# M3D-C1 ZOOM Meeting

09/13/2021

Announcements

CS Issues

1. LBL Report and Discussion
2. Mesh adaptation update
3. NERSC Time
4. Changes to github master since last meeting
5. Regression tests
6. Update:  Code hangs when writing slice #60 – Yao Zhou

Physics Studies

1. Mesh for DIII-D TQ Studies -- Strauss
2. Request for new parallel thermal conductivity -- Strauss
3. LP Coupling status – Samulyak, Lyons
4. Publications update:  Zhao, Clauser, Liu,
5. 1/1 mode with negative loop voltage on EAST – Liqing Xu
3\. Poloidal Velocity Boundary Conditions -- Lyons

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

# In attendance

Steve Jardin
Hank Strauss
Jin Chen
Adelle Wright
Nate Ferraro
Chen Zhao
Brendan Lyons
Chang Liu
Cesar Clauser
Andreas Kleiner

Mark Shephard
Seegyoung Seol
Usman Riaz
Morteza SIboni

Sam Williams
Sherry Li
Nan Ding
Yang Liu

# Announcements

- NERSC ERCAP requests due Oct 4
  - Separate requests for CPU (cori and Perlmutter) and GPU(Perlmutter)

- /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

# LBL Presentation and Discussion

**From Sam Williams, 9/13/21**

Both the CPU and GPU partitions use AMD EPYC CPUs. Each CPU has 256GB and 64 cores (4GB/core). Core for core, these are comparable to Haswell cores and much faster than a KNL core. Node for node, this is 4x the cores, 4x the memory, 1x the memory/core. In aggregate, there is ~150GB/s of bandwidth per CPU (about 2x haswell and 1/2x bandwidth per core) There are 2 CPUs per Perlmutter CPU node and 1 CPU+4GPUs per Perlmutter GPU node.

So, running flat MPI means each process will likely be 0.5x-1x a Haswell (depending on whether bandwidth- or compute-bound). However, with 4x the processes per CPU, per-node performance should increase by 2-4x. Cutting the #cores in half will double the performance of flat MPI memory-bound apps and double the bandwidth per core for all apps. However, this will hurt compute-bound apps.

GPUs have ~1.5x the FLOP/s, ~1.5x the bandwidth, and ~2.5x the memory. However, you can only use the former if you have massive parallelism on the GPU. Node for node, a Perlmutter GPU vs. Summit node should be ~1x performance and ~1.5x the memory (but its more concentrated in fewer GPUs)

# Mesh Adaptation Update

RPI?

Jardin and Strauss to send request for capability for refinement at rational surface and wall.

# NERSC Time

**mp288**

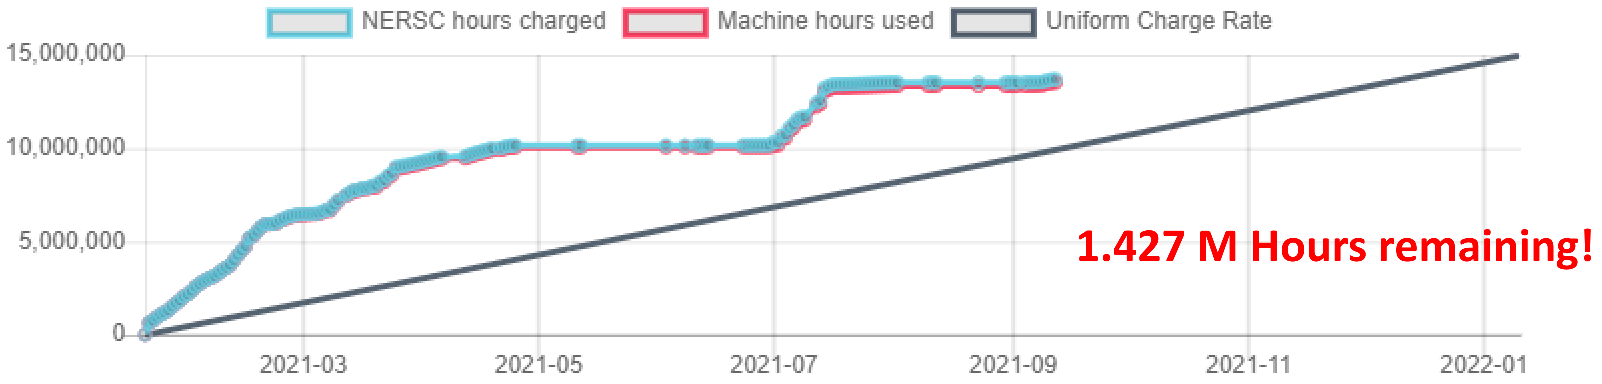


- mp288 received 10M Hrs for CY 2021, Initial allocation exhausted by May 1
- John Mandrekas (DOE) added 5M Hrs additional
- More time may be possible if this is exhausted
- Pearlmutter time will not be charged for this FY
- We are NESAP Tier 2.  Machine not yet ready.  Phase-I w GPUs
- FY2022 ERCAP now open

# Changes to github master since 08/09/21

**S. Jardin:**
**8/24/21:** re baselined RMP_nonlinear regression test due to kappar change

**Seegyoung Seol:**
**8/30/21**: adding partitioned mesh files for regression tests
**9/01/21**: minor changes in readme.stellar and makefile
**9/07/21:** adding makefile and readme for greene (openmpi-4.0.3)

**Nate Ferraro:**
**9/03/21:** changed stellar regtest batch scripts to “source” part_mesh.sh, to fix problem restrictions from intel-mpi
**9/09/21**: Corrected hmn diagnostic for itor=0

**Jin Chen:**
**9/08/21:** add ‘-sub-mat_mumps_icntl_14 50’ in options_bjacobi.type_mumps in regtests/base

# Local Systems

- PPPL centos7(09/12/21)
  - 5 regression tests **PASSED** on centos7:
  - ADAPT **failed**
- PPPL greene (09/12/21)
  - 4 regression tests **PASSED** on greene (m3dc1)
  - No batch file found for pellet
  - ADAPT **failed** (in same way as on centos7)
- STELLAR (09/12/21)
  - 6 regression tests **PASSED** on stellar
- TRAVERSE(09/12/21)
  - Code compiles
  - Regression test failed: split_smb not found in PATH

# Other Systems

- Cori-KNL (9/12/2021)
  - 5 regression tests **PASSED** on KNL
  - ADAPT **failed**
- Cori-Haswell (9/12/2021)
  - 5 regression tests **PASSED** on cori
  - ADAPT **failed**
- MARCONI
  - All regression tests PASSED on MARCONI (J. Chen, 9/04/20)

- **ADAPT only passes on stellar. Fails in same way on other machines**

# Difference in C1ke files for adapt

```
<      0  0.0000E+00  1.6085E-01  1.5011E-10  0.0000E+00  1.6085E-01  0.0000E+00  4.5714E+00  3.2116E+04  5.9487E+00   3.21267E+04
---
>      0  0.0000E+00  1.6085E-01  3.6410E-11  0.0000E+00  1.6085E-01  0.0000E+00  4.5726E+00  3.2125E+04  5.9487E+00   3.21360E+04
```

| | | |
|---|---|---|
| 1.5011E-10 | 4.5714E+00 | 3.2116E+04 |
| 3.6410E-11 | 4.5726E+00 | 3.2125E+04 |
| gamma_gr | emagp | emagt |

gamma_gr does not have any meaning for the 0 time step … can ignore
emagp is the integrated energy in the poloidal field
emagt is the integrated energy in the toroidal field

# Update: Code hangs when writing slice #60—Yao Zhou

9/8/21

Here's an update: I found that the issue was related to the scalar write chunk_size (100). I was using 5 time steps per slice and at the 60th slice (300th step) something strange must have been triggered and I still do not understand how.

My 'fix' is changing the chunk_size from 100 to 101, a prime number, which could lower the probability that N*(chunk size) coinciding with a time slice write. This change seems to have temporarily solved the problem. It is included in the merge request that I just generated. If similar problem reappears we will revisit this.

# DIII-D Mesh request -- Strauss

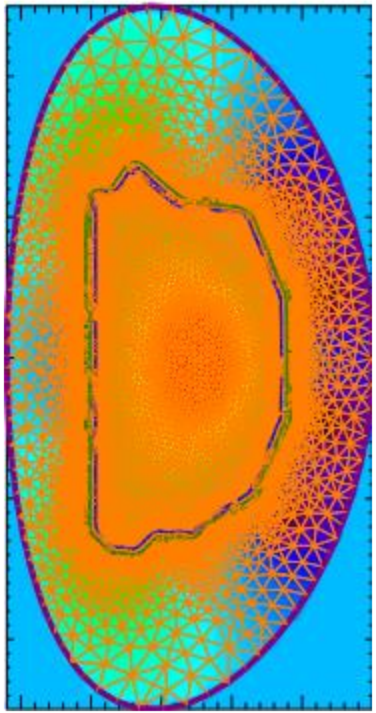

( 1a) Simmetrix

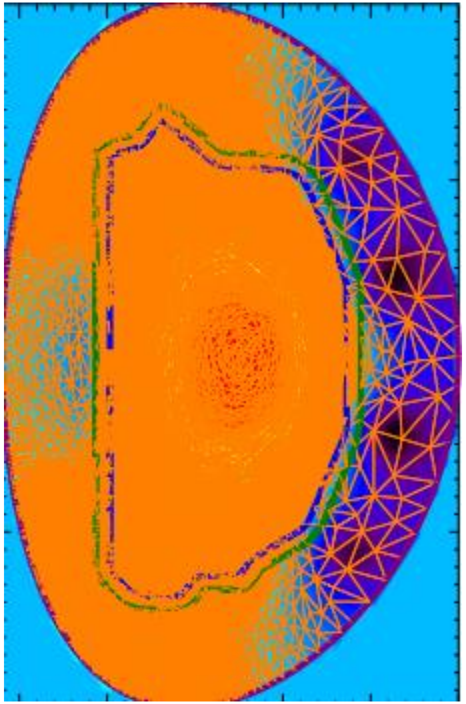

( 1b) adapted

The mesh in Fig.1 is superimposed on a plot of the magnetic flux ψ. The Simmetrix mesh in Fig.1(a) has high resolution in the resistive wall, lower resolution inside the wall, and lowest resolution between the resistive wall and the outer boundary. The adapted mesh in Fig.1(b) was initialized with a coarser mesh than in Fig.1(a), and adapted with with the sizefieldParam file

1.35 2 1 .05 .5 .05 .5 .1 .01 5. 5. 0.3 0.148

The resistive wall is not a contour of the normalized flux function $\tilde{\psi} = 1$. Note that the first parameter is set a1 > 1. This was not enough to cover the resistive wall, and it fails if it is a1 > 1.38. The sizefieldParam adaptation should be modified so the adaptation has an option to be centered on the resistive wall.

# Parallel thermal conduction and thermal quench

**H. Strauss**

Parallel thermal conduction depends strongly on edge temperature $T$, causing TQ time to vary by orders of magnitude. This is because $\chi_\parallel \propto T^{5/2}$ for small $T$, making a transport barrier as $T \to 0$. Model $\chi_\parallel$ has collisional and collisionless limits (mean free path $>$ connection length). In stochastic magnetic field, parallel transport time is $\tau_\parallel$. There might also be a RWTM with TQ time $\gamma^{-1}$. $b_n$ measures edge magnetic perturbations.

$$\chi_\parallel = \frac{\pi R v_e}{1 + \pi R/(2.1 v_e \tau_e)}, \qquad \tau_\parallel = \frac{a^2}{\chi_\parallel b_n^2}, \qquad \frac{1}{\gamma \tau_A} = \frac{1}{c_0} S^{1/3} S_{wall}^{4/9} \tag{1}$$

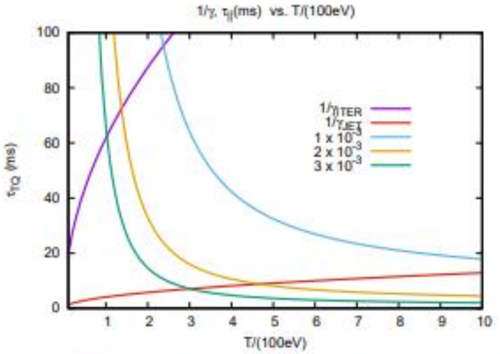


TQ time (ms) $\tau_{TQ}$ as function of T(100eV) with ITER parameters. $1/\gamma$ for ITER and JET, $\tau_\parallel$ [Strauss, 2021] with model (1), $b_n = 10^{-3}$ from simulations and [Devries, 2016], $b_n = 2 \times 10^{-3}$ island width $w = 0.3a$ model, and $b_n = 3 \times 10^{-3}$ [Paz-Soldan, 2020] If edge is collisional, there is no need for pellets, no REs.

**It would be desirable to have model $\chi_\parallel(T)$, analogous to Spitzer resistivity: $\chi_\parallel \propto [(T_1/T)^{5/2} + 1]^{-1}$, where $T_1 \sim 300 eV$.**

# M3D-C1 – LP Code coupling

- ZOOM call held 9/8/21 to discuss M3D-C1 stand-alone and LP loose coupling simulations of a DIII-D shot where a 0.66mm radius pure-neon pellet was injected with a velocity of 179.5 m/s (Experiment by Daisuke .Shiraki)

- Reasonable agreement: within 30% depending on modeling assumptons (regarding where to get the temperature and density of the background plasma)

- Brendan wrote to Daisuke on 9/10/2021 requesting more experimental information

# Chen Zhao paper in preparation

Simulation of the runaway electron plateau formation during current quench

C. Zhao[1], C. Liu[1], S. C. Jardin[1], N. M. Ferraro[1], B. C. Lyons[2], V. Bandaru[3], M. Hoelzl[3]

[1] Princeton Plasma Physics Laboratory, Princeton, NJ, United States of America
[2] General Atomics, San Diego, CA, United States of AmericaGeneral Atomics, San Diego, CA, United States of America
[3] Max Planck Institute for Plasma Physics, Boltzmannstaße, Garching, Germany

E-mail: czhao@pppl.gov

- Now only contains formulation and 2 test problems (1 cylindrical and 1 with JOREK)
- No section on experimental comparisons or on sawtooth
- Need some discussion on validity of Dreicer model (from Chang)
- Add section on comparison with characteristics model of advancing runaways?

# Clauser paper accepted by Nuclear Fusion

## Modeling of carbon pellets disruption mitigation in an NSTX-U plasma

C F Clauser[1,2], S C Jardin[2], R Raman[3], B C Lyons[4], N M Ferraro[2]

[1]Lehigh University, Bethlehem, Pennsylvania 18015, USA
[2]Princeton Plasma Physics Laboratory, Princeton, New Jersey 08543, USA
[3]University of Washington, Seattle, Washington 98195, USA
[4]General Atomics, San Diego, California 92121, USA

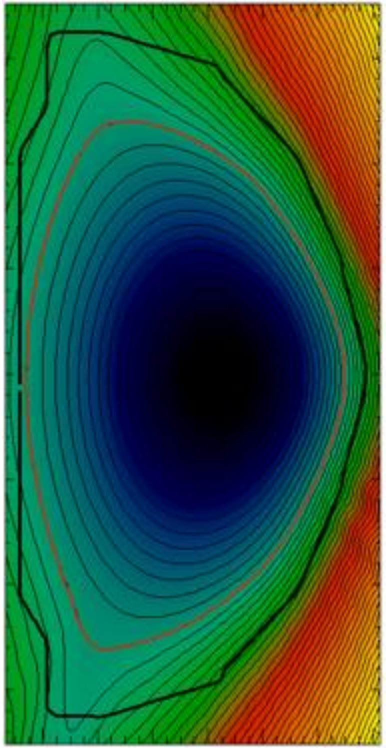

- May be some follow-on work as Roger Raman has gotten approval to start some Electro-magnetic Pellet Injection (EPI) work at PPPL

- Suggested looking at NSTX-U #203679

# Chang Liu paper accepted by PPCF

Self-consistent simulation of resistive kink instabilities with runaway electrons

Chang Liu, Chen Zhao, Stephen Jardin, N. Ferraro, Carlos Paz-Soldan, Yueqiang Liu, Brendan Lyons

# 1/1 mode with negative vloop on EAST

- Email from Liqing Xu on 8/31/21
- Recently, we found some interesting data in EAST toroidal loop voltage feedback control experiment.
- The target plasma has a very strong central electron heating by LHW and ECRH. R/L_{Te}~10 in the core. The plasma current is almost non-inductive, say I_{ohmic}~0.
- In Vloop FB experiment, a 1/1 mode is present when Vloop slightly below zero, and disappear when Vloop above zero, as shown in the attached figure.
- I don’t know the possible mechanism about Vloop~0 and 1/1 mode. Do you have any suggestion about this? Can we work together about it?

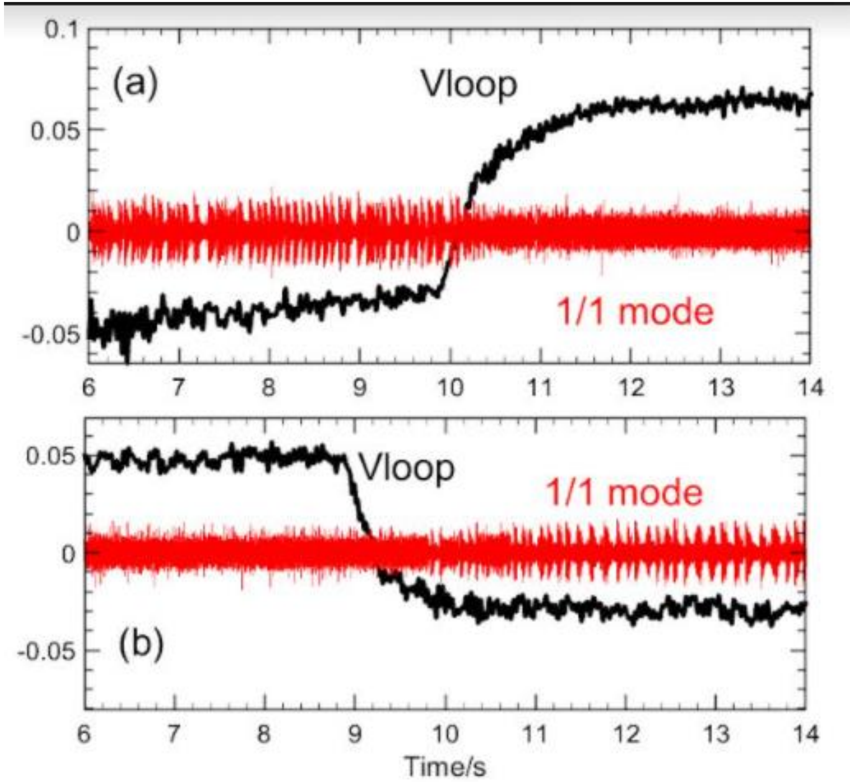

# Poloidal Velocity Boundary Conditions

**No normal flow**

should be

$$R^2\nabla U \times \nabla \phi \bullet \hat{n} + R^{-2}\nabla_\perp \chi \bullet \hat{n} = 0$$

we now have

$$R^2\nabla U \times \nabla \phi \bullet \hat{n} = 0$$

$$R^{-2}\nabla_\perp \chi \bullet \hat{n} = 0$$

**No slip**

should be

$$R^2\nabla U \times \nabla \phi \bullet \hat{t} + R^{-2}\nabla_\perp \chi \bullet \hat{t} = 0$$

we now have

$$R^2\nabla U \times \nabla \phi \bullet t = 0$$

$$R^{-2}\nabla_\perp \chi \bullet \hat{t} = 0$$

Brendan has found that in some pellet injection cases, this can lead to a large reversed flow at the edge, in disagreement with the NIMROD result

# Toroidal flow with density source

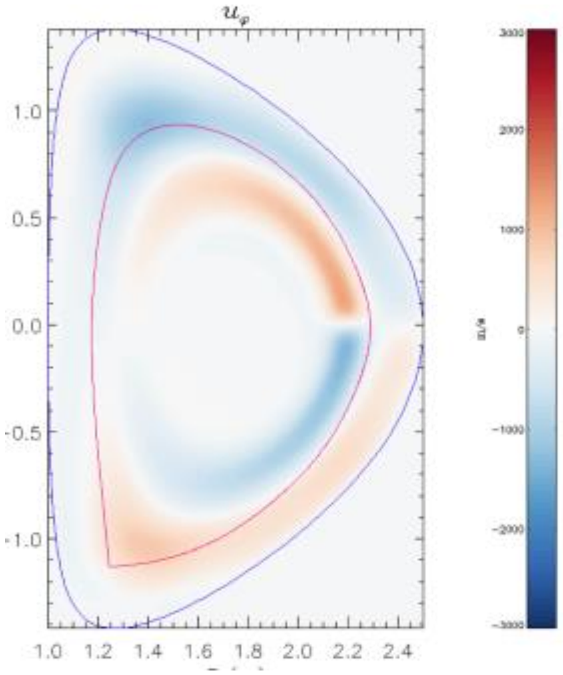


# That’s All I have

Anything Else ?