Large Scale Computing and Storage Requirements for
Fusion Energy Sciences
Aug. 3-4, 2010, Washington DC

Plasma Turbulence and Transport

C.S. Chang[a)], J. Candy[b)], S. Ethier/W. Wang[c)], and Z. Lin[d)]

[a)] Courant Institute of Mathematical Sciences, New York U.
[b)]General Atomics
[c)]Princeton Plasma Physics Laboratory, Princeton U.
[d)] U. California, Irvine

Introduction

- Tokamak confinement in the low confinement mode (L-mode) is dominated by turbulent transport.
- In the high confinement mode (H-mode), the ion turbulent transport is subdued, at least, in some layers (internal and edge transport barriers); and the magnetic curvature-driven neoclassical transport determines the ion confinement.
- Electron transport is found to be always turbulent.
- Toroidal rotation appears to be generated/transported by ion-scale turbulence, and to improve tokamak confinement.
- **Gyrokinetic turbulence codes push the limit of** Large Scale Computing and Storage Requirements → Presented here.

Gyrokinetic codes for fusion (1)

- **Gyrokinetic:** Reduce 6D (x,y,z, v1,v2,v3) to 5D ($x,y,z,v_{||},v_{\perp}$) by assuming that the gyrofrequency is much faster and that the gyroradius is much shorter than the scales of physical dynamics of interest.
- **Two complimentary approaches** in solving the Vlasov equation df/dt=C(f) and the Maxwell's equations
  - **Continuum:** solve the whole PDE system on 5D grid (**GYRO**)
    - Not optimized for large-scale parallel computing (CFL limit)
    - Optimized for fast production runs on smaller number of processors with larger memory → many users on local clusters!
    - Running many copies on many flux tubes have been developed (**TGYRO**)
  - **Particle-in-cell:** solve the original marker particle dynamics in 5D space, solve the Maxwell's equations on 3D position grid **(GTC, GTS, XGC1)**
    - Optimized for large-scale parallel computing
    - Smaller memory requirement for high resolution calculation--random sampling
    - Statistical particle noise 1/Sqrt(N) → Needs smoothing or large enough N

Gyrokinetic codes for fusion (2)

- **Full-f and Delta-f**
  - **Full-f:**

    Mean and perturbed physics are not separated (**XGC1, unstructured**) → Requires large & extreme-scale computing
  - **Delta-f=f_{full} – f_0 ,** assuming conserved system**:**

    Gyrokinetic codes calculate perturbed turbulence physics only (**GYRO, GTS, GTC**). The mean part is evaluated by an external mean-plasma code (**TGYRO**, GTC-neo, XGC0) → Efficiency computing.

Science case: GYRO on ECH heated DIII-D core plasma (512 cores)

Black, 2500 ms; Red, 2525 ms

- 40 toroidal modes
- 500 radial and 10 poloidal gridpoints
- 128 v-space gridpoints
  - ➢8 energy x 8 pitch angle x 2 signs
- 3 kinetic species (D^{+1}, C^{+6}, e^{-1})
- $\Delta x/\rho_s = 0.3$ and $0 < k_\theta \rho_s < 2.5$

- 178,600 MPP hours on 2,560 cores
- Local simulations were ~5X faster.

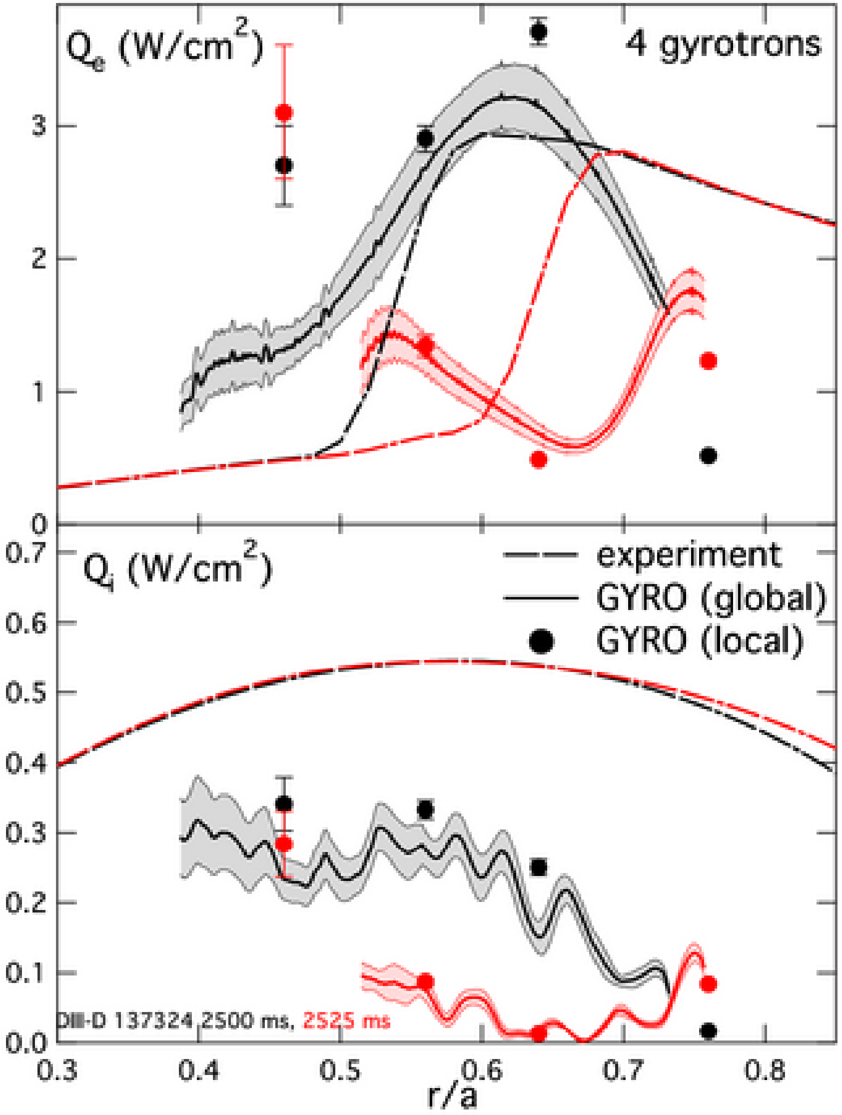

Study performed by
C. Holland, July 2010

Some remarks on GYRO

- GYRO was heavily optimized for single-core and vector (Cray X1) systems, but performance is non-optimal on current multi-core platforms.
- Users are increasingly targeting multi-scale simulations (resolving ion and electron scales simultaneously). GYRO functions, but has not been optimized, for this dramatically more challenging regime.
- SciDAC project (CSPM fall 2010) includes significant plans for performance analysis and re-optimization on multi-core and for multi-scale cases.
- GYRO is used by many researchers and most fusion labs worldwide. It has been the basis for numerous Ph.D. thesis projects* which are based on experimental data analysis. Re-optimization is critical for these users with limited CPU resources.

  *L. Lin on C-Mod, Casati in Tore Supra, Hein on Asdex, Pusztai on DIII-D, etc

Science Case: TGYRO (500x20=10K cores)

TGYRO manages execution of multiple instances of the kinetic neoclassical code NEO and the gyrokinetic code GYRO. Equilbrium profiles of n and T are modified by Newton iteration until measured losses from collisions (NEO) and turbulence (GYRO) balance experimental power and density sources.

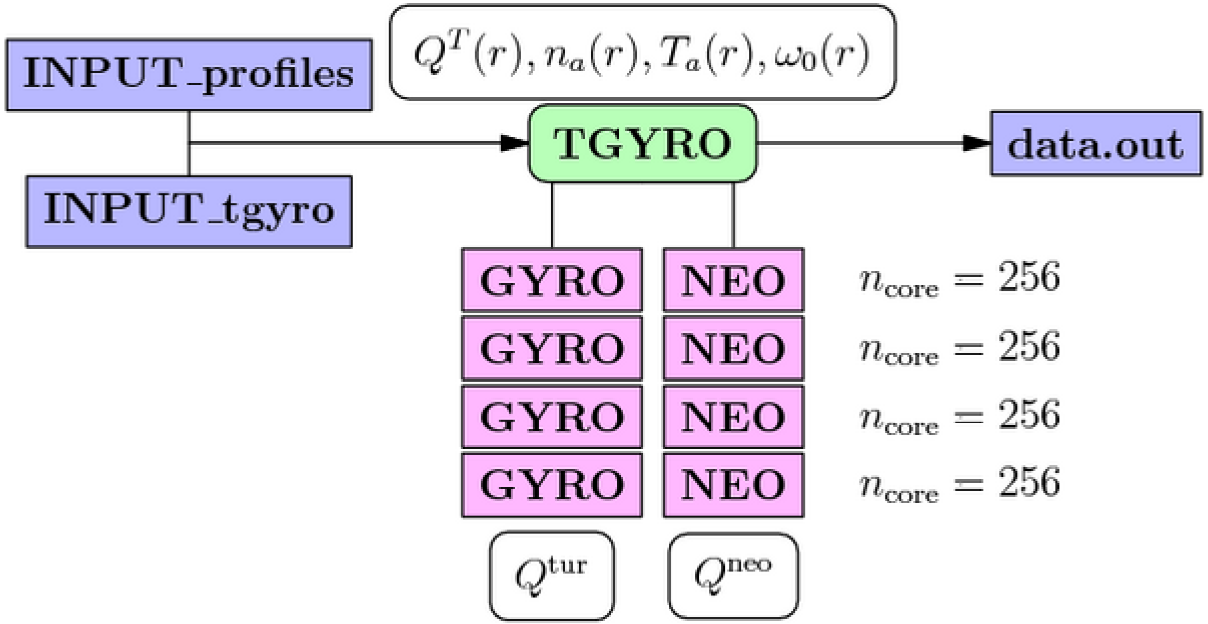

For a simple 4-radius case, the resource requirement jumps from a few minutes on 4 cores to 12 or more hours on 1024 cores. Increasing resolution in both GYRO and TGYRO quickly increases the core demand to greater than 10,000 cores for 24 or more hours.

TGYRO results (red curve) compare well with experiment (blue curve)

For this case, 10 simulation radii (10 instances of flux-tube GYRO) were used.

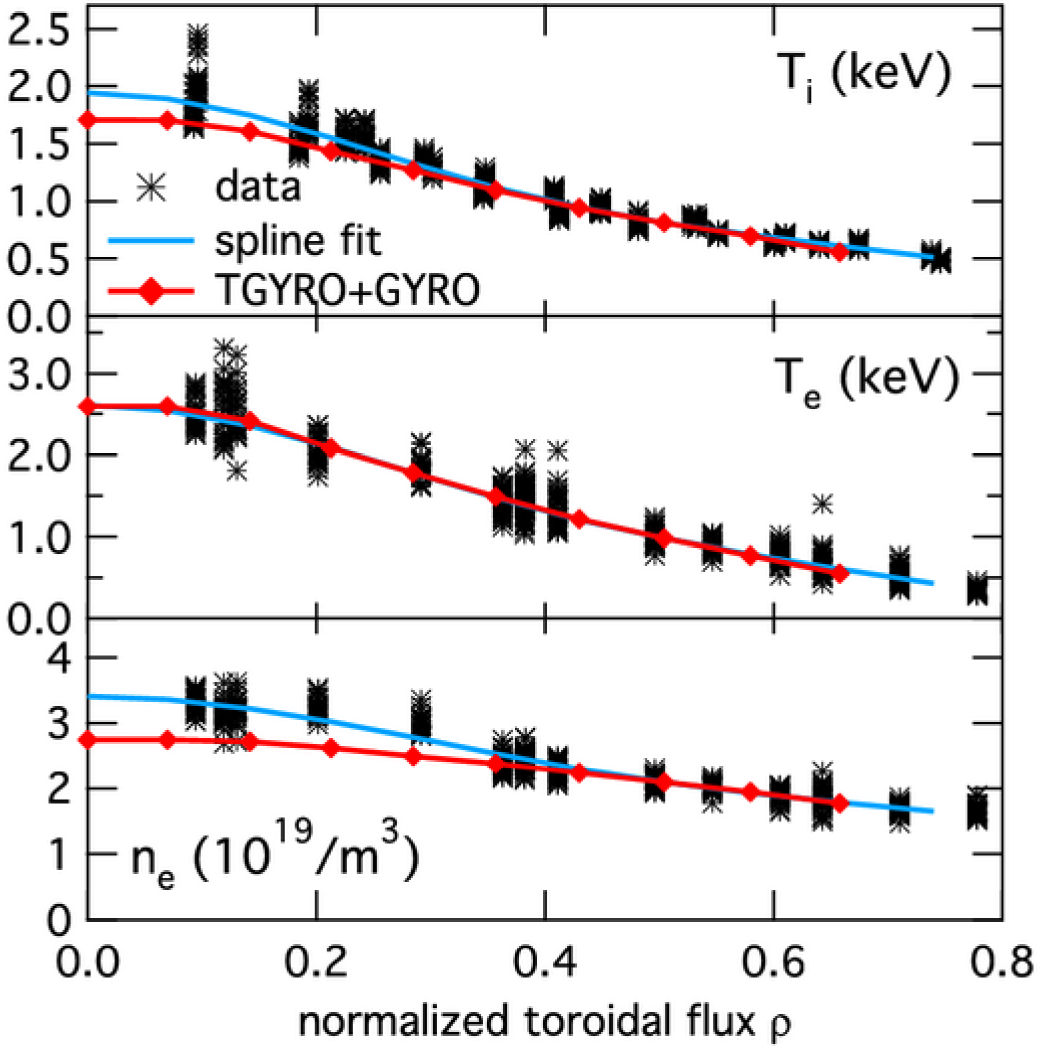

Science Case: GTS study of toroidal momentum generation and transport (8K-98K cores)

Control macroscopic stability; reduce micro-turbulence and energy loss

Very complex transport phenomena: anomalous, non-diffusive, non-local

- Toroidal plasmas can self-develop rotation without external torque!

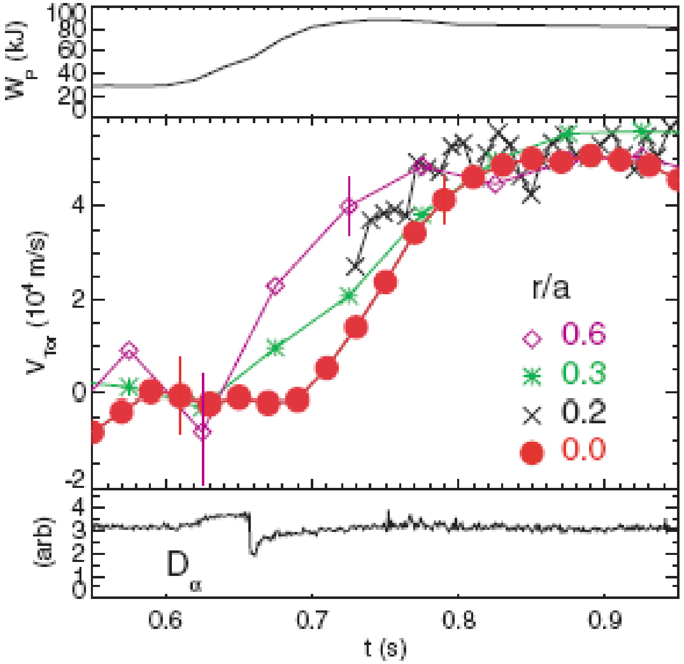

(Rice et al. '04)

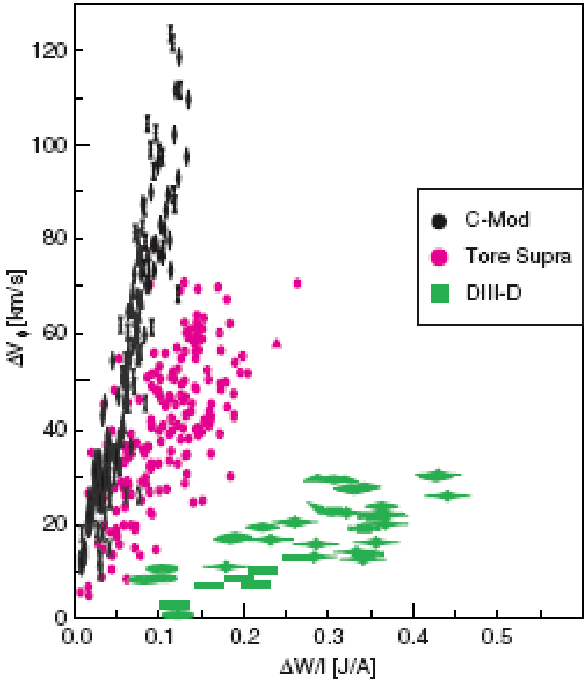

(Rice et al. '07)

Particularly important for ITER – intrinsic rotation will dominate

GTS finds Nonlinear Residual Stress can drive momentum efficiently –CTEM case

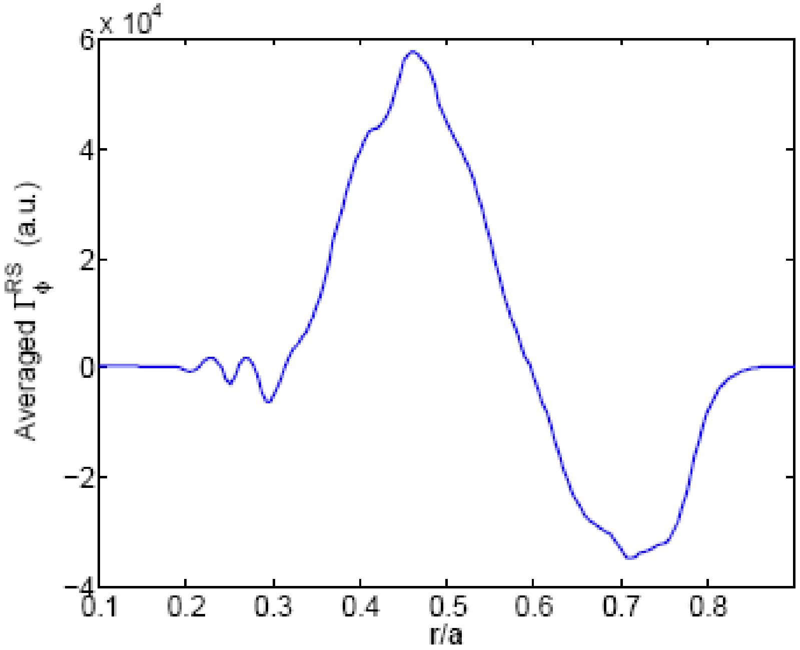

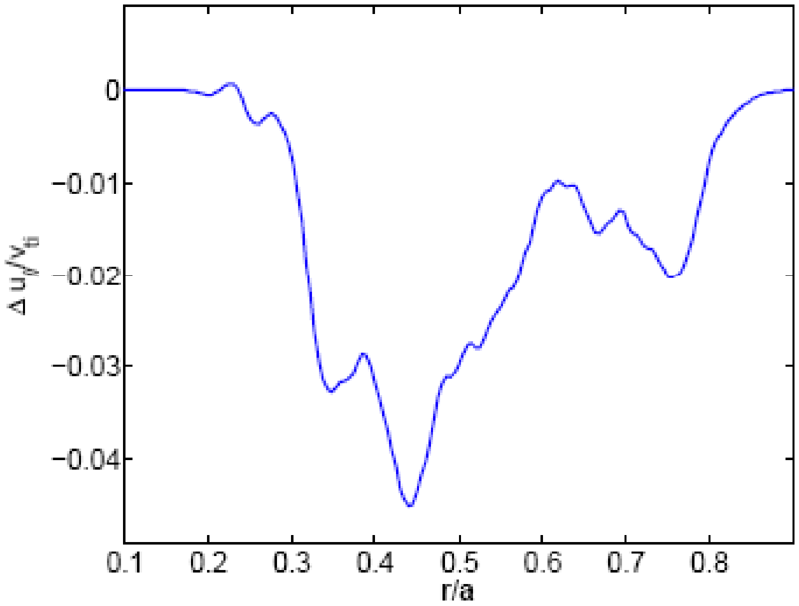

- Plasma initially rotation-free and momentum-source-free
- a net toroidal rotation produced in whole turbulence region
- in co-current direction, consistent with experimental trend
- $u_{\parallel} \sim 5\% \times v_{ti}$ at end of simulation
- Via momentum transfer from waves to particles

GTS elucidate origin of intrinsic rotation in tokamaks

ITG

case

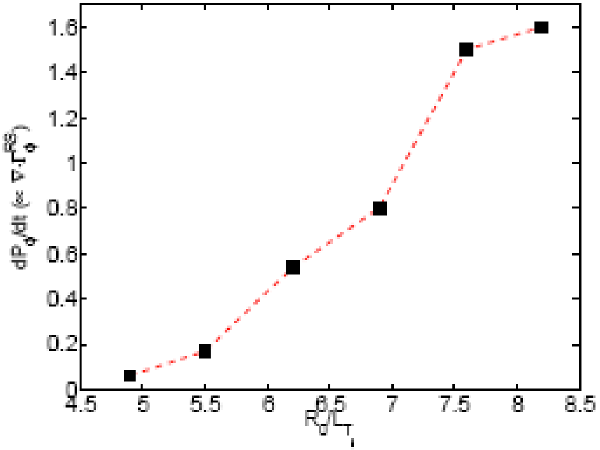

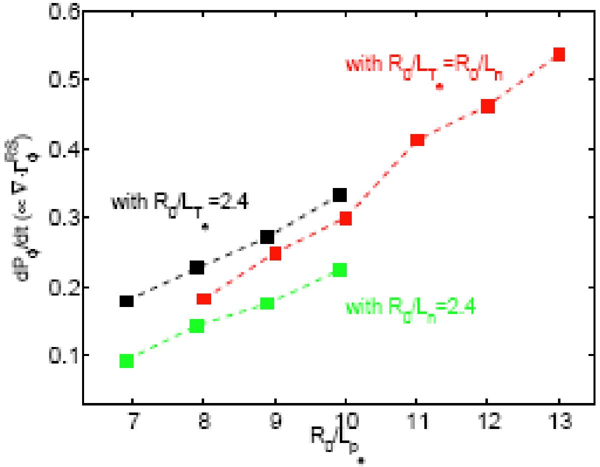

CTEM

case

- Turbulence driven intrinsic rotation scales linearly with pressure gradients
- Originated from dependences of fluctuation intensity and zonal flow shear on turbulence drives
- Reproduce empirical scalings obtained in experiments
  - $\Delta V_\phi \sim \Delta W_p / I_p$ in H-mode plasmas of multiple devices (Rice et al. '07)
  - $V_{\phi,\mathrm{central}} \sim \nabla P_{\mathrm{edge}}$ in C-MOD (Rice, APS-DPP'09)
  - intrinsic rotation increases with pressure gradient

    in JT-60 (Yoshida et al. '08); in LHD (Ida et al. '10); $\cdots$

"Flow pinch" phenomenon found in CTEM turbulence reproduces experiments

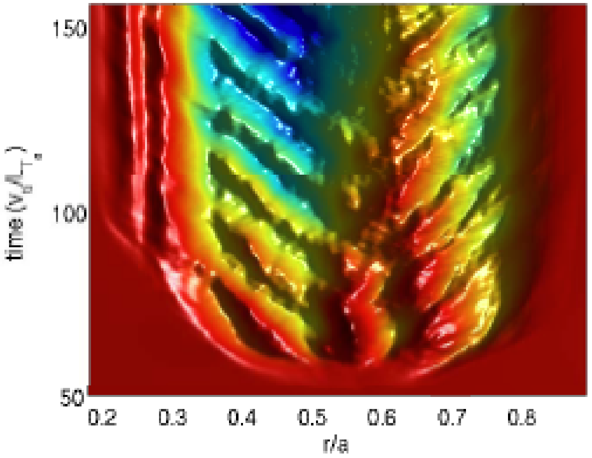

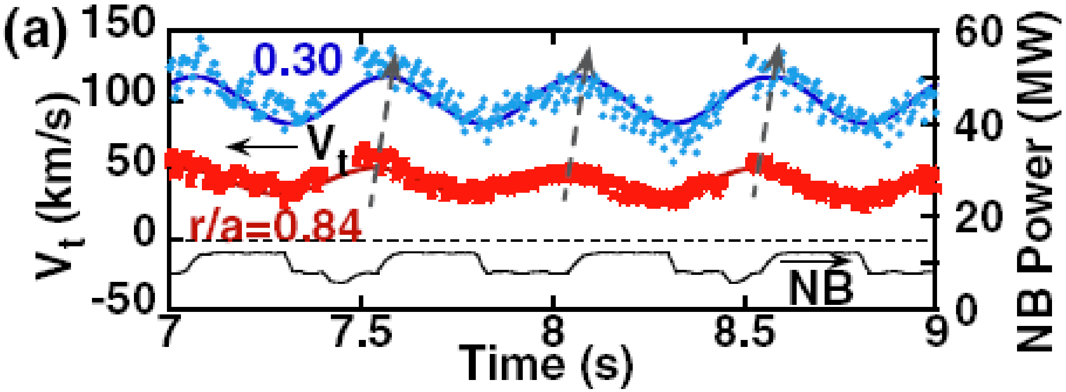

(From perturb. experiment, Yoshida et al. '08)

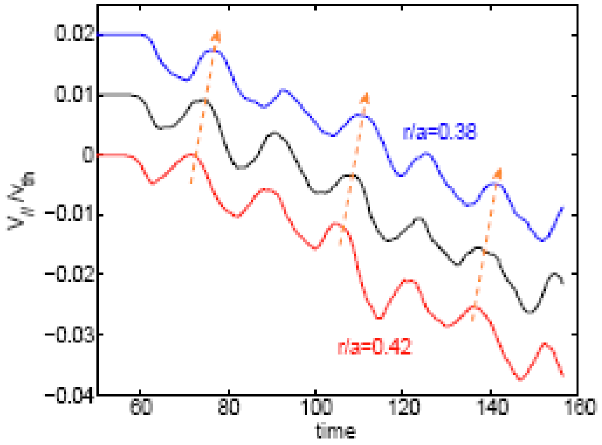

- Highly analogous to perturb. experiments
- $V_p \sim 7 \times 10^{-3} c_s, \quad f_p \sim 0.1 c_s/a$
- Flow perturbations generated locally in center

- Illuminate underlying dynamics governing the radial penetration of modulated flows in experiments

Science Case with XGC1: Edge pedestal is an urgent problem in tokamak research (100K-220K cores)

- Plasma near material wall must stay cold (~100eV)
- Plasma in the central core must be hot (>10 keV)
- Temperature-slope is limited by turbulence
  - T_i is too low in fusion core if in L-mode (<1980)
- ITER assumes H-mode pedestal
  - **Strong core-heating is necessary**
  - **Short propagation time (<< τ_{conf}) of the edge→core confinement properties**
  - **Stiff T_i profile**
- This physics must be understood

(instability from steep local profile not discussed here).

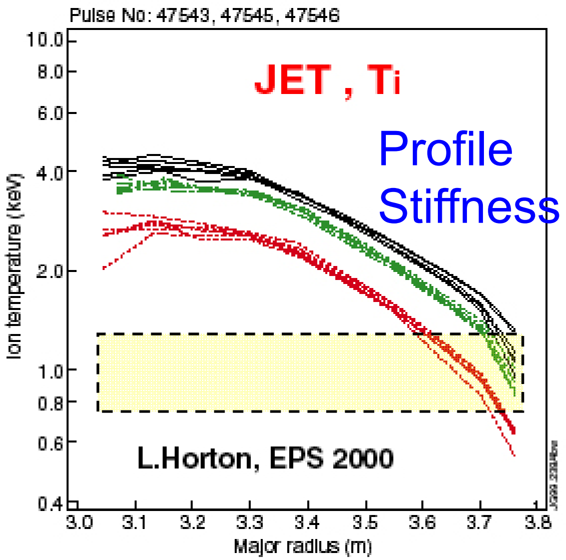

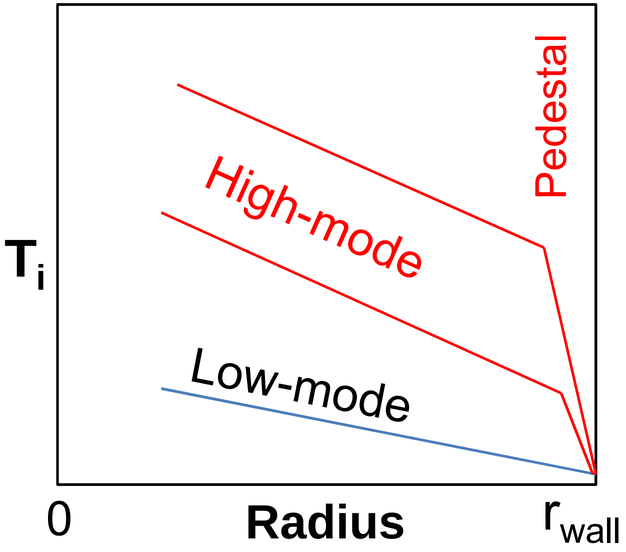

CORE TURBULENCE RAPIDLY SUPPRESSED AT LH TRANSITION

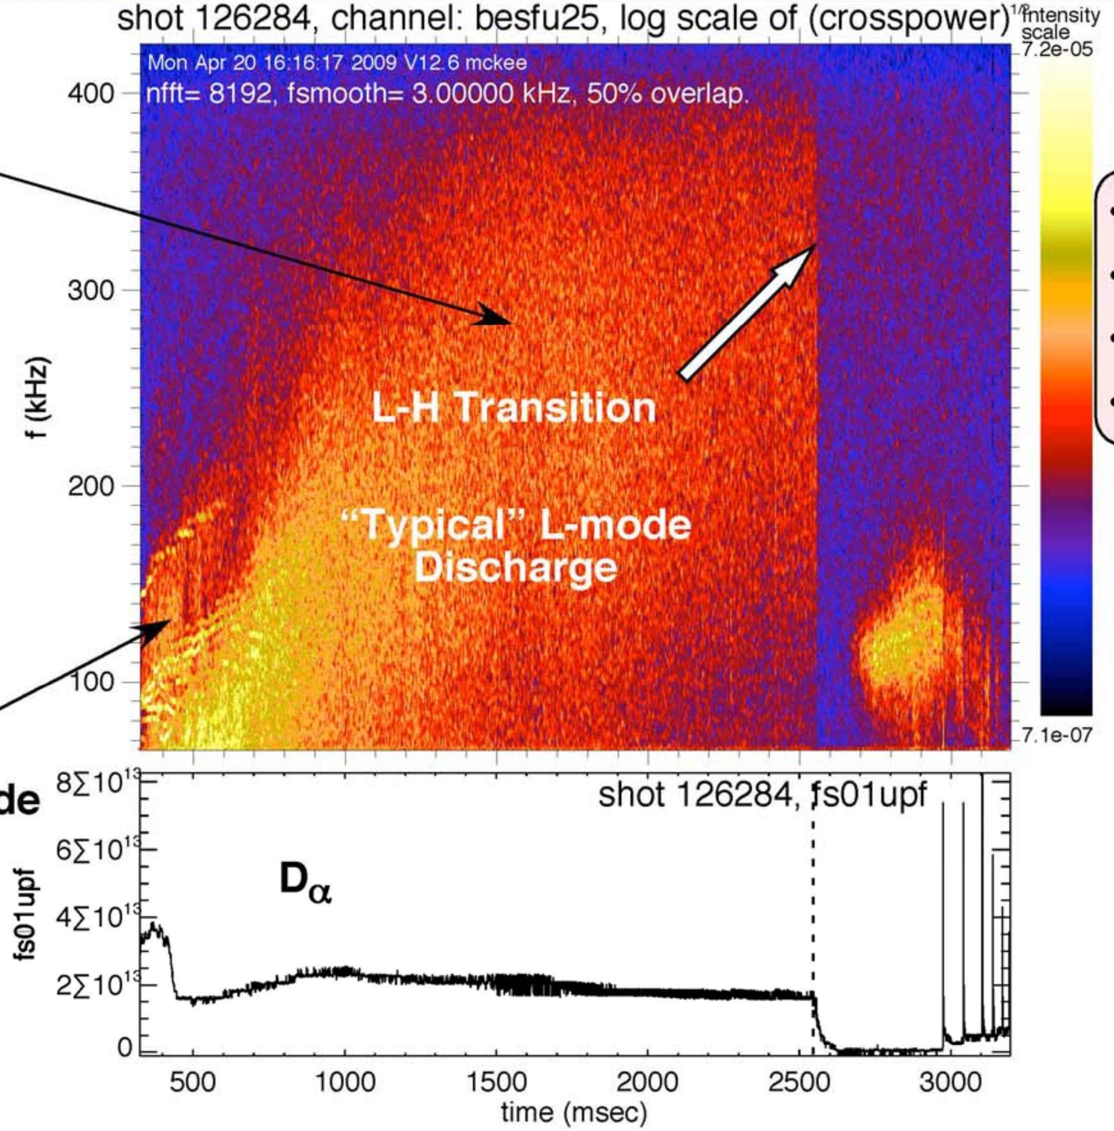

Broadband Turbulence

Suppression within several ms of LH transition

Alfven Eigenmode Activity

- **r/a=0.60**
- **Δr=1.2 cm**
- **ΔZ=1.5 cm**
- **f = 1 MHz**

Edge only simulation: Nonlocal ITG turbulence across the pedestal

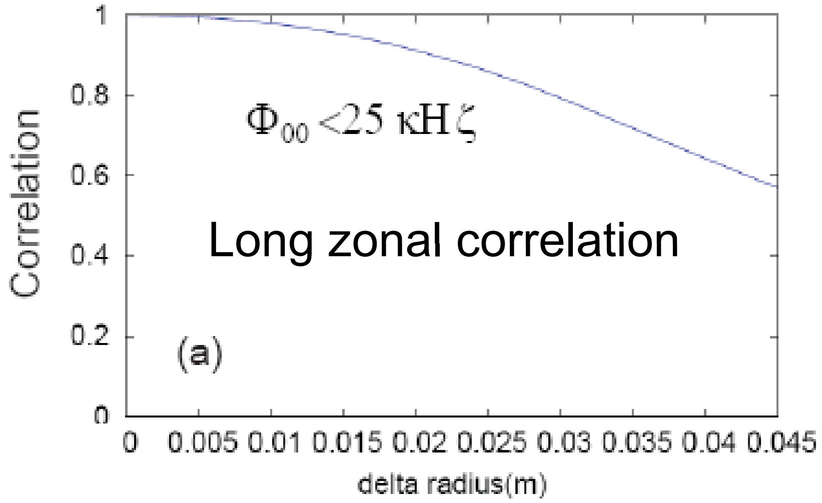

Long zonal correlation

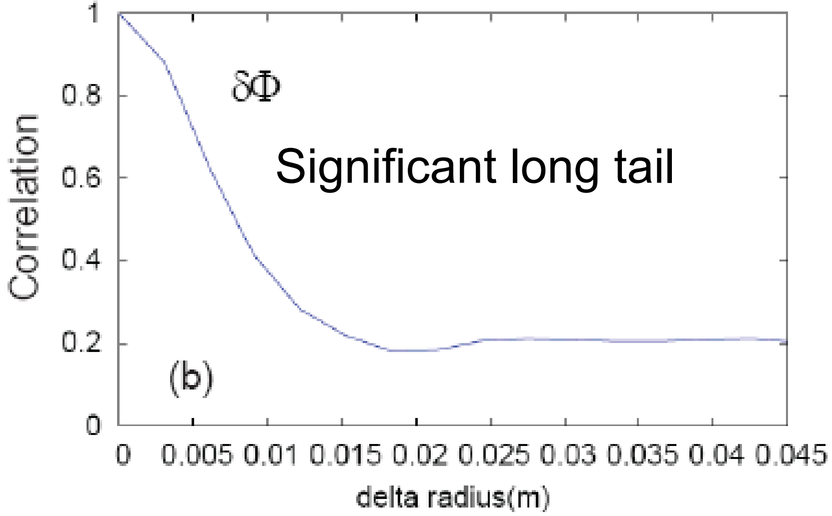

Significant long tail

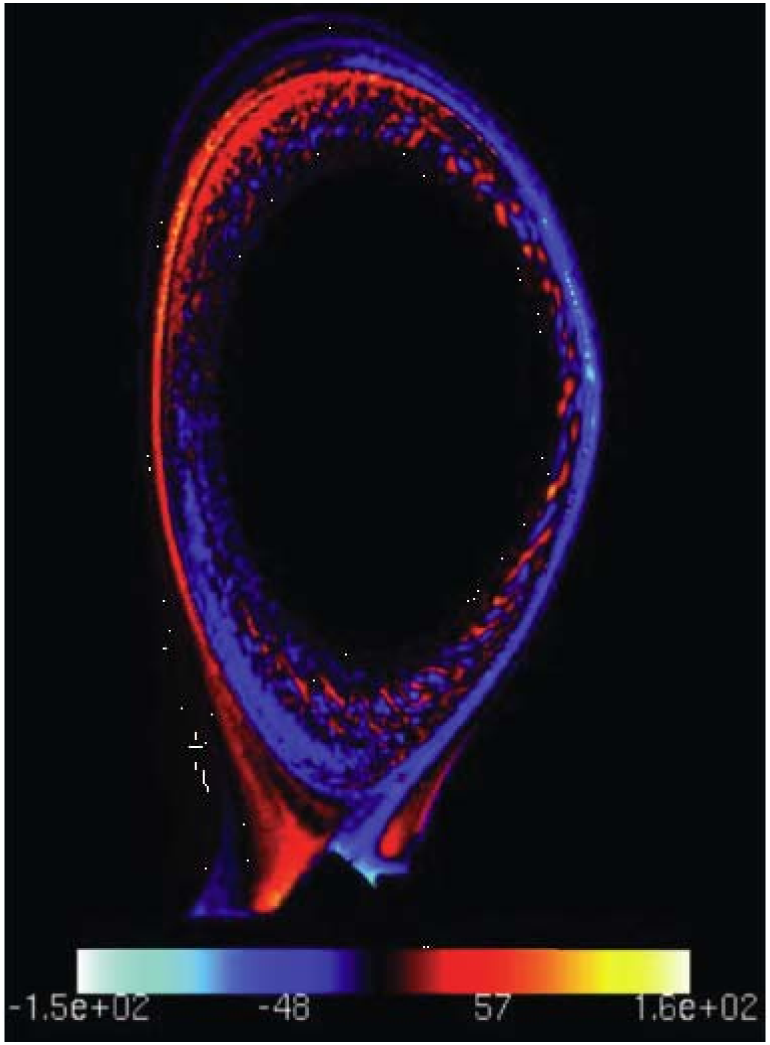

[Chang, et al, PoP 2009]

Edge ITG turbulence in XGC1 with and without the the mean field interaction

(Chang, Ku, et al, PoP 2009)

Full-function Multiscale simulation, with consistent mean field and turbulence, is important.

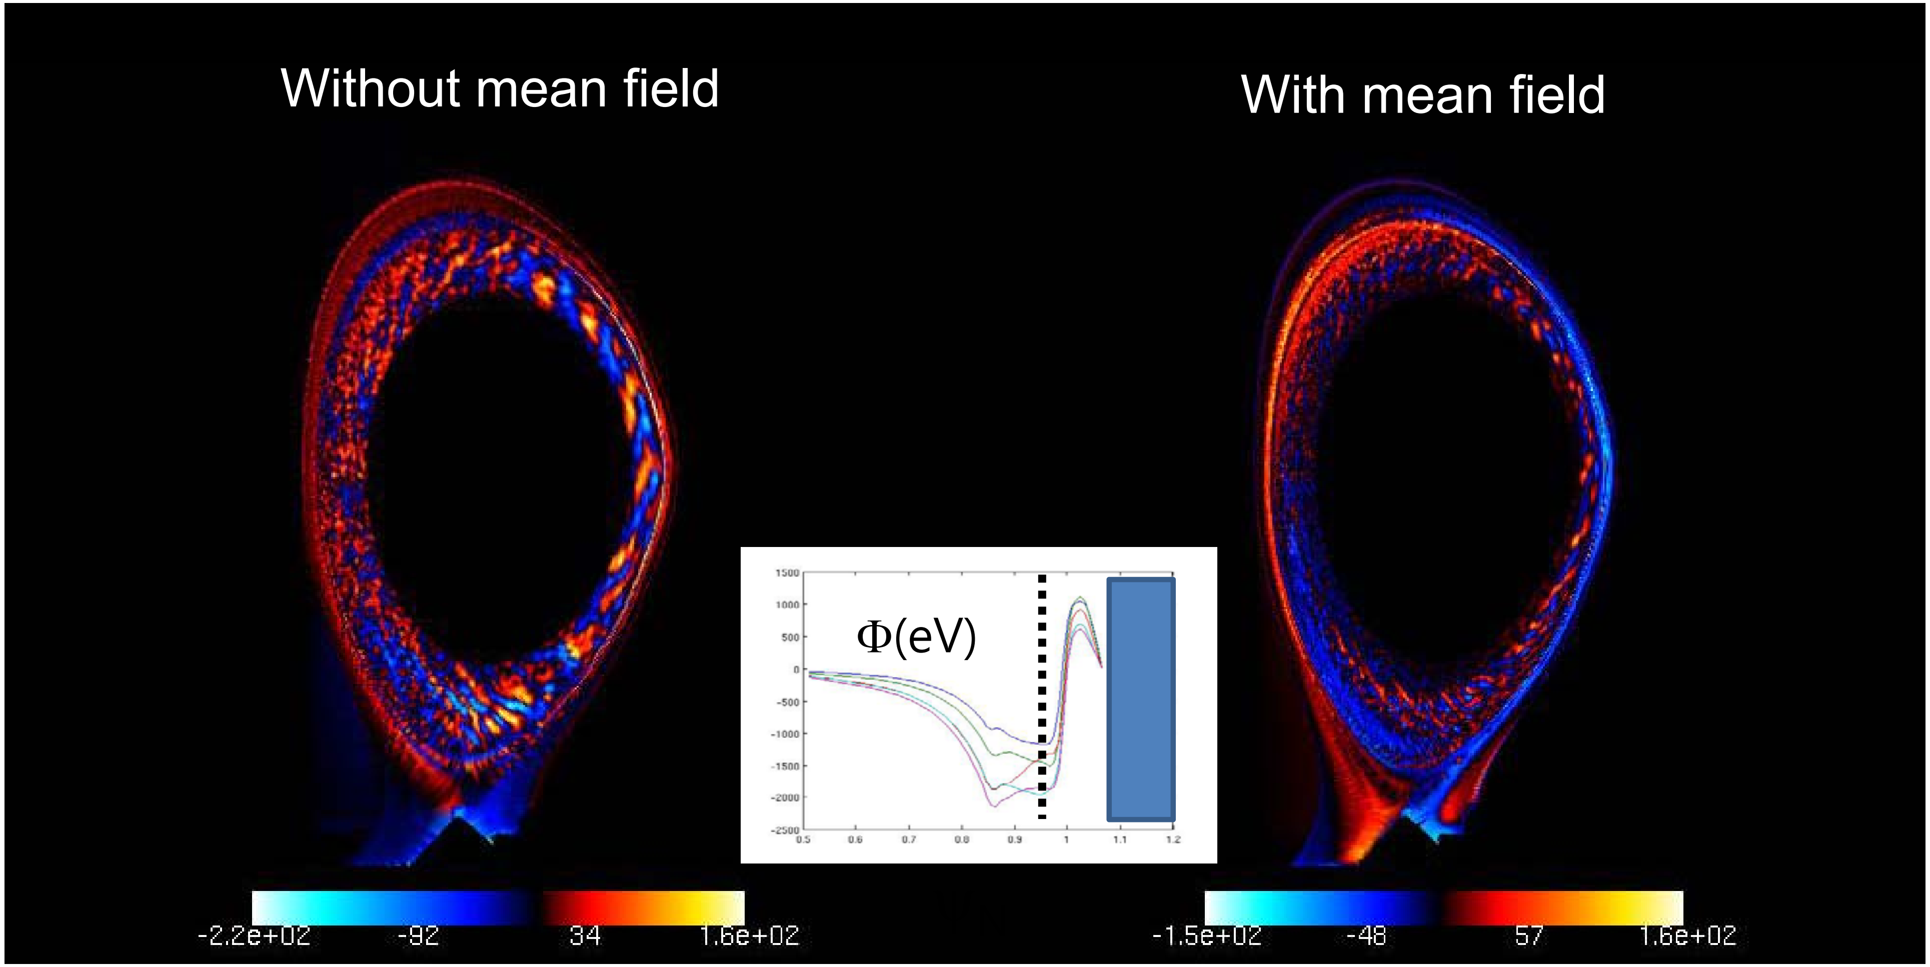

Sensitivity study to core-edge boundary location:

Edge turbulence solution is different when we remove the inner-radius boundary.

→ Core and edge simulations need to work together

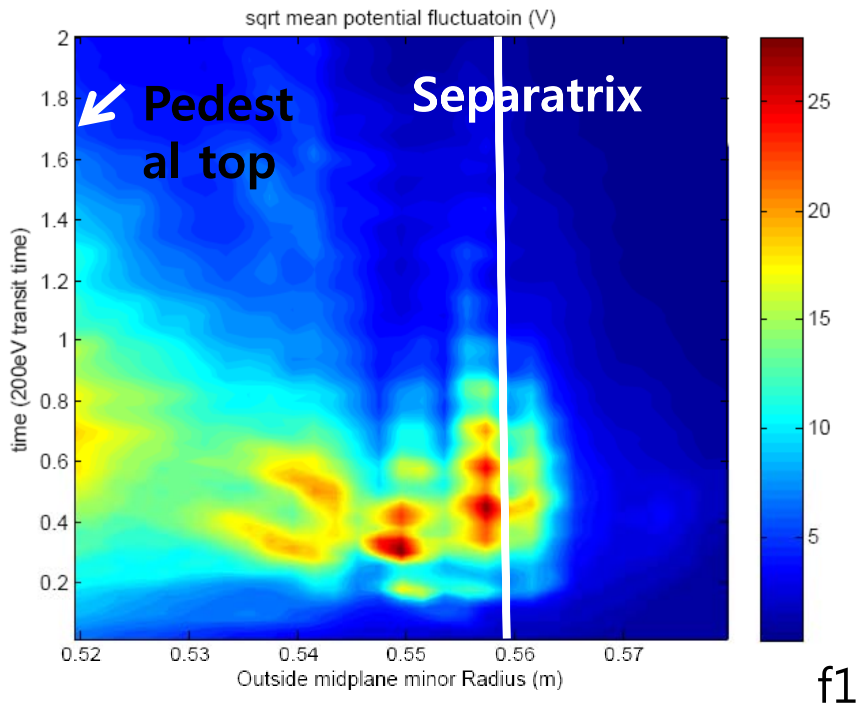

f119

With core-edge boundary

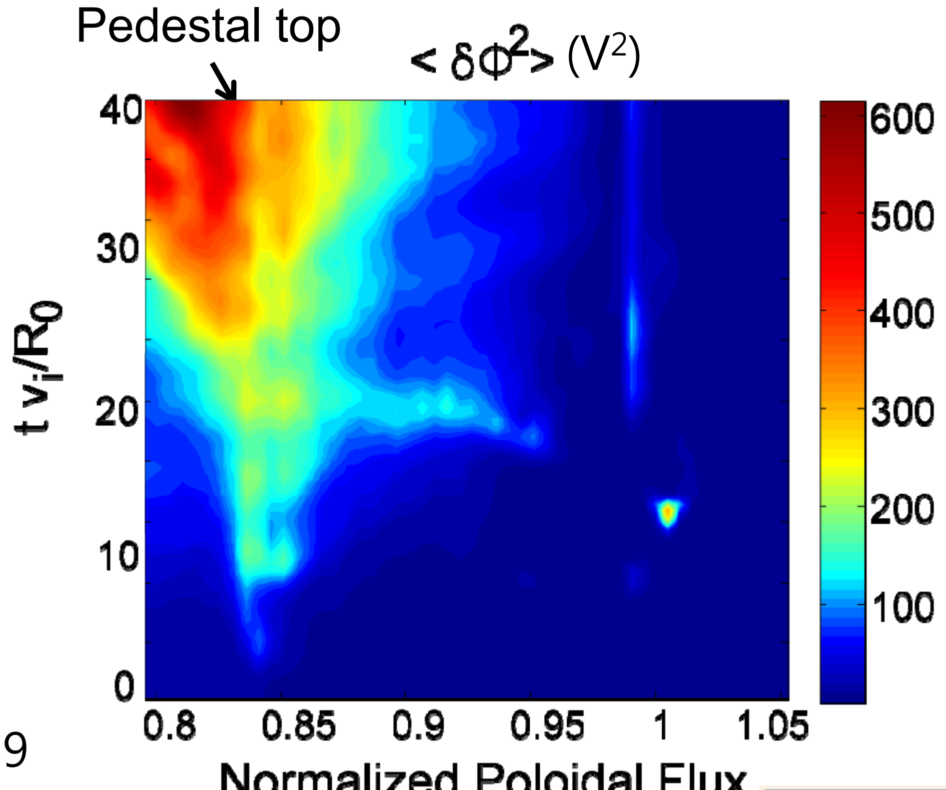

Without core-edge boundary

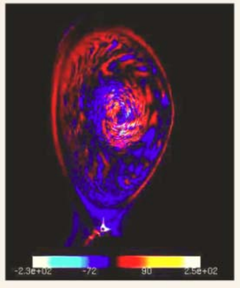

XGC1 with strong turbulence drive at pedestal top:

Inward-propagating turbulence controls core turbulence to self-organize with the **Outward heat bursts**

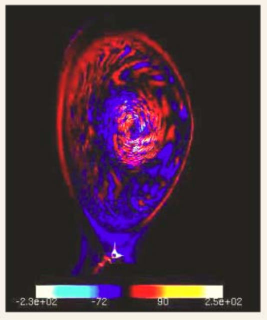

- **Role of the self-organizing ExB shearing is important**
- **Global turbulence and T_i profile settle down to SOC in several ms**

Heat flux

ExB Shearing rate

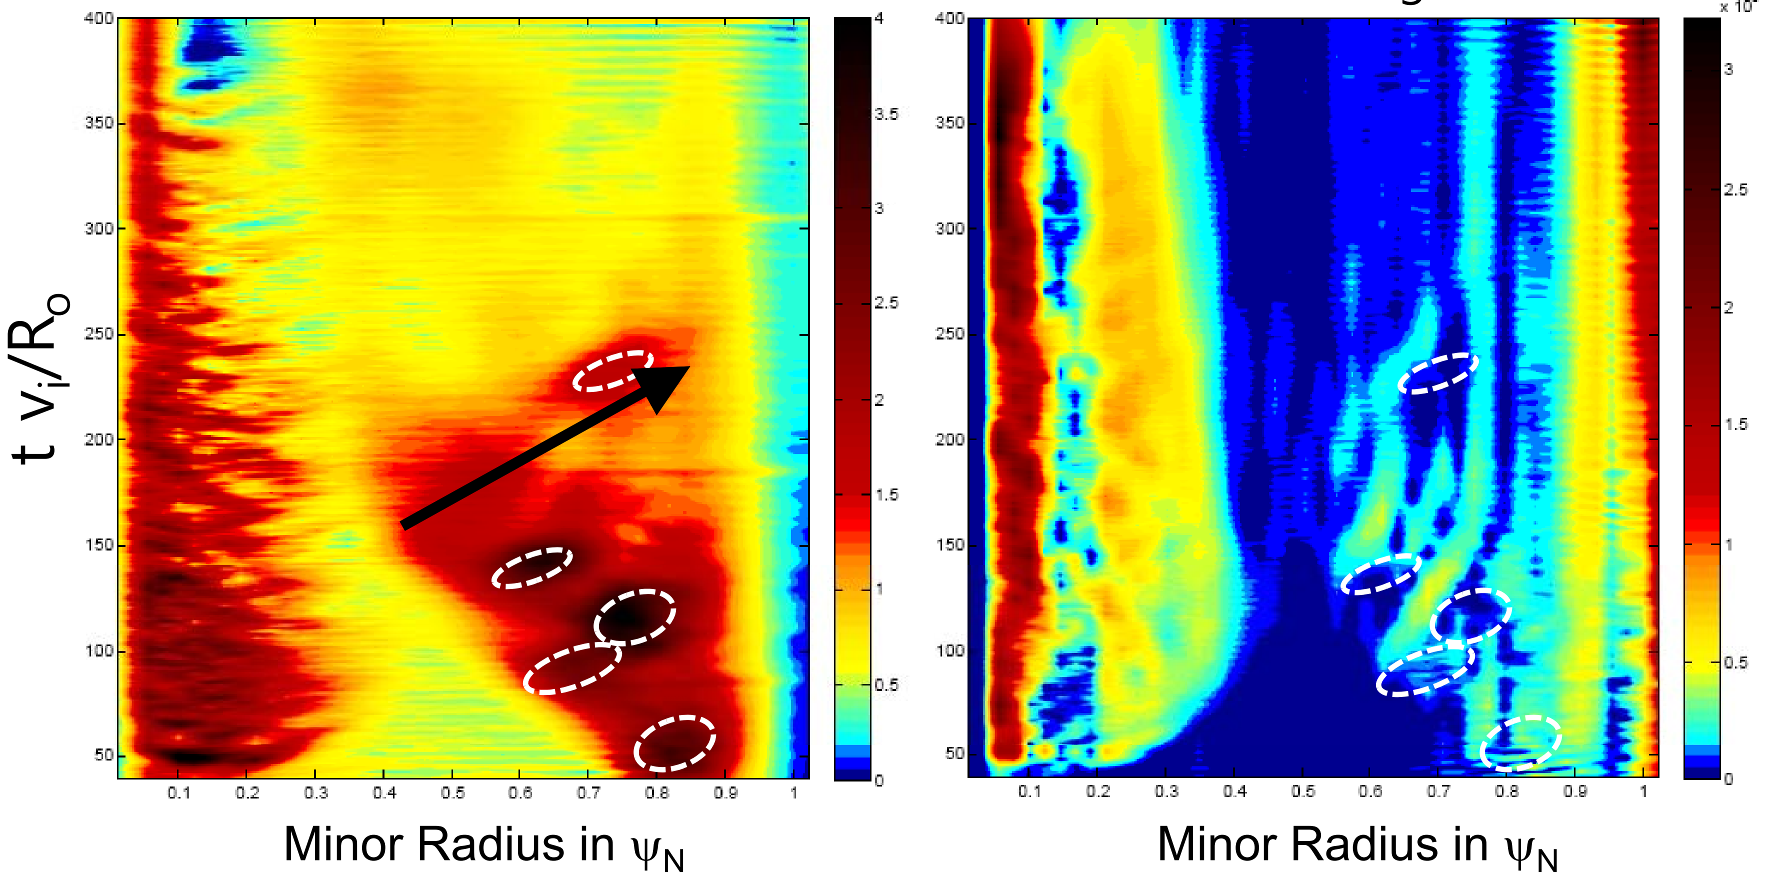

Particle scaling study of GTS and XGC1 on Jaguarpf (Cray XT5)

Number of particles moved 1 step in 1 second

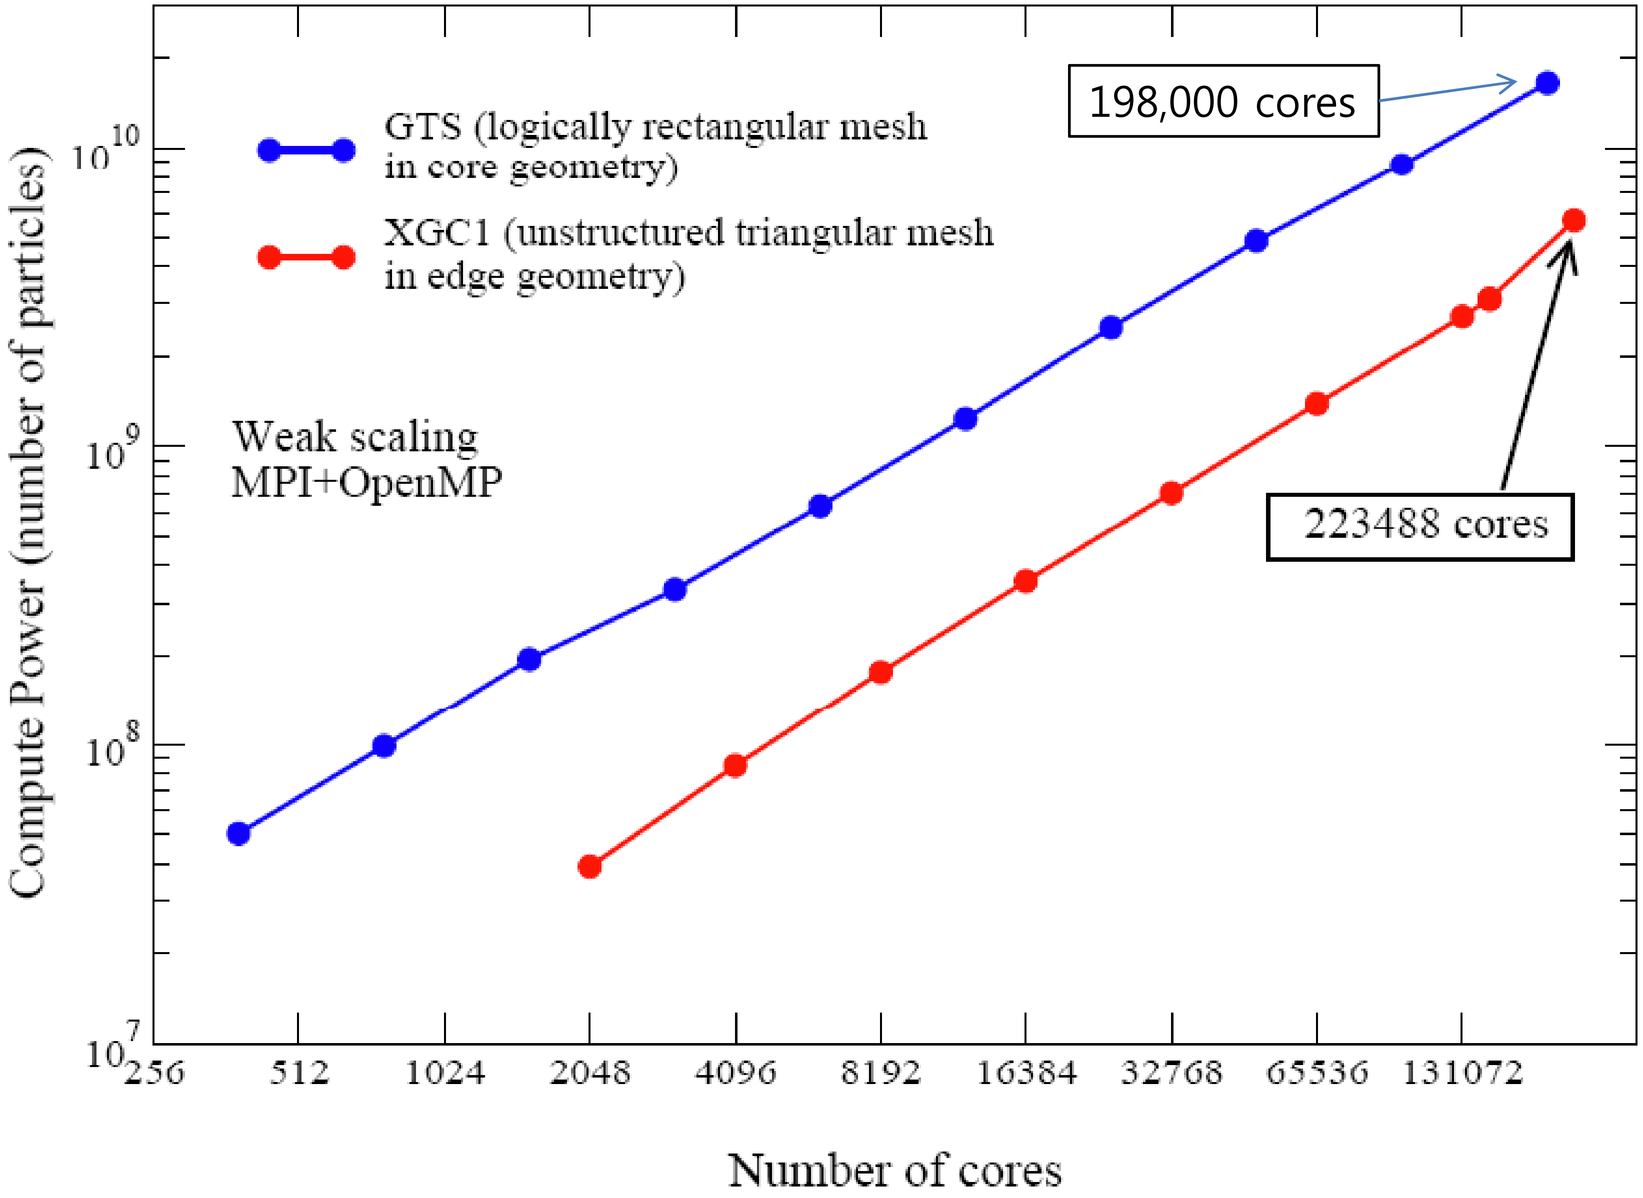

Weak scaling
MPI+OpenMP

198,000 cores

223488 cores

XGC1 Wallclock Time: MPI vs. OpenMP

300K particles/thread, 12 cores per node, 2010(C) expts. only

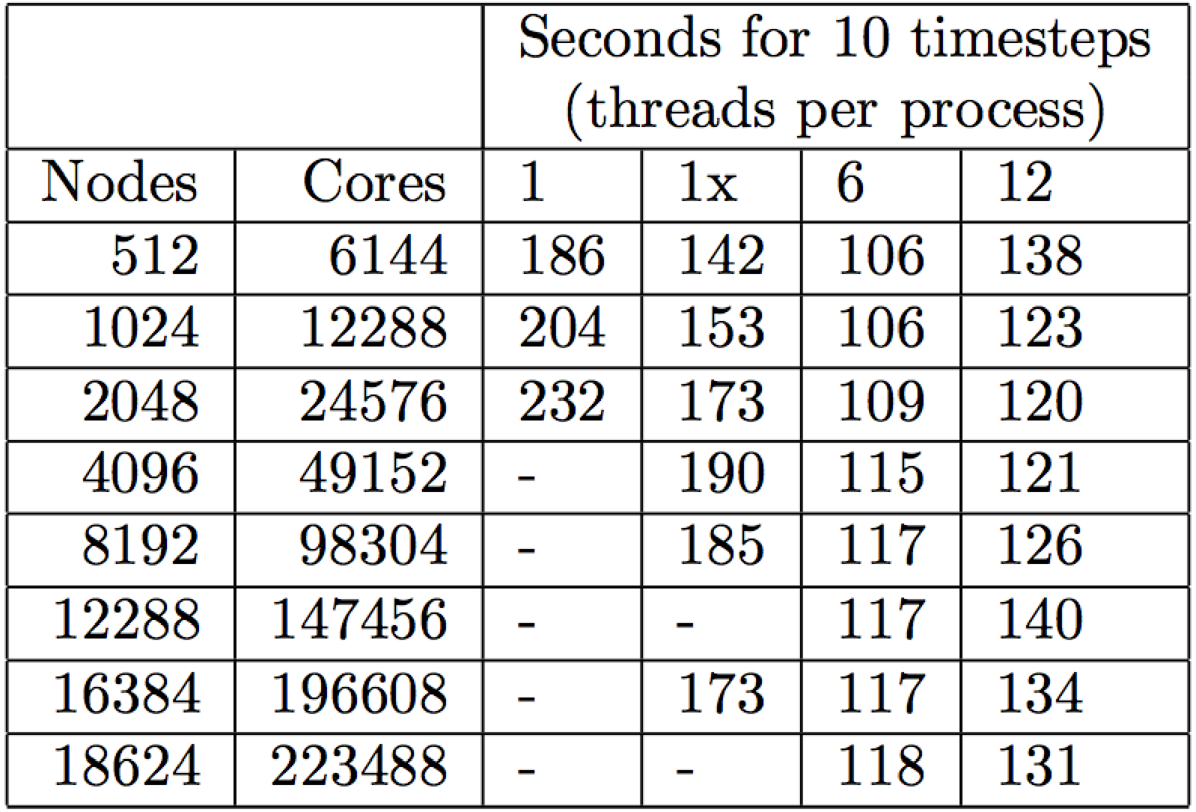

		Seconds for 10 timesteps (threads per process)			
Nodes	Cores	1	1x	6	12
512	6144	186	142	106	138
1024	12288	204	153	106	123
2048	24576	232	173	109	120
4096	49152	-	190	115	121
8192	98304	-	185	117	126
12288	147456	-	-	117	140
16384	196608	-	173	117	134
18624	223488	-	-	118	131

- “1x”: using only 8 cores pre node, so problem size only 0.67 that of other data.
- MPI-only not scaling well, and never competitive when using 6144 or more cores.
- **6-way OpenMP best performer in these experiments.**

Computing and Storage Resources

	GYRO		GTS		XGC1*	
Facilities	NERSC/OLCF		NERSC/OLCF		NERSC/OLCF	
Architectures	XT5,Power,Cluster		XT5		XT5	
Years	Present	In 5 yrs	Present	In 5 yrs	Present	In 5 yrs
Hrs used/year	30M	50M	24M	50M	65M	500M
NERSC'09 used	1.2Mhrs		~2Mhrs		~8M hrs	
#Cores per run	512	512	8-98K	32-130K	10-223K	1M
Wall clock/run	12	24	72 Hrs	72 Hrs	20-100hrs	20-100hrs
Memory/run	512GB	1.024TB	16-100T	32-160TB	40 TB	100 TB
Min Memory/core	1GB	2GB	1GB	1GB	0.3GB	0.1GB
Read/Write data			2.5TB	8TB	5TB	25TB
Checkpoint size	4GB	8GB	1-8GB	1-10 GB	1TB	5TB
Data in/out nersc			5GB/run	10GB/run	10GB/day	50GB/day
On-line storage			4TB/10K	8TB/10K	4TB/3K	5TB/3K
Off-line storage			25GB	100GB	1TB/30	10TB/100

***Unstructured mesh**

Conclusion

- Variety of gyrokinetic fusion codes are used for capability computing at NERSC
  - from 512 cores (continuum: GYRO) to maximal number of cores (particle: GTC→GTS, XGC1)
- Continuum and particle approaches complementary to each other, in numerical technique, physics, and Cluster/Cloud/HPC usage at NERSC.
- Continuum code GYRO will be re-optimized for multi-cores for higher efficiency.
- Some of the codes which require extreme computing (**XGC1**, GTS, GTC) are aggressively moving into "localized" computing and GPU, for higher fidelity simulation with more complete physics on more number of cores.

GTS global simulation uncovers momentum generation process by turbulence

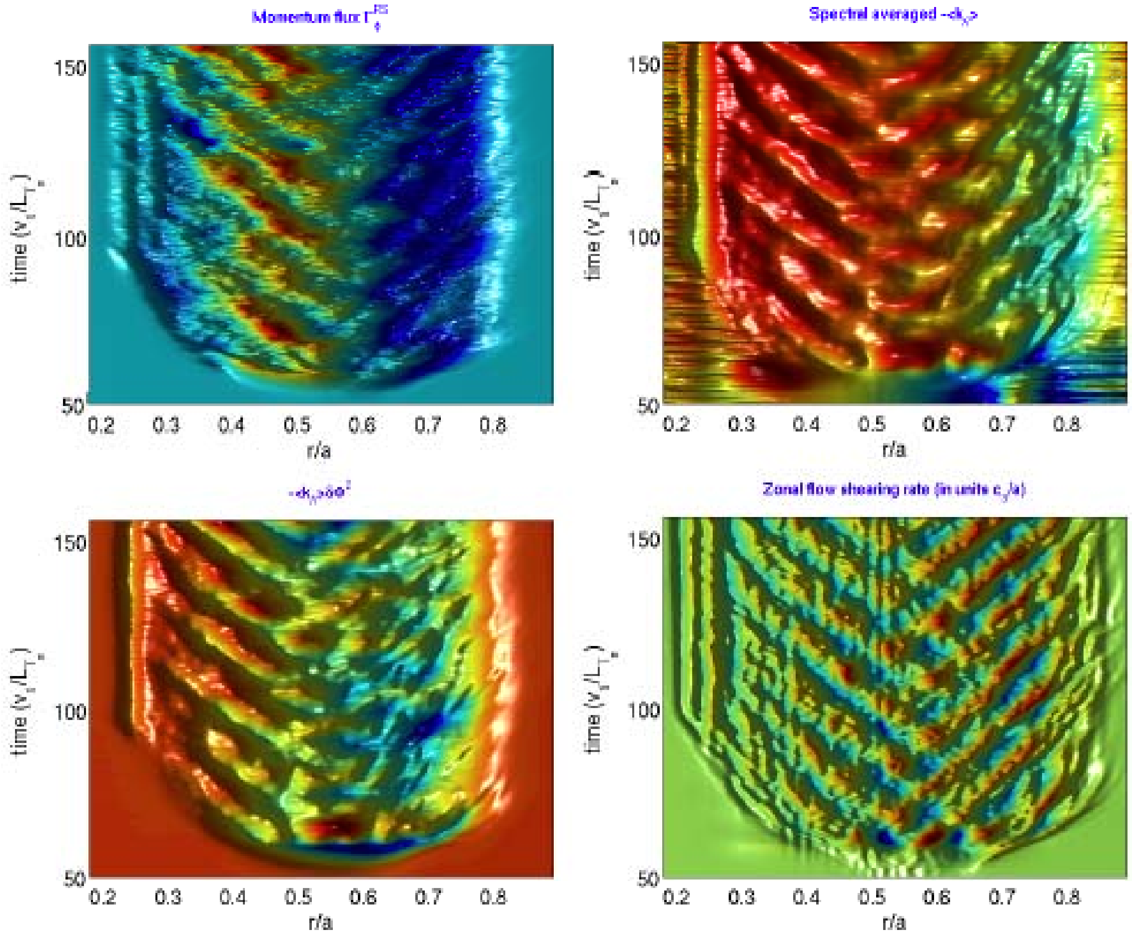

- typical DIII-D parameters
- real DIII-D geometry
- CTEM turbulence
- initially rotation free
- by GTS code

(Wang et al., PRL'09 & PoP'10)

- Residual stress driven by fluctuation intensity & intensity gradients
- Self-generated low frequency zonal flow shear plays key role
- Acts as "internal torque", transferring momentum from wave to particles