SAND2016-5004C

Evaluation of the performance of smoothers for fully-coupled algebraic multigrid preconditioners for implicit variational multiscale finite element resistive magnetohydrodynamics

ECCOMAS 2016

Crete, Greece

June 7, 2016

Paul Lin, John N. Shadid, Jonathan J. Hu
Center for Computing Research
Sandia National Labs
Paul Tsuji
Lawrence Livermore National Laboratory
Livermore CA

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.

Need for large-scale simulations

- Goal: high-fidelity solutions of transport phenomena for large-scale problems with complex physics and geometry
  - CFD, magnetohydrodynamic (MHD) simulations
  - US DOE interest: pulsed fusion reactors (e.g. z-pinch), magnetically confined fusion (e.g. ITER tokamak)
- One approach: FEM on unstructured meshes
- Fully-implicit Newton-Krylov solution approach
  - robust (promising for complex physics and chemistry)
  - preconditioner critical: robustness, efficiency, scalability
  - large-scale problems: multilevel/multigrid
- Talk focus: our fully-coupled Newton-Krylov multigrid preconditioned approach for large-scale FEM simulations
  - Drekar CFD/MHD application code (resistive MHD)
  - Trilinos linear solvers

Resistive MHD model

(J. Shadid, R. Pawlowski, E. Cyr, L. Chacon)

Navier-Stokes + Magnetic Induction

$$\rho\frac{\partial \mathbf{u}}{\partial t} + \rho(\mathbf{u}\cdot\nabla\mathbf{u}) - \nabla\cdot(\mathbf{T} + \boxed{\mathbf{T_M}}) - \rho\mathbf{g} = 0$$

$$\mathbf{T} = -(P + \frac{2}{3}\mu(\nabla\cdot\mathbf{u}))\mathbf{I} + \mu[\nabla\mathbf{u} + \nabla\mathbf{u}^T]$$

$$\boxed{\mathbf{T_M} = \frac{1}{\mu_0}\mathbf{B}\otimes\mathbf{B} - \frac{1}{2\mu_0}\|\mathbf{B}\|^2\mathbf{I}}$$

$$\frac{\partial \rho}{\partial t} + \nabla\cdot(\rho\mathbf{u}) = 0$$

$$\rho C_p\left[\frac{\partial T}{\partial t} + \mathbf{u}\cdot\nabla T\right] + \nabla\cdot\mathbf{q} - \boxed{\eta\|\mathbf{J}\|^2} = 0$$

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

Steady-state MHD generator
- Flow with external cross-stream B field
- 8 DOFs/mesh node

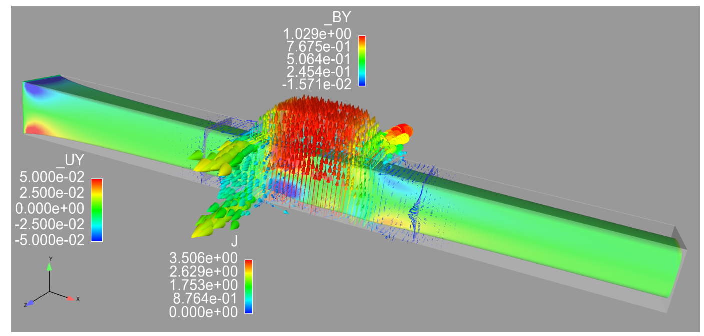

Drekar implicit/IMEX FE application

(J. Shadid, R. Pawlowski, E. Cyr, T. Smith, T. Wildey, E. Phillips, etc.)

- Navier-Stokes, MHD, LES, RANS
- stabilized FEM, unstructured hexahedral meshes
- fully-coupled multigrid preconditioned Newton-Krylov solve (Trilinos)
- Currently MPI-only; being transitioned to Kokkos for threading+GPUs

Numerical Solution Approach

- Resistive (incompressible) MHD model
- Most MHD turbulence simulations employ spectral methods
- Stabilized FEM: variational multiscale (VMS) (Hughes 1995)
  - many authors employ VMS for CFD
- MHD turbulence with VMS (Shadid et al, Sondak+Oberai)
- Unstructured meshes
- Fully-coupled Newton-Krylov (multigrid preconditioned)
  - Robust (promising for complex physics and chemistry)
  - but depends on efficiency of sparse linear solver
  - preconditioner critical: robustness, efficiency, scalability
  - large-scale problems: multilevel/multigrid

Smoothers for multigrid preconditioned approach

- Effective smoothers critical for multigrid
  - Need to efficiently damp high frequency errors
  - Standard relaxation not sufficiently robust for our MHD problems
  - Our standard smoother ILU(0) overlap=1 is expensive
  - Interested in other smoothers
- Krylov smoothers
  - Lots of previous work for SPD problems (Notay, Vassilevski, etc.)
  - Some previous work for Helmholtz (Elman)
  - Far less previous work for nonsymmetric systems
    - Birken, Bull, Jameson Thu 17:20
  - Setup much cheaper than ILU, but solve can be expensive
  - ~~Issues with GS or ILU for problems with zeroes on diagonal~~
    - ~~Krylov smoothers could be advantageous~~
- Evaluate smoothers
  - Transient (our main interest)
  - Steady-state (MHD pumps, MHD generator)

Brief Trilinos overview

- Classic Trilinos (Epetra-based) (Heroux et al.):
  - Limited by 32-bit integer global objects
  - Most packages employ flat MPI-only; future architectures?
- Trilinos solver stack (Tpetra-based):
  - No 32-bit limitation on globals (employs C++ templated data types)
  - Path forward for future architectures: Trilinos Kokkos (Edwards, Trott, Sunderland; not part of this talk)

Functionality	Classic stack	Newer solver stack
Distributed linear algebra	Epetra	Tpetra (Hoemmen,Trott, etc.)
Iterative linear solve	Aztec	Belos (Thornquist,Hoemmen,etc.)
Incomplete factor	Aztec, Ifpack	Ifpack2 (Hoemmen,Hu,Siefert, etc.)
Algebraic multigrid	ML	MueLu (Hu,Prokopenko,Wiesner,Siefert,Tuminaro,etc.)
Partition & load balance	Zoltan	Zoltan2 (Devine,Boman,Rajamanickam,Wolf,etc.)
Direct solve interface	Amesos	Amesos2 (Rajamanickam,etc.)

- PETSc is another well-known solvers library (ANL; Smith, Gropp, Knepley, Brown, McInnes, Balay, Zhang, et al.); 2015 SIAM/ACM CSE prize winner

Trilinos MueLu Library: algebraic multigrid preconditioners

(J. Hu, A. Prokopenko, J. Gaidamour, T. Wiesner, C. Siefert, R. Tuminaro)

- Smoothed aggregation; aggregates to produce a coarser operator
  - Create graph where vertices are block nonzeros in matrix A_k
  - Edge between vertices i and j added if block $B_k(i,j)$ contains nonzeros
  - Uncoupled aggregation
- Restriction/prolongation operator; $A_{k-1} = R_k A_k P_k$
- Repartition coarser level matrices (MueLu+Zoltan2) to reduce communication
- Coarsest level: serial direct solve (KLU; T. Davis) on 1 MPI process

Other approaches: LLNL Hypre (R. Falgout, U. Yang, T. Kolev, A. Baker, E. Chow, C. Tong, et al.), MLBDDC (S. Badia, A. Martín, J. Principe, et al.), etc.

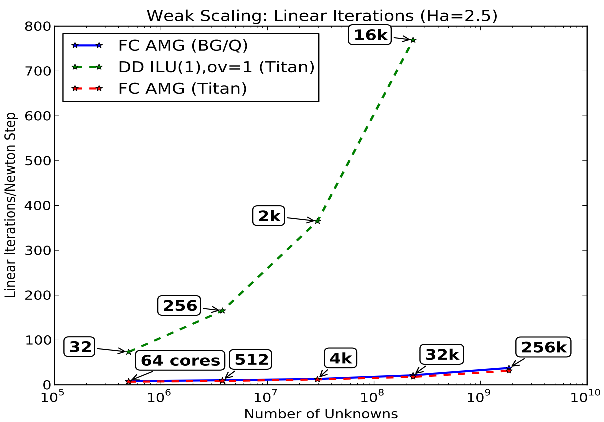

- Weak scaling: MHD generator
- Re = 500, Re_m = 1, Ha = 2.5
- Cray XK7, IBM Blue Gene/Q

Additive Schwarz domain decomposition does not scale
Multigrid critical for performance and scaling

Strong scaling: Poisson equation

(with J. Hu, J. Shadid, A. Prokopenko, E. Cyr, R. Pawlowski)

- 3D Poisson (1 DOF/mesh node)

(Image courtesy of LLNL)

- Simple cube geometry, near uniform mesh
- Fixed problem size (2.4b DOFs); 1 MPI task/core BG/Q
- Optimal iteration count to 1.6 million cores (full-scale Sequoia BG/Q)

MPI process	CG iterations	Solve t (s)	MG setup (s)	DOFs/MPI
131,072	6.3	1.17	7.67	~18,800
262,144	6.0	1.08	12.35	~9400
524,288	6.3	1	25.43	~4700
1,048,576	7.3	0.91	53.04	~2400
1,572,864	7.0	0.94	128.9	~1500

Weak scaling: fully-coupled multigrid MHD

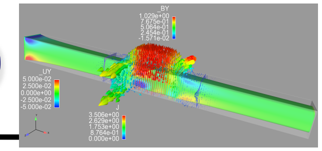

(with J. Hu, J. Shadid, A. Prokopenko, E. Cyr, R. Pawlowski)

MPI	DOFs	GMRES iterations /Newton	Time/Newton step (s)		
			Multigrid setup		Solve
			Hier+smoo	Smoother	
128	845,000	14.0	12.4	11.0	4.7
1024	6,473,096	20.0	14.7	13.0	6.6
8192	50,658,056	30.8	16.9	14.2	10.1
65,536	400,799,240	53.4	20.3	16.1	17.9
524,288	3,188,616,200	98.7	45.3	19.1	40.1

- BG/Q: 1 MPI/core
- Multigrid prec setup time/Newton step
- Smoother: ILU(0) overlap=1

- Drekar 3D MHD generator on BG/Q (simple geometry)
- Algorithmic scaling challenging for nonsymmetric matrices
  - 4096x increase in size: 6.0x iterations, 7.3x time
  - Petrov-Galerkin or energy minimization approaches promising
  - Need better aggregation, better smoothers, etc.
- Another challenge: sparse matrix-matrix multiply (A_c=R*A*P)
- Employ reuse of construction of hierarchy and smoothers (Prokopenko)
  - Application dependent (e.g. cannot reuse for adaptive mesh)
  - Critical for transient simulations (10^4 or 10^5 time steps)

Results (Transient)

Transient Taylor-green MHD vortex decay (resistive MHD VMS turbulence) (8 DOF/mesh node)

- Performance comparison of different smoothers (CFL ~0.5)
  - 3 problem sizes: 16.8M, 134M, 1.07b DOFs
  - GMRESR "outer" Krylov solve
  - Jacobi, block Jacobi, GS, block GS failed for even small problems and were not included in following tables
  - ILU(0) overlap=0,1
  - GMRES (no prec, Gauss-Seidel, block Jacobi, block GS)
  - Schwarz/domain decomposition GMRES smoother
    - no prec, Gauss-Seidel, block Jacobi, block GS
    - Overlap=0,1
- Robustness for higher CFL study
  - 2 problem sizes: 16.8M,1.07b DOFs
  - Focus on more promising smoothers

Transient Taylor-Green MHD
Smoother Comparison (16.8M DOFs)

- Transient Taylor-green MHD vortex decay (VMS resistive MHD)
- Cube domain, 128^3 elem, 20 time steps, dt=0.025 (CFL ~0.5)
- 256 MPI; linux cluster dual-socket SNB, IB fat-tree (TLCC2)

smoother		iters/dt	Prec setup(s)	Solve(s)	Prec+solve(s)	Total(s)	Mem(MB)
SGS		87.7	18	730	748	1108	1050
ILU	overlap=0	20.9	127	110	237	606	1157
	overlap=1	14.2	259	97	356	725	1436
GMRES	noprec	15.4	23	260	283	696	917
	ptGS	13.6	24	413	437	828	917
	bkJac	13.1	35	238	273	665	927
	bkGS	12.0	34	539	573	950	930
DD-GMRES	noprec ov0	21.4	35	530	565	929	917
	noprec ov1	15.9	125	486	611	995	1107
	ptGS ov0	20.2	51	986	1037	1396	917
	ptGS ov1	12.1	142	759	901	1263	1102
	bkJac ov0	20.2	45	535	581	942	925
	bkJac ov1	15.5	151	512	663	1069	1129
	bkGS ov0	20.3	69	1170	1239	1597	924
	bkGS ov1	11.9	162	902	1063	1422	1129

- DD-GMRES smoother not competitive (particularly GS and block GS prec)

Note: have multiple Newton steps per time step

Transient Taylor-Green MHD Smoother Comparison (134M DOFs)

- Transient Taylor-green MHD vortex decay (VMS resistive MHD)
- Cube domain, 256^3 elem, 20 time steps, dt=0.0125 (CFL ~0.5)
- 2048 MPI; linux cluster dual-socket SNB, IB fat-tree (TLCC2)

smoother		iters/dt	Prec setup(s)	Solve(s)	Prec+solve(s)	Total(s)	Mem(MB)
SGS		Failed					
ILU	overlap=0	28.85	138.3	154	292	680	1164
	overlap=1	20.15	311.1	134	445	862	1440
GMRES	noprec	18.65	23.5	328	352	741	920
	ptGS	16.85	26.84	516	542	921	920
	bkJac	15.65	41.33	306	348	749	933
	bkGS	17.05	40.94	954	995	1376	936
DD-GMRES	noprec ov0	31.6	43.25	842	885	1272	926
	noprec ov1	23.95	160	1094	1254	1657	1111
	ptGS ov0	27.25	61.01	1655	1716	2083	939
	ptGS ov1	16.9	185	1383	1568	1948	1111
	bkJac ov0	28	51.29	789	840	1218	931
	bkJac ov1	17.8	178.3	872	1050	1455	1134
	bkGS ov0	27.5	74.66	1881	1956	2319	939
	bkGS ov1	16.85	201.9	1582	1784	2166	1134

- DD-GMRES smoother not competitive (particularly GS and block GS prec)

Transient Taylor-Green MHD Smoother Comparison (1.07b DOFs)

- Transient Taylor-green MHD vortex decay (VMS resistive MHD)
- Cube domain, 512^3 elem, 20 time steps, dt=0.0625 (CFL ~0.5)
- 16384 MPI; linux cluster dual-socket SNB, IB fat-tree (TLCC2)

smoother		iters/dt	Prec setup(s)	Solve(s)	Prec+solve(s)	Total(s)	Mem(MB)
ILU	overlap=0	68.1	251	411	662	1380	1287
	overlap=1	37.3	407	280	687	1280	1519
GMRES	noprec	31.0	92	610	702	1326	1002
	ptGS	21.5	94	728	822	1374	1002
	bkJac	21.9	60	458	518	1028	1017
DD-GMRES	noprec ov0	69.0	61	2040	2101	2689	1031
	bkJac ov0	69.7	135	2227	2362	3088	1040

- DD-GMRES smoother not competitive
- GMRES smoother with either no preconditioner or block Jacobi is most competitive to ILU
  - GMRES can significantly lower iterations, but cost/iteration expensive
  - Trade-off between expensive ILU factorization for setup vs. solve of GMRES smoother
  - ILU smoother requires more memory

Transient Taylor-Green MHD
Smoother Comparison: Robustness with CFL (16.8M DOFs)

- Transient Taylor-green MHD vortex decay (VMS resistive MHD)
- Cube domain, 128^3 elem, 10 time steps
- 256 MPI; linux cluster dual-socket SNB, IB fat-tree (TLCC2)

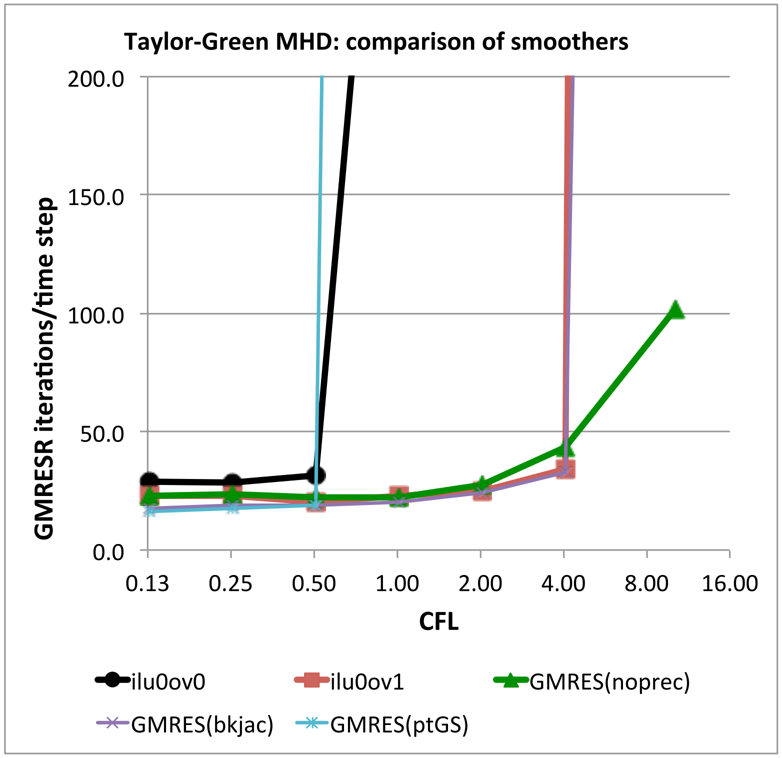

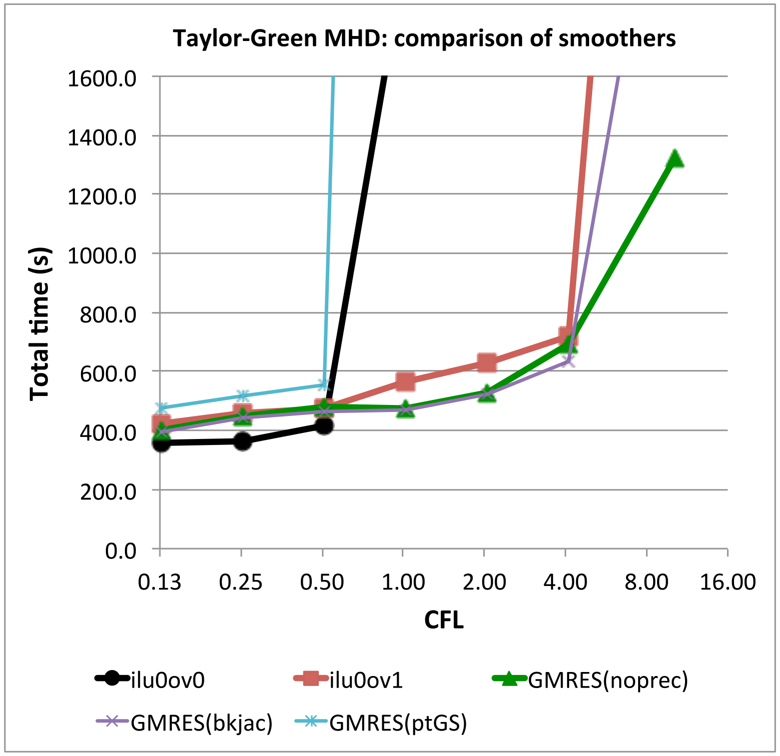

- ILU(0) overlap=0 and GMRES(bkGS prec) smoothers not sufficiently robust
- Need to compare ILU(0) overlap=1 with GMRES smoothers

U.S. DEPARTMENT OF ENERGY
NNSA National Nuclear Security Administration
Sandia National Laboratories

Transient Taylor-Green MHD
Smoother Comparison: Robustness with CFL (1.07b DOF)

- Transient Taylor-green MHD vortex decay (VMS resistive MHD)
- Cube domain, 512^3 elem, 10 time steps; 16384 MPI (BG/Q)

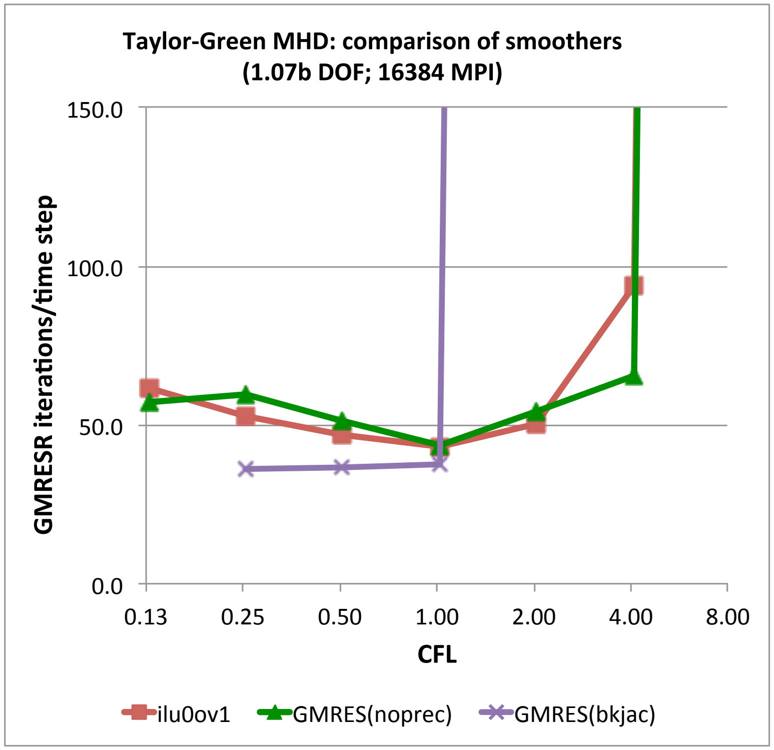

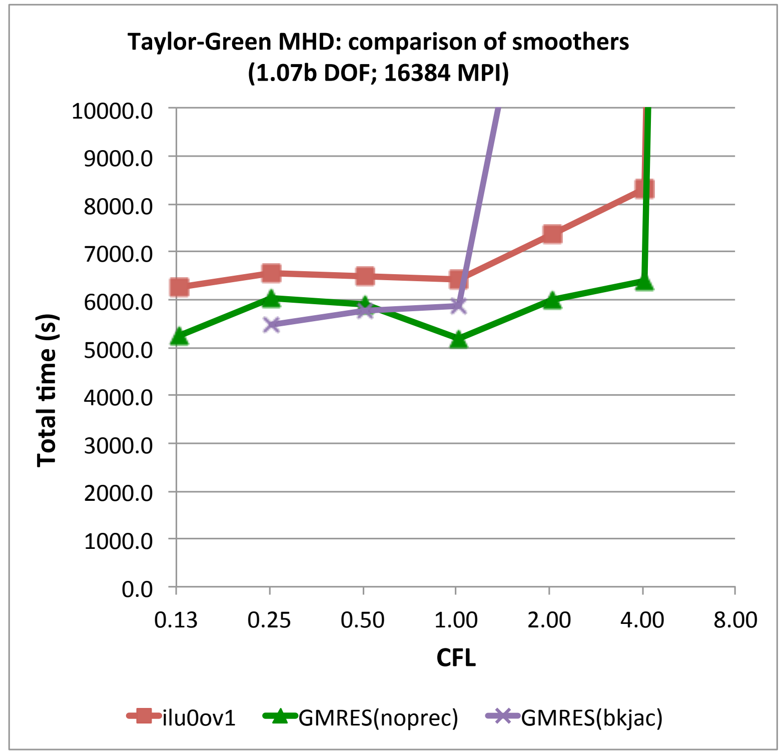

- Did not even run ILU(0) overlap=0 because fails at low CFL (unfortunate because ILU overlap=1 requires ~40% more memory than GMRES smoother)
- GMRES(bkGS prec) fails CFL ~2
- ILU(0) overlap=1, GMRES(no prec) fail CFL ~8

Sandia National Laboratories

Results (Transient): Island Coalescence

Transient resistive MHD (8 DOF/mesh node)

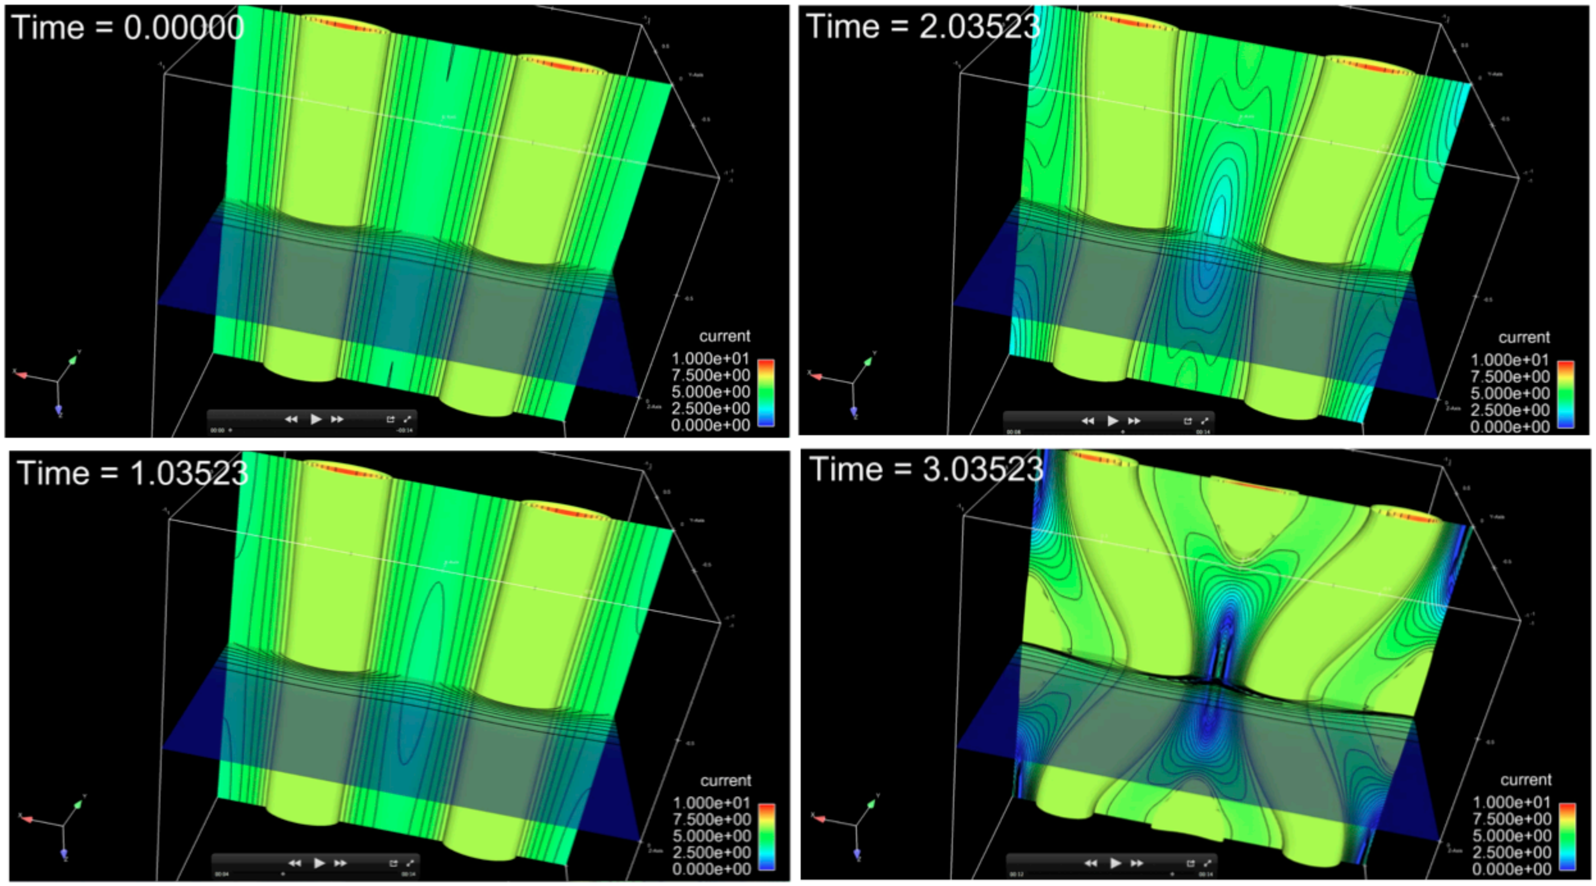

Results (Transient): Island Coalescence

- Fixed time step of 0.1 (CFL increases as refine mesh)
- Smoothers: ILU, GMRES (GS, block GS stalled)

64^3 elem
(2.1M DOF)
64 MPI

smoother		iters/dt	Prec setup(s)	Solve(s)	Prec+solve(s)	Mem(MB)
ILU	overlap=0	14.1	200.6	63.0	263.6	773
	overlap=1	6.8	351.2	36.3	387.5	914
GMRES	noprec	8.7	20.7	122.2	142.9	630
	bkJac	6.3	33.5	101.5	135.0	642

128^3 elem
(16.9M DOF)
512 MPI

smoother		iters/dt	Prec setup(s)	Solve(s)	Prec+solve(s)	Mem(MB)
ILU	overlap=0	12.2	210.2	61.9	272.1	773
	overlap=1	7.7	382.9	50.8	433.7	917
GMRES	noprec	11.1	27.4	188.9	216.3	632
	bkJac	7.0	42.2	137.1	179.3	645

256^3 elem
(135M DOF)
4096 MPI

smoother		iters/dt	Prec setup(s)	Solve(s)	Prec+solve(s)	Mem(MB)
ILU	overlap=0	13.3	241.2	73.4	314.6	792
	overlap=1	10.1	442.6	70.9	513.5	933
GMRES	noprec	15.1	58.9	275.3	334.2	670
	bkJac	8.9	87.5	194.3	281.8	666

For island coalescence test case, GMRES smoother is faster than our standard ILU smoother and requires less memory

Concluding remarks and future work

- Performed an initial evaluation of GMRES as an alternative smoother to our current standard (ILU)
  - Initial empirical study for two test cases
  - Shows promise: can solve an initial class of relevant problems (appears competitive; expensive, but so is ILU)
  - Memory usage benefits (ILU requires ~40% more)
  - ~~Krylov smoother should be advantageous for nonsymmetric matrices with zero diagonal entries (but still need to demonstrate)~~
- Need to look at many more test cases
- Issues
  - Need to go back and try to analyze method more carefully
    - Which Krylov method better at killing off high frequency modes?
  - AMG notorious for too many adjustable parameters
    - Krylov smoothers exacerbate this issue (e.g. precondition? sweeps? tolerance? Etc.)

Additional future work

- Many challenges for multigrid-preconditioned linear solve
  - algorithmic scaling for nonsymmetric problems
  - multigrid preconditioner setup (sparse mat-mat)
- Trilinos Kokkos for manycore and accelerators (“X” for MPI+X)
  - Tpetra with Kokkos implementation is ongoing work
  - other Trilinos packages in process of being implemented with Kokkos: MueLu setup, additional Ifpack2 smoothers, etc.
- Drekar progress depends on above Kokkos work
  - Trilinos components for assembly ported (I. Demeshko)
  - Related: Bettencourt 6:40pm performance portability for multi-fluid plasma assembly
- Additional physics and discretization issues include, e.g.
  - strong convection effects, hyperbolic systems
  - non-uniform FE aspect ratios

Thanks For Your Attention!

Paul Lin (ptlin@sandia.gov)
John Shadid, Jonathan Hu, Paul Tsuji

Acknowledgments:

- MueLu team: Andrey Prokopenko, Ray Tuminaro, Chris Siefert, Tobias Wiesner
- Drekar team: Roger Pawlowski, Eric Cyr, Edward Phillips. Tom Smith, Tim Wildey
- Mark Hoemmen, Eric Phipps
- LLNL LC BG/Q team (especially John Gyllenhaal and Scott Futral)

Funding Acknowledgment:
The authors gratefully acknowledge funding from the DOE NNSA Advanced Simulation & Computing (ASC) program and ASCR Applied Math Program