M3D-C1 ZOOM Meeting

12/06/2021

CS Issues

1. LBL Report
2. Perlmutter status
3. Multigrid in Toroidal angle preconditioner discussion
4. Mesh adaptation update (Morteza email)
5. ASCR Workshop on the Science of Scientific-Software Development and Use
6. NERSC Time
7. Changes to github master since last meeting
8. Regression tests
9. Software License discussion
10. Request for Python postprocessor (Dingyun Liu)

Physics Studies

1. Helical Coil Summary (sent to IPP)
2. Update on Soft beta limit study

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

In attendance

Steve Jardin
Adelle Wright
Jin Chen
Andreas Kleiner
Brendan Lyons
Chen Zhao
Nate Ferraro
P Sinha
Chang Liu
Seegyoung Seol
Usman Riaz
Morteza Siboni
Sam Williams
Sherry Li
Yang Liu

LBL Report

Perlmutter status

From Jin Chen (10/25/21)

- PETSc, SCOREC Library, PUMI Library and M3DC! Have been compiled
- Code fails at runtime with segfault
  - Comes from scorec library when it tries to allocate memory for matrices
- Seegyoung now has access also

Progress in multi-grid in φ preconditioner

Sam Williams (LBL) has shown some interest in developing a 1D (in φ) preconditioner for the M3D-C1 solves.

- He has asked if we could construct a matrix for the coarsened grid from the matrix we have for the finer grid?

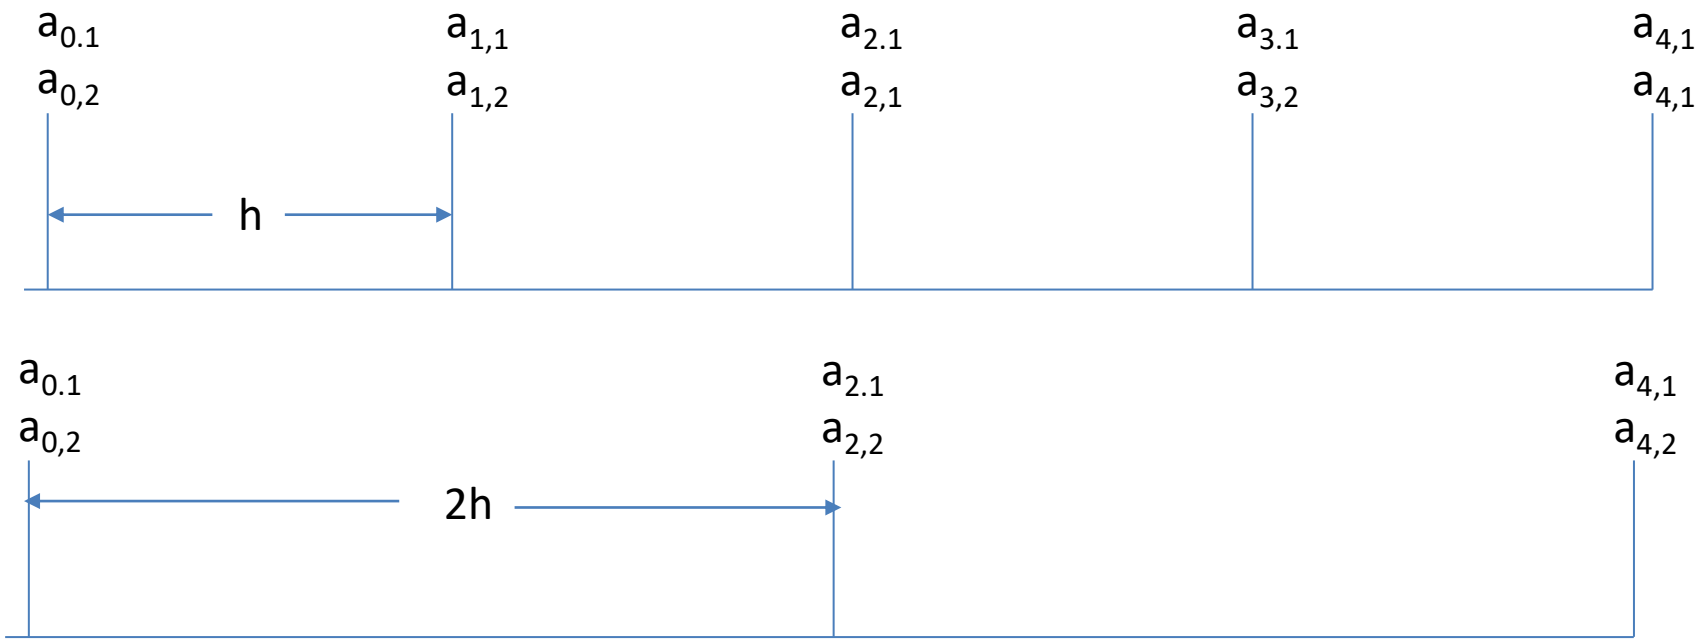

Consider Hermite cubics in 1D for a PDE

$$\frac{\partial \Phi}{\partial t} + V\frac{\partial \Phi}{\partial x} = \alpha\frac{\partial^2 \Phi}{\partial x^2} - \varepsilon\frac{\partial^4 \Phi}{\partial x^4}$$

Finite difference in time using the θ-implicit method:

$$\Phi^{n+1} = \Phi^n - \delta t\, V\left[\theta\frac{\partial \Phi^{n+1}}{\partial x} + (1-\theta)\frac{\partial \Phi^n}{\partial x}\right] + \delta t\,\alpha\left[\theta\frac{\partial^2 \Phi^{n+1}}{\partial x^2} + (1-\theta)\frac{\partial^2 \Phi^n}{\partial x^2}\right] - \delta t\,\varepsilon\left[\theta\frac{\partial^4 \Phi^{n+1}}{\partial x^4} + (1-\theta)\frac{\partial^4 \Phi^n}{\partial x^4}\right]$$

Expand Φ in Hermite cubic elements and apply the Galerkin method

$$\left[\mathbf{M} + \delta t\theta\left[V\mathbf{N} + \alpha\mathbf{P} + \varepsilon\mathbf{Q}\right]\right]\bullet\mathbf{Y}^{n+1} = \left[\mathbf{M} - \delta t(1-\theta)\left[V\mathbf{N} + \alpha\mathbf{P} + \varepsilon\mathbf{Q}\right]\right]\bullet\mathbf{Y}^{n}$$

How do the matrices **M, N, P, Q** depend on the mesh spacing *h*?

$$\left[\mathbf{M}+\delta t\theta\left[V\mathbf{N}+\alpha\mathbf{P}+\varepsilon\mathbf{Q}\right]\right]\bullet\mathbf{Y}^{n+1}=\left[\mathbf{M}-\delta t\left(1-\theta\right)\left[V\mathbf{N}+\alpha\mathbf{P}+\varepsilon\mathbf{Q}\right]\right]\bullet\mathbf{Y}^{n}$$

$$\mathbf{M}=\begin{bmatrix} .1286 & .0310 & .7428 & 0.0 & .1286 & -.0310 \\ -.0310 & -.0071 & 0 & .0190 & .0310 & -.0071 \end{bmatrix}\times h$$

$$\mathbf{N}=\begin{bmatrix} -.5000 & -.1000 & 0 & .2000 & .5000 & -.1000 \\ .1000 & .0167 & -.2000 & 0 & .1000 & -.0167 \end{bmatrix}$$

$$\mathbf{P}=\begin{bmatrix} -1.2 & -.100 & 2.4 & 0. & -1.2 & .100 \\ 0.100 & -.0333 & 0 & .2666 & -.100 & -.0333 \end{bmatrix}/\,h$$

$$\mathbf{Q}=\begin{bmatrix} -12 & -6 & 24 & 0. & -12 & 6 \\ 6 & 2 & 0 & 8 & -6 & 2 \end{bmatrix}/\,h^{3}$$

Since M3D-C1 does not store separate matrices for each operator, it is not feasible to calculate the coarse-grid matrice from the fine-grid matrix.

Mesh adaptation update

From Morteza email 12/3/21.

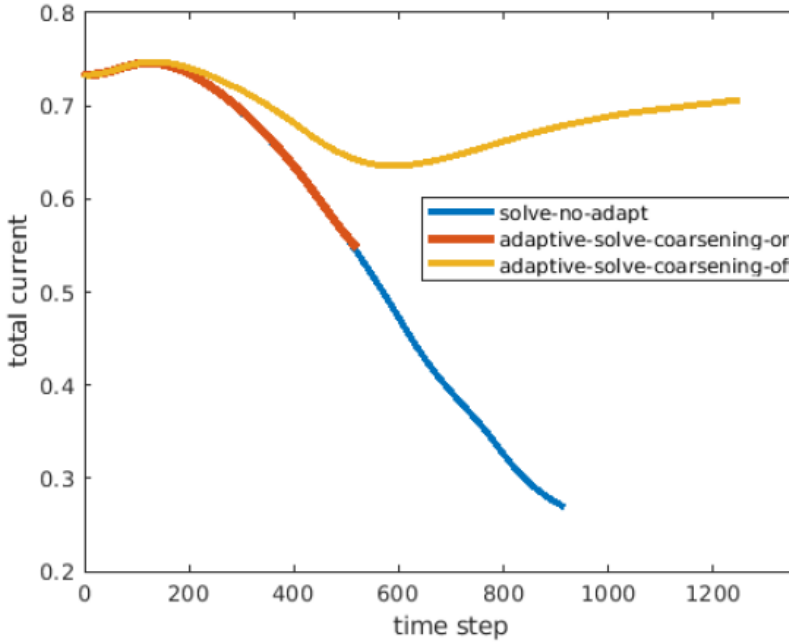

Figure 1: Total current as a function of time step (dt=1 for all ca

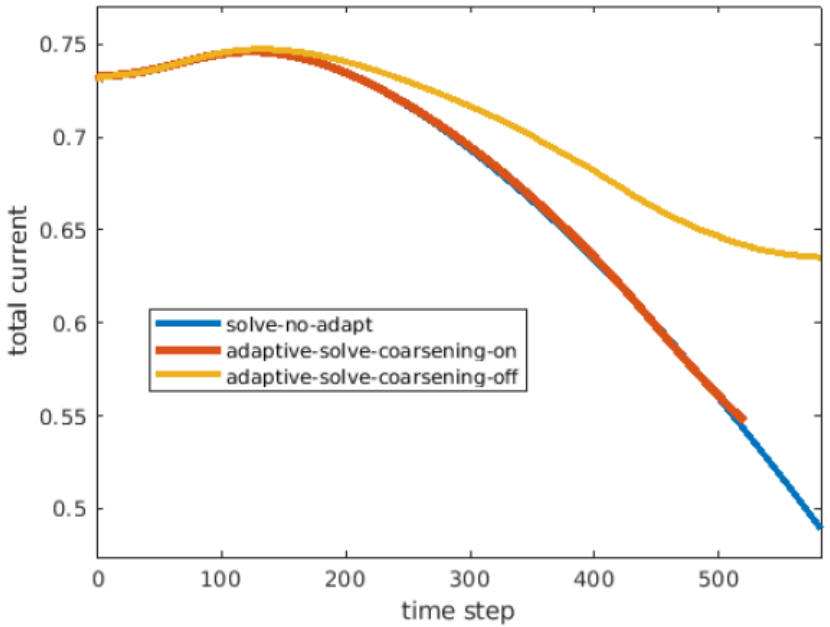

Figure 2: Zoomed-in version of figure 1.

Why does case with coarsening off diverge from the others?

Mesh adaptation update-2

I also tried to run a few cases where I use finer initial meshes, to begin with, and no mesh adapts to find the baseline solutions. I have attempted meshes with average sizes L0, L0/2, L0/4, etc (where L0 is the average size of the initial mesh Brendan provided).Even with a mesh that is only L0/2 the solution diverges pretty quickly. I have tried reducing the time-step as suggested by Stephan (going from dt = 1 to dt =0.1) and still was not able to run the simulations on the finer meshes past a couple of hundred time steps (the simulation on the original (L0) mesh runs to 900+ time steps). Any input on this would be appreciated. I am thinking about using a different solver (MUMPS for example). Do you think this would be a good idea?

SCJ: different solver unlikely to help. Can you try increasing amu?

ASCR Workshop on the Science of Scientific-Software Development and Use

- December 13-15, 2021: 12-5 PM (ET)
- Jin Chen is invited to participate based on her Position Paper
- Others can register as observers for the plenary sessions

Accepted position paper: J. Chen, N. Ferraro, S. Jardin:
"On the need for efficient and scalable solvers for ill-conditioned sparse matrix equations"

Posted on m3dc1.pppl.gov on 11_01_21

NERSC Time

mp288

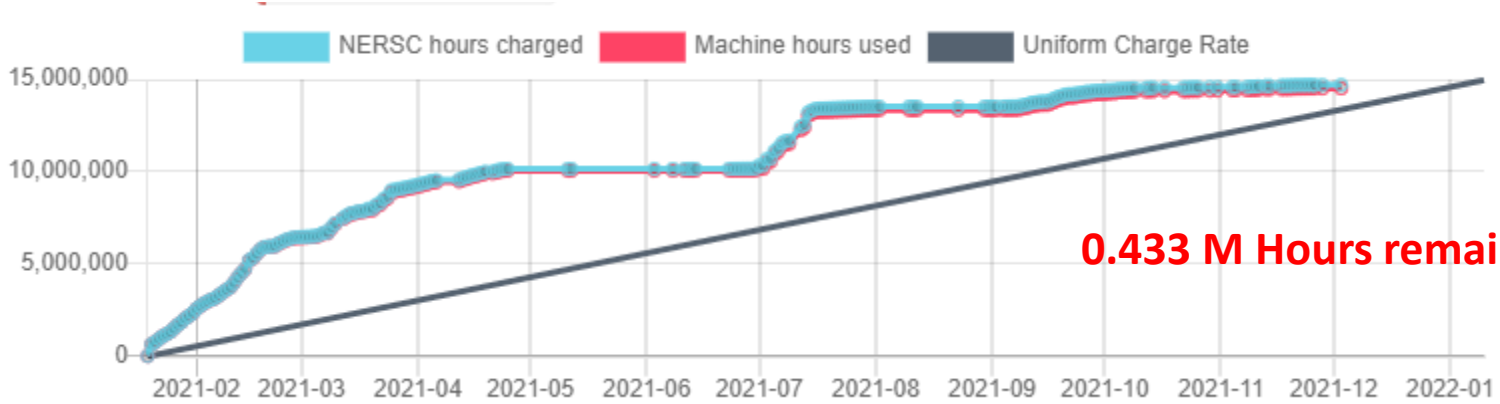

0.433 M Hours remaining!

- mp288 received 10M Hrs for CY 2021, + 5M Hrs additional
- Pearlmutter time will not be charged for this FY
- We are NESAP Tier 2. . Phase-I w GPUs We have been given a repo m3984
- N9ES-N2 M3D-C1: J. Chen , C. Liu, S. Seol are early users

Changes to github master since 11/21/21

No Changes!

Local Systems

- PPPL centos7(11/26/21)
  - 7 jobs PASSED
- PPPL greene (11/26/21)
  - 5 jobs PASSED
- STELLAR (11/26/21)
  - 6 regression tests **PASSED** on stellar
  - adapt **FAILED** field energies off by 0.02%
- TRAVERSE(11/26/21)
  - 6 regression tests **PASSED**
  - adapt **FAILED** should have passed. Energies are ok. Only gr_rate off

Other Systems

- Cori-KNL (11/27/2021)
  - 7 regression tests **PASSED** on KNL
- Cori-Haswell (11/27/2021)
  - 7 regression tests **PASSED** on cori
- MARCONI
  - All regression tests PASSED on MARCONI (J. Chen, 9/04/20)

Possible M3D-C1 Software License

The source code of the M3D-C1 project is available to the world-wide fusion community. This is normally accomplished by granting "Read-only" access to the M3D-C1 Princeton University github repo. The M3D-C1 code requires the installation of several 3rd party software packages including MPI, SCOREC and PETSc. The M3D-C1 team may be available to install these packages at the users facility upon request, but this will require granting user privileges to one or more M3D-C1 team members and is not guaranteed. The following restrictions apply

1) The M3D-C1 code (binaries and source) will not be passed to third parties.
2) Publications or figures made with the M3D-C1 code will acknowledge the M3D-C1 code and include references to one or more of the original publications.
3) Any paper being submitted to a journal for publication using M3D-C1 results must be sent to the M3D-C1 team prior to submission.
4) Modifications made to M3D-C1 by the users of the code will be provided to the M3D-C1 team upon request, and will be added to the M3D-C1 github repo at the discretion of the M3D-C1 team
5) All input files and auxiliary data analyzed using M3D-C1 in a publication are to be provided to the team upon request.
6) There is no guarantee that the M3D-C1 code is free of errors. Any errors or "bugs" that the user may find will be reported immediately to the M3D-C1 team.
7) It is understood that the M3D-C1 code is still under development and thus may not contain the physics or features that users need/want for their problem of interest.
8) It is understood that the M3D-C1 team does not guarantee that support will be available to users of the code and that extensive support from a M3D-C1 team member(s) generally implies that published work will include that team member(s) as co-author(s) .

I understand and agree to the above conditions: User name, institution, and date

Email request for Python postprocessor

Email 12/06/21 @ 9:20 AM
This equilibrium has been running very well so far. And I'm trying to read the C1.h5 file with python to get the density and temperature fields. For example, I loaded the 'C1.h5' as 'h5', opened h5['time_033']['fields']['ne'], and got an array of (206128 by 80). I am trying to understand how the field is stored in the array. Does 206128 correspond to the mesh and 80 correspond to the time?
Do you know if there is any document from which I can know the location of each node in the mesh? Is each time slice divided evenly to 80 time points?
I would really appreciate your help.

Dingyun Liu
Will Fox
Sayak Bose

Helical coils subject to uniform loop voltage

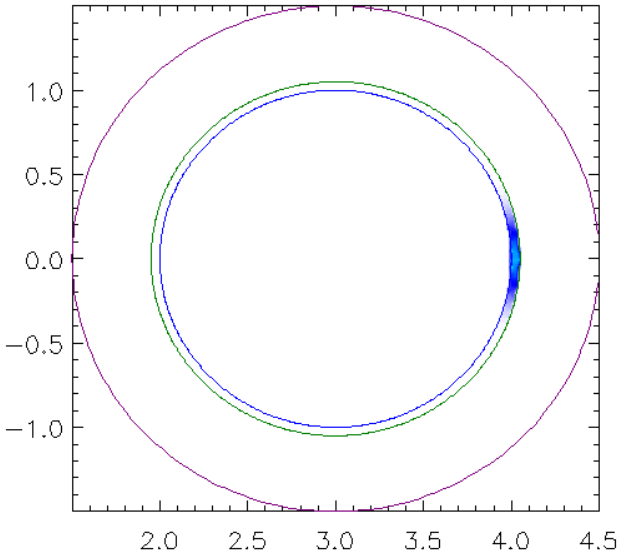

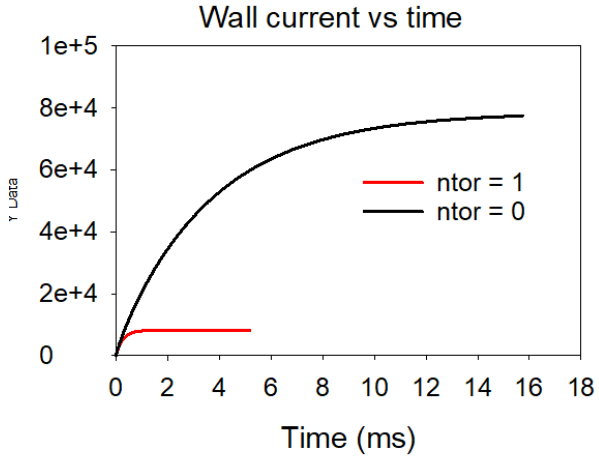

minor radii = 1.0, 1.05, 1.5
major radius = 3.0
V_loop = 1.e-4 Volts

'/projects/M3DC1/sjardin/helical5g/C1.h5'
'/projects/M3DC1/sjardin/helical5f/C1.h5'

f = cos(ntor*phi - mpol*theta)
f = exp((f-1.)/sigma**2)
wall_resistivity = 10.**(log10(wall_resistivity)*(1.-f) + log10(eta_rekc)*f)

sigma = 0.5 eta_rekc = 1.e-6 etar = eta_outer_wall = 0.2
wall_resistivity = 0.001 ntor = 0,1 mpol=1

To get SI resistivity multiply by μ_0 = 1.256 E-6

Update on Soft-Beta-Limit Study

Initial

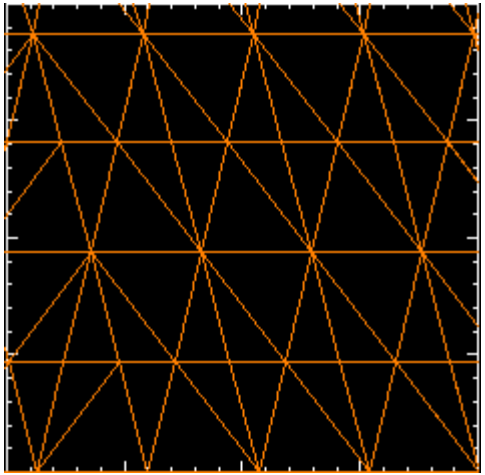

Coarse (C)

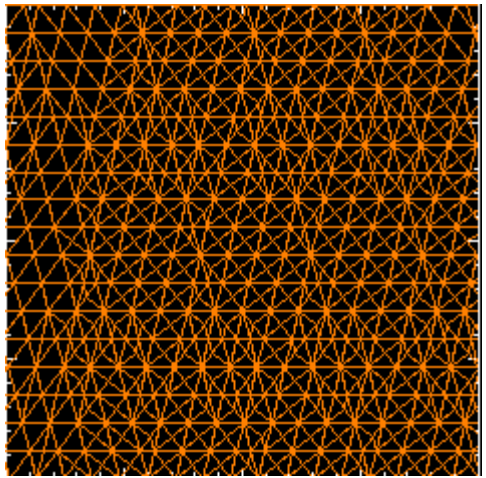

Fine (F)

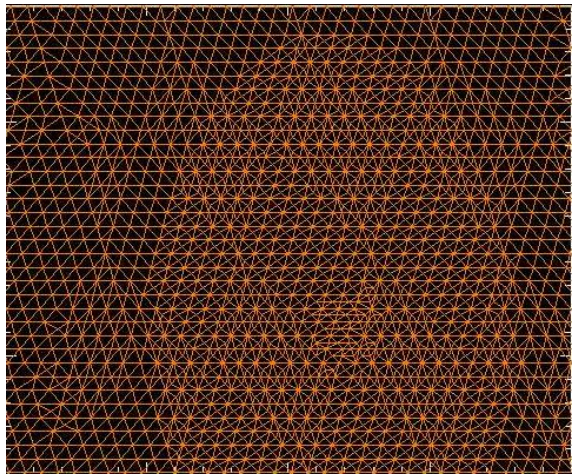

Result of Convergence Study

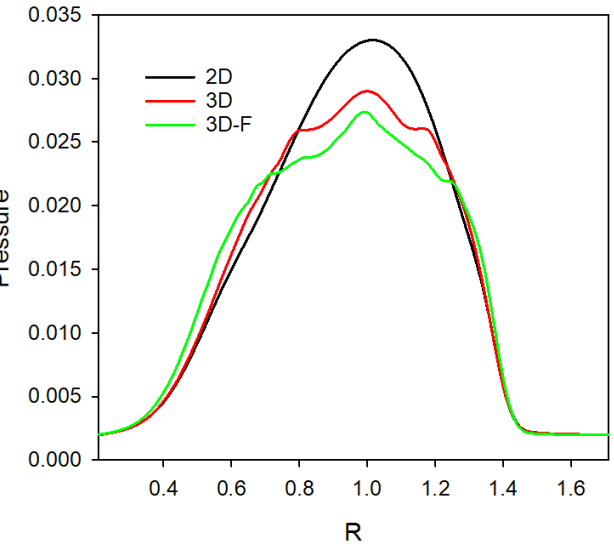

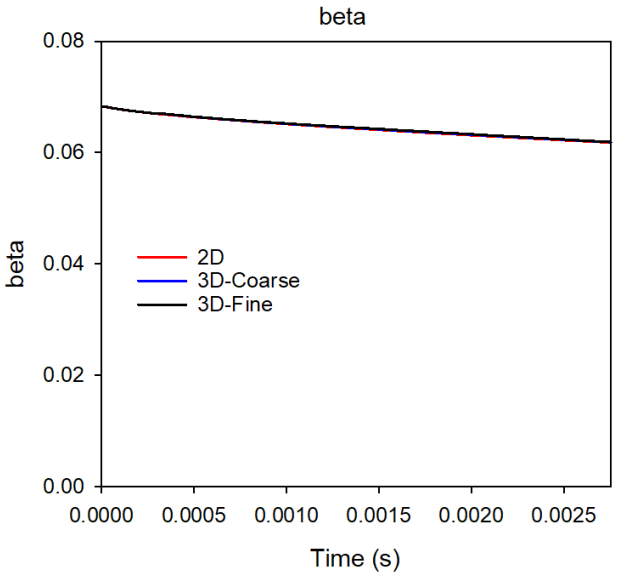

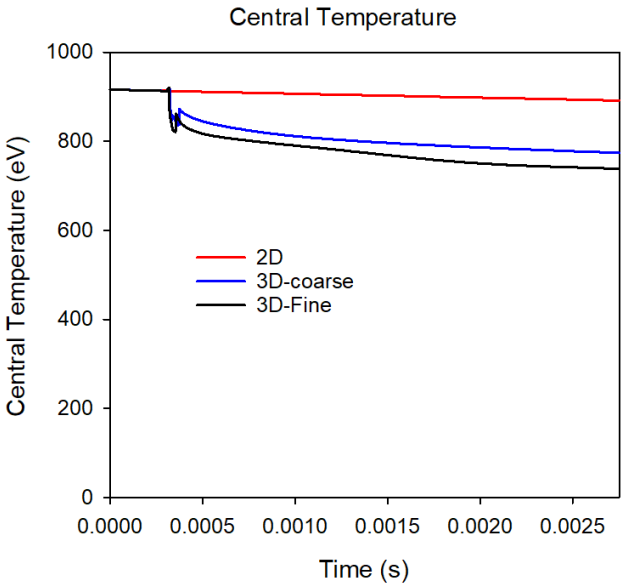

Not clear this has converged on Fine Grid. Will try one more round of refinement

However, it is clear original 2D case was unstable, saturated 3D case is stable. Will ask Sabbagh group to analyze initial case to see if they agree it is unstable

That’s All I have

Anything Else ?