M3D-C1 ZOOM Meeting

05/09/2022

Upcoming meetings and deadlines

CS Issues

1. Mesh adaptation and make update -
2. NERSC Time
3. Changes to github master since last meeting
4. Regression tests
5. PETSC error on cori Haswell and timing tests

Physics Studies

1. Chen RE paper
2. Bateman scaling app

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

In attendance

Steve Jardin
Adelle Wright
Hank Strauss
Nate Ferraro
Chang Liu
Anders Kleiner
Brendan Lyons
Jin Chen
Chen Zhao
Mark Shephard
Seegyoung Seol

Upcoming Meetings and Deadlines

- IAEA Technical Meeting on Plasma Disruptions and their Mitigation 19-22 July
  - In person at ITER HQ in France
  - Abstract submission by May 31

- INSITE requests for FY 23 now open for Frontier(EF), AORORA(EF), Polaris(44PF), Summit(200PF), Theta(12PF)

Mesh adaptation update (and make update)

Email from Brendan to Seegyoung 05/04/22

The idea is that by typing **make scorec** and **make scorec COM=1** from unstructured, it would go into m3dc1_scorec and build the libraries needed for whatever system is specified by $M3DC1_ARCH.
It just utilizes the existing make infrastructure in m3dc1_scorec, so all we need to do is specify in the .mk file which config.sh file to use for that system.
Then, instead of linking to precompiled libraries on each system, we just link to these newly built libraries.
This should greatly simplify maintaining builds across systems.
While you and/or Jin would still need to build the underlying libraries (like PETSc), you wouldn't need to build the m3dc1_scorec libraries every time a change is made.

NERSC Time

mp288

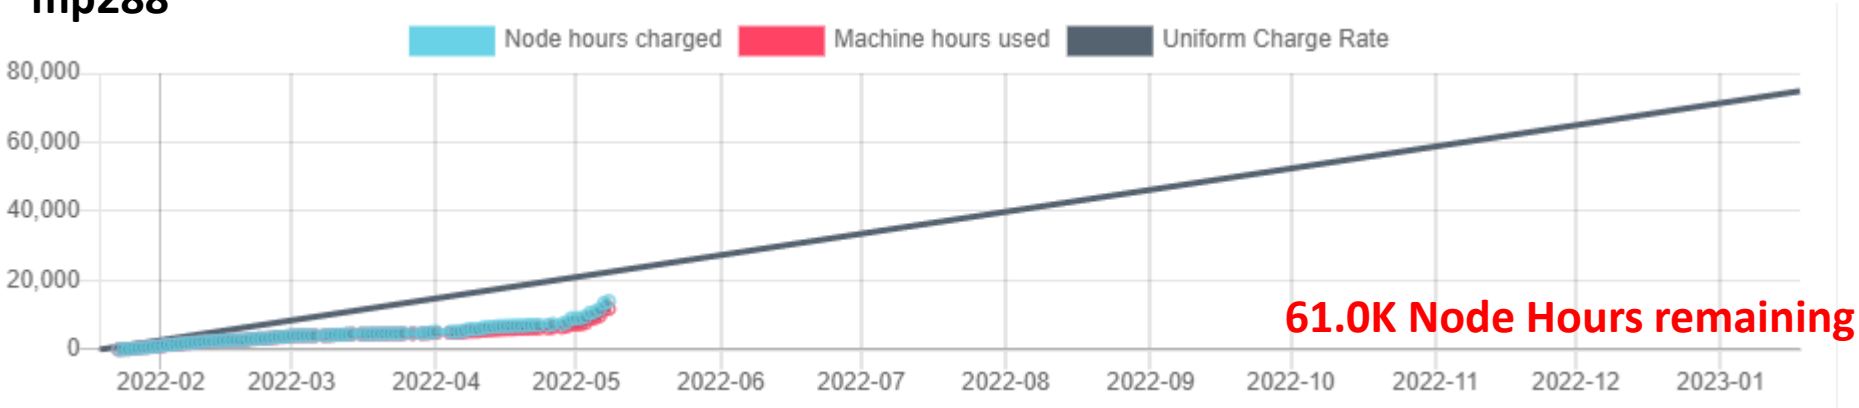

- We are NESAP Tier 2 for Pearlmutter. . Phase-I w GPUs We have been given a repo m3984 with a small allocation. Presently we are not being charged.
- N9ES-N2 M3D-C1: J. Chen , C. Liu, S. Seol are early users

- We are under-utilizing our time on cori. We will lose 17K Node Hours on May 19th unless we have used more than 5% of initial allocation. **(Have now used 18% !)**
- Next Reduction is Sept 15, when we should have used 25%

Changes to github master since 4/25/22

Seegyoung Seol
04/25/22: adding m3dc1_ent_measure
04/28/22: elaborating readme.Perlmutter
05/04/22: updating config.sh for PPPL greene & centos7
05/04/22: updating config.sh for Princeton machcines

Nate Ferraro
04/29/22: Disabled writing wall region data in 3D. This hangs for some reason; maybe creating the different groups is expensive with many processes.

Local Systems

- PPPL centos7(05/08/22)
  - 7 jobs PASSED
  - NCSX FAILED on first try (0.3% difference in KE), then passed
- PPPL greene (05/08/22)
  - 5 jobs PASSED
- STELLAR (05/08/22)
  - 7 regression tests **PASSED** on stellar
- TRAVERSE_gpu(05/08/22)
  - 5 regression tests **PASSED**
  - KPRAD_2D, KPRAD_restart FAILED due to 0.001 fractional diff in C1ke
  - NCSX FAILED on first try (0.17% difference in KE), then passed

NERSC

- Cori-KNL (05/08/2022)
  6 regression tests **PASSED**
  NCSX **FAILED** with segmentation fault

- Cori-Haswell (05/08/2022)
  7 regression tests **PASSED**
  NCSX failed on first try with DIVERGED iteration

- Perlmutter (05/08/2022)
  6 regression tests **PASSED**
  RMP_nonlin failed first time, but passed on resubmission
  NCSX **FAILED** with Segmentation violation

PETSc error on cori Haswell

I tried rerunning a job I did last year on Haswell with 32 tasks per node. Last year, it ran fine for thousands of timesteps. Now it fails at a random time (one case 6, another 30+ timesteps)j. It fails with the error:

2337]PETSC ERROR: PetscAbortErrorHandler: MatFactorNumeric_MUMPS() at /global/cfs/cdirs/mp288/jinchen/PETSC/petsc.20220107/src/mat/impls/aij/mpi/mumps/mumps.c:1688 Error reported by MUMPS in numerical fa >
To prevent termination, change the error handler using PetscPushErrorHandler()

/global/cfs/cdirs/mp288/jardin/m3dnl/Brendan/Brendan32f/Run01b

Jin found that this error goes away by switching from mumps to superlu_dist

Runs on Cori

Good runs:

vertices	partitions	V/P	planes	nodes	t-per-n	sec/TS	n-sec/TS	Solver
9358	128	73	32	128	32	60	7680	superlu
19965	32	623	24	96	8	200 (70)	19200	mumps
19965	64	312	24	96	16	170(100)	16320	superlu
19965	64	312	24	192	8	140(75)	26880	
19965	128	156	24	192	16	110(70)	21120	
19965	256	78	24	192	32			
No good								
19965	32	623	24	48	16			
19965	128	156	24	96	32			

16 tasks-per-node better than 8 (for the same # of partitions) 32??

Fewer partitions better than more for the same number of tasks-per-node (if it runs)

Papers in Preparation

- Chen Zhao, C. Liu, et al, “Simulation of DIII-D disruption with pellet injection and runaway electron beam”
  - ***Brendan:*** Can you include the equations and text for pellet ablation that were used in the paper status?
  - ***Chang:*** Can you look at the description of the runaway sources in the paper and correct? Status?

  - ***Chen:*** Is paper ready for another read?

Bateman Scaling App

Email from Brendan 05/04/22:

I implemented the Bateman scaling as an OMFIT function operating on any gEQDSK file. You can give it an arbitrary BCENTR and CURRENT (in gEQDSK parlance) and it will scale the equilibrium accordingly

. For example, here's a case where I took g163303.03170 (blue in attached), which has BCENTR = **-2.0101074** & CURRENT = **1182404.4**, and asked for BCENTR=**2.5** & CURRENT=**-2E6** (orange). I then loaded the new equilibrium into EFIT and resolved the equilibrium with the new field, current, p', and FF'. The result (green) is unchanged from the scaling.

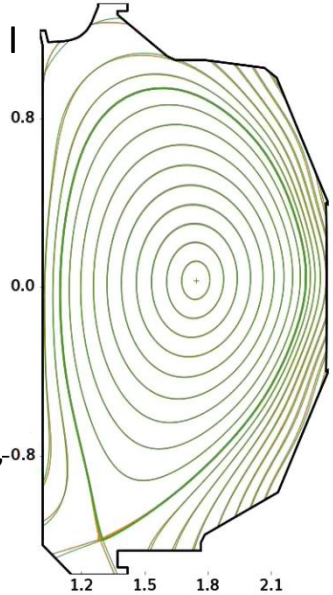

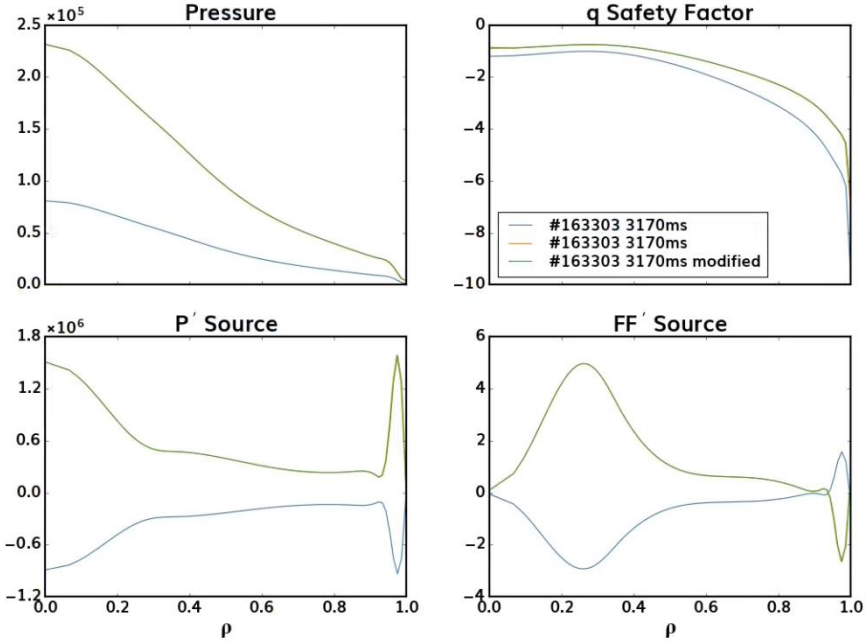

That’s All I have

Anything Else ?

Next Meeting (with LBL) June 6

I will be on vacation May 11-16
Monday May 23-25 is Exascale workshop
Monday May 30 is Memorial Day