National Laboratories
SAND2021-2441C

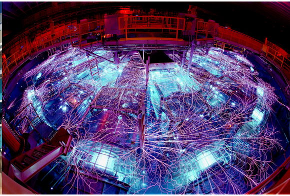

An Algebraic Monolithic Multigrid Method

for Volume-coupled Multiphysics Problems

Peter Ohm, T. Wiesner, R.S. Tuminaro, J.J. Hu, J. Shadid, E.C. Cyr

Center for Computing Research, Sandia National Laboratories

March 1, 2021

Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy as National Nuclear Security Administration under contract DE-NA0003525.
SAND NO. 2021-00000

Motivation

Sandia National Laboratories

What is the problem?

- **Multiphysics problems:** *application-specific* with *increasing complexity* through new types of physics and/or discretizations
- No resource-efficient *general* black-box solver available

What can we do about this?

Flexible software framework for iterative solvers/preconditioners

- Modular design for application-specific solver layouts
- Usability through simplified user interfaces for non-experts

Representation of fully-coupled multiphysics problems

Blocked linear operator

- Single-field problems represented by diagonal blocks

Example: 3×3 blocked operator

$A =$

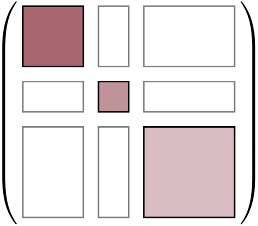

Representation of fully-coupled multiphysics problems

Blocked linear operator

- Single-field problems represented by diagonal blocks
- Coupling represented by off-diagonal blocks

Example: 3×3 blocked operator

$A =$ 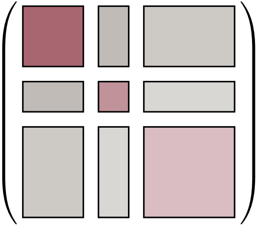

Representation of fully-coupled multiphysics problems

Blocked linear operator

- Single-field problems represented by diagonal blocks
- Coupling represented by off-diagonal blocks
- Nested blocked operators for hierarchical dependencies

Example: 2×2 blocked operator with nested 2×2 blocked operator

$A =$

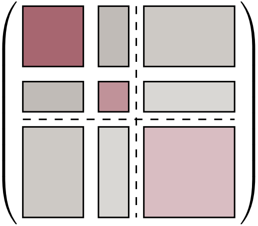

Representation of fully-coupled multiphysics problems

Blocked linear operator

- Single-field problems represented by diagonal blocks
- Coupling represented by off-diagonal blocks
- Nested blocked operators for hierarchical dependencies

Example: 2×2 blocked operator with nested 2×2 blocked operator

$A =$

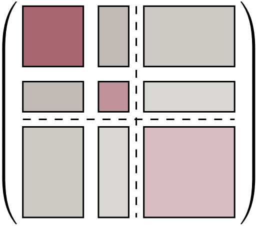

How to design efficient multigrid preconditioners for multiphysics problems?

Idea of Algebraic Multigrid Methods

Multigrid method

Transfer operators + Level smoothers

- Generate coarse representations $A^{(i)}$ of fine level problem $A^{(0)}$

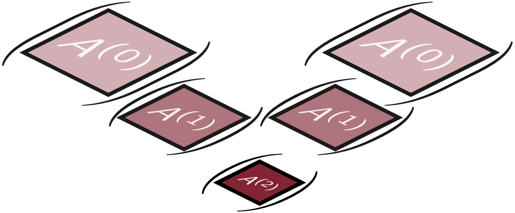

Idea of Algebraic Multigrid Methods

Multigrid method

Transfer operators + Level smoothers

- Generate coarse representations $A^{(i)}$ of fine level problem $A^{(0)}$ using $A^{(i+1)} = R_{i\to(i+1)}A^{(i)}P_{(i+1)\to i}$
- Rectangular transfer operators $P_{(i+1)\to i}$ and $R_{i\to(i+1)}$

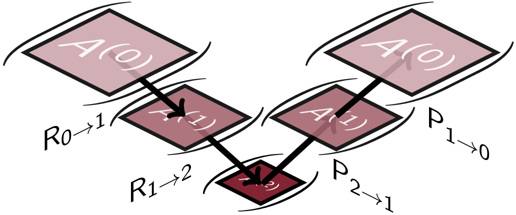

Idea of Algebraic Multigrid Methods

Multigrid method

Transfer operators + **Level smoothers**

- Generate coarse representations $A^{(i)}$ of fine level problem $A^{(0)}$ using $A^{(i+1)} = R_{i\to(i+1)}A^{(i)}P_{(i+1)\to i}$
- Rectangular transfer operators $P_{(i+1)\to i}$ and $R_{i\to(i+1)}$
- Level smoothers $\mathcal{S}^{(i)}$ damp high-oscillatory error modes

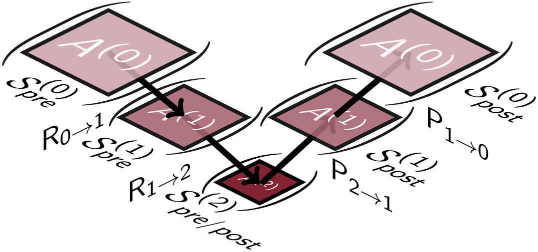

Multilevel approach for fully-coupled multiphysics problems

How to adapt AMG methods to block systems?

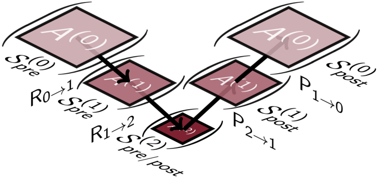

$\overset{?}{\Longleftrightarrow}$

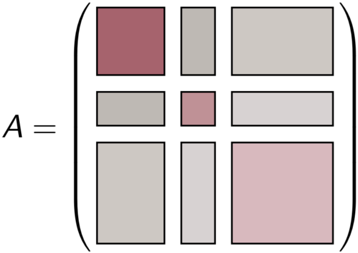

Multilevel approach for fully-coupled multiphysics problems

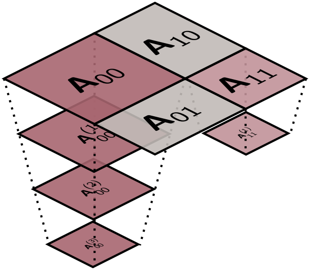

Physics-based approach

- Block factorization, approximate Schur complement
- Not-so-obvious approximate factorizations

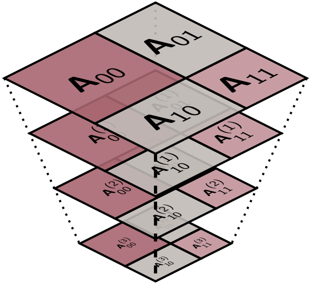

Monolithic AMG approach

- Multilevel representation of inter-field interactions
- Not-so-obvious AMG smoothes & grid transfers

MueLu – The Trilinos Multigrid framework

MueLu multigrid framework:

- Extensible software layout
  - Modularity: Preconditioner layout defined by small building blocks
  - Logic: Building blocks connected through logical data dependencies
- Flexible user-friendly XML-based user input system
- Designed for next-generation HPC systems

Example: Building blocks for transfer operators and level smoother

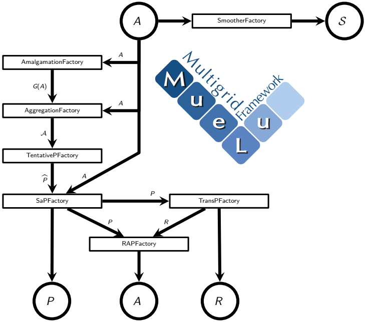

www.trilinos.org/packages/muelu

Multigrid preconditioner layout for multiphysics

Multigrid for multiphysics

Transfer operators + Level smoothers + Coupling

- Segregated transfer operators P and R to keep algebraic blocks separate on coarse levels

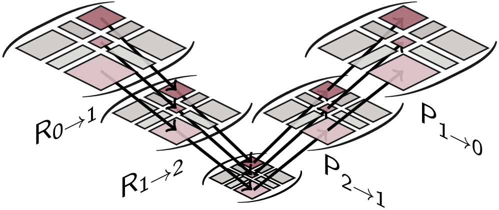

Multigrid preconditioner layout for multiphysics

Multigrid for multiphysics

Transfer operators + **Level smoothers + Coupling**

- Segregated transfer operators P and R to keep algebraic blocks separate on coarse levels
- Nested block smoothers consider coupling of different fields

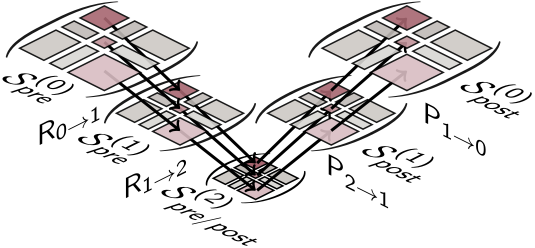

Segregated transfer operators

$A_i =$ 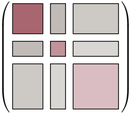

Transition from level i to $i+1$:

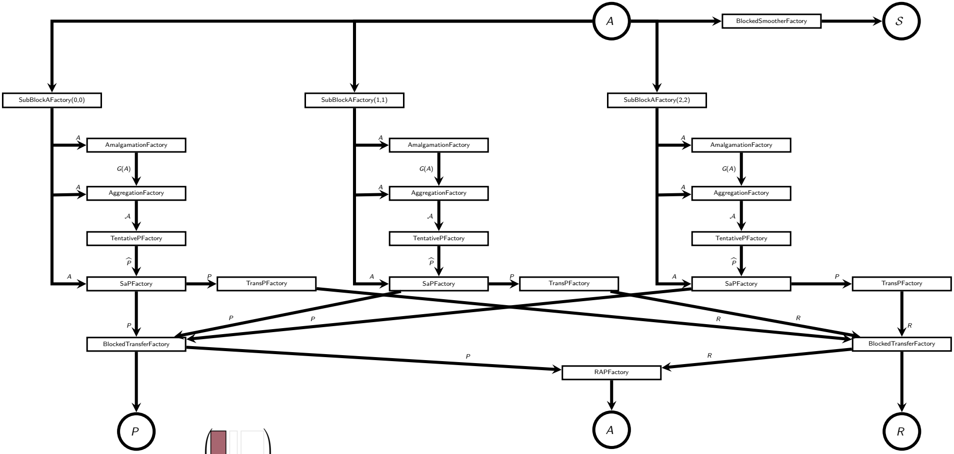

$P_{(i+1)\to i} =$ 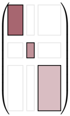

$A_{i+1} =$ 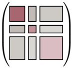

$R_{i\to(i+1)} =$ 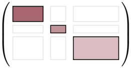

Block relaxation

Pool of block relaxation methods

- General $n \times n$ block systems: Blocked Gauss-Seidel smoother
- General 2×2 block systems: SIMPLE, Uzawa, Braess-Sarazin
- Physics-based block smoothers from the Teko package

Build your application-specific block relaxation

- Consider the coupling blocks when designing the block relaxation

$A =$

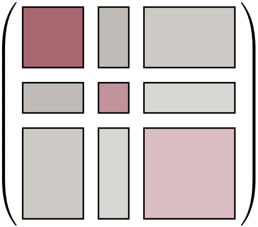

Block relaxation

Pool of block relaxation methods

- General $n \times n$ block systems: Blocked Gauss-Seidel smoother
- General 2×2 block systems: SIMPLE, Uzawa, Braess-Sarazin
- Physics-based block smoothers from the Teko package

Build your application-specific block relaxation

- Consider the coupling blocks when designing the block relaxation
- Use nested block relaxation:
  1 BGS (0.5)
  - 1 SIMPLE (0.8)
    - ILU(0), ov=1
    - ILU(0), ov=1
  - ILU(0), ov=1

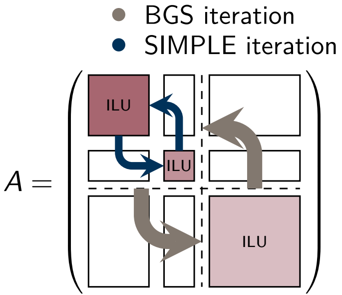

Resistive Magnetohydrodynamics

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

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

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

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

with appropriate initial and boundary conditions.

Discretization

Stabilized discretization of MHD equations using equal order piecewise bilinear elements on hexahedrons. Stabilizes [**u**,P], [**B**,r] saddle point problems, and adds a streamline upwind Petrov-Galerkin (SUPG) stabilization of convective transport.

Reference solver

Discretization

Stabilized discretization of MHD equations using equal order piecewise bilinear elements on hexahedons

⇒ Collocated solution unknowns ($\boldsymbol{u}_x, \boldsymbol{u}_y, \boldsymbol{u}_z, p, \boldsymbol{B}_x, \boldsymbol{B}_y, \boldsymbol{B}_z, r$) on each mesh node

Reference solver

- Preconditioned GMRES (from Belos or AztecOO package)
- Fully-coupled MueLu multigrid preconditioner
  - 8 DOFs per node
  - Level smoother: Additive Schwarz (overlap=1) with ILU(0)
  - Non-smoothed transfer operators

P.T. Lin, J.N. Shadid, R.S. Tuminaro, M. Sala, G.L. Hennigan, R.P. Pawlowski; *A parallel fully coupled algebraic multilevel preconditioner applied to multiphysics PDE applications: Drift-diffusion, flow/transport/reaction, resistive MHD*; Int. J. Numer. Meth. Fluids, 64,1148-1179; 2010

J.N. Shadid, R.P. Pawlowski, E.C. Cyr, R.S. Tuminaro, L. Chacon, P.D. Weber; *Scalable implicit incompressible resistive MHD with stabilized FE and fully-coupled Newton-Krylov-AMG*; Comput. Methods Appl. Mech. Engrg., 304, 1-25; 2016

Block multigrid preconditioner

Incompressible resistive MHD equations:

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

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

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

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

Block structure of linear systems after discretization:

$$\begin{pmatrix} \begin{matrix} J_u & B^T \\ D & L_p \end{matrix} & \begin{matrix} Z & 0 \\ 0 & 0 \end{matrix} \\ \begin{matrix} Y & 0 \\ 0 & 0 \end{matrix} & \begin{matrix} J_I & B^T \\ B & L_\phi \end{matrix} \end{pmatrix}$$

- 2×2 block system with 4 DOFs per node each block
- Solution variables: $(\boldsymbol{u}_x, \boldsymbol{u}_y, \boldsymbol{u}_z, p)$ and $(\boldsymbol{B}_x, \boldsymbol{B}_y, \boldsymbol{B}_z, r)$

Block multigrid preconditioner for MHD

Design principles:

- Preserve coincidence of MHD unknowns on coarse levels
- Reduce memory footprint by avoiding global ILU smoothers
- Performance through ILU as single-field smoothers

Solver layout:

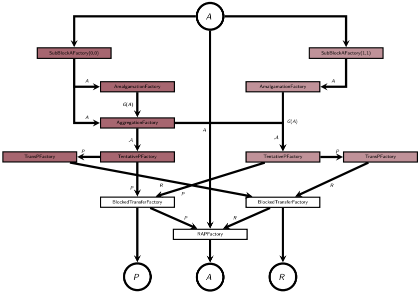

- Reuse aggregates $\mathcal{A}$ from Navier-Stokes part for magnetics part
- Non-smoothed transfer ops.
- Block smoother:
  n BGS(ω)
  - ILU(0), ov=1
  - ILU(0), ov=1

Block multigrid preconditioner for MHD

Design principles:

- Preserve coincidence of MHD unknowns on coarse levels
- Reduce memory footprint by avoiding global ILU smoothers
- Performance through ILU as single-field smoothers

Solver layout:

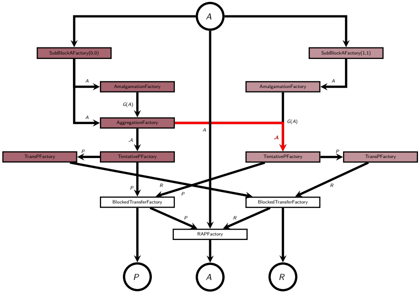

- Reuse aggregates $\mathcal{A}$ from Navier-Stokes part for magnetics part
- Non-smoothed transfer ops.
- Block smoother: n BGS(ω)
  - ILU(0), ov=1
  - ILU(0), ov=1

Block multigrid preconditioner for MHD

Design principles:

- Preserve coincidence of MHD unknowns on coarse levels
- Reduce memory footprint by avoiding global ILU smoothers
- Performance through ILU as single-field smoothers

Block smoother:

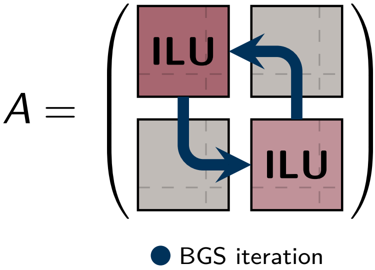

- Reuse aggregates $\mathcal{A}$ from Navier-Stokes part for magnetics part
- Non-smoothed transfer ops.
- Block smoother:
  n BGS(ω)
  - ILU(0), ov=1
  - ILU(0), ov=1

Block multigrid preconditioner for MHD

Design principles:

- Preserve coincidence of MHD unknowns on coarse levels
- Reduce memory footprint by avoiding global ILU smoothers
- Performance through ILU as single-field smoothers

Block smoother:

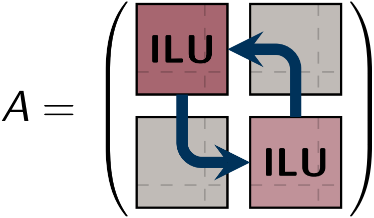

● BGS iteration

- Reuse aggregates $\mathcal{A}$ from Navier-Stokes part for magnetics part
- Non-smoothed transfer ops.
- Block smoother:
  n BGS(ω)
  - ILU(0), ov=1
  - ILU(0), ov=1

Reuse of aggregates

Colocated degrees of freedom allow for reuse of aggregates.

- more efficient setup
- one-to-one relationship between coarse Maxwell DoFs and associated Navier-Stokes DoFs

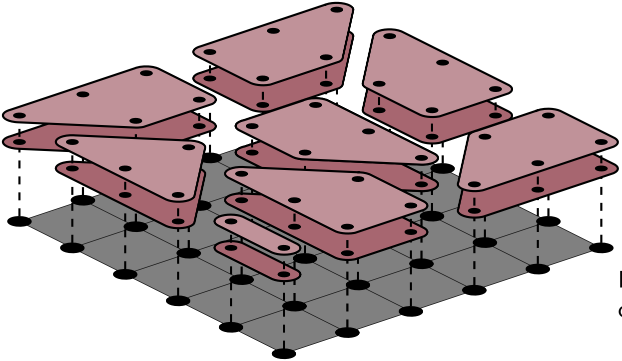

Cloned aggregates for Maxwell equations.

Aggregates built from Navier-Stokes equations.

Node discretization of current level.

MHD generator problem

- Stabilized discretization of MHD equations using equal order piecewise bilinear elements on hexahedrons.
- 3D channel geometry of size $(15, 1, 1)$

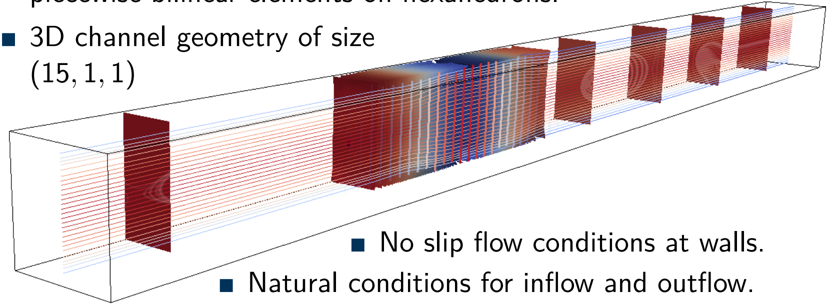

- No slip flow conditions at walls.
- Natural conditions for inflow and outflow.
- Dirichlet BCs for magnetic field at top and bottom

$$\begin{bmatrix} B_x \\ B_y \\ B_z \end{bmatrix} = \begin{bmatrix} 0 \\ \frac{1}{2} B_0(\tanh(x - x_0)/\Delta - \tanh(x - x_f)/\Delta) \end{bmatrix}$$

with $x_0 = 4.0, x_f = 6.0, \Delta = 0.5$ and $B_0 = 3.354$.

MHD generator problem – results

Sandia National Laboratories

240 × 16 × 16 mesh on 32 processors

Averaged timings after 5 runs using same nodes on cluster

Iterations:

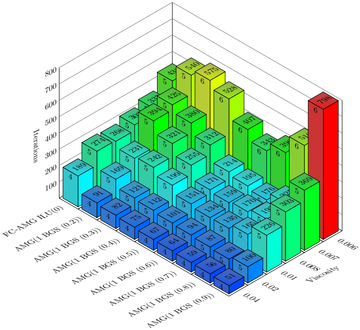

Timings:

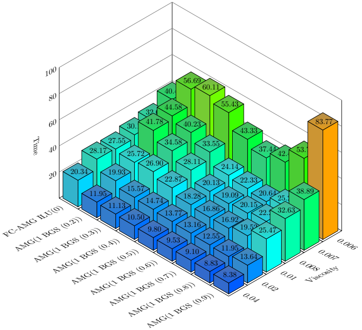

Relative linear solver tolerance: $\varepsilon = 1 \cdot 10^{-5}$

Multigrid hierarchy:

ℓ	rows	nnz	nnz/row	c ratio	procs	smoother
0	491,520	97,234,432	197.82		32	1 BGS (ω)
1	23,040	3,024,896	131.29	21.33	32	1 BGS (ω)
2	1,280	95,872	74.90	18.00	5	direct

MHD generator problem – results

$480 \times 32 \times 32$ mesh on 256 processors

Iterations:

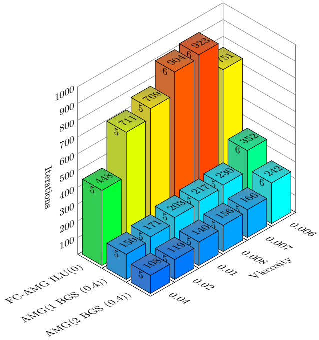

Timings:

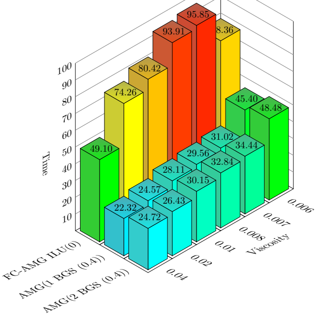

Multigrid hierarchy:

ℓ	rows	nnz	nnz/row	c ratio	procs
0	3,932,160	813,194,752	206.81		256
1	184,320	27,105,344	147.06	21.33	256
2	10,240	1,127,616	110.12	18.00	40
3	1,160	78,792	67.92	8.83	4

MHD generator problem - Mixed FE

- Stabilized discretization of MHD equations using **Q2 for hydrodynamics** and **Q1 for electromagnetics**.

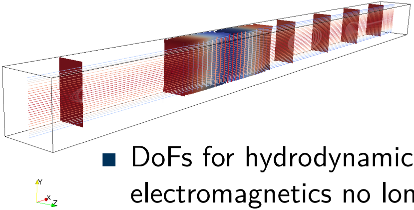

- DoFs for hydrodynamics and electromagnetics no longer colocated
- Hard for standard MueLu AMG kernels
  - Treat PDE system as if if is a scalar PDE
- Blocked approach with individual aggregation

MHD generator problem – results

Sandia National Laboratories

	Preconditioner	Proc	visc	n_N	n_L	n_L/n_N	Setup Time	Solve Time
$64 \times 32 \times 32$	AMG(1 BGS (0.6))	32	0.04	5	458	91.6	52.4024	327.15104
	AMG(1 BGS (0.6))	32	0.01	5	421	84.2	52.2752	295.11249
	AMG(1 BGS (0.6))	32	0.006	6	970	161.7	62.7376	709.46071
	ILU	32	0.04	7	3016	430.9	323.842	913.46691
	ILU	32	0.01	11	5016	456.0	510.683	1514.91137
	ILU	32	0.006	21	1016	477.0	978.112	3029.31644
$128 \times 64 \times 64$	AMG(1 BGS (0.6))	256	0.04	5	663	132.6	56.9724	519.58234
	AMG(1 BGS (0.6))	256	0.01	5	475	95.0	56.8077	357.89469
	AMG(1 BGS (0.6))	256	0.006	6	723	120.5	68.187	556.24512
	ILU	256	0.04	15	7024	468.3	758.517	2200.64326
	ILU	256	0.01	19	9024	474.9	963.582	2831.3441
	ILU	256	0.006	8	3524	440.5	404.601	1098.83657
$256 \times 128 \times 128$	AMG(1 BGS (0.6))	2048	0.04	5	1080	216.0	59.9361	1003.43002
	AMG(1 BGS (0.6))	2048	0.01	6	993	165.5	71.4519	856.84345
	AMG(1 BGS (0.6))	2048	0.006	6	944	157.3	71.1785	805.97758
	ILU	2048	0.04	21	10032	477.7	1103.45	3554.11075
	ILU	2048	0.01	22	10532	478.7	1153.6	3729.02421
	ILU	2048	0.006	11	5032	457.4	573.599	1760.38893

MHD generator problem - Mixed FE

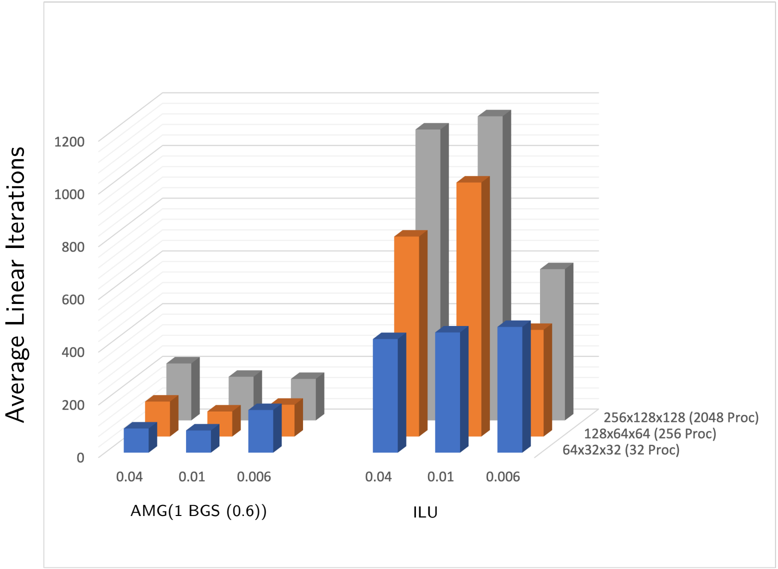

Questions?

Sandia National Laboratories

In collaboration with

- John N. Shadid (SNL)
- Eric C. Cyr (SNL)
- Tobias Wiesner (Leica Geosystems AG)
- Jonathan J. Hu (SNL)
- Raymond S. Tuminaro (SNL)

Acknowledgement
This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Scientific Discovery through Advanced Computing (SciDAC) program and Applied Mathematics program.