Pushing the Envelope of Fusion Edge Physics on Extreme Scale HPCs

S. Ku[1], S. Parker[2], C.S. Chang[1], R. Hager[1], J. Lang[1], and the EPSI Team

PPPL[1], U. Colorado Boulder[2]

Institutional Leaders

OFES
U. Colorado: S. Parker
Lehigh U.: Arnold Kritz
MIT: Martin Greenwald
UCSD: George Tynan

OASCR
LBNL: Mark Adams (FASTMath)
ORNL: Scott Klasky (SDAV), P. Worley, E. D'Azevedo (SUPER)
RPI: Mark Shephard (FASTMath)
Rutgers U.: M. Parashar (SDAV)
U. Texas: Bob Moser (QUEST)

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

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

"Toroidal" Tokamak Geometry

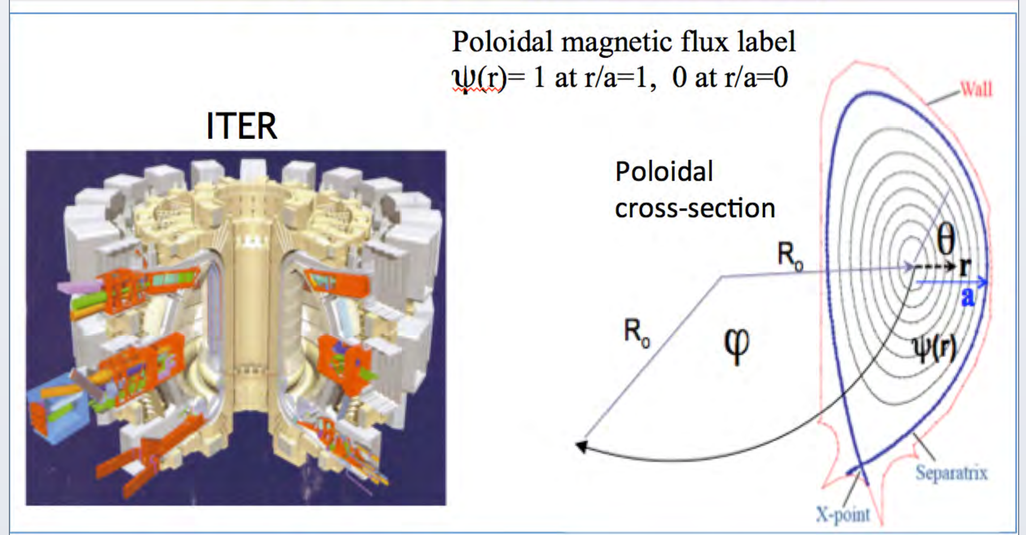

Torus, not a straight cylinder: physics becomes more complicated through the magnetic inhomogeneity and the toroidal mode coupling.

What science are we studying?

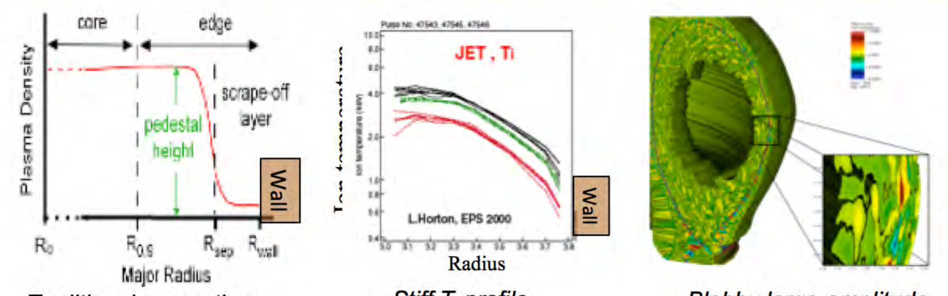

Traditional separation: Edge/pedestal/scrape-off — *Stiff T_i profile T_i(center) follows T_i(edge)* — *Blobby large-amplitude edge turbulence*

- Edge plasma self-organizes into a steep pedestal shape (H-mode). → Smaller & cheaper tokamak, by allowing a hot plasma at plasma edge.
- Challenging physics: Scale inseparable plasma, turbulence and neutral particle dynamics across steep gradient, and in contact with material wall
  - **Non-equilibrium thermodynamics**, hence non-Maxwellian
    → "Fluid closure" is difficult.
  - Large amplitude nonlinear coherent turbulence structures, "blobs."

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.

XGC1 Performance: Weak Particle Scaling on DIII-D grid

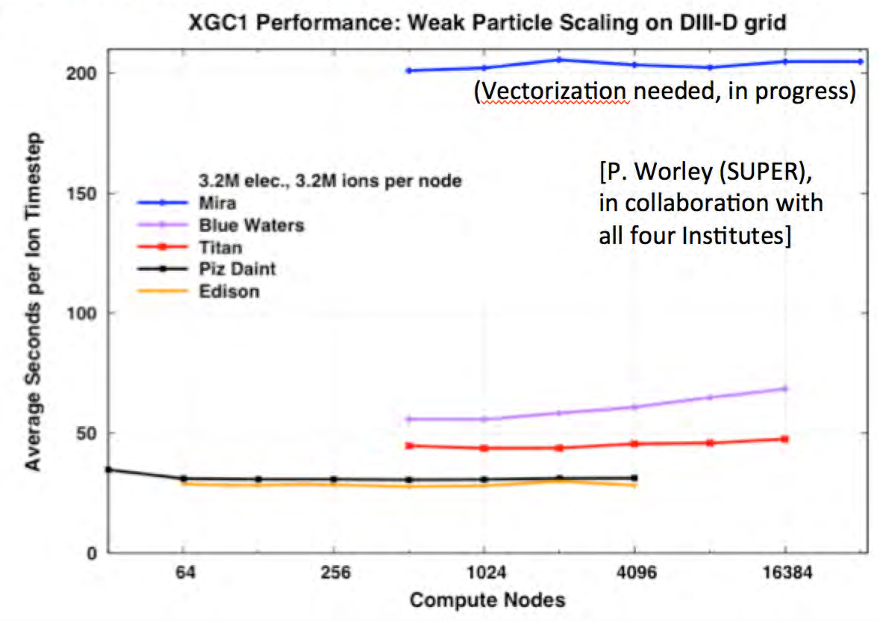

Computing Resources for XGC in 2015

- INCITE: 270M hours
  - Titan: 170M hours (Extreme scale jobs with full physics, 10-20PFs, usually 90% capability computing)
  - Mira: 100M hours (Large scale jobs with partial physics 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

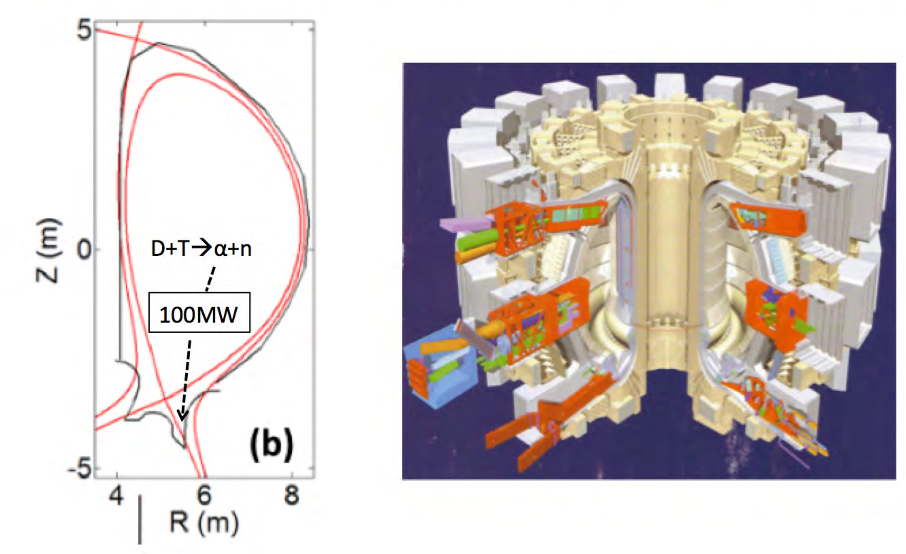

--A representative science case-- Divertor heat-load width: a serious issue for ITER

- 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 λ_q≈1mm → need research when the technological limit ~10MW/m²
  - ✧ 100 MW on an extrapolated strip → ≈ 20MW/m² 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 in non-equilibrium kinetic state: non-Maxwellian, and scale-inseparable multiscale.
  → **Extreme scale computing**

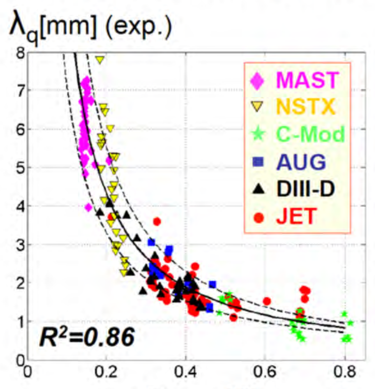

T. Eich et al., NF 2013

Ability to produce the nonlinear "blobby" edge turbulence is a pre-requisite for heat-load study

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

- Experiments: edge plasma is in "blobby" form
  - Large amplitude density and electric potential blobs (~50%)
  - Only simple "models" existed
  - Computationally difficult and expensive
- Titan and ASCR collaboration enabled gyrokinetic blob production in XGC1, for the first time (reported in SciDAC-14)
  - Ab initio understanding of blobs together with background kinetic physics
  - In realistic diverted geometry
  - OLCF Featured Highlight, 2/2014
  - Stage was set for the divertor heat-load footprint prediction

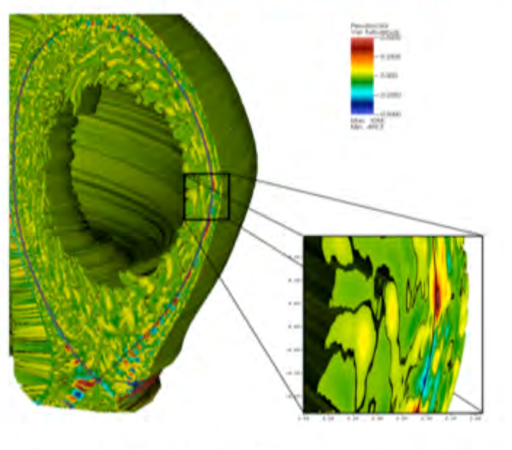

Simulation by S. Ku, Visualization by D. Pugmire

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

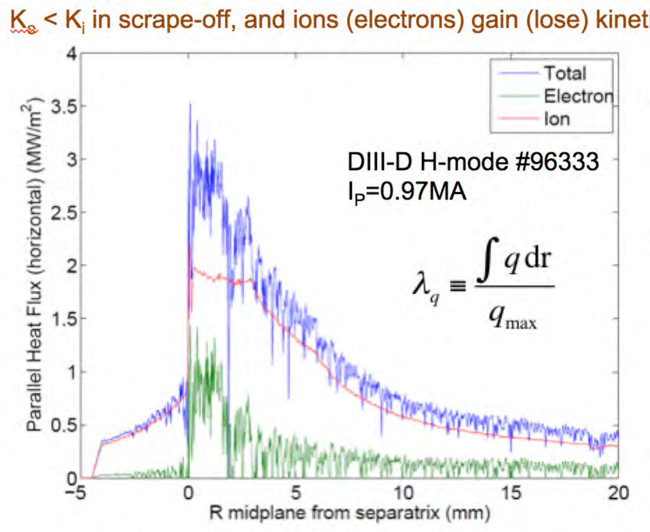

$$\lambda_q = \frac{\int q\,dr}{q_{max}}$$

- 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 radial blob size (>1cm)

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

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

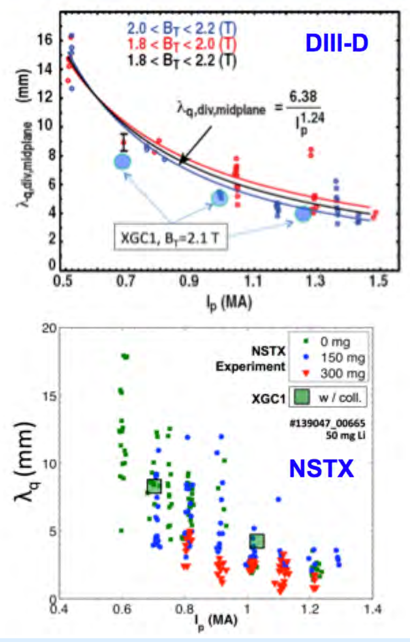

- $\lambda_{q,mid}$ from XGC1 agrees with experimental values from two very different tokamak devices
  - ✧ DIII-D for conventional aspect ratio
  - ✧ NSTX for tight aspect ratio
- Broadening of $\lambda_{q,mid}$ by blobs (width >1cm) is found to be insignificant in present machines.
- Will the blobs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

Turbulence driven radial drift is much weaker than the ion orbital drift in present-day tokamaks → $1/I_p$ scaling

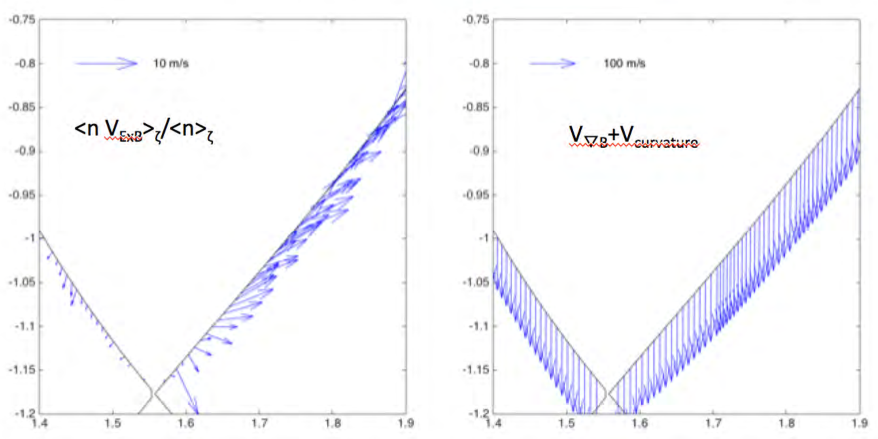

These roles may reverse in ITER because $V_{\nabla B}+V_{curvature}$ is much smaller, while the blob effect may not decrease, compared to the present-day tokamaks.

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.
  - **XGC1 prediction does not agree with a simple extrapolation from the present-day experimental data**
  - **The heat-load width in ITER from XGC1 is ~1cm for both electrons and ions, instead of ~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.**
  - A rigorous verification activity will follow.

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

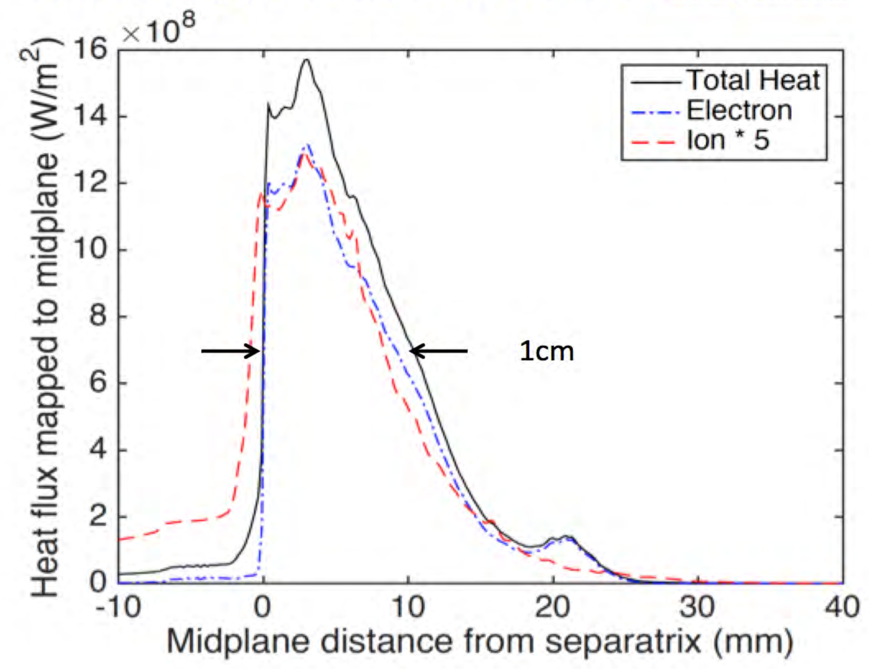

Representative Science on Mira (using 33% capability): Bootstrap current study in edge pedestal

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 and new physics understandings, a new analytic formula has been created.

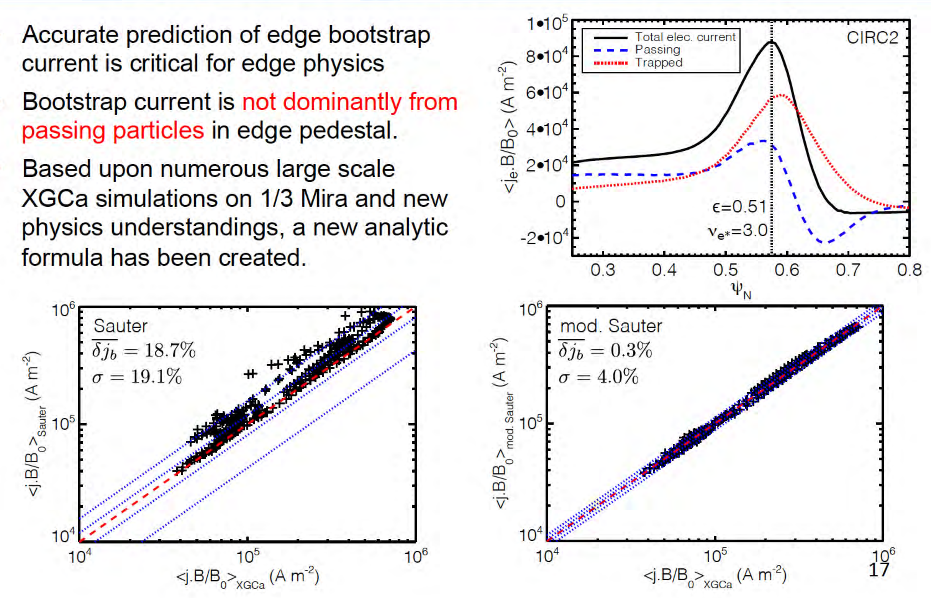

Further development of XGC1, with SciDAC Institutes and HPC Centers

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 good 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, …

Cross-verification of E&M modes between XGC1 and GEM

The toroidal mode number (n) scan for the linear growth rate shows a similar transition from (kinetic) peeling to kinetic ballooning modes between XGC1 and GEM

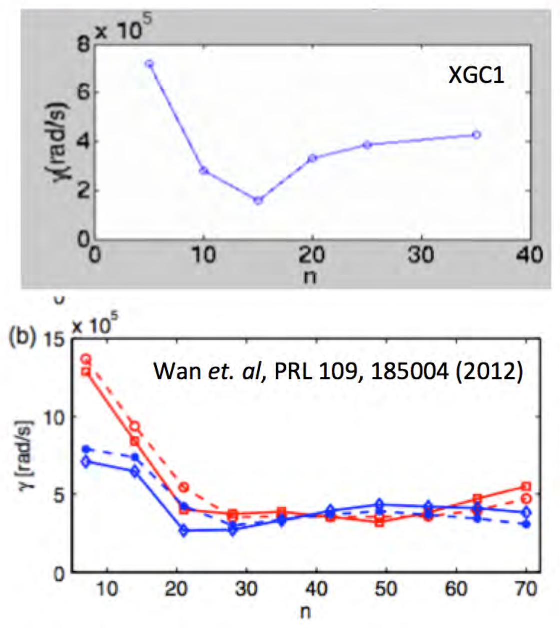

ITG-KBM Transition Verification

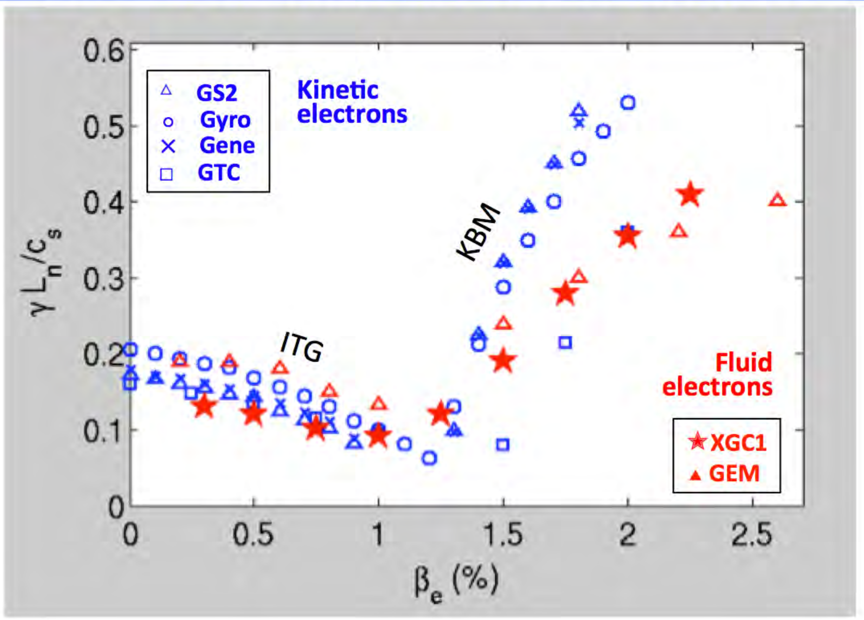

6D particle dynamics for validating 5D gyrokinetic equations (U. Colorado)

- 6D Lorentz ions yield the lowest level ab initio simulation
- 2nd order implicit orbit averaging and sub-cycling algorithm
  - GPU is utilized efficiently: 50X speedup over single CPU
- 6D simulation verified by analytic theory

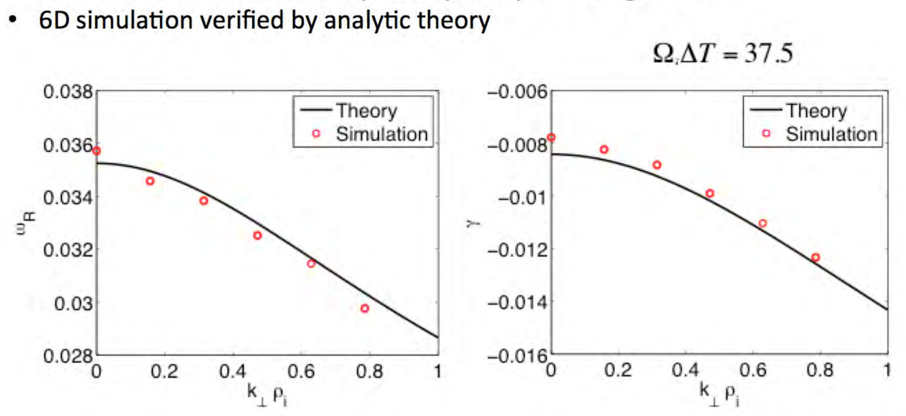

Kinetic-Kinetic Multiscale Integration

- $G(f^F,\Phi^F;b)=0$
- $G(f^C,\Phi^C;\Phi^F,b)=0$
- Uses the same eqn G
- b=B.D. condition, heating profile, etc.
- $\Phi(x,t)$ is E&M field
- Tighter than Heterogeneous Multiscale Method

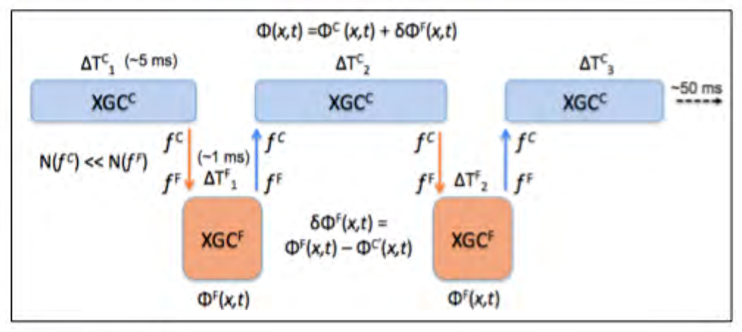

- We have built a basic CS framework (Adios+Data Spaces)
- Next challenge is mostly in math and physics
  - ✧ How do we steer the multiscale path to the correct one?

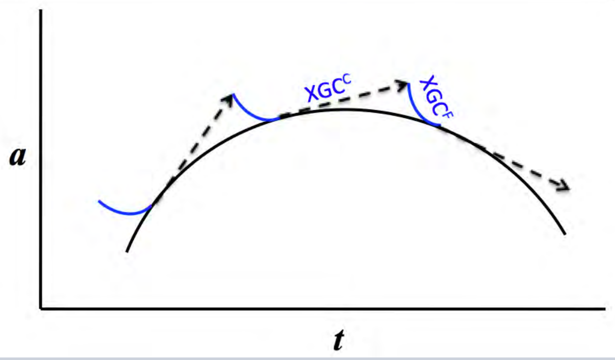

Having the correct multi-scale dynamics is very important in the multi-scale code coupling

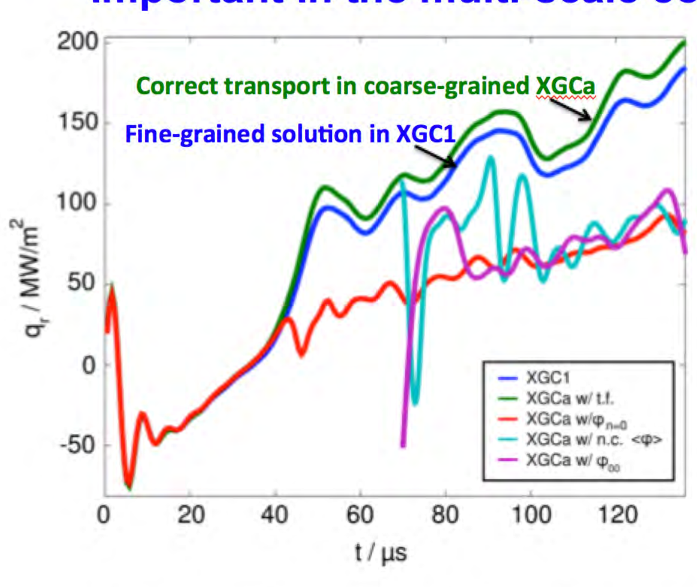

- Any modification of either the neoclassical or turbulent large scale dynamics yields incorrect transport.
- Neoclassical dynamics, consistently with turbulence, is important.

We are approaching the point where we need help from applied math.