Speeding up multi-physics simulations through reuse of multigrid components

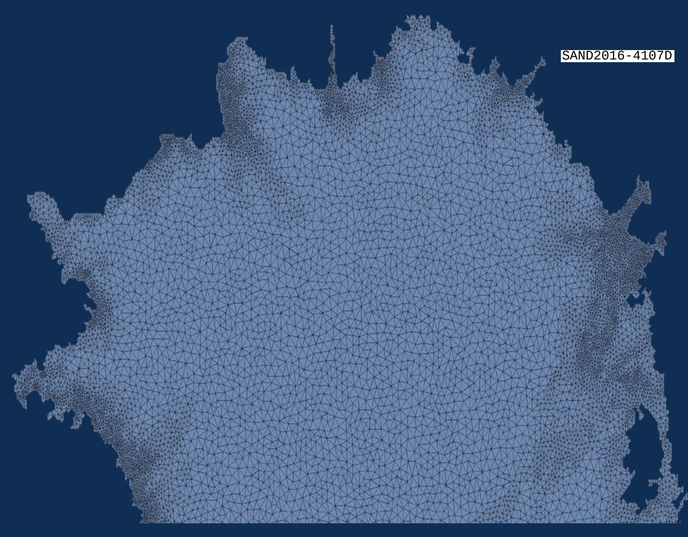

Sandia National Laboratories

A. Prokopenko[1], J. Hu[2]

[1]Sandia National Laboratories, New Mexico, 87185 [2]LSandia National Laboratories, California, 94551

Problem

Motivation: Transient simulations may become prohibitively expensive due to a large number of time steps or nonlinear solver iterations. In many cases, most of the time is spent in setting up linear solvers such as algebraic multigrid (AMG). Our research concentrates on reusing components of existing multigrid hierarchy for subsequent setups, particularly for same mesh connectivity structure.

Goal: Reduce multigrid hierarchy setup by reusing multigrid components

Funding: "FastMATH" SciDAC (DOE ASCR)

MueLu

MueLu is a multigrid solver library in the Sandia Trilinos project. MueLu is a part of the templated solver stack that is based upon Kokkos (intra-node parallelism) and Tpetra (inter-node parallelism).

- Aggregation-based algebraic multigrid algorithms (AMG) for Poisson, elasticity, convection-diffusion, Maxwell
- Design facilitates new algorithm development
- Scalar, local/global ordinals, and node template types. Permits architecture-specific algorithms & optimizations

MueLu is based on the *factory design pattern:*

- Each building block is a factory
- Each factory has input and output
- All input and output data is (temporarily) stored in a container for each multigrid level
- The hierarchical data dependencies are automatically handled by MueLu

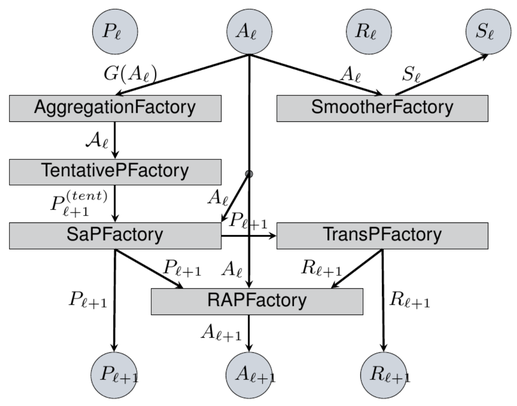

Reuse types

No reuse

The default scenario where everything is recomputed every step (transient or Newton).

Reuse of smoothers (S)

Recompute everything except level smoothers (only numerical factorization)

Reuse of tentative prolongators (TP)

Recompute smoothed prolongators (reusing matrix graphs), coarse level matrices (reusing matrix graphs) and smoothers (only numerical factorization)

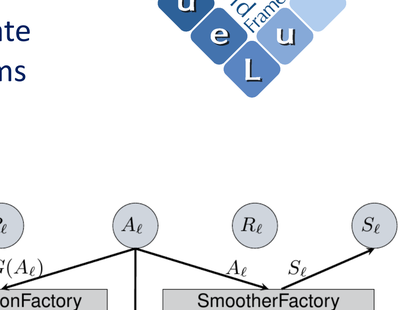

Reuse (TP)

Reuse of smoothed prolongators (RP)

Recompute coarse level matrices (reusing matrix graphs) and smoothers (only numerical factorization).

May negatively affect convergence.

Reuse of coarse levels (RAP)

Recompute only fine level smoother (only numerical factorization)

May negatively affect convergence.

Reuse of the full hierarchy

Recompute nothing.

Least likely to converge.

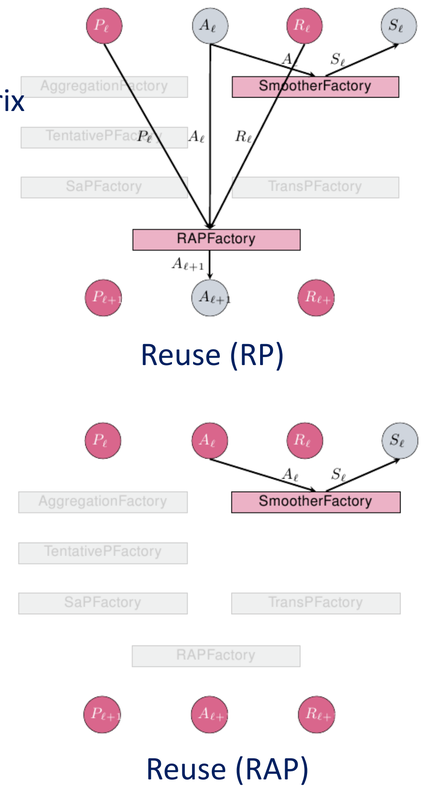

Reuse (RP)

Reuse (RAP)

Results: MHD application

- Steady-state 3D MHD generator
- Resistive MHD model
- Stabilized FE
- Single Newton-Krylov solve
- 8 DOFs/mesh node
- Monolithic preconditioner
- Unsmoothed prolongator
- Heavy smoother (DD/ILU(0))

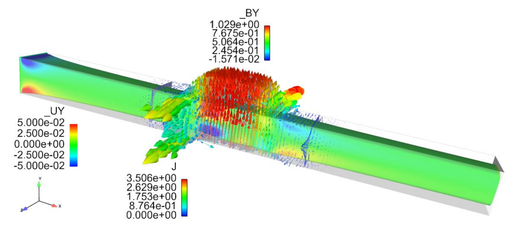

Reuse	Setup	Solve	Total
No reuse	3.79	2.65	6.44
S	3.27	2.63	5.91
RP	3.14	2.61	5.75
RAP	2.74	42.80	45.53

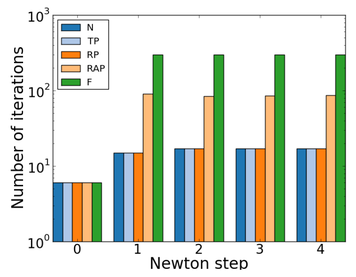

Results: CFD application

- 3D jet, Re=106, CFL=0.25
- Symmetric V(3,3) Gauss-Seidel
- Setup cost almost entirely
  - Smoothed prolongator
  - Galerkin product

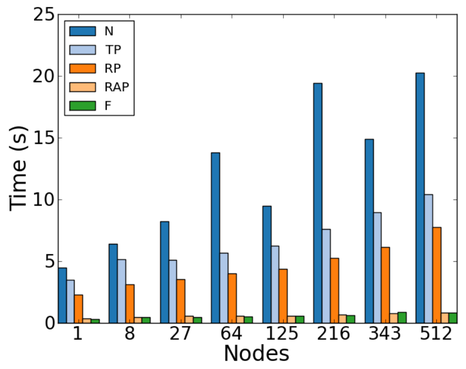

Ongoing Development

- Different reuse strategies for reuse within Newton solver and across transient steps
- Better information exchange between nonlinear and linear solvers
- Ability to enable/disable reuse for particular time steps

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.

SAND 2010-xxxP