M3D-C1 ZOOM Meeting

02/07/2022

Upcoming Meetings

CS Issues

1. LBL Report
2. Perlmutter status
3. Traverse status --
4. Mesh adaptation update
5. NERSC Time
6. Timing tests
7. Changes to github master since last meeting
8. Regression tests

Physics Studies

1. Hybrid simulation with kinetic thermal ions – Chang Liu
2. Update on Soft beta limit study

Note: **meeting minutes posted on m3dc1.pppl.gov**

In attendance

Steve Jardin
Adelle Wright
Dingyun Liu
Hank Strauss
Nate Ferraro
P Sinha
Chang Liu
Anders Kleiner
Brendan Lyons
Jin Chen
Chen Zhao
Mark Shephard
Seegyoung Seol
Morteza Siboni
Usman Riaz
Sam Williams
Sherry Li
Yang Liu

Upcoming Meetings

- Sherwood April 4-6 Santa Rosa, CA (in person)
  - CTTS Sunday April 3
- ITPA: MHD, Disruptions, Control April 4-8
- EPS Conference on Plasma Physics June 27—July 1 (online)
  - Abstract submission by February 25
- IAEA Technical Meeting on Plasma Disruptions and their Mitigation 19-22 July
  - In person at ITER HQ in France
  - Abstract submission by May 31

LBL Report

Perlmutter status

- Perlmutter is available for testing, but is not in a production mode
- Maintenance of the cooling system for the next 6 weeks means that only 500 nodes will be available to users
- I have done some timing tests and find it to be unreliable…code crashes after some timesteps for no apparent reason
- No file /usr/common/usg/bin/nersc_host
- Problem with installing PETSc with NVIDIA…Seegyoung

TRAVERSE GPU Status (Jan 27 from Chang Liu)

I have pushed the code to do gpu matrix assembling to the master branch. They are named ludef_t_gpu.f90 and metricterms_new_gpu.f90. They should not affect the code compiling on non-GPU machines. You can test it on traverse using ARCH=traverse_gpu, or do the following

export M3DC1_CODE_DIR=<your path to M3DC1 directory>
cd $M3DC1_CODE_DIR
module use $M3DC1_CODE_DIR/unstructured/modules/traverse
module load m3dc1/devel-gpu

You can then compile the code on traverse and run regression tests. Currently, only part of the code is ported, and it seems that only the RMP test can pass. I will continue working on it and make more tests pass.

Note: to compile a GPU version you need to enable ACC=1 when compiling.

Mesh adaptation update

- Pellet case …. Brendan
- Interest in Soft-Beta-Limit case ?

NERSC Time

mp288 Remaining %: 97.7%

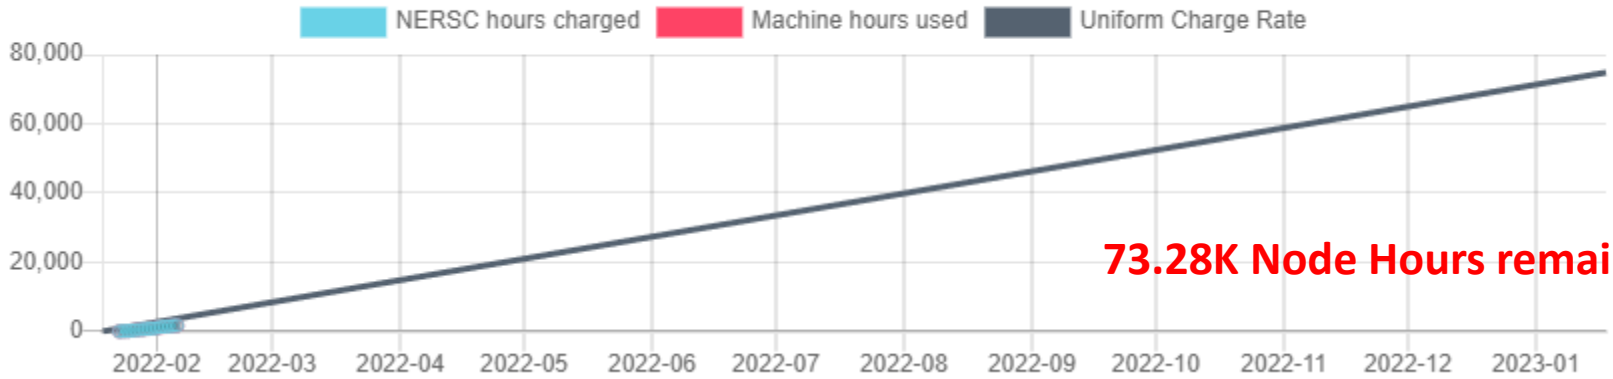

73.28K Node Hours remaining

- New award period began Jan 19
- We are NESAP Tier 2 for Pearlmutter. . Phase-I w GPUs We have been given a repo m3984 with a small allocation. Presently we are not being charged.
- N9ES-N2 M3D-C1: J. Chen , C. Liu, S. Seol are early users

FY22 allocation (started Jan 19)

Project name: mp288
CPU Node Hours Award: 75,000
GPU Node Hours Award: 7,000
Archive Storage Award (TB): 157
Project CFS Award (TB): 20

75,000 = .34 x 365 x 24 x 25
=> we can use an average of 25 nodes on Haswell continuously

➔ For 2022, the Machine Charge Factors are:

➔ Perlmutter CPU Nodes: 1.0
➔ Cori KNL Nodes: 0.20
➔ Cori Haswell Nodes: 0.34
and the charge units are "CPU Node Hours"

Timing Tests

	P	T	N	ludef -s	solve-s	rate/task
Cori Haswell	32	16	48	112	180	4.45×10^{-6} s^{-1}
	64	32	48	72	95	3.89×10^{-6} s^{-1}
Perlmutter (CPU)	32	32	24	85	160	5.37×10^{-6} s^{-1}
Stellar	32	96	8	125	200	4.04×10^{-6} s^{-1}

P -- # of partitions per plane
T -- # of tasks per node
N --# of nodes

Note: Perlmutter CPU-only nodes will have twice the memory of present nodes

Rate = $(\text{ludef-s} + \text{solve-s})^{-1}$

28792 = 170^2 triangles, 24 planes

Changes to github master since 01/10/22

Nate Ferraro:

02/03/22: Fixed bug in plot_surfmn introduced when plot_surfmn was extended to 3D data

02/03/22: Improvements to plot_br for 3D data

02/03/22: Fixed error in field_data.pro parsing kprad_particle_source ionization state

02/04/22: Added ntor parameter to plot_bmn for compatibility with plotting 3D data

02/04/22: Fixed bug plotting linear bmn data

02/04/22: Corrected sign of toroidal field perturbation with irmp=1 in 3D version

02/04/22: Updated field_spectrum.pro to use only positive ntor values

Updated schaffer_plot.pro to correctly normalize 3D data

Jin Chen:

01/24/22 : ST=1 option on Perlmutter

Steve Jardin:

02/04/22: Added kappari_fac; ion parallel thermal conductivity is kappari_fac x electron value

Chang Liu:

01/27/22: Enable matrix assembling using GPU on traverse

01/31/22: Fix a type in ludef_t.f90

Seegyoung Seol

02/06/22: minor change in mesh collapse not to mandate a model file

New Variable

kappari_fac

Ion parallel thermal conductivity is kappari_fac x electron value
(only for numvar=3 and ipres=1)

Local Systems

- PPPL centos7(02/07/22)
  - 7 jobs PASSED
- PPPL greene (02/07/22)
  - 5 jobs PASSED
- STELLAR (02/07/22)
  - 7 regression tests **PASSED** on stellar
- TRAVERSE(02/07/22)
  - 7 regression tests **PASSED** on (01/24/22)
  - 7 regression tests **FAILED** on (02/07/22)

Other Systems

- Cori-KNL (02/07/2022)
  7 regression tests **PASSED**

- Cori-Haswell (02/07/2022)
  7 regression tests **PASSED**

- Perlmutter (01/29/2022)
  - 6 regression tests **PASSED**
  - NCSX timed out
  - “make all” does not work
  - Had to modify the PATH command, replacing ‘pwd” with actual directory

Hybrid simulation with kinetic thermal ions

Chang Liu to present

Update on Soft Beta Limit Study

Comparison of the 3 scaled equilibria

BS=1.1
β=5.8%

BS=1.0
β=6.8%

BS=0.9
β=8.2%

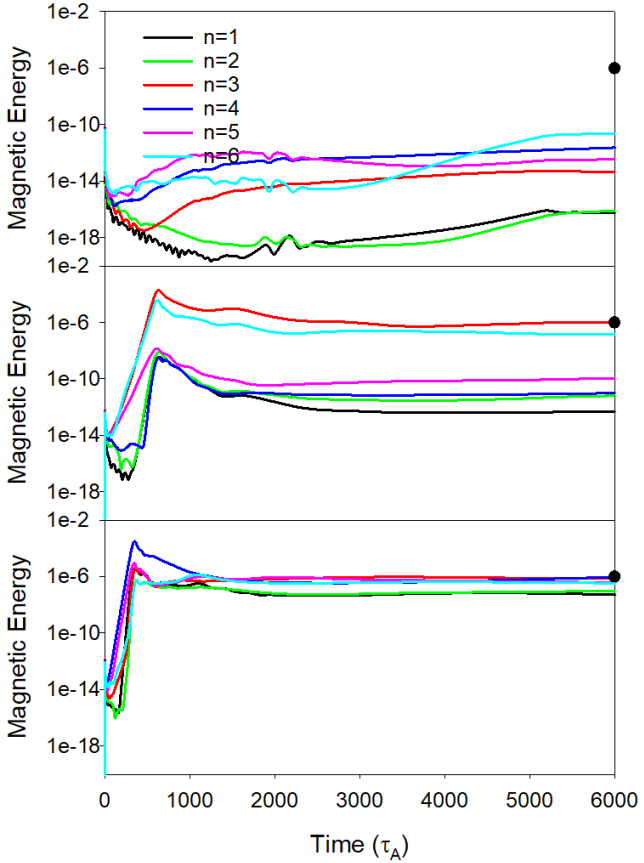

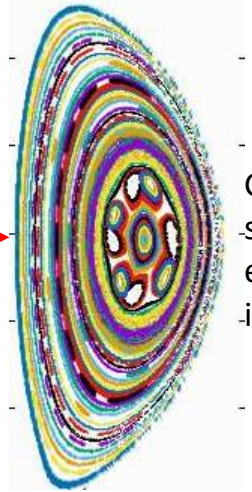

Good surfaces except n=3 in center

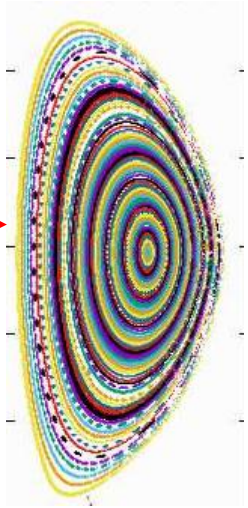

Good surfaces everywhere

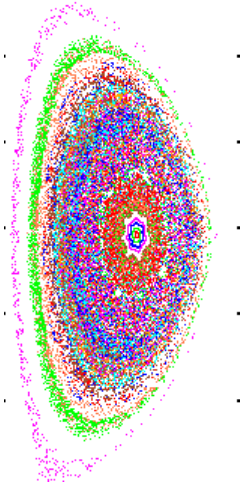

Poor surfaces with multiple n-modes

Trend is similar to experiments on NSTX

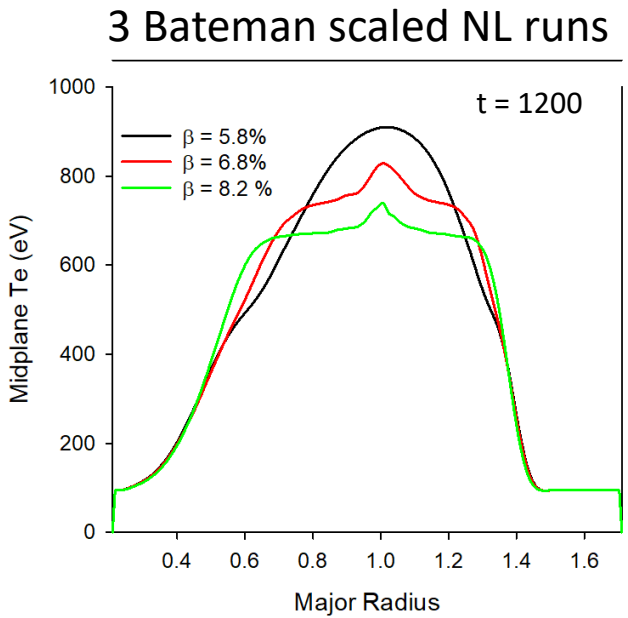

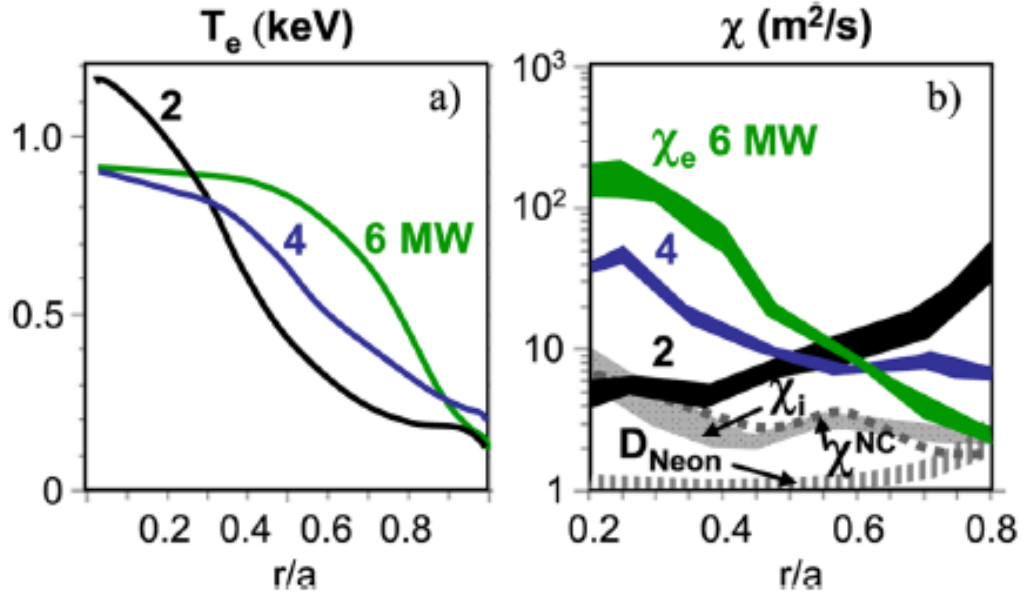

- M3D-C1: Central temperature decreases with β
- Exp data: Central transport increases with β

Stutman, et al. PRL (2009)

M3DC1 shows similar scaling with B_T as experiment

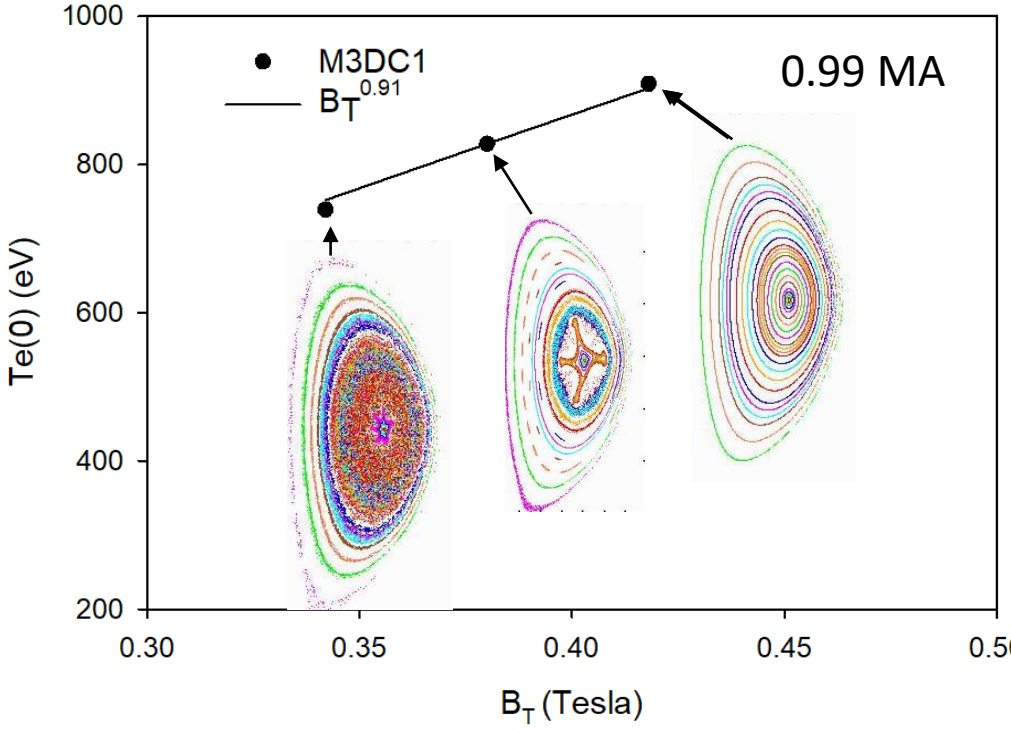

Note: Plot on left is Te(0). On the right is τ_E

Kaye, et al, PRL (2007)

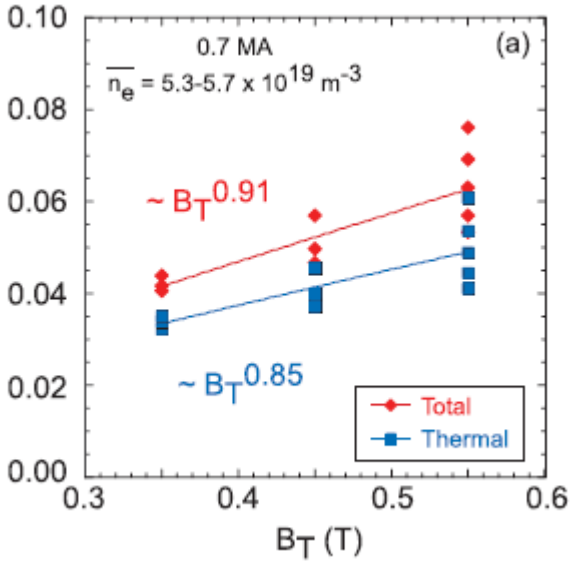

me of the discharges in this study exhibit both low amplitude low-n MHD activity as well as the fast ion driven Alfven eigenmode (AE) activity,"

More realistic: Start with stable equilibrium and apply heating power: First in 2D

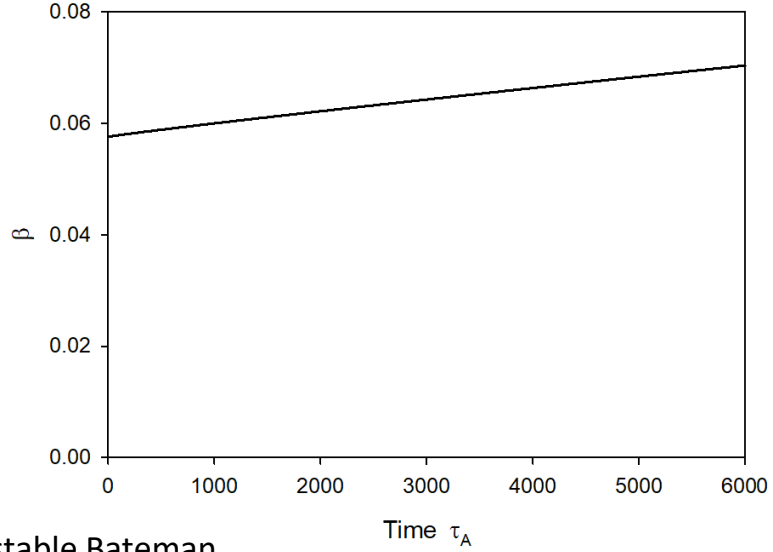

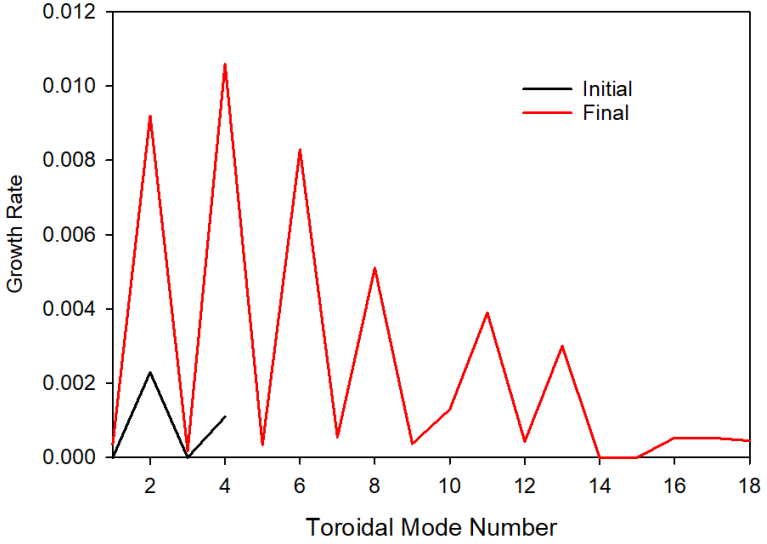

- Start with stable Bateman scaled equilibrium with β = 5.8%
- Run in 2D with heating source, increasing β to 7%
- Linear analysis shows final equilibrium unstable to many modes (shown on right)
- Now repeat with 3D run. Do these saturate nonlinearly?

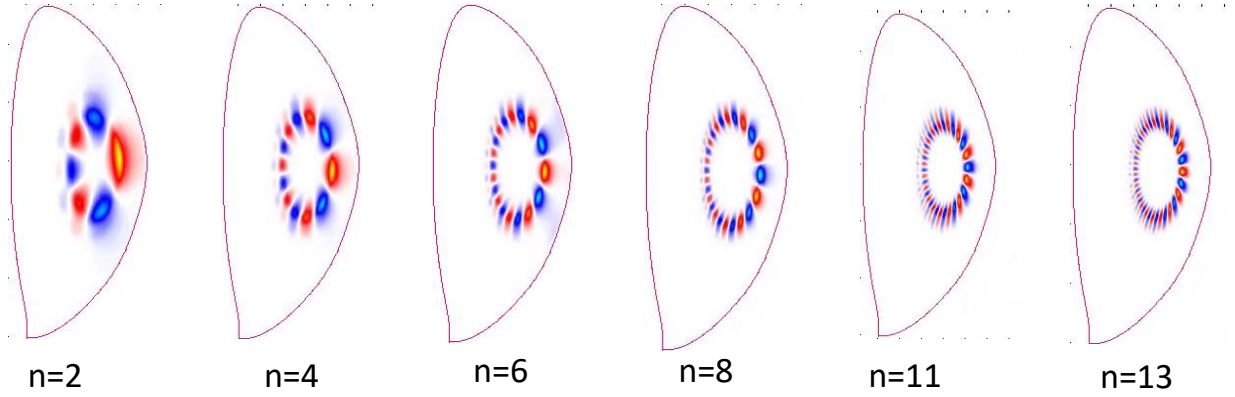

Summary of 0 < t < 3000 τ_A

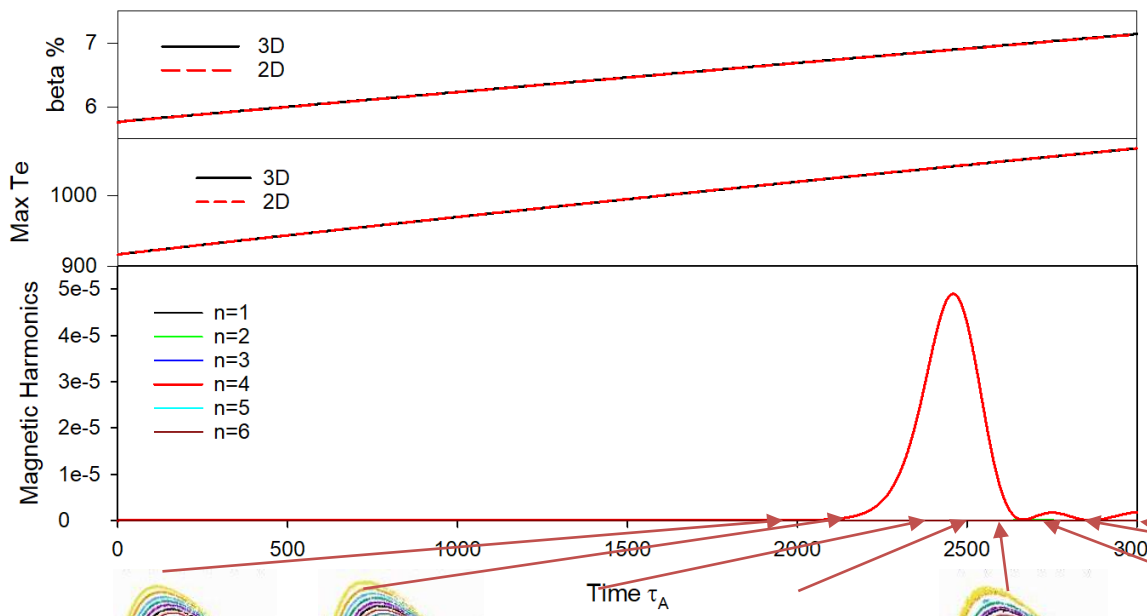

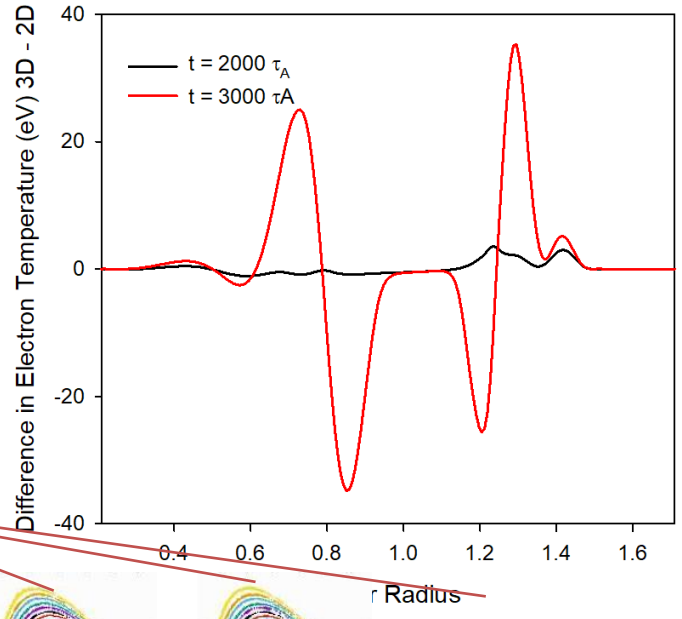

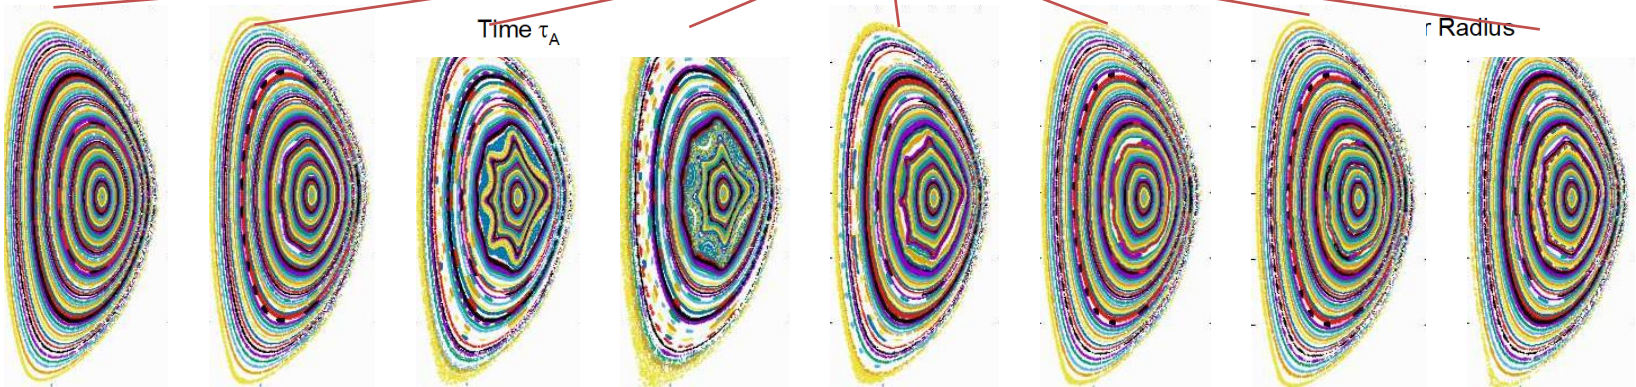

3000 < t < 6000

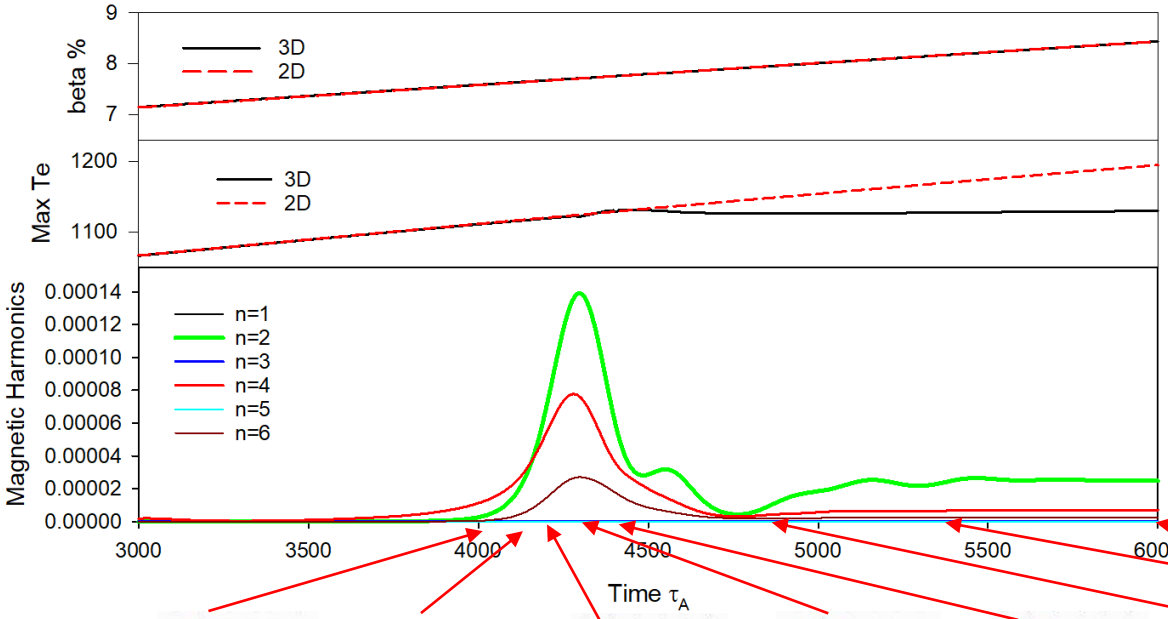
beta %
3D
2D
Max Te
3D
2D
Magnetic Harmonics
n=1
n=2
n=3
n=4
n=5
n=6
Time τ_A

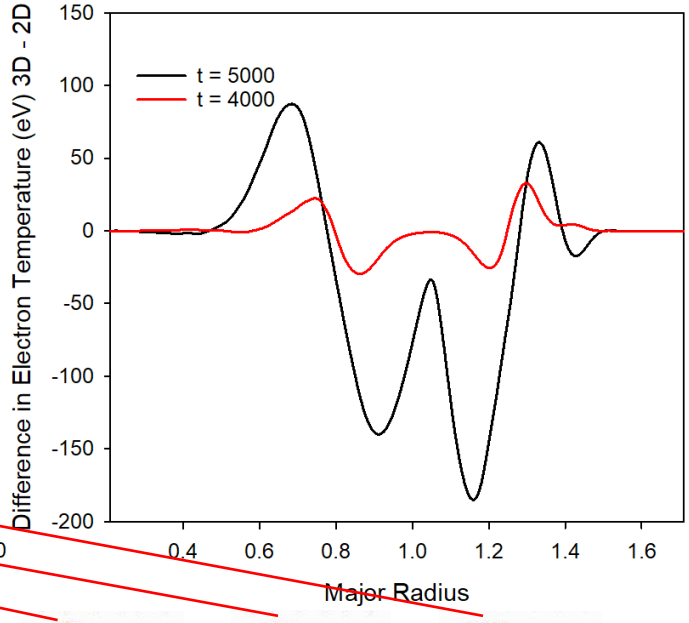
Difference in Electron Temperature (eV) 3D - 2D
t = 5000
t = 4000
Major Radius

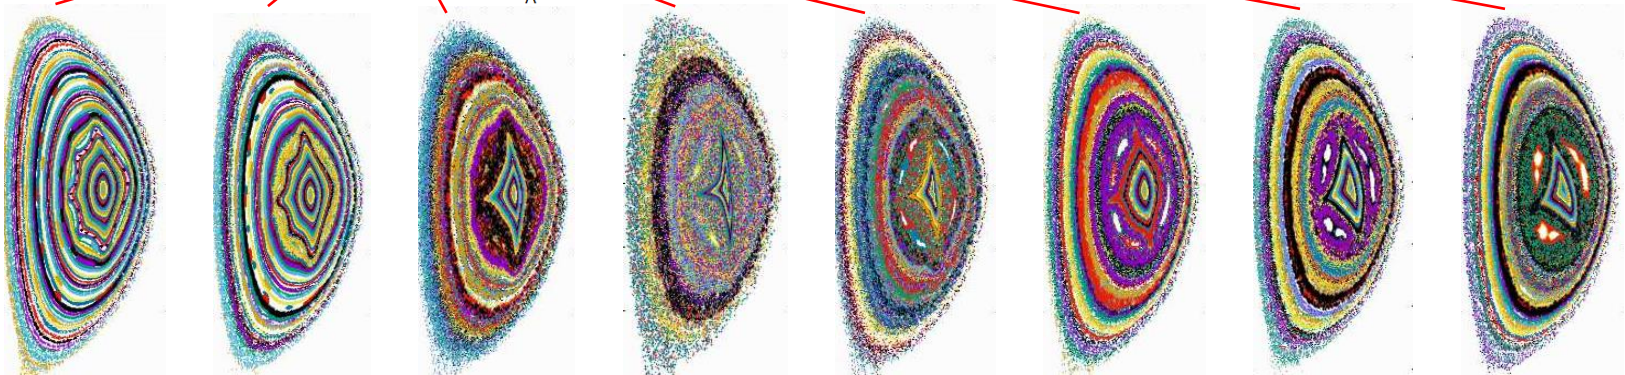

Effect of Sheared Rotation ~ 20 kHz in Center

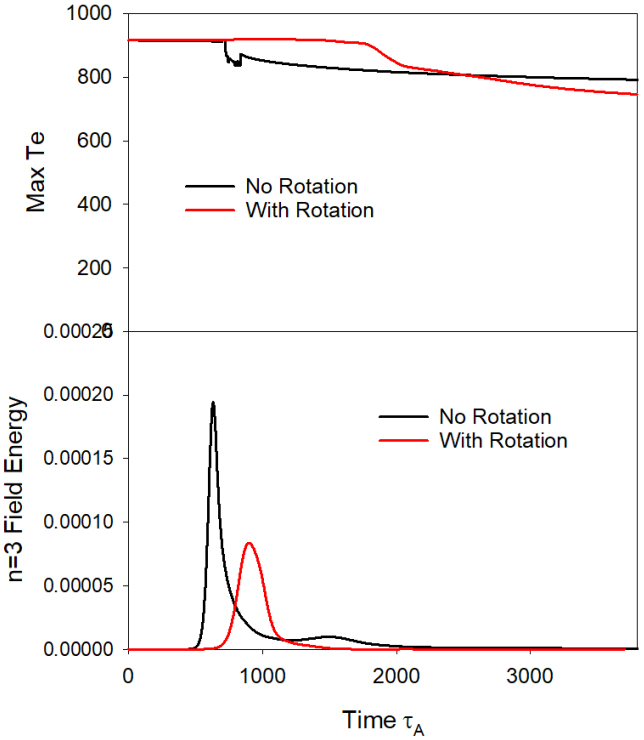

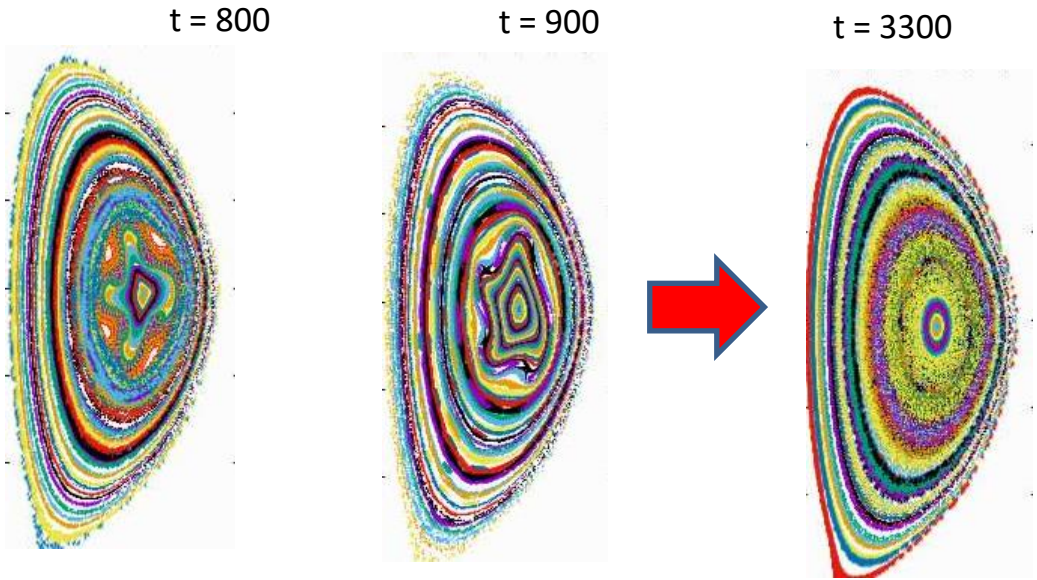

Results are similar, but instability growth rates are less and tend to symmetrize final configuration

Comparison of midplane temperatures at t = 3200 τ_A (1.90 ms)

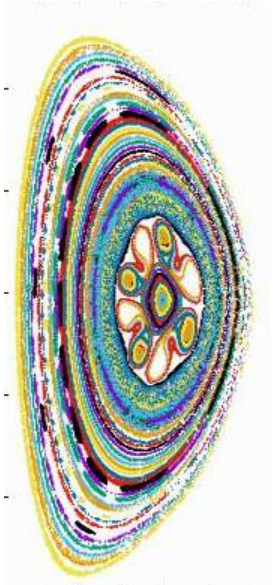

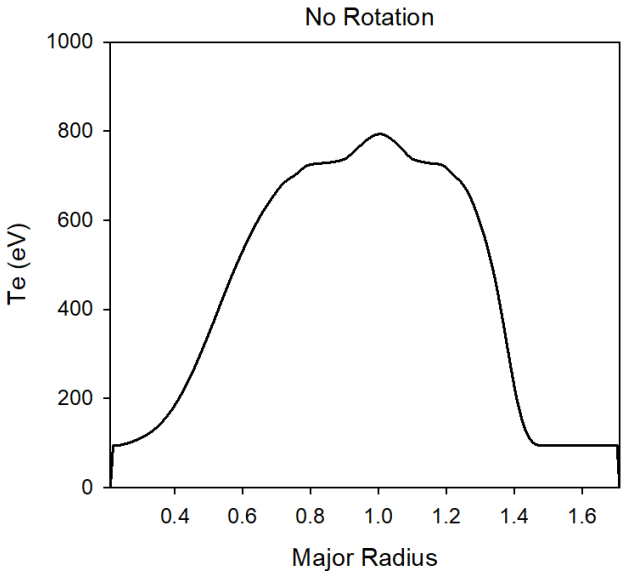

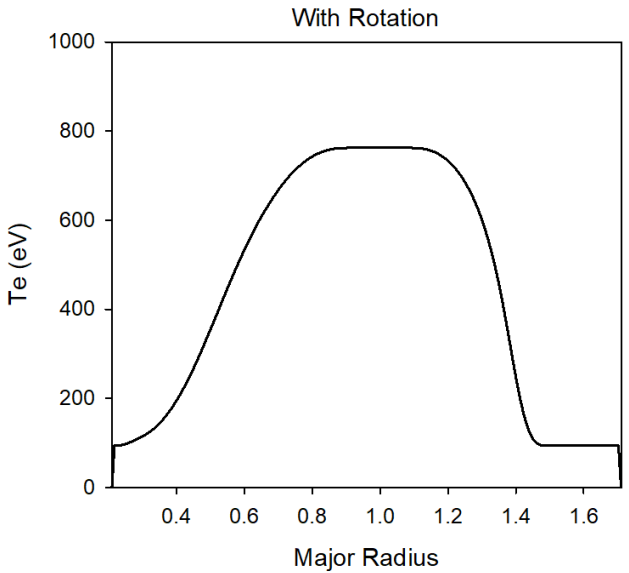

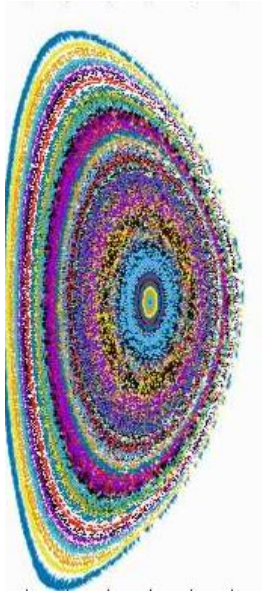

Effect of 3D on τ_E

NSTX experimentalists want to know the effect of the 3D instabilities on τ_E

To attempt to answer this, we are running a case with Beams and Transport, first in 2D and then in 3D to see the difference

S = 36.3 MW, $\kappa_\perp$ = 21.8 m^2/s, $\tau_E \sim$ 7.76 ms

Separate ion and electron temperatures

3D in progress

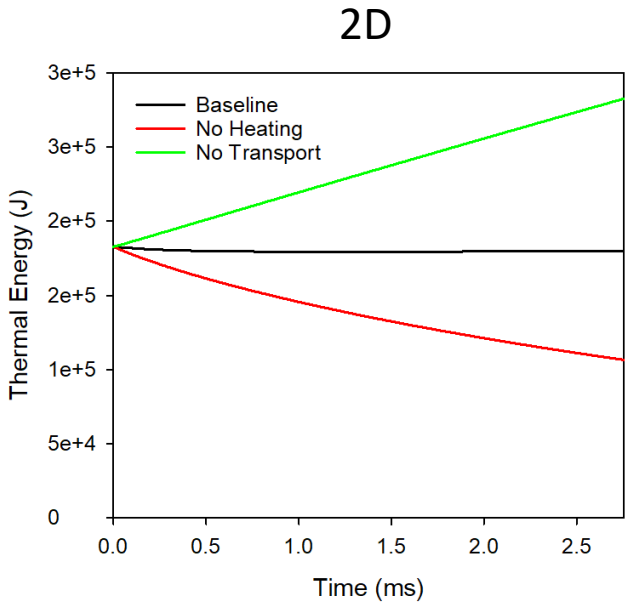

That’s All I have

Anything Else ?

Next Meeting in 3 weeks: Feb 28

Update on Soft-Beta-Limit Study

0.8 cm

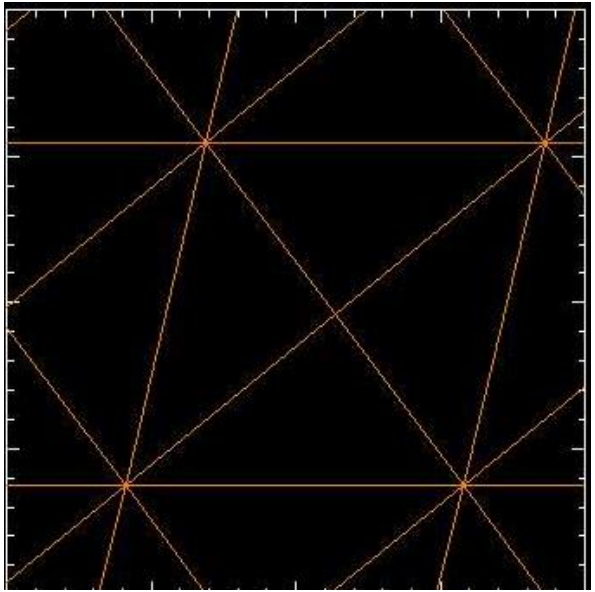

0.4 cm

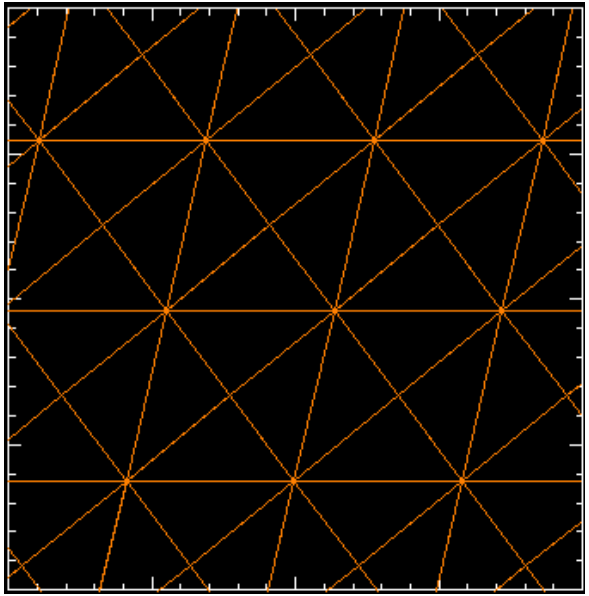

0.2 cm

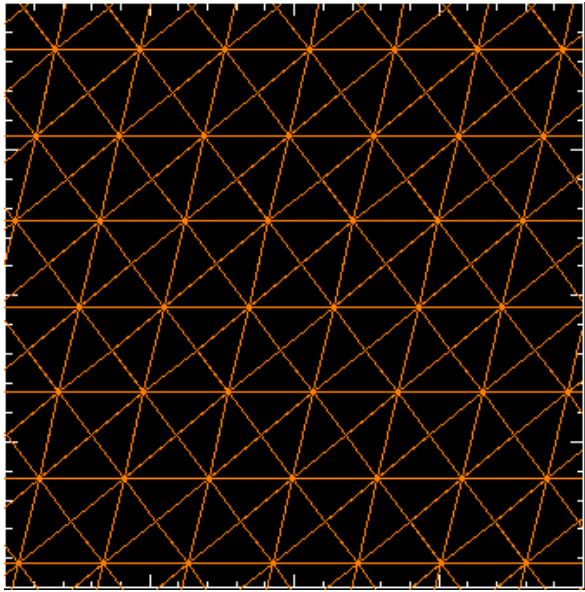

These are close-ups in center of grid (near magnetic axis)

Result of Convergence Study

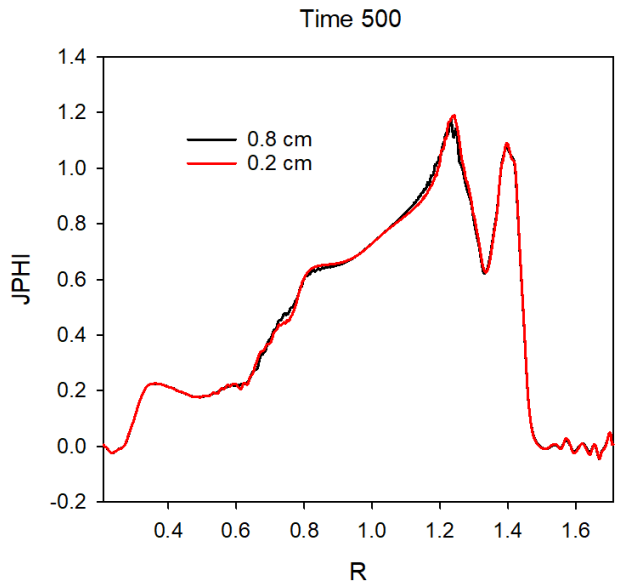

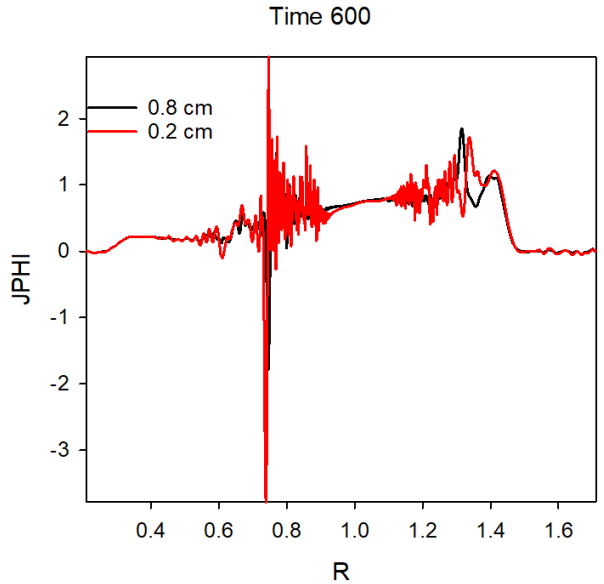

Solution for jphi still very noisy in region 0.5 < R < 0.9, even for the finest grid with 0.2 cm in center

Grid was not refined where J gets jagged

0.8 cm

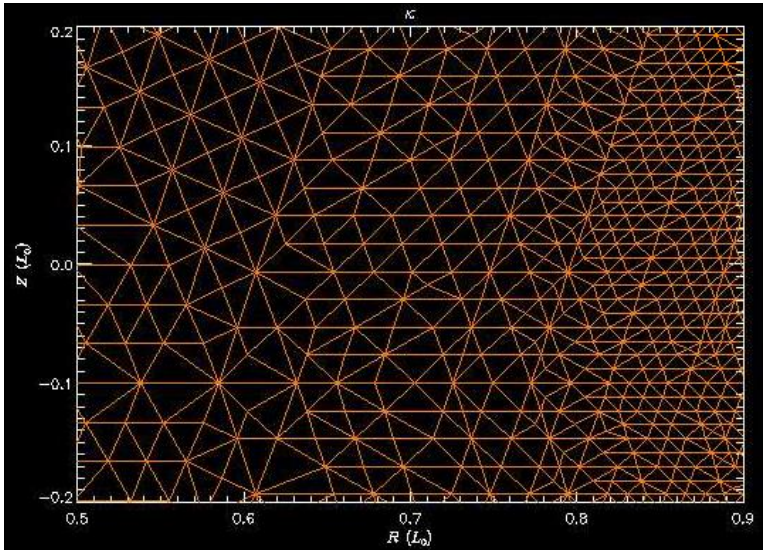

0.2 cm

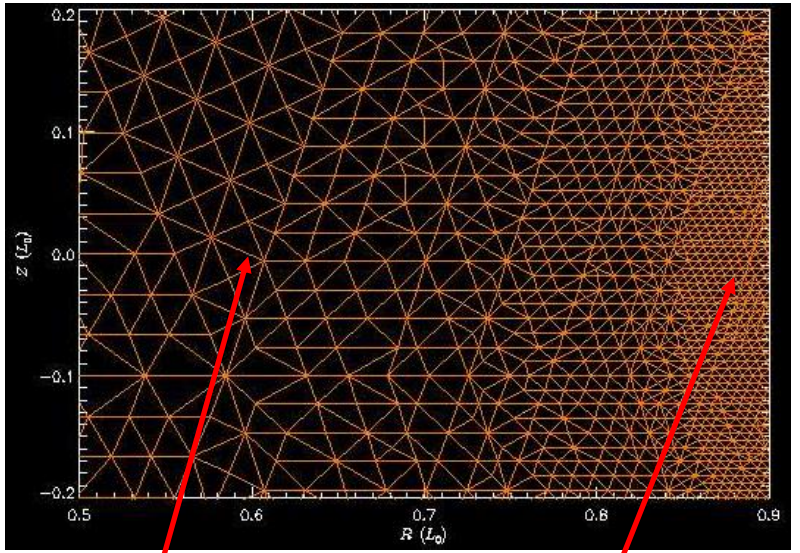

Resolution increased here

But not here

Now producing better grids that are refined where the current gets jagged