Extreme-Scale Preconditioners for Fully-Coupled Magnetohydrodynamics

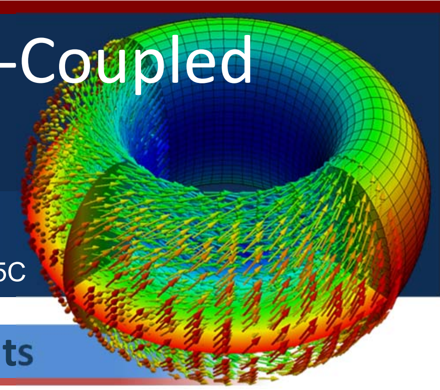

Sandia National Laboratories
Eric C. Cyr, Paul T. Lin, John N. Shadid
Sandia National Laboratories, New Mexico 87123

SAND2013-6415C

Problem Statement

The magnetohydrodynamics (MHD) model describes the dynamics of charged fluids in the presence of electromagnetic fields. MHD models describe important phenomena in scientific applications of critical importance to DOE. The coupled nonlinear resistive MHD PDE model is:

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u}\cdot\nabla\mathbf{u} - \nabla\cdot\nu\nabla u + \nabla p + \nabla\cdot\left(-\frac{1}{\mu_0}\mathbf{B}\otimes\mathbf{B} + \frac{1}{2\mu_0}\|\mathbf{B}\|^2 I\right) = \mathbf{0}$$

$$\nabla\cdot\mathbf{u} = 0$$

$$\frac{\partial \mathbf{B}}{\partial t} + \nabla\times(\mathbf{u}\times\mathbf{B}) - \frac{\eta}{\mu_0}\nabla\times\nabla\times\mathbf{B} = \mathbf{0}$$

$$\nabla\cdot\mathbf{B} = 0$$

Scientific understanding through computational simulations relies on accurate and robust long-time simulations. Implicit strongly-coupled MHD methods can provide the required accuracy and robustness; our approach is to develop implicit fully-coupled Newton-Krylov solvers[1]. Efficient preconditioners[1,2] are key to good scalability and performance at extreme scales.

Using a stabilized 3D Lagrange multiplier B-field MHD formulation, we present:

- Two new approximate block factorization preconditioners that address the character of the coupling inherent in MHD
- A fully-coupled algebraic multigrid preconditioner is used to characterize solver performance at the extreme scale

Preconditioning Approach

Fully Coupled Algebraic Multigrid (FC AMG)

(Tuminaro, Hu, Siefert, Gee, Sala)

- All 8 PDEs in a single node of the graph
- Nonsmoothed aggregation; uncoupled aggregation
- Rebalance coarser levels
- ILU Smoothers
- FC AMG previously applied to 2D MHD formulation[1]

Block Preconditioners:

- Segregate linear system into physics components

$$\mathcal{J} = \begin{bmatrix} F & B_p^T & Z & \\ B_p & C_u & & \\ Y & & D & B_r^T \\ & & B_r & C_B \end{bmatrix} \Rightarrow \begin{bmatrix} F & B_p^T & \hat{Z} \\ B_p & C_u & \\ \hat{Y} & & \hat{D} \end{bmatrix}$$

- Good block preconditioners attack essential coupling
- Two preconditioners: extension of previous work on 2D MHD[2]

(1) $$\mathcal{J} \approx \mathcal{M}_{GS} = \begin{bmatrix} F & B^T & \hat{Z} \\ B & C & \\ & & \hat{D} \end{bmatrix}$$

Fluid Schur Complement: $\hat{S} = -Q_p F_p^{-1} A_p \ (\approx C - BF^{-1}B^T)$

(2) $$\mathcal{J} \approx \mathcal{M}_{Split} = \begin{bmatrix} F & & Z \\ & I & \\ Y & & \hat{D} \end{bmatrix} \begin{bmatrix} F^{-1} & & \\ & I & \\ & & I \end{bmatrix} \begin{bmatrix} F & B^T & \\ B & C & \\ & & I \end{bmatrix}$$

Velocity-Pressure: $\hat{S} = -Q_p F_p^{-1} A_p \ (\approx C - BF^{-1}B^T)$

Magnetics-Velocity: $\hat{P} = D - Y\,\mathrm{diag}(F)^{-1} Z \ (\approx D - YF^{-1}Z)$

Preliminary Results

Two Challenging Test Problems

1) Steady MHD Generator:
- 3D flow with external cross-stream B field

2) Transient Hydromagnetic Kelvin-Helmholtz
- 3D flow initialized with a U_x shear layer and a Harris-sheet B_x profile

Simple meshes, challenging coupling
⇒ good test problems
(simplistic internal mesh load balancing)

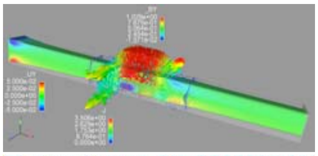

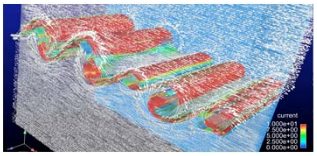

3D Steady MHD Generator

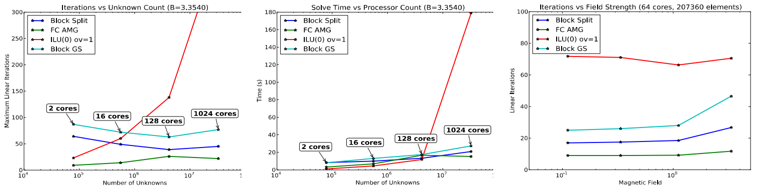

- Weak scaling of FC AMG and block preconditioners reasonable to 1024 cores
- Both suffer some performance degradation on this capacity machine

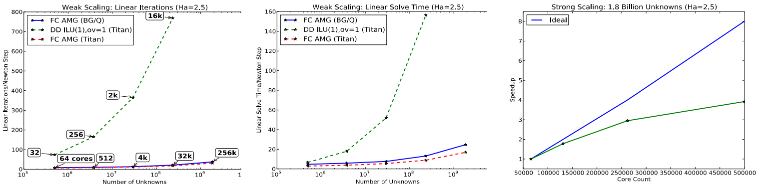

Weak scaling to 1.8 billion unknowns: Blue Gene/Q (BG/Q) and Cray XK7 (Titan)
- Algorithmic and CPU time scaling: encouraging

Strong scaling: 1.8 billion unknowns on 500,000 cores (BG/Q)
- Scaling tails off at ½ million cores (pursuing significant improvements)

Hydromagnetic Kelvin-Helmholtz

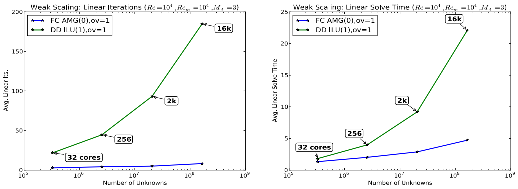

Weak scaling to 165 million unknowns Cray XK7 (Titan)
- Initial algorithmic and CPU time scaling: reasonably encouraging, need work on fine grid smoother efficiency and load balance.

Conclusions

- Demonstrated results on 2 leadership-class architectures
- Demonstrated reasonable scaling to very-large core counts (128K, 256K)
- Demonstrated ability to run near extreme-scale (½ million cores)
- Identified issues to improve scaling
  - Improve efficiency and scaling of ML preconditioner setup
  - Optimize load balance for linear algebra / ML kernels

References:
1. J. N. Shadid, R. P. Pawlowski, J. W. Banks, L. Chacón, P.T. Lin, R. S. Tuminaro, **Towards a scalable fully-implicit fully-coupled resistive MHD formulation with stabilized FE methods**, JCP, 2010.
2. E. C. Cyr, J. N. Shadid, R. S. Tuminaro, R. P. Pawlowski, L. Chacon, **A New Approximate Block Factorization Preconditioner for 2D Incompressible (Reduced) Resistive MHD**, SISC, 2013

Sandia National Laboratories

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Office of Science