Nonlinear MHD 2007 Q4 Milestones Report

J. Breslau, J. Chen, E. Held, S. Jardin[1], S. Kruger, A. Pankin, P. Snyder, C. Sovinec, H. Strauss, and the M3D, NIMROD, and CEMM Teams

November 19, 2007

1. Improve understanding of the present discrepancy between NIMROD and M3D and move to new CDX-U-relevant cases with more realistic parameters and sources.

Q4: Perform a new nonlinear comparison.

The status of this has now been submitted for publication in Computer Physics Communications. A preprint is available on the "Project Page" of the CEMM site. Look for "An Improved Tokamak Sawtooth Benchmark for 3D Nonlinear MHD (J. Breslau, C. Sovinec, S. Jardin, 2007).

The primary results from that paper are presented here in Figures 1-3. Figure 1 shows the time history of kinetic energy by toroidal mode number during the first three sawtooth crashes in the latest iteration of the nonlinear CDX benchmark.

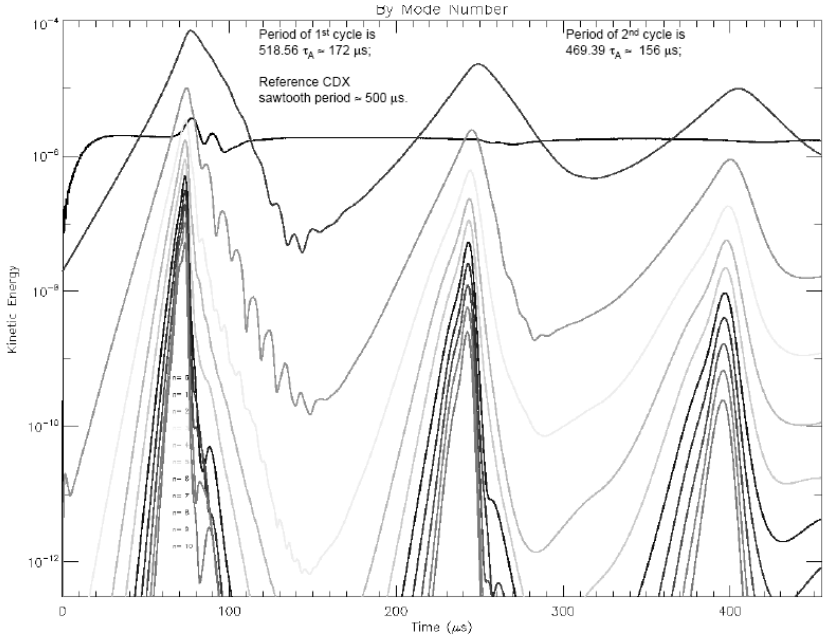

Figure 1a. M3D result for the kinetic energy (energy in normalized units, time in μs).

[1] Editor: jardin@pppl.gov

a.

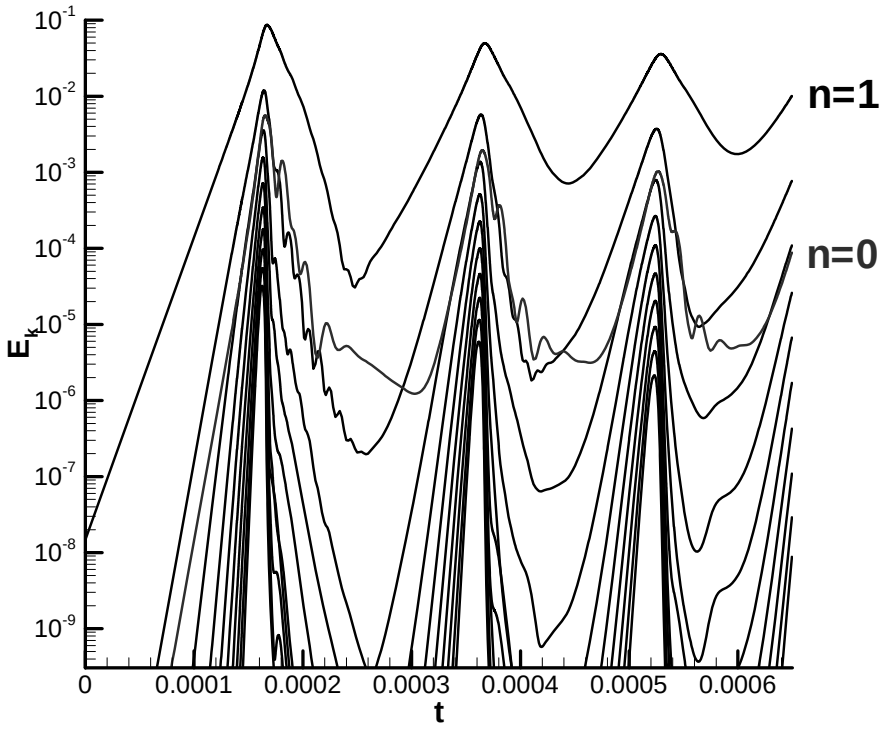

Figure 1b: NIMROD result for the kinetic energy (mks units).

It is clear from the figures that the codes are now in substantial, detailed agreement. Like NIMROD, M3D now conserves q_0 in the absence of mode activity, and thus shows a constant linear growth rate for the n=1 mode until just before the first crash. The codes also agree in the relative magnitudes of the various toroidal modes before, during, and after the crashes; on the detailed time behavior of the low-n modes; on the degree of damping of the oscillation in successive cycles, and on the cycle period of ~200 μs. It has also been confirmed that the M3D result is now converged toroidally.

When we investigate the actual plasma state at various corresponding times in the two runs, we also find detailed agreement (Figs. 2-3).

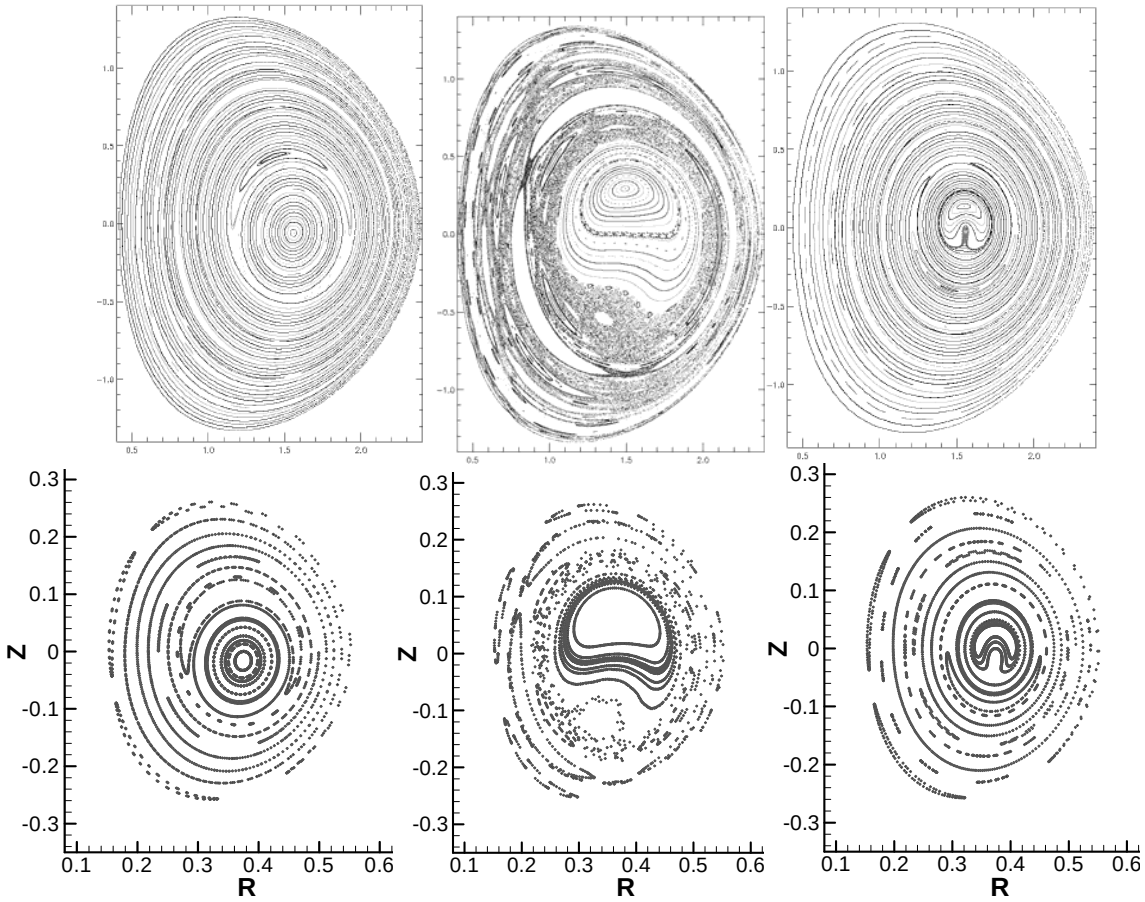

Figure 2: Poincaré sections showing magnetic flux surfaces at several time instants during the M3D (top) and NIMROD (bottom) sawtooth cycles. Left: late linear phase. Center: at culmination of crash. Right: Early recovery phase following crash.

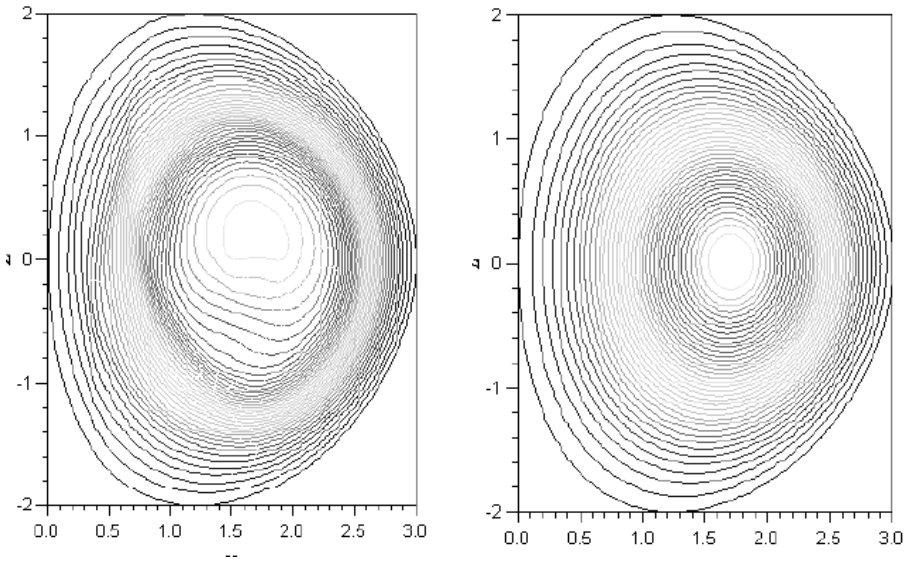

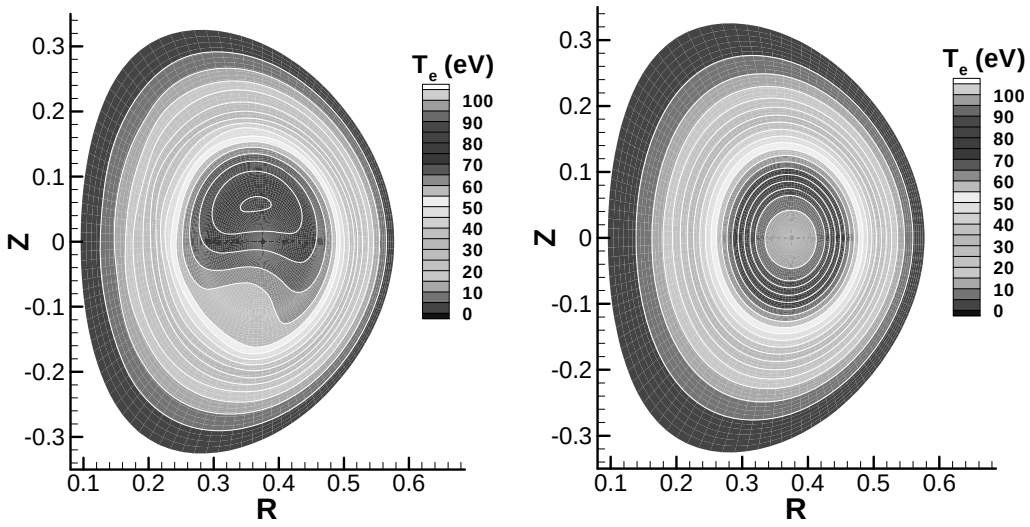

Figure 3: Temperature contours at two times during the M3D (top) and NIMROD (bottom) sawtooth cycles. Left: at culmination of crash. Right: Early recovery phase following crash.

This agreement constitutes a successful verification of the two codes. However the modeling is insufficient to produce quantitative agreement with the experimental results for the two predictions that can be compared directly with soft X-ray data from the experiment: the sawtooth period and the crash time . The predicted period of 200 μs is significantly less than the observed 500 μs sawtooth period in CDX, and the predicted crash time is a much larger fraction of the total cycle time than is observed in the device. Hence this study cannot be considered a successful validation of the model. A more refined model and a set of initial and boundary conditions that show greater fidelity to the experimental conditions are required.

2. Perform a linear edge stability calculation in a non-diverted equilibrium with a resistive code, and compare results with the linear ideal MHD code ELITE.

Q4: Repeat Q2 and Q3 using a two-fluid model.

Report:
New comparisons between ELITE and NIMROD are shown below in figure 4:

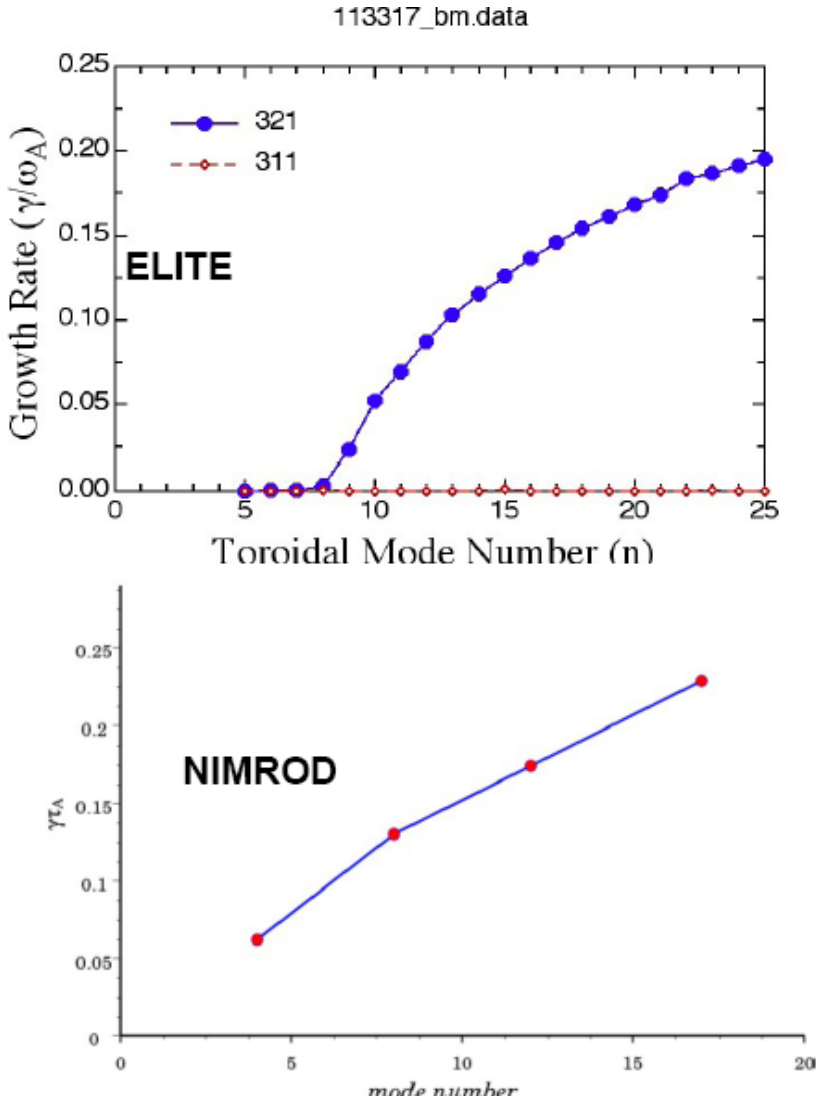

Figure 4: Comparison of ELITE and NIMROD for DIII-D discharge 113317.

In this new benchmark, performed in collaboration with A. Pankin from the CPES project, the TEQ code was used to alter experimentally based equilibrium computed with the EFIT code for the DIII-D discharge 113317 with well resolved pedestal area. The growth rates computed in ELITE and NIMROD codes agree reasonably well for large toroidal mode numbers. The ELITE code yields complete stabilization of modes with n<8, while the NIMROD code yields nonzero growth rates for these modes. Except for this difference, the growth rates agree remarkably well in this case.

3. Extend the 2D GEM nonlinear benchmark to non-zero guide field and more extreme parameters.

Q4: Document results and publish as warranted

We are trying to better understand the results presented in the Q3 report and put them in context before publication. The present status of our understanding is as follows:

Introduction:
Modern tokamaks normally operate in what we will call the "modern tokamak" regime:

$$\delta \ll \rho_S \ll d_i$$
$$\beta \ll 1 \sim \beta_P$$

Here, $\delta \sim \eta^{1/2}$ is the Sweet-Parker reconnection width, $\rho_S = c_s / \Omega_i = \sqrt{\beta} d_i$ is the ion Larmor radius, $c_S = (T/M_i)^{1/2}$ is the ion sound speed, $\Omega_{ci} = eB/M_i c$ is the ion cyclotron frequency $d_i = c/\omega_{Pi}$ is the ion skin depth, where ω_{Pi} is the ion plasma frequency, and β and β_P are the plasma pressure normalized to the toroidal and poloidal magnetic field pressures, respectively. In addition, the reconnecting (B_P) and background (B_T) magnetic fields normally satisfy the condition $\mu_P \equiv (m_e / m_i) B_T^2 / B_P^2 \ll 1$, so that kinetic Alfven Waves are not expected to contribute to fast reconnection[1]. Our study seeks to clarify under what conditions we can expect to see fast reconnection in the modern tokamak regime, if it can be related to the presence of Whistler waves alone, and the associated question of what role the ion Larmor radius plays.

Several numerical studies have appeared in the literature that show that fast reconnection can occur in the presence of a guide magnetic field for sufficiently large values of β_P. [1,2,7,8,12,14]. Other studies show that for β_P values near or less than unity, the guide field will tend to reduce the reconnection rate, but has a relatively small effect. [3,5,6,18]

In the present study, we include electron viscous dissipation but not electron inertia. This has been shown to be a good approximation for global reconnection studies in the presence of a guide field.[2,4,11,13,14,19]

Dependence of Reconnection Rate on Guide Field and plasma β

Here we present a series of runs performed with the M3D-C^1 code in the Spring of 2007 and presented at the Sherwood CEMM meeting. There was some concern because the M3D-C^1 code looked to have a large error in the energy balance. We now feel that the results are basically correct, but the energy conservation is very poor due to non-conservation of toroidal flux due to a boundary condition.(see below) Fixing this boundary condition leads to essentially the same results, except the energy is now converged

The geometry is basically the same as used in the GEM reconnection study.[see J. Birn, J. F. Drake, M. A. Shay, et al., J. Geophys. Res [Space Phys.] 106 (2001) 3715. We repeat the calculation with different values of guide field (B0) and record the maximum kinetic energy and the time at which it occurs.

Table I: Data from M3D-C^1 121 x 121 Guide Field Cases From Spring 07

Facden	mu	B0	run#	KE	time	Energy	Error
1.0	.05	0	5	6.97	32.4		
		0.2	30	6.41	32.8		
		1.0	35/37	1.73	42.		
		2.0	38/39	.081	62.		
	.005	0.2	44	7.84	32.0	-3.03	-3.2%
		1.0	41	2.50	38.1	1.01	2.9%
		2.0	42/43	.279	53.3	6.60	12%
		5.0	45/46	.250	13.5	7.98	16%
0.1	.05	1.0	36	3.72	10.9	1.56	2.3%
		2.0	48	1.59	13.8	6.73	13%
		5.0	49	0.30	20.	4.52	15%

The variable “facden” multiplies the initial density, so that a value of facden=1 is the original GEM specification where the initial density varies from 0.2 far away to 1.2 at the reconnection layer. By lowering facden to 0.1, the density varies only from 0.2 to 0.3 at the reconnection layer. The variable “mu” is the viscosity.

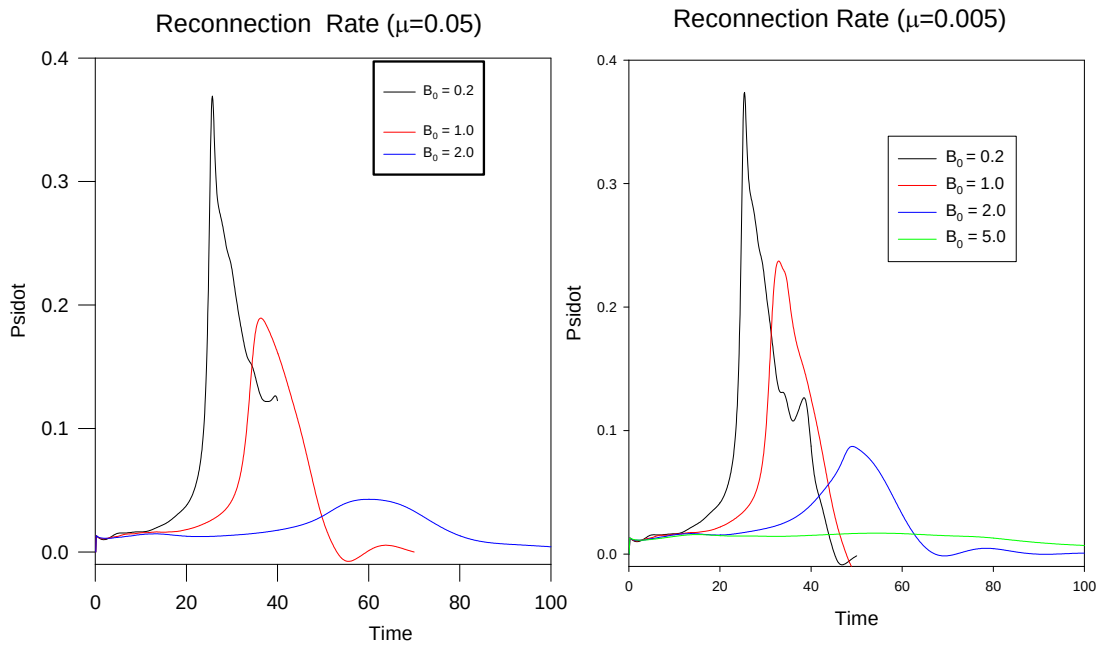

Figure 5: GEM reconnection rate for different guide fields and viscosities

It is clear that raising the toroidal field sharply curtails the maximum reconnection rate (which correlates with the kinetic energy (KE) and delays the time. This effect is partially offset by lowering the viscosity (Figure 5), and by lowering the initial density (Figure 6).

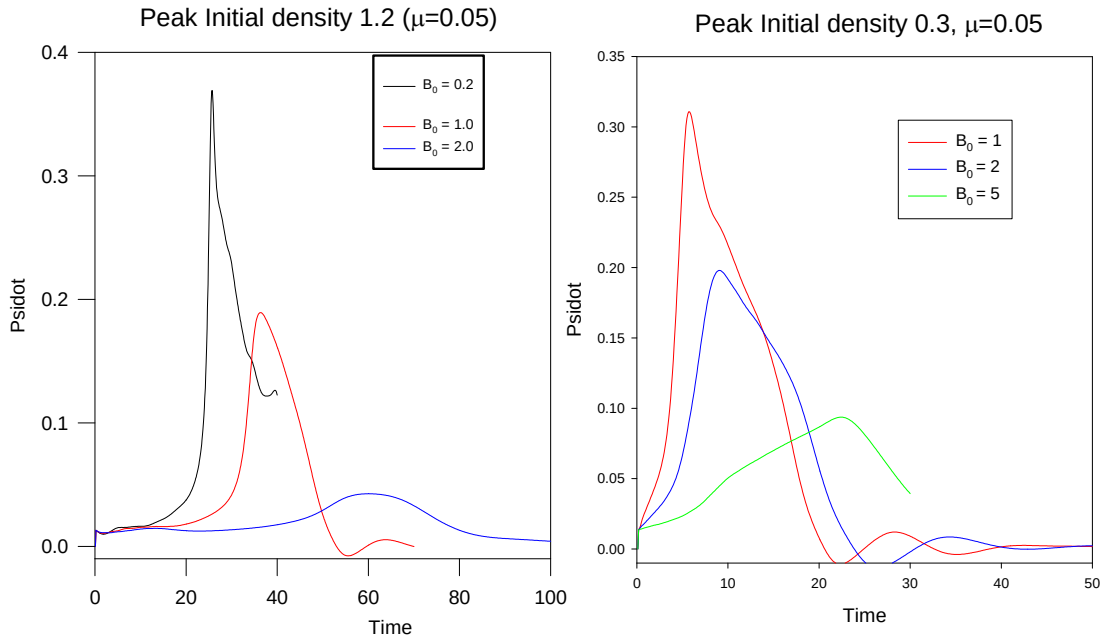

Figure 6: The effect of initial density on the GEM guide field reconnection study.

On 10/05/07 we implemented a new version of the M3D-C^1 code with $\theta(\theta\text{-}1) \rightarrow \theta^2$ in the implicit time advance. (Implemented with IADV=0, THIMP=0.55, TH2MP=1.0). This

leads to a more stable code, that converges much better when calculating steady states, and is also now stable for the largest grids used (211x211 elements). However, the results are basically the same as above.

On 11/02/07: The effect of toroidal field conservation on energy conservation was confirmed. We added a TF correction feature that relaxes the value of the edge toroidal field value, "gbound", towards the value that conserves toroidal flux, and the energy error decreases greatly while the KE stays about the same. We are experimenting with the coefficient, fc, used in that relaxation.

Table II: Effect of energy error on toroidal field correction coefficient

Run	fc	% error
gf002n	0.0	18.1
gf005n	1.0	4.6
gf006n	10.	1.7

Note that because of the normalizations used, our studies always assume the ion collision depth is unity, di = 1.

Discussion

There is some literature on the possible importance of diamagnetic effects [15,16,17]. We have experimented with increasing the initial pressure from 0.5 up to 5.0, but have not seen much effect on the reconnection rate. A plot showing the effect of gyroviscosity on reconnection rate is given in Figure 7.

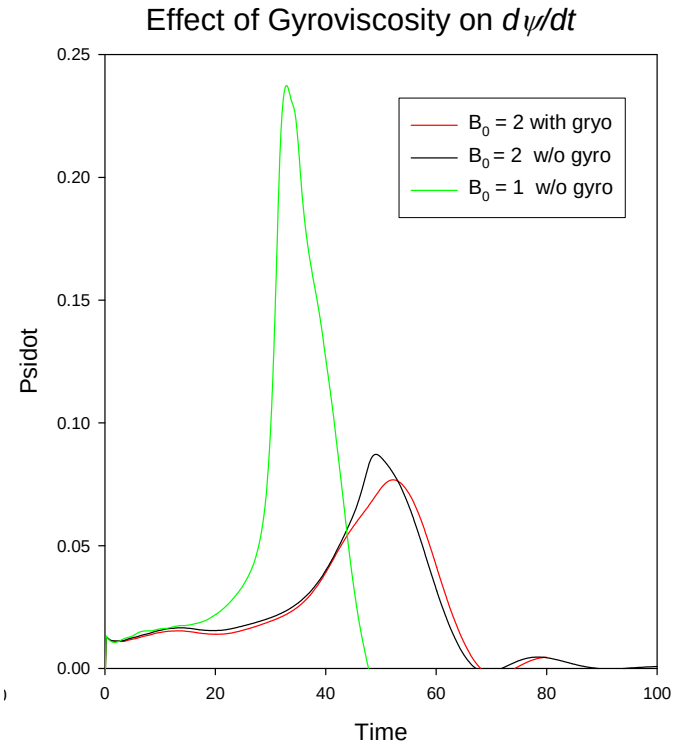

Figure 7: Effect of gyroviscocsity on guide field reconnection.

A possible explanation of what we are seeing in these studies is the effect of density depletion on the reconnection rate. At low values of the guide field, the density tends to deplete in the reconnection region, making the effect of the Hall term more pronounced. As we raise the value of the guide field, the flow becomes more incompressible and there is less density depletion. We are performing further studies to better understand the importance of this effect, and if it can totally explain the studies.

Bibliography

[1] B. N. Rogers, R. E. Denton, J.F. Drake, and M. A. Shay, "Role of Dispersive Waves in Collisionless Magnetic Reconnection", PRL 87 p 195004 (2001)
[2] P. A. Cassak, J. F. Drake, and M. A. Shay, "Catastrophic onset of fast magnetic reconnection with a guide field", Phys. Plasmas 14, 054502 (2007)
[3] J. D. Huba, "Hall magnetic reconnection: Guide field dependence", Phys. Plasmas 12, 012322 (2005)
[4] M. Hesse, M. Kuznetsova, K. Schindler, J. Birn, "Three-dimensional modeling of electron quasiviscous dissipation in guide-field magnetic reconnection", Phys. Plasma 12, 100704 (2005).
[5] M. Swisdak, J. Drake, and M. Shay, "Transition from antiparallel to component magnetic reconnection", J. Geophysical Research, 110, A05210, (2005)
[6] P. Ricci, J. Brackbill, W. Daughton, G. Lapenta, "Collisionless magnetic reconnection in the presence of a guide field", Phys. Plasma 11 (2004) 4102
[7] R. Fitzpatrick, "Scaling of forced magnetic reconnection in the Hall-MHD Taylor problem with arbitrary guide field", Phys. Plasmas 11 (2004) 3961
[8] X. Wang, A. Bhattachargee, and Z. Ma, "Collisionless reconnection: Effects of Hall current and electron pressure gradient", J. Geo. Res. 105 no. A12 (2000) 27,633-27,648
[9] D. Grasso, F. Califano, F. Pegoraro, and F. Porcelli, "Ion Larmor Radius Effects in Collisionless Reconnection", Plasma Physics Reports, 26, no 6 (2000) 512-518.
[10] D. Biskamp, E. Schwarz, and J. Drake,"Two-fluid theory of collisionless magnetic reconnection", Phys. Plasmas, 4, (1997) 1002
[11] B. Rogers and L. Zakharov, "Collisionless m=1 reconnection in tokamaks", Phys. Plasmas 3 (6) (1996) 2411
[12] R. Kleva, J. Drake, and F. Waelbroeck, "Fast reconnection in high temperature plasmas", Phys. Plasmas 2 (1995) 23
[13] M. Ottaviani and F. Porcelli, "Fast nonlinear magnetic reconnection", Phys. Plasmas 2 (11) (1995) 4104
[14]A. Aydemir, "Nonlinear studies of m=1 modes in high-temperature plasmas", Phys. Fluids B 4 (1992) 3469
[15] B. Rogers and L. Zakharov, Phys. Plasmas 2, 3420 (1995)
[16] M. Swisdak, J. F. Drake, M. A. Shay, and B. N. Rogers, J. Geophys. Res. 108, 1218 (2003)
[17] K. Germaschewski, A. Bhattachargee, to be published.
[18] B. Sullivan, B. Rogers, and M. A. Shay, "The scaling of forced collisionless reconnection", Phys. Plasmas 12, 122312 (2005)
[19] M. A. Shay and J. F. Drake, "The role of electron dissipation on the rate of collisionless magnetic reconnection", Geophys. Research Lett, 25, p. 3759-3762 (1998)

4. Scalability studies on leading edge computers.

Q4: Document results and present at conference as warranted.

Report:

Improvements in the use of SuperLU by NIMROD:
A difficulty in scaling problem size when using parallel sparse direct solvers, such as SuperLU, is the growth of required memory. In this past quarter, we revised the parallel communication used within NIMROD to complete matrix rows before interfacing SuperLU. The routines had previously used MPI allreduce operations, which are easy to program for this purpose but use a substantial amount of buffer space. The revised routines use point-to-point communication to complete distinct sets of matrix rows. This allows us the option of calling either of two different interfaces to SuperLU. The first interface has each processor holding a distinct set of matrix rows. The second interface has each processor holding a full copy of the matrix, and the NIMROD routine now completes the full copies with MPI allgather operations after the new point-to-point communication, using less memory than allreduce operations. Results of a weak scaling performed with a single Fourier component are shown in the figure below. The problem size scales from one to sixty-four blocks of 100 high-order elements each. The same scaling with the previous implementation was not able to run the 64-processor case due to memory limitation. The new communication is also being used to interface to the PETSc library, which allows many solver choices. So far, we have successfully used the SuperLU and MUMPS libraries through the PETSc interface. For the next quarter, we will add Fourier-layer parallelism to our weak-scaling study with the expectation of running tests to thousands of processors for the fluid algorithm.

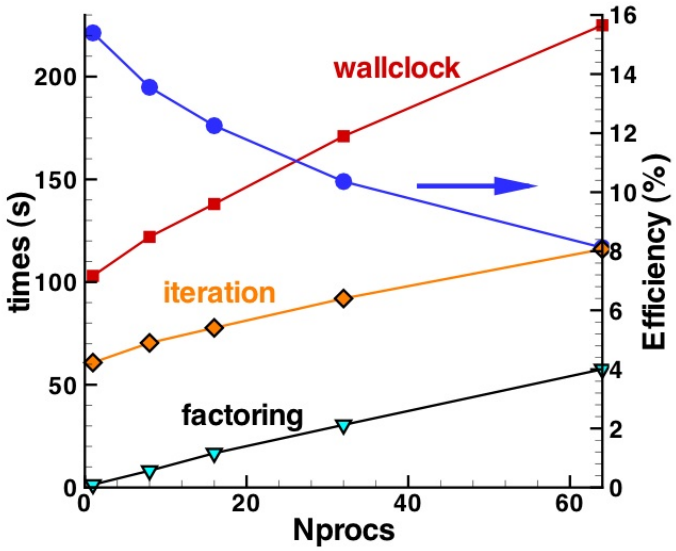

Figure 8: Results of a weak-scaling study performed for a two-dimensional NIMROD (one Fourier component) fluid computation. Revised communication is used with the full-storage interface to SuperLU. Results with the distributed-memory interface are virtually identical.

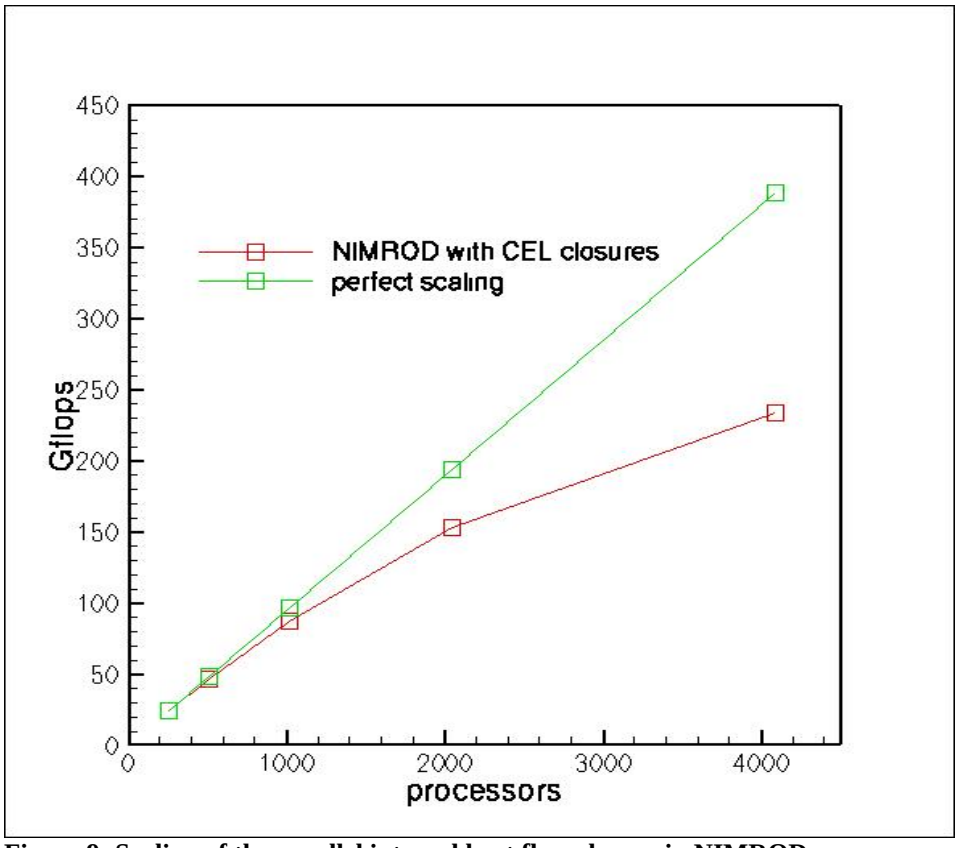

Figure 9: Scaling of the parallel integral heat flow closure in NIMROD

As part of NERSC's ERCAP process, further scaling runs were performed using NIMROD in conjunction with the parallel, integral heat flow closure. In the interest of understanding electron heat confinement in the presence of stochastic magnetic fields for SSPX, temperature was evolved to steady state keeping the magnetic field frozen. The core temperature results from balancing parallel and perpendicular transport against resistive heating. The parallel heat flow closure calculation entails solving the steady state electron drift kinetic equation approximately 1 million times throughout the computational domain before the temperature is advanced. The plot below shows good scaling up to 2048 processors with an efficiency of about 75%. Up to 1024 processors, efficiency is nearly perfect at 90%. One possible reason for the fall off beyond 1000 processors is the fact that processors added beyond the 1024 mark were solely used for the closure calculation. Although this part of the problem scales perfectly, the total time (summed over processors) spent waiting for the updated temperature from the fluid processors increases. Improving the scaling beyond 2000 processors most likely requires increasing the fluid processors by the same factor as the closure processors and being in regime where the fluid part scales perfectly. Future improvements of scaling with the integral parallel closures are planned for the next quarter including performing scaling studies on Franklin and possibly Jaguar.

3D Strong Scaling with M3D:
We have performed 3D strong scaling studies with M3D. In these studies, the problem size was held fixed as we increased the number of processors from 624 to 4992. As can be seen from the figure, the solvers are scaling even better than ideal, but the data copy is bringing the entire scaling down. Since these studies were performed, Jaquar has been taken down for an upgrade and the new Franklin computer at NERSC has begun commissioning. We expect to resume these scaling studies on Jaguar and Franklin in FY08 as the machines allow.

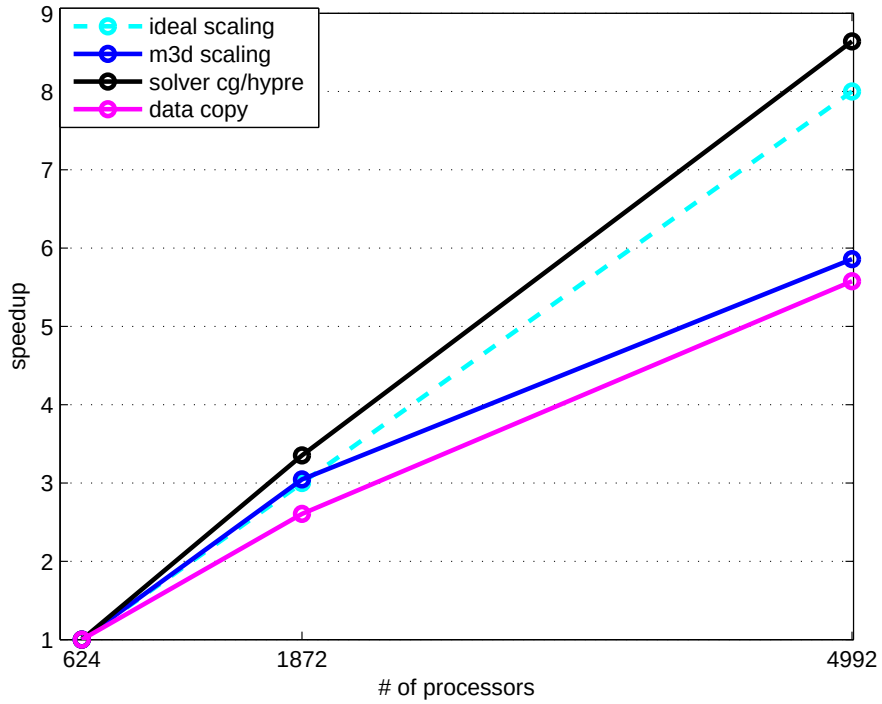

Figure 10: M3D strong 3D scaling on Jaguar