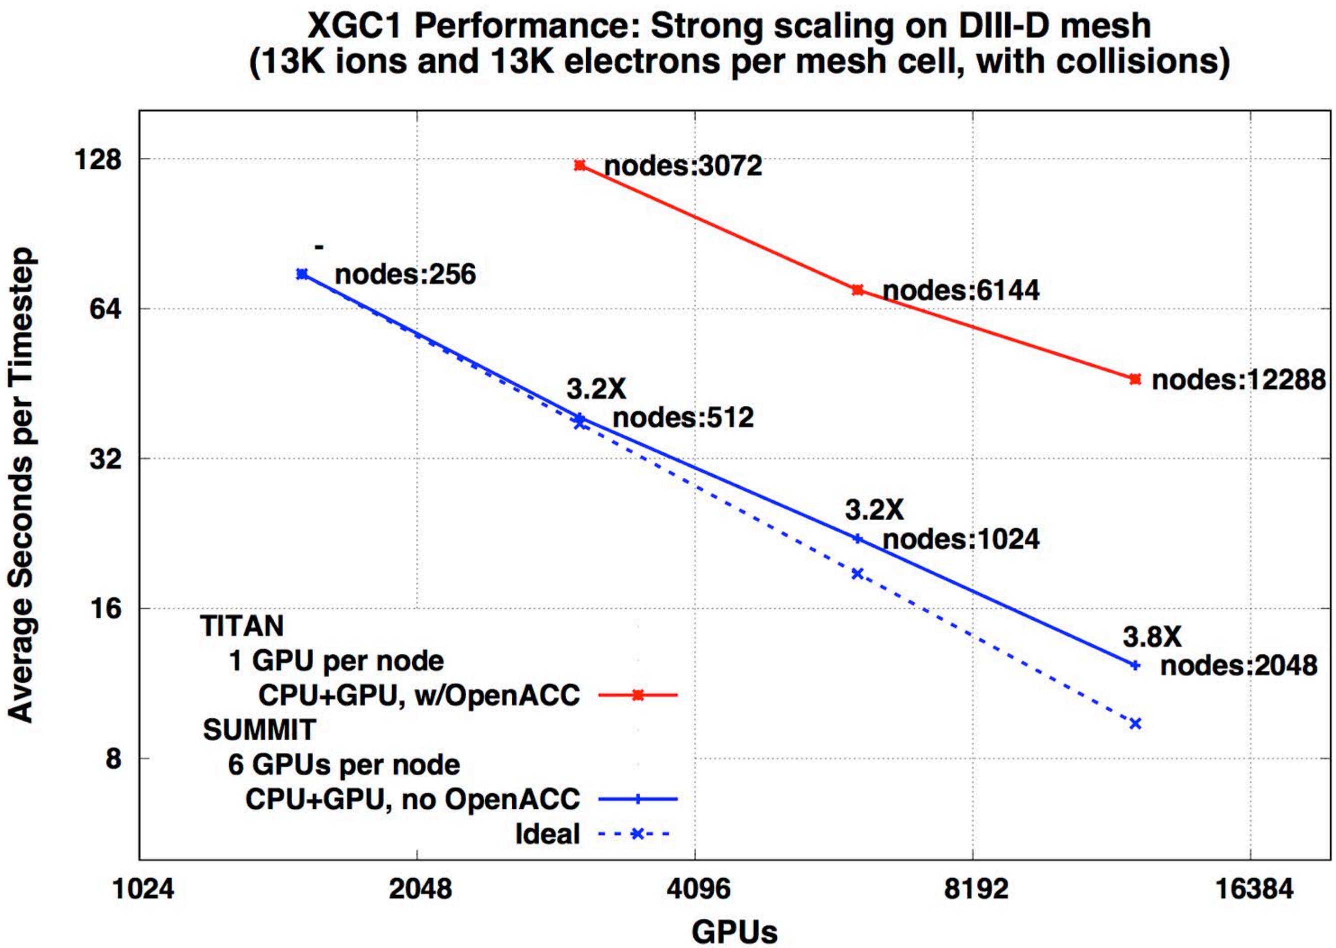

We will continue the scalability study to the full Summit machine and execute our early science study on Summit.

HBPS Data Management: I/O speed became a non-issue on all the modern major HPCs

- Extensive study of writing checkpoint-Restart data (weak scaling) and physics data (strong scaling) using ADIOS on Summit, Cori, Theta and TSUBAME3 shows
  – XGC's write time using NVRAM (Burst Buffers) are reduced to a few seconds on all major platforms → Not an issue.
- I/O time changes from 200 GB/s (Titan, Luster) to
  – 400 GB/s on Cori NVRAM
  – 300 GB/s (32 nodes on Summit-dev NVRAM) EST: to over 50 TB/s on Summit NVRAM
  – 3 TB/s on Theta NVRAM
  – 90 GB/s on TSUBAME-64 nodes NVRAM

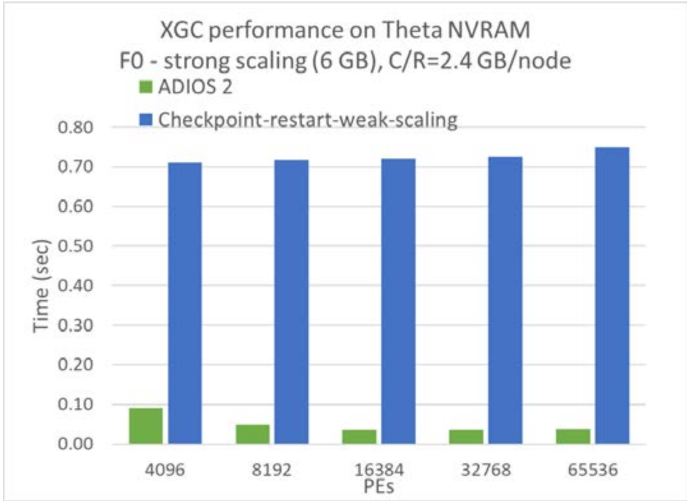

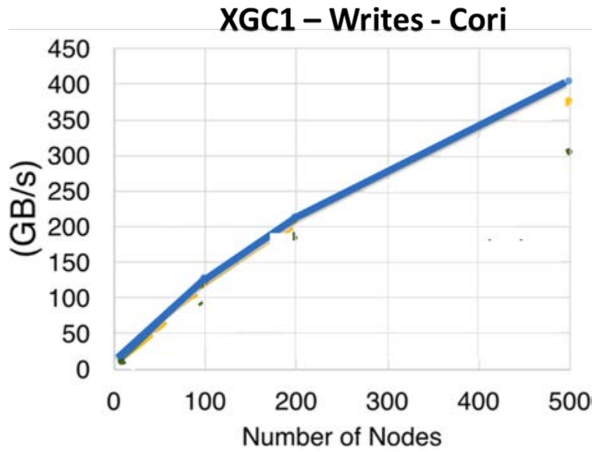

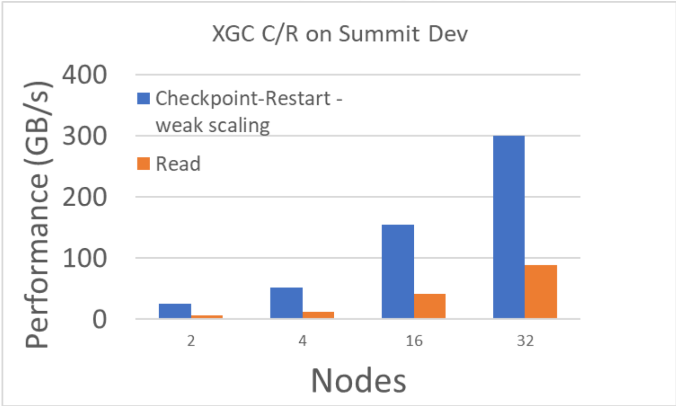

Tools from RAPIDs are used for real-time on-memory data analysis/filtering, reduction and visualization (see poster)

- XGC physics data is becoming too big for the file system.
- Realtime, on-memory **machine-learning tools** are to be used, in collaboration with **RAPIDS** and **FASTMath**
  - Fusion SciDAC ML Workshop held at PPPL, June 6-7, 2018

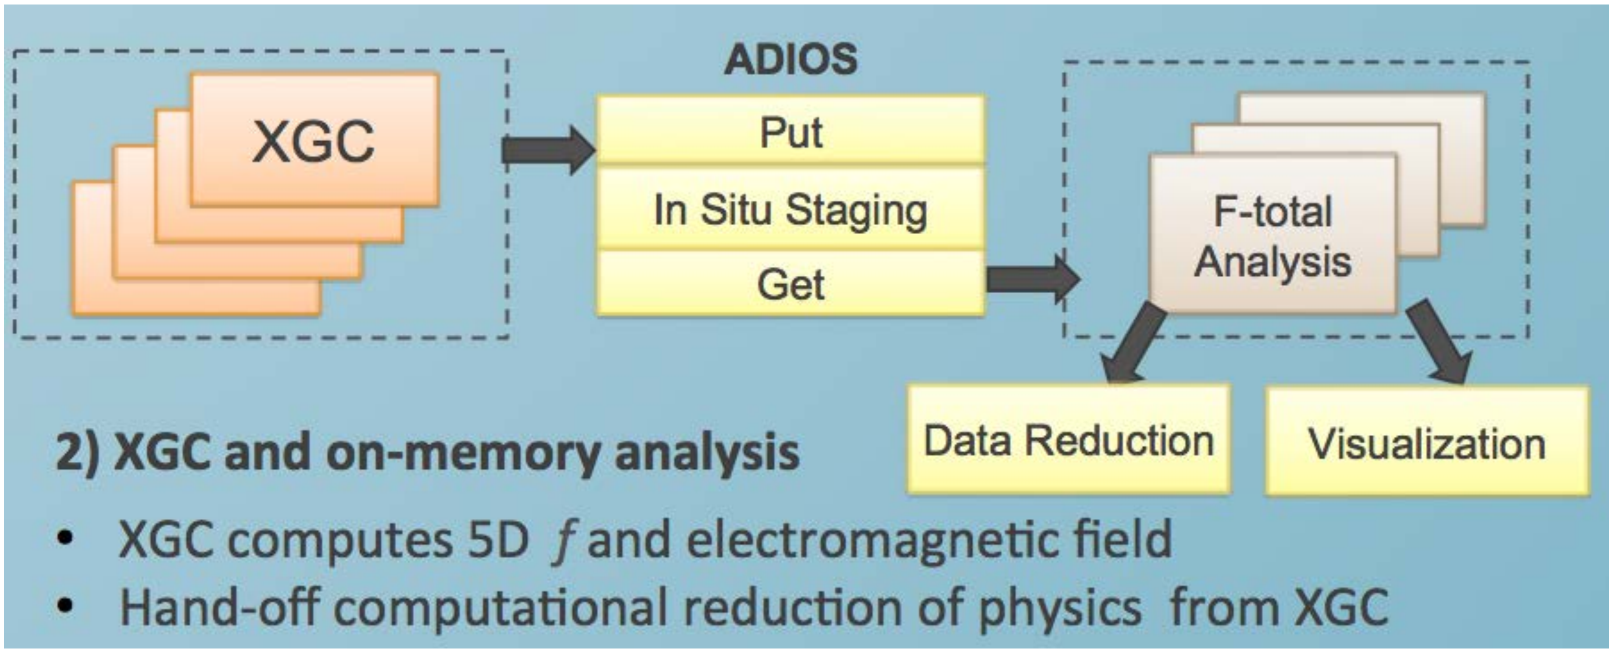

Applied Math is another central theme in HBPS (See poster)

- Solvers, various PIC algorithms, UQ, meshing, and their interactions
- Present focus is on a few game-changing algorithms
  - Fully implicit, kinetic EM algorithm has been successfully implemented (Chacon): The notorious “cancellation issue” not seen at the longest wavelengths.
  - ML to optimize pre-conditioner: in collaboration with FASTMath and RAPIDS
  - Improvement of DG algorithm for Gkeyll (Hauck)
  - Particle compression and resampling, ML for PDF reconstruction (Carey, Chacon)
  - Parallel Unstructured Mesh PIC (PUMIpic): particle migration and load balancing with minimal data movement (Shephard’s talk)
  - Multifidelity Monte Carlo UQ for extreme-scale PIC codes (Moser)

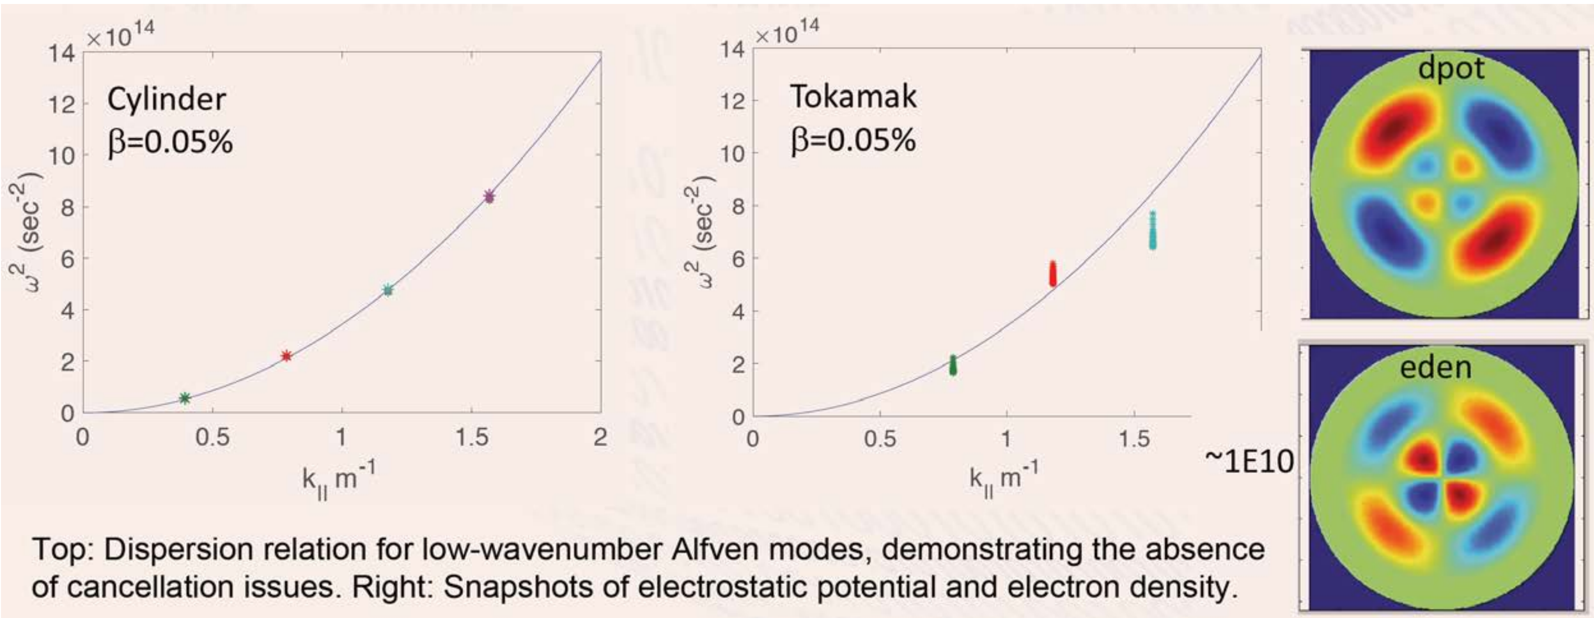

Top: Dispersion relation for low-wavenumber Alfven modes, demonstrating the absence of cancellation issues. Right: Snapshots of electrostatic potential and electron density.

Outline$^{+,*}$

- Leadership in HBPS and the project OV
- The XGC gyrokinetic code
- Example scientific discoveries
  - L-H transition
  - Divertor Heat-Flux Width
- Enabling Techonolgies
  - Performance optimization
  - Data management
  - Applied mathematics
- **Integration with other Fusion SciDAC codes for WDM**
- Summary

Integration with other Fusion SciDAC codes for WDM

- HBPS will produce a boundary plasma module for **WDM**
- HBPS plasma module need to interact with the **material module from PSI-2**
  - HBPS is using hPIC 6D Debye sheath code to provide the ion angle-energy distribution to PSI-2 material module [D. Curreli]
  - HBPS module will in-return accept the recycled and sputtered neutral particles
  - RAPIDS and FASTMath technologies will be used.
- HBPS module needs to couple with the **energetic particle and RF modules**
  - XGC is a total-f code that can handle energetic and non-Maxwellian particles
  - We are generalizing the Fokker-Planck solvers to include energetic particles [Adams, Chacon]

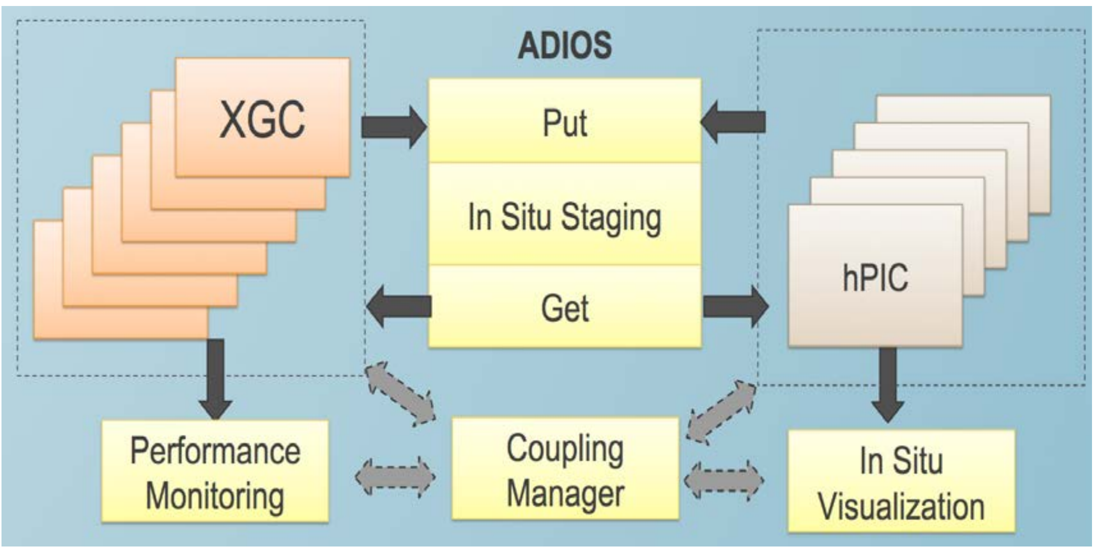

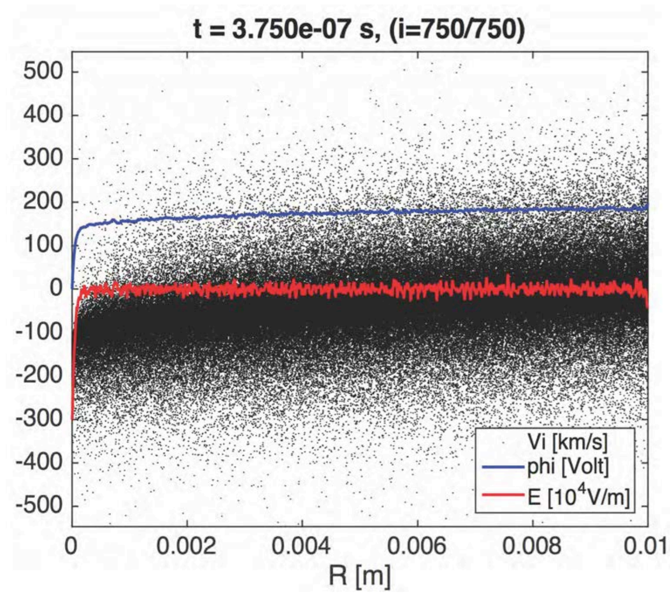

Summary

- HBPS is making scientific discoveries that would not have been possible without the SciDAC framework and US Leadership Class Computers
  - Invited Talks at major scientific conferences, including APS-DPP2017, Sherwood2017, IAEA-FEC 2016, and IAEA-FEC 2018
  - Cover story on 2018 Physics of Plasmas
  - Editor's pick by Physics of Plasmas
  - 1 Physical Review Letter 2017
- XGC is in good standing for the WDM integration
- XGC is in all three Exascale or pre-exascale programs
  - Argon-ESP, NESAP, and CAAR
  - Applied for Summit ESP
  - Scales well in all the US leadership class computers, including Summit to the maximal available # nodes (almost half of the full Summit #nodes)
- XGC is in the INCITE program (on Titan and Theta in 2018)
- The existing strong collaboration with RAPIDS and FASTMath is exptected to grow even further.
  - Collaboration in Machine Learning is being initiated
- **HBPS.pppl.gov**

SciDAC Fusion Machine-Learning Workshop 2018

Princeton Plasma Physics Laboratory, June 6-7, 2018