ALCF

SciDAC-PI Meeting, July 22-24, Bethesda, MD

Extreme Scale Computing for Fusion Edge Physics

C.-S. Chang
Princeton Plasma Physics Laboratory

for

SciDAC-3 Partnership for Edge Physics Simulation

EPSi Edge Physics Simulation

Outline

- Introduction to the Project and the Fusion Edge Gyrokinetic Code XGC1
- Example for XGC1 achievement at OLCF: ITER heat-load
- Example for XGCa achievement at ALCF: Edge bootstrap current
- Achievement examples by liaisons with four Institutes

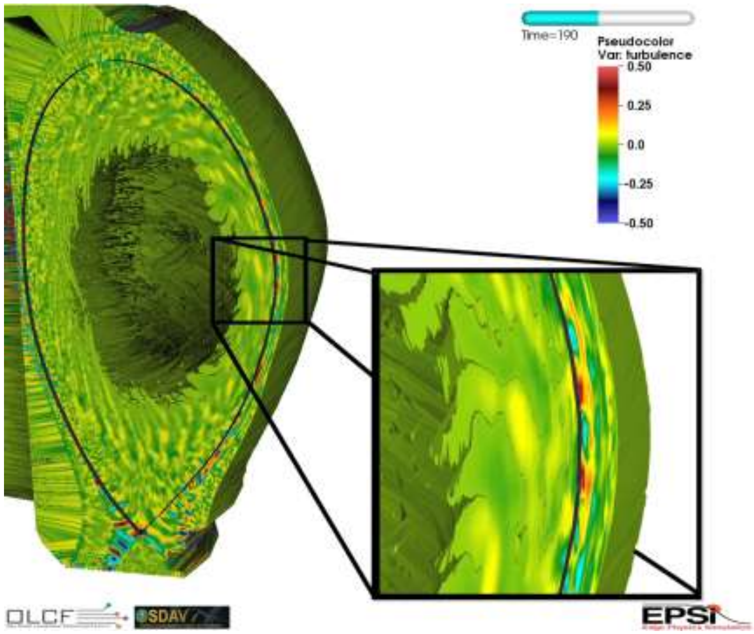

Visualization by D. Pugmire

Acknowledgements

The material presented is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Science and Office of Advanced Scientific Computing Research. This research used resources of OLCF, ALCF and NERSC, which are DOE Office of Science User Facilities. Awards of computer time was provided by INCITE (OLCF, ALCF) and ERCAP (NERSC) programs. This work was also supported by a computing-time grant from the Swiss National Supercomputing Centre (CSCS).

We also thank the HPC center liaisons.

Institutional Pis and SciDAC Institute Liaisons

OFES	OASCR
U. Colorado: S. Parker	LBNL: Mark Adams (FASTMath)
Lehigh U.: Arnold Kritz	ORNL: Scott Klasky (SDAV), Ed Dazevedo and Pat Worley (SUPER)
MIT: Martin Greenwald	RPI: Mark Shephard (FASTMath)
	Rutgers U.: M. Parashar (SDAV)
UCSD: George Tynan	U. Texas: Bob Moser (QUEST)

PPPL: C.S. Chang (OFES and OASCR)

Additional Visualization member: K. Ma, UC Davis (SDAV)

"Toroidal" Tokamak Geometry

Poloidal magnetic flux label

$\psi(r) = 1$ at $r/a=1$, 0 at $r/a=0$

ITER

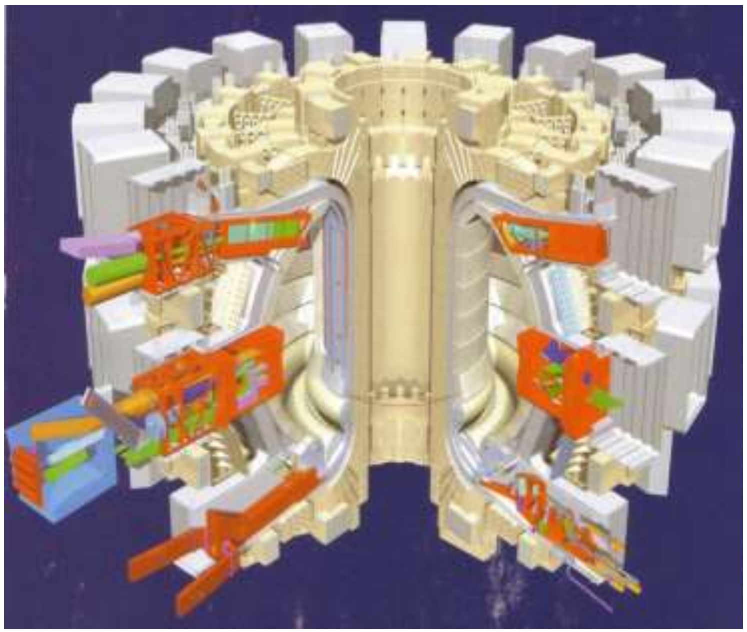

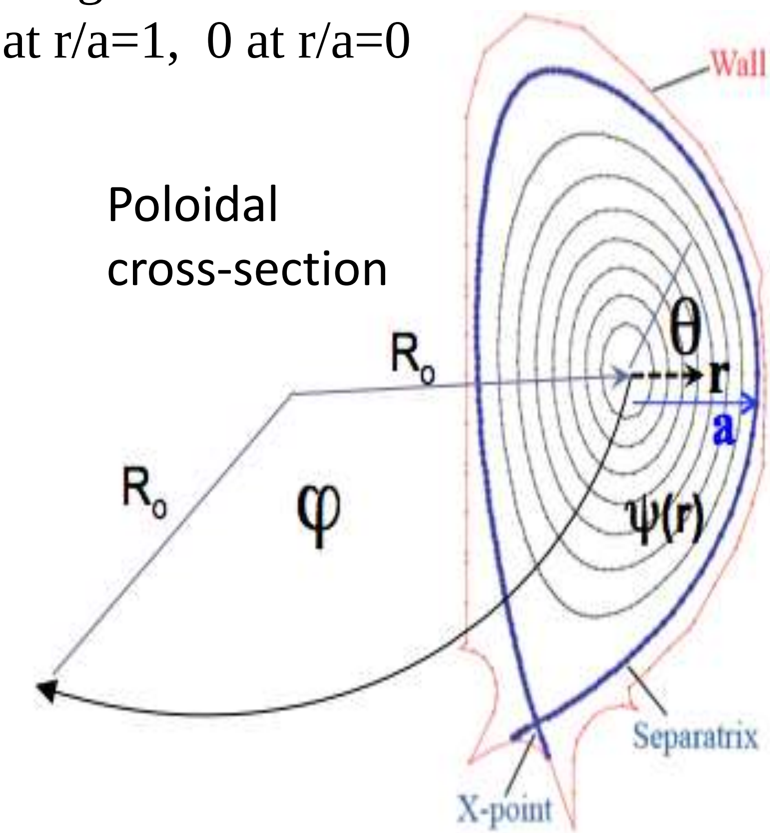

Torus, not a straight cylinder: physics becomes more complicated through the magnetic inhomogeneity and the **toroidal mode coupling**.

XGC1, with its excellent portability, could take advantage of all the LCFs tested so far

- A result from close collaboration between ASCR and OFES Scientists.

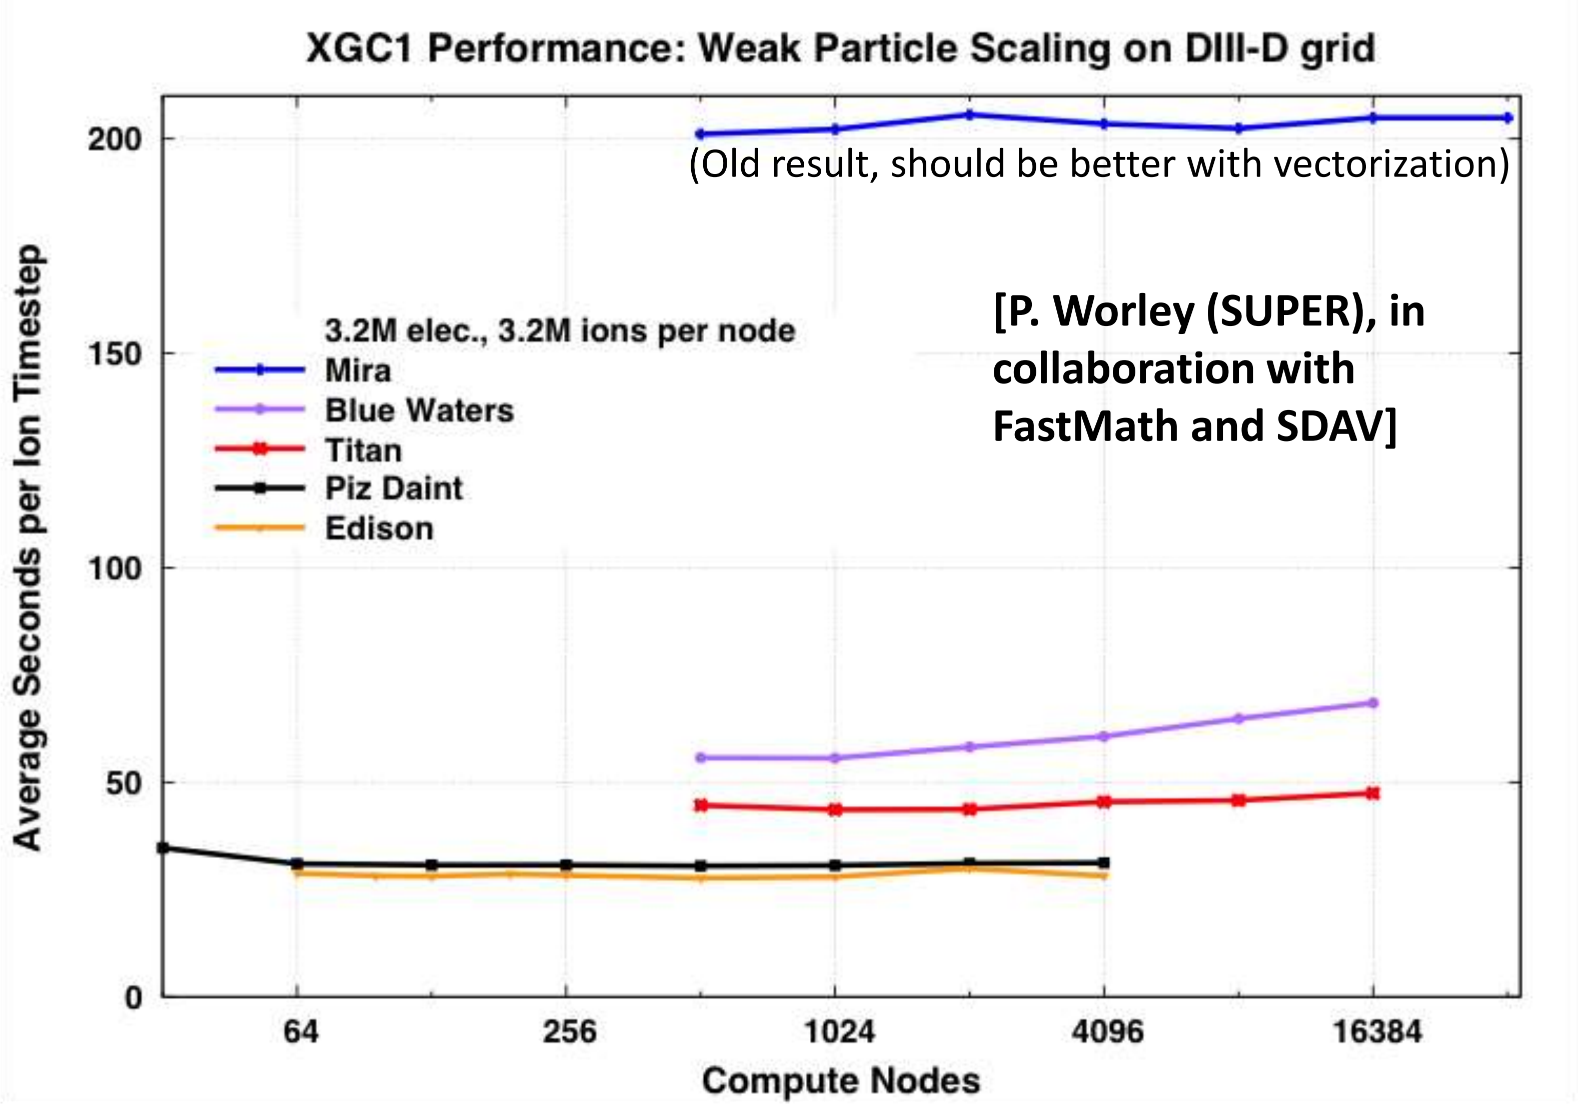

In XGC1, strong scaling is also decent unless #particles/GPU becomes too small

- Maximal #particle/node is set by the size of node and GPU memory
- If #particles/node becomes too little in the strong scaling study, the GPU's compute intensity is diluted and the communication cost gets exaggerated.
- PIC Code: more physics (bigger mesh #) requires more particles:
  more compute nodes → weak scaling dominant.

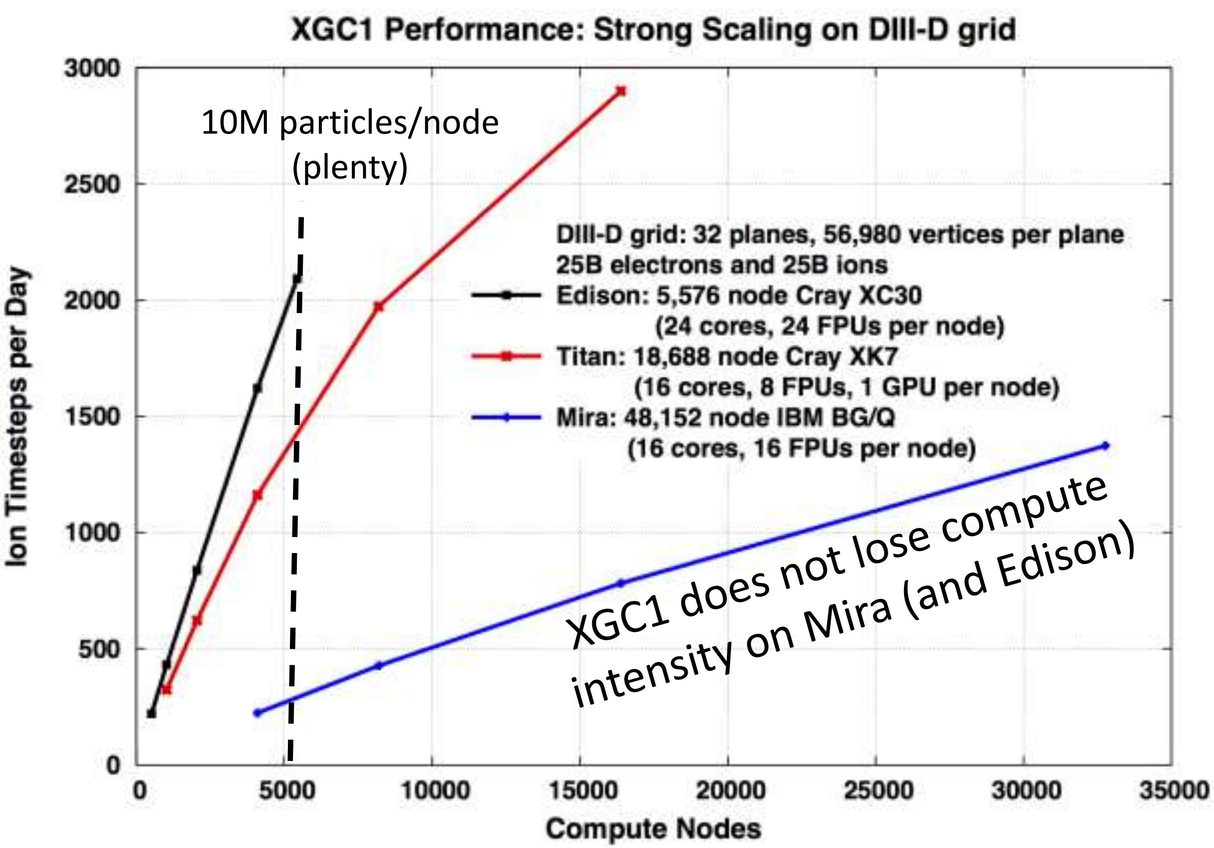

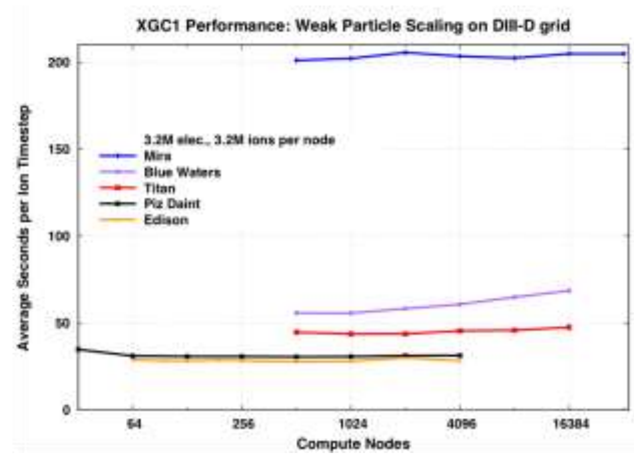

Computing Resources for XGC in 2015

- INCITE: 270M hours (Capability computing)
  - Titan: 170M hours (Extreme-scale jobs with full physics, 10-20PFs, usually 90% capability computing)
  - Mira: 100M hours (Next level extreme-scale jobs at 3.3 PFs, ~1/3 capability computing)
- NERSC: 70M hours (capacity computing on Edison, <1.5PFs)

Pre-Exascale Program

- CAAR at OLCF: postdoc support
- NESAP at NERSC: Tier 1, postdoc support

--A representative scientific discovery case--
Divertor heat-load width: a serious issue for ITER

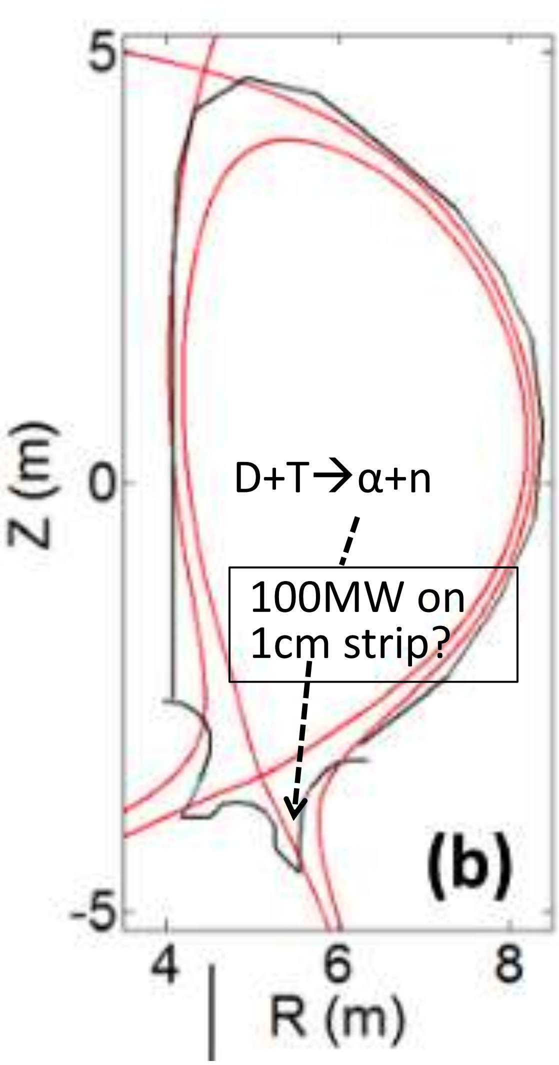

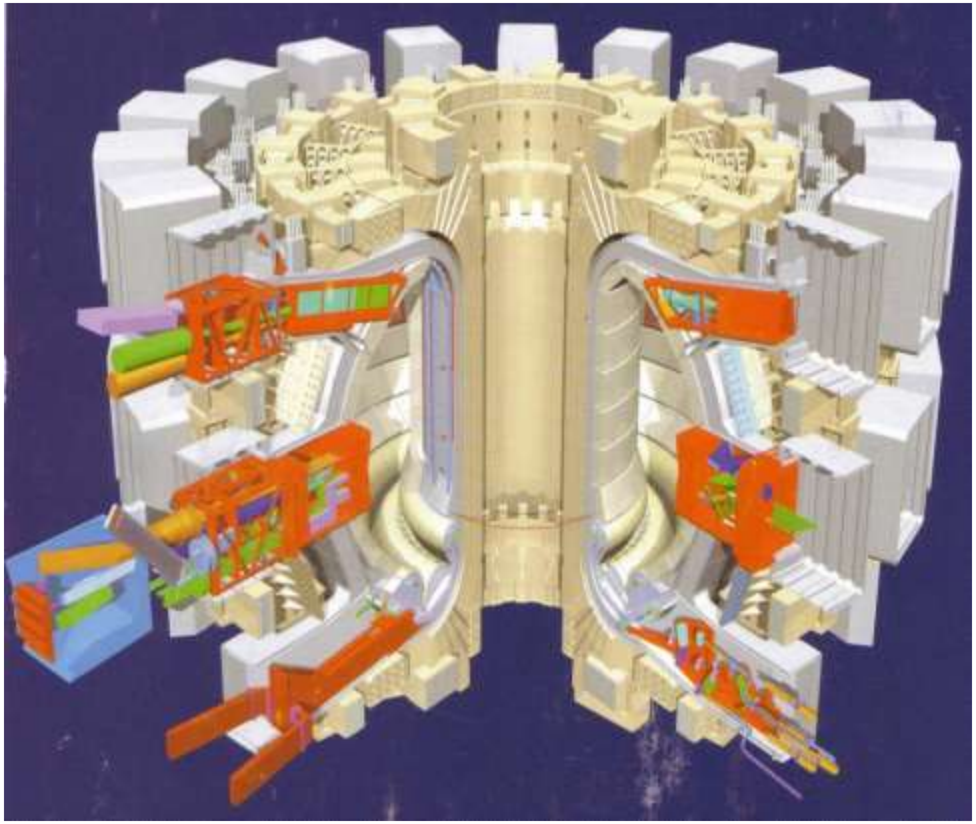

--A representative science case--

Divertor heat-load width: a serious issue for ITER

- If extrapolated from the present-day trend ($\lambda_q \propto 1/I_p^\gamma$, γ~1),
  - ✧ Divertor heat-load width in ITER, when mapped back to outer midplane would be $\lambda_q \approx 1$mm → serious issue when the technological limit ~10MW/m^2
  - ✧ 100 MW on an extrapolated strip →

    ≈ 20MW/m^2 steady, plus pulsed heat
- Non-turbulence dominant models, by XGC0 and Goldston, have shown

  $\lambda_q \propto 1/I_p^\gamma$, γ~1
- Unanswered critical questions:
  - ✧ Will the $1/I_P$ trend applicable to ITER?
    - Extrapolation is too far
  - ✧ How can we widen λ_q?
  - ✧ **Physics understanding needed.**
- Edge plasma is in non-equilibrium kinetic state: non-Maxwellian.

  → **Extreme scale computing**

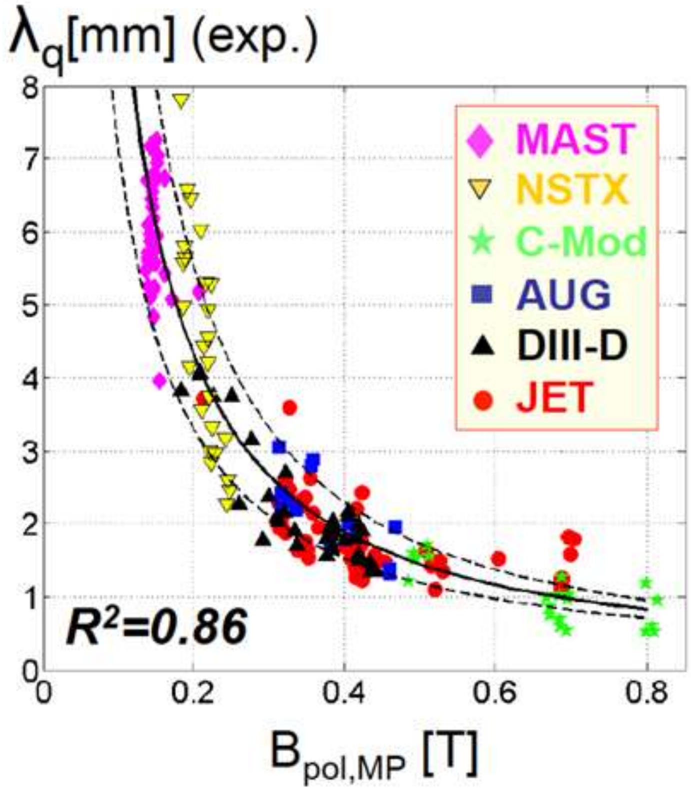

T. Eich et al., NF 2013

Ability to produce the nonlinear "blobby" edge turbulence + orbit dynamics is a pre-requisite

2013-2014 INCITE, using 90% (16,384+ nodes~20pF) maximal heterogeneous Titan

- Experiments: edge plasma is in "blobby" nonlinear form
  - Large amplitude density and electric potential blobs (~20%)
  - Only simple "models" existed
  - Kinetic, non-theraml equilibrium: Computationally difficult and expensive
- Titan and ASCR collaboration enabled gyrokinetic blob production in XGC1, for the first time (reported in SciDAC-14)
  - Kinetic understanding of blobs In realistic diverted geometry
  - OLCF Featured Highlight, 2/2014
- **Stage was set for the divertor heat-load footprint prediction**

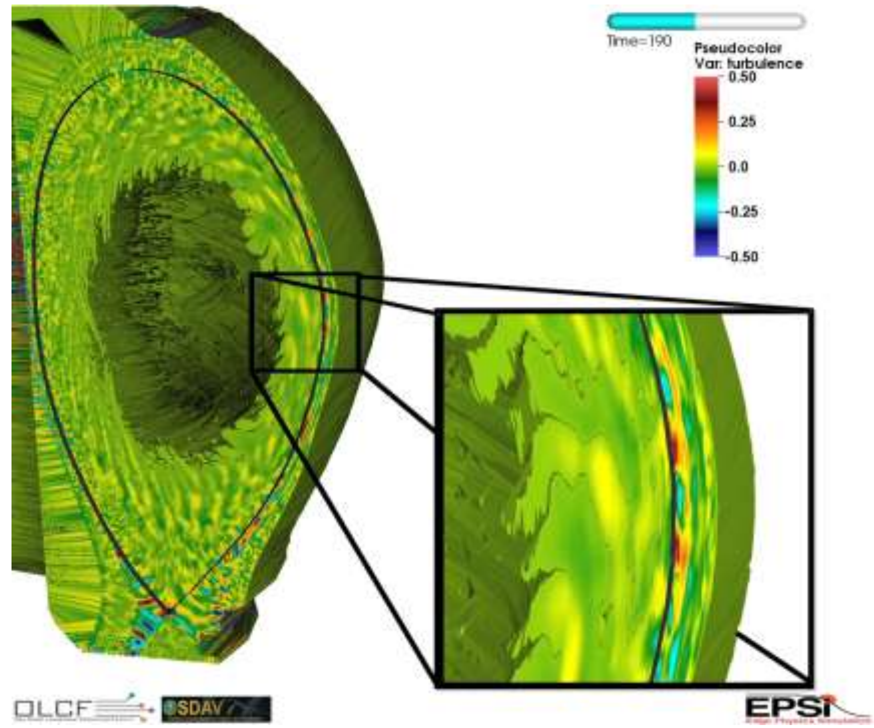

Simulation by S. Ku, Visualization by D. Pugmire

Gyrokinetic XGC1 simulation of edge blobs in DIII-D plasma

λ_q is dominated by ions in this DIII-D plasma

$K_e < K_i$ in scrape-off, and ions (electrons) gain (lose) kinetic energy in the pre-sheath

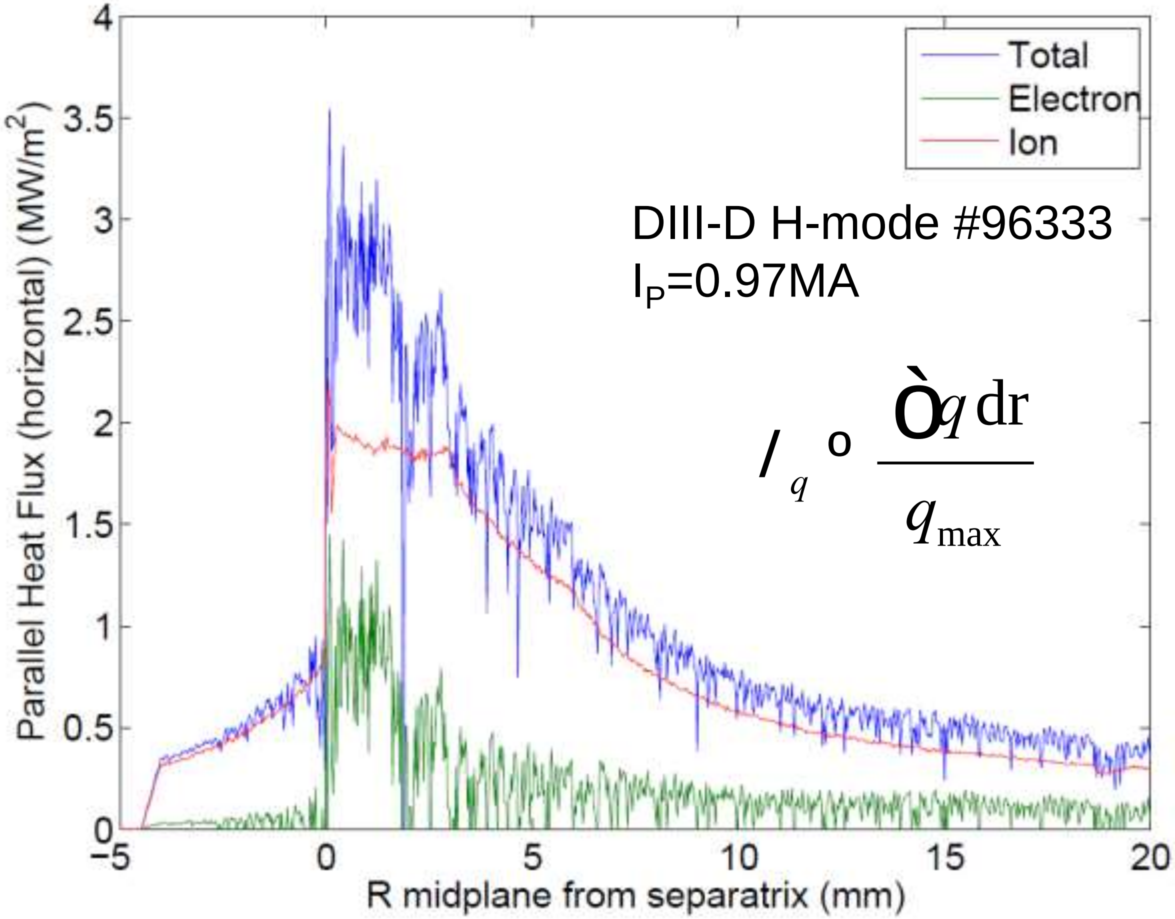

- λ_q =5.1 mm at I_P=0.97MA
  - Neutral particle effect is only ~10%
- λ_q is closer to ion orbit spreading width than the turbulent blob size (≳1cm)

Heat-load spreading by **blobs** (represented by λ_{qe} ~2mm in the electron channel) is masked by the ion orbital spreading.

Predictions for DIII-D & NSTX are in the right ballpark

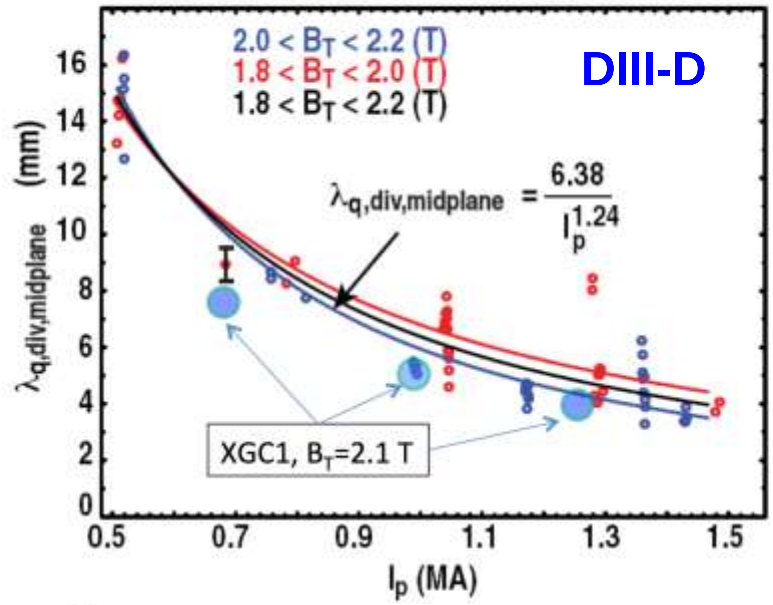

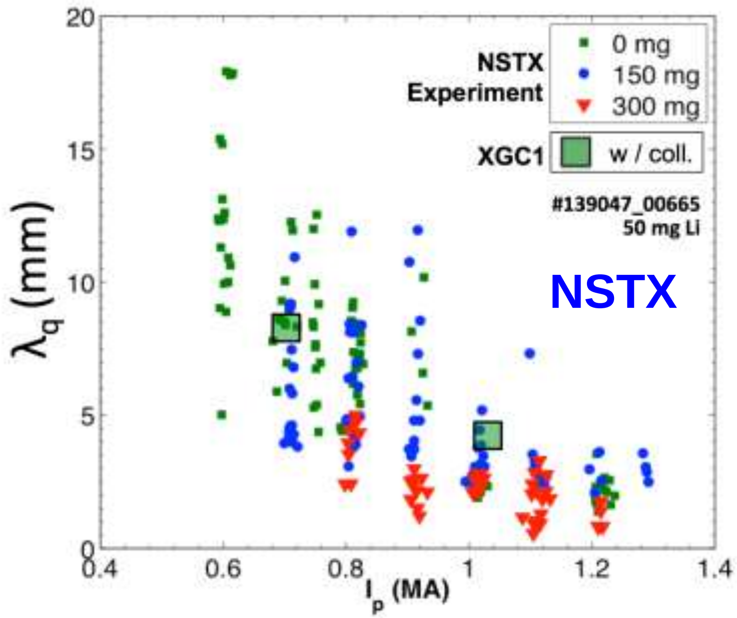

- $\lambda_{q,mid}$ from XGC1 agrees with experimental values from two very different tokamak devices
  - ✧ DIII-D for conventional aspect ratio
  - ✧ NSTX for tight aspect ratio
- Broadening of $\lambda_{q,mid}$ by $\gtrsim$1cm blobs is found to be insignificant in present-day machines.
- Will the blobs survive and saturate the $1/I_p$ scaling when the ion orbit width becomes ~ 1mm (<< blob width) in ITER?

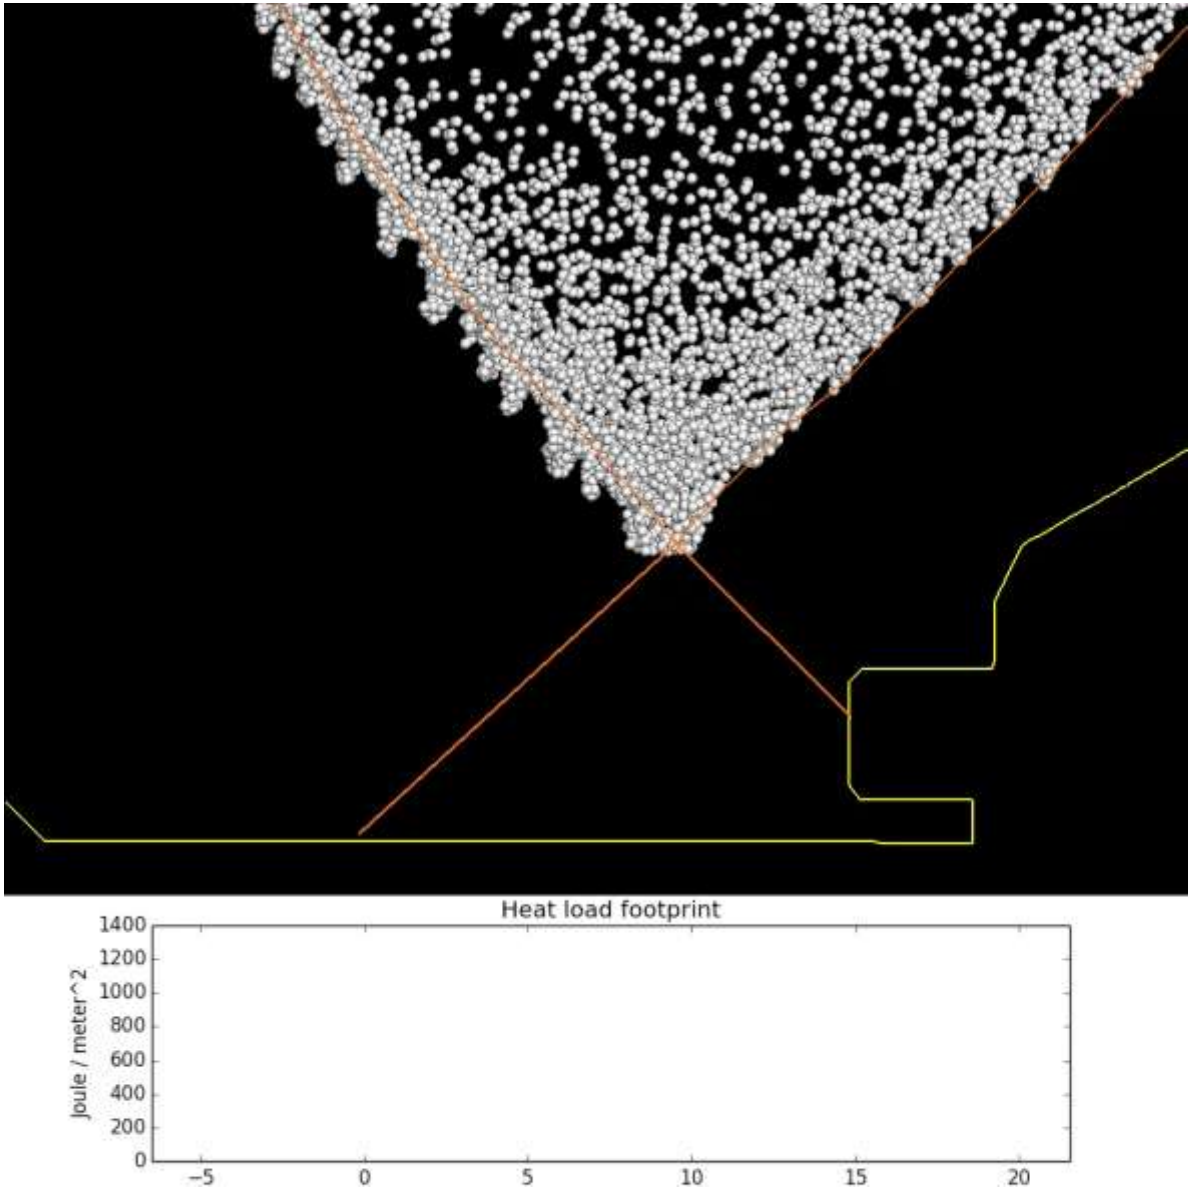

XGC1 reveals the heat-load footprint physics in unprecedented details

Vis. by Dave Pugmire

Toroidally averaged turbulent and orbital drifts of ions

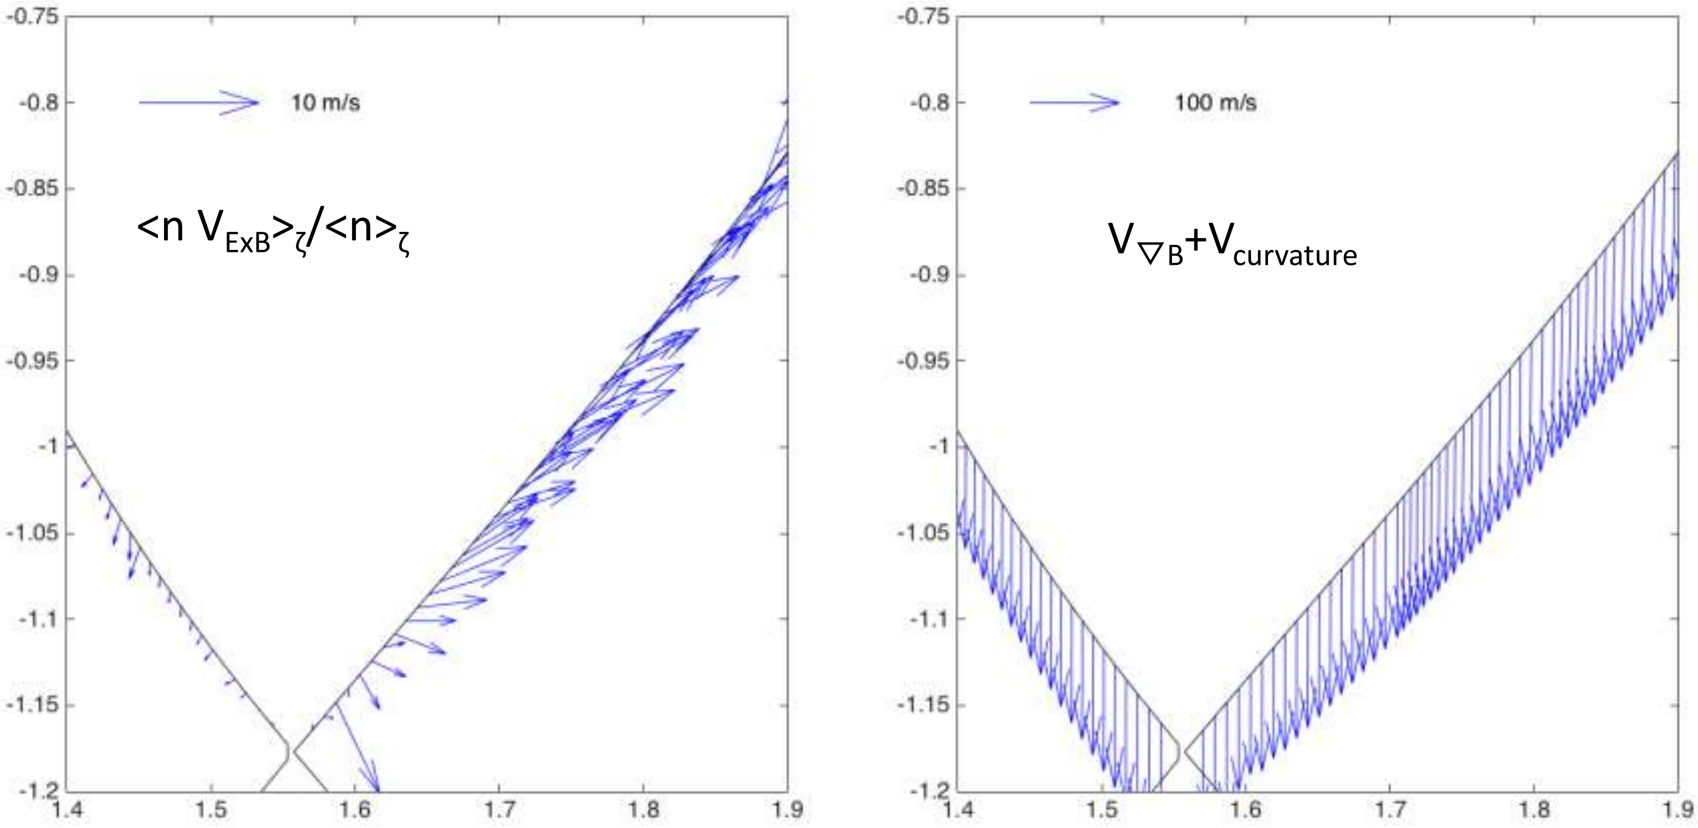

- Kinetic ion orbit width effect dominates over turbulence effect in present-day tokamaks. If this remains to be true for ITER, there is an expensive technological challenge remaining.
- **Roles may reverse in ITER due to neglibibly small ion orbit width effect: importance of including both neoclassical and turbulence physics → Gyrokinetic simulation**

Heroic runs for the ITER heat-load width prediction

Built upon XGC1's success on the heat-load width prediction on two representative present-day tokamaks, we have proceeded with the long-awaited ITER simulation.

- 90% of Titan has been utilized for ~3 days + fault-damaged ~2 days.
  - XGC1 became more fault resilient.
  - 0.53ms of physics time, already reached saturated edge turbulence
- Preliminary result shows that the nonlinear "blobby" turbulence dominates the heat-load width physics in ITER.
  - **Extrapolation from the present-day experimental data is not consistent with the XGC1 prediction**
  - **The heat-load width in ITER from XGC1 is ~1cm for both electrons and ions, instead of the ion-orbit width ~1 mm as predicted by extrapolation of the present-day data.**
  - **With 1cm heat-load width (mapped value to outside midplane), ITER will not have much problem with the divertor wall material issue.**
  - If verified, **this could be one of the best news ITER could have** since its funding agreement.

Heat-load footprint from the "standard ITER plasma," when mapped back to outboard midplane.

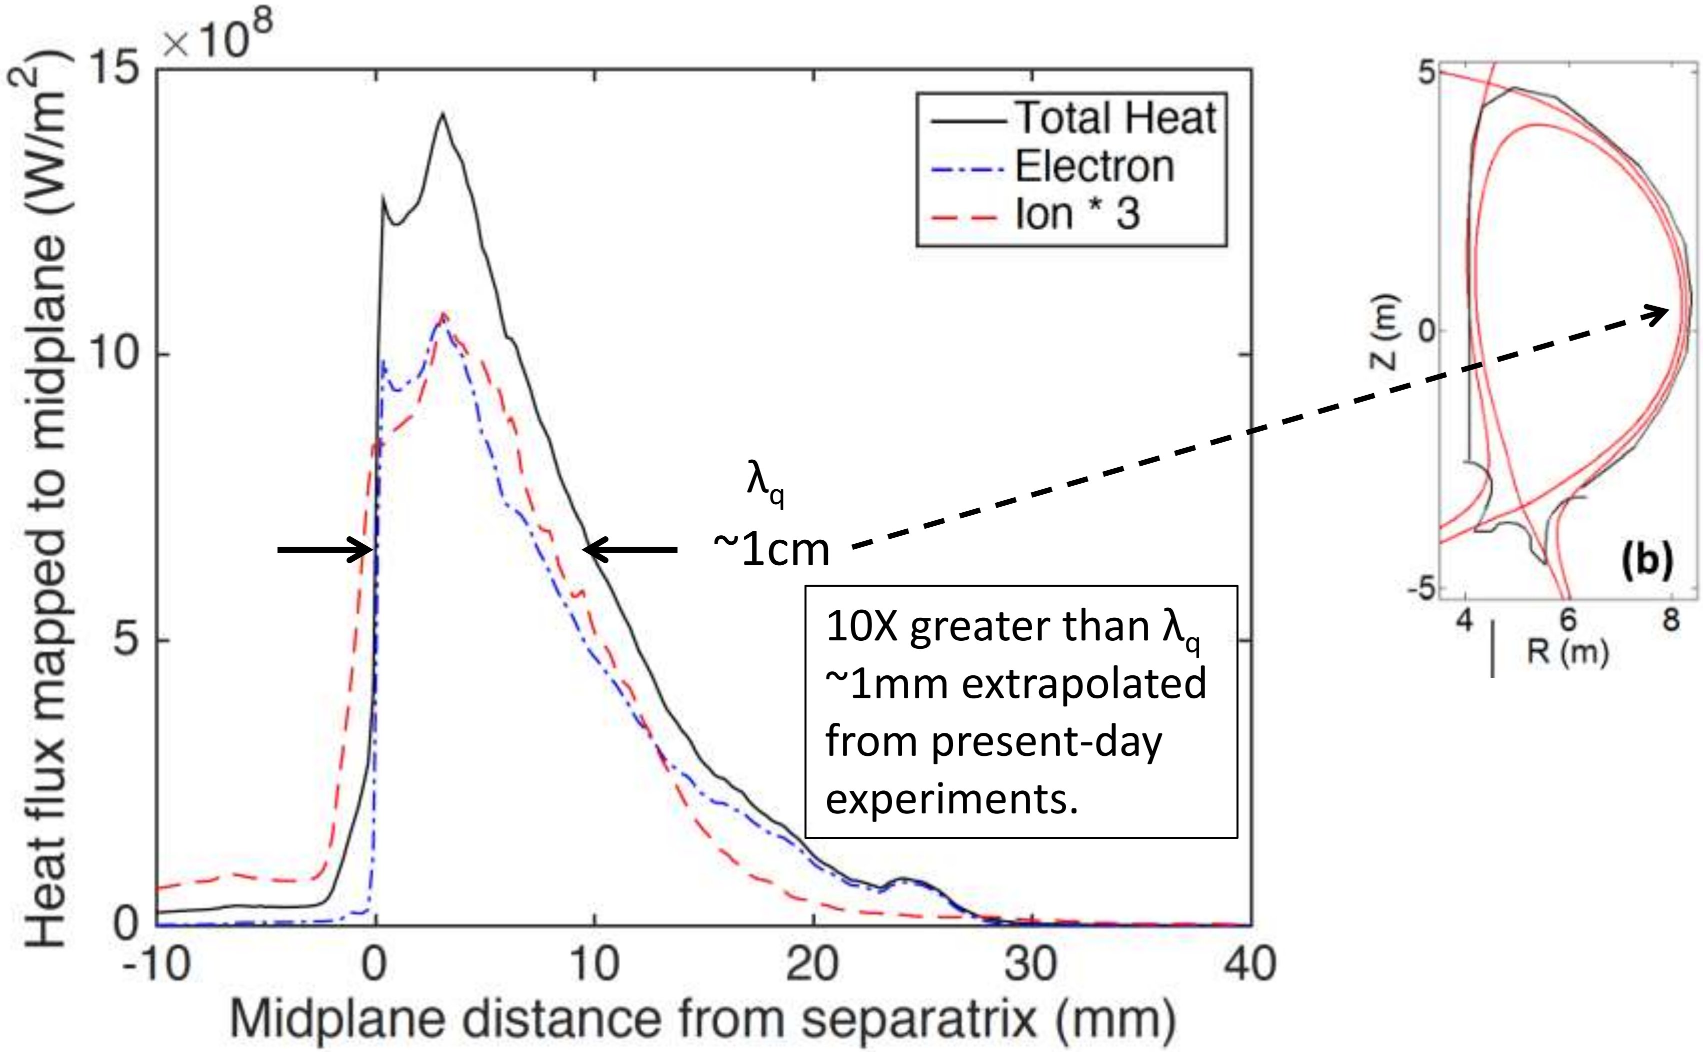

We may need to re-consider the divertor design for ITER and fusion reactors.

Ion particle-loss in ITER from inside the separatrix surface (Titan, 90% capability)

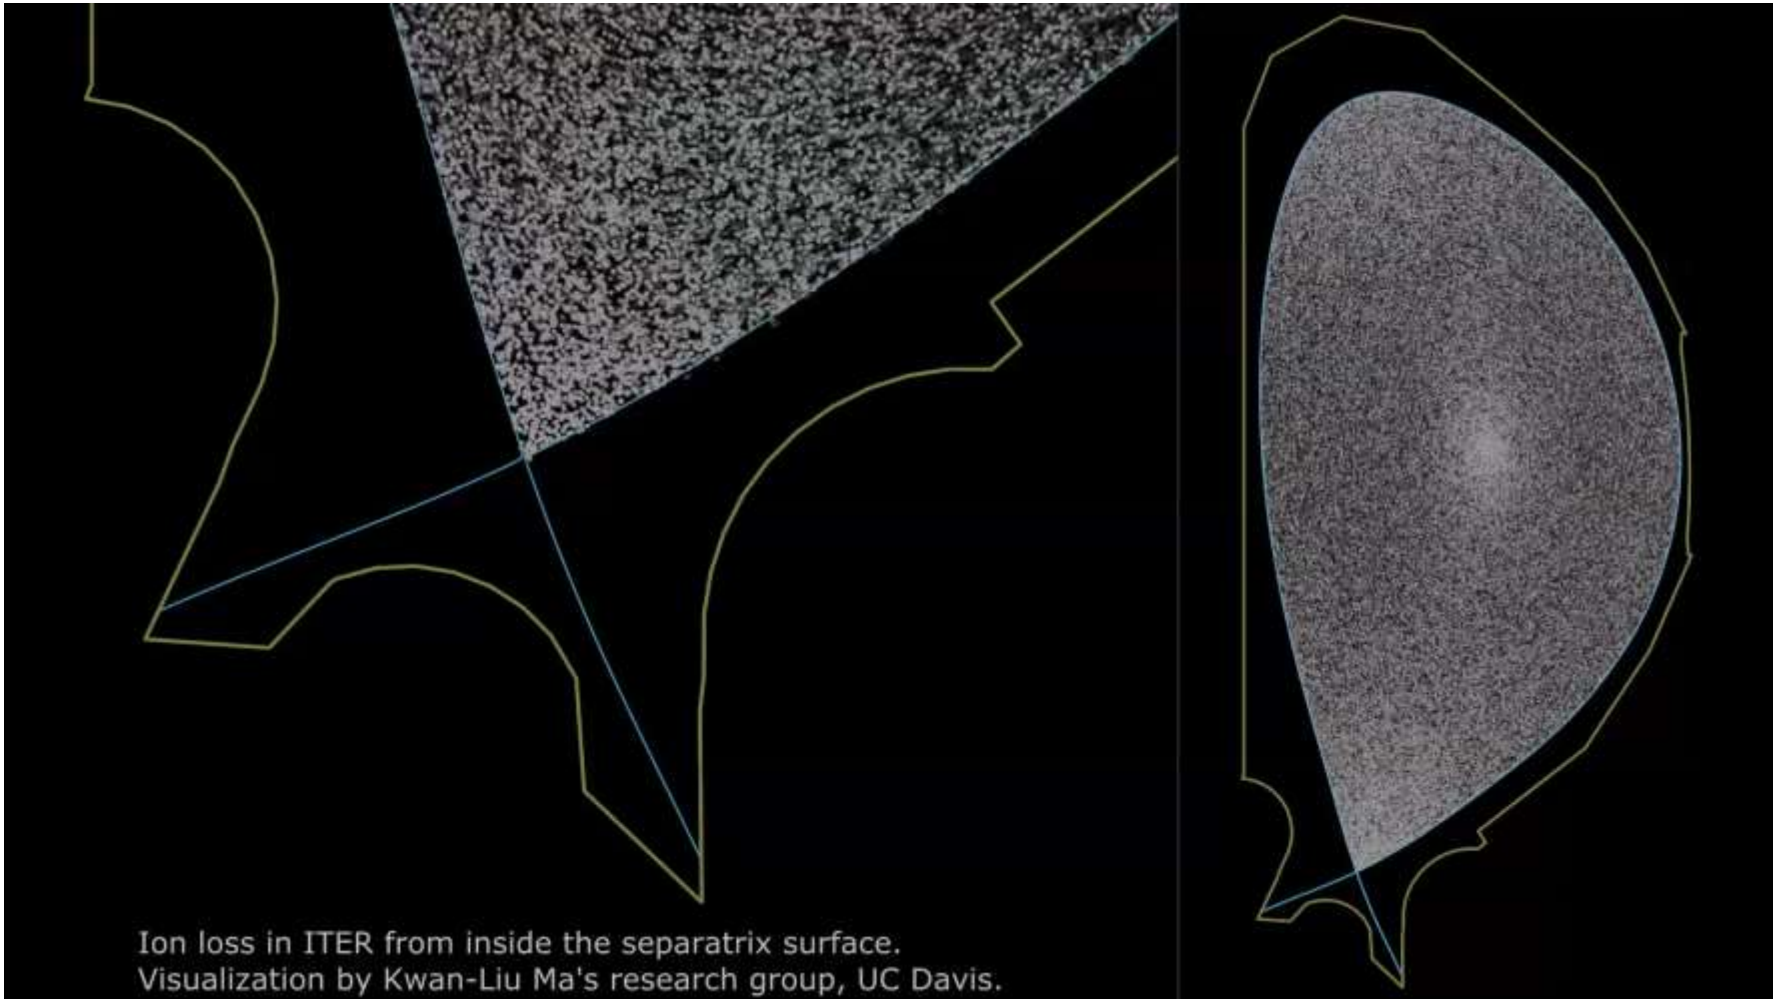

Representative Science on Mira (using 33% capability):
Bootstrap current study in edge pedestal
XGCa finds that textbooks need to be modified.

- Accurate prediction of edge bootstrap current is critical for edge physics
- Bootstrap current is not dominantly from passing particles in edge pedestal.
- Based upon numerous large scale XGCa simulations on 1/3 Mira and new physics understandings, a new analytic formula has been created.

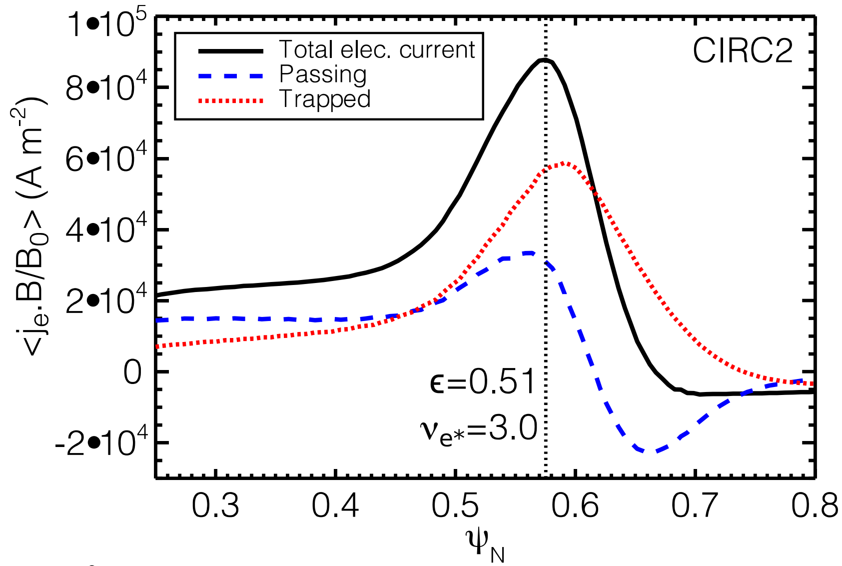

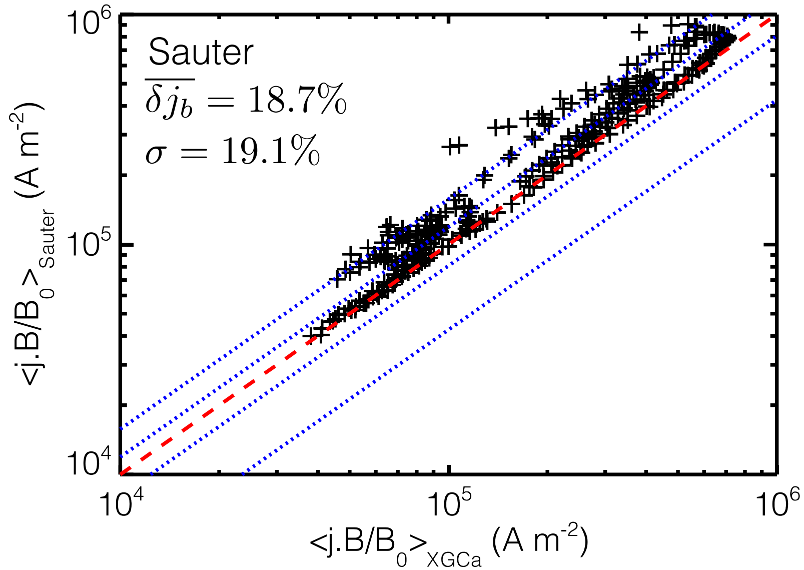

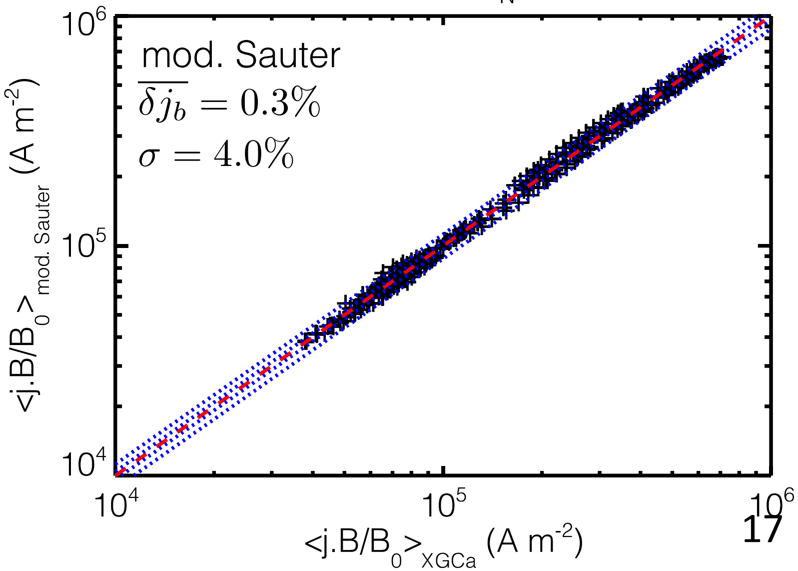

P. Worley and E.D'Azevedo

Up to 40% Performance Improvement from New Hybrid Load Balancing

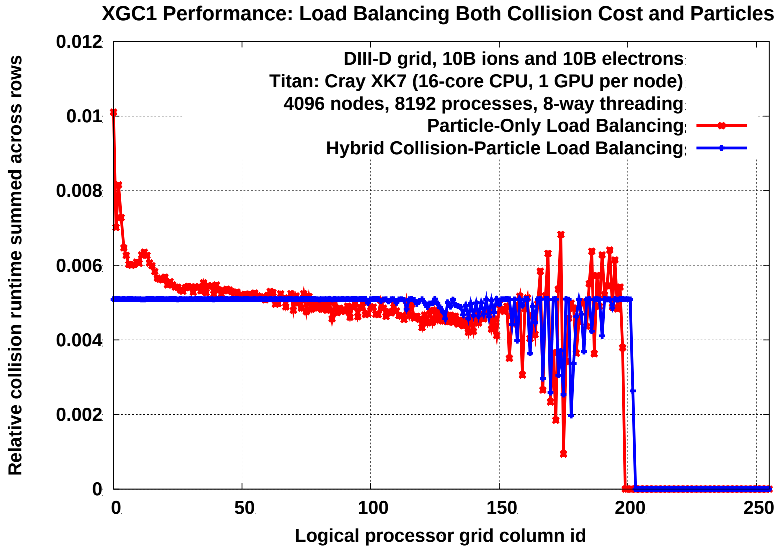

Example comparison of load-balancing only particle distribution (red) with also load-balancing collision cost (blue) across columns of logical 2D processor grid. Cost is summed over rows of grid. Full model performance improvement is 30% for this example.

ITER simulation
Whole code

Timestep cost (secs)	Step 1: Particle-only load balance	Step 2: Particle +Collision
1 outer thread	2230 →0.19X	1072
8 outer threads	793 →0.53X	422

Challenge

- Existing particle load balancing algorithm does not adequately equi-distribute the collision cost in parallel decomposition.

Solution

- Two level automated optimization strategy: (a) balance collision cost subject to constraint on particle load imbalance, (b) optimize XGC1 performance by varying constraint periodically, converging to the optimum if distributions are static and adapting to the changing distributions otherwise.

Result

- 10%-40% improvement for production runs.
- Could be generalized to other similar codes.

M. Adams

Fully implicit E&M Solver for Fluid Electrons in XGC1

- Fast Alfven oscillations put a sever CFL limit to Δt, while physics of interest has much slower dynamics.
- PETSc is used.
- When completed, cost of E&M simulation (even with kinetic electron closure) will not be much greater electrostatic simulation cost.
- Verification work in progress toward realistic physics parameters

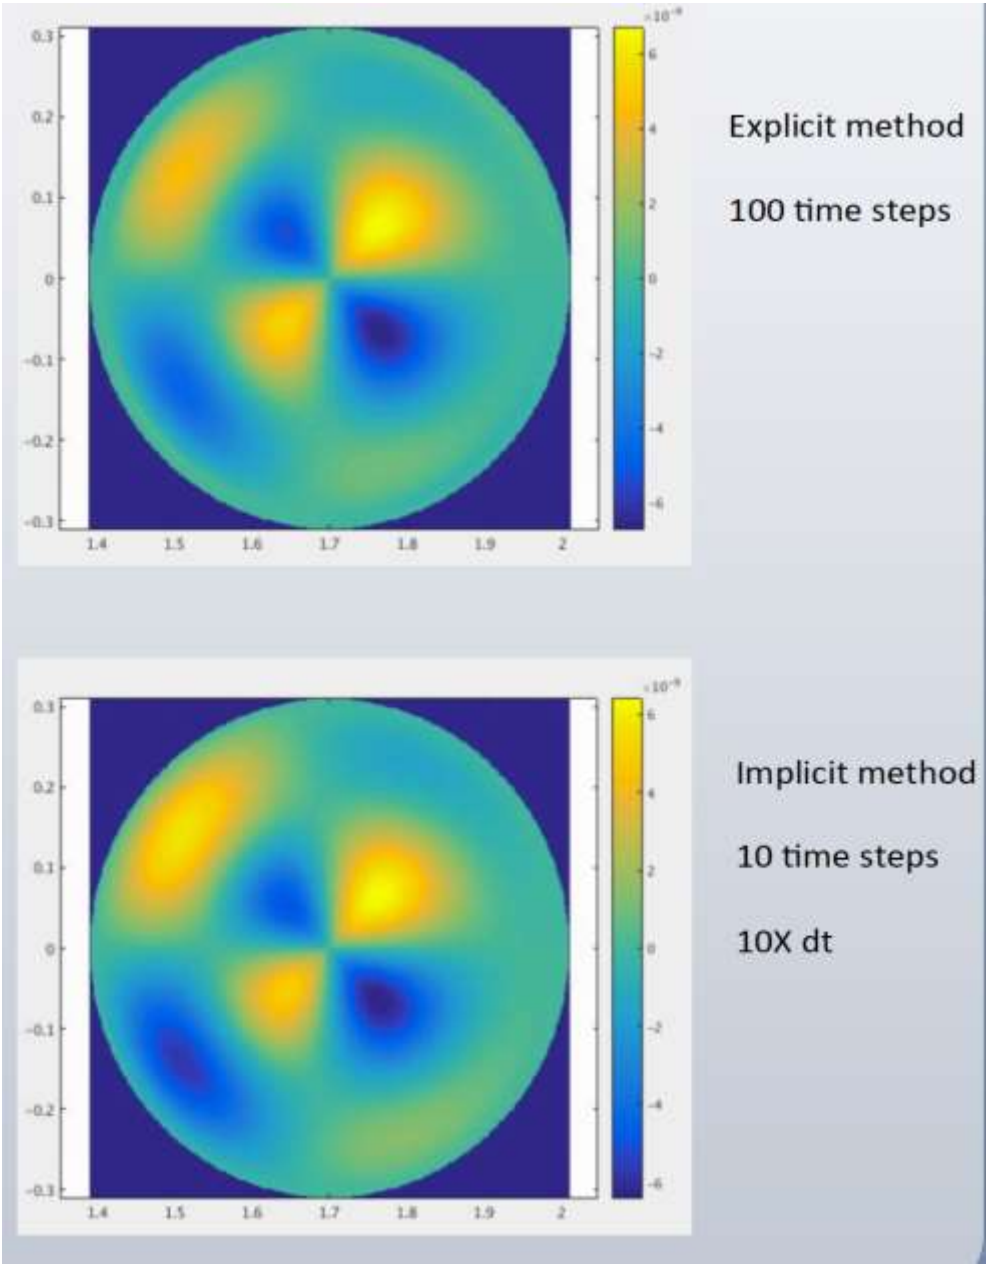

M. Shephard

Improve Fokker-Planck collision eq. solver and meshing

- Due to non-Maxwellian nature of plasma ions in edge, fully non-linear Fokker-Planck collision is important
- The original two species (ion+electron) Fokker-Planck scheme was expensive
  - up to 80% of non-turbulent XGCa time on Mira
  - up to 50% of turbulent XGC1 on Titan
- Used the translational symmetry of Fokker-Planck operator.
  - 4X (2X) improvement in single (multi) species operator
  - Became the base for the hybrid load-balancing scheme (SUPER)
- Variable time-step scheme developed
  - Homework to SUPER: Combining with the hybrid load-balancing scheme [FASTMath+Super]

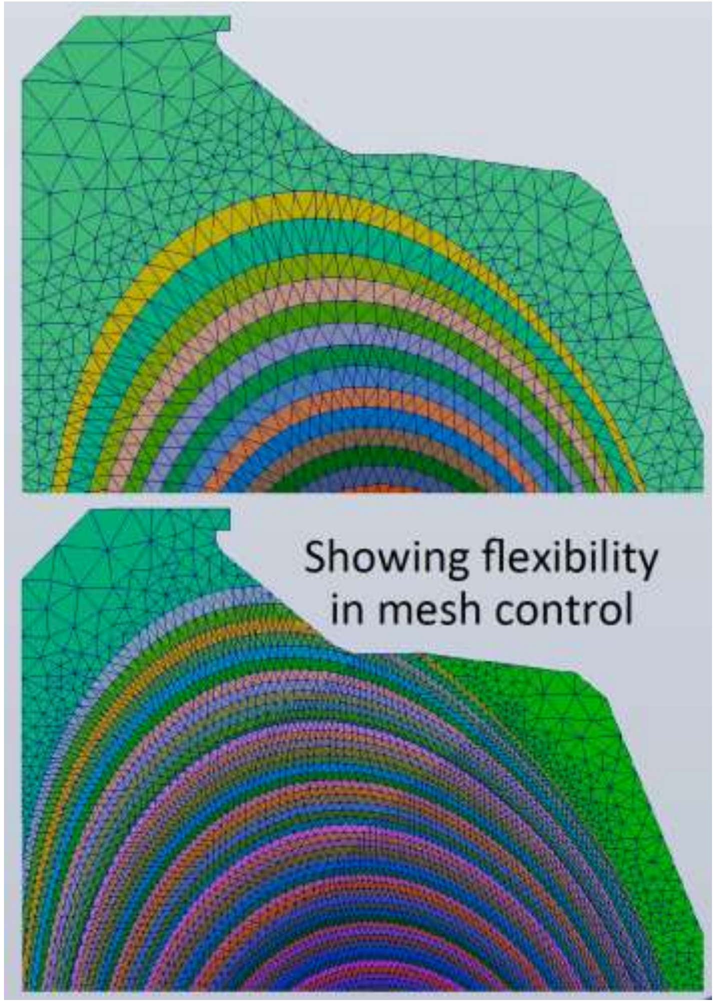

Highly flexible meshing tool to physics and geometry requirements

Algorithm improvement to the Nonlinear Collision Operator

- **O(N2 x N2) operations in 2-dimensional velocity space.**
- **To improve the performance, translational symmetry of the collision operator for coefficient calculations applied to reduce the number operations (~O(1/N) reduction).**
- **Results in up to 35% reduction in total simulation time**

Performance test of XGCa with and without improvements in collision algorithm

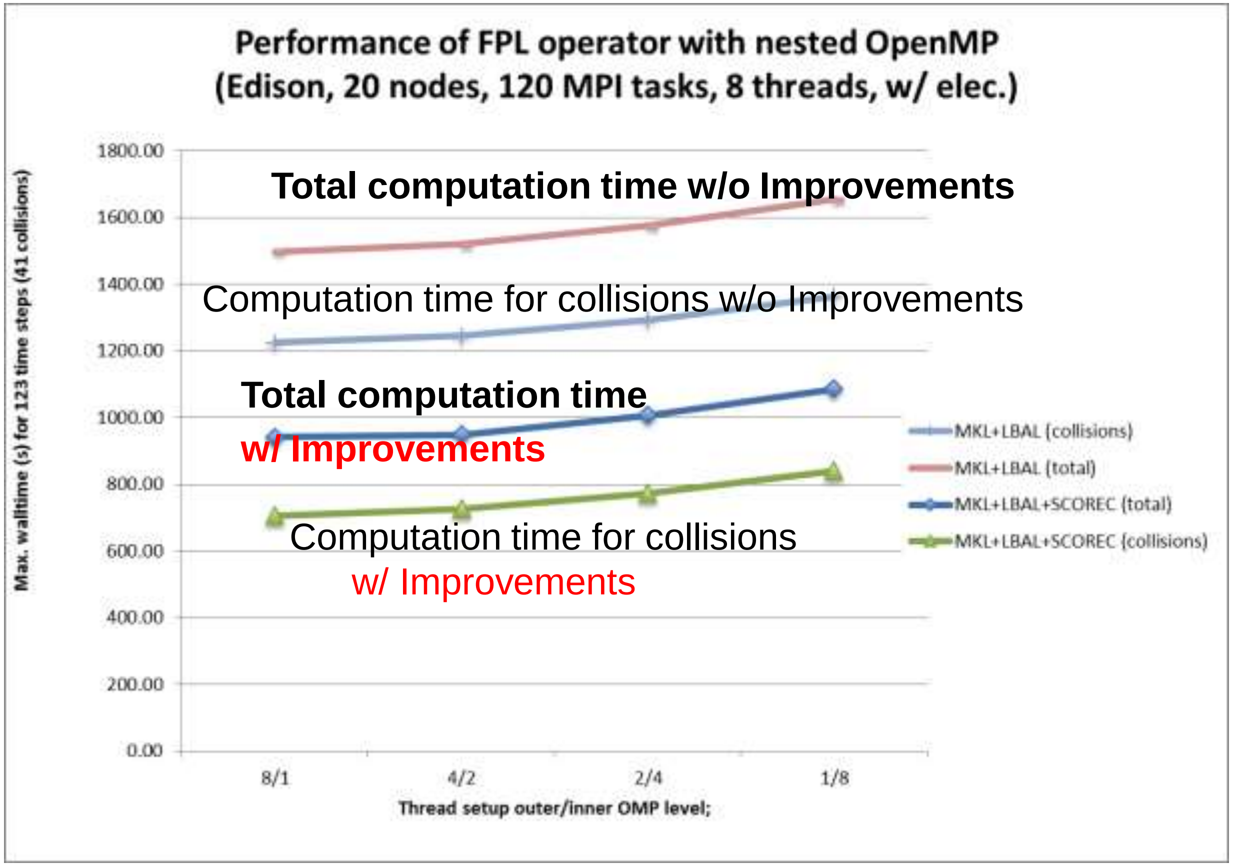

Extreme-Scale DM for EPSI

S. Klasky
M. Parashar

- Staging with DataSpaces has been applied in various scenarios
  - XGC1-XGCa memory-based kinetic-kinetic multiscale time integration
  - XGC1 checkpoint writing: 10x improvement
  - Hierarchical data management (in progress)
- DM performance
  - 10x-200x improvements on turbulence and particle data exchange in the XGC1-XGCa in-memory coupling over the file-exchange coupling
  - Emergency responses: Check-point data size issue for ITER simulation, Contaminated Luster file, Visualization, etc.

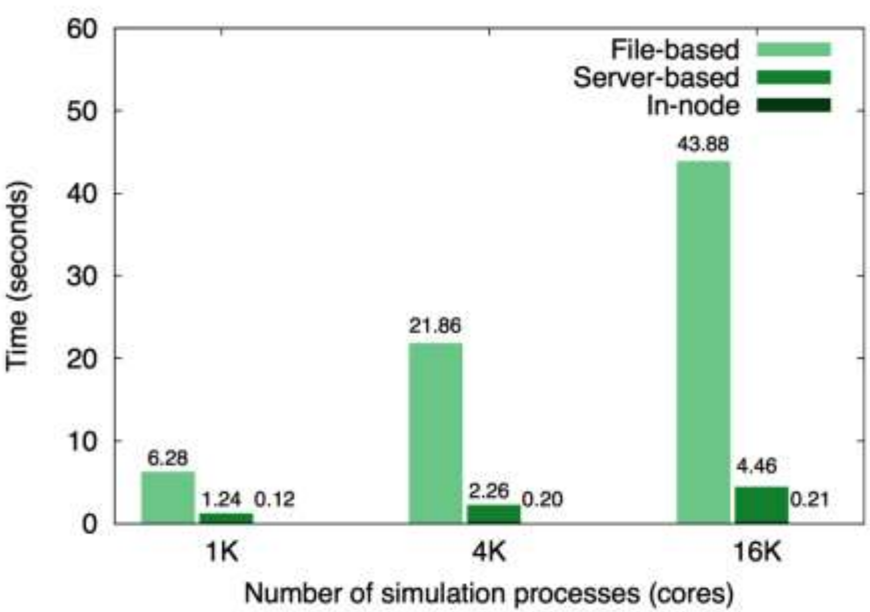

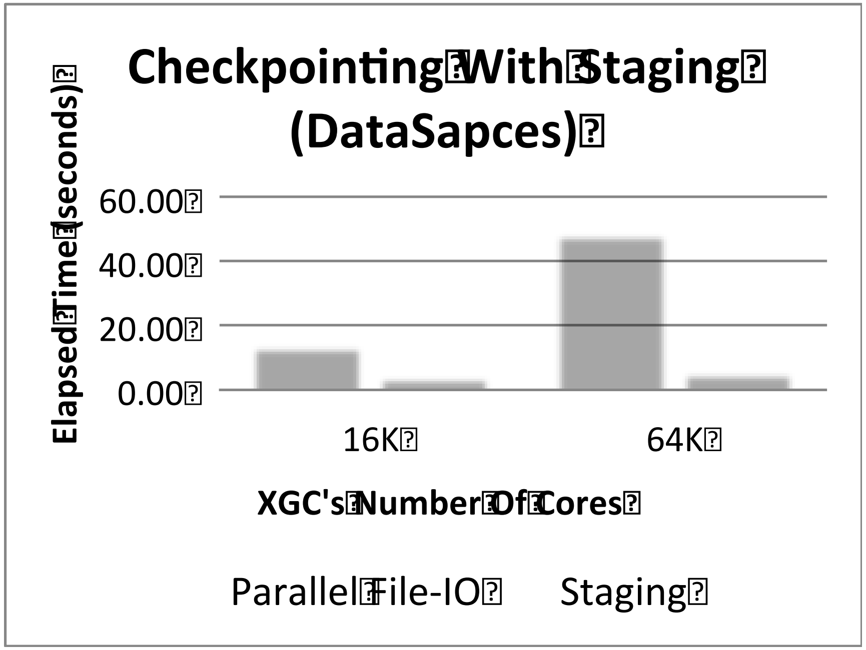

$G(f^F, \Phi^F; b) = 0$
$G(f^C; \Phi^C; \Phi^F, b) = 0$

- Uses the same eqn G
- b=B.D. condition, heating profile, etc.
- Φ(x,t) is E&M field
- Tighter than Heterogeneous Multiscale Method

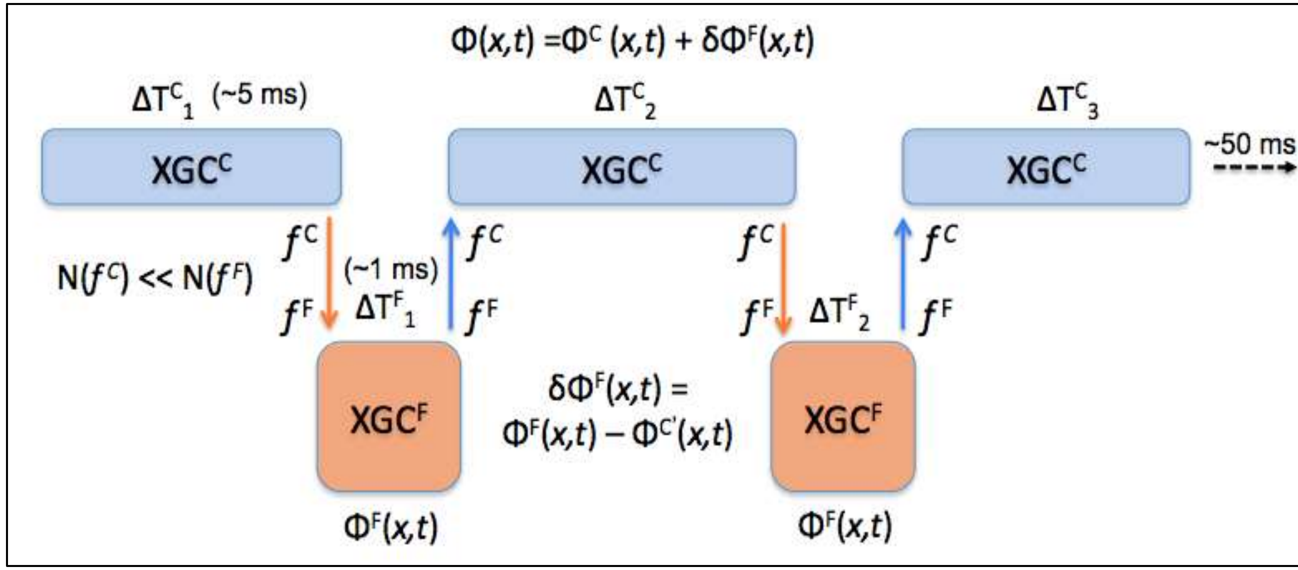

We have built a basic CS framework (Adios+Data Spaces)

Next challenge is mostly in math

- How do we steer the multiscale path to the correct path?

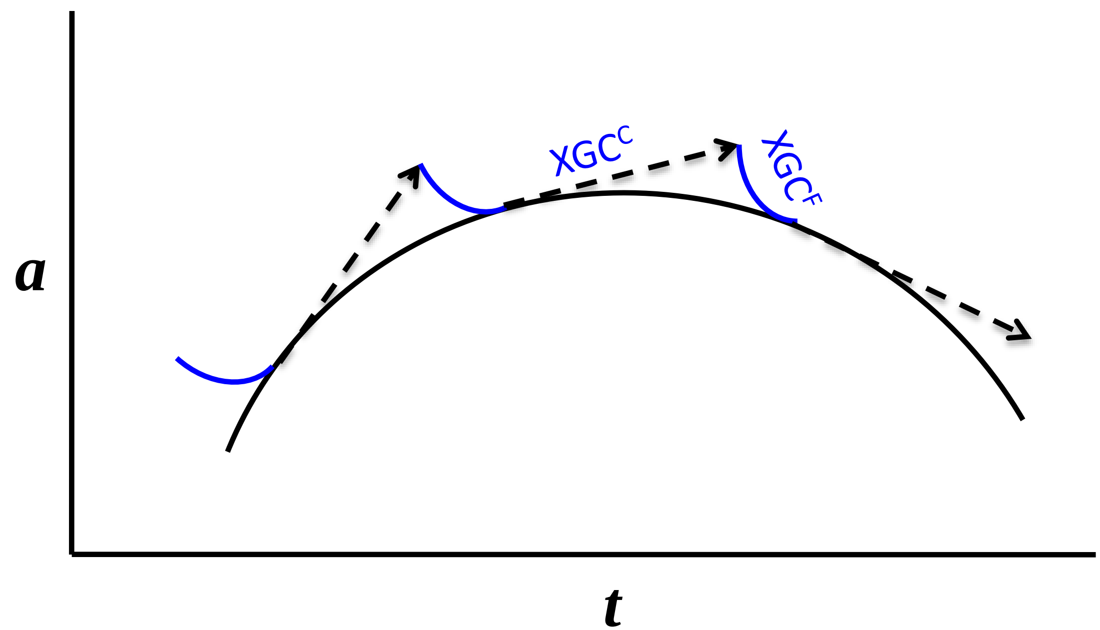

Adios in EPSI

S. Klasky, M. Parashar

Integrated data-centric execution environments for memory-to-memory code coupling, staged data process, and monitoring with a support of dynamic workflow system for leadership class computing.

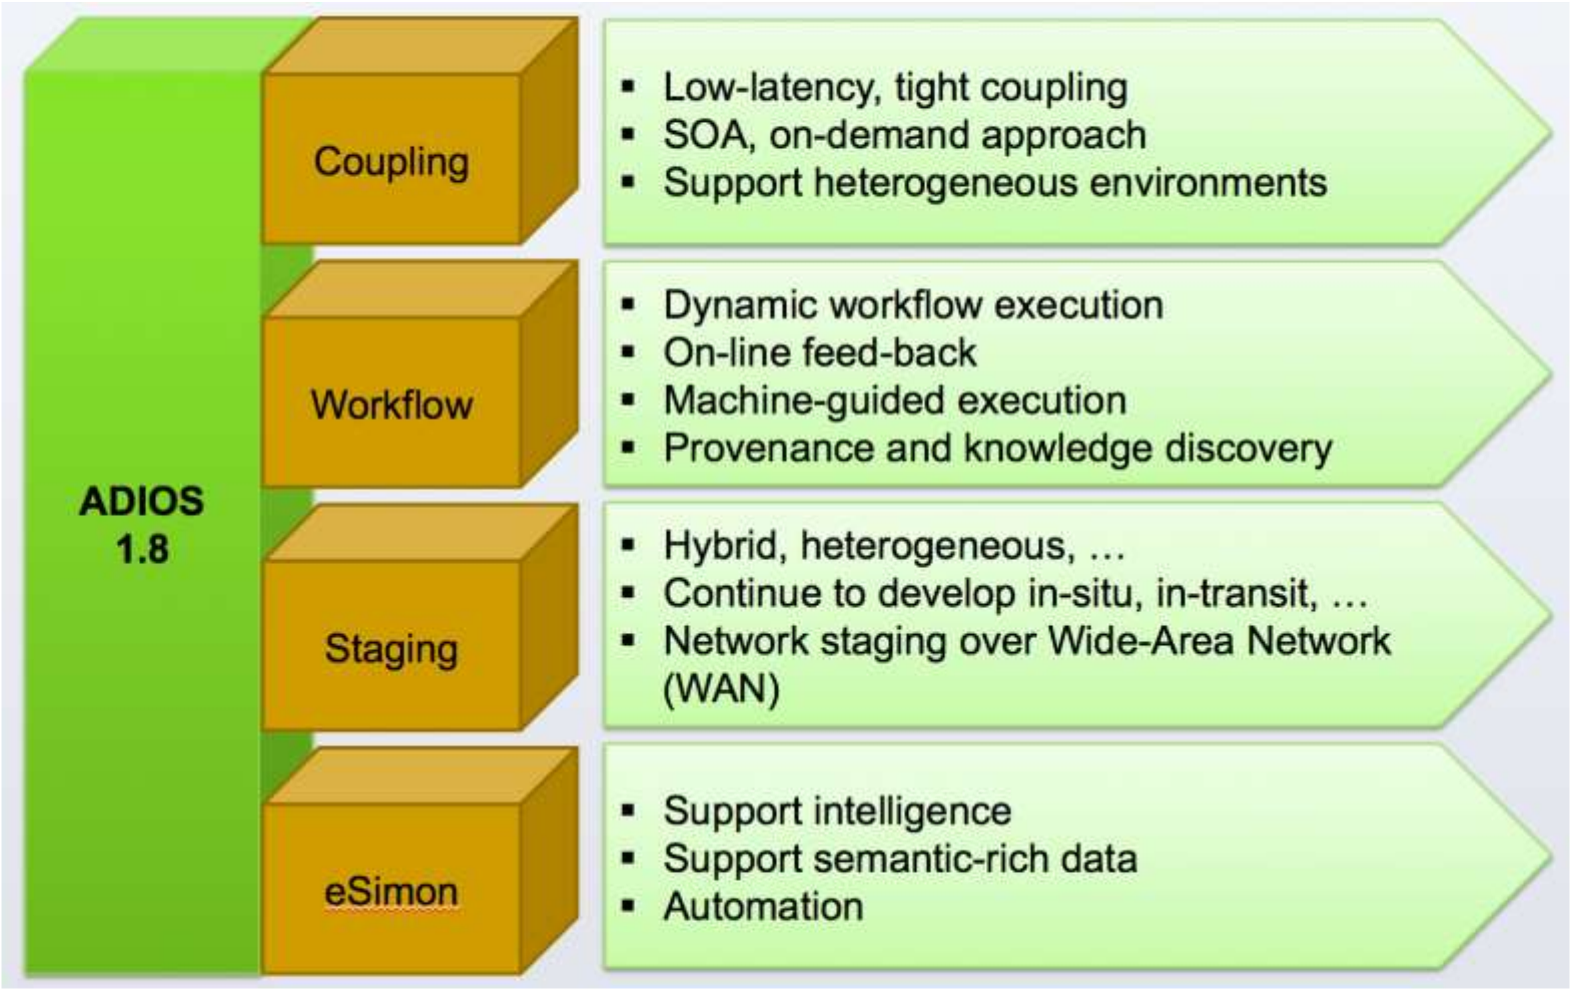

The UQ “telescoping” approach in EPSI

R. Moser

1. UQ in reduced size tokamaks -- more details and samples to construct response surface: Combine into DM
2. Scaling-up the UQ results to larger sizes (*telescoping*): limited number of studies
   - Calibration/enrichment of the telescoped response surface.
   - Validation of predictions against present-day experiments (including UQ on experimental data), and inform surrogate model.
   - Telescope further to ITER-scale, compare prediction against ITER-scale XGC1 simulations

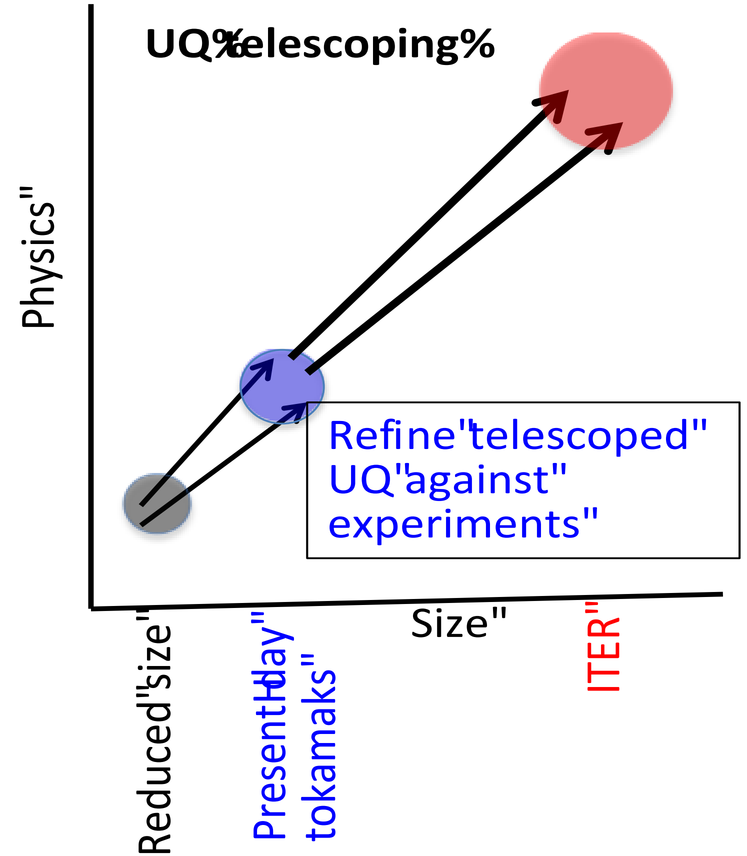

Telescoping and calibration considerations:

- Separate consideration of scale-independent and –dependent quantities.
- Physics guidance important -- response surface should inform, not dictate.
- Negative telescoping results also useful -- identify key regimes in parameter space for high-fidelity simulation (e.g. bifurcations in parameter space), compatible with Expected Information Gain (EIG) base approaches.

Further development of XGC1, with SciDAC Institutes and HPC Centers

Physics capability

Electromagnetic turbulence

- Edge electrons can be more like fluid: Gyrokinetic ions + fluid electrons.
  - This choice removes the "cancellation issue" in the kinetic ion E&M
  - Kinetic electron physics can be added later in the form of closures
- **Utilize the guide work by GEM** (delta-f, core plasma) for technology transfer to XGC1, including the 6D verification work.

Kinetic-kinetic multi-scale integration

Computational capability

- Pre-exascale programs (in CAAR and NESAP)
  - Vectorization
  - Cuda Fortran → OpenACC for easier portability
  - Heterogeneous memory management
  - Multiple GPUs in a node, …….
- **Fault tolerance**
- Implicit and variable time stepping
- In-memory DM, analysis, UQ, …

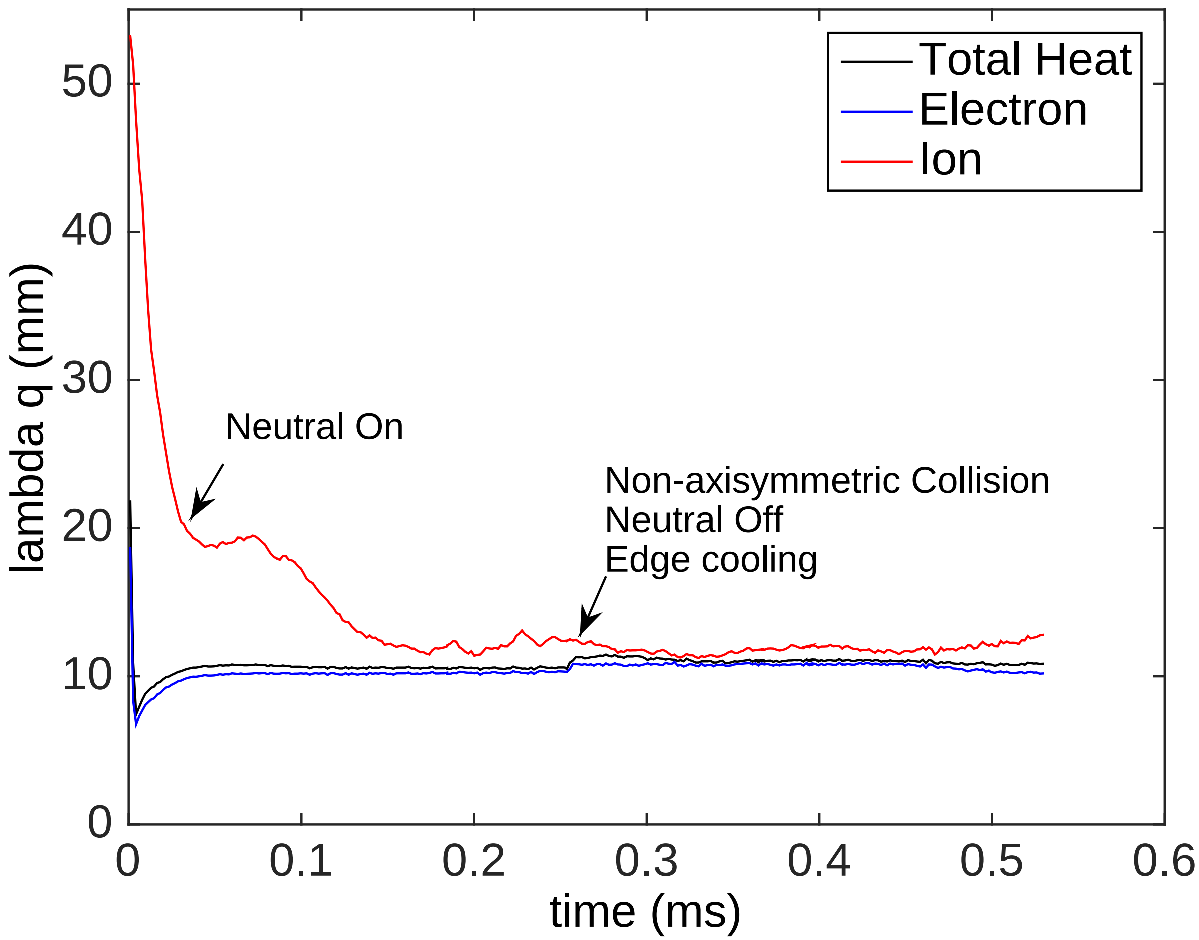

Total Heat
Electron
Ion
Neutral On
Non-axisymmetric Collision
Neutral Off
Edge cooling
lambda q (mm)
time (ms)
0
10
20
30
40
50
0
0.1
0.2
0.3
0.4
0.5
0.6