SAND2014-15818PE

Towards Scalable Fully-Coupled Algebraic Multigrid Preconditioners for Magnetohydrodynamic Simulations

MOST2014

Moselle Valley

July 30, 2014

Paul Lin, John Shadid, Jonathan Hu, Eric Cyr
Sandia National Labs

U.S. DEPARTMENT OF ENERGY
National Nuclear Security Administration

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.

Sandia National Laboratories

Need for Large-scale Simulations

- High fidelity solutions of transport phenomena for large-scale problems with complex physics
  - Magnetohydrodynamic (MHD) simulations
- Fully implicit Newton-Krylov solution approach
  - 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)
- Talk focus: our multigrid approach for large-scale FEM simulations

Brief Trilinos Overview

Object-oriented software framework for the solution of large-scale, complex multi-physics engineering and scientific problems

- Current Trilinos (Epetra based):
  - Limited by 32-bit integer global objects
    - need for increasing fidelity for challenging problems at SNL
  - Most packages employ flat MPI-only; future architectures?

Trilinos ML Library: Algebraic Multigrid Preconditioners

(R. Tuminaro, J. Hu, C. Siefert, M. Sala, M. Gee, C. Tong)

- 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 matrix (ML+Zoltan)

Another well-known AMG library: LLNL Hypre (Falgout, Yang, Kolev, Baker, etc.)

Drekar (JN Shadid, RP Pawlowski, EC Cyr, TM Smith, T Wildey)

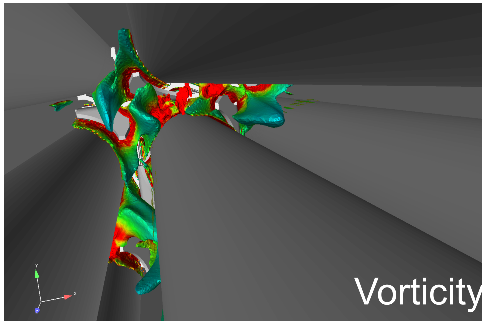

LES: Flow over spacer grid

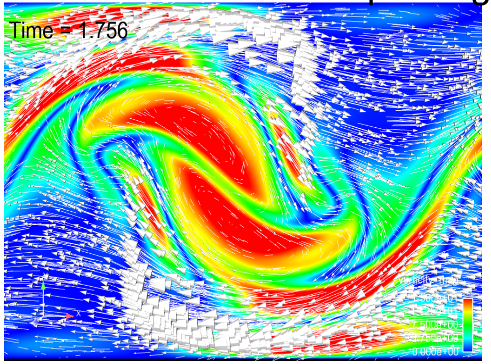

MHD: Hydromagnetic Kelvin-Helmholtz

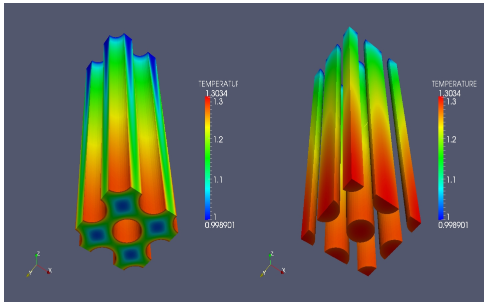

Conjugate Heat Transfer

Scalable parallel implicit FE code

- Includes: Navier-Stokes, MHD, LES, RANS
- Architecture admits new coupled physics
- Support of advanced discretizations
  - mixed, compatible and high-order basis functions
  - multi-physics capable (conjugate heat transfer)
- Advanced UQ techniques
  - Embedded stochastic Galerkin
  - Adjoint based sensitivities and error-estimates
- Advanced solution methods
  - Parallel solvers from SNL's Trilinos framework
  - Physics-based preconditioning
  - Fully-coupled multigrid for monolithic systems

Resistive MHD Model for MHD Generator Problem

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

Resistive MHD Model

Navier-Stokes + Electromagnetics

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

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

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

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

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

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

- Steady-state MHD Generator
- 3D flow with external cross-stream B field

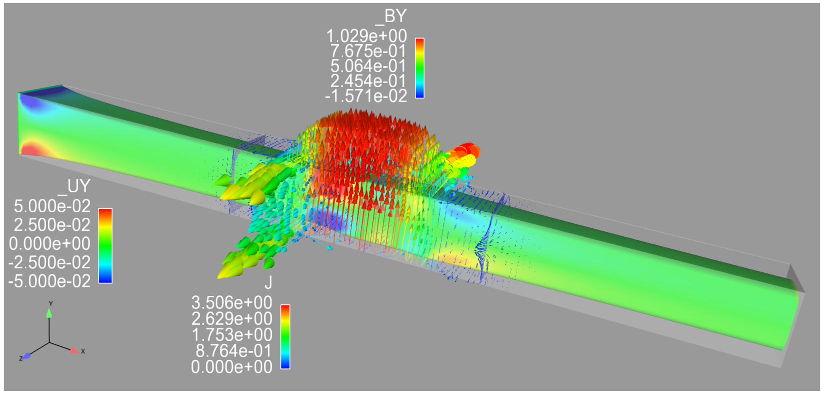

Initial Weak Scaling Studies: 3D MHD Generator

Re = 500, Re_m = 1, Ha = 2.5
Cray XK7 and IBM BG/Q

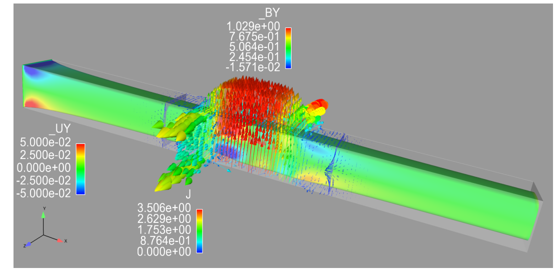

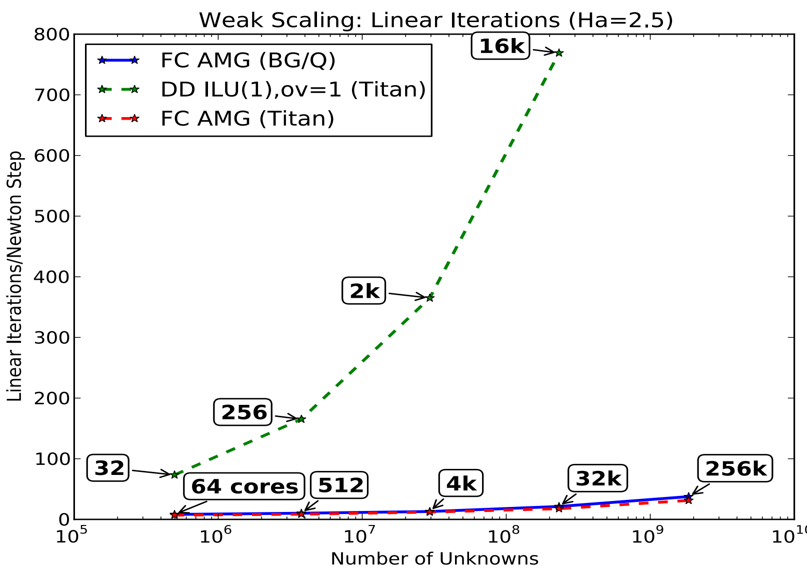

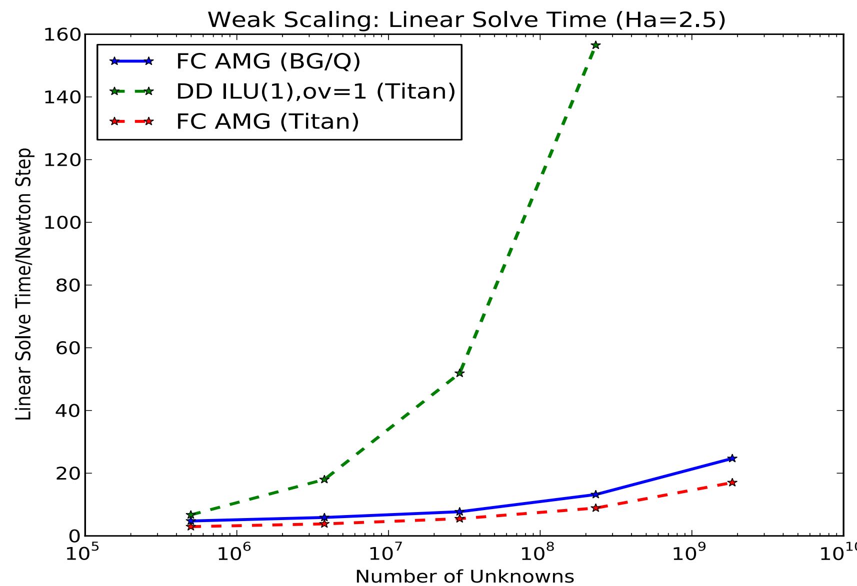

Initial Strong Scaling: 3D MHD Gen to 500,000 Cores BG/Q

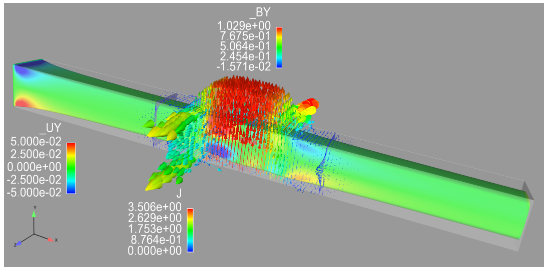

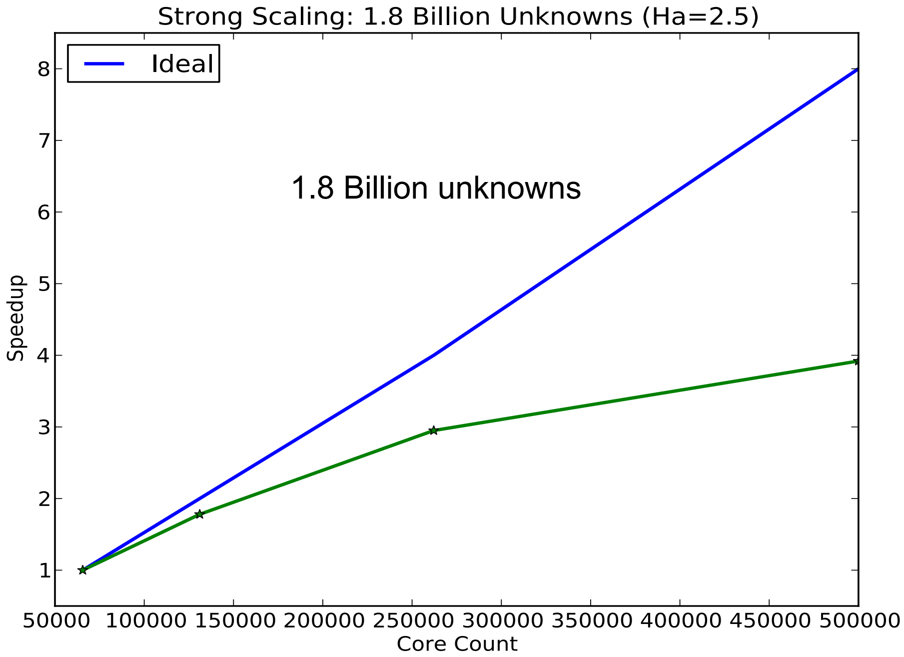

[Preliminary strong scaling of Krylov linear solver + preconditioner (ML: FC – AMG), Tuminaro, Hu, Siefert et. al.]

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

Scaling and Performance for Drekar

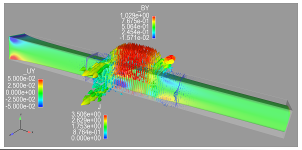

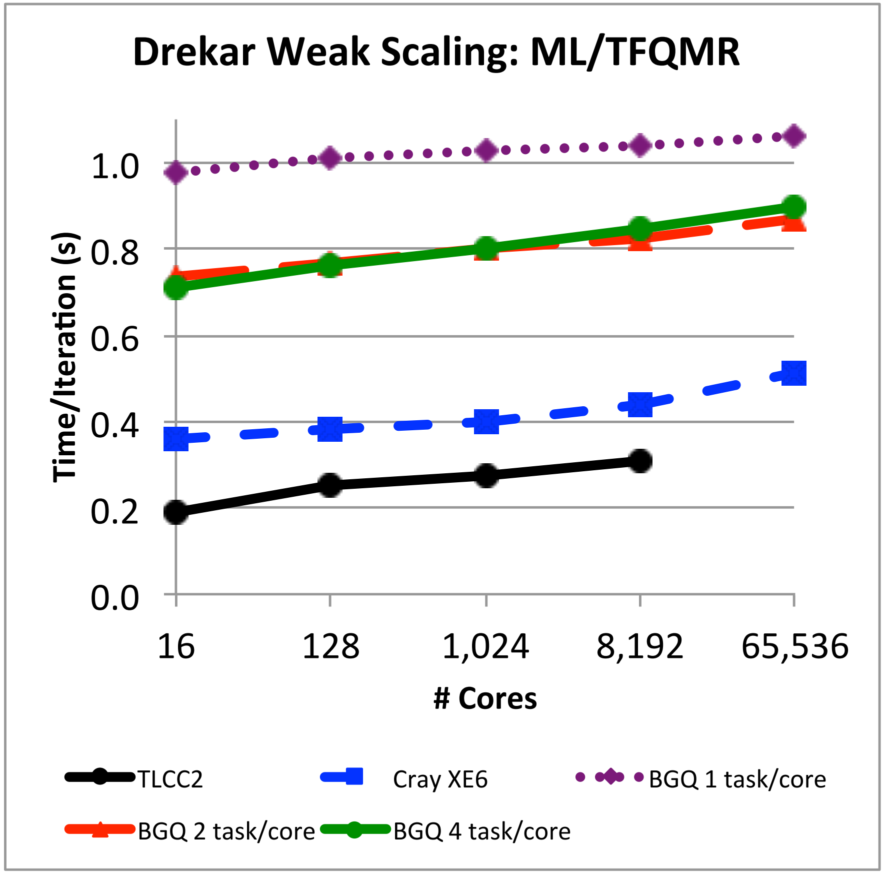

- 3D MHD Generator (~30k DOF/core)
- GNU 4.7.2 compiler on BGQ
- Matrix solve time, ML prec TFQMR
- 1 compute node (normalized time):
  - TLCC2: 1
  - XE6 (Cielo): 1.9
  - BGQ: 5.2 (1 task), 3.9 (2 tasks), 3.8 (4 tasks)
- 512 nodes (8192 cores; norm. time):
  - TLCC2: 1
  - XE6 (Cielo): 1.4
  - BGQ: 3.4 (1 task), 2.65 (2 tasks), 2.74 (4 tasks; total 32,768)
- 512 node est. power (Top500 list)
  - TLCC2: ~190kW
  - XE6: ~230kW
  - BGQ: ~40kW
  - real FLOPS/W: BGQ/TLCC2 ~ 1.8
- 4096 nodes (65,536 cores; normalize to XE6)
  - BGQ: 1.7 (2 task), 1.8 (4 task; 262144 tot)
  - Est power: XE6 ~1830kW, BGQ ~329kW
  - Real FLOPS/W: BGQ/XE6 ~3.2

Initial Study: 3D Hydromagnetic Kelvin-Helmholtz Instability

$Re = 10^4$, $Re_m = 10^4$, $M_A = 3$; CFL_{max} ~5
Cray XK7

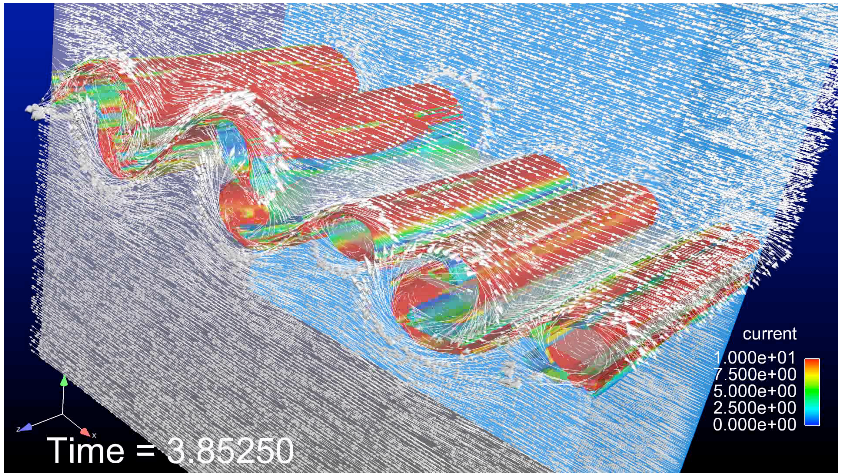

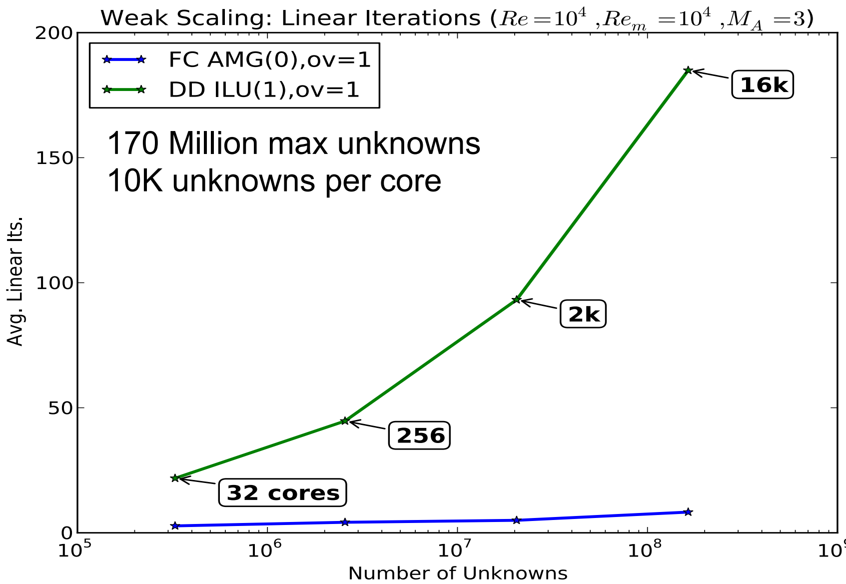

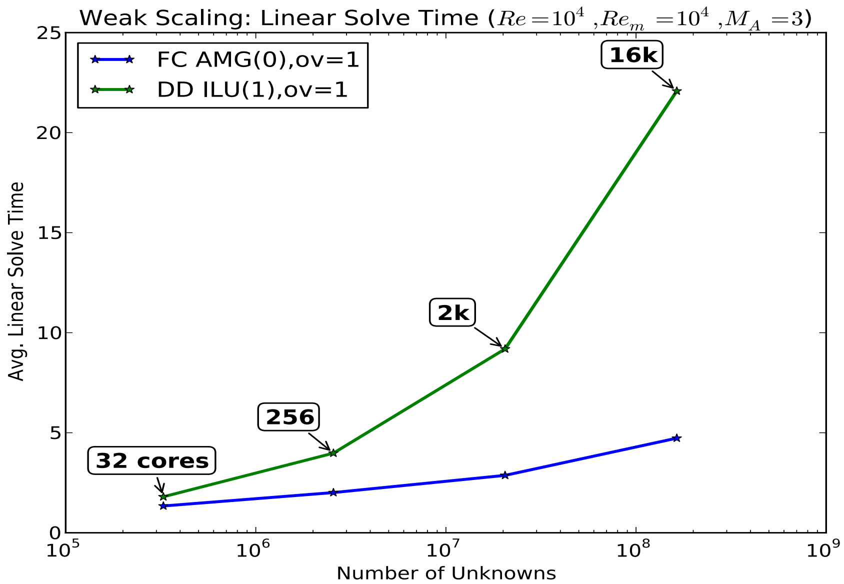

U.S. DEPARTMENT OF ENERGY
NNSA National Nuclear Security Administration
Sandia National Laboratories

3D Island Coalescence

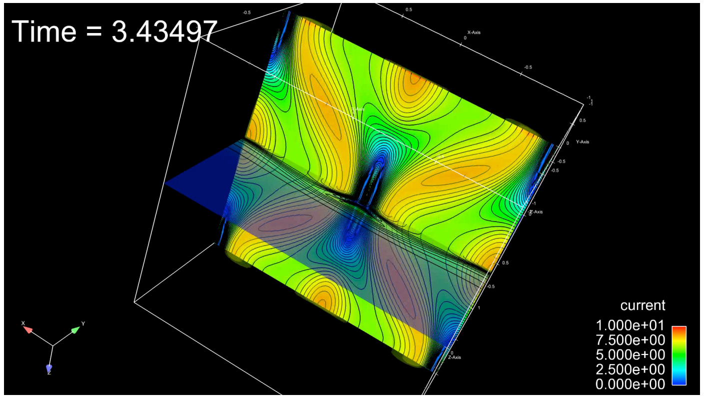

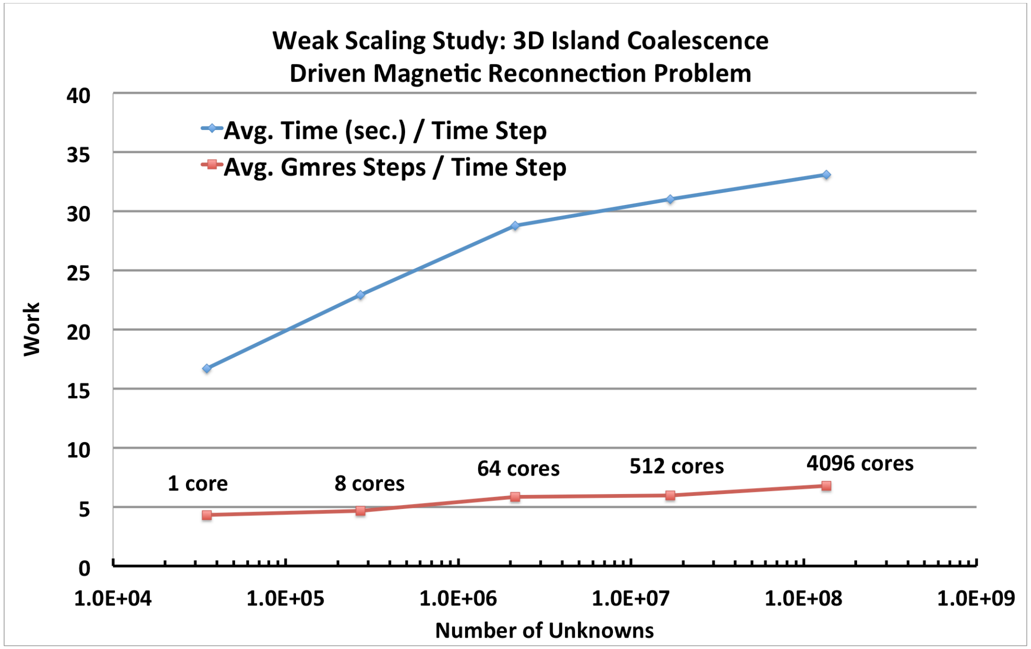

Scaling with Lundquist No.

Lundquist No. S	Newt. Steps / dt	Gmres Steps / dt
1.0E+03	1.36	5.2
5.0E+03	1.43	5.7
1.0E+04	1.51	6
5.0E+04	2	9.8
1.0E+05	2	12
5.0E+05	2	8.4
1.0E+06	2	8.4

Mesh: 128x128x128, dt = 0.0333.

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

Brief Second GenerationTrilinos Overview

Object-oriented software framework for the solution of large-scale, complex multi-physics engineering and scientific problems

- Original Trilinos (Epetra based):
  - Limited by 32-bit integer global objects
    - need for increasing fidelity for challenging problems at SNL
  - Most packages employ flat MPI-only; future architectures?
- Second generation Trilinos solver stack
  - No 32-bit limitation on global objects (employs C++ templated data types)
  - Path forward for future architectures (Kokkos; not part of this talk)
- Want Trilinos to impact production engineering applications
  - Significant effort to mature second generation Trilinos

Functionality	Current solver stack	New solver stack
Distributed linear algebra	Epetra	Tpetra
Iterative linear solvers	AztecOO	Belos
Incomplete factorizations	AztecOO, Ifpack	Ifpack2
Algebraic multigrid	ML	MueLu
Partition & load balance	Zoltan	Zoltan2
Direct solvers interface	Amesos	Amesos2

Trilinos MueLu Library: Algebraic Multigrid Preconditioners

(J. Hu, A. Prokopenko, J. Gaidamour, P. Tsuji, C. Siefert)

- 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 matrix (MueLu+Zoltan2)

3D CFD: Kelvin-Helmholtz Jet

VMS FE Simulation of 3D Kelvin-Helmholtz Unstable Jet Shear Layer with a Secondary Cross-stream Instability, Re = 10^8

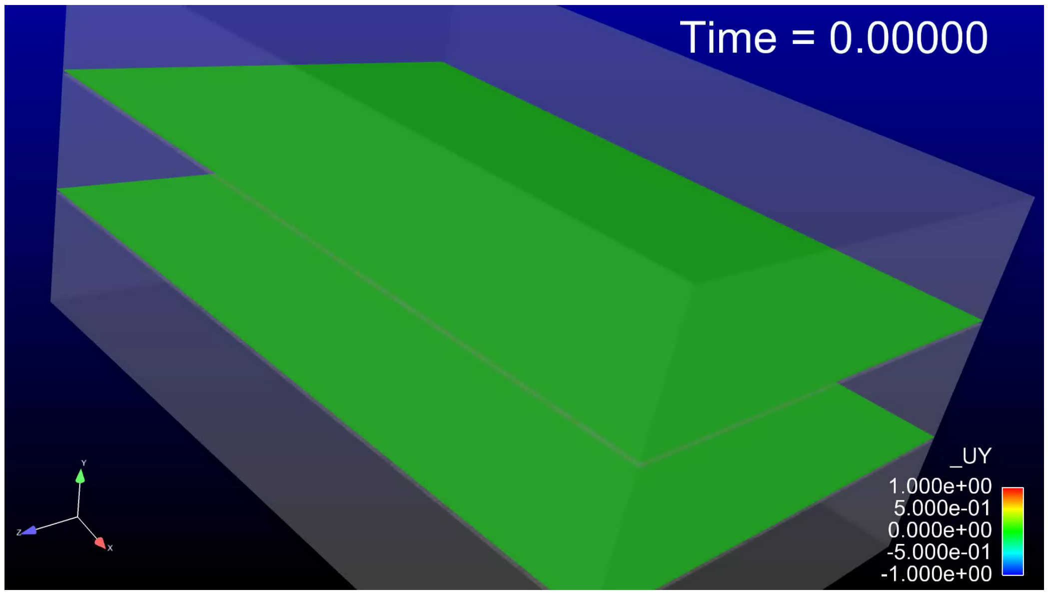

Preliminary Weak Scaling BG/Q: 3D Jet (CFL ~0.25)

Drekar/Epetra/Aztec/ML/Ifpack SGS

cores	DOF	Newt /dt	Iter/ Newt	Iter/ dt	Time/Newt (sec)		
					Prec	Solve	Jac
32	901056	3.40	8.41	28.6	1.81	8.54	31.44
256	6931504	3.50	10.8	37.8	1.95	10.5	31.57
2048	54,723,004	3.60	15.72	56.6	2.96	14.8	31.64
16,384	434,886,004	3.60	22.42	80.7	3.06	21	31.69
131,072	3,467,532,004	Cannot run problems with DOF > 2.1b					

Drekar/Tpetra/Belos/MueLu/Ifpack2 SGS

cores	DOF	Newt /dt	Iter/ Newt	Iter/ dt	Time/Newt (sec)		
					Prec	Solve	Jac
32	901,056	3.40	8.44	28.7	4.31	7.29	37.38
256	6,931,504	3.60	11.92	42.9	4.49	9.88	37.42
2048	54,723,004	3.70	16.32	60.4	4.71	13.9	37.46
16,384	434,886,004	3.90	24.59	95.9	6.35	20.2	37.46
131,072	3,467,532,004	4.00	36.3	145	163	30.4	37.55

Tpetra/MueLu enables simulations > 2.1b (32-bit)

MueLu for 3D Poisson Problem up to 1 Million Sequoia Cores

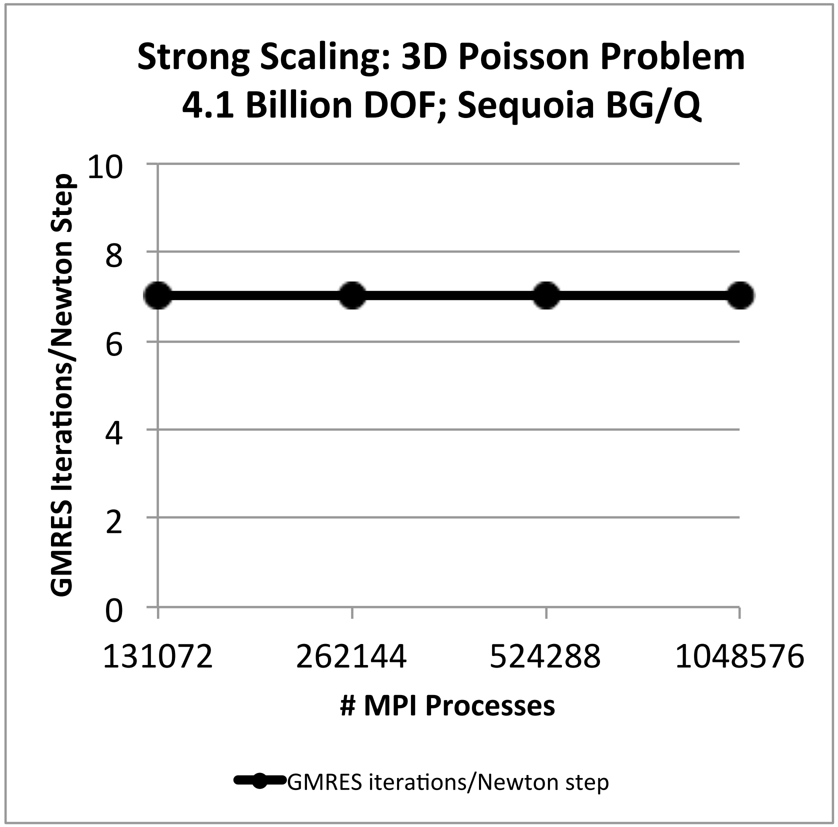

- Drekar 3D Poisson problem with MueLu preconditioned GMRES
- MueLu maintains optimal iteration count for 4.1 billion row matrix (4.1 billion elements) as increase cores to 1 million

Concluding Remarks and Future Work

- Newton-Krylov/AMG methods are promising for large-scale simulations (CFD/MHD)
- Scalable linear solvers critical to scalability and efficiency for large-scale simulation
- New Tpetra/Belos/MueLu stack
  - Allows > 2.1 billion DOFS
  - Demonstrated solution for 3.5 billion DOF CFD problem
  - Demonstrated solution for 4 billion element Poisson problem on 1 million cores
  - Need to improve preconditioner setup and iteration count
- Issues
  - Strong convection effects, hyperbolic systems
  - Highly non-uniform FE aspect ratios
- Path forward for future architectures: Trilinos Kokkos

Thanks For Your Attention!

Paul Lin (ptlin@sandia.gov)

Acknowledgments:

- Roger Pawlowski on Drekar Team
- Andrey Prokopenko on MueLu team
- Ray Tuminaro on ML team
- LLNL BG/Q team

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

Extra Slides

Why Newton-Krylov Methods?

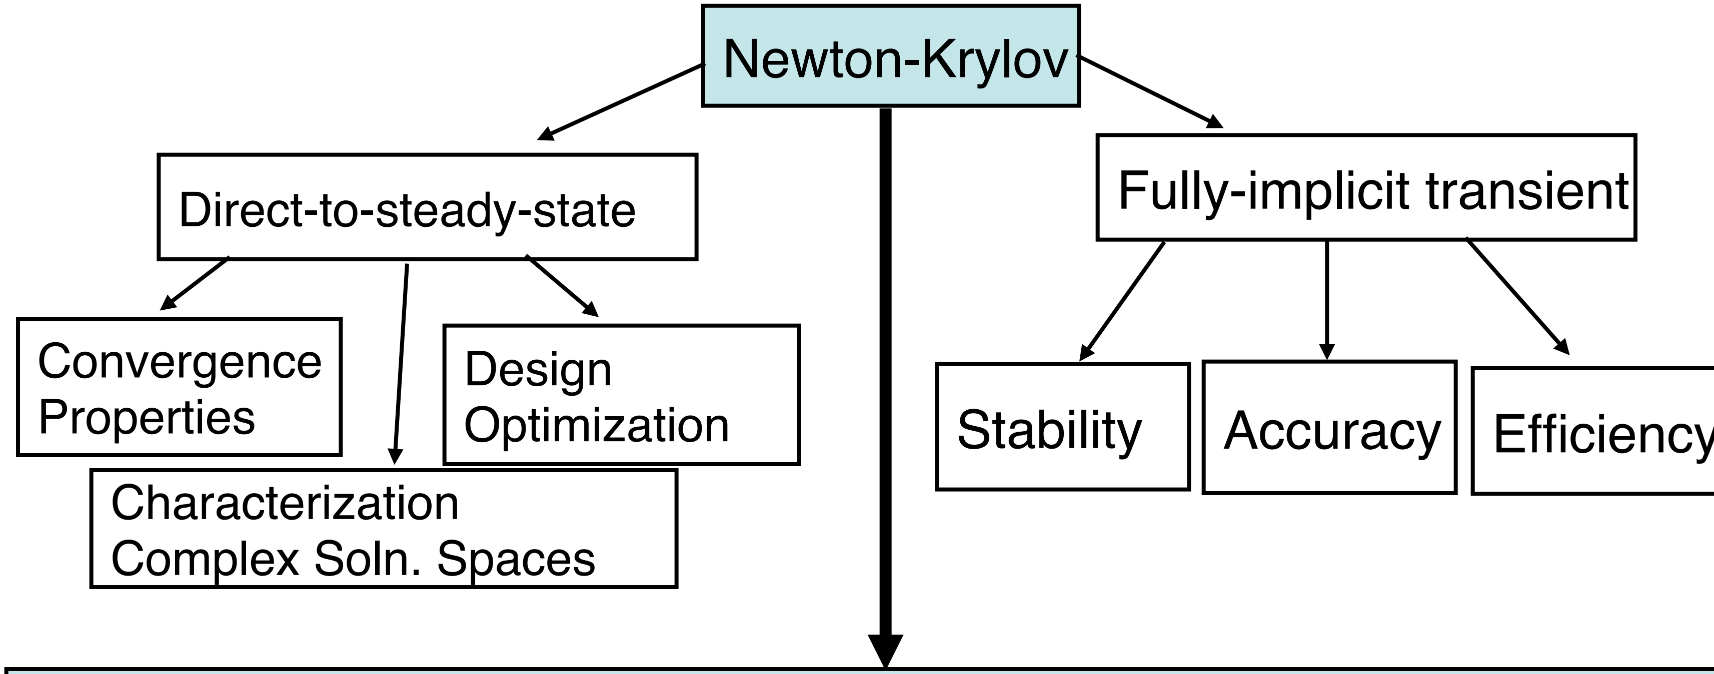

Very Large Problems -> Parallel Iterative Solution of Sub-problems

Krylov Methods - Robust, Scalable and Efficient Parallel Preconditioners

- Approximate Block Factorizations
- Physics-based Preconditioners
- Multi-level solvers for systems and scalar equations

Preliminary Weak Scaling: Kelvin-Helmholtz

Drekar (CFD/MHD): R. Pawlowski, E. Cyr, J. Shadid

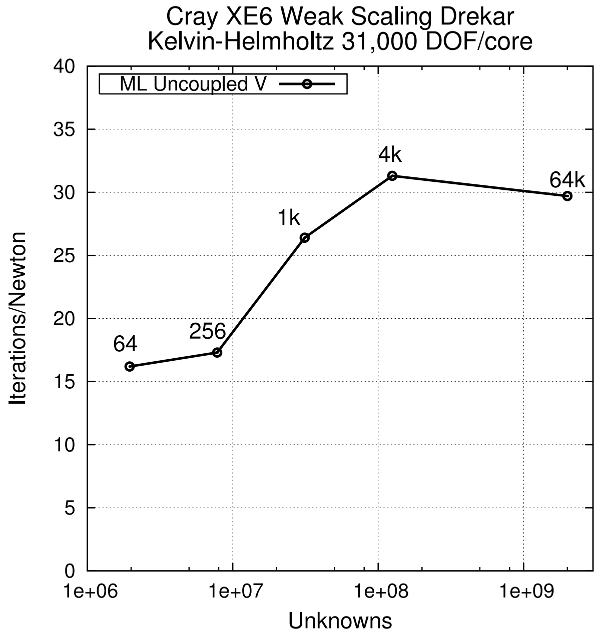

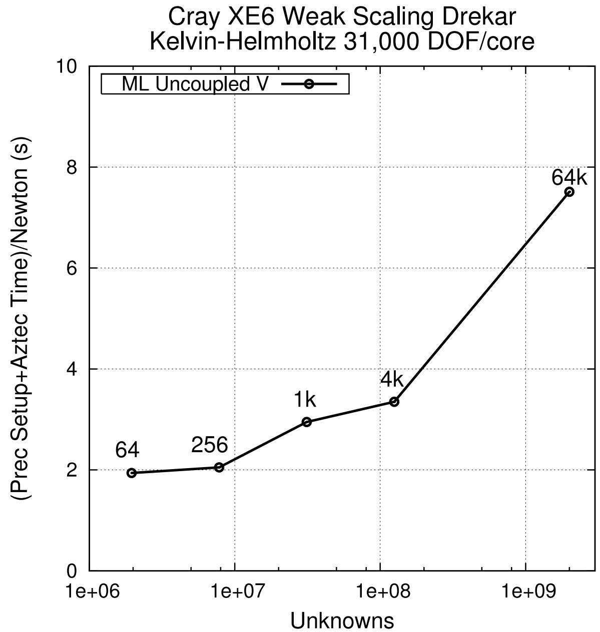

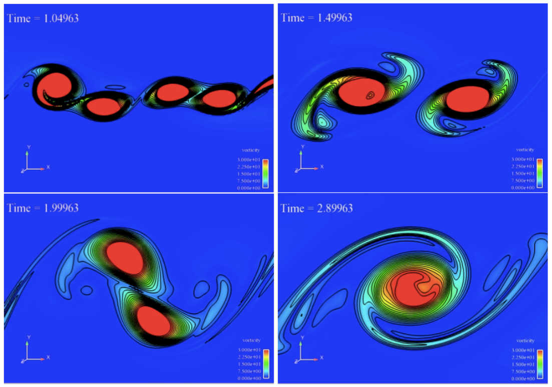

- 2D Transient Kelvin-Helmholtz instability
- Re=5000 shear layer, CFL ~2
- Cray XE6

Preliminary Weak Scaling: 3D CFD (Fixed CFL)

Drekar (CFD/MHD): R. Pawlowski, E. Cyr, J. Shadid

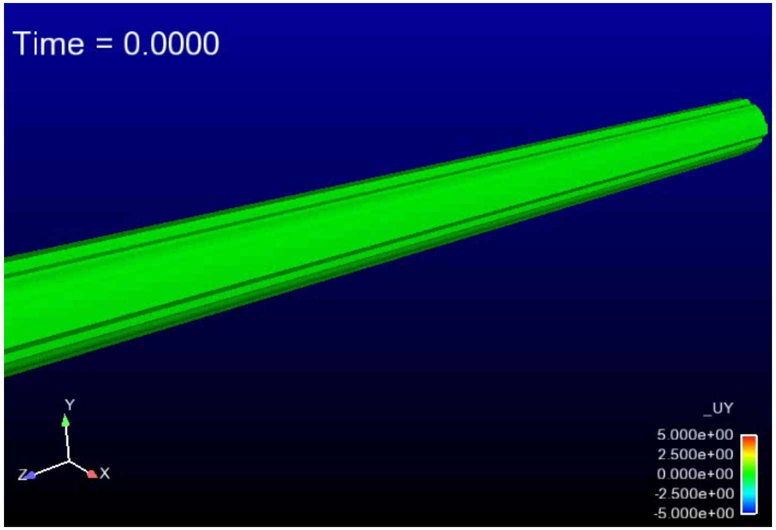

- 3D transient swirling jet RANS Model, fixed CFL (CFL ~1)
- GMRES, ML prec
- Cray XE6

cores	Newt/dt	Iter/Newt step	Iter/dt	Time/Newton step (sec)	
				Prec	Aztec
256	3.7	14	52	1.3	1.0
2048	4.0	20	80	1.8	1.6
16384	4.0	30	120	2.8	3.0
131072	3.8	34	129	5.4	3.5

Preliminary Weak Scaling: 3D CFD (Fixed dt)

Drekar (CFD/MHD): R. Pawlowski, E. Cyr, J. Shadid

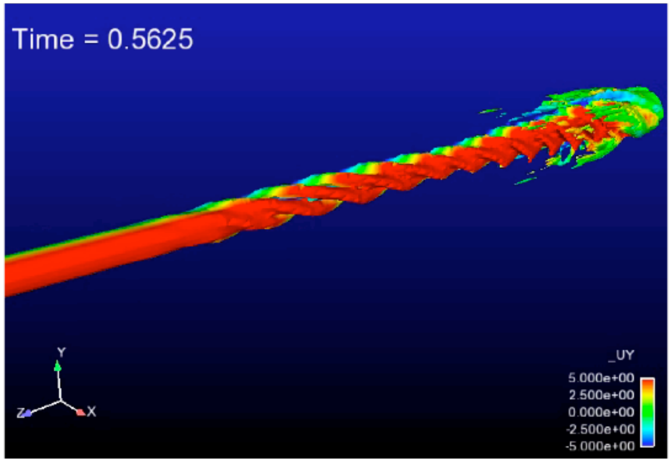

- 3D transient swirling jet, fixed dt=0.001
- GMRES, ML prec
- Cray XE6

cores	Newt/dt	Iter/Newt step	Iter / dt	Time/Newton step (sec)		TFQMR time/iter
				Prec	Aztec	
2048	2.3	22	51	1.9	1.0	0.15
16384	3.6	27	97	2.2	2.4	0.18
131072	3.8	34	129	5.4	3.5	0.19

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

Sandia National Laboratories

Need to Reduce Iteration Count

- 1-level preconditioners (e.g. additive Schwarz) do not scale due to lack of global coupling
  - For 2D semiconductor drift-diffusion FEM GMRES iteration count scales by sqrt(DOF)
- Need methods with global coupling: multilevel/multigrid

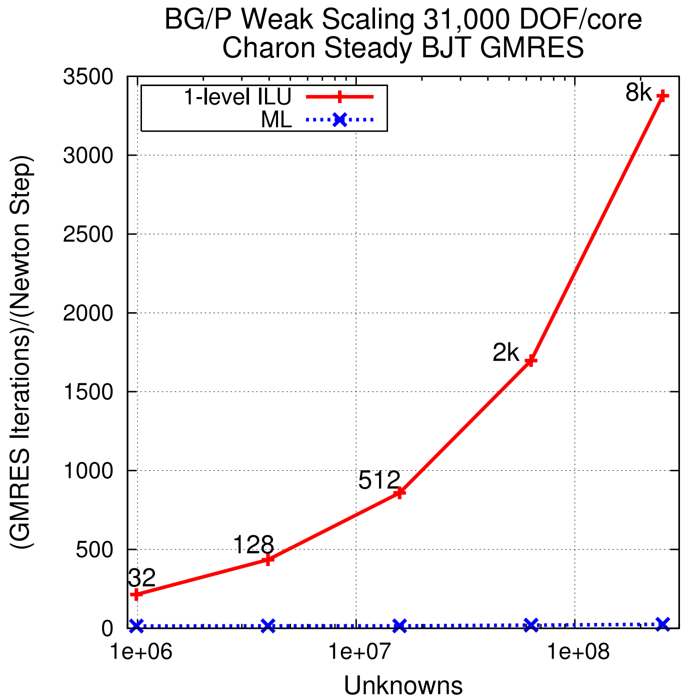

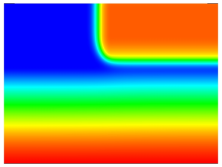

IBM Blue Gene/P

- Compared with 1-level preconditioner for 8192-core, 252 million DOF, ML reduced iterations by 182x, time by 412x

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

MHD Generator: Iterations Increase as Weak Scale

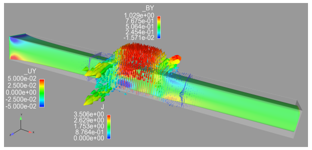

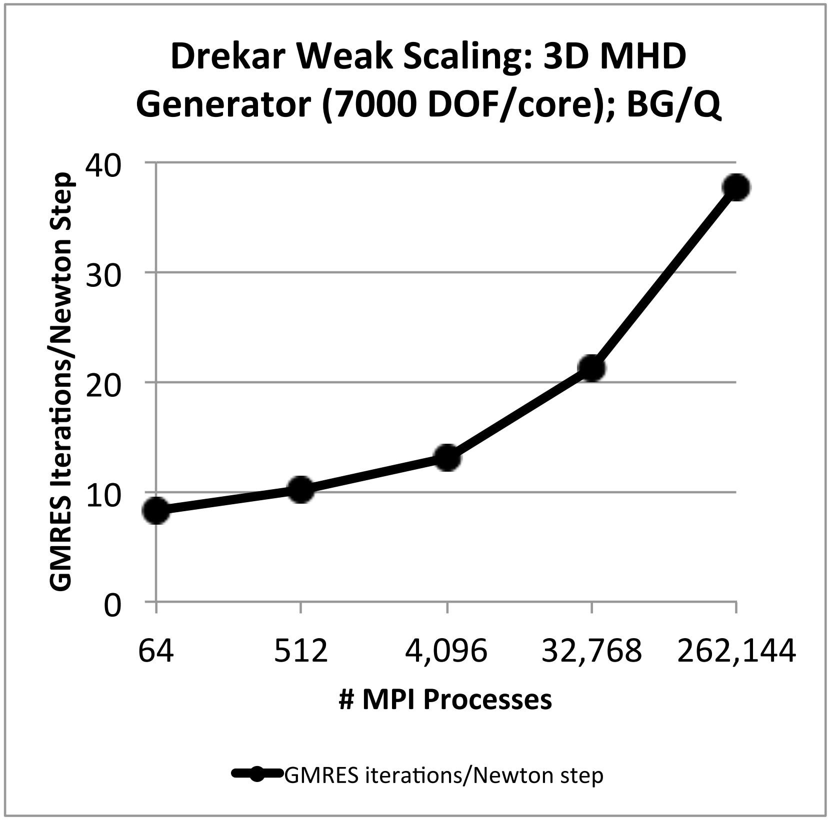

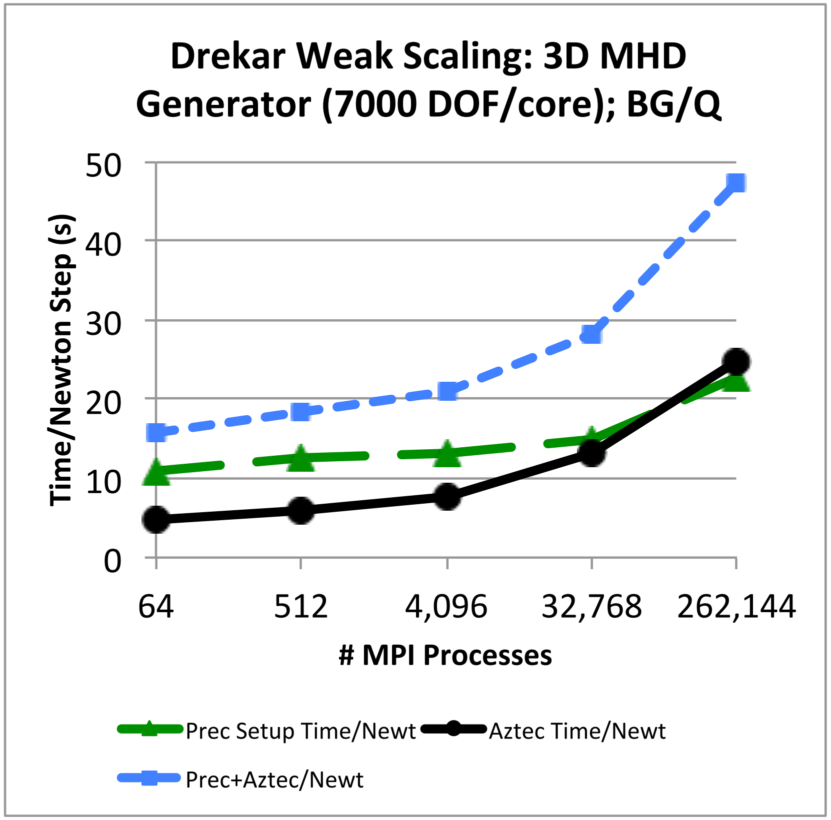

- GMRES iterations increase -> solve time rapidly increases
- Key: how to reduce iteration count? Better aggregation, smoothers

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

Sandia National Laboratories