Solving the Multiscale Multi-Physics Puzzles in ITER Edge

CS Chang, ***Princeton Plasma Physics Laboratory***

PI: INCITE PHY122, "Exascale Gyrokinetic Study of ITER Challenge on Power-Exhaust and ELM-Free Edge"

– **on behalf of the National/International XGC team**

Funded by DOE FES & ASCR

Outline

XGC is an edge-optimized total-f EM gyrokinetic code

Designed to include **as much realistic physics and geometry as possible***

- **Total-f whole-volume EM,** not a usual perturbative code that assumes scale-separation or local physics
  - From the magnetic axis to material wall
  - Include all-scale kinetic physics
- **Wall-recycling of MC neutral particles** with atomic interactions with plasma
- **Numerical Debye sheath + pre sheath at material wall** as subgrid quantity
- **Heating and cooling** models
- **Unstructured triangular mesh** to handle complicated geometry
- PIC scheme to allow for the **non-Maxwellian particles near wall** and to handle **Lagrangian sub-grid physics**
- **Nonlinear, conservative Fokker-Planck collision**
- Externally driven magnetic perturbation

*Incomplete capabilities

- Ambipolar wall loss is assumed
- Limited detached divertor capability: radiation transport
- Large amplitude, fast compressional MHD is not included

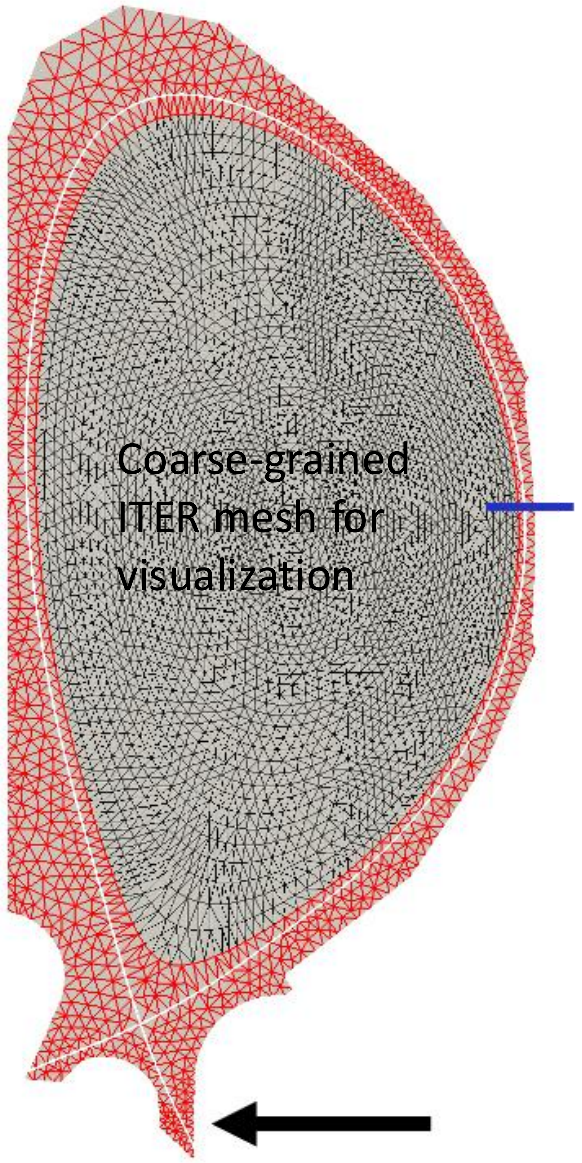

The magnetic-fusion plasma physics code XGC and OLCF

Development of the modern kinetic code XGC is deeply indebted to OLCF/NCCS

- XGC was a FES/ASCR SciDAC code since 2005
  - After a significant discovery on NERSC SEABORG: see next page
- The first OLCF INCITE award to XGC was in 2007: beginning of large-scale simulation
- In 2009, XGC became an ASCR Joule Milestone code on Jaguar (led by D. Kothe)
- Since 2007, XGC has been an Early Science Code for all the subsequent new HPCs at OLCF
- In 2012, XGC became one of the first massive GPU codes on Titan
- In 2016, XGC became an ECP application and co-design code: WDMApp and CoPA
  - Became one of the most successful exascale codes on Frontier (and Sunspot)
- XGC has made several discoveries that could only be made on capability computers
- XGC is now used by four FES SciDAC-5 projects

SciDAC Scientific Discovery through Advanced Computing

ECP EXASCALE COMPUTING PROJECT

kokkos

H-mode, discovered in 1980 in ASDEX-U experiment: puzzling

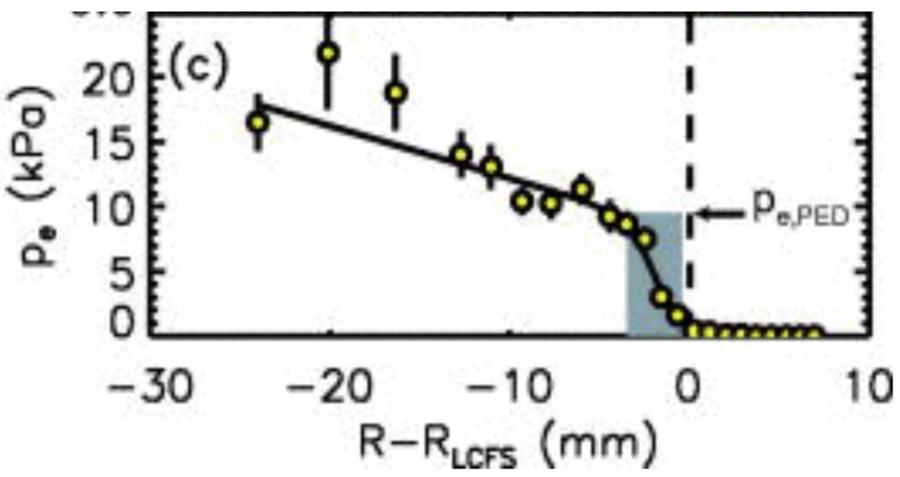

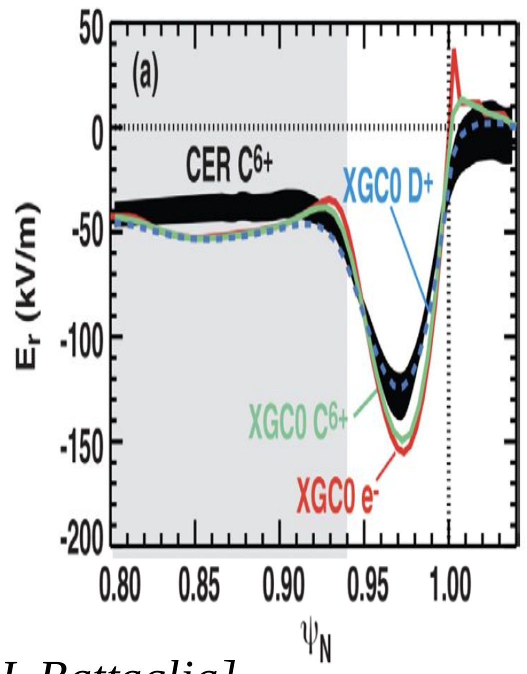

[D.J. Battaglia]

- Upon strong heating, tokamak edge spontaneously bifurcates into H-mode (high confinement mode)
  - allowed an economical size reactor to be possible
- Negative electric field well was found to provide the ion transport barrier at edge
  - How it formed was a mystery
- XGC0 discovered the ion **X-point orbit loss** physics (on SEABORG)
  - One of most hottest theoretical discoveries in magnetic fusion
- But, **XGC had to wait for a capability HPC before studying the low to high confinement mode bifurcation itself**

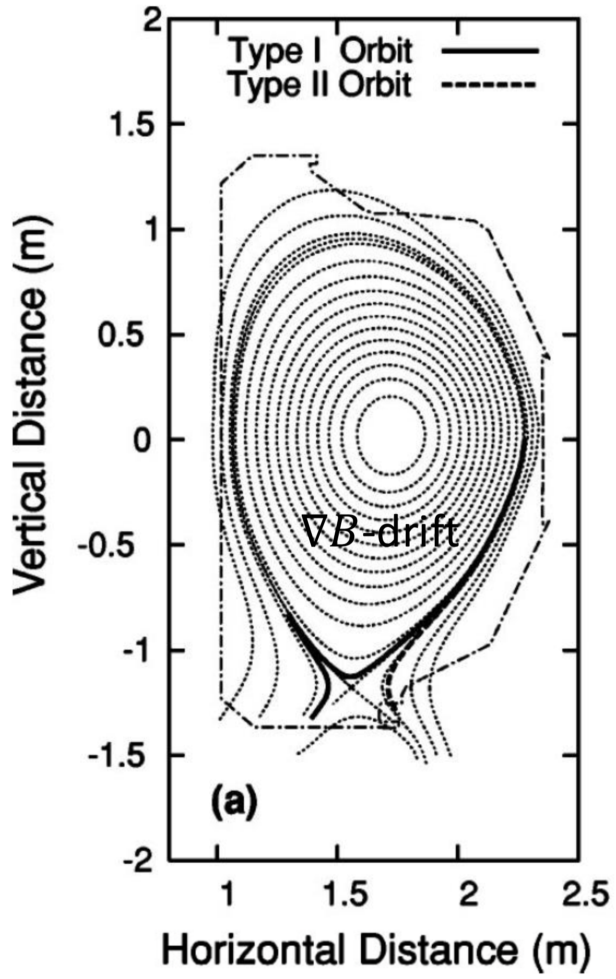

[Chang, Ku et al., Phys. Plasmas 2002, 2004]
[S. Ku, Chang et al., Phys. Plasmas 2004]

→ SciDAC ('05) and INCITE ('07)

OLCF INCITE allowed XGC to utilize then the #1 Titan (2012-2019) to achieve the Low to High operation mode bifurcation (L-H bifurcation)

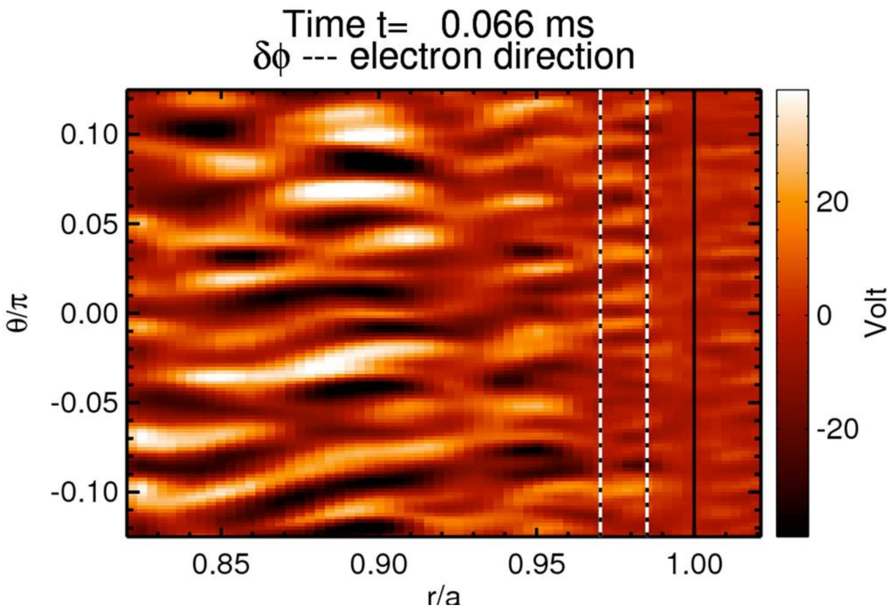

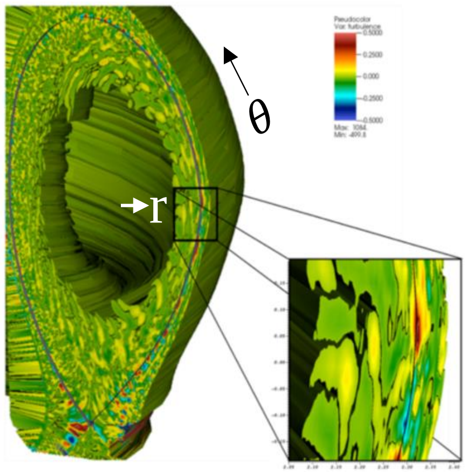

- Getting the L-H bifurcation is a necessary condition for ITER
  - First-principles understanding of the L-H bifurcation physics did not exist
- XGC found that L-H bifurcation is from multiscale self-organization
  - Reynolds stress transfers turbulence energy to small-scale sheared flow and starts the turbulence suppression
  - Global sheared flow from X-point orbit loss finishes up the bifurcation
  - **A hot achievement in fusion theory/computation: still unique**

[Chang, Ku et al., Phys. Re. Lett. 2017]
[Ku, Chang et al., Phys. Plasmas 2018]

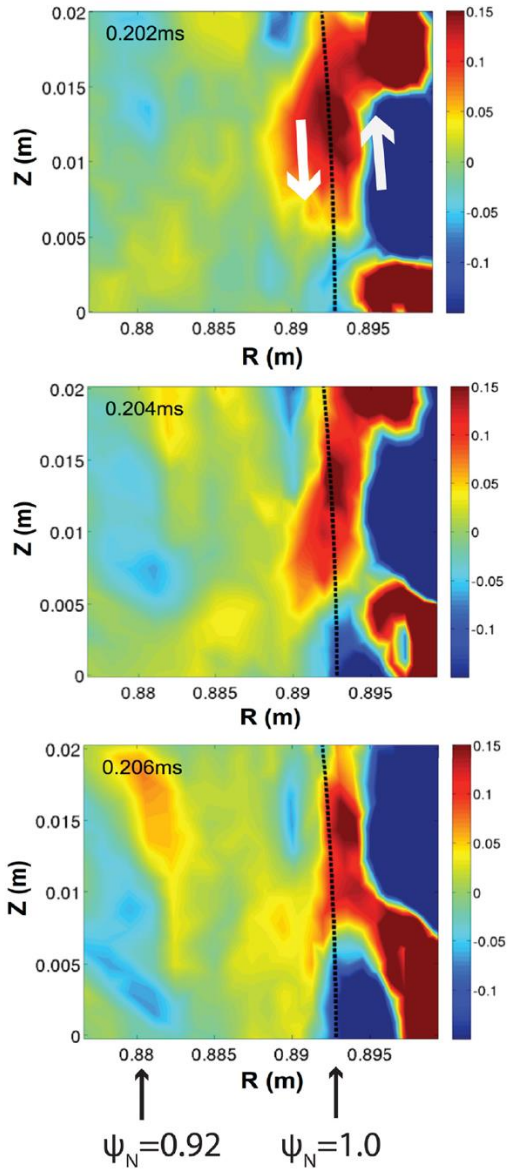

Titan then enabled XGC to predict, for the first time, the electrostatic-turbulence widening of the exhaust heat-load on the divertor plates of ITER

- Data from present-day tokamaks showed extremely narrow heat-load channel width: another puzzle
- XGC simulations agreed with all the experimental data
  - Capacity computing at NERSC was stretched for these studies on present-day tokamak-edge
  - One puzzle solved: dominated by (neo)classical physics
- Data extrapolation to ITER gave heat-load width 10^{-1} of the design value: against engineers' intuition; bigger puzzle

  → burning of the divertor material in << 1s: big issue!!!
  - **ITER edge simulation could not be done at NERSC**
- **Titan enabled first ITER edge simulation** (took ~2 weeks)
  - XGC found that ITER edge is in a different physics regime: turbulence-dominance regime
  - 12x wider heat-load width than data extrapolation
    - INCITE time allowed only one data-point
  - ITER is already accepting this prediction

[Chang, Ku et al., Nucl. Fusion 2017]

→ Invited talks at major conferences and highlights

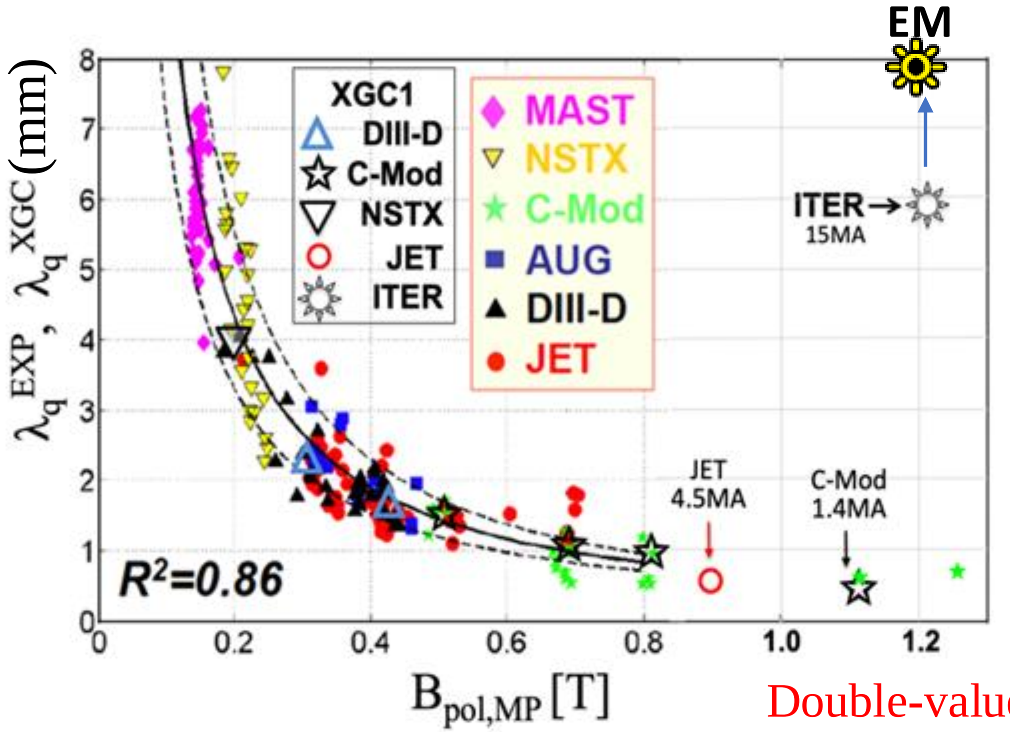

Double-valued λ_q^{XGC}

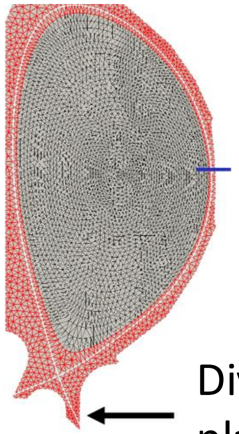

Divertor plates

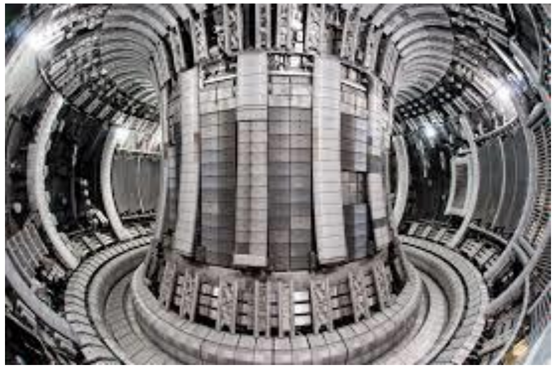

Summit enabled more comfortable simulations of ITER edge (~2 days, electrostatic)
More data points enabled an AI program to produce a simple surrogate model

→ **"Simulation anchored predictive machine learning"** [CS Chang et al., Phys. Plasmas 2021]

→ *Invited talks at major conferences and highlights*

Physics-informed supervised AI workflow using Eureqa

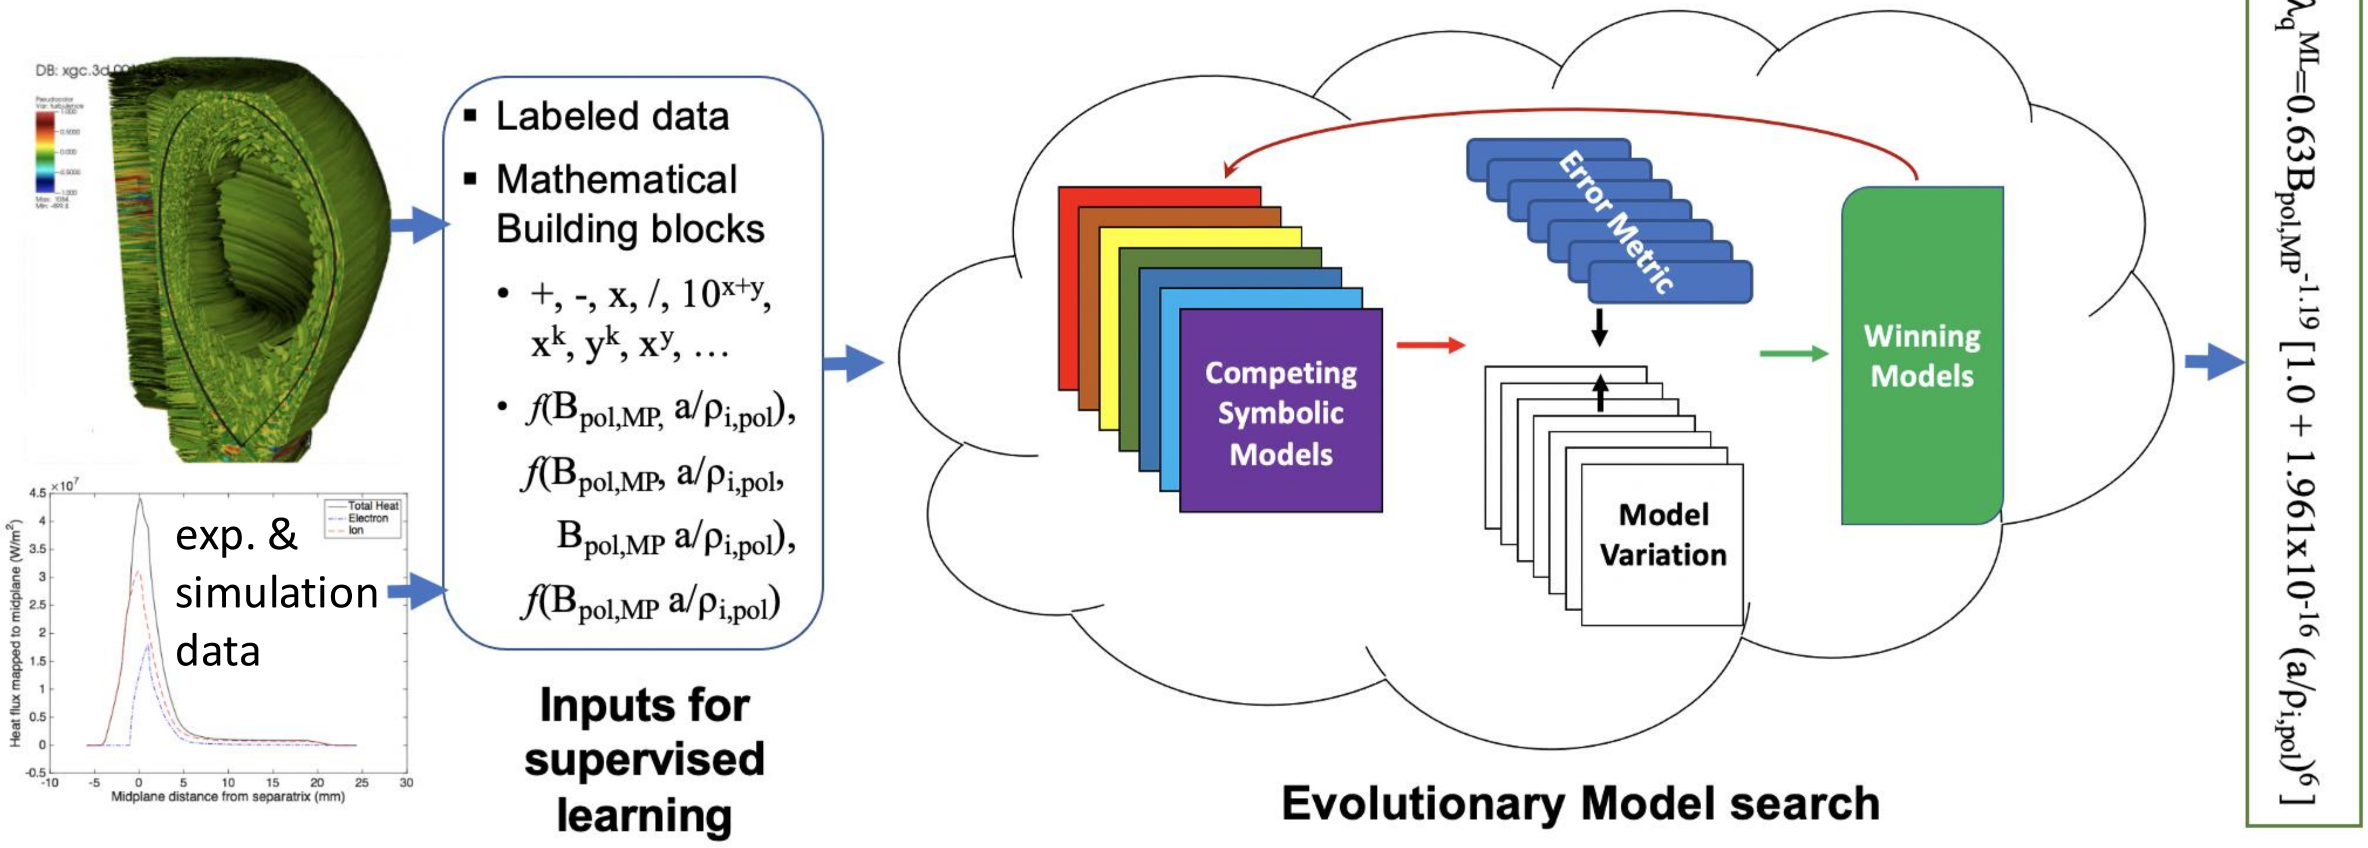

#An AI-powered modeling engine by Nutonian (https://www.nutonian.com/products/eureqa/); acquired later by DataRobot

Summit then enabled us to make another significant discovery. EM turbulence this time: Last confinement surface in ITER is leaky!!!

[Chang, Nucl. Fusion 2024, invited talk, IAEA-FEC2023]

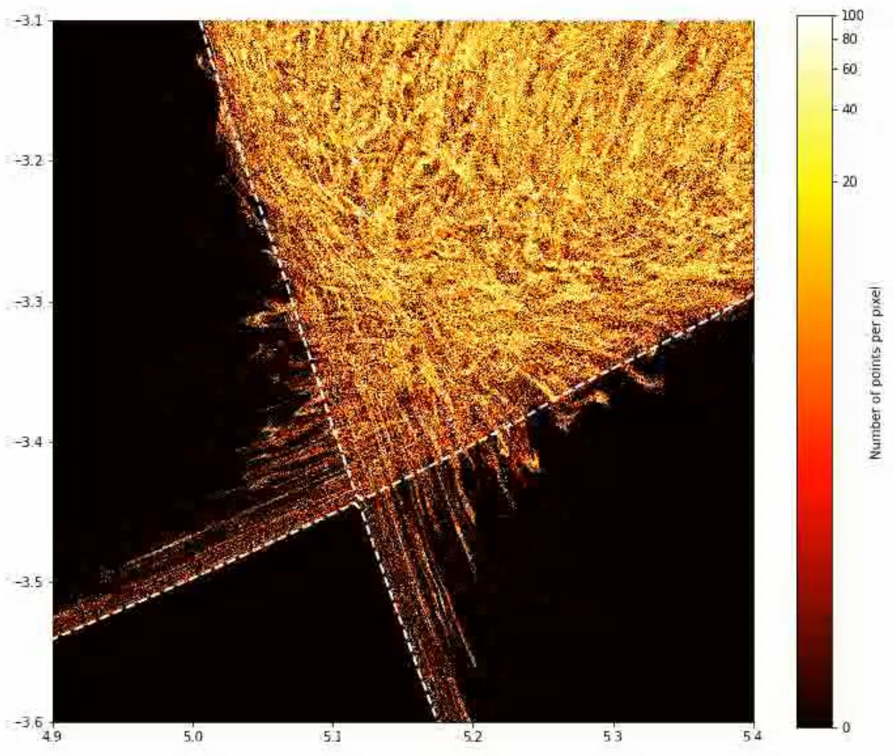

Poincare puncture plot of magnetic field lines in ITER edge, making the separatrix magnetic surface to be stochastically tangled up by the intrinsic EM turbulence, with thin and long lobes connecting core and divertor plasmas.

- Fusion researchers have assumed that the magnetic separatrix surface provides the last confinement to steady fusion-producing plasma (unless intentionally broken by external coils)
- However, experimental data often shows that the plasmas inside the separatrix surface and in the divertor area are somehow connected: puzzled
- XGC has now discovered that the separatrix in ITER (and in the high-performance present-day tokamaks) is stochastically tangled up by intrinsic EM turbulence
  - With the thin and long lobes connecting core and divertor plasmas
- We can utilize this discovery to open a new road to control the edge plasma
  - By using a rf actuator to enhance the tangles?
  - A further broadening of divertor heat-load width has already been observed

[Simulation credit: S. Ku, Vis credit: J. choi & D. Pugmire]

What is homoclinic tangle?

[First discovery by Poincare in 1881]

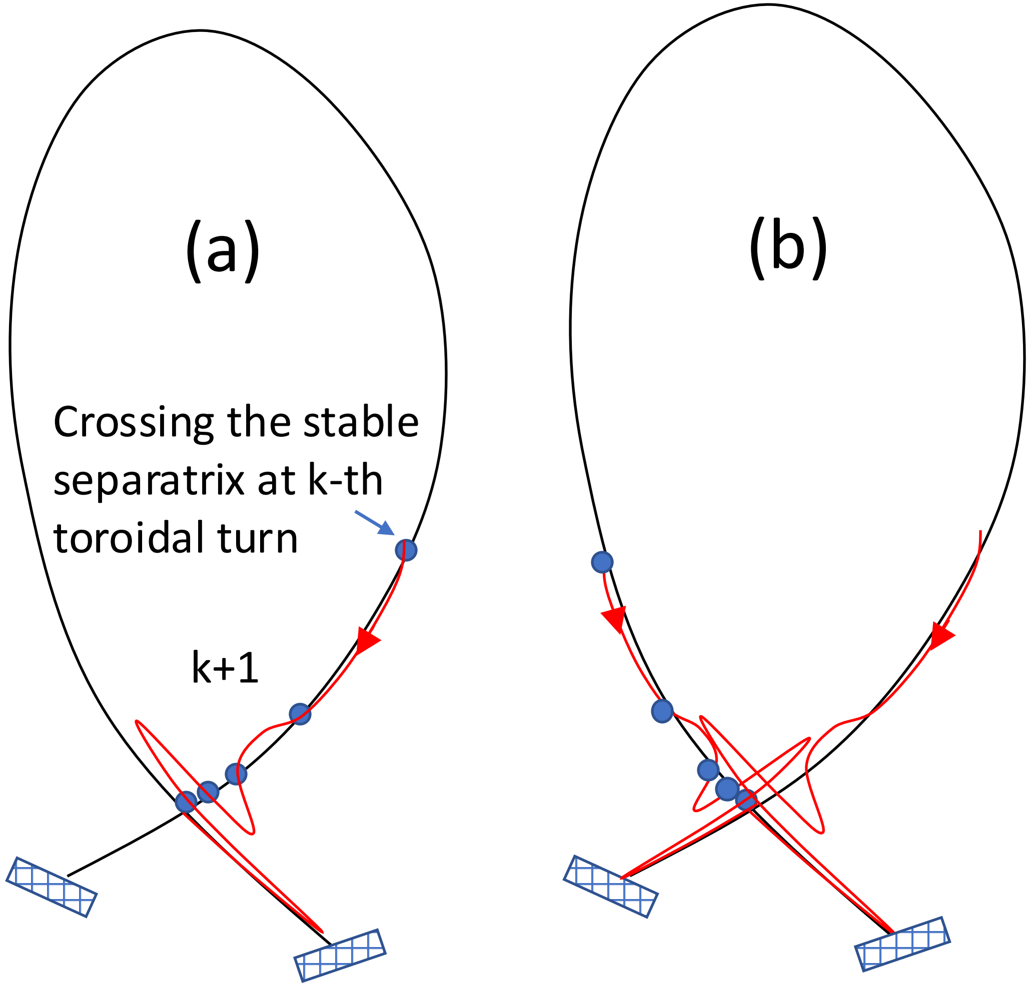

- Homoclinic tangle is formed due to magnetic-flux conservation upon perturbation of the separatrix B-field, which has a hyperbolic fixed point (= X-point)
- Tangled $\vec{B}$ avoids B-field crossing ($\nabla \cdot B = 0$): poloidal crossing at different toroidal locations
- Homoclinic tangles were known to be only driven by external field or MHD, see
  - T. E. Evans et al., *Contrib. Plasma Phys.* (2004)
  - A. Punjabi and A. Boozer, *Phys. Lett.* (2014), and others
- XGC now finds that the space-time fluctuating homoclinic tangle is an intrinsic property of diverted tokamaks, at much finer scale.

Predictions for pre-fusion ITER on Frontier: X-point radiation from Neon

- ITER will inject Ne to radiate away a significant portion of exhaust heat in the divertor region; hence to protect the divertor plates: called “high-recycling divertor operation”
- However, present-day magnetic fusion experiments using ITER-like W-wall finds that the Ne or N radiation happens around X-point and improves fusion performance: puzzle
- ITER wanted to know if Ne will accumulate around the X-point in pre-fusion ITER plasma, in which the kinetic physics could be different from present-day tokamaks
- Frontier allowed us to perform D^+, Ne^{2+} & e^- simulation in XGC for a pre-fusion ITER plasma edge
- **XGC finds that the pre-fusion ITER will indeed have Ne ions to be accumulated near the X-point**
- **Physics mechanism: Grad-B drift**
- Improvement of core performance: longer simulation is needed

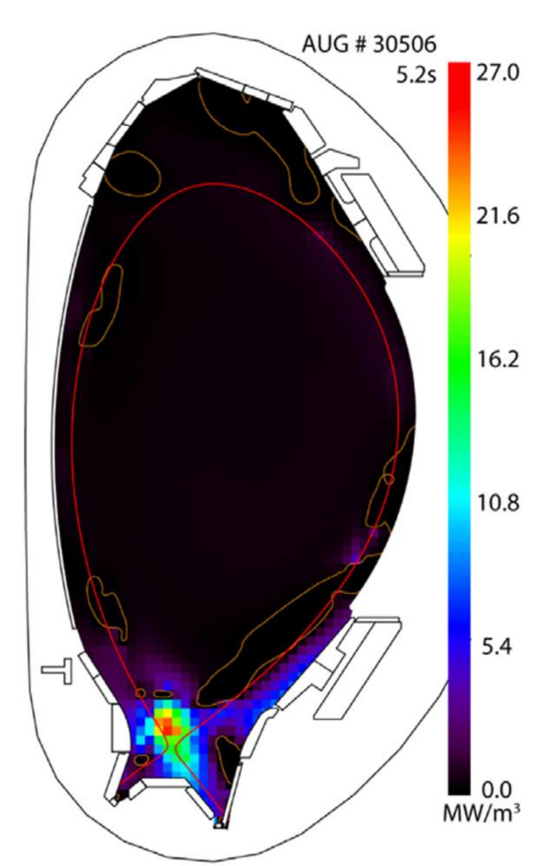

(a) AUG

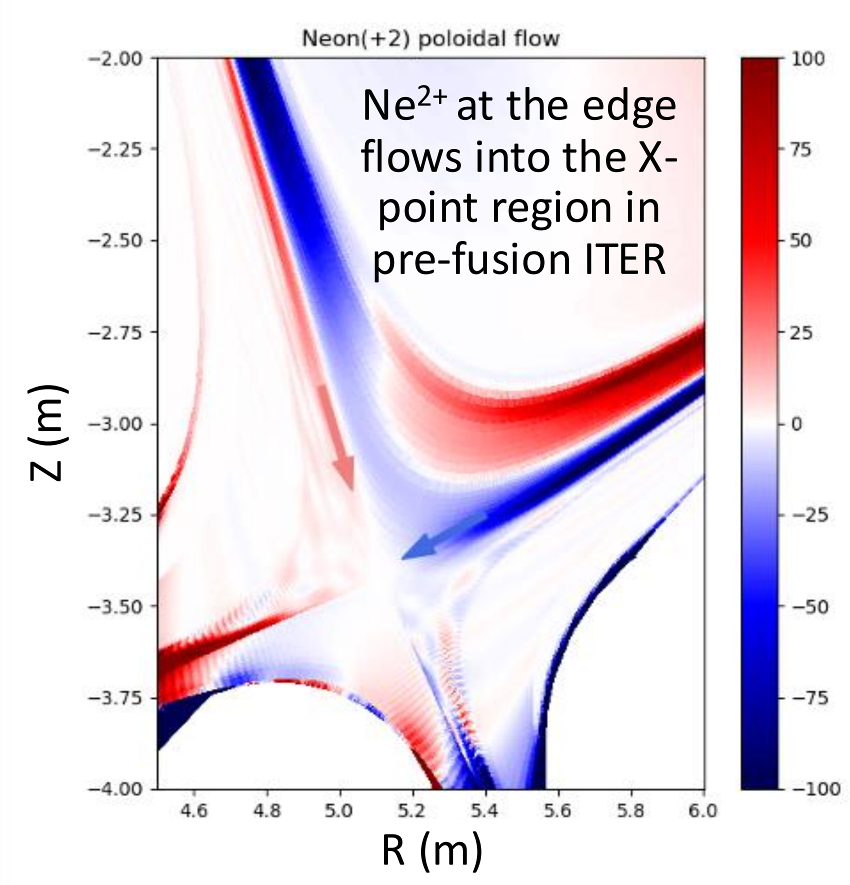

(b) Pre-fusion ITER

[Credit: J. Wilkie]

Prediction for divertor heat-load width in pre-fusion ITER on Frontier: Scattering of particle motions by atomic interaction could spread divertor heat-load?

- Existing data-base was on low-recycling divertor plasmas
  - Q: Will the high-recycling plasma in ITER make difference?
- Our preliminary result indicates that the atomic physics scattering in high-recycling divertor could be strong enough to widen the heat-load width further
  - Requires more simulations before being conclusive
  - If verified/validated to be true, this could make the ITER operation much easier
- Higher collisionality in low temperature divertor plasma requires smaller timestep size
- Needs half (4,096) of available Frontier nodes for two days (~0.1M node hours per simulation)

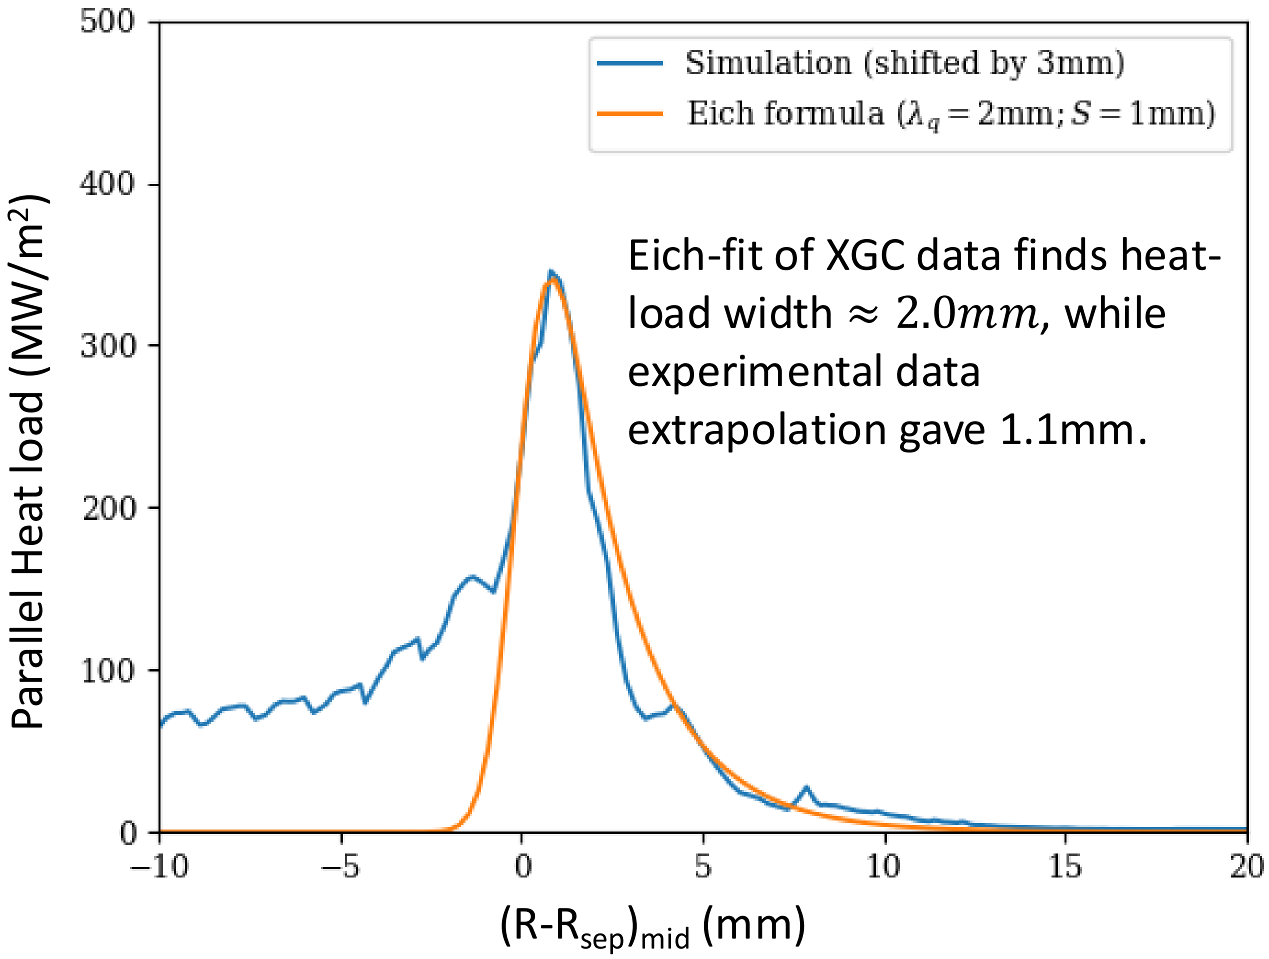

Eich-fit of XGC data finds heat-load width $\approx 2.0mm$, while experimental data extrapolation gave 1.1mm.

[Credit: G. wilkie]

How do we analyze the subgrid (in x-v space) particle transport data in extreme scale simulation? → Streaming DM and VIS on analysis nodes.

- Post-processing of large-scale particle data requires coarse-graining in particle number and time
- In-line analysis heavily consumes compute memory and also slows down the main computation
- These issues are solved by sending the “target data” asynchronously to analysis nodes.

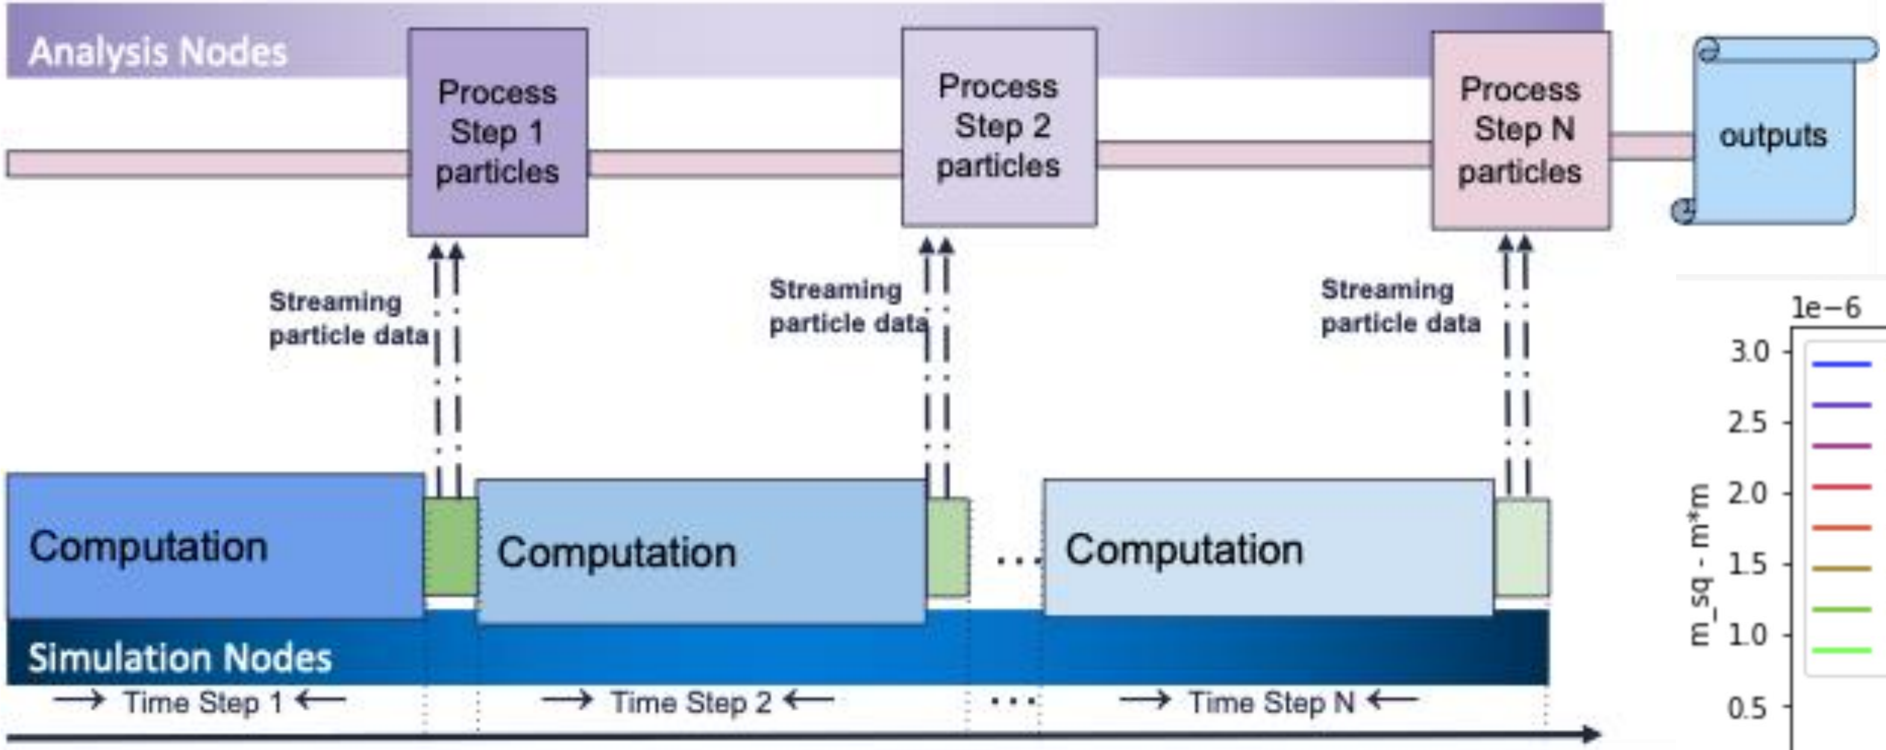

Use ADIOS2 Framework

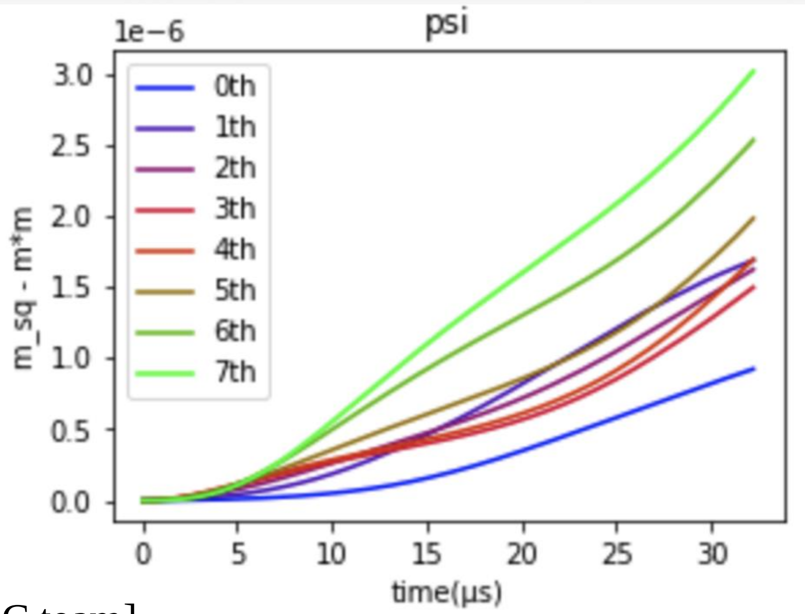

[Credit: Junmin Ku, John Wu, Paul Lin, S. Ku, C.S. Chang and the XGC team]

XGC's Engineering Approaches

- Performance portability with Kokkos and Cabana
- Major focus on encapsulation/modularity
- Templating
  - Easy experiment/swap out options
- Stand-alone kernels
  - Most major code components can be run independently
  - Use the same code base (no copies!):
    - No outdated version
    - Don't require extra maintenance
    - Improvements immediately benefit the full code
- Testing/CI
  - Unit tests, kernel regression tests, and run test on every pull request
  - Automated for biweekly testing of small number of representative physics problems

kokkos CoPA-CABANA

ECP EXASCALE COMPUTING PROJECT

[XGC performance engineering team has weekly meeting]

XGC scales well on Frontier (ECP activity)

Most computation is now on GPU, thus little “low-hanging fruit” remains.

Future improvement for XGC will mostly be

- Load balancing
- Overlapping communication and computation
- Identifying alternatives to communication-intensive algorithms
- Fine tuning the existing GPU kernels (stand-alone components great for this)

[Credit to A. Scheinberg, K. Huck, S. Abbott with Frontier COE team, S. Ku, R. Hager et al.]

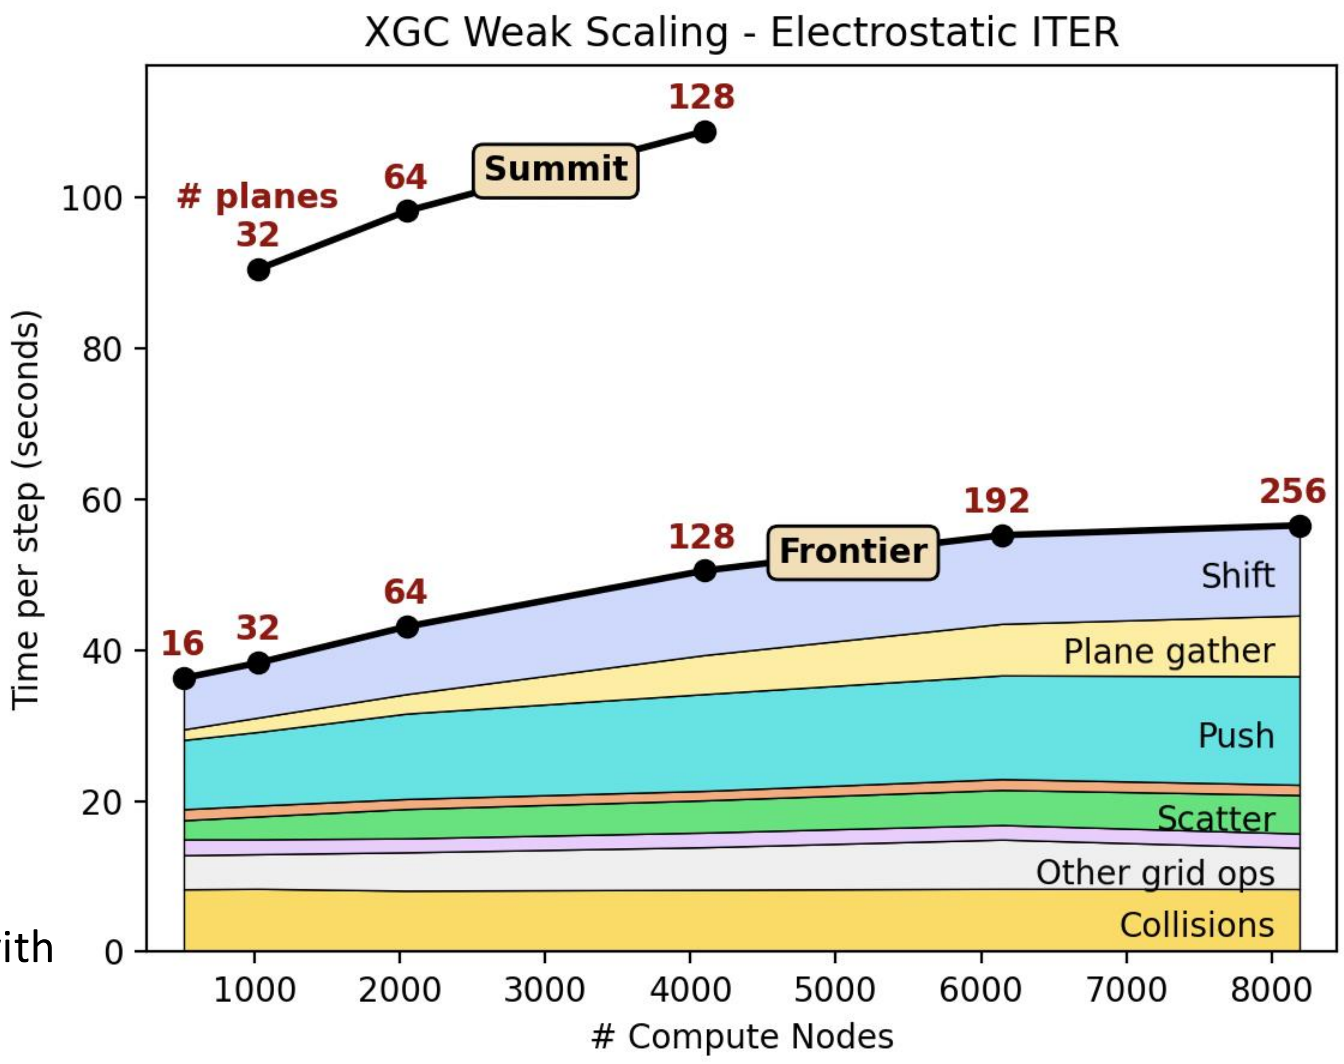

At the commissioning time of Frontier

Example for further improvement: FFT filtering

- For EM simulation of MHD modes, filtering for target mode numbers are often needed.
- Total-f EM simulation of ITER using realistic number of particles (10k ptl/vertex) and desired toroidal mode number fits on 8,192 Frontier nodes → Good weak scaling is needed.

N_N = #nodes per poloidal plane
N_P = # poloidal planes
$N_L = N_N/N_{ranks\_per\_plane}$ = #nodes assigned to each rank

Instead of assigning compute nodes to each poloidal plane, nodes are now assigned to each rank ($N_L \ll N_N$).

→ $N_P x N_N$ operation became $N_P x N_L$ operation

→ A big win

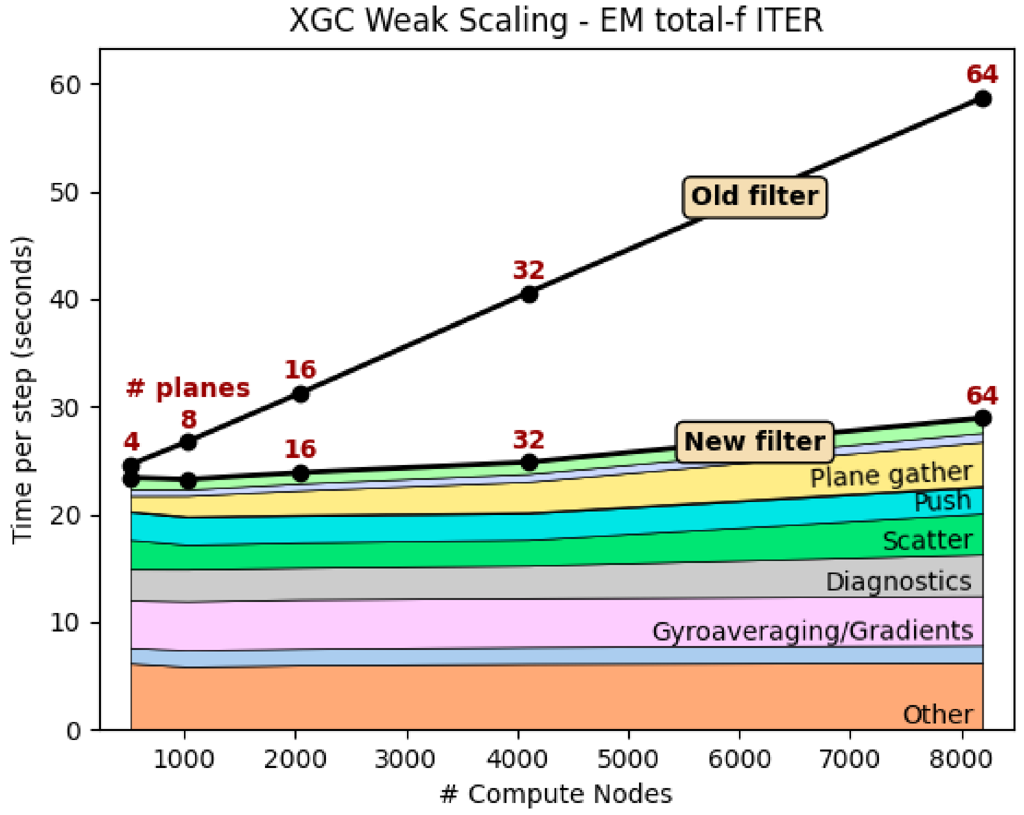

[Credit to Aaron Scheinberg]

XGC's weak-scaling performance is a bit better on Frontier than Aurora at low node counts, but worse at high node counts

- Could be due the higher ratio of computing speed to networking speed on Aurora (6 GPUs/node) than Frontier (4 GPUs)?
  - Under investigation

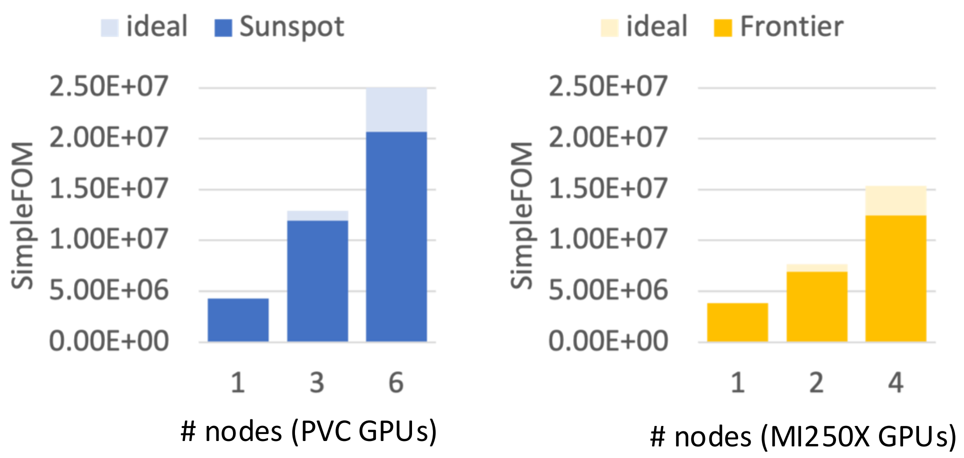

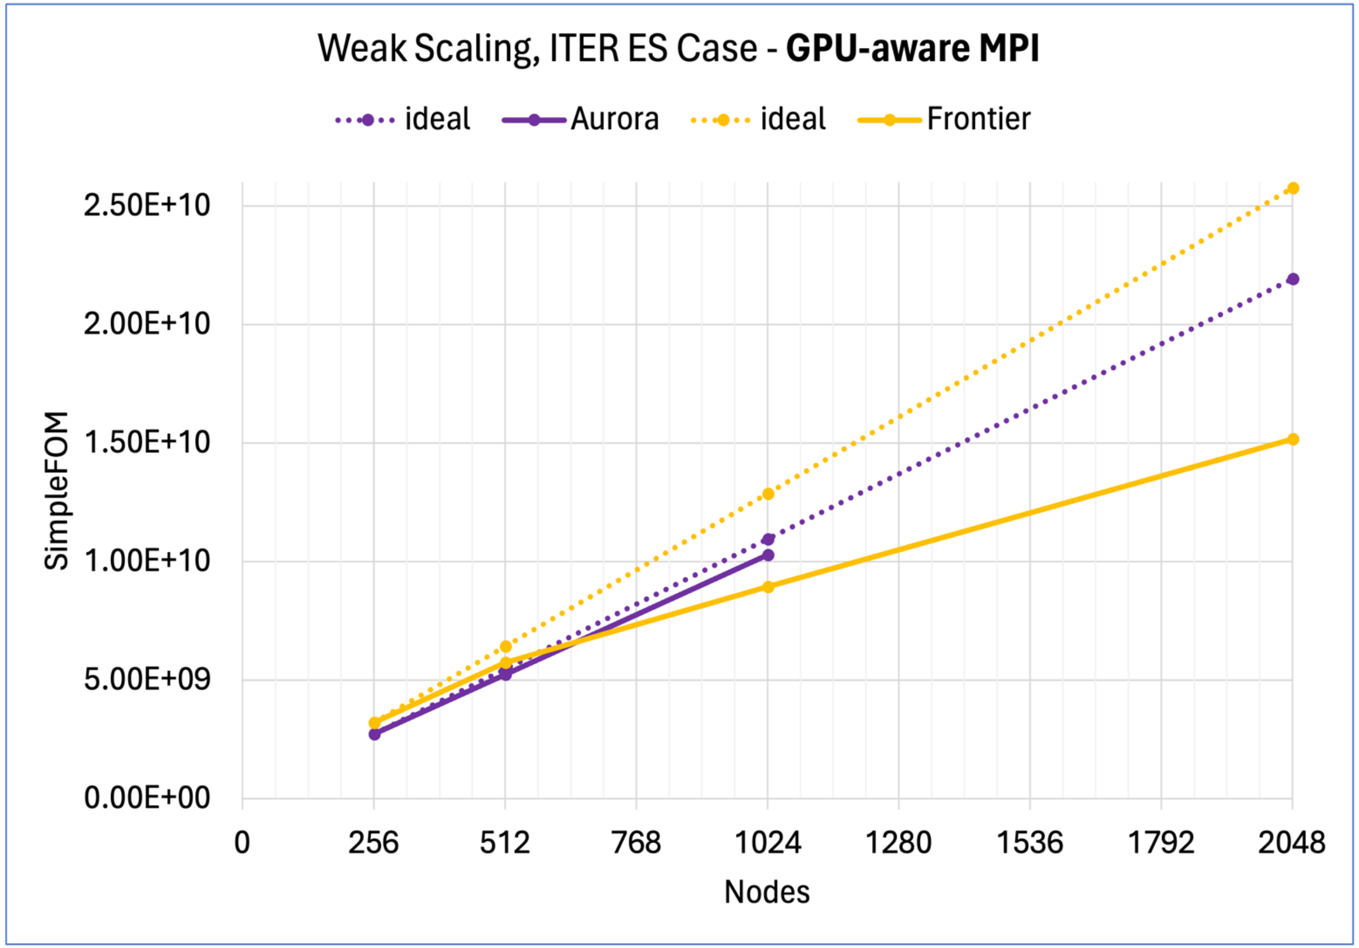

[Credit to Tim Williams]

Conclusion

- Capability computing on leadership class computers at OLCF enabled XGC to
  - solve several long-time unanswered puzzles in tokamak edge plasmas using first-principles-based kinetic code XGC
  - predict some crucial ITER edge behaviors that have not been foreseen according to data from present-day experiments or by lower-dimensional codes
- XGC will utilize exascale HPCs to include as much complete first-principles physics as possible in predicting the performance of ITER and Fusion Pilot Plants
  - Near term goal is to include W, Ne, D+T fuel, Helium ash and high recycling divertor, and to study the power exhaust and ELM-free issues, which are essential for successful ITER
  - Aims to be an **exascale digital-twin**, also producing fusion reactions (loading of α –particles and 14MeV neutrons on the plasma and wall (under development by UKAEA collaborators).
  - Can be coupled to engineering codes for a complete **"exascale fusion reactor digital-twin"**
- Performance engineering will continue to utilize the exascale computers more efficiently
  - By working with facility staff and computational scientists
- We will continue to adopt AI/ML in XGC's UQ, streaming data mining/analysis, and data federation with fusion reactors (ITER and Fusion Pilot Plants)