

M3D-C1 ZOOM Meeting

08/01/2022

Upcoming meetings and deadlines

CS Issues

1. LBL Report
2. Solver problem on Perlmutter_cpu
3. New Latex documentation
4. New meshing capabilities requested
5. Mesh adaptation update -
6. Perlmutter request and NERSC Time
7. Changes to github master since last meeting
8. Regression tests

Physics Studies

1. Status of Chen Zhao paper
2. Upgrade to impurity radiation model – Brendan Lyons
3. Problem with BC in ITER problem
4. Stellarator version convergence
5. Anything else

Upcoming Meetings

OpenACC and Hackathons Summit 2022

August 02-04

C. Liu to present talk on M3D-C1

J. Chen to attend

International Conference on Numerical Simulation of Plasma

August 30-September 2 2022 – online only

APS-DPP

Oct 17-21 In Spokane Washington

Jardin & Kleiner to give invited (only ones in PPPL theory?)

26th Workshop on MHD stability Control is on Oct 14-15

CTTS SciDAC meeting is on Sunday Oct 16 (TBA)

LBL Progress on Solvers

Solver Problems on Perlmutter_cpu

I have a large mesh (98K faces) with 128 partitions that I would like to run on Perlmutter_cpu. The base directory is:

`/global/cfs/cdirs/mp288/Jardin/m3dnl/Perl_cpu/128-K`

- Runs ok with mumps with 4 planes with 32 CPUs/NODE (Run01) and 64 CPUs/NODE (Run02)
- Same runs FAIL with superlu_dist (Run01a and Run02a)
- The following runs all failed with PETSc error

Run	cpu/N	planes	solver
Run03	64	36	mumps
Run04	64	36	superlu
Run05	64	8	mumps
Run06	32	8	mumps
Run07	32	32	mumps

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".

Updated since 7/11/22

mesh-gen.tex

doc.tex

8.3 Grad-Shafranov Solver

8.4 Transport Coefficients

8.5 Hyper-Diffusivity

8.6 Normalizations

8.7 Boundary Conditions (partial)

Other C1input sections being worked on

New Meshing Capabilities Requested

3/28/22: Nate requested extension of M3DC1_meshgen with more regions

6/21/22: Mark suggested extending ToMMs (which XGC uses) instead

Nate noted that we need different info than XGC for electromagnetics

6/27/22: Mark proposed an extension to m3dc1_meshgen for an arbitrary # regions :

- # of nested closed loops
- The geometric definition of the first being the vacuum loop that is an indicated analytic expression and its parameters
- The geometry of the remaining loops will be a set of discrete points one loop at a time going from the outside to the inside
- # of island loops
- For each island loop indicate the two nested loops it lies between. Give its geometry in terms of a discrete set of points

Mesh adaptation update (and make update)

Any update?

Perlmutter Request

What application codes does your project use today?

M3D-C1, M3D-C1-S, M3D-C1-K

Is your project ready to run those codes productively and efficiently on Perlmutter's GPU-accelerated nodes?

Yes, M3D-C1-K (kinetic MHD) runs efficiently on Perlmutter_GPU

Do you have additional codes that you plan to run on Perlmutter GPU? What are they?

We may run M3D-C1-S (stellarator version) but it requires development

Are there additional barriers that are preventing you from using Perlmutter's GPUs effectively?

We use PETSc solver GMRES with block Jacobi PC using SUPER_LU_DIST.

This is not fully GPU optimized

Are there things NERSC could do to help you be more productive?[1]

Can queue show estimated start times?


Do you have any other general comments, questions, or suggestions?

IRIS does not accurately reflect our GPU allocation. (7000 vs 0 N-H)

Sent by NERSC on 7/25, request reply by 8/05

NERSC Time

mp288



- MP288 usage is on track. Both value and rate are ok.
- All users now have access to Perlmutter_cpu. There is no charge for the year!
 - Some issues for large problems...see future slide
- NERSC Checkpoint/Restart Requirements Gathering Workshop, July 12-13
 - Jin Chen attended. Any comments?
- We should have 7000 GPU node hours as well. (Doesn't show in iris)

Changes to github master --after 2022-7-11

Steve Jardin:

07/18-25/22: Updated Sections 8.3,8.4 of Latex documentation

Nate Ferraro

07/12/22: Updated Perlmutter reg test scripts to use mp288

07/12/22: Increased mumps memory allocation in pellet regtest on Perlmutter

07/14/22: Updated plot_field_vs_phi to work correctly now that we can read multiple toroidal angles in plot_field at once

Seegyong Seol

07/11/22: Added polar_meshgen in Users Guide

07/16/22: Adding config files for SCOREC RHEL7

07/18/22: Adding mesh generation program for multi-region and polar

07/18/22: Adding test case for multi-region mesh generation

Local Systems

- PPPL centos7(07/29/22)
 - 7 jobs **PASSED**
- PPPL greene (07/29/22)
 - 5 jobs **PASSED**
- STELLAR (07/29/22)
 - 7 regression tests **PASSED** on stellar
- TRAVERSE_gpu(08/01/22)
 - 5 regression tests **PASSED**
 - KPRAD_2D, KPRAD_restart **FAILED** due to 0.001 fractional diff in C1ke

NERSC

- Cori-KNL (07/29/2022)
7 regression tests **PASSED**
- Cori-Haswell (07/29/2022)
7 regression tests **PASSED**
- Perlmutter (07/29/2022)
6 regression tests **PASSED**
NCSX **FAILED** with “PC failed due to SUBPC_ERROR”
- Perlmutter_cpu (07/29/22)
6 regression tests **PASSED**
NCSX **FAILED** with “PC failed due to SUBPC_ERROR”

Papers in Preparation

- Chen Zhao, C. Liu, et al, “Simulation of DIII-D disruption with pellet injection and runaway electron beam”
 - ***New Version circulated 7/10/22 --- Jardin proofed and sent back to Chen on 7/28/22. Shared file (too big for email). Will send corrections for final section separately.***

Upgrade to impurity radiation Routines

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

Any Response?


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?



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.

That's All I have

Anything Else ?

Next Meeting Aug 15

No Meeting Sept 5 (Labor Day)

In attendance

Steve Jardin

Jin Chen

Hank Strauss

Andreas Kleiner

Chang Liu

Chen Zhao

Nate Ferraro

Brendan Lyons

Mark Shephard

Seegyong Seol

Morteza

Usman

Sam Williams

Sherry Li

Yang Liu

Nan Ding