

M3D-C1 ZOOM Meeting

06/05/2023

Upcoming Meetings

CS Issues

1. LBL Report -- LBL
2. Preliminary Frontier results – Chang Liu
3. Adaptation update --RPI
4. Reduced precision SuperLU ...Jin Chen
5. Perlmutter_cpu update
6. NERSC Time
7. Changes to github master since last meeting
8. Regression tests
9. Debug for Dingyun's application

Physics Studies

1. Double Tearing Mode in NSTX
2. Anything else

In attendance

Steve Jardin

Hank Straus

Chang Liu

Jin Chen

Brendan Lyons

Cesar Clauser

Priyanjana Sinha

Chen Zhao

Andreas Kleiner

Nate Ferraro

M. Yoo

Mark Shephard

Seegyong Seol

Usman Riaz

Nan Ding

Yang Liu

Sherry Li

Hans Johansen

Upcoming Meetings

EPS	July 3-7	Bordeaux, France
TSDW	July 19-21	Princeton, NJ
IAEA	Oct 16-21	London, UK
APS	Oct 30 – Nov 1	Denver, CO
AAPPS-DPP	Nov 12-17	Nagoya, JP

LBL Update

LBL

Preliminary Frontier Results

Chang Liu

Adaption Update

RPI?

Reduced Precision SuperLU

On June 5, Sherry writes:

Jin,

Did you use runtime option: `-mat_superlu_dist_single_precision`

The code looks correct, i.e., correctly using `sizeof(float)` in the memory tracker for single precision. I also did a standalone test, using single and double examples in `EXAMPLE/` dir:

```
$ mpiexec -n 1 pddrive big.rua
```

```
** NUMfact space (MB): (sum-of-all-processes)
```

```
L\U :      1.98 | Total : 10.24
```

```
$ mpiexec -n 1 psdrive big.rua
```

```
** NUMfact space (MB): (sum-of-all-processes)
```

```
L\U :      1.03 | Total : 5.44
```

You can see that single precision uses about half of the memory.

I suspect that you did not use the single precision version.

Sherry

Jin replies (same day):

I did use `"-mat_superlu_dist_single_precision"`. To be sure it's passed to PETSc, I have used `options_left` to check it. PETSc didn't report that it was unused


perlmutter_cpu update

- Large jobs with 73 K vertices(N) and 98 K vertices (L) with 380 vertices/partition
 - 192-N-09 with 192 partitions, 380 v/p, 64 planes, 64 cpu/node: runs ok!
 - 192-N-10 with 192 partitions, 380 v/p, 32 planes, 64 cpu/node: runs ok!
 - 128-K-09 with 128 partitions, 386 v/p, 64 planes, 64 cpu/node: runs ok!
 - 128-K-10 with 128 partitions, 386 v/p, 32 planes, 64 cpu/node: runs ok!
 - 256-L-09 with 256 partitions, 384 v/p, 64 planes, 64 cpu/node: segmentation v
 - 256-L-10 with 256 partitions, 384 v/p, 32 planes, 64 cpu/node: runs ok!
- Smaller meshes with 9 K vertices (B) and 19 K vertices (H)
 - Normally run ok with 128 cpu/node and 150-200 vertices/partition
 - Also run ok with 64 cpu/node and 300-400 vertices/partition

Thanks to Jin for correcting my options_bjacobi file with `-sub_mat_mumps_icntl_14 100`

NERSC Time 2023

mp288



- MP288 usage rate is a bit high but leveling off
- Also, 8.0k GPU node hours
- Cori is gone!
- I have contacted DOE to see the likelihood of getting more time – no response to date

Changes to github master --after 2023-05-21

Yao Zhou:

06/02/23: Fixed bug on tepsifkappar

Nate Ferraro:

05/23/23: Some changes to the random perturbation initialization routine to use the correct values of ψ when external coils are present

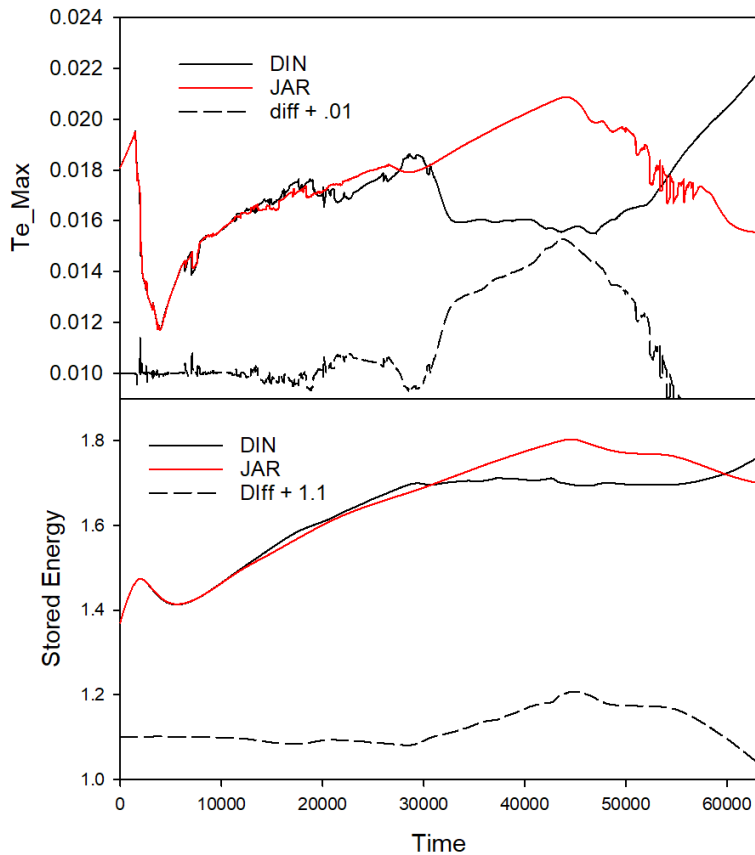
Local Systems

- PPPL centos7(06/04/23)
 - 7 jobs **PASSED**
- PPPL greene (06/04/23)
 - 5 jobs **PASSED**
- STELLAR (06/04/23)
 - 7 regression tests **PASSED** on stellar
- TRAVERSE_gpu(11/04/22)
 - Compilation error (being looked at by Seegyound , Jin, and Chang)

NERSC

- Perlmutter_cpu (06/04/23)
6 jobs **PASSED**
NCSX failed with very small difference in C1ke
- Perlmutter_gpu (02/19/2023) **06/04/23: will not compile**
 - pellet, RMP, & RMP_nonlin, adapt all **PASSED**
 - KPRAD_2D, KPRAD_restart, NCSX all failed with very small differences
 - adapt not submitted

Dingyun is trying to reproduce a case in my 2020 paper



Need to find when these differences first occur:

```
git log --after 2018-06-30 > logout
```

```
git clone
```

```
https://github.com/PrincetonUniversity/M3DC1
```

```
git checkout 8a6a036
```

Had to:

- remove `wrrestartglobal` from `output.f90` and `restart.f90`
- Comment out `MPI_Comm_split` in `newpar.f90`
- Copy `stellar.mk` from current version

Progress on debug

Differences first appear on 11/17/20

- Mostly concerned with bf to bfp change $\text{bf} = f$, $\text{bfp} = f'$



$$\mathbf{B} = \nabla \psi \times \nabla \varphi - \nabla_{\perp} f' + F \nabla \varphi$$

$$R^2 \nabla \bullet \nabla_{\perp} f = F - F_0$$



$$R^2 \nabla \bullet \nabla_{\perp} f' = F'$$

- The two versions agree exactly in 2D !
- The two versions agree fairly well in 3D if $f = f' = 0$, and these results can be very different from f and f' nonzero
- Differences can appear in this and other problems at long time
- These differences likely due to differences in truncation error. Convergence studies should be performed



Run49, fp, op=1
11_17_20: 17:00
Commit: b572e1e7



$t=4500$



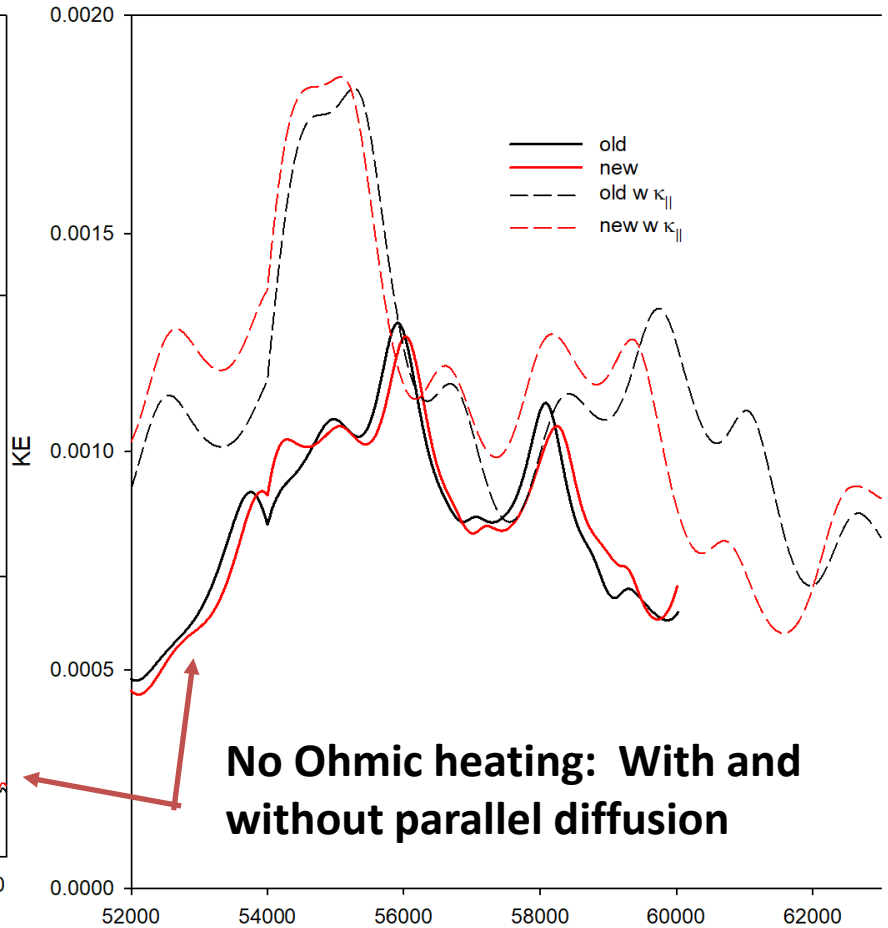
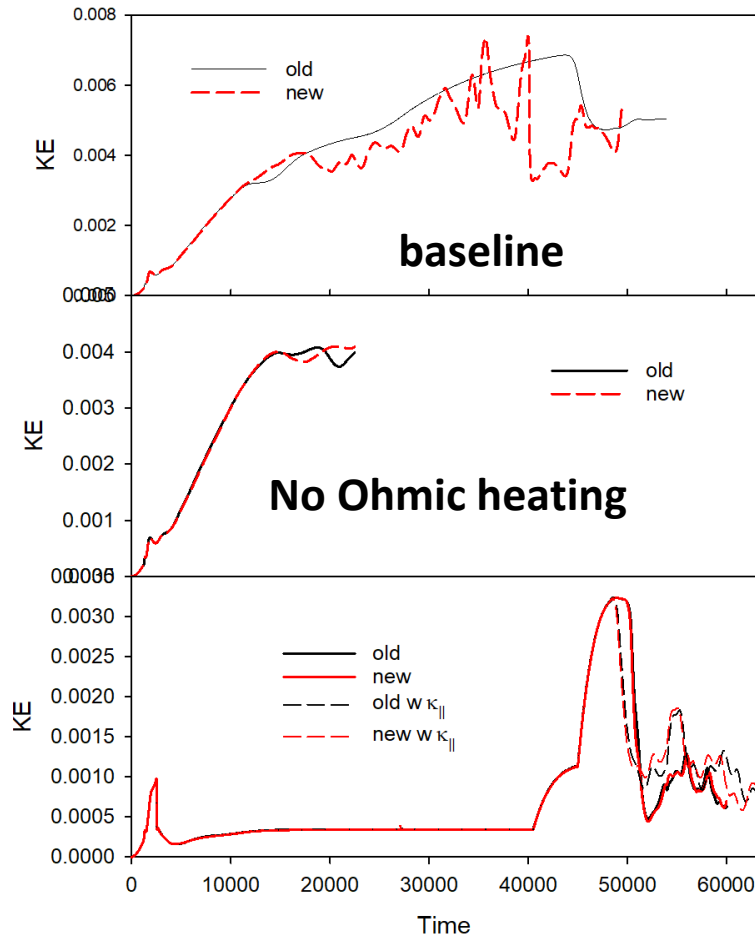
$t=9000$




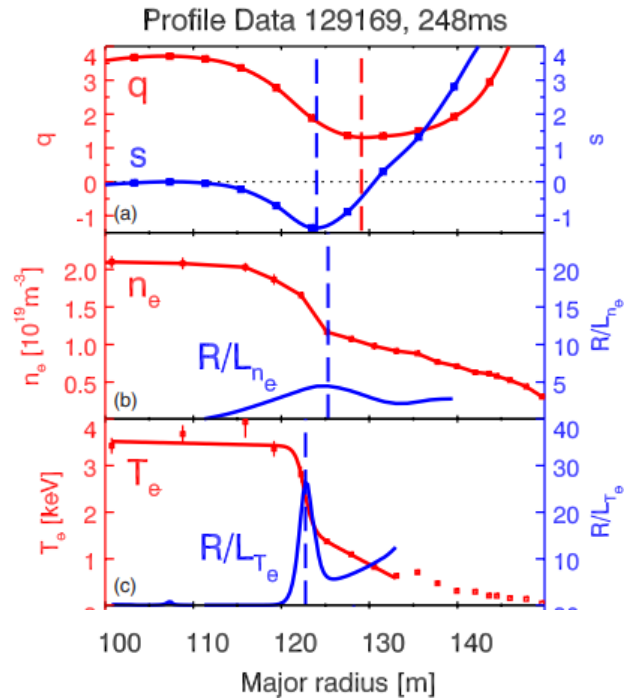
$t=13500$

Run48, f, op=11
11_17_20: 14:48
Commit: 43061986

Try turning off certain terms




Double Tearing Mode in NSTX




From: Yuh, et al, Phys. Plasma (2009)

Some MHD activity starting at about 270 ms


DTM Eigenfunction and growth rate



Linear χ at 247 ms













q -profile in geqdsk file is smooth. However, when we recompute the equilibrium in M3D-C1, oscillations appear

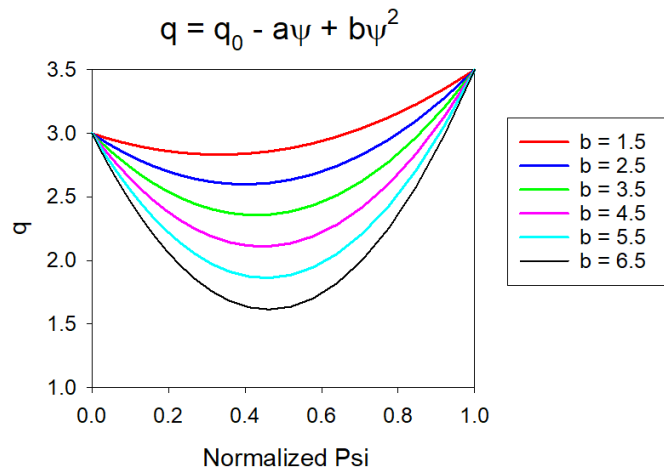


Linear growth rate as a function of linear mesh size for spitzer η and spitzer x 10

NSTX Reversed Shear shot 129169 @ t=247 ms



Systematic study of DTM in circular torus has begun



- This should recover analytic DTM results at low pressure
- However, at higher pressure, there will be some competition with infernal modes
- Stay tuned for coming results.

That's All I have

Anything Else ?