

M3D-C1 ZOOM Meeting

10/04/2021

Announcements

CS Issues

1. Questions for LBL group
2. NERSC Time
3. Changes to github master since last meeting
4. Regression tests
5. Update to `adapt_by_field` -- Morteza

Physics Studies

1. New code version became unstable on some problems
2. Test of `ikapparfunc=1` ..Hank Strauss
3. Planning for next 5 years

Note: [meeting minutes posted on m3dc1.pppl.gov](https://m3dc1.pppl.gov)

In attendance

Steve Jardin

Adelle Wright

Jin Chen

Andreas Kleiner

Nate Ferraro

Chen Zhao

Brendan Lyons

Chang Liu

Hank Strauss

Mark Shephard

Seegyoung Seol

Usman Riaz

Morteza Siboni

Sam Williams

Sherry Li

Nan Ding

Yang Liu

Announcements

- ITPA Meeting 11-15 October
- John Mandrekas requested a 90 min presentation from all SciDACs
 - CTTS now scheduled for Oct 22 at 1:00 PM ET
- APS Nov 8-12
 - Meeting will be IN PERSON with virtual option. Will DOE allow travel?
- EPS 6/27 – 7/1 2022 in Maastricht, Netherlands
 - Nominate invited speakers by 29 October 2021

Questions for LBL group

Chang Liu:

I wonder whether there is a plan to do MatSolve on GPU. In my tests, I found that when setting the matrix type to mpiaijcuspars, then superlu_dist can do the factorization using GPU, but the iterative solver is still running on CPU, which can involve CPU-GPU communications and is slow.

2. Related to the previous question, I wonder if it is possible to offload the matrix completely after the factorization, so that communication during the iterative solving is not needed.


3. I tried to use superlu_dist on GPU with MPI. I found that if the mpi ranks is equal to the number of GPUs, then each rank will occupy a GPU. For more MPI ranks, multiple ranks will share a GPU. I wonder what is the best strategy then, should I set the rank to be equal to # of GPUs?

Jin Chen

How does superlu_dist interact with ptscotch ? What would happen without it?

NERSC Time

mp288



- mp288 received 10M Hrs for CY 2021, + 5M Hrs additional
- Pearlmuter 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 and C. Liu will be early users

Changes to github master since 09/27/21

Nate Ferraro:

9/27/21: Updated centos7.mk to include link to netcdf fortran library and module

9/29/21: Minor change to debug output formatting to account for fact that curv has 3 elements
Declared gsl functions as “external” to avoid compiler warnings
Changed BC of fp back to fp=0 unless ST=1

Jin Chen:

9/28/21: stellar readme file and regtests option files update

Brendan Lyons:

9/30/21: Allow for coarser mesh refinement above/below certain Z values (useful for reducing unnecessary mesh refinement near Z-points)

Chang Liu:

09/30/21: Increase time limit of regtests for traverse

10/01/21: fix the wrong compiler flags in traverse_gpu.mk

10/02/21: set OMP_NUM_THREADS to fix timeout error

Yao Zhou

10/02/21: Removed quadruple precision in read_vmec

Local Systems

- PPPL centos7(10/01/21)
 - 6 regression tests **PASSED** on centos7
 - adapt **FAILED**
- PPPL greene (10/01/21)
 - All tests **FAILED** --- error partitioning the mesh
- STELLAR (10/01/21)
 - 7 regression tests **PASSED** on stellar
- TRAVERSE(10/04/21)
 - 6 regression tests **PASSED**
 - adapt **failed**

Other Systems

- Cori-KNL (10/01/2021)
 - 6 regression tests **PASSED** on KNL
 - ADAPT **failed**
- Cori-Haswell (10/01/2021)
 - 5 regression tests **PASSED** on cori
 - ADAPT **failed**
 - NCSX failed
- MARCONI
 - All regression tests PASSED on MARCONI (J. Chen, 9/04/20)
- **ADAPT fails in same way on Haswell, KNL, centos7, traverse**


New adapt option

Author: Morteza H. Siboni <hakimm2@rpi.edu>

Date: Tue Sep 21 16:16:15 2021 -0400



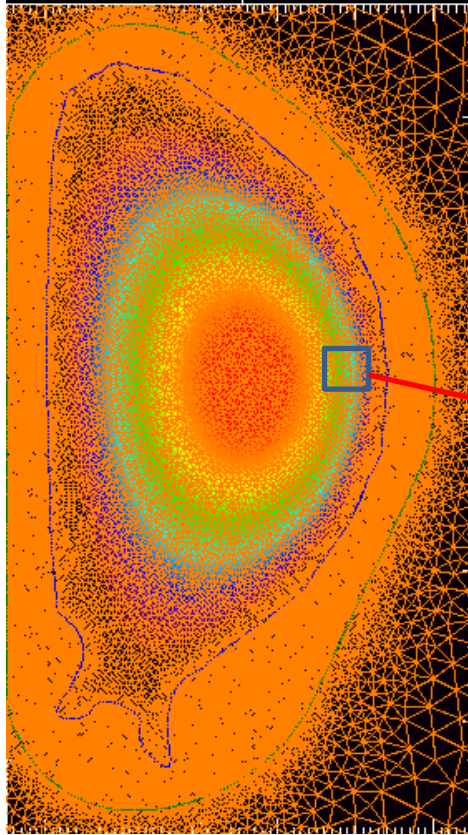
Updates the new logic for `adapt_by_field`

The 14th parameter in `sizefieldParam` (if exists) should be either 0 or 1 and with this change the following behaviours can be expected

- (1) if there are 13 parameters things will work as before
- (2) if there are 14 parameters the last parameter should be either 0 or 1 (any other value will cause an error).
 - (2a) value of 1 will leave coarsening "on"
 - (2b) value of 0 will turn coarsening "off"
- (3) if there are more than 14 or less than 13 parameters in "sizefieldParam" this will cause an error. 

1. 2. 2. .01 .4 .01 .4 .1 .1 .01 .02 .05 .5 0

Testing on ITER equilibrium -- 1






/p/tsc/m3dnl/ITER/NewMesh/Eq2

and .../Adapted

Refines plasma region ok

Testing on ITER equilibrium -- 2

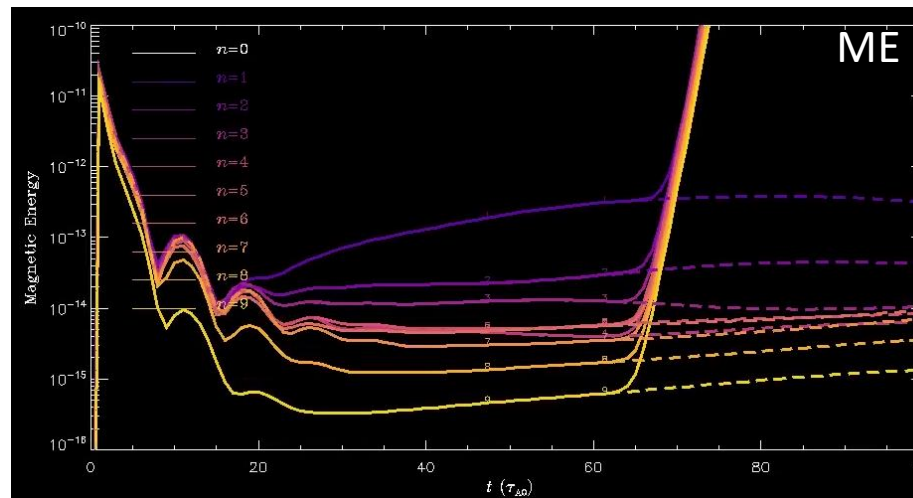



/p/tsc/m3dnl/ITER/NewMesh/Eq2

and .../Adapted

Leaves wall zones untouched!


New Code version becomes unstable on some problems



Solid: /home/sjardin/data/Ideal/3D-2F6
Dashed: /home/sjardin/data/Ideal/3D-2F6b
Also fails at NERSC

Sept 22 git master version
Sept 16 code version

Another unstable case



/home/sjardin/data/ITER/Run05NM-redo-2-3D good (Sept 16 version)
/home/sjardin/data/ITER/Run05NM-redo-2-3DF fails (Sept 22 version)

Debugging status

1. Tried swapping out lundef_t.f90 ... did not help
2. Next routines to swap out: (changed after 9/16...in time loop)
basic_mesh.f90 scorec_mesh.f90
biharmonic.f90 scorec_vector.f90
boundary.f90 spline.f90
element.f90 transport.f90
m3dc1_nint.f90
metricterms_new.f90
model.f90
newpar.f90
petsc_vector.f90

update

3. Next routines to swap out: (changed after 9/16...in time loop)


use new versions

scorec_mesh.f90 model.f90
scorec_vector.f90 newpar.f90
spline.f90 petsc_vector.f90
basic_mesh.f90 element.f90
boundary.f90

use 9/16 versions

lundef_t.f90
metricrms_new.f90
transport.f90
m3dc1_nint.f90

Unstable!



changes

- scorec_mesh.f90: in subroutine _local_dofs:
changed: call rotate_dofs(temp(j:k), dof(j:k), norm, 0., -1)
to: call rotate_dofs(temp(j:k), dof(j:k), norm, curv, -1)
- scorec_vector.f90: change curv to curv(3) in 4 routines only
- spline.f90: in subroutine get_hermite_coefs
changed: a(1) = s%y(s%n)
to: if(ex.eq.0) a(1) = s%y(s%n)
- basic_mesh.f90: in subroutine load_sms_mesh
changed: nodes(i)%curv =
to: nodes(i)%curv(1) =
- boundary.f90: chane curv to curv(3) in 14 routines only

Changes (cont)

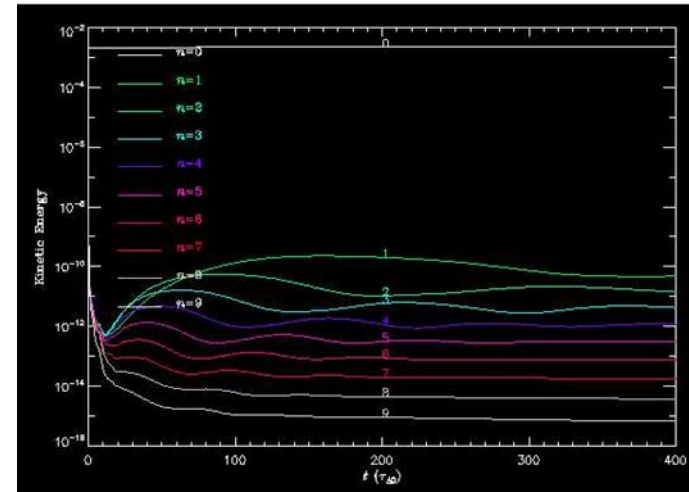
model.f90 in subroutine boundary_mag
changed: temp = 0
to: call get_node_data(bfp_field(1), I, temp)

} Reversed this and the
problem went away!


newpar.f90 curv → curv(1) in 6 places

petsc_vector.f90 changed curv to curv(3) in 4 routines

element.f90 changed curv to curv(1) in 30 places



Strauss test of ikapparfunc=1



Initial temperature profiles is unphysically jagged, goes negative

Planning for next 5 years


- We should start thinking about what should be proposed for the next SciDAC
- Both from the Physics and CS sides
- We could start a discussion via email

That's All I have


Anything Else ?

DIII-D Pellet injection case goes unstable (without RE)

plot_hmn



plot_field, 'kprad_totden'



Could the impurity density be going negative?
See /scratch/gpfs/cz12/kprad2_test

Chen Zhao

Chen Zhao paper in preparation

Simulation of the runaway electron plateau formation during current quench

C. Zhao¹, C. Liu¹, S. C. Jardin¹, N. M. Ferraro¹, B. C. Lyons²,
V. Bandaru³, M. Hoelzl³

¹ Princeton Plasma Physics Laboratory, Princeton, NJ, United States of America

² General Atomics, San Diego, CA, United States of America
General Atomics, San Diego, CA, United States of America

³ Max Planck Institute for Plasma Physics, Boltzmannstraße, Garching, Germany

E-mail: czhao@pppl.gov

- Now only contains formulation and 2 test problems (1 cylindrical and 1 with JOEKE)
- 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?