Scalable Multilevel Preconditioners

Ray Tuminaro, Jonathan Hu, Chris Siefert, Pavel Bochev
Sandia National Laboratories*

Summary

Our goal is to enable new simulation capabilities by extending the applicability of scalable algebraic multigrid solvers. This includes basic algorithm research and direct interaction with applications. FY2007 accomplishments center on a new algebraic multigrid method for compatible discretizations of the Maxwell's equations.

We recently developed a new reformulation of the eddy current equations. This new reformulation is consistent with compatible discretization philosophy and gives an identical discrete solution to a standard edge element formulation of the eddy current equations. The advantage is that the reformulated system is amenable to multigrid methods. A corresponding algebraic multigrid method was devised for the reformulated system which leverages standard algebraic multigrid software. The new method demonstrates mesh independent convergence rates and relative insensitivity to coefficient jumps.

The eddy current equations are given below:

$$\begin{aligned} \nabla\times\frac{1}{\mu}\nabla\times E+\sigma E &= f \quad in\ \Omega \\ n\times E &= 0 \quad on\ \Gamma \\ n\times\frac{1}{\mu}\nabla\times E &= 0 \quad on\ \Gamma^* \end{aligned}$$

where μ is magnetic permeability, σ is electric conductivity, and E is the unknown electric field. An edge element compatible discretization leads to a linear system of the form

$$(\mathcal{M}_1 + \mathcal{D}_1^T \mathcal{M}_2 \mathcal{D}_1)x = b$$

where $\mathcal{D}_1$ is the curl operator, $\mathcal{M}_1$ is the edge mass matrix (using a σ weighted inner product), and $\mathcal{M}_2$ is the face element mass matrix (using a $1/\mu$ weighted inner product).

Standard algebraic multigrid methods typically fail due to the large null space of $\mathcal{D}_1$. To avoid this, we reformulate the equations to the entirely equivalent system given by

$$\begin{pmatrix} \mathcal{M}_1+\mathcal{D}_1^T\mathcal{M}_2\mathcal{D}_1+\tilde{\mathcal{M}}_1\mathcal{D}_0\mathcal{M}_0^{-1}\mathcal{D}_0^T\tilde{\mathcal{M}}_1 & \mathcal{M}_1\mathcal{D}_0 \\ \mathcal{D}_0^T\mathcal{M}_1 & \mathcal{D}_0^T\mathcal{M}_1\mathcal{D}_0 \end{pmatrix}$$

where $\mathcal{D}_0$ is the gradient operator, $\tilde{\mathcal{M}}_1$ is an edge mass matrix (using an unweighted inner product) and $\mathcal{M}_0$ is a nodal mass matrix (using a $1/\mu$ weighted inner product). The advantage of the above system is that the (2,2) block is a scalar Laplace operator and the (1,1) block is essentially a vector Laplace operator plus a mass matrix. In principle, a multigrid preconditioner can be developed for each of these blocks and then combined in a block-Jacobi fashion.

The (2,2) block is easily handled using any standard algebraic or geometric multigrid method. While the (1,1) block is essentially a vector Laplacian, it still presents two

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL85000.

difficulties that must be overcome. The first is that the new compatible gauge term involving the $\mathcal{D}_0$'s has a significantly large number of non-zeros. The second is that edge element basis functions have directionality and orientation features which must be considered within the solver.

To resolve these difficulties a special smoother and special grid transfer are used only on the finest level operator. The grid transfer is based on piecewise constant interpolation and properly addresses orientation features. In addition to coarsening, it also transforms the (1,1) block on the fine grid from an edge element basis to a coarse level nodal basis. This means that a standard algebraic multigrid method can be used for all subsequent coarser levels. The finest level smoother is actually applied to the original eddy current formulation

$$(\mathcal{M}_1 + \mathcal{D}_1^T \mathcal{M}_2 \mathcal{D}_1)x = b.$$

This is much less expensive as it avoids using the gauge term during the smoothing step and in fact avoids the need to form the gauge term on the finest level. The resulting algebraic multigrid method is inexpensive to setup and to apply within the V cycle.

The table below illustrates the number of iterations required to reduce the residual by 10 orders of magnitude for a 2 dimensional model problem corresponding to 2 materials

		σ_2				
Grid	cmplx	10^0	10^{-2}	10^{-4}	10^{-6}	10^{-8}
9^2	1.07	7	7	7	7	7
27^2	1.20	12	12	12	12	12
81^2	1.25	15	16	16	16	16
243^2	1.27	17	18	18	18	18

shown in Fig.1. The table correspond to using smoothed aggregation for the (2,2) block and the coarse levels of the (1,1) block in conjunction with 1 pre and 1 post smoothing sweep of Gauss-Seidel. 'cmplx'.

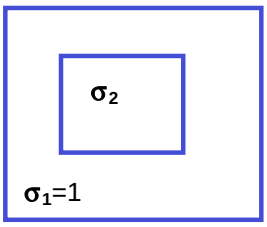

Figure 1. Problem setup for varying σ.

measures the costs of the V-cycle. It is the sum of matrix nonzeros on all levels divided by the matrix nonzeros on the finest level. The table illustrates relative insensitivity to variations in σ, and mesh independent convergence rates.

We have observed similar behavior when the algorithm is run on problems arising from Z-pinch modeling. Figure 2 illustrates a simple model of a Z-pinch, containing conductivities ranging over six orders of magnitudes. This one million unknown problem was solved with 38 conjugate gradient iterations using the newly developed multigrid preconditioner.

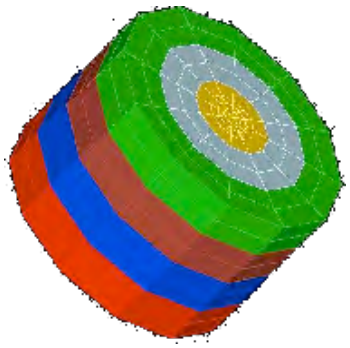

Figure 2. A protoype z-pinch geometry.

Finally, the algorithm has been run on more than 2000 processors with efficiencies of about 70%.

For further information on this subject contact:

Dr. Anil Deane, Program Manager
Mathematical, Information, and Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: 301-903-1465
e-mail: deane@mics.doe.gov