Iterative and direct linear solvers in fully implicit magnetic reconnection simulations with inexact Newton methods

Xuefei (Rebecca) Yuan[1], Xiaoye S. Li[1], Ichitaro Yamazaki[1], Stephen C. Jardin[2], Alice E. Koniges[1] and David E. Keyes[3,4]

[1]LBNL (USA), [2]PPPL (USA), [3]KAUST (Saudi Arabia), [4]Columbia University (USA)

The work was supported by the Petascale Initiative in Computational Science at National Energy Research Scientific Computing Center (NERSC) . Additionally, we gratefully acknowledge the support of NERSC for export advice and time on the new Cray XE6 system (Hopper). This research was supported in part of the Director, Office of Science, Office of Advanced Scientific Computing Research, of the U.S. DOE under Contract No. DE-AC02-05CH11231.

MATHEMATICAL MODEL: FOUR-FIELD EXTENDED MHD EQUATIONS

The reduced two-fluid MHD equations in two-dimensions in the limit of zero electron mass can be written as:

$$
\begin{cases}
\frac{\partial}{\partial t}\nabla^2\varphi + \vec{V}\cdot\nabla(\nabla^2\varphi) = \left[\nabla^2\psi,\psi\right] + \mu\nabla^4\varphi \\
\frac{\partial}{\partial t}V + \vec{V}\cdot\nabla V = [B,\psi] + \mu\nabla^2 V - \mu h\nabla^4 V \\
\frac{\partial}{\partial t}\psi + \vec{V}\cdot\nabla\psi = d_i[\psi,B] + \eta\nabla^2\psi - \nu\nabla^4\psi \\
\frac{\partial}{\partial t}B + \vec{V}\cdot\nabla B = [V,\psi] + d_i\left[\nabla^2\psi,\psi\right] + \eta\nabla^2 B - \nu\nabla^4 B
\end{cases}
$$

- the ion velocity: $\vec{V} = \nabla\varphi\times\hat{z} + V\hat{z}$
- the magnetic field: $\vec{B} = \nabla\psi\times\hat{z} + B\hat{z}$
- the out-of-plane current density: $j = -\nabla^2\psi$
- the Poisson bracket: $[f,g] \equiv \nabla f\times\nabla g\cdot\hat{z}$
- the electrical resistivity: η
- the collisionless ion skin depth: d_i
- the fluid viscosity: μ
- the hyper-resistivity (or electron viscosity): ν
- the hyper-viscosity: h

- the computational domain $\Omega = [0, 0.5L_x]\times[0, 0.5L_y]$: the first quadrant of the physical domain (finite diff., (anti-)symmetric fields);
- boundary conditions: Dirichlet at the top, anti-symmetric in φ, B and symmetric in ψ, V at other three boundaries;
- initial conditions: a Harris equilibrium and perturbation combination for ψ, and other three fields are zeros

$$\psi(\xi,\eta,0) = \frac{1}{2}\ln\cosh 2\eta + \varepsilon\cos k_x\cos k_y,$$

$$k_x = \frac{2\pi}{L_x},\quad k_y = \frac{2\pi}{L_y},\quad \varepsilon = 0.1,\quad L_x = 25.6,\quad L_y = 12.8.$$

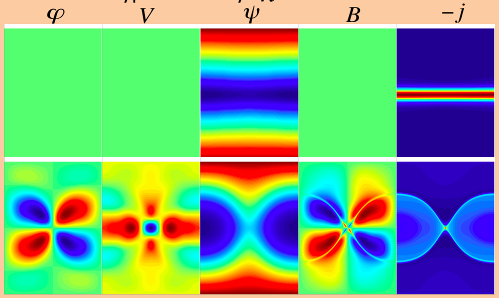

Four fields and the negative out-of-plane current: top (t=0), bottom (t=40): $\eta = 0.005$, $\mu = 0.05$, $d_i = 1.0$.

NUMERICAL DIFFICULTY: LARGER VALUE OF SKIN DEPTH

The MHD system applied to strongly magnetized plasma is inherently ill-conditioned because there are several different wave types with greatly differing wave speeds and polarizaiton. This is especially troublesome when the collisionless ion skin depth is large so that the Whistler waves, which cause the fast reconnection, dominate the physics.

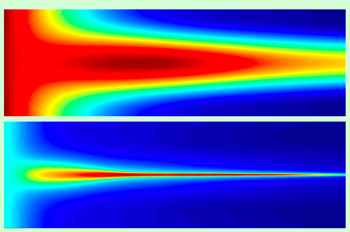

The mid-plane current density vs. time: vertical axis is $-j$ along the mid-plane $y = 0$, and the horizontal axis is time $t = 0\sim 40$. Top: $d_i = 0.0$, bottom: $d_i = 1.0$.

grid size: 256X256, dt=0.5, nt=80.

$linear\ solver \backslash d_i$	0.0	0.2	0.4	0.6	0.8	1.0
bj_ilu	O	O	✗	✗	✗	✗
bj_lu	✓	✓	✓	✓	✓	✓
asm_ilu	✓	✓	✓	✗	✗	✗
asm_lu	✓	✓	✓	✓	✓	✓
SuperLU	✓	✓	✓	✓	✓	✓

grid size: 512X512, dt=0.2, nt=200.

$linear\ solver \backslash d_i$	0.0	0.2	0.4	0.6	0.8	1.0
bj_ilu	✗	✗	✗	✗	✗	✗
bj_lu	✗	✗	✗	✗	✗	✗
asm_ilu	✓	✓	✓	✗	✗	✗
asm_lu	✓	✓	✓	✓	✓	✓
SuperLU	✓	✓	✓	✓	✓	✓

Five different linear solvers are tested for different skin depth from 0.0 to 1.0 in 256X256 and 512X512 problem size: iterative GMRES solvers (bj_ilu, bj_lu, asm_ilu, asm_lu) and direct solver (SuperLU). As the skin depth increases, iterative solvers need a good preconditioner to converge, while the direct solver converges for all cases. The block Jacobi (bj) has not applied the freedom to vary the block size, which would enhance the linear convergence for the higher skin depth case. The additive Schwarz methods (asm) has overlap numbers n>1.

NUMERICAL EXPERIMENTS: SCALABILITY STUDIES

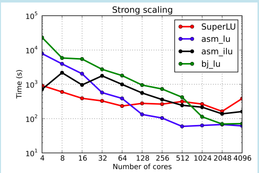

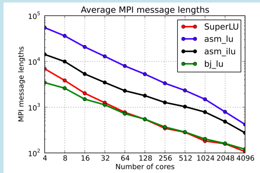

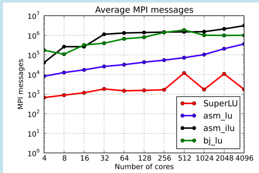

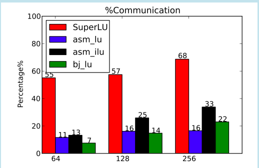

- Three iterative solvers (bj_lu, asm_ilu, asm_lu) and the direct solver (SuperLU) for a 256X256 size problem for di=0.2, dt=0.5, nt=10;
- SuperLU and bj_lu have lower MPI message lengths;
- the communication percentage of SuperLU is over half of the wall time and increases as the number of cores increases;
- IPM and PETSc profiling tools are used;
- the SuperLU uses sequential ordering algorithm and symbolic factorization, and this time doesn't decrease with increasing # of cores.

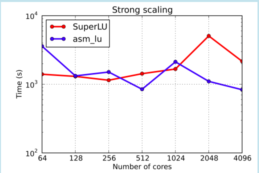

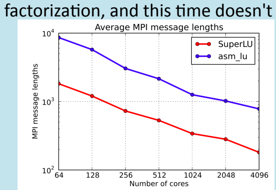

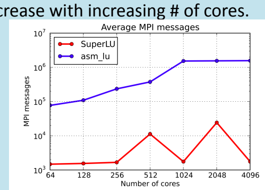

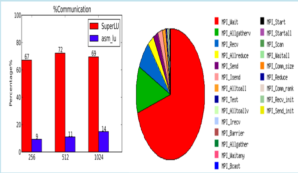

- For a very challenge case where the skin depth number di=1.0, the problem size is 512X512, only asm_lu and SuperLU provide converged solutions;
- the wall time does not decrease when number of cores increases;
- the 70% MPI time is MPI_Wait for SuperLU;
- the communication percentage of SuperLU is over half of the wall time and increases as the number of cores increases;

$lin.\ solver \backslash dt$	0.2	0.4	0.6	0.8	1.0
asm_ilu	4.6|245.6	4.8|388.7	4.9|497.6	4.9|615.7	4.9|676.9
asm_lu	4.6|245.2	4.7|372.5	4.9|485.2	4.9|559.8	4.9|628.9
SuperLU	2.5|2.5	3.1|3.1	2.9|2.9	3.1|3.1	2.9|2.9

The average nonlinear and linear iteration numbers for 512X512 grid size problem per time step, where di=0.2, and $t = 0\sim 40$.

- The Newton iteration numbers do not increases as dt increases;
- the linear iteration numbers for iterative solvers increase as dt increases;
- the nonlinear/linear iteration numbers do not increase as dt increases for SuperLu.

SIMULATION ARCHITECTURE: NERSC CRAY XE6 "HOPPER"

- 6384 nodes, 24 cores per node (153,216 total cores)
- 2 twelve-core AMD 'MagnyCours' 2.1 GHz processors per node (NUMA)
- 32 GB DDR3 1333 MHz memory per node (6000 nodes), 64 GB DDR3 1333 MHz memory per node (384 nodes)
- 1.28 Petaflop/s for the entire machine
- 6 MB L3 cache shared between 6 cores on the 'MagnyCours' processor
- 4 DDR3 1333 MHz memory channels per twelve-core 'MagnyCours' processor

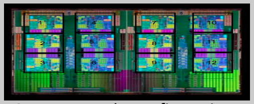

Compute node configuration

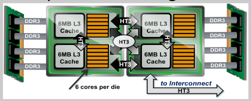

MagnyCours processor