Advanced Algebraic Multilevel Preconditioned Newton-Krylov Solution of Transport/Reaction Systems

John N. Shadid, Ray S. Tuminaro, Paul T. Lin,
Roger P. Pawlowski, Gary L. Hennigan
Sandia National Laboratories
Marzio Sala
ETH Zurich

Summary

Our goal is to develop new robust, efficient and scalable solution methods for computational simulation of complex fluid flows, thermal energy and mass transfer systems with non-equilibrium chemical kinetics. Our work includes basic algorithm research, software development and numerical experimentation in applications. FY2006 accomplishments centered on the continued development and comprehensive evaluation of an aggressive-coarsening block algebraic multilevel preconditioner for systems of equations.

Transport/Reaction (TR) systems with fluid flow, thermal energy transfer and chemical species transport, with non-equilibrium chemical reactions, are common in the natural physical world and highly important for many advanced technology applications of importance to DOE. The computational simulation of these systems requires the solution of strongly-coupled interacting physics on high resolution unstructured meshes. The discretization of the governing PDE equations produces large non-linear systems of equations, for which robust and efficient parallel iterative solution methods are necessary. Our work focuses on the development of scalable algebraic multilevel domain decomposition preconditioners for systems of equations. These preconditioners can be effectively used in a robust fully-coupled Newton-Krylov solution of these large-scale systems.

An additive Schwarz domain decomposition (DD) approach partitions the original domain into overlapping sub-domains and approximately solves the discrete problem corresponding to the individual sub-domains in parallel. A known disadvantage is that this method (referred to as a one-level Schwarz preconditioner) does not scale as the number of sub-domains increases. That is, the number of iterations rises as the number of subdomains (and/or total size of the problem) increases. To remedy this situation, multigrid ideas can be employed. The general geometric multigrid philosophy is to use a sequence of meshes to damp errors at different frequencies thereby accelerating the convergence to the solution on the next mesh. In a typical multigrid method, many coarse meshes are used in conjunction with fairly lightweight smoothers. In a domain decomposition setting, only a few coarse meshes are usually employed and these meshes are significantly coarser than the finest mesh.

Briefly, our new algebraic multilevel preconditioner is based on developing a few very coarse operators that approximate the

* 925-294-2564, tuminaro@ca.sandia.gov

behavior of the fine grid operator on coarser meshes. The sequence of coarse operators are developed by using the fine mesh non-zero block structure of the Jacobian matrix as a graph. This graph essentially represents the nodal connectivity of the finite element mesh. We then employ serial and parallel graph partitioning techniques to aggressively coarsen this graph by an aggregation procedure. This process, along with an assumption about the zero-energy modes of the original continuous PDE operator, is used to define the projection operators to produce the required sequence of coarse operators.

These aggressive coarsening block oriented algebraic multilevel (acbAMG) preconditioners have been evaluated on a number of important challenging applications. These include incompressible and compressible fluid flow, fluid flow with thermal buoyancy effects, turbulent incompressible flow modeling with LES methods, incompressible flow with thermal energy and mass transfer along with chemical reactions, and recently on drift diffusion formulations for semi-conductor device modeling.

Figure 1 displays the algorithmic scaling of the acbAMG preconditioners relative to a one-level DD and a two-level geometric DD method. The results are for a 3D thermal buoyancy driven flow problem. This figure clearly indicates that the acbAMG (two- and three-level) methods achieve optimal convergence in terms of a constant iteration count as the problem size is increased by over three orders of magnitude. In contrast the one-level method scales as $\sqrt{N}$. These parallel results are for up to nearly 100 million unknowns and up to 2048 processors. Our results have demonstrated that the multilevel methods can be as much as 100x faster than the one-level methods for 2D and 10x faster in 3D for large problems. Currently we have not demonstrated strictly optimal performance in CPU time as the cost of the coarse grid solve can become important in a number of applications. Current work includes evaluation of these methods on highly convected flows and strong chemical reactions.

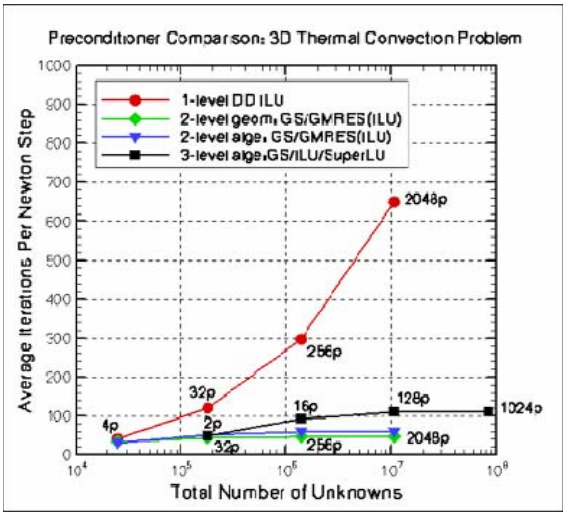

This new multilevel preconditioner is available in the ML package in Trilinos. http://software.sandia.gov/trilinos

[1] *Sala, Shadid & Tuminaro.* ***An Improved Convergence Bound for Aggregation-Based Domain Decomposition Preconditioners****,* SIMAX, Vol 27, 744-756, 2005.

[2] *Shadid, Salinger, Pawlowski, Lin, Hennigan, Tuminaro & Lehoucq.* ***Large-Scale Stabilized FE Computational Analysis of Nonlinear Steady State Transport/Reaction Systems,*** *CMAME, Vol. 195, 1846-1871, 2006.*

[3] *Lin, Sala, Shadid, Tuminaro.* ***Performance of Fully-Coupled Algebraic Multilevel Decomposition Preconditioners for Incompressible Flow and Transport,*** *IJNME, Vol 67, 2, 208-225*, 2006..

For further information on this subject contact:
Dr. Anil Deane, Program Manager
Mathematical, Information, and Computational Sciences Division
Office of Advanced Scientific Computing Research
deane@mics.doe.gov