SciDAC GSEP & CSEP: Gyrokinetic Simulation of Energetic Particle

Zhihong Lin

University of California, Irvine

NERSC FES Requirements for 2017
March 19-20, 2013
Rockville, MD

1. Project Description: GSEP

Zhihong Lin, University of California, Irvine
Ronald Waltz, General Atomics
Donald Spong, Oak Ridge National Laboratory
Chandrika Kamath, Lawrence Livermore National Laboratory

- Summarize your project(s) and its scientific objectives through 2017

extend the capability of highly scalable gyrokinetic (GK) turbulence codes based on the complementary particle-in-cell GTC and continuum GYRO to study the energetic particle (EP) driven Alfvén modes and their interaction with background thermal plasmas

1. Project Description

- Our present focus is …

2008-2011: extend GK from microturbulence (micro-scale) to EP Alfven modes (meso-scale); V&V.

2011-2016: study EP turbulence & transport; extend GK to MHD (macro-scale).

First-principles, integrated, peta-scale GK simulation of kinetic MHD: microturbulence, EP, MHD, neoclassical...

- By 2017 we expect to …

2016-?: study EP interaction with background thermal plasmas (microturbulence & MHD); extend global code from kinetic MHD to include heating/current drive.

Exa-scale simulation of multi-physics processes.

Linear Simulations of TAE, RSAE, BAE Verified

- DIII-D shot # 142111, GTC, GYRO, and TAEFL frequency up-sweeping, mode structures, and transition of RSAE to TAE in good agreement with experiments.

- BAE *[H. Zhang et al, PoP2010]*
- RSAE *[W. Deng et al, PoP2010]*
- TAE *[E. Bass & R. Waltz, PoP2010; W. Zhang et al, PoP2012]*

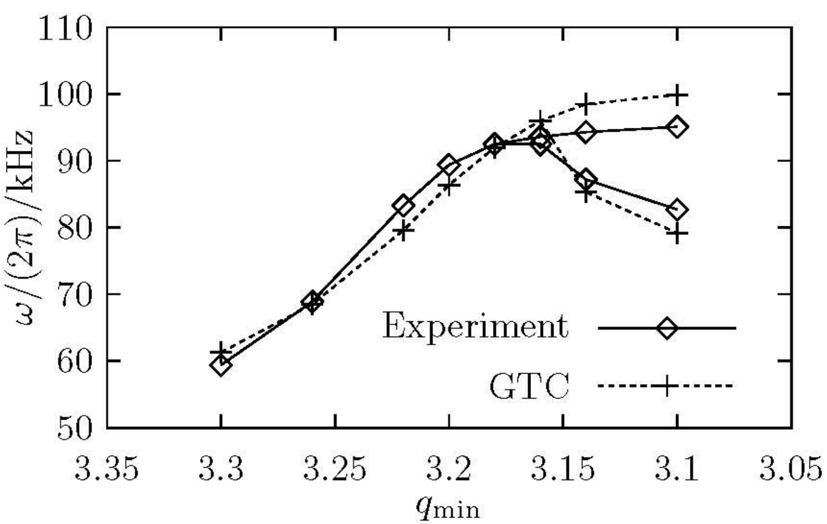

GTC simulation of RSAE to TAE transition with real geometry & kinetic electrons

[W. Deng et al, NF2012]

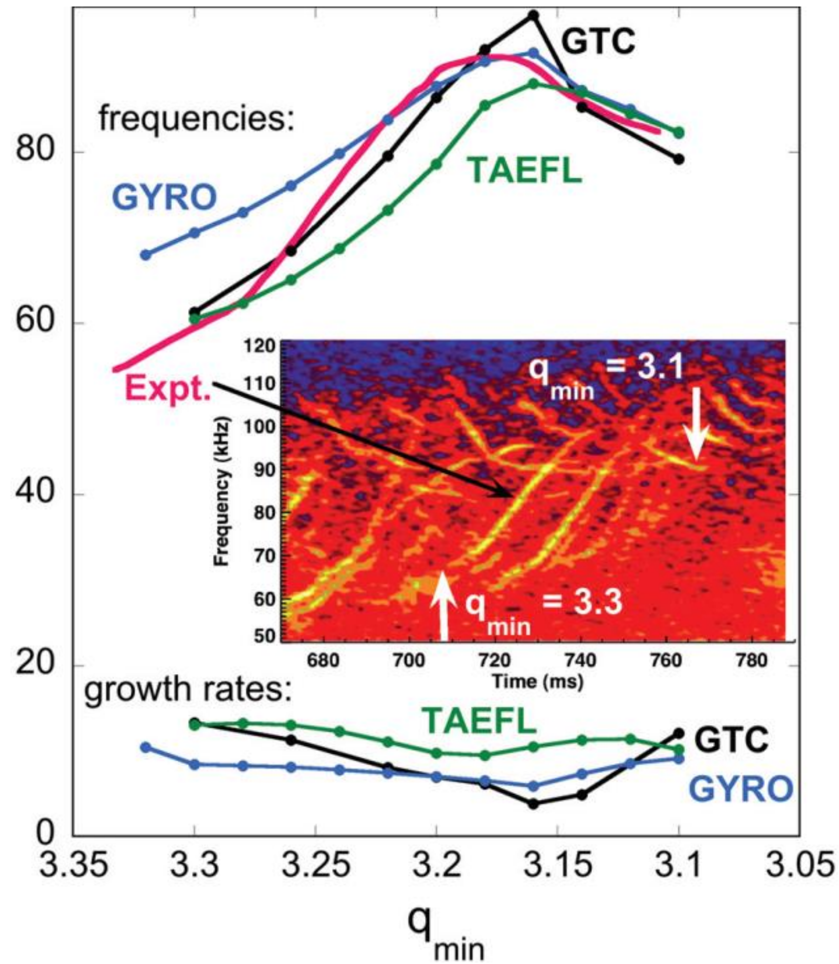

RSAE in DIII-D shot # 142111 at 725ms

[D. Spong et al, PoP2012]

GTC Nonlinear Simulations of BAE Find Fast Chirping

- Fast, repetitive, mostly downward chirping, sub-millisecond period
- 90^o phase shift between intensity oscillation and frequency chirping
- Simulation features observed in recent NSTX TAE, ASDEX BAE
  - EP transport enhanced by chirping
- No sources and sinks. Chirping mechanism? Universal dynamics?

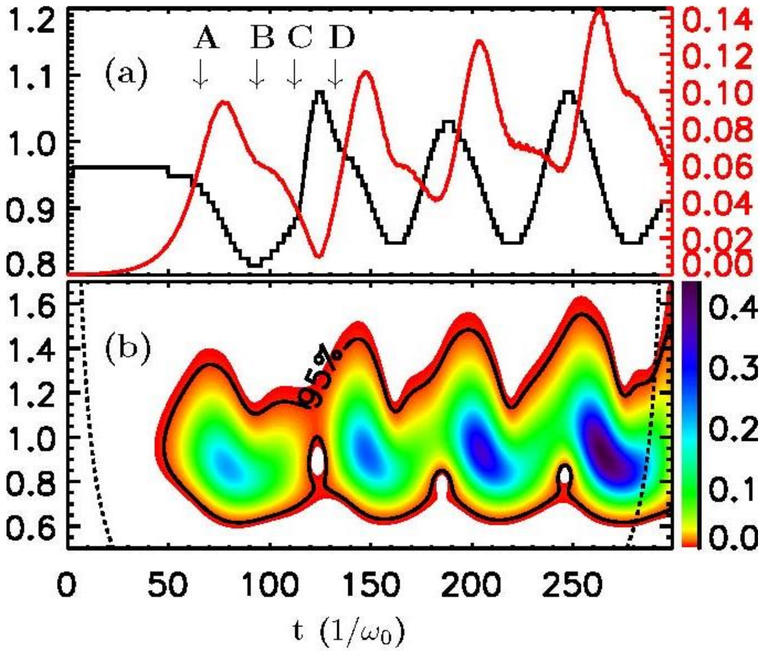

[H. S. Zhang, Z. Lin, I. Holod, PRL2012]

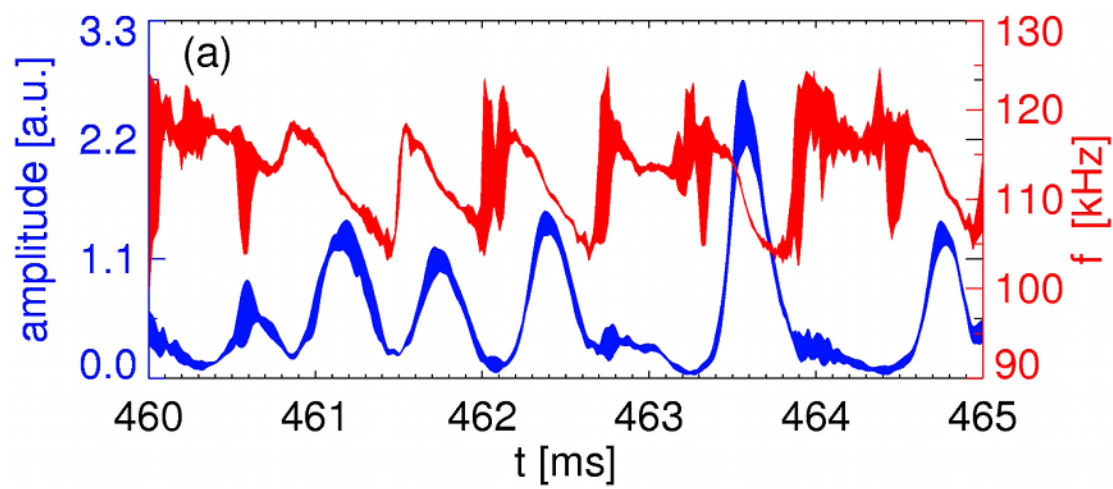

[M. Podesta et al, NF2011; PPPL-4719]

2. Computational Strategies

- We approach this problem computationally at a high level by …
  GK particle-in-cell GTC and continuum GYRO
- The codes we use are …
  GTC, GYRO, TAEFL, HMGC
- These codes are characterized by these algorithms: …
  PIC, continuum, gyrofluid, MHD-PIC hybrid
- Our biggest computational challenges are …
  Improve efficiency while maintaining scalability
- Our parallel scaling is limited by …
  linear solvers, I/O
- We expect our computational approach and/or codes to change (or not) by 2017 in this way …

Gyrokinetic Toroidal Code (GTC)

- *Confinement and stability properties of fusion plasmas depend on nonlinear interaction of multiple physical processes*
  - ▸ Microturbulence, energetic particle (EP), magnetohydrodynamic (MHD) modes, heating/current drive using radio-frequency (RF), ….

- GTC Physics Integration: first-principles simulations of microturbulence + EP + MHD + RF
  - ▸ General geometry & experimental profiles
  - ▸ Kinetic electrons & electromagnetic fluctuations
  - ▸ Gyrokinetic or fully kinetic ions
  - ▸ Equilibrium current; δf or full-f
  - ▸ Scalable to 10^5 cores; GPU acceleration

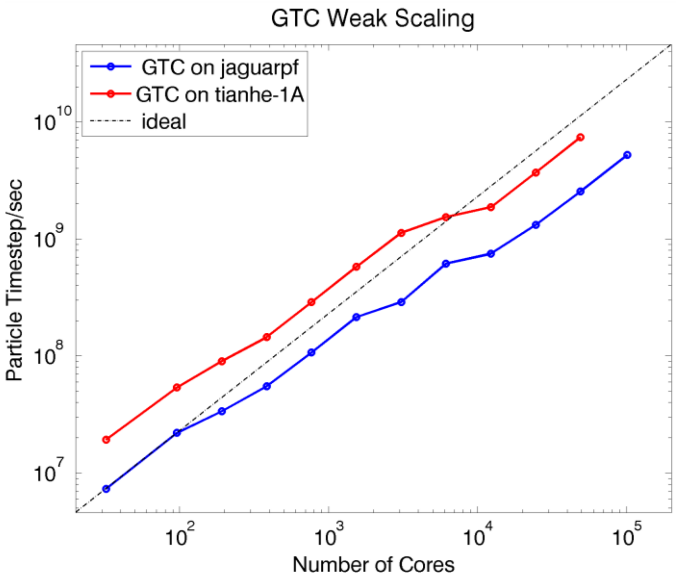

GTC scalability using parameters from actual physics simulations

[Z. Lin et al, Science1998]

[I. Holod et al, PoP2009]

[http://phoenix.ps.uci.edu/GTC]

3. Current HPC Usage (see slide notes)

- Machines currently using

  Hopper, Jagarpf, Tianhe-1A, Titan, Edison
- Hours used in 2012 (list different facilities)

  Hopper: 20M; Jaguarpf: 5M/24M (ALCC); Tianhe-1A: 5M (ITER-CN)
- Typical parallel concurrency and run time, number of runs per year

  10^3-10^5 cores, 1-20 hours, hundreds of runs per year
- Data read/written per run

  restart files: 0.01 to 0.1 of total node memory
- Memory used per (node | core | globally)

  0.1 to 1 of node memory
- Necessary software, services or infrastructure

  linear solvers, I/O
- Data resources used (HPSS, NERSC Global File System, etc.) and amount of data stored: HPSS: 7unit

4. HPC Requirements for 2017

(Key point is to directly link NERSC requirements to science goals)

- Compute hours needed (in units of Hopper hours)

  1 billions hours
- Changes to parallel concurrency, run time, number of runs per year

  parallel concurrency increases with core #, run time/number of runs stay unchanged
- Changes to data read/written

  size increases with core #
- Changes to memory needed per (core | node | globally)

  Memory stays unchanged per core, increases with code # per node/globally
- Changes to necessary software, services or infrastructure

  linear solvers (e.g., petsc by ANL), I/O by ORNL

5. Strategies for New Architectures

- Our strategy for running on new many-core architectures (GPUs or MIC) is …

  Collaboration with CS.
- To date we have prepared for many core by …

  Multi-level parallelism: MPI/OpenMP

  Production version: GTC-GPU via CUDA Fortran & OpenACC (NSCC-TJ & Nvidia)
- We are already planning to do …

  MIC optimization (NSCC-TJ).
- To be successful on many-core systems we will need help with

  Linear solver (ANL), I/O (ORNL)

5. Summary

- What new science results might be afforded by improvements in NERSC computing hardware, software and services?

  First-principles, integrated simulation of cross-scale interaction between multiple physics processes.
- Recommendations on NERSC architecture, system configuration and the associated service requirements needed for your science
  heterogeneous, multi-level concurrency.
- NERSC generally refreshes systems to provide on average a 2X performance increase every year. What significant scientific progress could you achieve over the next 5 years with access to 32X your current NERSC allocation?

  Coupling between microturbulence, EP, MHD, neoclassical, e.g., Neoclassical tearing mode, resistive wall mode, sawtooth etc.
- What "expanded HPC resources" are important for your project?

  Linear solvers with multi-level parallelization
- General discussion