M3D-C1 ZOOM Meeting

11/21/2022

CS Issues

1. New Block Multigrid Preconditioner – J. Chen
2. RPI Update on Meshing and Adaptation
3. Update on Perlmutter_cpu problems
4. New Latex documentation -
5. NERSC Time
6. Changes to github master since last meeting
7. Regression tests
8. isurface

Physics Studies

1. Update on NSTX shot 124379
2. Chen Zhao, Brendan Lyons updates
3. Anything else

In attendance

Steve Jardin
Dingyun Liu
Cesar Clauser
Andreas Kleiner
Jin Chen
Chen Zhao
Chang Liu
Brendan Lyons
Adelle Wright
Cesar Clauser
Hank Strauss
Mark Shephard
Seegyoung Seol
Usman

New Block MG preconditioner

J. Chen

RPI update on Meshing Capabilities and Adaptation

Update on large job problems on Perlmutter_cpu

My large mesh (98K faces) with 128 partitions now runs on Perlmutter_cpu. The base directory is: /global/cfs/cdirs/mp288/jardin/m3dnl/Perl_cpu/128-K
This has 768 faces per partition and uses 64 processors per node

Cases using mumps all work:

Case	planes	nodes	sec/TS
Run02	4	8	220
Run05	8	16	200
Run08	18	36	250
Run03	36	72	300+

Cases using SuperLU_dist all hang in first GS solve (or produced segmentation fault)
(However, smaller 4-plane case with 32cpu/node did work with SLU)

Case	planes
Run02a	4
Run04	36

New LaTex Documentation

A cleaned version of M3DC1 user's guide is uploaded to M3DC1/doc and this version compiles on all machines using the command "pdflatex M3DC1.tex".

.pdf version is available on m3dc1.pppl.gov

All C1input varialbes are now documented (sectin 6)

Should Appendix B (SCOREC API) be eliminated or revised.?

I will continue to (slowly) review sections for completeness and correctness.

NERSC Time

mp288

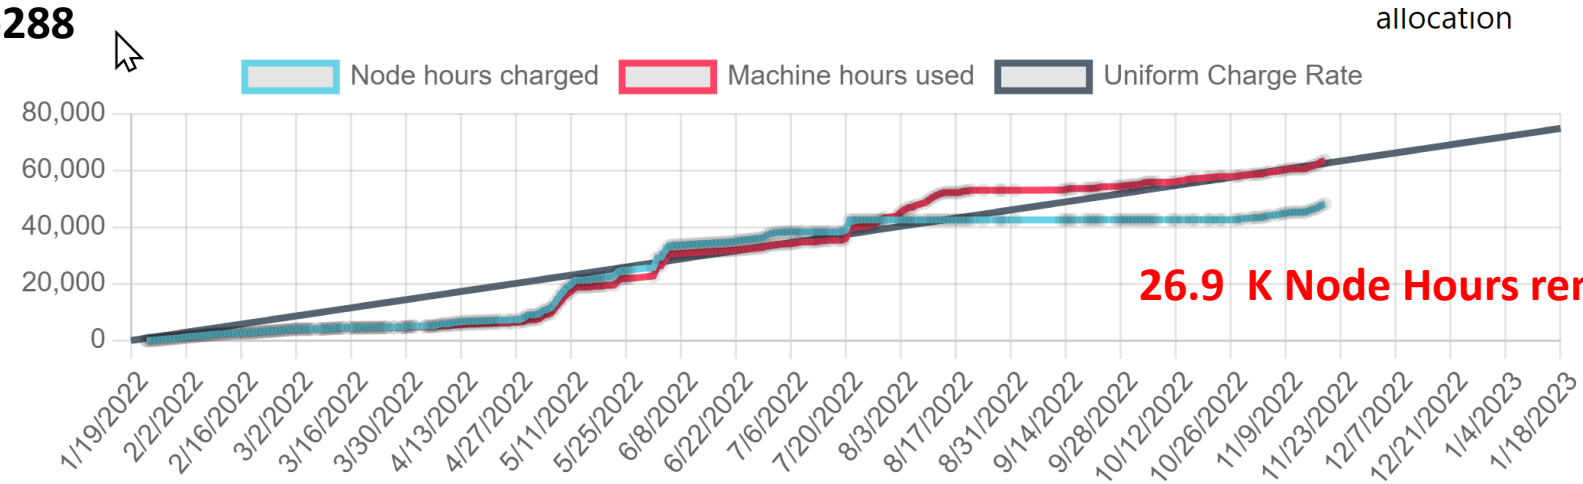

26.9 K Node Hours remaining!

- MP288 usage is on track. Both value and rate are ok.
- We are now being charged for Perlmutter and Perlmutter_cpu.
- We have 2100 (out of 7000) GPU node hours as well.

- Cori is scheduled to go away March 2023

Changes to github master --after 2022-10-03

Steve Jardin:

10/12/22: Fixed bug preventing density control

Brendan Lyons

10/10/22: Allow changes to colorbar charsize

10/11/22: Add Braginskii temperature scaling for parallel thermal conduction

* implemented as ikapparfunc=2
* Created bound with kappar_min and kappar_max
* Fixed ikapparfunc=1 (needs testing)
* Use Te field for kappar definition

11/01/22: Turn on paralled conduction when ikapparfunc=2 but kappar=0

Yao Zhou

10/10/22: Include bfp when changing eqsubtract from 0 to 1

Changes to github master --after 2022-10-03

Seegyoung Seol
10/12/22: Config files for Perlmutter_cpu added
10/20/22: Mesh generation programs installed in Princeton Stellar
10/31/22: adding mesh conversion from simmetrix to pumi
11**/17/2**2: adding config and makefile for perlmutter cudatoolkit 11.7

Jin Chen:
10/12/22: Updates for both Perlmutter_cpu and gpu nodes after extended maintenance
10/24/22: modified Perlmutter_cpu.mk
11/01/22: perllmutter_cpu.mk update

Local Systems

- PPPL centos7(11/19/22)
  - 7 jobs PASSED
- PPPL greene (11/19/22)
  - 5 jobs PASSED
- STELLAR (11/19/22)
  - 7 regression tests **PASSED** on stellar
  - NCSX failed on first try due to a .0012 fractional difference in gamma_gr
- TRAVERSE_gpu(11/04/22)
  - Compilation error (being looked at by Seegyound , Jin, and Chang)

NERSC

- Cori-KNL (10/01/2022)
  7 regression tests **PASSED**

- Cori-Haswell (10/01/22)
  7 regression tests **PASSED**

- Perlmutter (11/19/2022)
  6 jobs **PASSED** (had to run make clean)
  NCSX failed with very small difference in C1ke

- Perlmutter_cpu (11/19/22)
  6 jobs **PASSED**
  NCSX failed with very small difference in C1ke

isurface

- We plan to remove this option

Ideall MHD Stability analysis of NSTX shot 124379

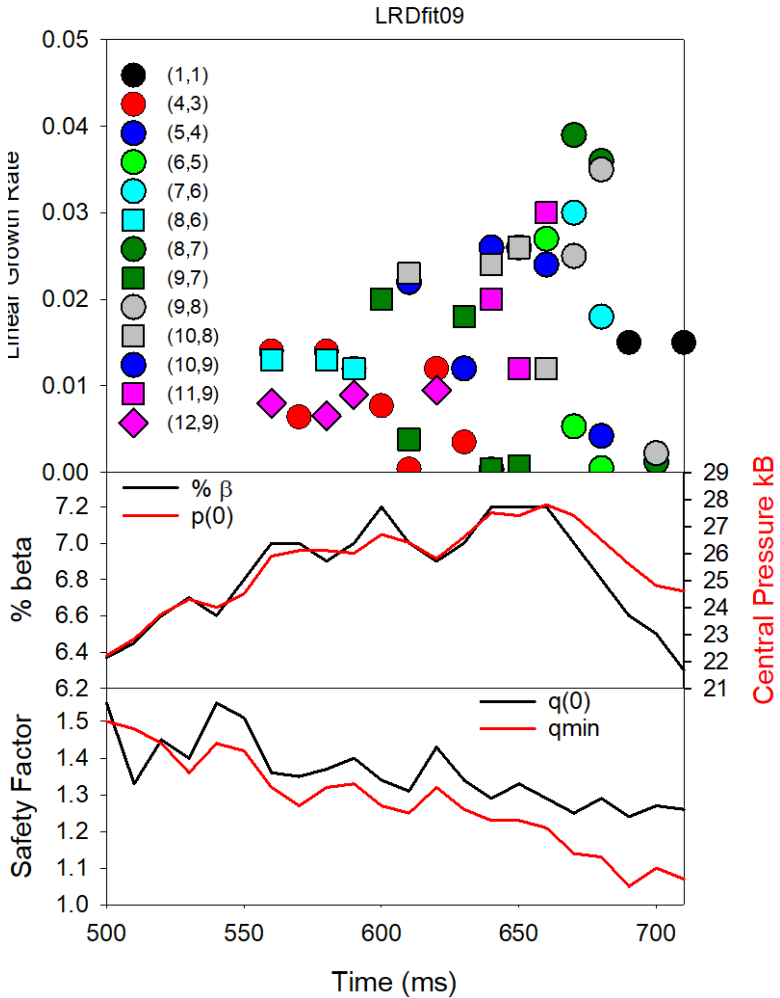

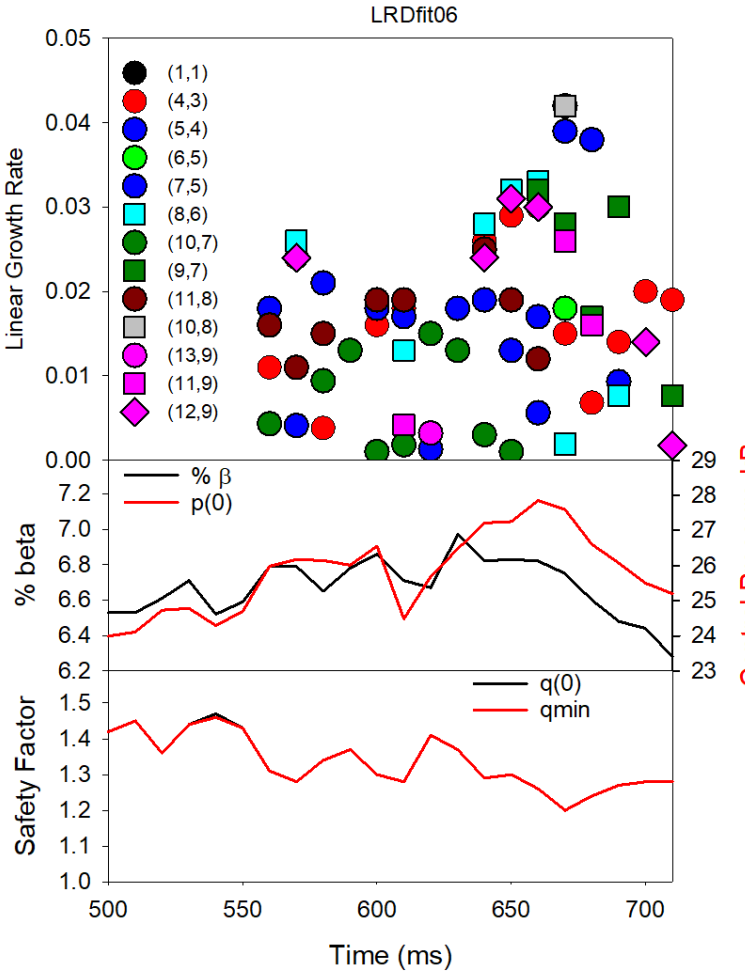

t= 560

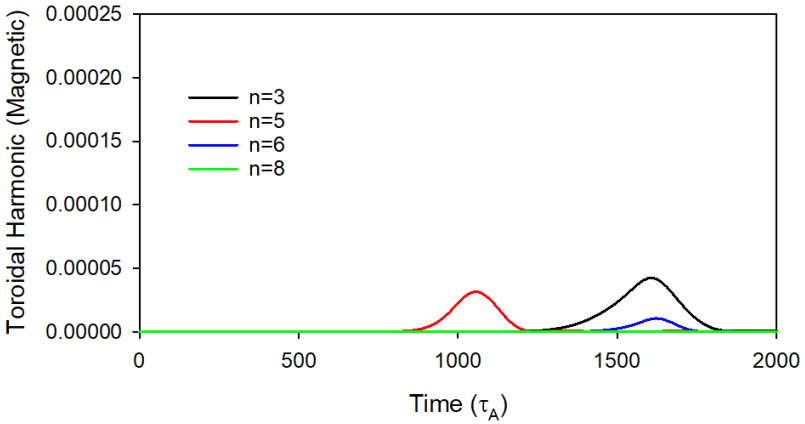
Toroidal Harmonic (Magnetic)
0.00025
0.00020
0.00015
0.00010
0.00005
0.00000
n=3
n=5
n=6
n=8
0
500
1000
1500
2000
Time (τ_A)

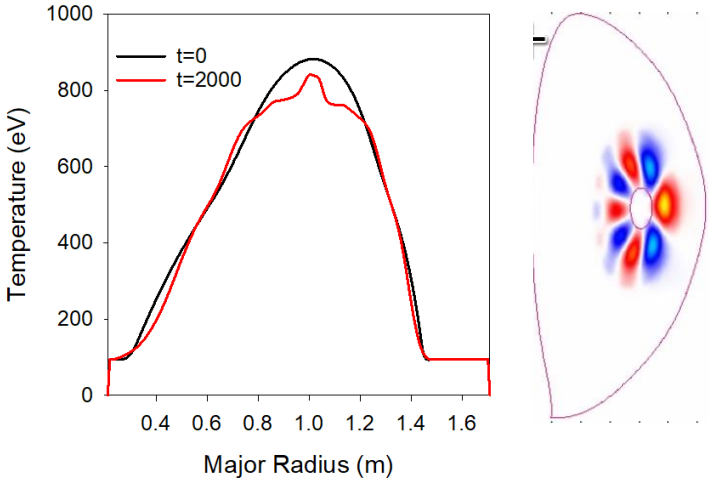
Temperature (eV)
1000
800
600
400
200
0
t=0
t=2000
0.4
0.6
0.8
1.0
1.2
1.4
1.6
Major Radius (m)

t= 640

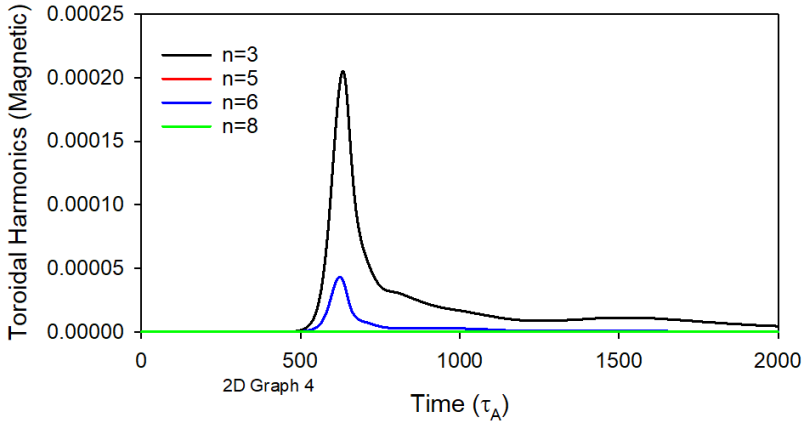
Toroidal Harmonics (Magnetic)
0.00025
0.00020
0.00015
0.00010
0.00005
0.00000
n=3
n=5
n=6
n=8
0
500
1000
1500
2000
2D Graph 4
Time (τ_A)

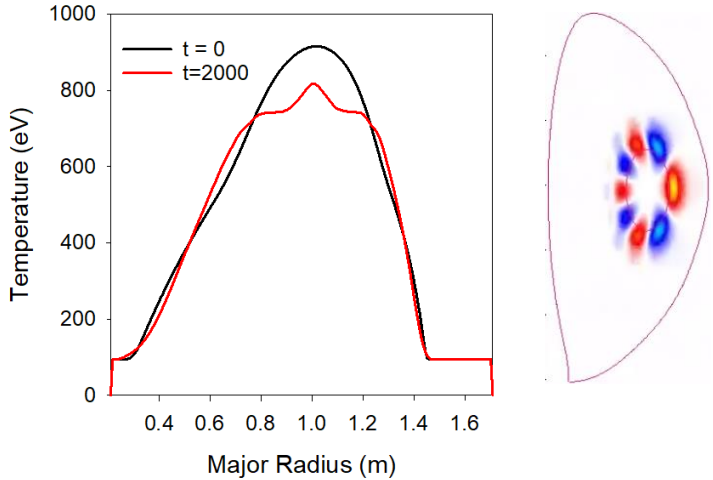
Temperature (eV)
1000
800
600
400
200
0
t = 0
t=2000
0.4
0.6
0.8
1.0
1.2
1.4
1.6
Major Radius (m)

Update on NSTX shot 124379

I have asked Kathreen Thome and Joey McClenaghan (GA) to look at these cases with the ideal MHD codes GATO and DCON.

So far they have not found instabilities!

Should we ask the MARS code?

M3D-C1 Convergence Tests

Focus on LRDFIT06, time=640ms, n=3
eta_fac=1., amu=1.e-6, kappat=1.e-6, kappar=1.e-6

Grid	size(b)	size(0)	faces	grate
A	5.0	0.8	11k	2.676E-02
B	5.0	0.4	19k	2.681E-02
C	5.0	0.2	34k	2.683E-02
D	4.0	4.0	6k	2.650E-02
E	2.0	2.0	23k	2.699E-02
F	4.0	1.0	10k	2.681E-02
G	2.0	1.0	28k	2.702E-02
J	1.5	1.5	43k	2.708E-02
K	1.0	1.0	98k	2.715E-02

Set eta_fac = 0.1

A	5.0	0.8	11k	2.673E-02

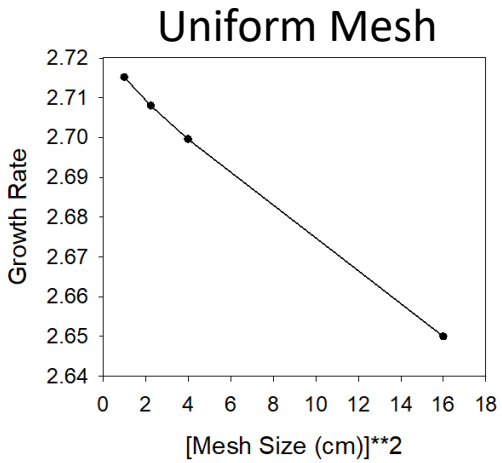

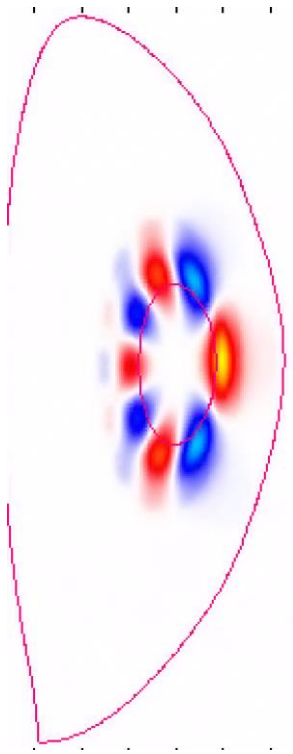

Benchmark from Manickam, et al NF (1987)

$q = q_0 + q_1 \Psi^{\alpha_q}$ $q_0 = 1.05,$ $q_1 = 2.05$ $R/a = 4$ circular

$p = p_0 \left(1 - \Psi^{\alpha_2}\right)^{\alpha_1}$ $\alpha_1 = 4,$ $\alpha_2 = 1.5$

Results somewhat ambiguous

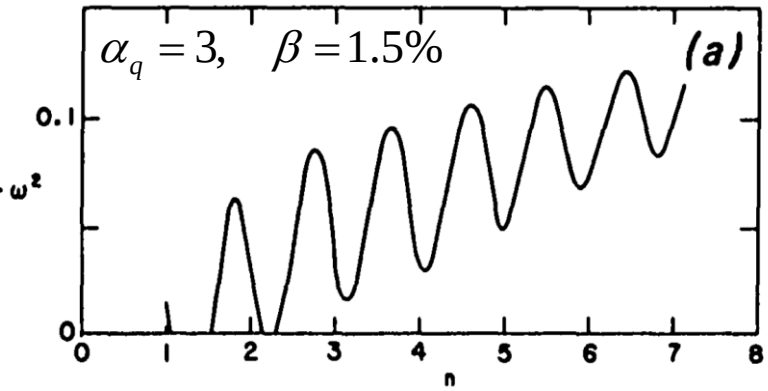

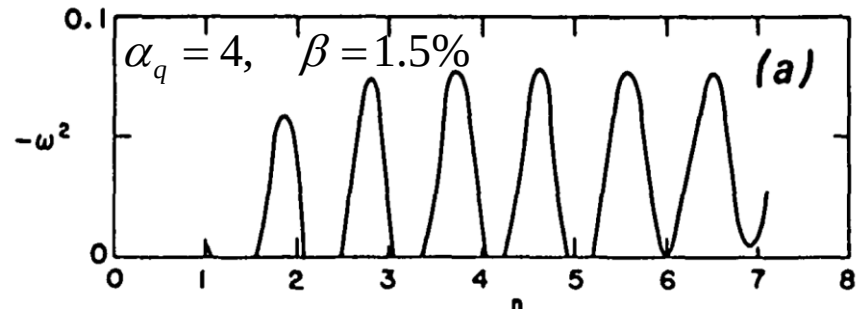

M3D-C1 Results

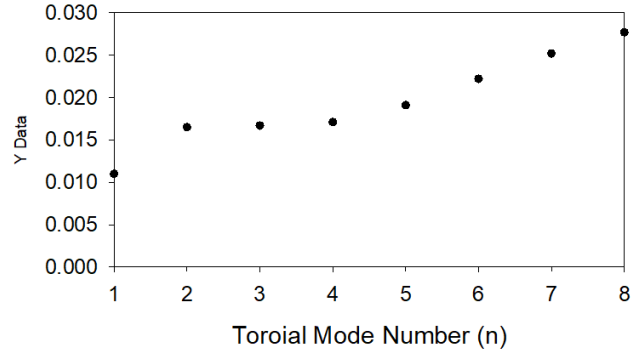

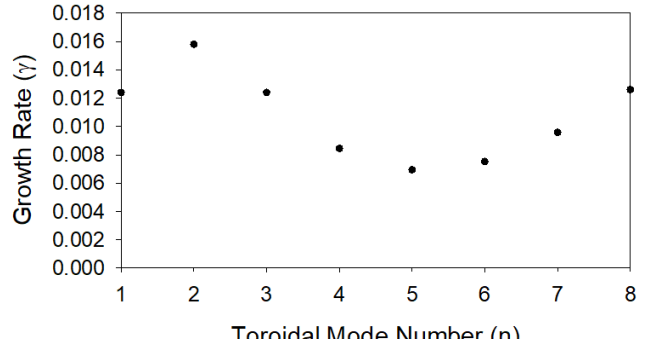

Update on other Jobs

- Chen Zhao produced a new revision of his paper “Simulation of DIII-D disruption with pellet injection and runaway electron beam”
  - I’m still going thru it
  - Why does the RE current first increase and then decrease? Lost at boundary?
  - Should we have an experimental co-author? Will this need to go through DIII review?

- Brendan, status of ITER DM run?

- Other

That’s All I have

Anything Else ?

Papers in Preparation

- Chen Zhao, C. Liu, et al, “Simulation of DIII-D disruption with pellet injection and runaway electron beam”
  - ***Chen is rerunning this now with a lower isotropic thermal conductivity***

Note: May want to plot total thermal energy vs time to compare with experiment….not Te_max

Chen DIII-D VDE case

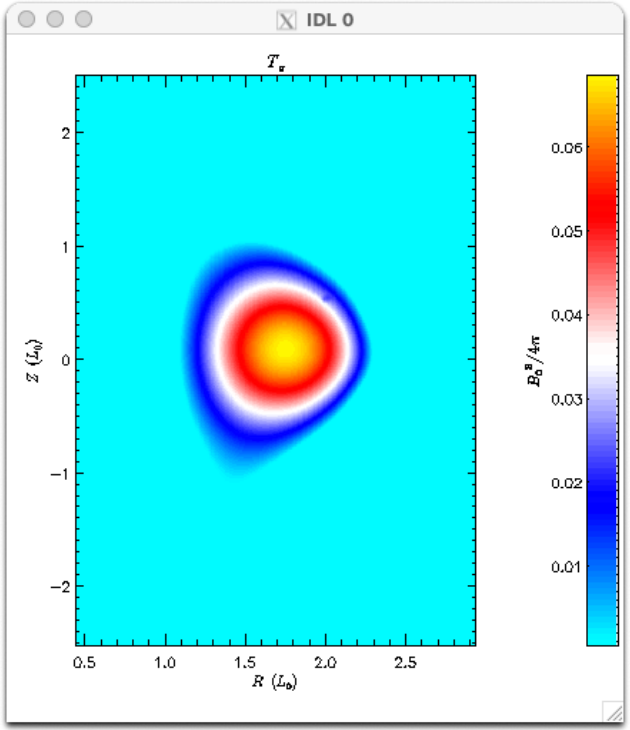

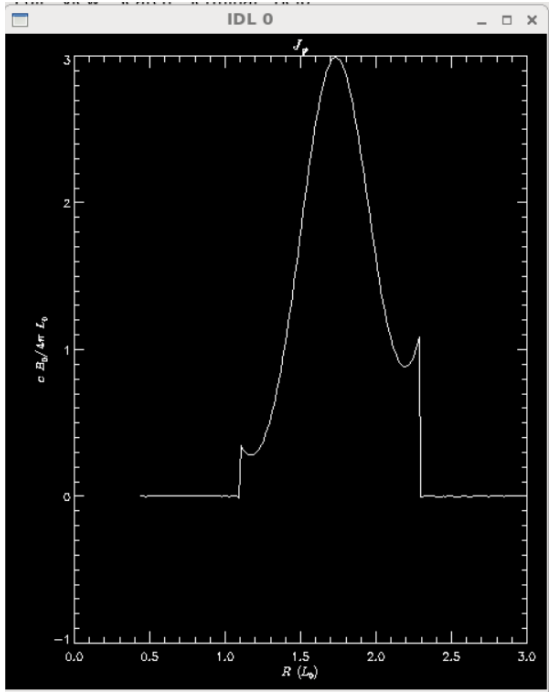

This current profile will not work for a VDE case.
Current density must go smoothly to zero!!

ITER 2D VDE

/scratch/gpfs/sjardin/ITER2D/ITER-03NM-H-64

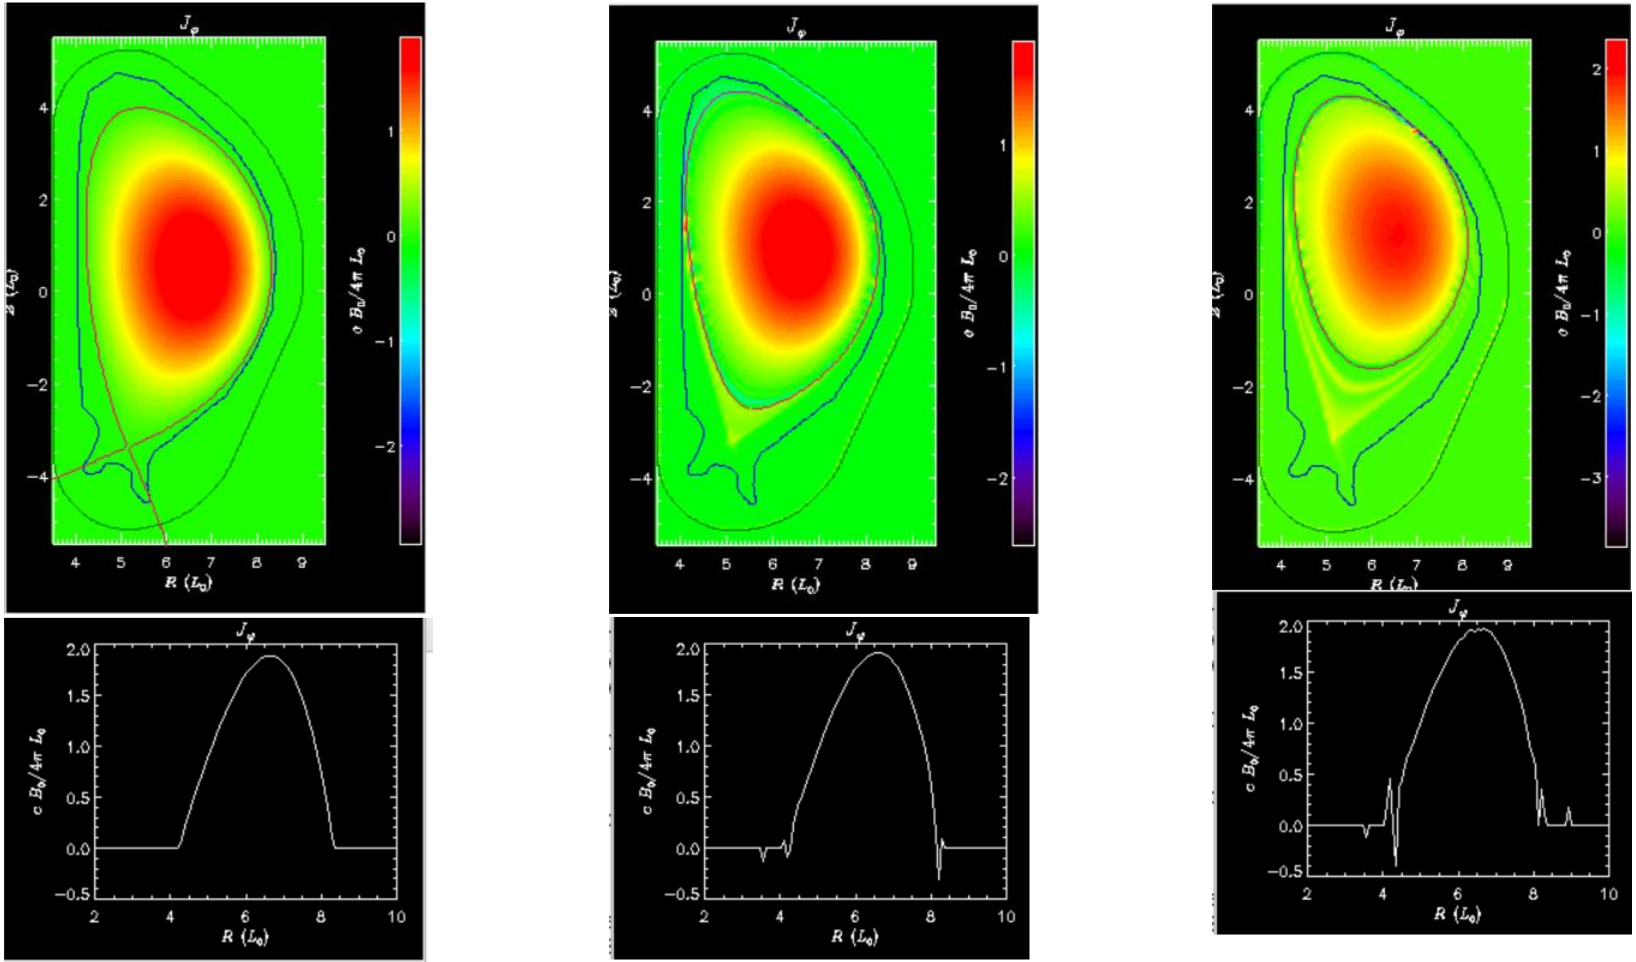

ITER Boundary Conditions

6/27/22 Brendan Lyons

I'm still struggling to get an ITER case to run with the new inoslip_pol=2 boundary condition. I can use it in DIII-D, JET, and KSTAR runs without too much of a problem, but ITER develops an instability at the boundary right away. This is hindering our ability to bring M3D-C1 to bear on ITER SPI modeling, …..

7/11/22 S. Jardin

Only solution I have found is to set inoslip_pol=1. Is this acceptable?

From Michael Lehnen to Brendan 8/24/22

The boundary condition problem is a severe one and maybe you can communicate with Steve Jardin about how to solve it.

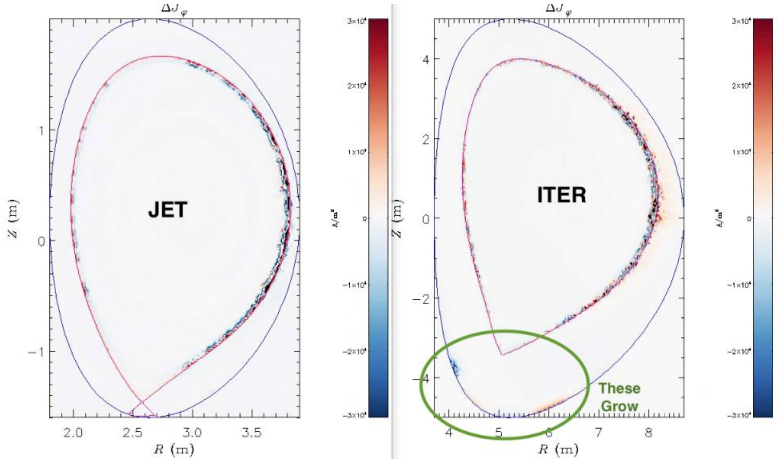

ITER 2D inoslip_pol=1 …cont

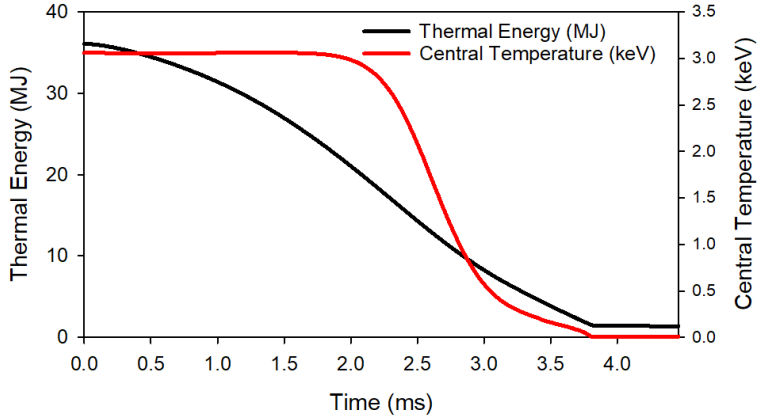

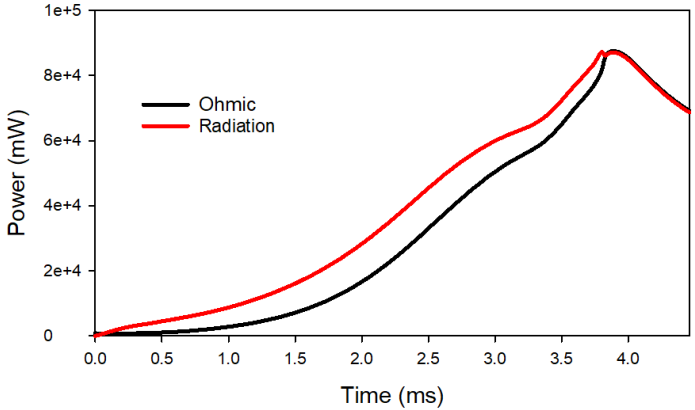

L-mode profiles

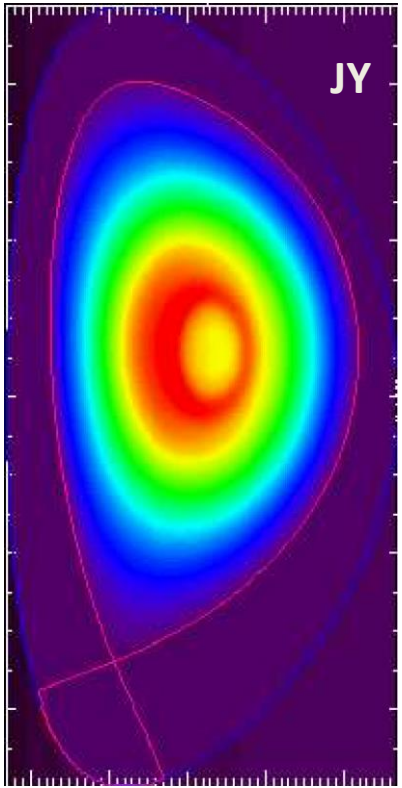

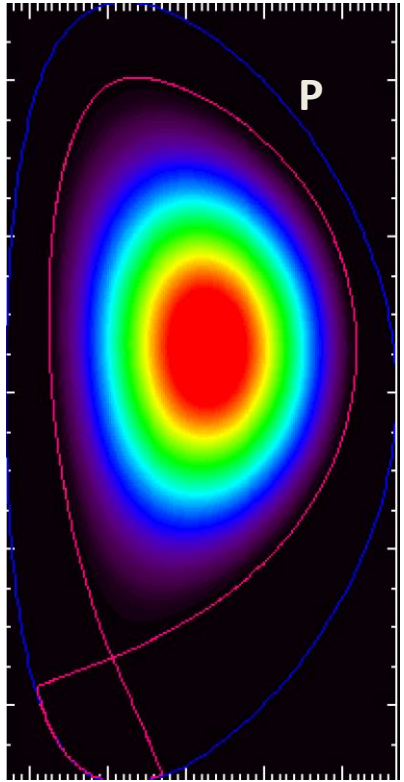

/scratch/gpfs/sjardin/BrendanP1

ITER 2D inoslip_pol=1

kprad_rad

0.89 ms 1.78 ms 2.67 ms 3.56 ms 4.46 ms

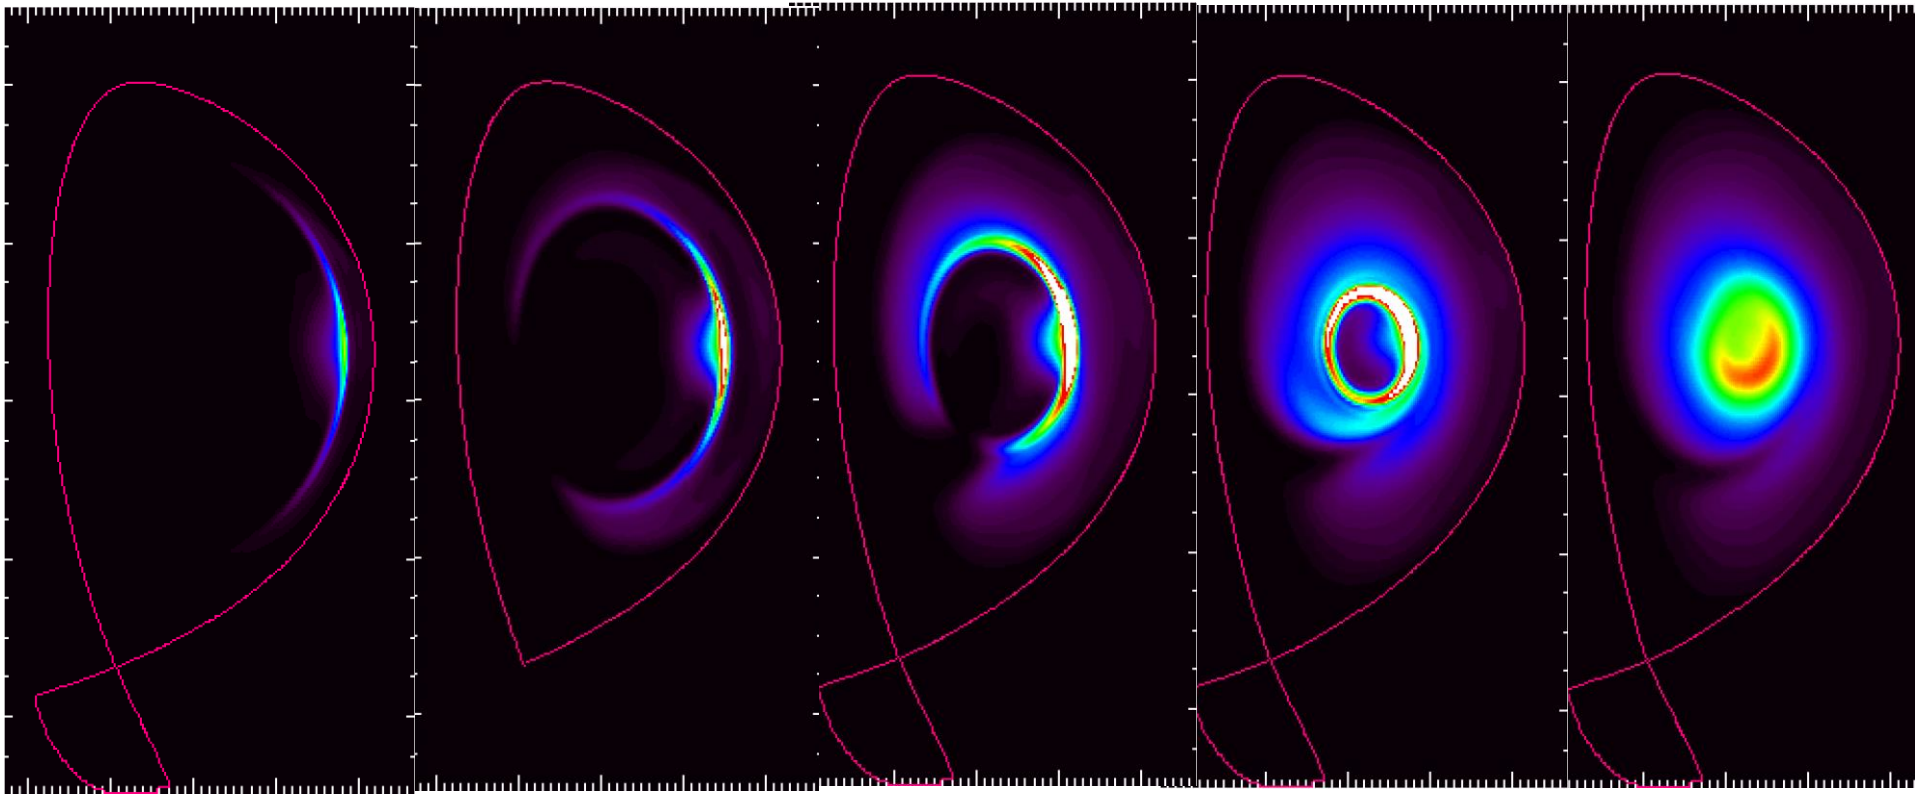

Stellarator Version Convergence

Adelle posted on SLAC on 7/14/22:

We're running into a bit of an issue with the stellarator version, and I was wondering if someone could please take a look.

- In the stellarator version, we have more difficulty getting the velocity solve to converge. In practice, this means we run at viscosities that are 1-2 orders of magnitude higher than in tokamaks. The thought is that the preconditioner is not as effective.
- Based on the fact that quasihelically symmetric (QH) cases are proving much more difficult to run than quasiaxisymmetric (QA) cases (although those are very challenging too), the issue might be specifically related to the non-quasiaxisymmetric terms in the velocity equation.
- If there is a way to improve the preconditioning, I think that would really improve the performance of the stellarator version, especially if it allows us to decrease the viscosity.

Upgrade to impurity radiation Routines

Brendan Lyons wrote to martin.omullane at ADAS on 6/22/22

Any Response?