M3D-C1 ZOOM Meeting

08/23/2021

Announcements

CS Issues

1. Intel-MPI on stellar
2. Mesh adaptation update
3. NERSC Time
4. Changes to github master since last meeting
5. Regression tests
6. Plans for optimizing the matrix assembling on GPUs – Chang Liu

Physics Studies

1. Energy conservation with itemp=0,1, ipres=0,1 -- Lyons
2. Jump in profiles in first timestep in benchmark run --Lyons
3. Velocity on open-field-lines in benchmark run -Lyons

Note: meeting minutes posted on m3dc1.pppl.gov

In attendance

Steve Jardin
Hank Strauss
Mark Shephard
Jin Chen
Adelle Wright
Nate Ferraro
Chen Zhao
Brendan Lyons
Chang Liu
Seegyoung Seol
Cesar Clauser
Usman
Morteza SIboni

Announcements

- No meeting next week, Aug 30: Possible meeting Tues Sept 7 (9/6 is a holiday)

- /p/tsc to be upgraded during Sept. maintenance period
  - 20 times faster and greatly expanded size

- John Mandrekas requested a 90 min presentation from all SciDACs
  - Progress over last 4 years
  - Plans for next year
  - Synergy between Physics and CS teams
  - Most important unsolved problems in our area

- APS Nov 8-12
  - Meeting will be IN PERSON with virtual option
  - M3D-C1 Invited talks by C. Liu, A. Wingen
  - CTTS meeting? Will DOE allow travel?

- EPS 6/27 – 7/1 2022 in Maastricht, Netherlands
  - Nominate invited speakers by 29 October 2021

stellar.princeton.edu

8/12/20: The default compilation has been changed to use intel-mpi instead of openmpi

Note : For Poincare plots and q-profiles with the new modules**,** you need to copy an updated version of “trace” from /home/nferraro/fusion-io/bin (see NEWDOC)

Mesh Adaptation Update

RPI?

NERSC Time

mp288

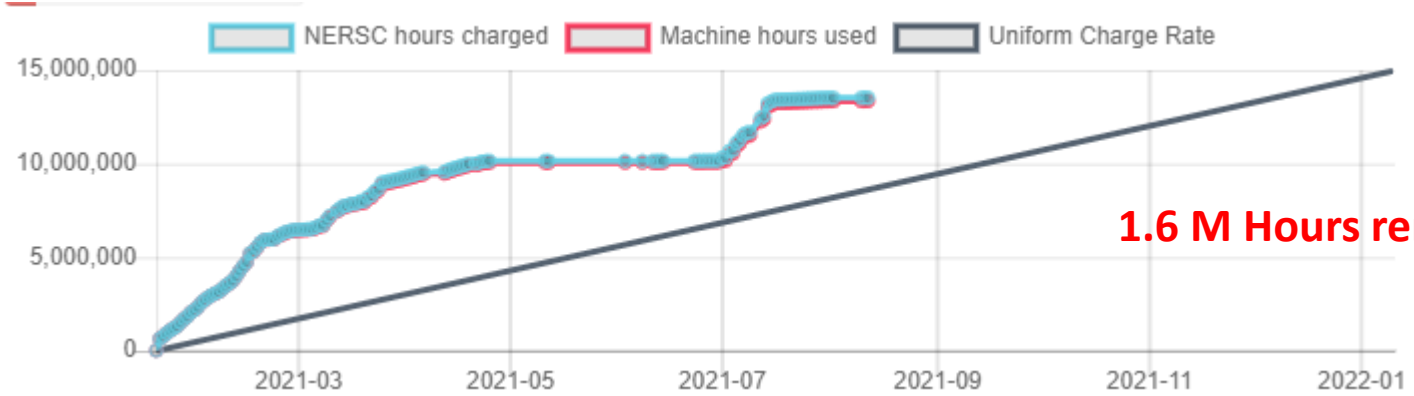

1.6 M Hours remaining!

- mp288 received 10M Hrs for CY 2021
- Initial allocation exhausted by May 1
- John Mandrekas (DOE) added 5M Hrs additional
- More time may be possible if this is exhausted
- Pearlmutter time will not be charged for this FY
- We are NESAP Tier 2. Machine not yet ready. Phase-I w GPUs

Changes to github master since 08/09/21

S. Jardin:
8/13/21: multiply kappar by [1+(1-pefac)/pefac] for ipres=0, numvar.ge.3

Seegyoung Seol:
8/13/21: adding solution transfer in adaptation unit test

Nate Ferraro:
8/12/21: changed the RMP regtest on stellar back to mumps, since superlu was failing
8/12/21: Changed stellar makefile to intel-mpi
8/12/21: updated README/readme.stellar and removed extra stellar makefile

NOTE: password authentication no longer valid after 08/13/21. Must use personal access token.

Local Systems

- PPPL centos7(08/23/21)
  - 4 regression tests **PASSED** on centos7:
  - RMP_nonlin failed (ekin 5.4056e-06 -> 5.4225e-06)
  - ADAPT failed (need to rebaseline?)
- PPPL greene (08/23/21)
  - 3 regression tests **PASSED on greene (m3dc1)**
  - RMP_nonlin failed (ekin 5.4056e-06 -> 5.4225e-06)
  - ADAPT failed (need to rebaseline?)
- STELLAR (08/23/21)
  - 5 regression tests **PASSED** on stellar
  - RMP_nonlin failed (ekin 5.4056e-06 -> 5.4225e-06)
- TRAVERSE(03/29/21)
  - Code compiles
  - Regression test failed: split_smb not found in PATH

Other Systems

- Cori-KNL (8/23/2021)
  - 4 regression tests **PASSED** on KNL
  - RMP_nonlin failed (ekin 5.4056e-06 -> 5.4225e-06)
  - ADAPT failed (need to rebaseline?)
- Cori-Haswell (8/23/2021)
  - 4 regression tests **PASSED** on cori
  - RMP_nonlin failed (ekin 5.4056e-06 -> 5.4225e-06)
  - ADAPT failed (need to rebaseline?)
- MARCONI
  - All regression tests PASSED on MARCONI (J. Chen, 9/04/20)

- **We need to re-baseline RMP_nonlin and ADAPT**

Progress on optimizing the matrix assembling on GPUs

Chang Liu email 8/20/21 (summary)

- M3D-C1 matrix assembling has been partially ported to GPU. Tested for a numvar=3 nonlinear run in toroidal geometry. Results exactly the same as CPU run

- Some terms (not used in the test) still need to be converted: gyro viscosity, parallel viscocsity, two-flouid, heating and cooling, temperature equation

- Plan for the Future:
  - Phase 1: can obtain GPU version for traverse or Perlmutter by compiling with ACC=1. People encouraged to compare results with CPU version
  - Phase 2: Once thoroughly tested, make GPU version the default (also runs on CPU)

- Coding details in the document "proposal of a new method for matrix element calculation in M3D-C1 with GPU optimization"

Discrepancies in Heat Flux Diagnostic

by

Brendan C. Lyons

August 2nd, 2021

GENERAL ATOMICS

Conclusions

- Something is off with the `flux_heat()` diagnostic
- Nature of discrepancy is dependent on the `ipres` and `itemp`
- There may be something wrong about how the heat flux itself, not just the diagnostic, is implemented in the code for `itemp=0`

Problem Solved !

1. Do not use itemp=0 if kappar .ne. 0

2. To make itemp=1, ipres=1 and itemp=1, ipres=0 identical

for ipres=0, replace (in temperature_lin)

$$\text{kappar} \rightarrow \left[1+\frac{(1-\text{pefac})}{\text{pefac}}\right]\text{kappar}$$

3. To get more precise energy balance for large kappar, may need to increase spacial resolution in SOL:

4. I committed the change but it makes one of the regression tests fail (RMP_nonlin)

Results with increased resolution

Low resolution (original)

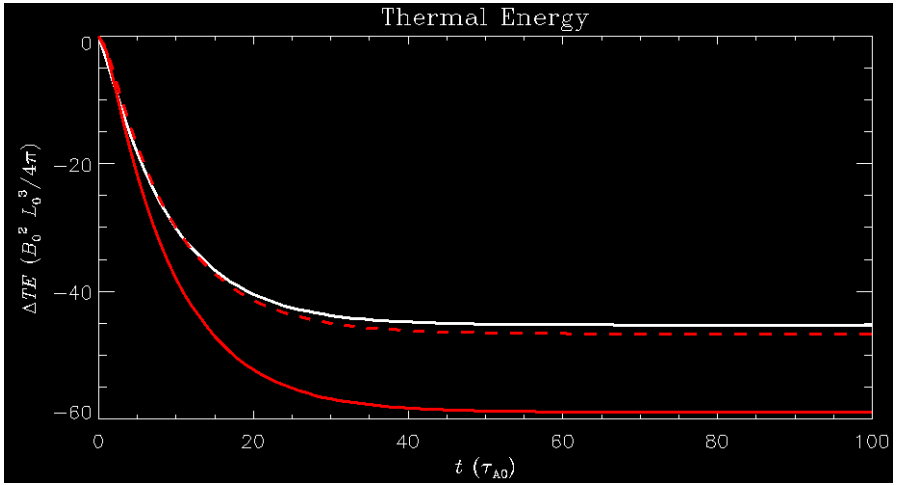

High resolution (#zones x 4)

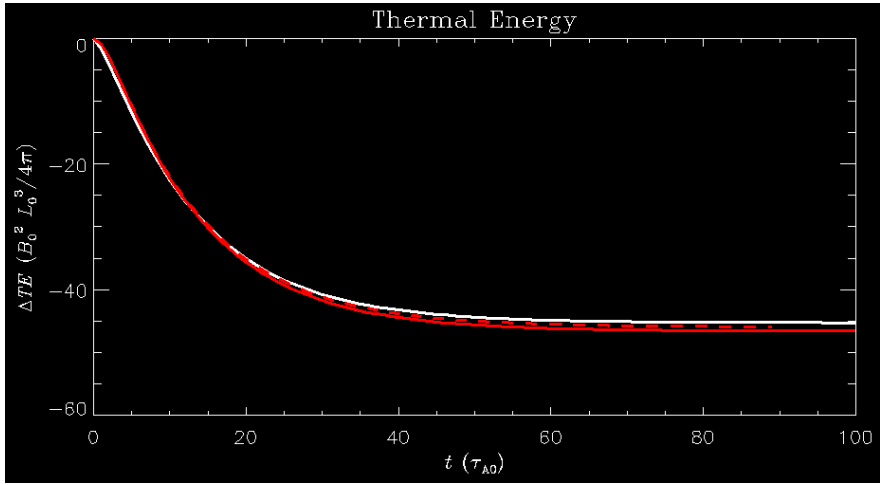

Odd Initial Temperature Relaxation

by
Brendan C. Lyons

August 6th, 2021

GENERAL ATOMICS

Conclusions

- **Something causes fast relaxation of pressure on grid scale in first time step**
- **Improved but not solved by increased resolution**
- **Noisy parallel conduction to blame?**
  - Surface averaged noise leads to net temperature change?
  - Initial temperature not a flux function?

Runs on Stellar

/scratch/gpfs/bclyons/C1_42891:
    10 steps with pellet
/scratch/gpfs/bclyons/C1_42893:
    10 steps without pellet
/scratch/gpfs/bclyons/C1_42944:
    10 steps with pellet at 1 us
/scratch/gpfs/bclyons/C1_44084
/scratch/gpfs/bclyons/C1_44248
/scratch/gpfs/bclyons/C1_44300:
    100 steps with pellet at 0.1 us
/scratch/gpfs/bclyons/C1_44832
    10 steps at 1 us with gEQDSK pressure
/scratch/gpfs/bclyons/C1_44836
/scratch/gpfs/bclyons/C1_45251
/scratch/gpfs/bclyons/C1_45254:
    100 steps at 0.1 us with gEQDSK pressure

Mid-resolution Mesh (44)
/scratch/gpfs/bclyons/C1_45365:
    10 steps at 1 us with gEQDSK pressure
/scratch/gpfs/bclyons/C1_45368:
    Turn off ohmic heating, little change
/scratch/gpfs/bclyons/C1_45718:
    16 planes, no change *deleted*

High-resolution Mesh (45)
/scratch/gpfs/bclyons/C1_45813
    No ohmic heating; better than 45368

Analysis: Initial configuration

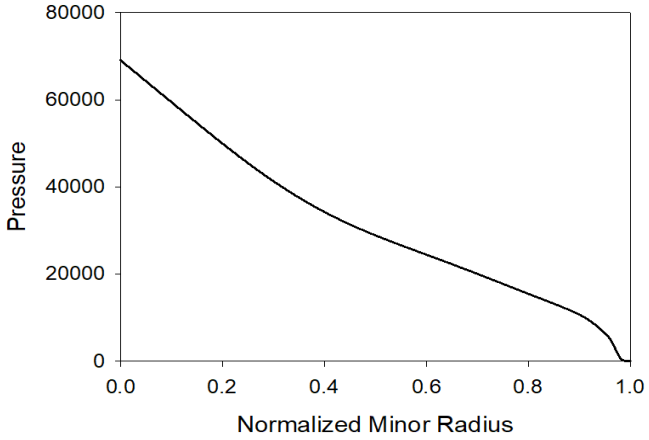

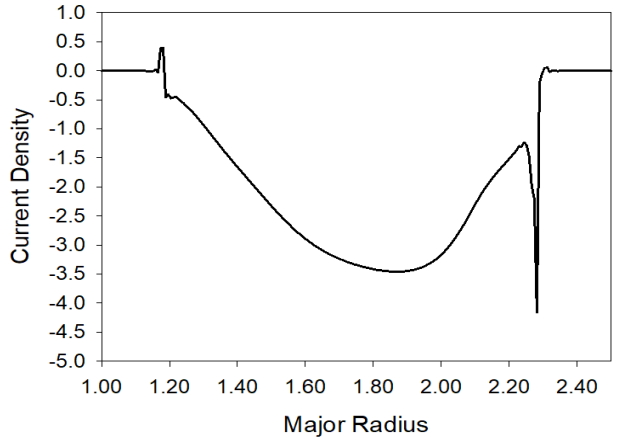

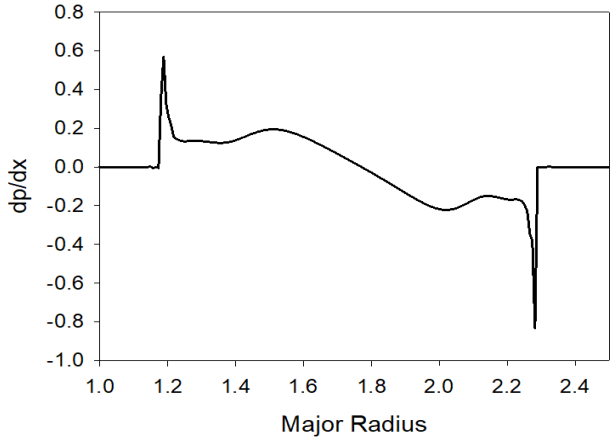

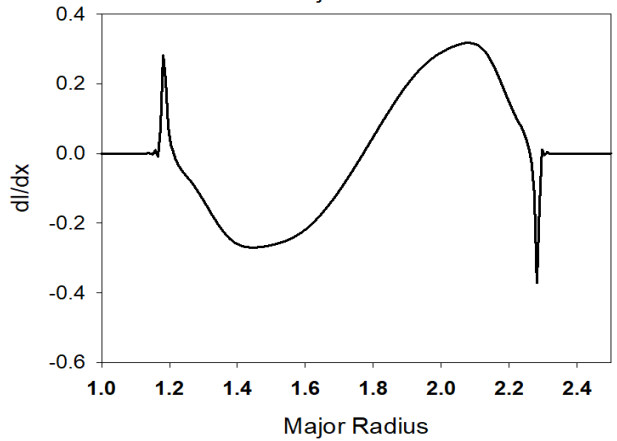

- Large narrow spike in current near edge (bootstrap ?)
- Large narrow spikes in dp/dx near edge (H-mode)

Change in first timestep

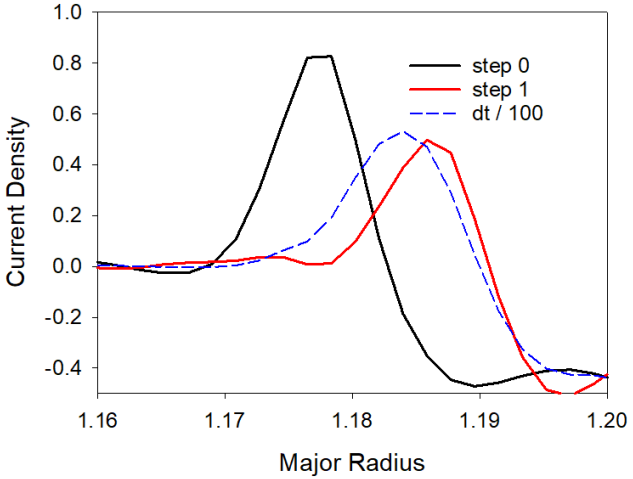

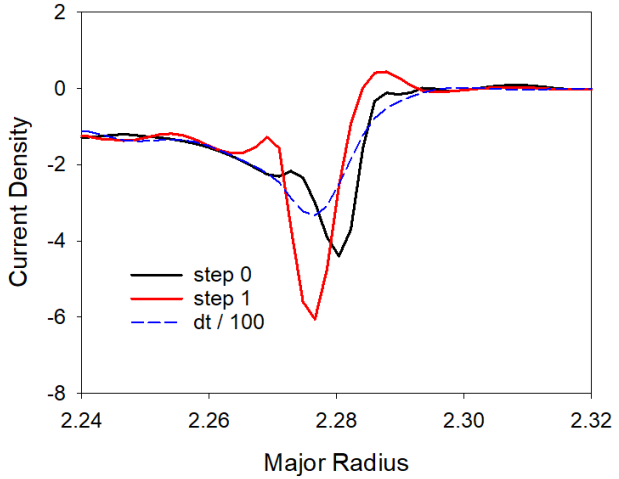

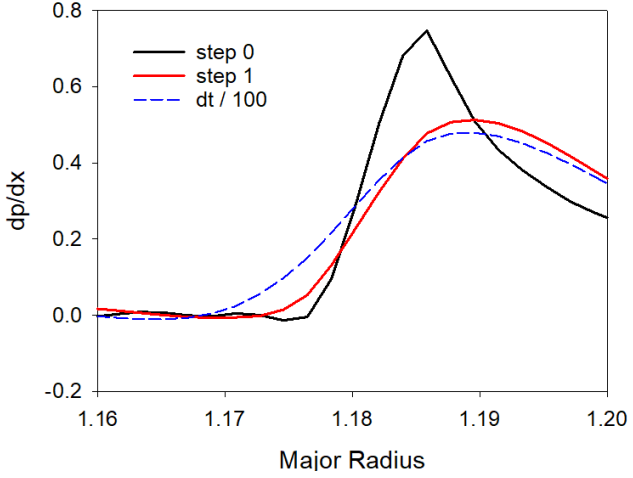

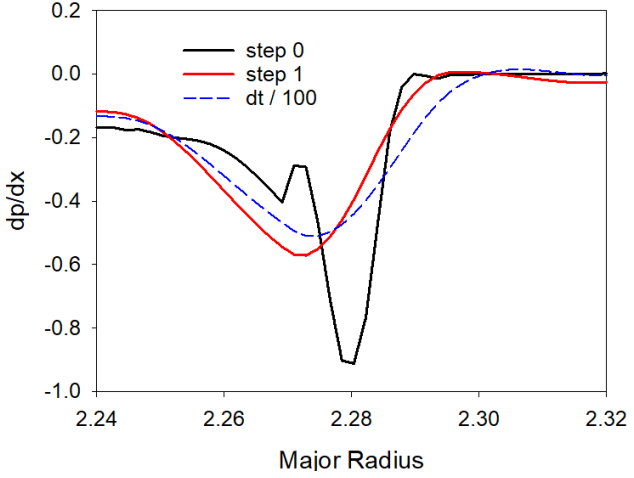

Significant changes in current spikes on first timestep

Less change with dt reduced by 100 (and 100 timesteps) but still significant

Also large change in dp/dx peaks in first timestep

Somewhat reduced with smaller dt

Effect of $\kappa_{||}$

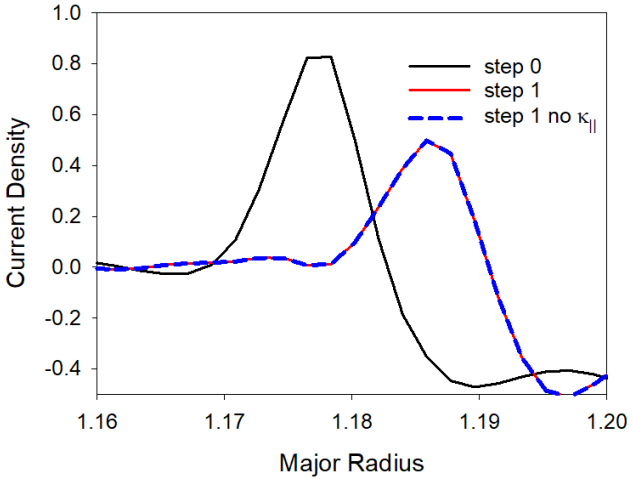

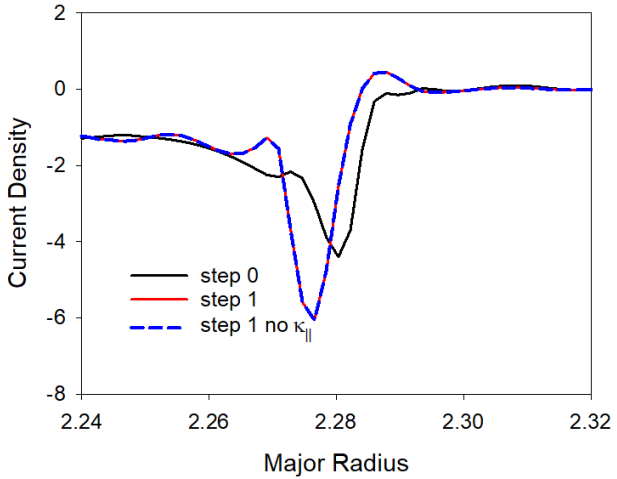

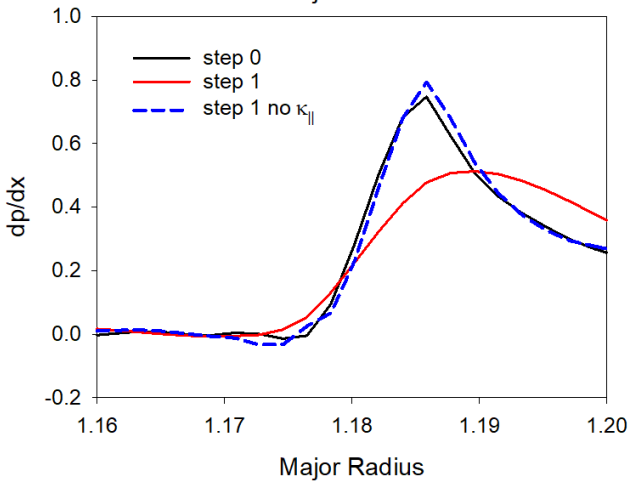

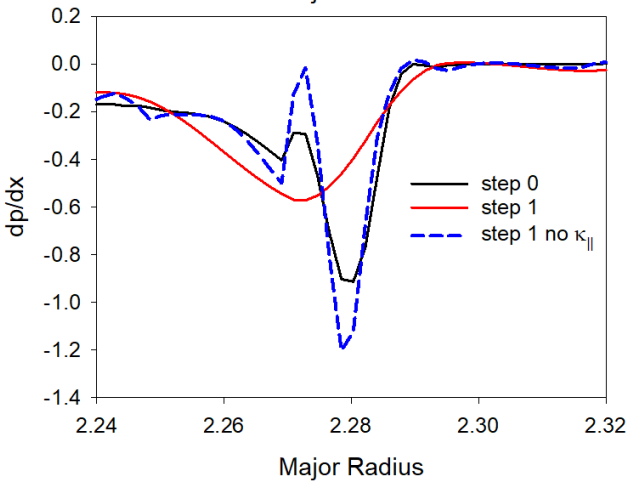

- Setting kappar=0 does not affect the jump in jphi for the first timestep but greatly reduces the change in dp/dx
- Why does jphi change on the first timestep?

Effect of η and V

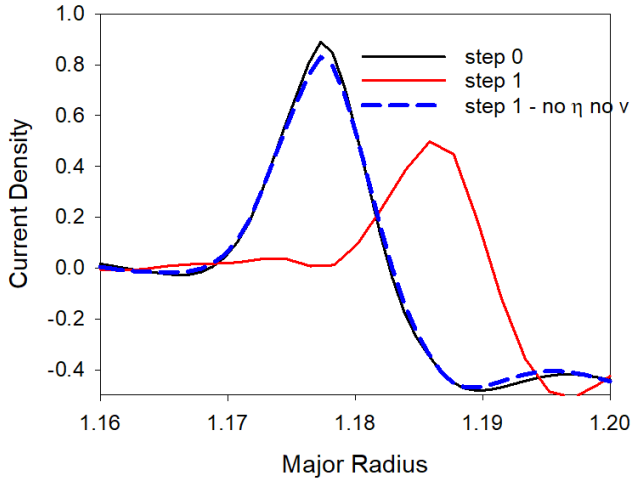

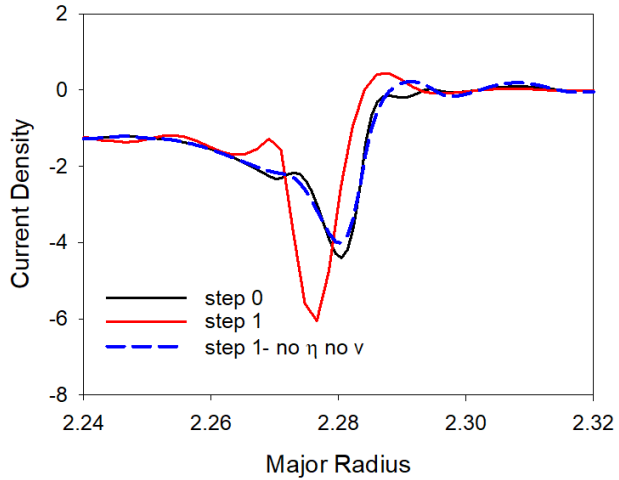

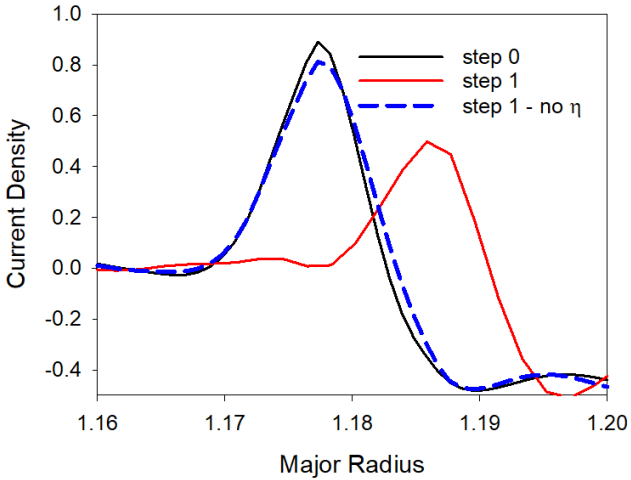

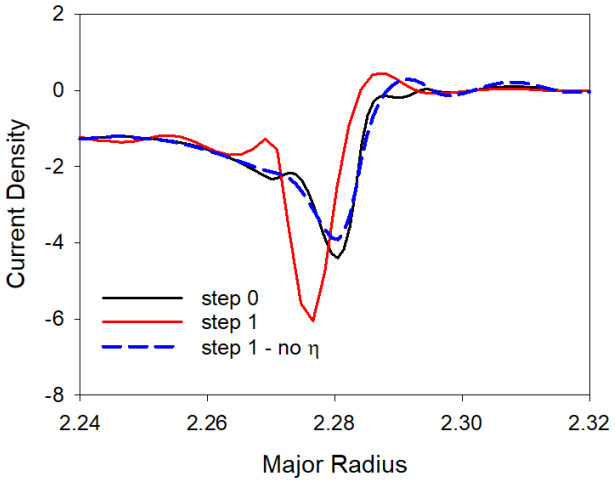

- Setting η=0 eliminates the jump in J in the first time-step
- Setting V=0 has very little additional effect

Effect of increasing Tedge

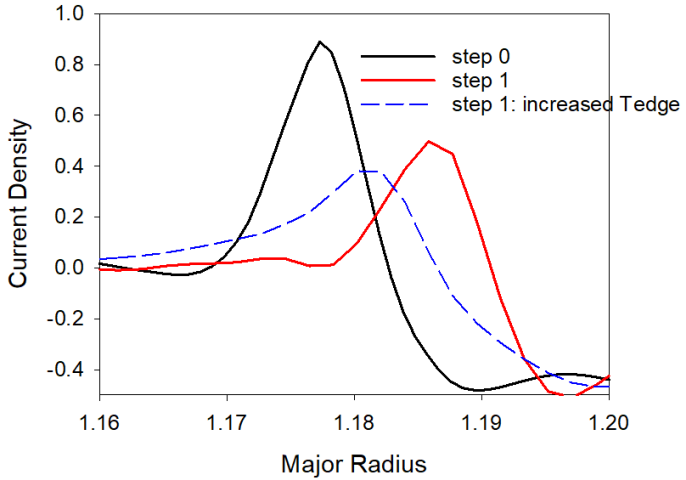

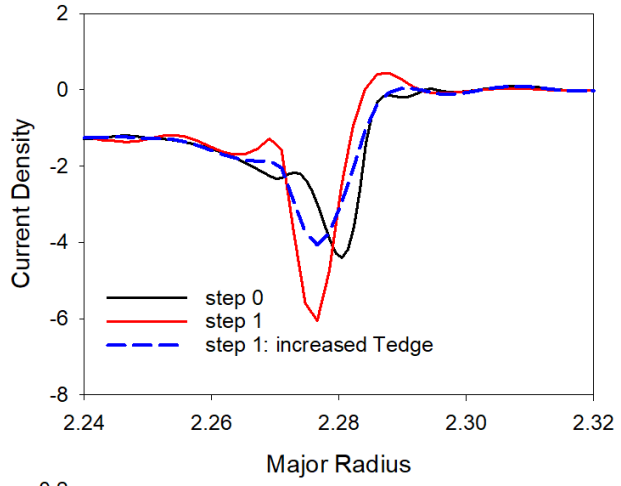

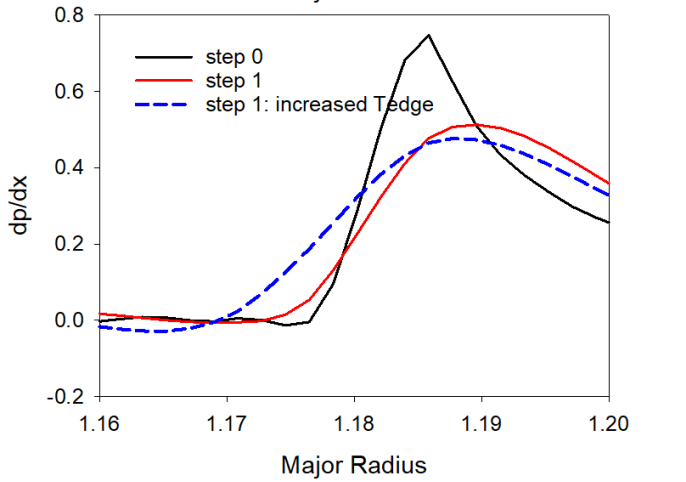

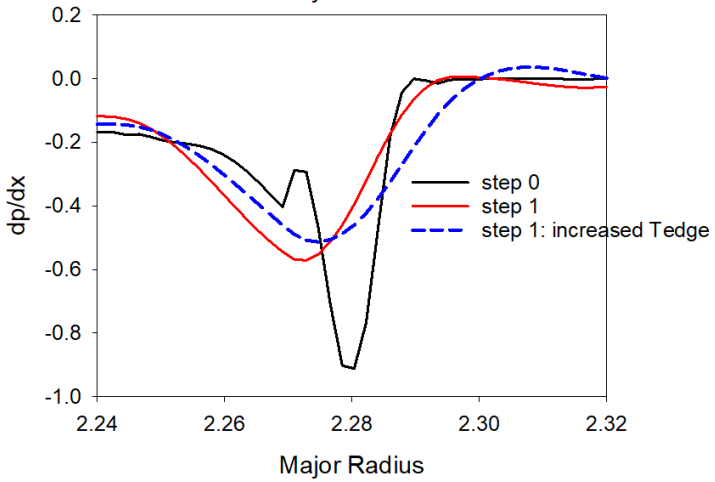

- Increasing Tedge from 0.5 eV to 10 eV lessons the jump in J the first timestep
- Increasing Tedge has very little effect on the jump in dp/dx the first timestep

More on no η no v (and no hyper)

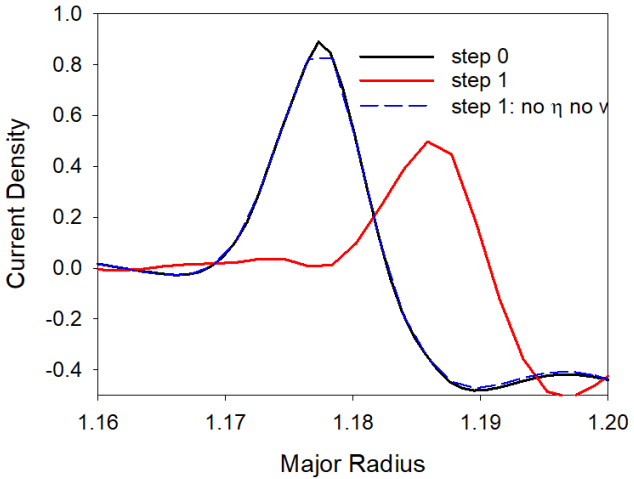

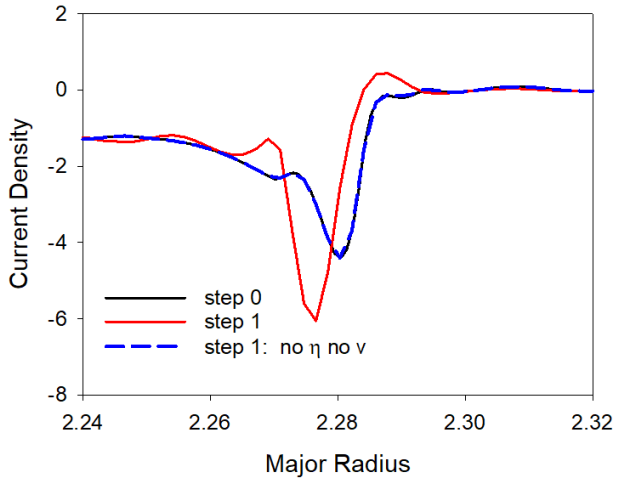

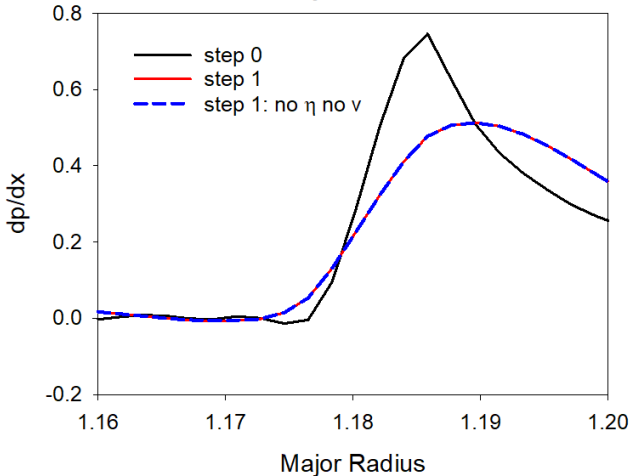

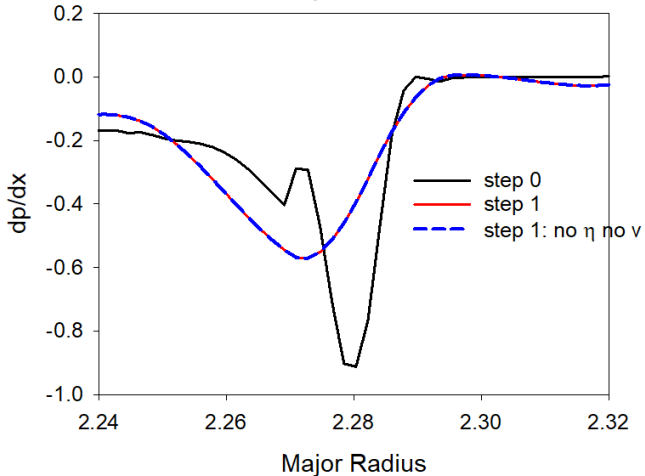

- If we zero out η, V, and hyper, J will not change the first time-step
- However, this has almost no effect on the pressure evolution

Summary of 1st time-step jump

- Jump in J and in p are seemingly unrelated
- Jump in J caused by large η
- Jump in p caused by large $\kappa_{||}$
- Changes somewhat reduced by reducing dt by 100 and running for 100 time-steps, but still significant
- Likely causes of jump in cycle 1:
  - Source terms missing, eg: $\mathbf{E} + \mathbf{V} \times \mathbf{B} = \eta(\mathbf{J} - \mathbf{J}_{BS})$
  - Thermal conductivity profile not consistent with H-mode profiles
  - Not enough spatial resolution for these extreme equilibrium profiles
- These should be circumvented by using eqsubtract=1

Results for eqsubtract=1

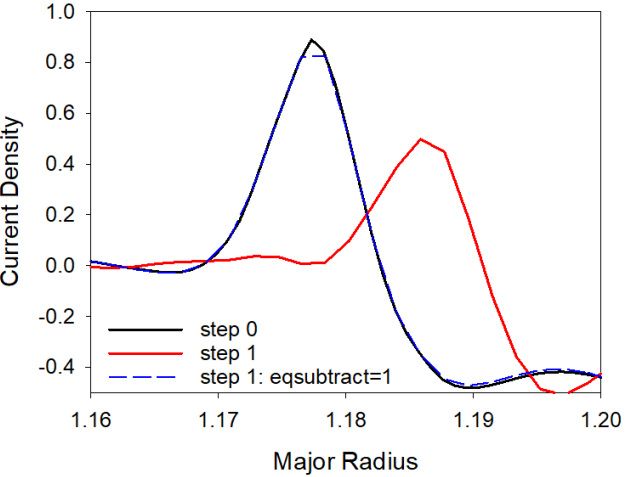

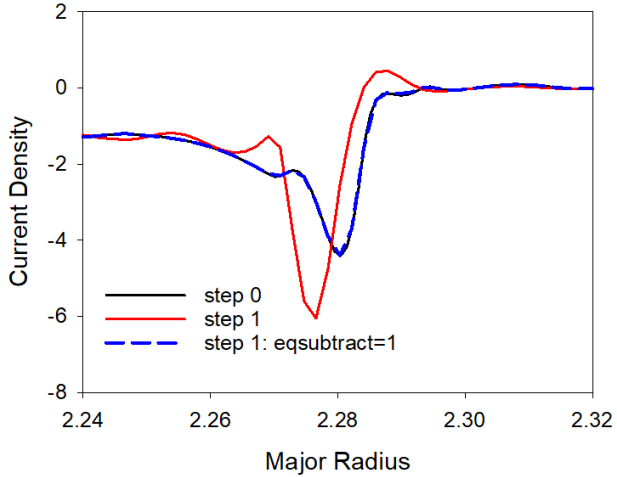

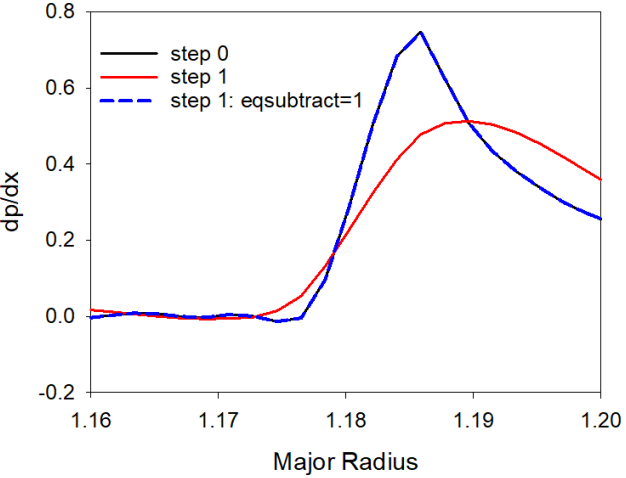

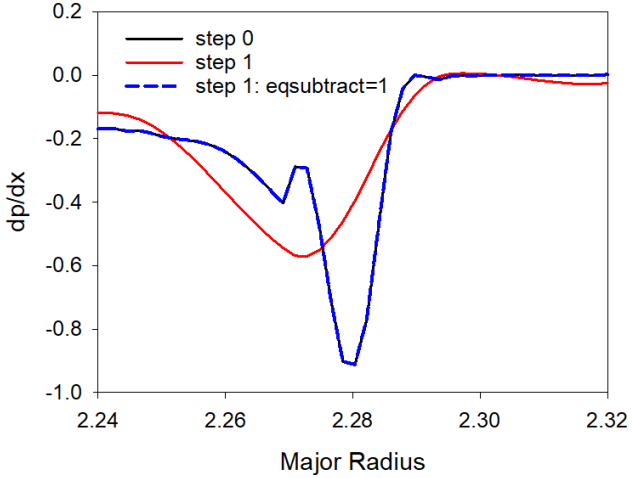

- With eqsubtract=1 both the current profile and the pressure profile change very little in the first timestep
- This is likely what NIMROD does (should check)
- **Problem solved!?**

Brendan #3 … flow on open-field-lines

Email 8/13/21:
One of the most persistent discrepancies that we observe in the NIMROD/M3D-C1 benchmarks is the flow in the open-field-line region (OFLR). NIMROD sees almost none, whereas as M3D-C1 sees large flows in the opposite direction of that seen in the pellet [plasma?]

See /scratch/gpfs/bclyons/C1_45331

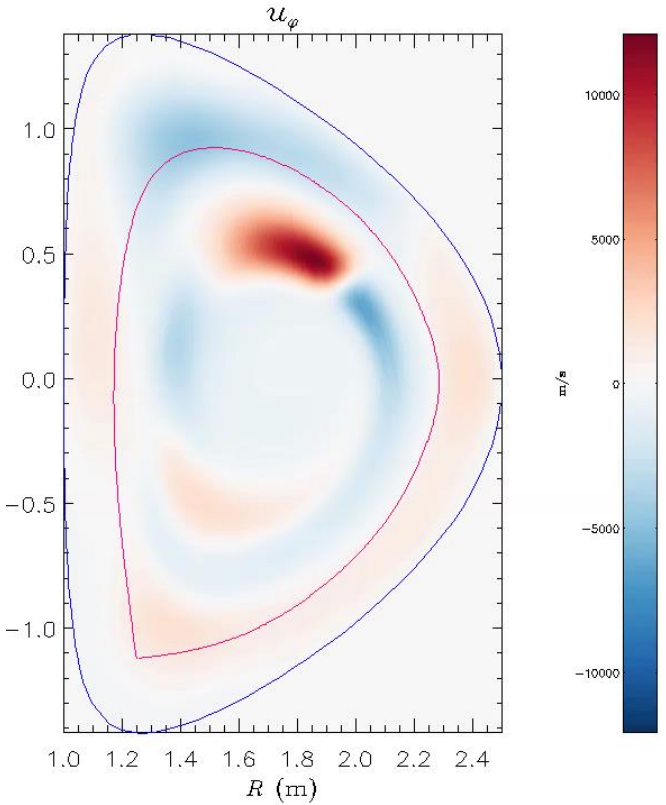

Uφ midplane and M3D-C1 2D plot

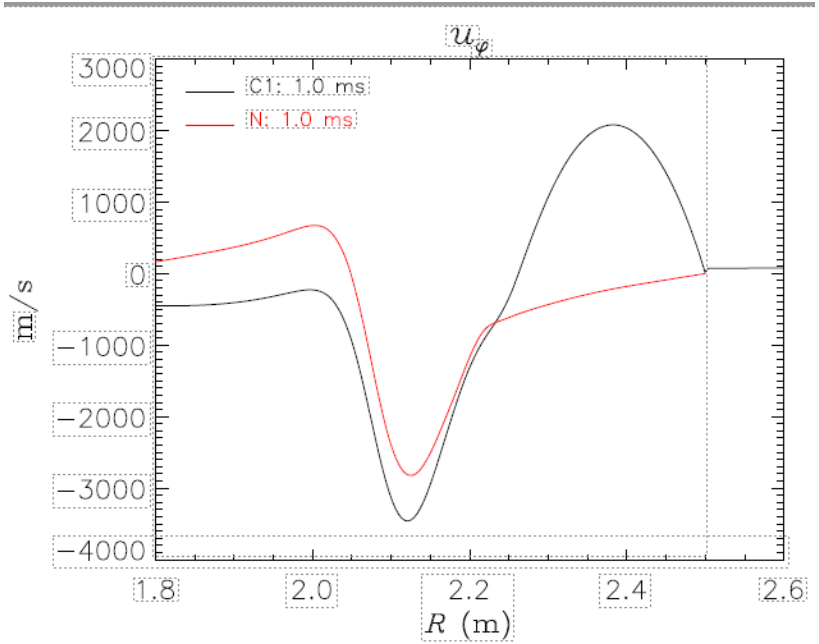

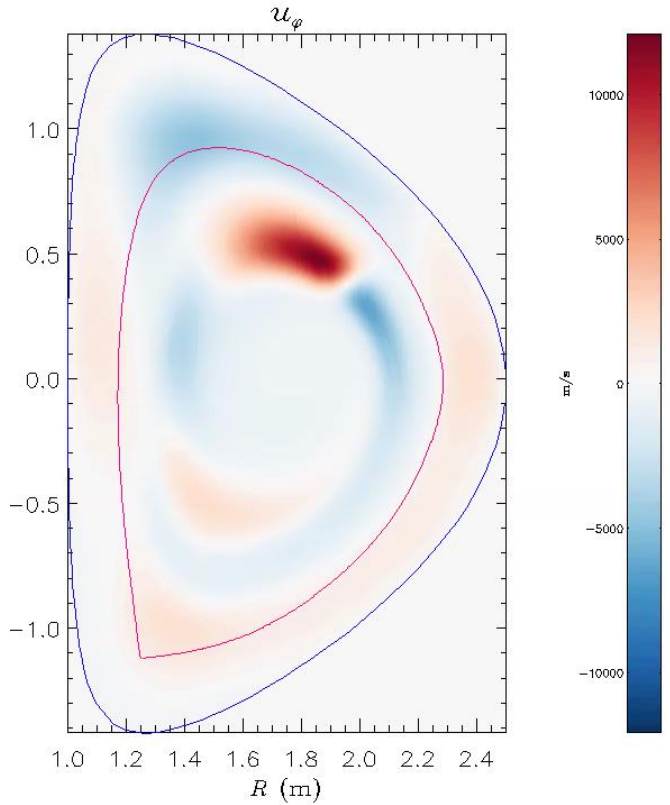

U_R and U_Z midplane

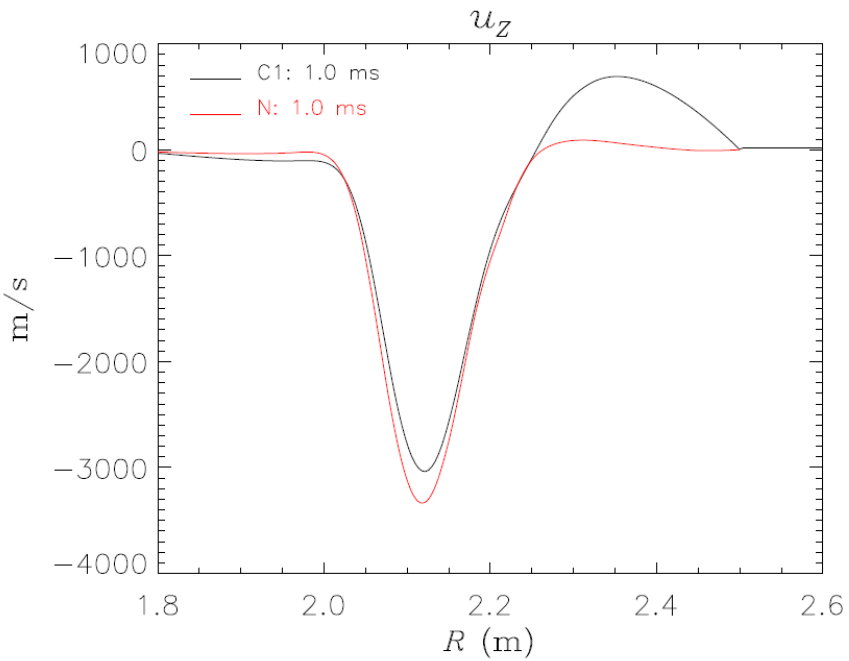

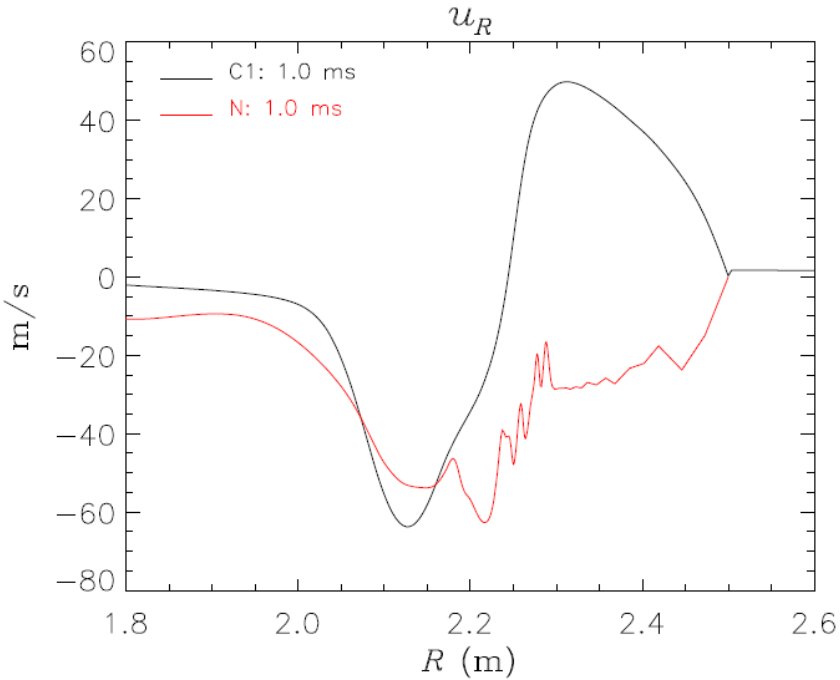

Flows in 2D and3D very similar

2D

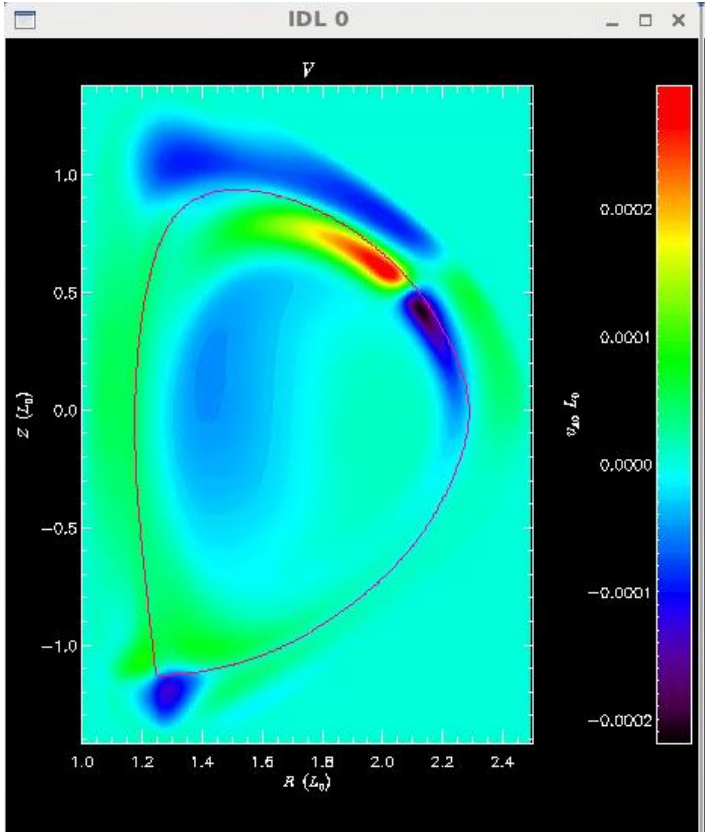

3D

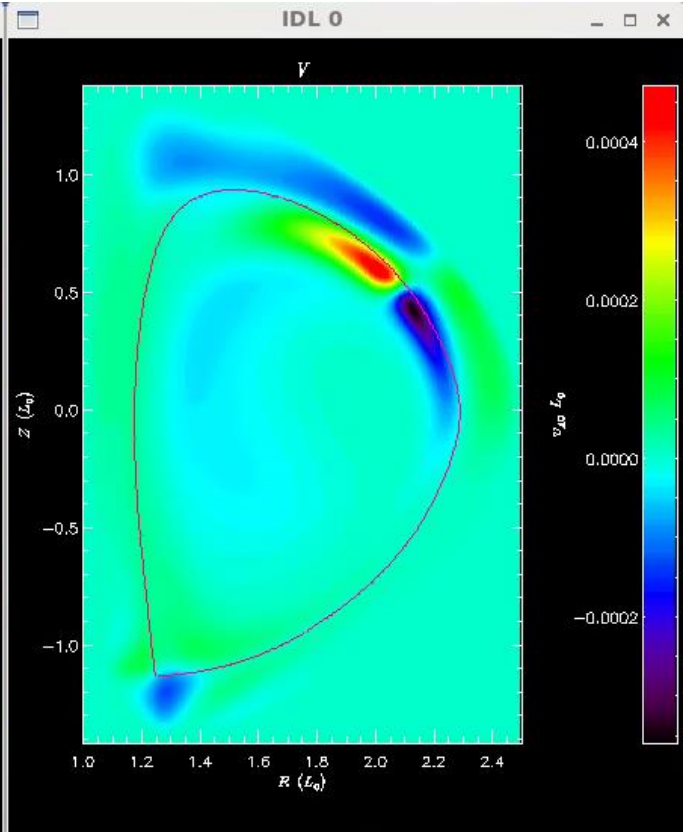

Pfirsch-Schluter flows driven by η and p’

$$\left.\nabla p = \vec{J}\times\vec{B}\right\} \quad \vec{J} = \left(\frac{\vec{B}\times\nabla p}{B^2}\right) + J_{\parallel}\frac{\vec{B}}{B}$$

$$\left.\begin{aligned}\vec{E}+\vec{V}\times\vec{B} &= \eta\vec{J}\\ \frac{\partial\vec{B}}{\partial t} &= -\nabla\times\vec{E} = 0\end{aligned}\right\} \quad -\nabla\Phi + \frac{1}{2\pi}V_L\nabla\varphi + \vec{V}\times\vec{B} = \eta\vec{J}$$

$$\vec{B} = \nabla\varphi\times\nabla\psi + g(\psi)\nabla\varphi$$

$$J \equiv \left[\nabla\psi\times\nabla\theta\bullet\nabla\varphi\right]^{-1}$$

$$\left.\nabla\bullet\vec{J} = 0\right\} \quad \frac{\partial}{\partial\theta}\left(\frac{J_{\parallel}}{B}+\frac{gp'}{B^2}\right) = 0\Bigg\} \quad \left(\frac{J_{\parallel}}{B}+\frac{gp'}{B^2}\right) = f(\psi)$$

$$\vec{V}\bullet\nabla\psi = -\eta p'\left[\frac{|\nabla\psi|^2}{B^2} + g^2\left(\frac{1}{B^2}-\frac{1}{\langle B^2\rangle}\right)\right] - \frac{1}{2\pi}V_L\left(1-\frac{\langle B_T^2\rangle}{\langle B^2\rangle}\right)$$

$$= -\eta p' R^2\left(1-\frac{B_T^2}{\langle B^2\rangle}\right) - \frac{1}{2\pi}V_L\left(1-\frac{\langle B_T^2\rangle}{\langle B^2\rangle}\right)$$

$$\nabla U\times\nabla\varphi\bullet\nabla\psi = -\eta p'\left(1-\frac{B_T^2}{\langle B^2\rangle}\right) - \frac{1}{2\pi R^2}V_L\left(1-\frac{\langle B_T^2\rangle}{\langle B^2\rangle}\right)$$

Comments

- Is density diffusion subtracted from momentum equation?

$$\frac{\partial}{\partial t}\left(n\mathbf{V}\right) \rightarrow n\frac{\partial \mathbf{V}}{\partial t} + \mathbf{V}\frac{\partial n}{\partial t}$$

- Is momentum conserved in toroidal direction? Should it be?

- Can we ask NIMROD to run the 2D case and check its scaling with η and p'?

That’s All I have

Anything Else ?

Resistive Wall Mode in Periodic Cylinder

H. Strauss, 8/9/21

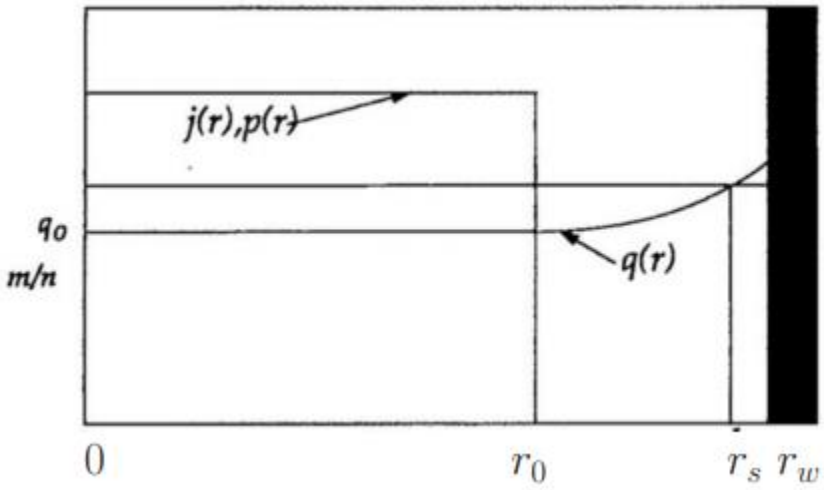

RWM

$$\gamma\tau_{wall} = -m\frac{1-(m-nq)}{1-(m-nq)-(r_0/r_w)^{2m}}$$

RWTM

$$\gamma\tau_A = \frac{c_0}{S^{1/3}S_{wall}^{4/9}}$$

$$c_0 = 2.46\left(\frac{q'r_s}{q}\right)^{2/9} f^{4/9} = 2.46 f^{4/9}$$

$$f = \frac{(r_s/r_w)^{2m}}{[1-(r_s/r_w)^{2m}]^2}$$

Plasma resistivity?

What is f ?

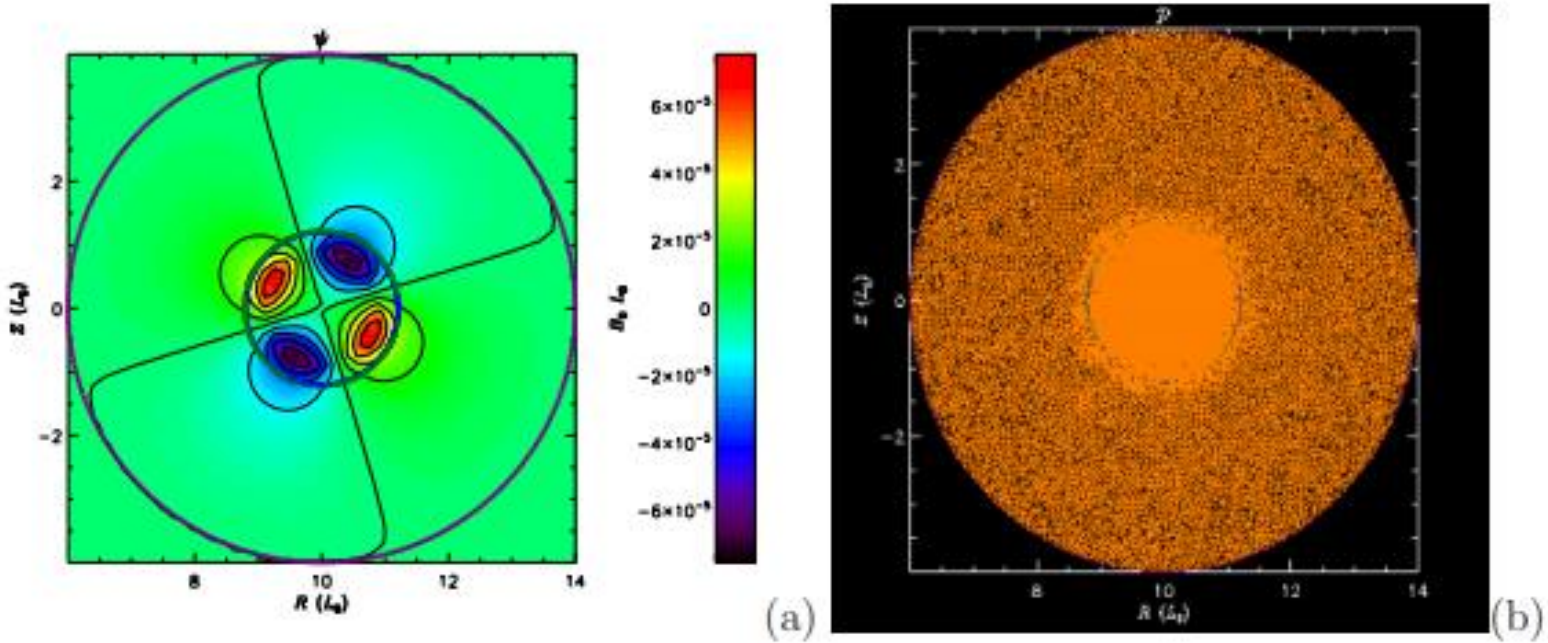
ψ
z (L0)
R (L0)
8
10
12
14
2
0
−2
6×10⁻⁵
4×10⁻⁵
2×10⁻⁵
0
−2×10⁻⁵
−4×10⁻⁵
−6×10⁻⁵
B0 L0
P
(a)
(b)

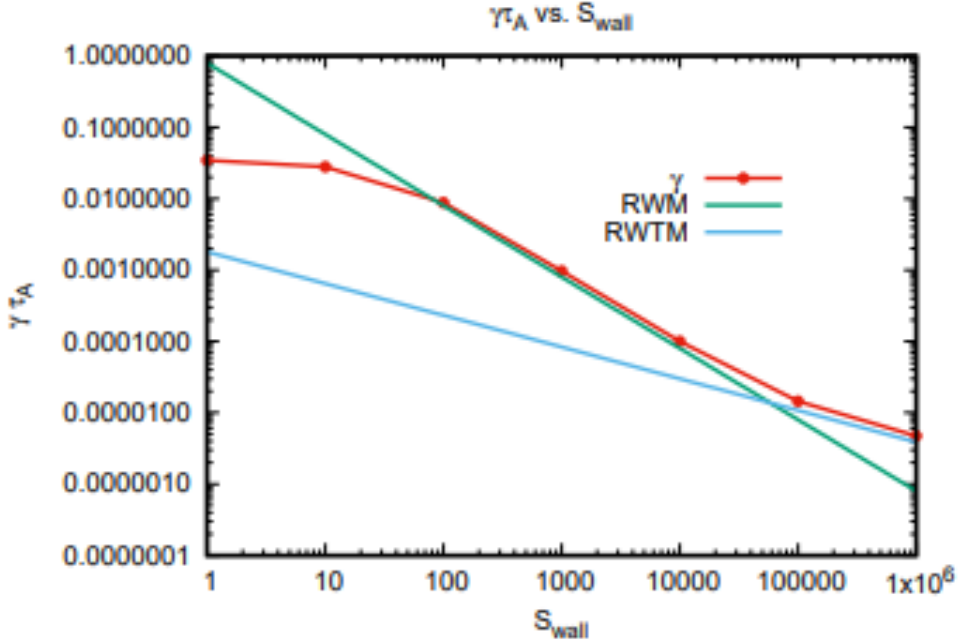

Fig.2(a) shows linear ψ for the case $S_{wall} = 10^4$. Fig.2(b) shows the adapted mesh used in all the cases. The mesh has a thin wall, $\delta = 0.02$. Fig.3 shows growth rate γ as a function of S_wall. For $S_{wall} \geq 10^5$, the most ustable mode appears to be a RWTM. The straight line fits are to S_{wall}^{-1} for a RWM and $S_{wall}^{-4/9}$ for a RWTM.

Dependence on η ?
Does mode structure change?

Sawteeth with 50% runaway current

08092021

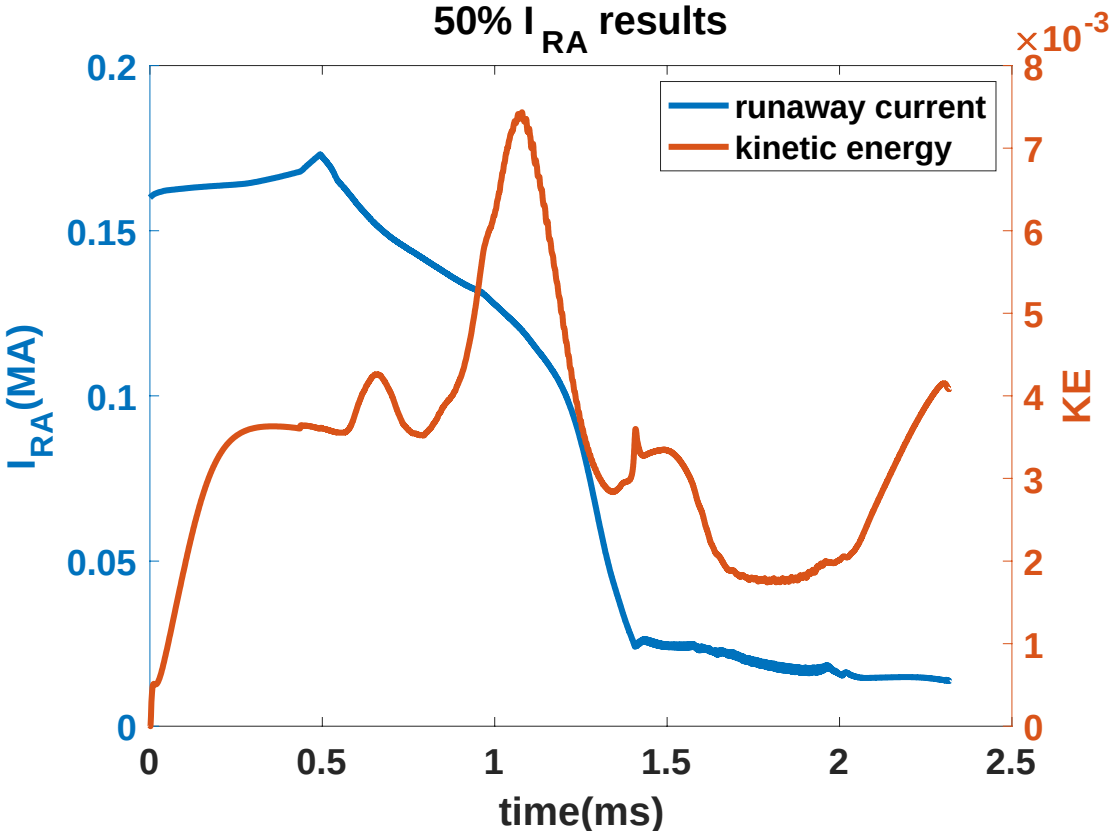

- The runaway electron current dropped dropped to nearly 0 at about t=1.1ms by the MHD instability during ST.

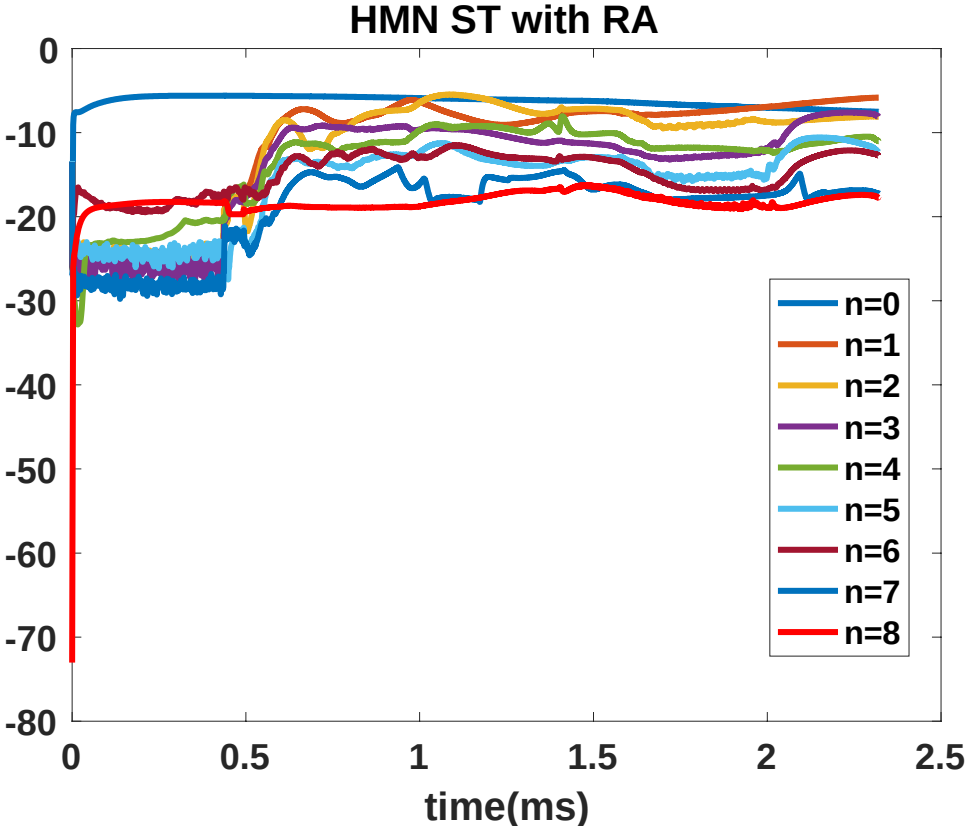

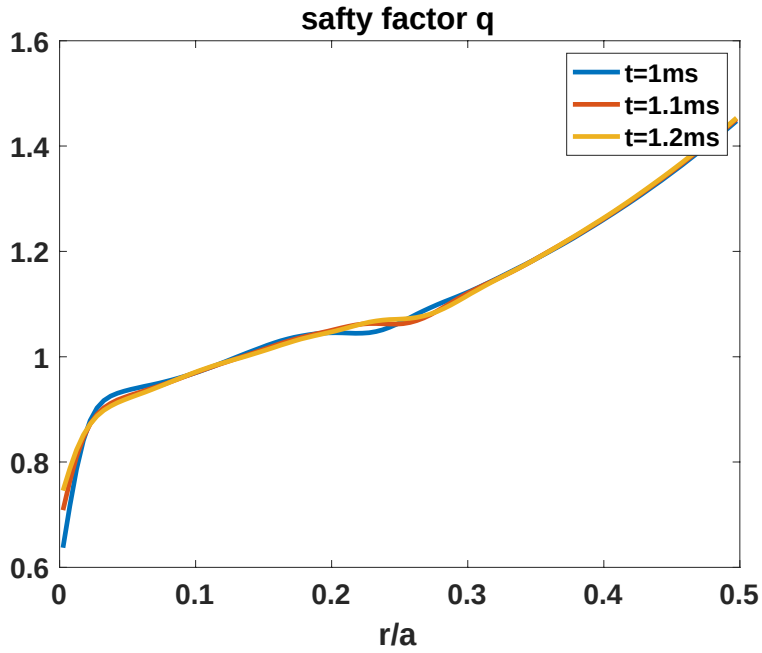

The highest is n=2 at t=1.1ms

Runaway current density profile at t=1.1ms

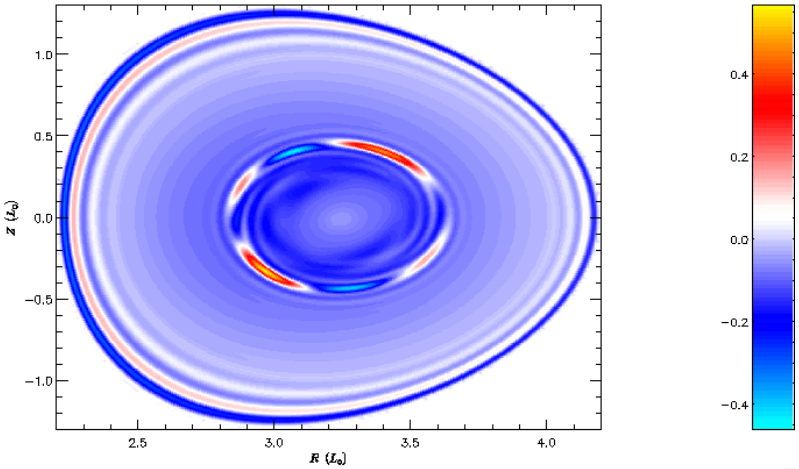

Maybe n=2 m=3 mode

Poincare plot at t=1.1ms

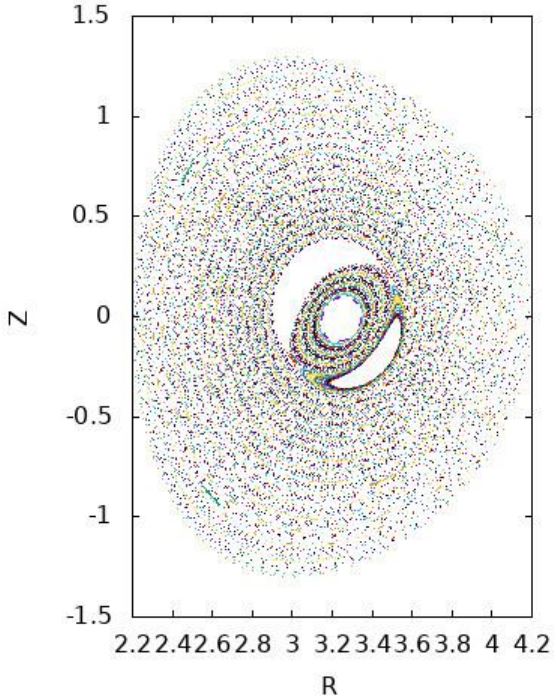

Expulsion of RE Seed - Sawteeth

From Umar Sheikh (7/22/21)

- Sawteeth expel Res
  - Seed population lost - HXR spikes
  - No RE beam after MGI
- Lower I_p removed Sawteeth
  - HXR spikes gone
- Stable RE beams produced

- Sharp drop in I_p indicates larger Ohmic current contribution
  - Sharp drop in I_p if fewer REs (green)
  - Significant RE current before MGI