Multiscale Gyrokinetic Analysis in the Tokamak Pedestal

By **E.A. Belli**[1], J. Candy[1], I. Sfiligoi[2]

[1] **General Atomics**
[2] UCSD San Diego Supercomputing Center

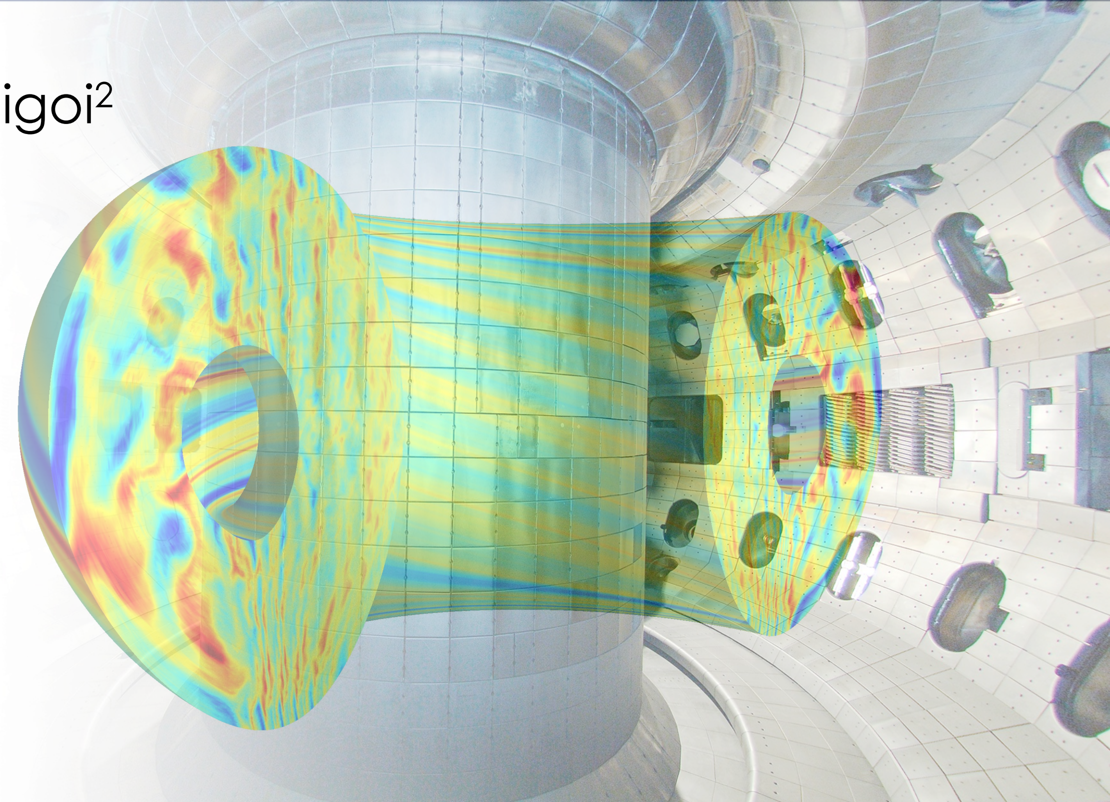

Presented at the Fusion HPC Workshop

December 2022

Supported by U.S. DOE under DE-FG02-95ER54309 and DE-FC02-06ER54873

Understanding transport in the H-mode pedestal can help to develop operating regimes for optimal confinement and fusion performance.

- The most promising operating scenario for achieving fusion in tokamaks is **H-mode confinement regime**.
- H-mode is characterized by the formation of edge transport
- Region of reduced transport leads to steeper gradients in density and temperature → "**pedestal**" structure at the edge of the plasma
- Pedestal plays a key role in **determining global energy confinement**.

Turbulent transport in pedestal is less well understood than in the core.

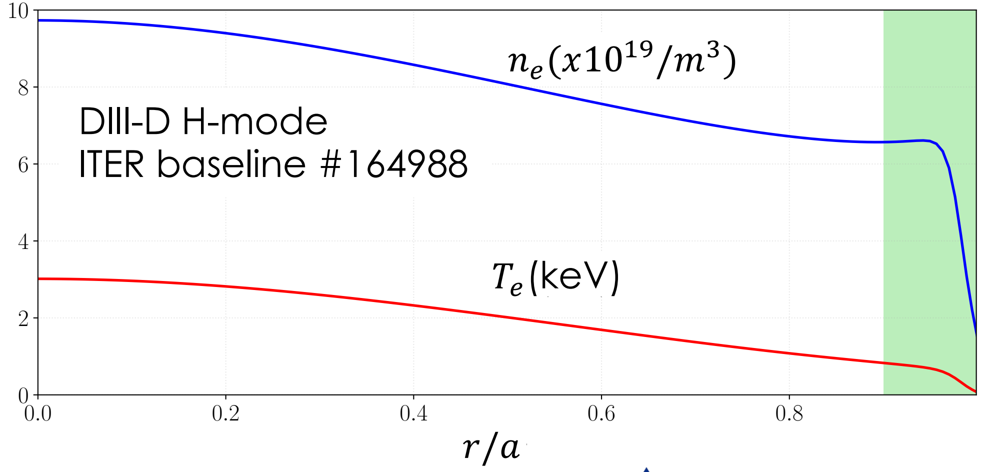

GENERAL ATOMICS

While ion-scale turbulence in the core has been modeled extensively, multiscale pedestal simulations are far more challenging.

- **Strongly shaped edge geometry**
  - Need 2-6X increase in resolution in parallel (field-line) direction θ

DIII-D H-mode ITER Baseline #164988

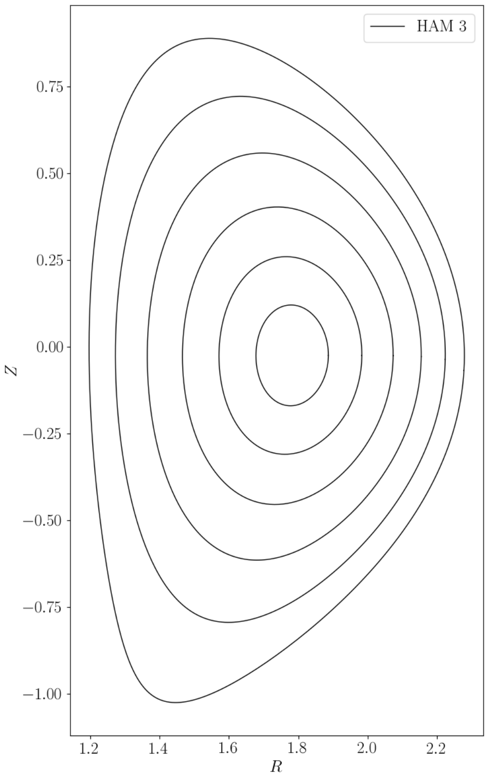

GENERAL ATOMICS

While ion-scale turbulence in the core has been modeled extensively, multiscale pedestal simulations are far more challenging.

- Strongly shaped edge geometry
  - Need 2-6X increase in resolution in parallel (field-line) direction θ
- **Weaker pitch of confining magnetic field (large safety factor q)**
  - Need large radial resolution

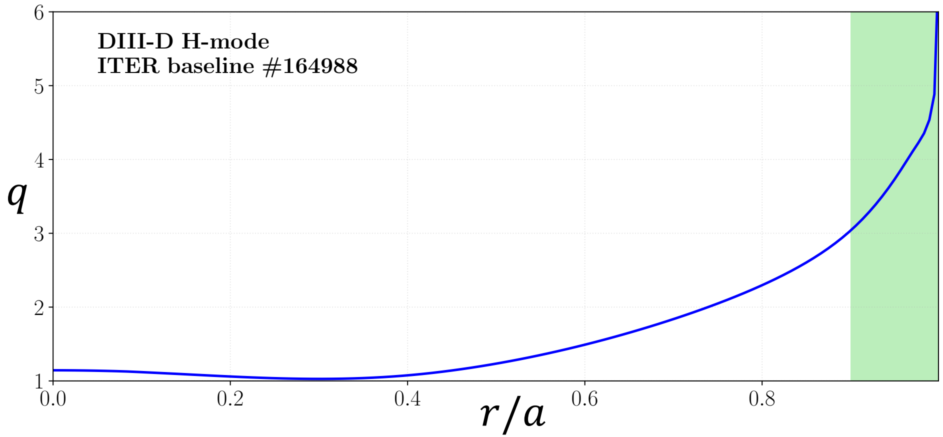

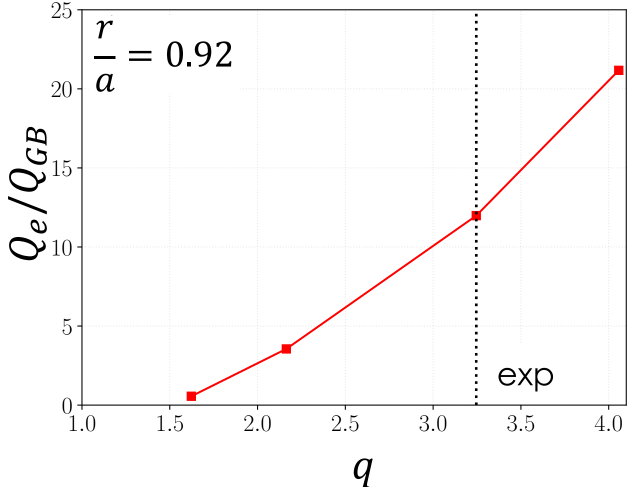

While ion-scale turbulence in the core has been modeled extensively, multiscale pedestal simulations are far more challenging.

- Strongly shaped edge geometry
  - Need 2-6X increase in resolution in parallel (field-line) direction θ
- Weaker pitch of confining magnetic field (large safety factor q)
  - Need large radial resolution
- **Large collisionality**
  - Need advanced collision models

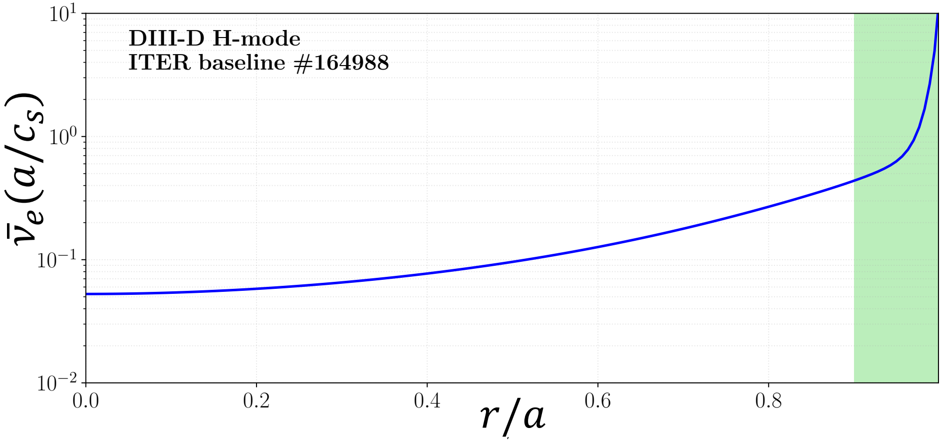

Standard e pitch angle scattering
+ energy diffusion
+ conservation (n,v,E)
+ inter-species (**multiscale in velocity space**)

GENERAL ATOMICS

While ion-scale turbulence in the core has been modeled extensively, multiscale pedestal simulations are far more challenging.

- Strongly shaped edge geometry
  - Need 2-6X increase in resolution in parallel (field-line) direction θ
- Weaker pitch of confining magnetic field (large safety factor q)
  - Need large radial resolution
- Large collisionality
  - Need advanced collision models
- **Large gradients drive multiple instab. across broad range of spatial scales**
  - **Ion-scales ($k_\theta \rho_i \lesssim 1$)**: ES modes (ITG, TEM), EM modes (MTM)
  - **Electron-scales ($k_\theta \rho_e \sim 1$)**: ETG

Scales Simulated in Different Simulation Types

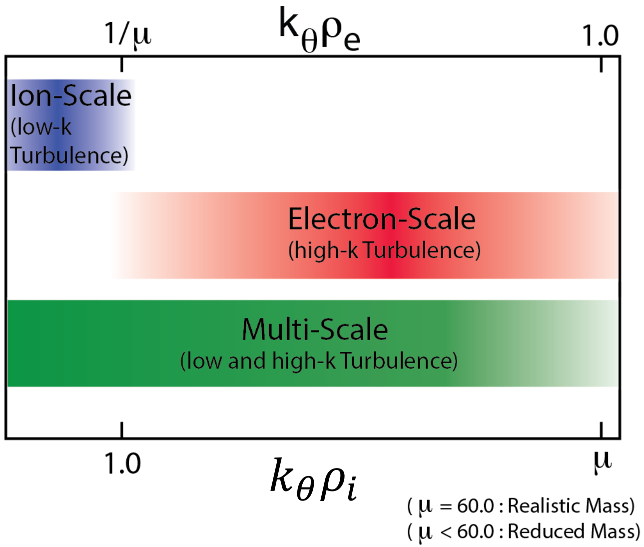

Electron heat transport will play a dominant role in reactors

→ Multiscale resolution needed

GENERAL ATOMICS

While ion-scale turbulence in the core has been modeled extensively, multiscale pedestal simulations are far more challenging.

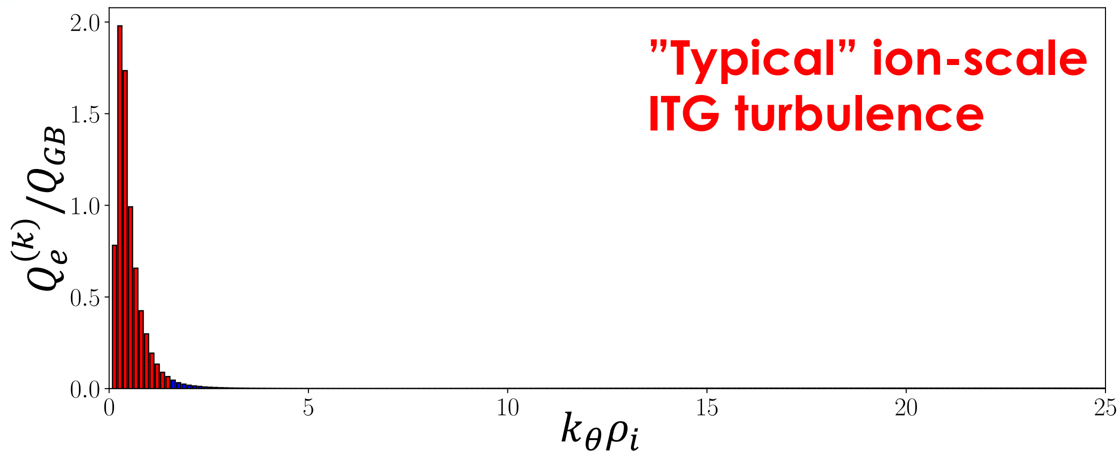

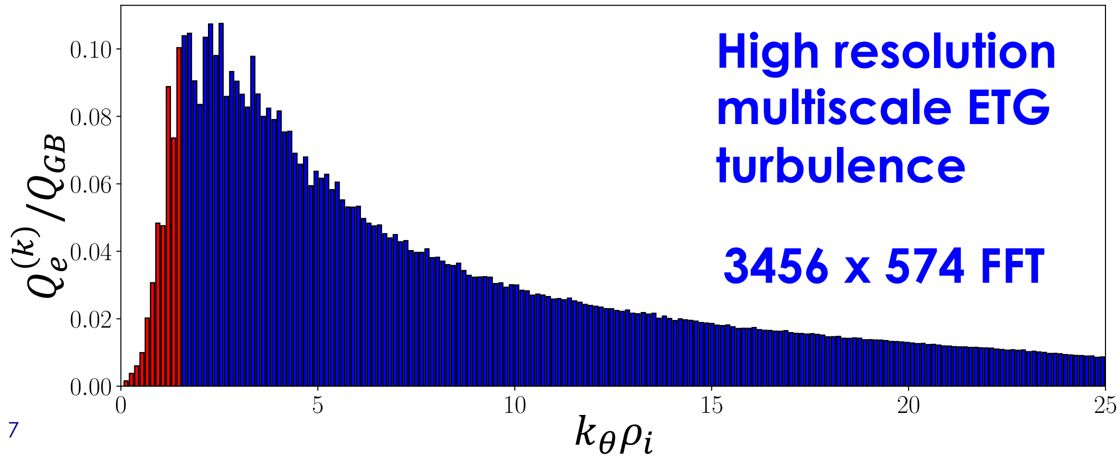

Requires **leadership-scale computing resources** and **highly optimized solvers.**

CGYRO: A multiscale-optimized gyrokinetic turbulence solver

- Solves the 5D (3 spatial+2 velocity) δf gyrokinetic-Poisson-Ampere equations using **Eulerian approach**
- Motivations are **accurate collisions** in H-mode pedestal and and to provide efficient nonlinear, electromagnetic **multiscale simulations**.
  - Complex nonlinear cross-scaling coupling requires extremely fine mesh in real space

    → **Specialized numerical schemes** are needed to prevent severe bottlenecks related to:
    - gyroaveraging
    - Maxwell field solve
    - ExB nonlinearity

GENERAL ATOMICS

CGYRO implements highly efficient spectral/pseudospectral numerical schemes optimized for multiscale simulations.

x	Radial	spectral
y	Binormal	spectral
θ	Poloidal	Finite diff

$$H_a(x, y, \theta, \xi, v)$$

ξ	Pitch angle	pseudospectral
v	Velocity	psuedospectral

- **Fully spectral in (x, α)** provides maximal multiscale efficiency
  - Ensures collision operator is algebraic in space
  - Allows for most efficient evaluation of gyroaverages
  - Evaluation of nonlinear term on GPUs (cuFFT) ensures maximum performance and scalability

- **Pseudospectral in (ξ, v)** provides optimal accuracy of collisions

Unlike most gyrokinetic codes, CGYRO uses velocity-space coordinates optimized for the collisional problem.

GYRO & GS2 use (λ,ε) coordinates

$$\lambda = \frac{v_\perp^2}{v^2 B} \qquad \varepsilon = \frac{m_a v^2}{2T_a}$$

Advantage:

No need for derivatives across trapped/passing boundary since θ discretization is aligned with particle orbits

Disadvantage:

Difficult to evaluate collision operator due to irregular grid in (ξ,θ) $\mathcal{L} = \frac{1}{2}\frac{\partial}{\partial \xi}(1-\xi^2)\frac{\partial}{\partial \xi}$

NEO has instead had great success with (ξ,v) coordinates, implementing spectrally-accurate collision operators.

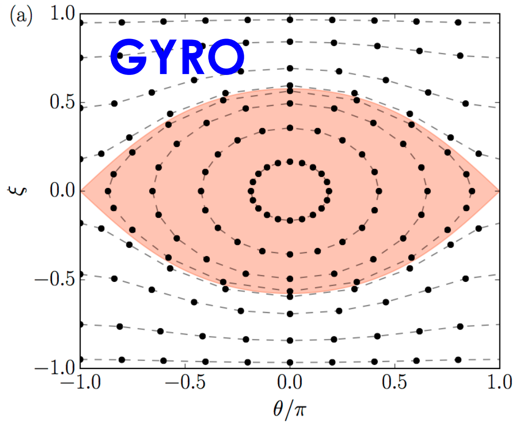

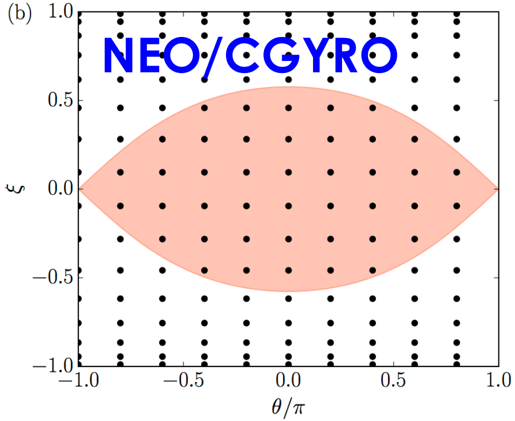

GENERAL ATOMICS

CGYRO has the first pseudospectral implementation of the collision operator in a gyrokinetic code.

- **Legendre polynomials in ξ**

$$\mathcal{L} = \frac{1}{2}\frac{\partial}{\partial\xi}(1-\xi^2)\frac{\partial}{\partial\xi}$$

- **Nonstandard orthogonal polynomials in v**

$$\xi = v_{\parallel}/v$$

$$u_a = v/\sqrt{2v_{ta}}$$

  - Accurate for energy integration and differentiation
  - Appropriate for semi-infinite energy domain

$$\int_0^b du\, e^{-u^2} Q_k(u) Q_l(u) = \gamma_k \delta_{kl}$$

$u \in [0, b]$

b→0: shifted monic Legendre

b→∞: half-range Hermite

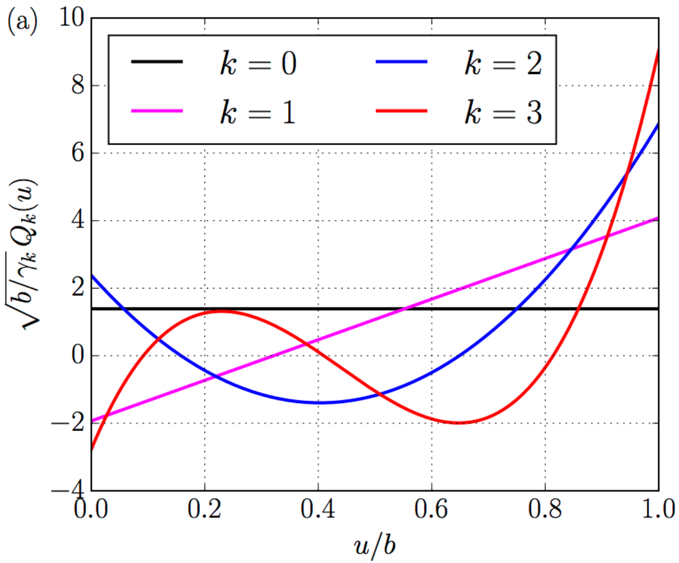

GENERAL ATOMICS

Spectral convergence in (ξ, v) velocity space is observed.

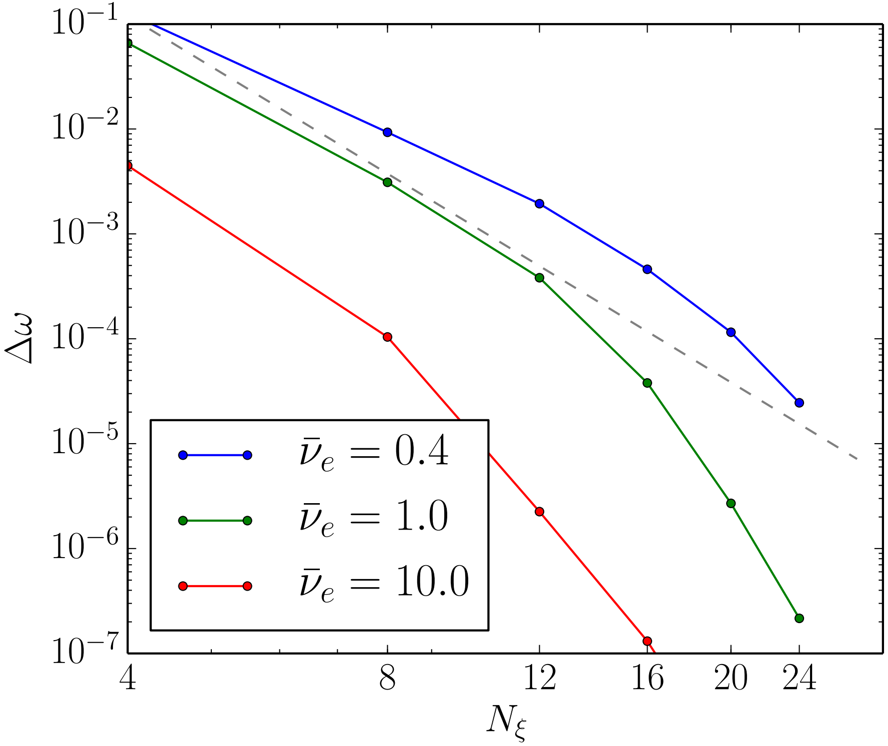

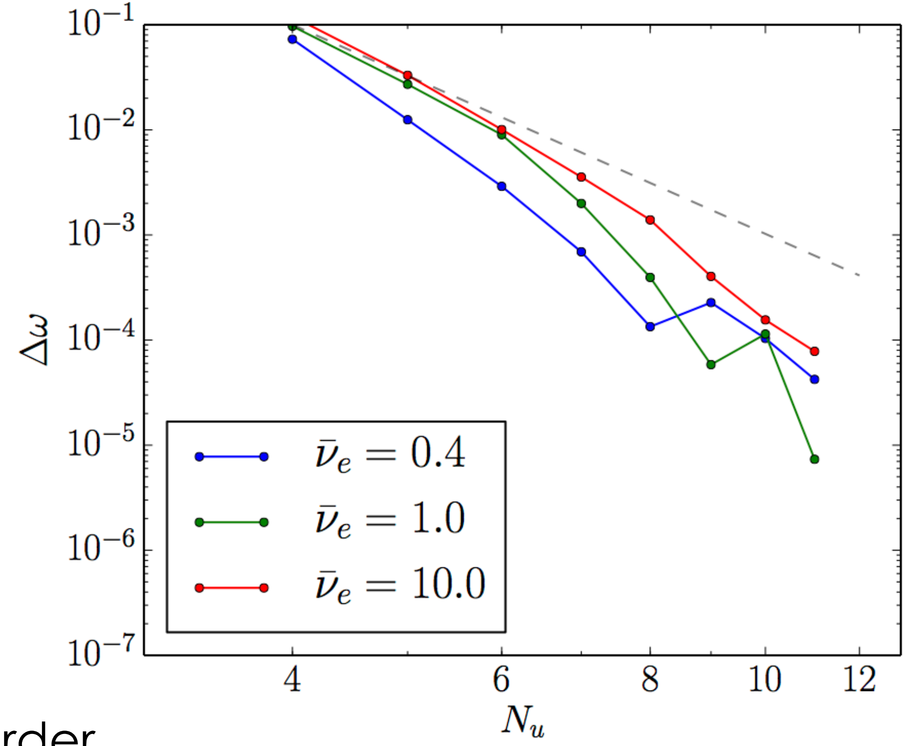

--- 4^{th} order

GENERAL ATOMICS

The GKE exhibits highly-collisional behavior at the lowest energies, transitioning to collisionless behavior at high energies.

|H_e| at various energy meshpoints

$\bar{\nu}_e = 1.0$

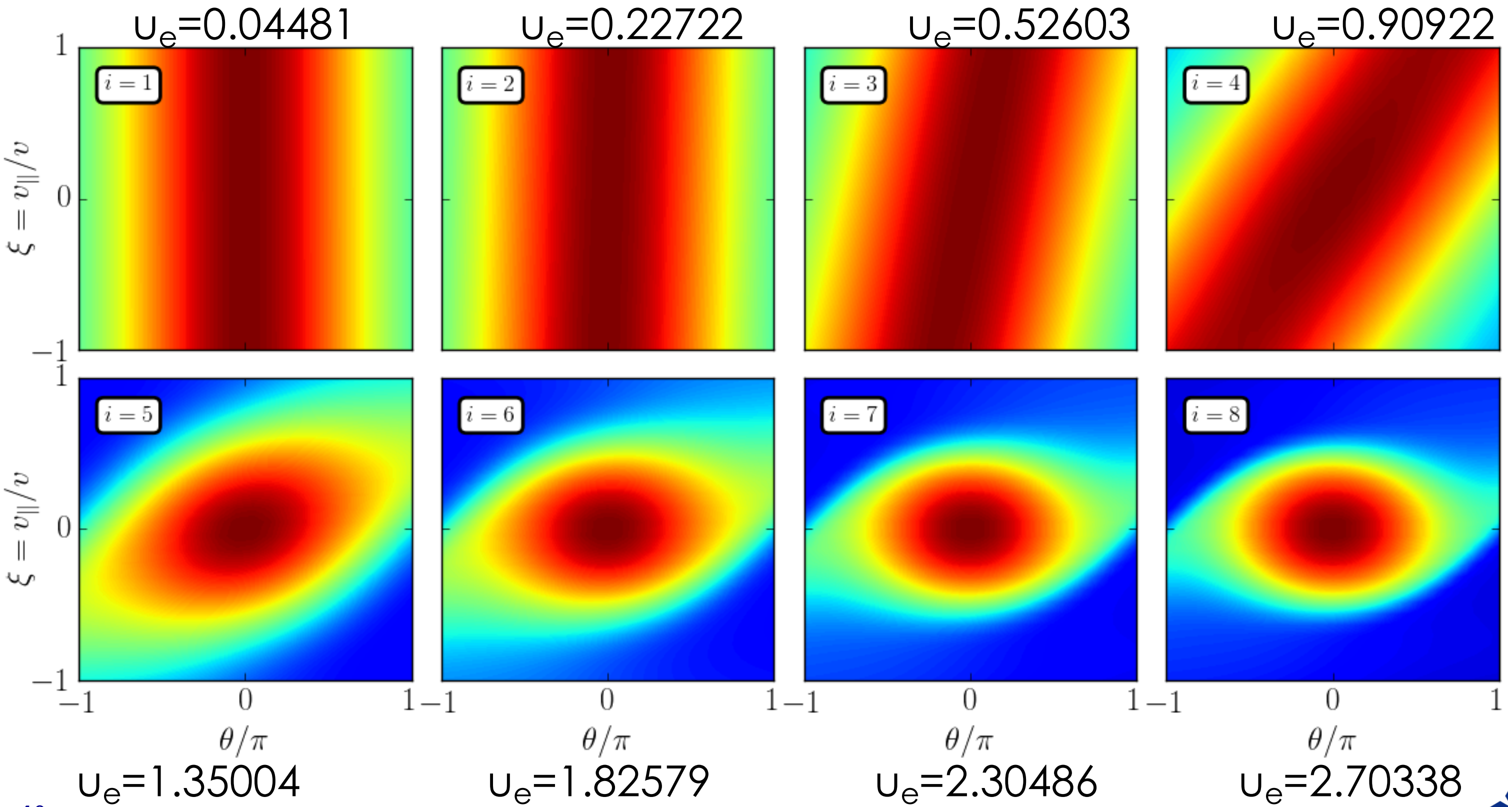

GENERAL ATOMICS

CGYRO operator splitting for time integration

$$\frac{\partial h_a}{\partial \tau} + A(H_a, \Psi_a) + B(H_a, \Psi_a) = 0$$

- **Nonlinear, Collisionless step**:
  - **Operates primarily in space** (distributed in velocity dimensions)
    - Spectral in x and y; finite difference in θ
  - Nonlinearity via **2D FFT with dealiasing** (**well-suited to GPUs** → cuFFT)
  - Explicit in time: **Adaptive embedded RK5(4)**
    - Time step restriction set by fastest Alfven wave
    - Efficient for nonlinear multiscale (large number of radial & binormal wavenumbers)
    - Adaptive algorithm gives faster solution for systems with impulse and oscillatory behavior

GENERAL ATOMICS

In the field line direction, a novel 5th order conservative algorithms is used to permit high-accuracy electromagnetic simulation.

- Focus on schemes suitable for **explicit advection**

$$\frac{\partial h_a}{\partial \tau} + \hat{v}_\parallel \frac{\partial H_a}{\partial \theta}$$

$$H_a = h_a - G_{0a}\tilde{v}_\parallel \delta A_\parallel + G_{0a}\delta\phi$$

- Implement a **5th order upwind scheme**:

$$\frac{\partial (h_a)_j}{\partial \tau} + \hat{v}_\parallel \boldsymbol{D}^{(6)}(H_a)_j + |\hat{v}_\parallel| \boldsymbol{S}^{(6)}(g_a)_j$$

$$g_a = h_a - G_{0a}\tilde{v}_\parallel \delta A_\parallel$$

6th-order finite difference (7-pt derivative)

6th-order filter (7-pt smoother)

$\boldsymbol{S}^{(6)} \sim (\Delta\boldsymbol{\theta})^5$: Continuum limit obtained as num gridpts increases

The conservative upwind scheme yields accurate discretization in the long-wavelength, high beta limit and for high-k ETG modes.

$$\frac{\partial (h_a)_j}{\partial \tau} + \hat{v}_\parallel \boldsymbol{D}^{(6)}(H_a)_j + \boldsymbol{S}^{(6)}\left[|\hat{v}_\parallel|(\boldsymbol{g}_a)_j - \langle(\boldsymbol{g}_a)_j\rangle\right]$$

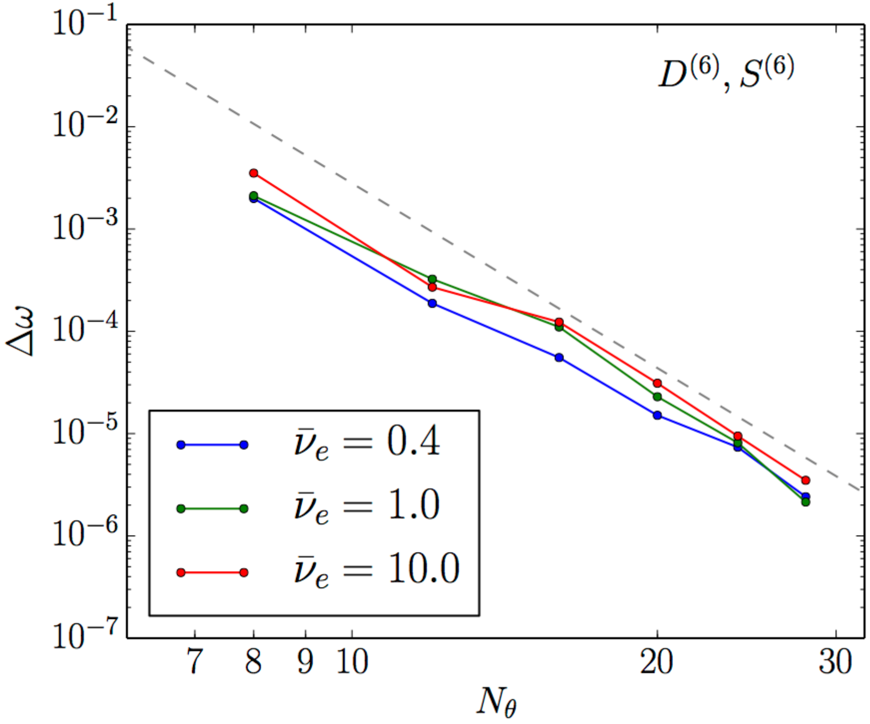

Density conservation

- **Dissipation** causes inaccuracy due to violation of number conservation
- **Project out the gyrocenter density distribution caused by the dissipation**
- Method **conserves gyrocenter number** with respect to the numerical dissipation

CGYRO operator splitting for time integration

$$\frac{\partial h_a}{\partial \tau} + \boldsymbol{A}(\boldsymbol{H}_a, \boldsymbol{\Psi}_a) + \boldsymbol{B}(\boldsymbol{H}_a, \boldsymbol{\Psi}_a) = 0$$

- **Collisional + trapping step**:
  - **Operates in velocity space** (distributed in spatial dimensions)
  - Implicit in time: **2nd order CN**
    - Required for stability due to scaling of ν_e with inverse powers of v
    - Matrix is large and well-suited to execution on GPUs

$$\begin{bmatrix} H_1^+ \\ H_2^+ \\ \vdots \\ H_{Na}^+ \end{bmatrix} = \mathbb{M} \begin{bmatrix} H_1^- \\ H_2^- \\ \vdots \\ H_{Na}^- \end{bmatrix}$$

Rank($\mathbb{M}$)=$N_\xi N_v N_a$

GENERAL ATOMICS

CGYRO uses a spatial discretization & array distribution scheme that targets scalability on next-generation HPC systems

- Operates on 5+1 dimensional grid
- Several steps in the simulation loop, where each step
  - Can cleanly partition the problem in at least one dimension
    - But no one dimension in common between all of them
  - All dimensions are compute-parallel
    - But some dimensions may rely on neighbor data from previous step

Easy to split among several CPU/GPU cores and nodes

Requires frequent transpose ops (i.e. MPI_ALLtoALL) → Communication heavy

CGYRO uses a spatial discretization & array distribution scheme that targets scalability on next-generation HPC systems

Communication happens on **2 orthogonal communicators**

Kernel	Data dependence	Dominant operation	Communication
Collisionless	$\boldsymbol{k_x^0}, \boldsymbol{\theta}$ $\quad [\boldsymbol{k_y}]_1, [\xi, \mathrm{v}, \boldsymbol{a}]_2$	Loop (linear)	MPI_ALLREDUCE
Nonlinear	$\boldsymbol{k_x^0}, \boldsymbol{k_y}$ $\quad [\boldsymbol{\theta}, [\xi, \mathrm{v}, \boldsymbol{a}]_2]_1$	FFT	MPI_ALLTOALL
Collisional	$\xi, \mathrm{v}, \boldsymbol{a}$ $\quad [\boldsymbol{k_y}]_1, [\boldsymbol{k_x^0}, \boldsymbol{\theta}]_2$	Matrix-vec multiply	MPI_ALLTOALL

- **All CGYRO kernels are ported to GPUs** using OpenACC and cuFFT
- **Critical use of GPUDirect MPI** minimizes cost of memory movement
  - Gives 30-40% reduction in comm timing on OLCF Summit
  - **Optimal for Frontier**

GENERAL ATOMICS

CGYRO uses a spatial discretization & array distribution scheme that targets scalability on next-generation HPC systems

- **For small to medium simulations:**
  - Comm1 is typically more "chatty"
  - Keep most of data inside the node to reduce network traffic
  - But increasing #MPI, increases data comm of Comm2
- **For multiscale:**
  - When #MPI is multiple of #species, can **exchange only per-species data** and comm2 data volume cut by #species
  - **Smarter, adapative time advance** reduces both compute time and data volume

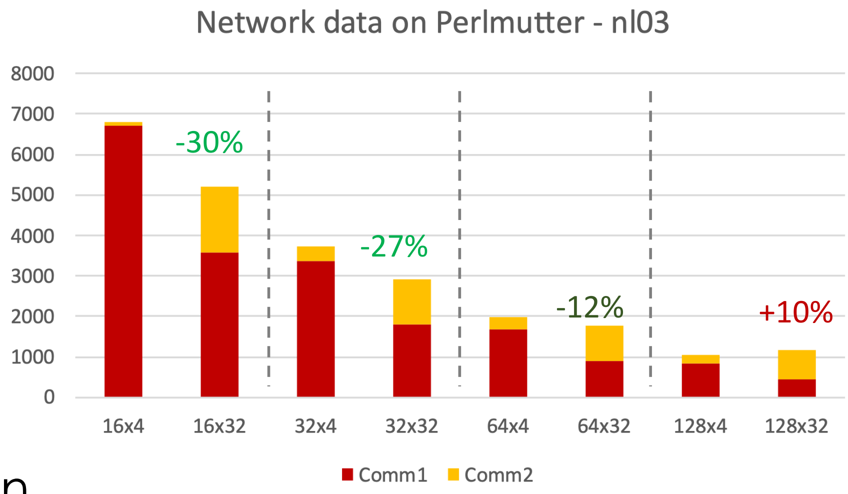

CGYRO strong scaling shows excellent performance on GPU systems on both a per node and maximum performance basis.

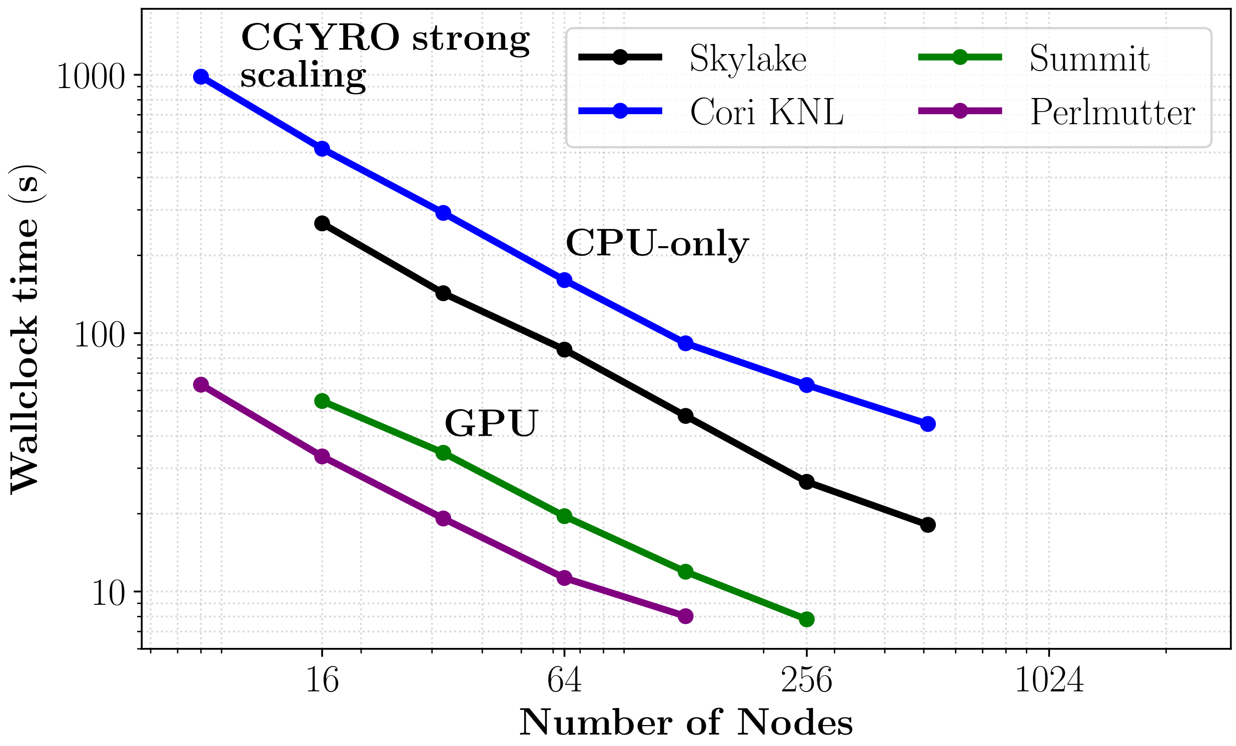

GENERAL ATOMICS

On CPU systems, compute time is dominated by nonlinear FFT and cost of communication:compute ratio is smaller.

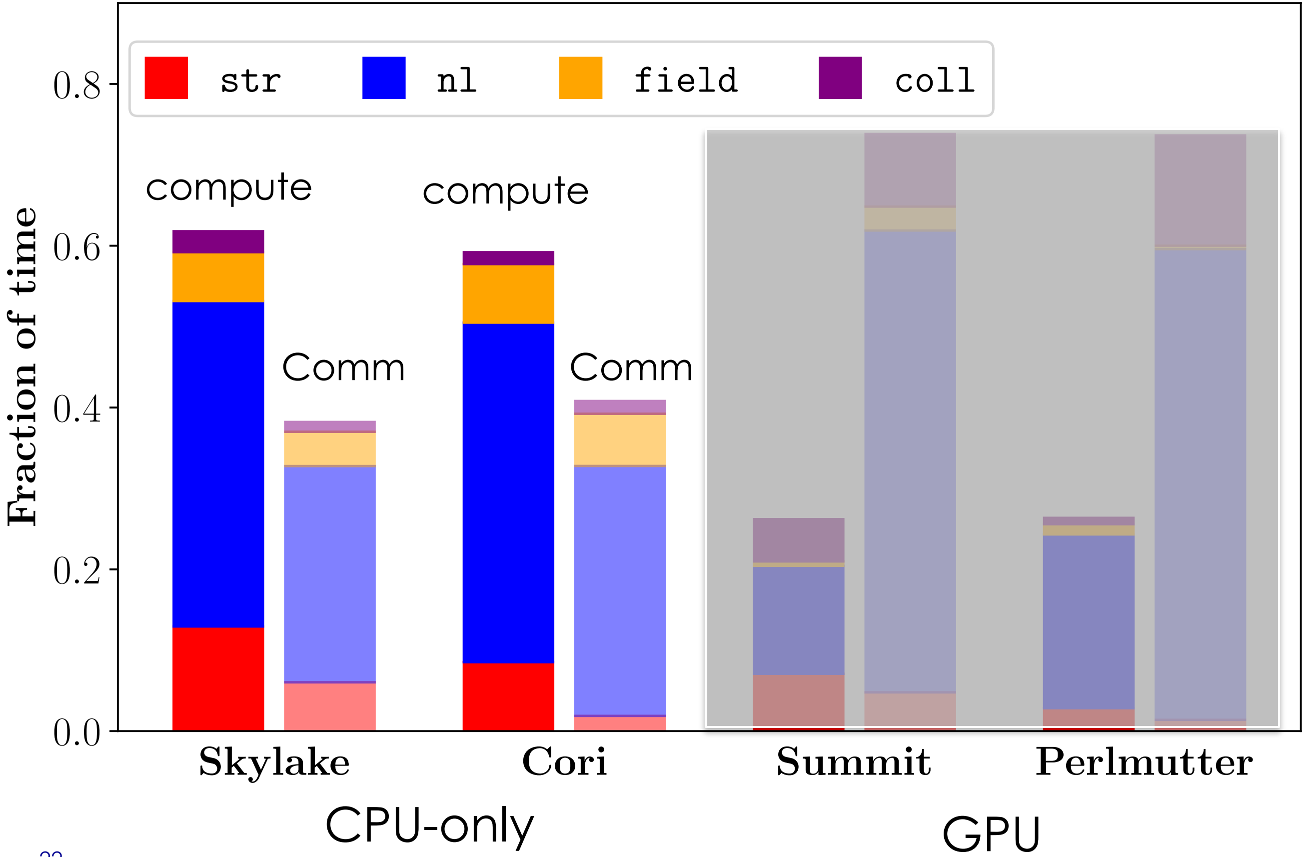

On GPU systems, high performance cuFFT means short time spent in nonlinear kernel & code is communication-intensive.

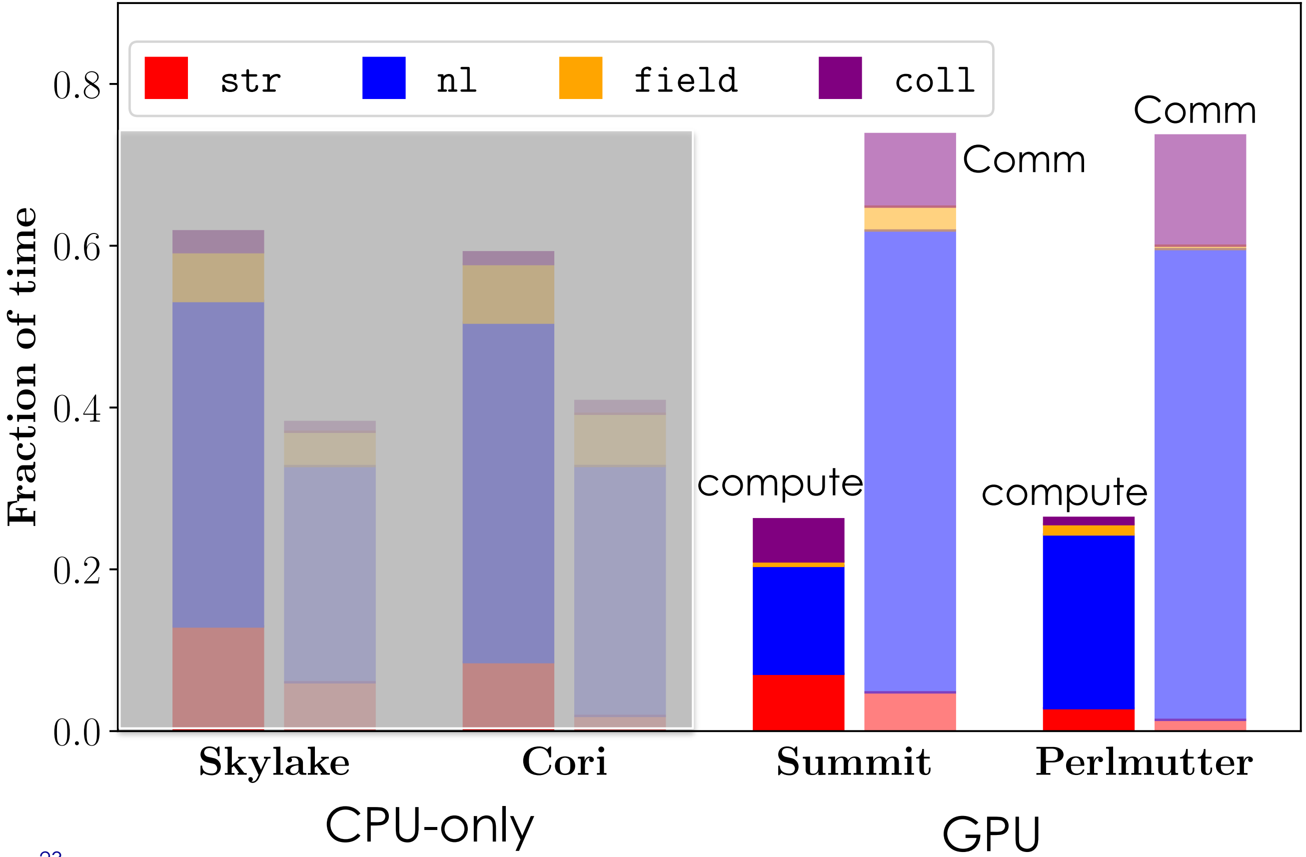

Reflects high absolute performance of GPUs rather than poor performance of interconnect

Requires 600 GB of GPU memory

CGYRO multiscale simulation is well-suited to capability simulation on accelerated systems like Summit/Frontier.

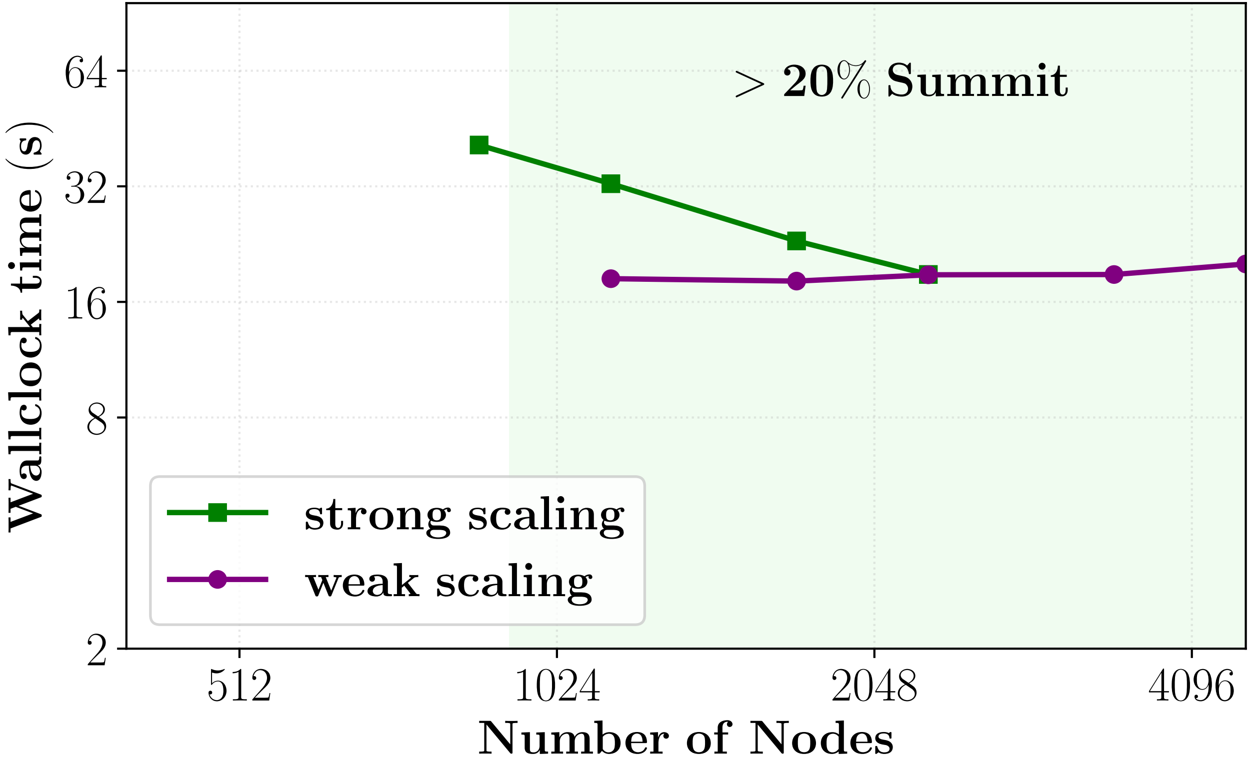

Strong scaling:
4-species

Weak scaling:
Varying $N_a = (2 - 8)$

Requires >10 TB of GPU memory

GENERAL ATOMICS

CGYRO multiscale gyrokinetic turbulence analysis in the tokamak pedestal

DIII-D ITER Baseline H-mode #164988, r/a=0.92

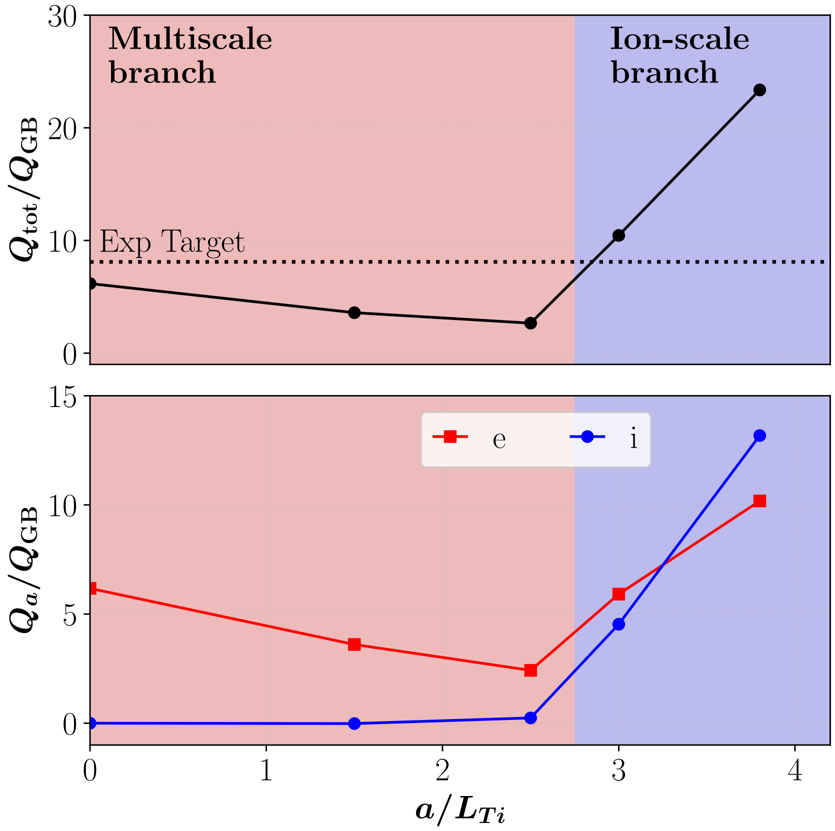

CGYRO: 250K Summit node-hrs

The experiment lies in a **bifurcation region** between **ion-scale-dominated** and **multiscale-dominated** turbulence regimes.

$$\frac{1}{L_{Ti}} = -\frac{dlnT_i}{dr}$$

GENERAL ATOMICS

Multiple drift modes are linearly unstable across a broad range of spatial scales from ion-scales to electron-scales.

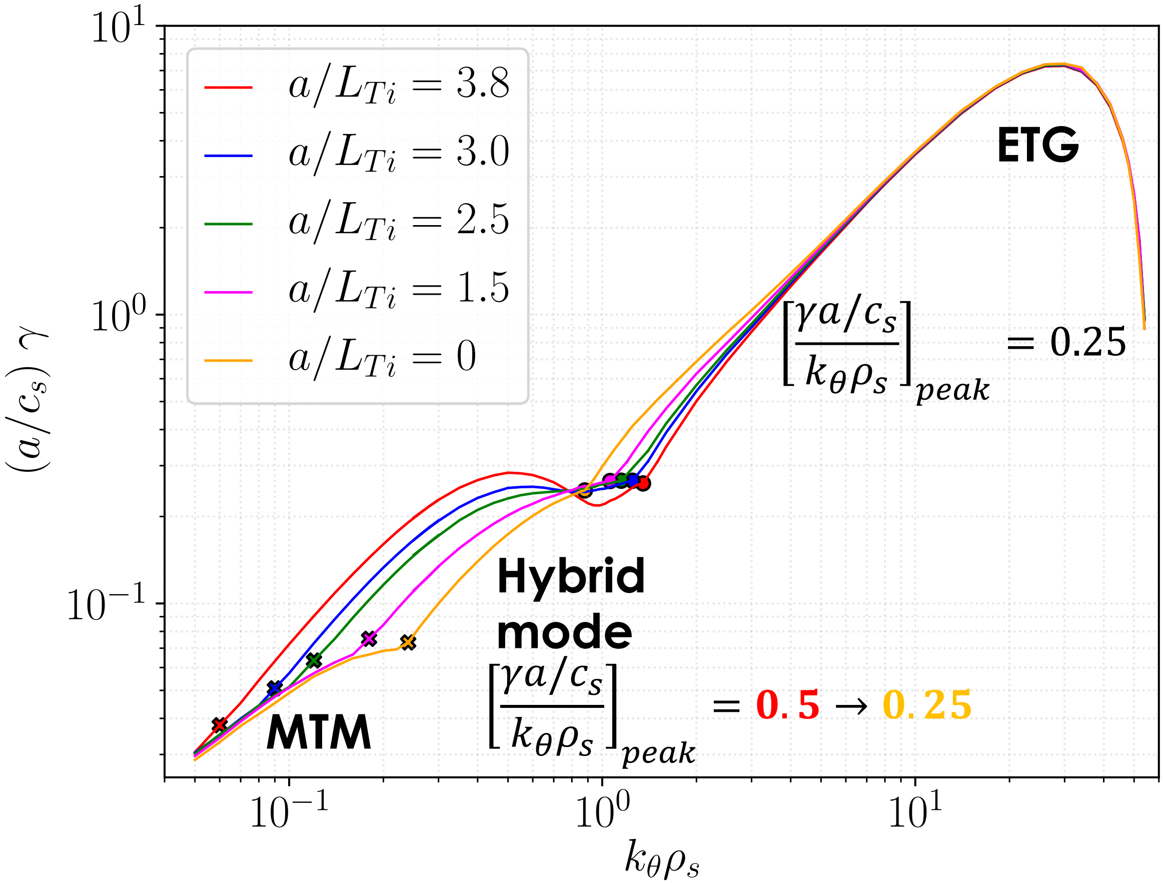

GENERAL ATOMICS

The electron-scale turbulence is reduced by ion-scale fluctuations through nonlinear mode-mode interaction & an increase in zonal flows.

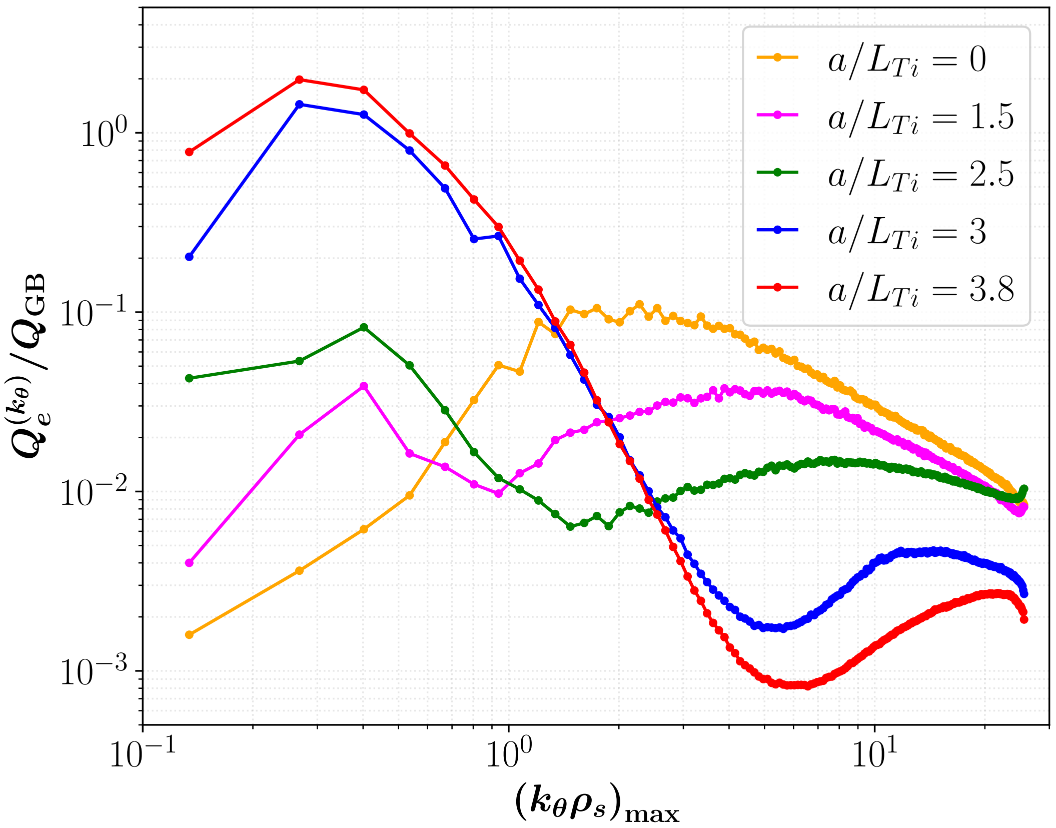

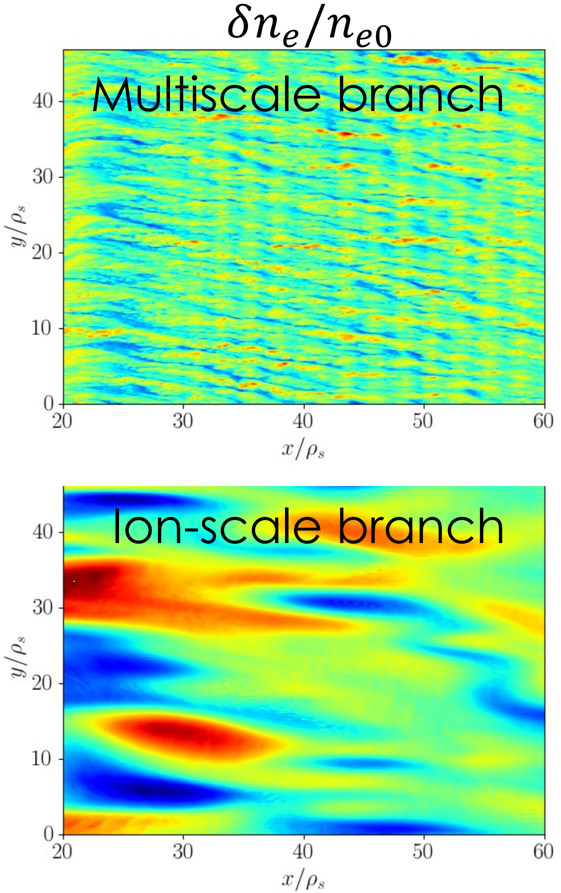

On the ion-scale branch, the fluctuating electrostatic potential intensity is peaked around k_x^0=0 and the total amplitude is large.

$$I(k_\theta, k_x^0) \doteq \left\langle \left| \delta\hat{\phi}(k_\theta, k_x^0) \right|^2 \right\rangle_t$$

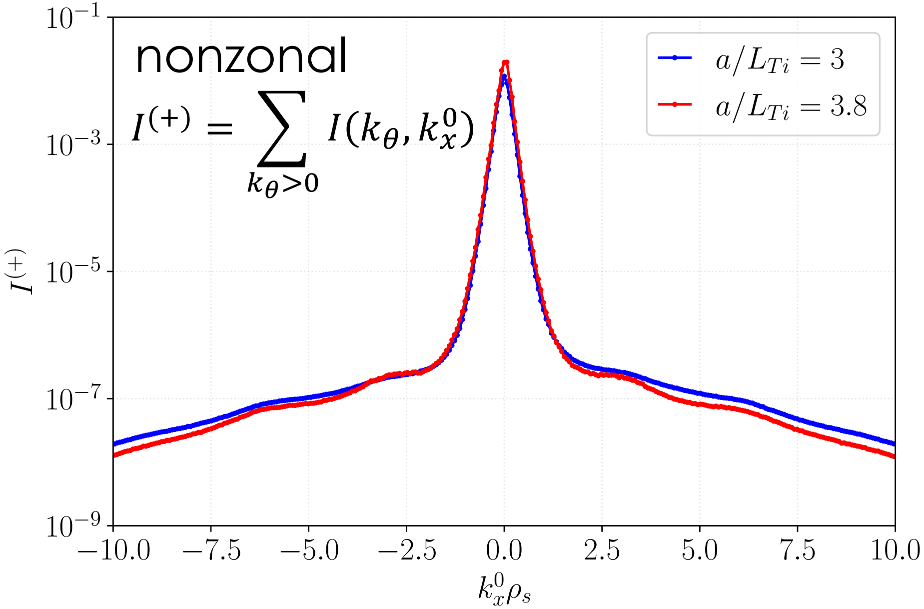

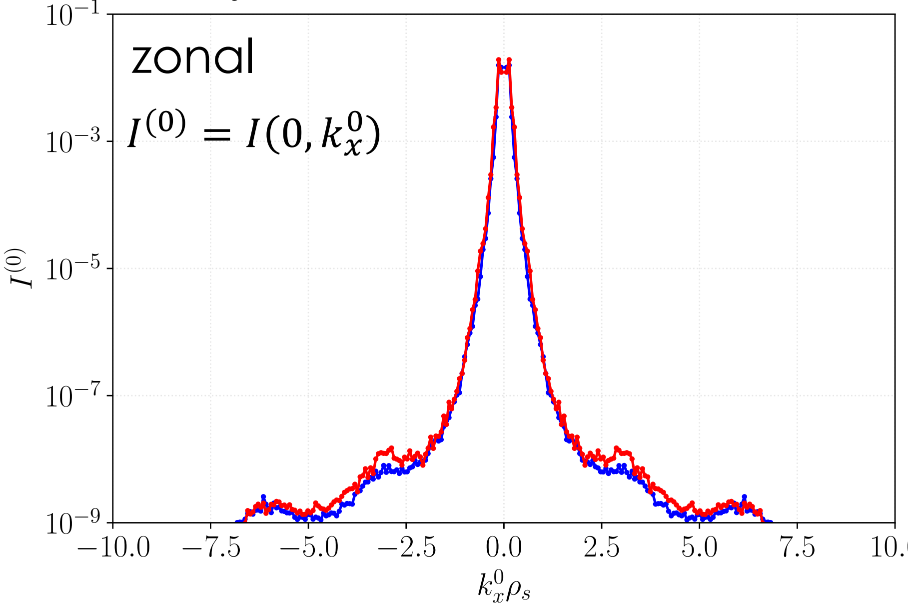

High-k_x nonzonal modes damped by ion FLR.

Zf potential is driven by low-k_θ modes and increases significantly with ion drive.

GENERAL ATOMICS

Increase in the nonzonal fluctuating intensity is well correlated with increase in the ion energy flux.

$$I(k_\theta, k_x^0) \doteq \left\langle \left| \delta\hat{\phi}(k_\theta, k_x^0) \right|^2 \right\rangle_t$$

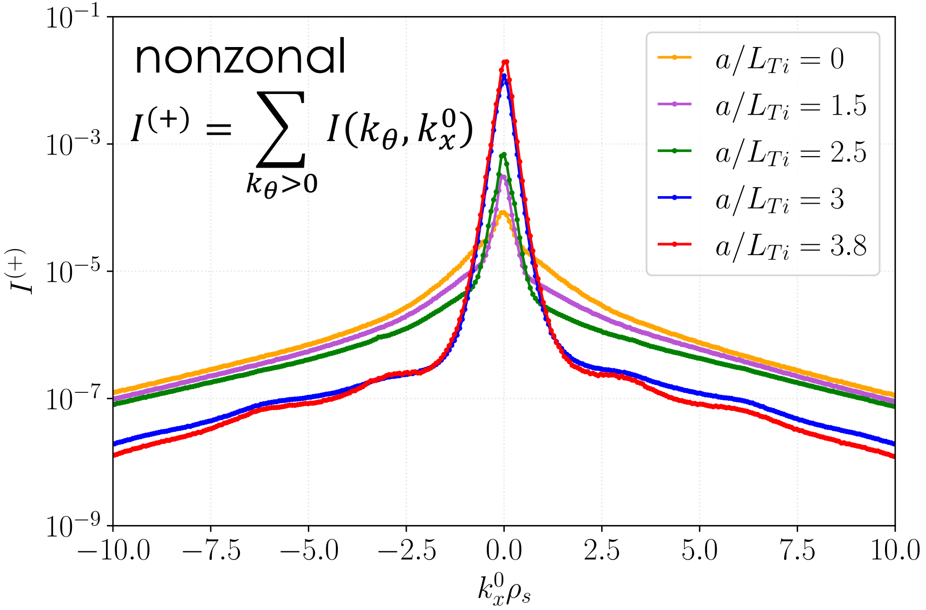

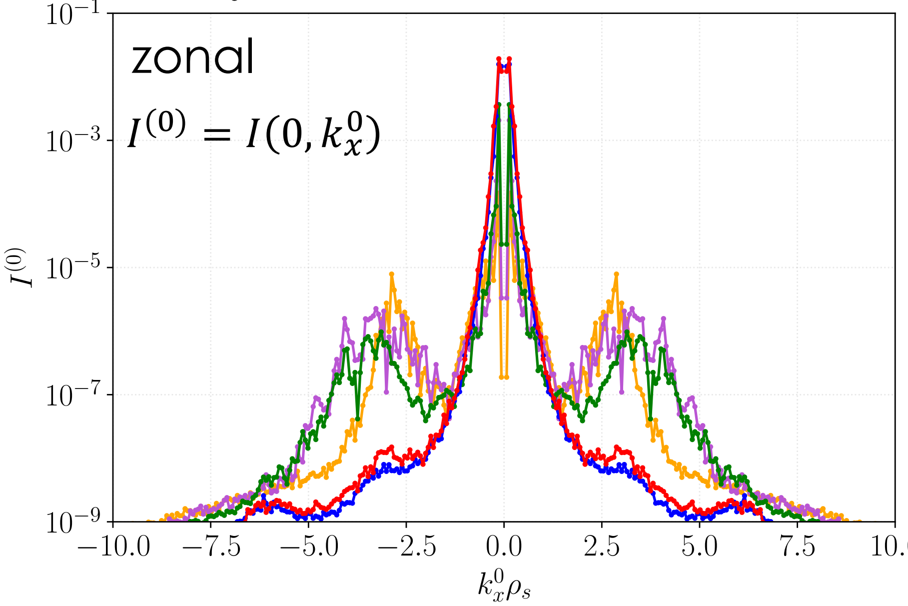

Long wavelength nonzonal fluctuations & low-k-driven zfs play a role in suppressing high-kETG turbulence.

ETG-driven zf's give secondary peaking & help to regularize multiscale turbulence.

The drift energy associated with the fluctuating ExB velocity is enhanced for the multiscale branch.

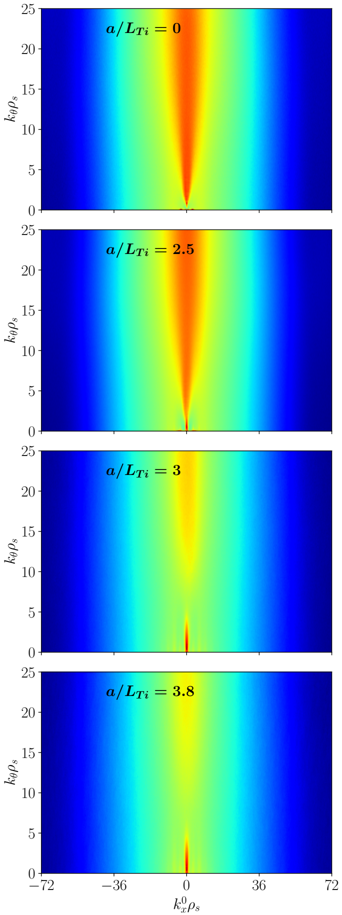

$$K(k_\theta, k_x^0) \doteq k_\perp^2 \rho_s^2 \left\langle \left| \delta\hat{\phi}(k_\theta, k_x^0) \right|^2 \right\rangle_t$$

$$\vec{v}_{ExB}(\vec{k}_\perp) = -i\frac{c}{B}\delta\hat{\phi}(\vec{k}_\perp)\vec{k}_\perp \times \vec{b}$$

In multiscale to ion-scale transition, the energy shifts from dominantly drift kinetic to potential intensity & is correlated with the energy flux.

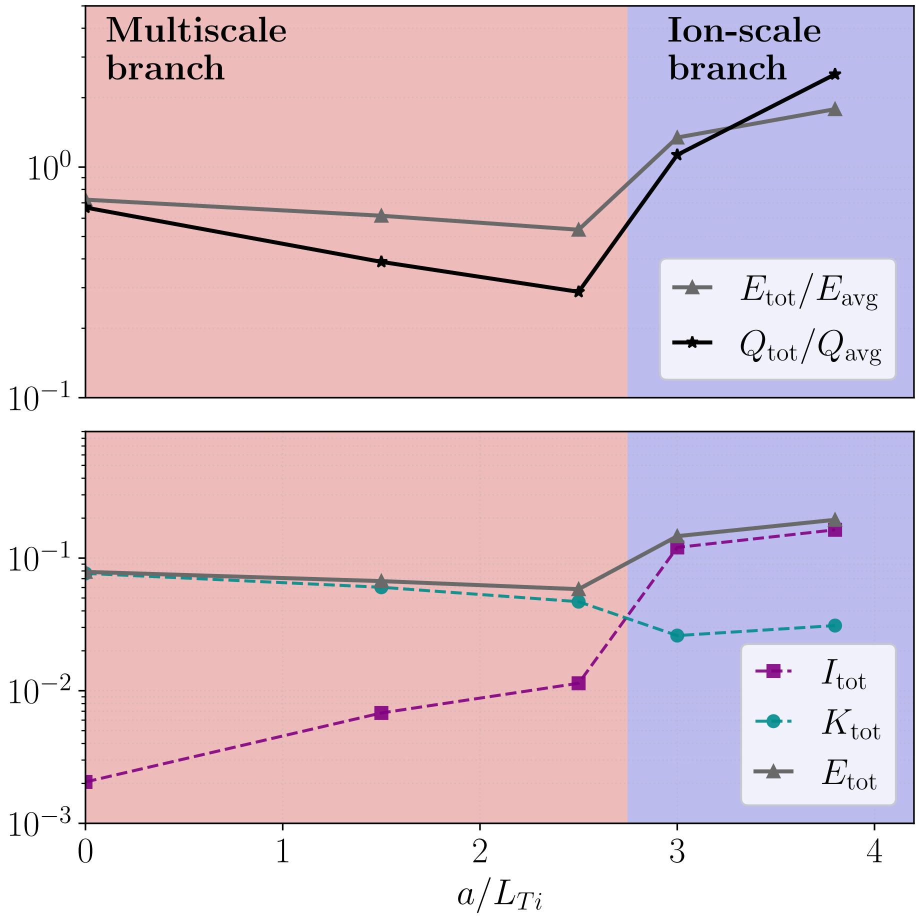

$$E_{tot} = K_{tot} + I_{tot}$$

Drift energy associated with ExB velocity:

$$K(k_\theta, k_x^0) \doteq k_\perp^2 \rho_s^2 \left\langle \left|\delta\hat{\phi}(k_\theta, k_x^0)\right|^2 \right\rangle_t$$

Fluctuating electrostatic potential intensity:

$$I(k_\theta, k_x^0) \doteq \left\langle \left|\delta\hat{\phi}(k_\theta, k_x^0)\right|^2 \right\rangle_t$$

Summary

- A **multiscale-optimized gyrokinetic turbulence solver** (CGYRO) was developed.
  - Uses highly efficient **spectral/pseudospectral** numerical schemes in 4/5 dims
  - Pseudospectral in velocity space gives optimal accuracy of collisions
  - Fully spectral in $k_\perp$ gives spectral gyroaverages → efficiency for multiscale
  - Nonlinear evaluation on GPUs (cuFFT) → maximum performance & scalability
  - Novel conservative upwind scheme in θ permits high accuracy EM simulation
  - Spatial discretization and array distribution scheme **targets scalability on next-generation, exascale HPC systems** (GPU-accelerated)
- Optimizations enabled a **first multiscale analysis of pedestal-like transport** with full ion-electron cross-coupling.
  - Experiment lies in bifurcation region btw multiscale-dominated & ion-scale-dominated turbulence regimes. In the transition, electron-scale transport is reduced by nonlinear mixing w/ ion-scale fluctuations & ion-scale-driven zfs.

First CGYRO Exascale simulations of multi-species burning plasmas on Frontier in 2023