M3D-C1 ZOOM Meeting

05/21/2023

Upcoming Meetings

CS Issues

1. Adaptation update …. RPI
2. IDL Postprocessing error … solved?
3. Reduced precision SuperLU …Jin Chen
4. Perlmutter_cpu update
5. NERSC Time
6. Changes to github master since last meeting
7. Regression tests
8. Debug for Dingyun's application

Physics Studies

1. NSTX TAE chirping simulation
    1. Numerical stability with toroidal rotation
2. Double Tearing Mode in NSTX
3. Progress on papers
4. Anything else

In attendance

Steve Jardin
Hank Straus
Chang Liu
Jin Chen
Brendan Lyons
Cesar Clauser
Priyanjana Sinha
Chen Zhao
Andreas Kleiner

Seegyoung Seol
Usman Riaz

Upcoming Meetings

EPS	July 3-7	Bordeaux, France
TSDW	July 19-21	Princeton, NJ
IAEA	Oct 16-21	London, UK
APS	Oct 30 – Nov 1	Denver, CO
AAPPS-DPP	Nov 12-17	Nagoya, JP

Adaption Update

RPI?

IDL Postprocessing error

From Cesar to Nate and myself on 4/17/23

Is this all cleared up?

Reduced Precision SuperLU

Jin Chen to present

Perlmutter_cpu problem

Cesar email May 7:
I'm having problems in getting the latest version of code to run on Perlmutter (cpu):

* I compiled the latest version of the master branch and the regression tests worked well (except the small error in NCSX).
* However, when trying to rerun previous cases that I had, they don't work any more. These are 2D cases (cold VDEs)
* Previously, I used to use superlu (srun -n 192 -c 2 m3dc1_2d), but now I see the readme has mumps (srun -n 192 -c 2 m3dc1_2d -pc_factor_mat_solver_type mumps). Both cases fail in different ways:
    - superlu fails with an error before even solving Grad-Shafranov, and the run just hangs there without stopping (consuming all your allocated time)
    - mumps solves GS and then but, after ~10 time step, it starts giving NaNs.
* The same run use to work before.
* I also tried the same run on Stellar using the latest version and it works well.

Perlmutter_cpu problem-2

Cesar email May 7:
I'm having problems in getting the latest version of code to run on Perlmutter (cpu):

Nate and Jin suggested:
For the 2d case with mumps, try adding **"-mat_mumps_icntl_14 100"**
to the "srun" line. This will make umps allocate more memory.

I suggested:
This (2D) run had 13993 vertices and was using 192 partitions with 96 cpu/node and 2 nodes. I suggest changing to 128 partitions, using all 128 cpus on a single node. This would be about 110 vertices/partition which should fit in a single node

perlmutter_cpu update

- Large jobs with 73 K vertices(N) and 98 K vertices (L) with 380 vertices/partition
  - 192-N-09 with 192 partitions, 380 v/p, 64 planes, 64 cpu/node: runs ok!
  - 192-N-10 with 192 partitions, 380 v/p, 32 planes, 64 cpu/node: runs ok!
  - 128-K-09 with 128 partitions, 386 v/p, 64 planes, 64 cpu/node: runs ok!
  - 128-K-10 with 128 partitions, 386 v/p, 32 planes, 64 cpu/node: runs ok!
  - 256-L-09 with 256 partitions, 384 v/p, 64 planes, 64 cpu/node: segmentation v
  - 256-L-10 with 256 partitions, 384 v/p, 32 planes, 64 cpu/node: runs ok!

- Smaller meshes with 9 K vertices (B) and 19 K vertices (H)
  - Normally run ok with 128 cpu/node and 150-200 vertices/partition
  - Also run ok with 64 cpu/node and 300-400 vertices/partition

Thanks to Jin for correcting my options_bjacobi file with -sub_mat_mumps_icntl_14 100

NERSC Time 2023

mp288

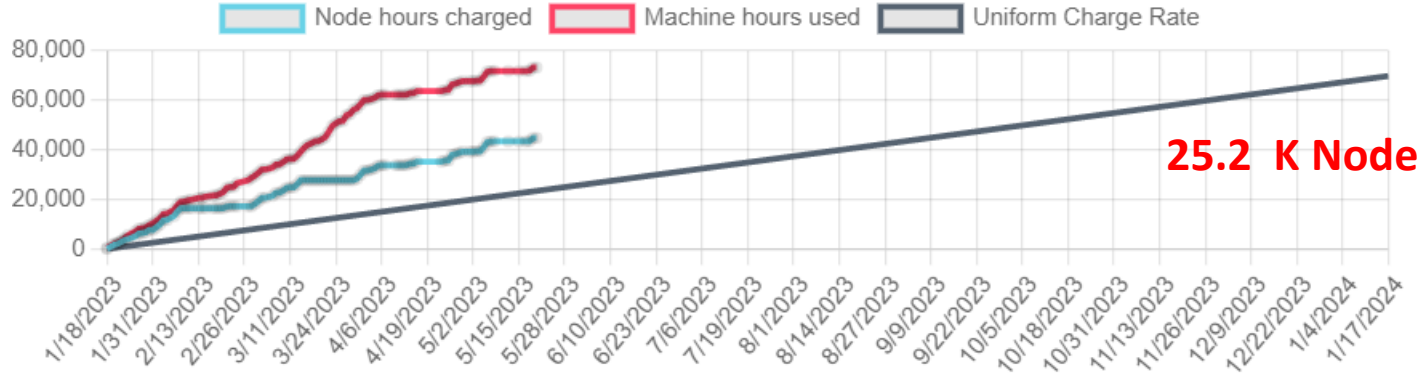

25.2 K Node Hours remaining!

- MP288 usage rate is a bit high but leveling off
- Also, 8.9k GPU node hours
- Cori to go away May 31 2023, 12:00 noon PT
- I will contact DOE to see the likelihood of getting more time

Changes to github master --after 2023-04-27

Jin Chen:
04/27/23: modified matrix dump time

Nate Ferraro:
05/01/23: Fixed issue with getgeomclass setting izone improperly when there are multiple conducting regions

Seegyoung Seol:
05/08/23: porting m3dc1_scorec to 64-bit PETSC

Andreas Kleiner:
05/18/23: Improvements to Python routines
-updated docstrings
-eigenfunction.py: Corrected error in calculation of Fourier spectrum, added data export
-fpylib.py: Bug fix in setup_sims
-Read growth rate files more robustly
- input_vs_t.py: Added options for msec, read files more robustly

Local Systems

- PPPL centos7(05/20/23)
  - 7 jobs PASSED
- PPPL greene (05/20/23)
  - 5 jobs PASSED
- STELLAR (05/20/23)
  - 7 regression tests **PASSED** on stellar
- TRAVERSE_gpu(11/04/22)
  - Compilation error (being looked at by Seegyound , Jin, and Chang)

NERSC

- Perlmutter_cpu (05/20/23)
  6 jobs **PASSED**
  NCSX failed with very small difference in C1ke

- Perlmutter_gpu (02/19/2023) **05/20/23: will not compile**
  - pellet, RMP, & RMP_nonlin, adapt all **PASSED**
  - KPRAD_2D, KPRAD_restart, NCSX all failed with very small differences
  - adapt not submitted

Dingyun is trying to reproduce a case in my 2020 paper

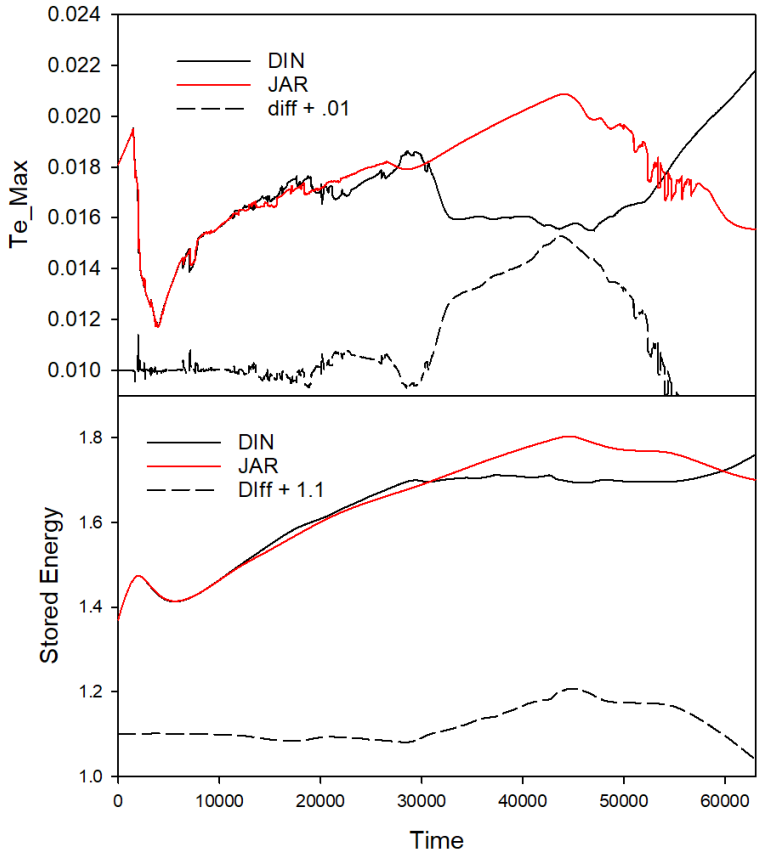

Need to find when these differences first occur:

git log --after 2018-06-30 > logout

git clone https://github.com/PrincetonUniversity/M3DC1

git checkout 8a6a036

Had to:

- remove wrrestartglobal from output.f90 and restart.f90
- Comment out MPI_Comm_split in newpar.f90
- Copy stellar.mk from current version

Progress on debug

Differences first appear on 11/17/20

- Mostly concerned with bf to bfp change $\mathrm{bf} = f, \quad \mathrm{bfp} = f'$

$$\mathbf{B} = \nabla\psi \times \nabla\varphi - \nabla_\perp f' + F\nabla\varphi \qquad R^2\nabla \bullet \nabla_\perp f = F - F_0$$

$$R^2\nabla \bullet \nabla_\perp f' = F'$$

- The two versions agree exactly in 2D !
- The two versions agree fairly well in 3D if $f = f' = 0$, and these results can be very different from f and f' nonzero
- Differences can appear in this and other problems at long time
- These differences likely due to differences in truncation error. Convergence studies should be performed

Run49, fp, op=1
11_17_20: 17:00
Commit: b572e1e7

Run48, f, op=11
11_17_20: 14:48
Commit: 43061986

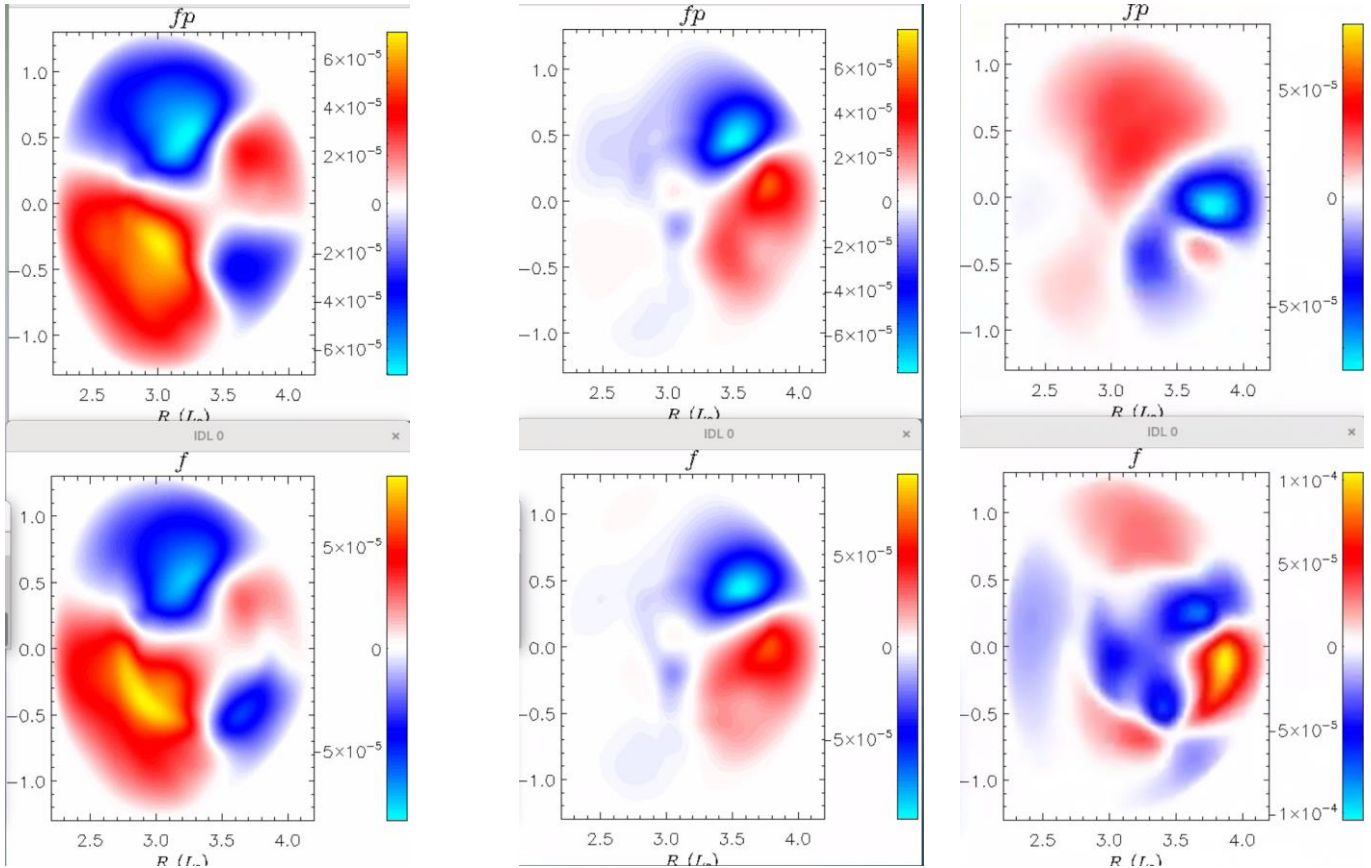

t=4500 t=9000 t=13500

Try turning off certain terms

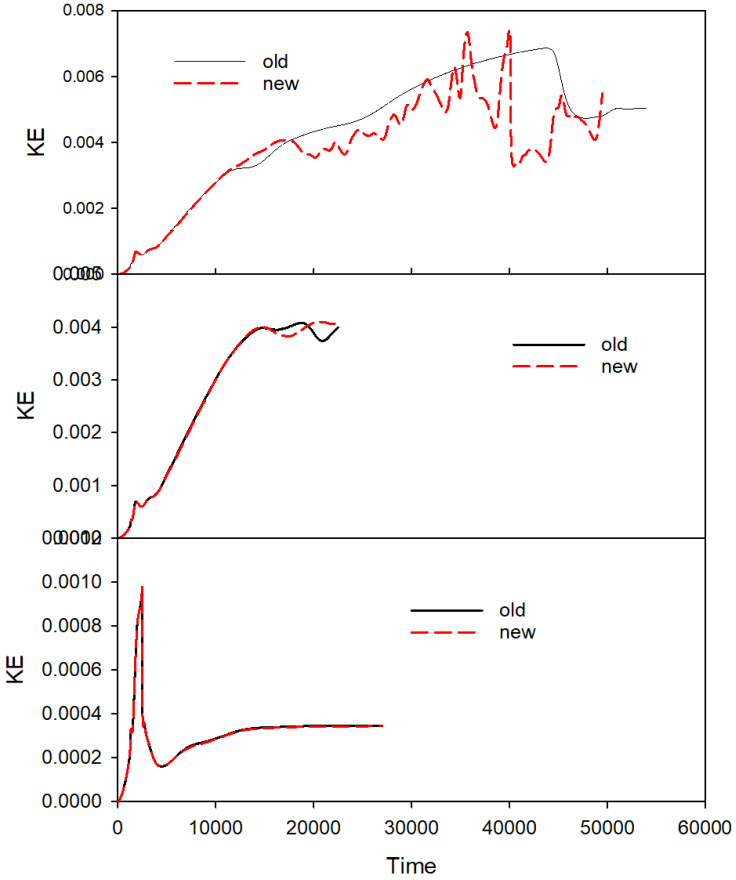

baseline

No Ohmic heating

No Ohmic heating or parallel diffusion

NSTX TAE chirping simulation

Chang Liu to present

Numerical Stability with toroidal rotation

Chang Liu writes on May 11:

. . . I am doing some NSTX MHD simulation using M3D-C1 and I got an issue related to the plasma rotation . . .

The linear MHD simulation (eqsubtract=1) is stable for this case. However, when I include the plasma toroidal rotation, I find the simulation becomes unstable. This is even true with a uniform rotation profile, and true with numvar=1 (reduced MHD).

- How do you include toroidal rotation when numvar=1 ? $\mathbf{V} = \nabla\phi \times \nabla\Phi$
- Try including iupstream = 1 or 2 and setting magus.

Double Tearing Mode in NSTX

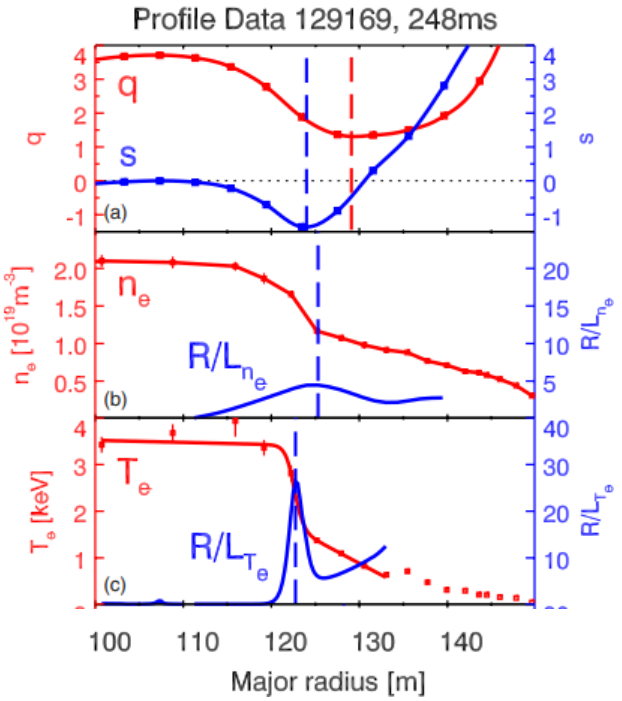

From: Yuh, et al, Phys. Plasma (2009)

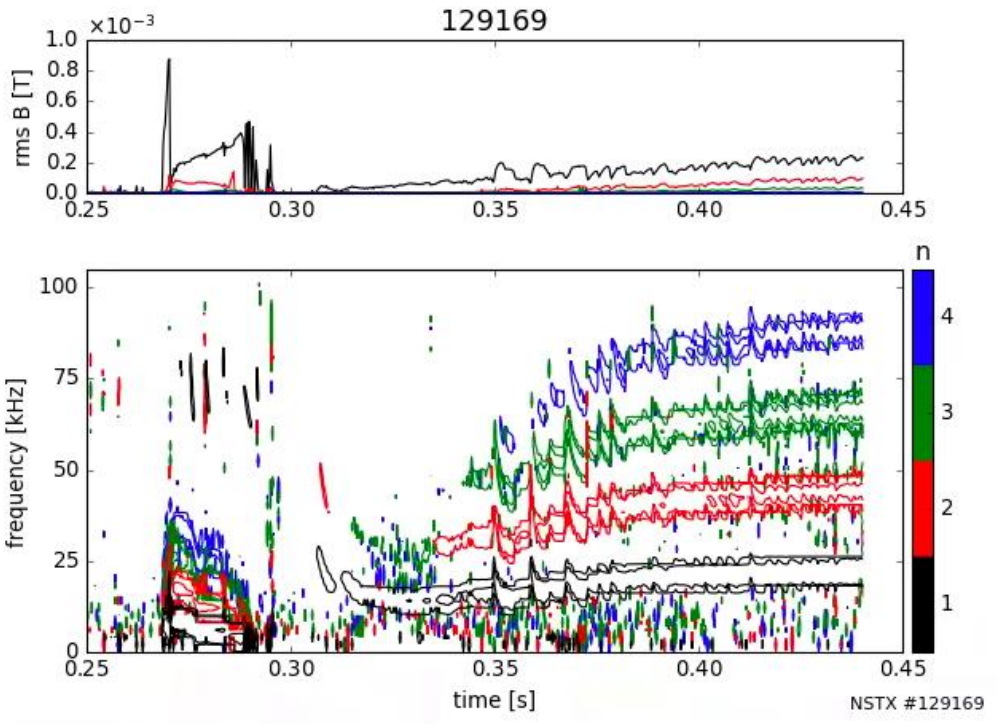

Some MHD activity starting at about 270 ms

DTM Eigenfunction and growth rate

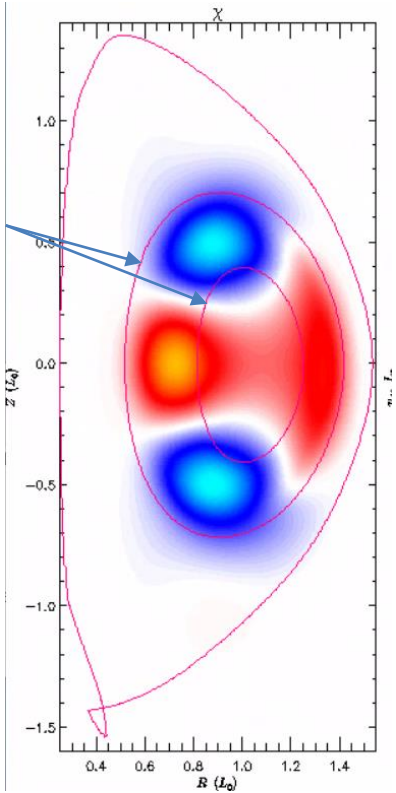

Linear χ at 247 ms

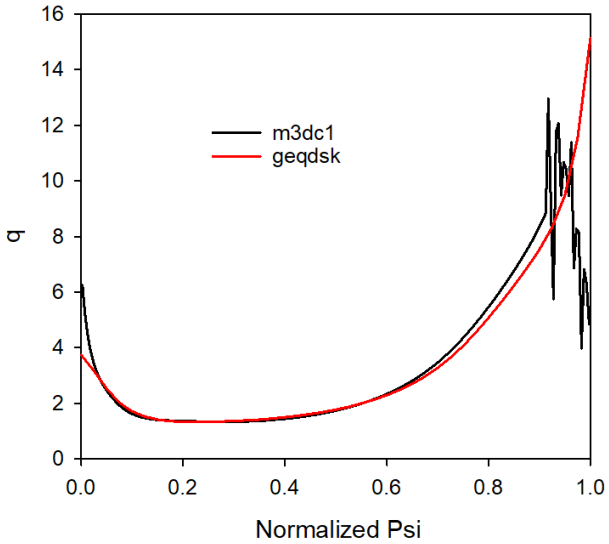

q-profile in geqdsk file is smooth. However, when we recompute the equilibrium in M3D-C1, oscillations appear

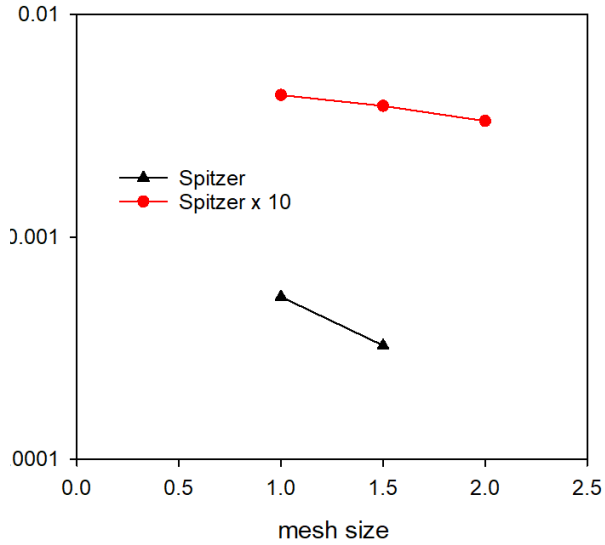

Linear growth rate as a function of linear mesh size for spitzer η and spitzer x 10

Chen paper on modeling the runaway plateau in DIII-D

Simulation of DIII-D disr[illegible]et injection and runaway electron beam

C. Zhao[1], C. Liu[2], S. C. Jardin[2], [illegible] C. Lyons[1]
1. *General Atomics, San Diego, [illegible]nerica*
2. *Princeton Plasma Physics Labo[illegible]ory, Princeton, NJ, United States of America*

Unfortunately, due to some mis-communication, A neon pellet was used in the simulation but the experiment had argon MGI. This is now being redone.

That’s All I have

Anything Else ?