Algebraic Multigrid (AMG) for Ground Water Flow and Oil Reservoir Simulation

Klaus Stüben[1], Patrick Delaney[2], Serguei Chmakov[3]
[1]*Fraunhofer Institute SCAI, Klaus.Stueben@scai.fhg.de, St. Augustin, Germany*
[2]*Waterloo Hydrogeologic Inc., PDelaney@flowpath.com, Waterloo, Ontario, Canada*
[3]*Waterloo Hydrogeologic Inc., Schmakov@flowpath.com, Waterloo, Ontario, Canada*

ABSTRACT

A serious bottleneck in performing large-scale numerical simulations is the speed with which the underlying sparse systems of equations can be solved. If these systems exceed a certain size, they can no longer be solved efficiently with standard numerical solvers simply because these solvers are not scalable. In this paper, we introduce a scalable state-of-the-art solver package based on algebraic multigrid, SAMG, which has been developed by the Fraunhofer Institute for Algorithms and Scientific Computing (FhG-SCAI).

INTRODUCTION

Computer simulations are becoming increasingly important. Generally, the finer the resolution of a discretization grid, the higher the accuracy of the numerical simulation. Unfortunately, however, increasing the grid resolution also increases the size of the corresponding (sparse) matrix equations which have to be solved numerically. Problems with millions of degrees of freedom (variables) are being tackled nowadays. Such large matrix equations can no longer be solved efficiently by standard off-the-shelf numerical methods such as conjugate gradient combined with some classical preconditioner. Instead, *numerically scalable* solvers are needed, that is, solvers for which the computational work depends only linearly on the number of variables. Scalability, in turn, necessarily requires *hierarchical* (multi-level) approaches which ensure a rapid reduction of both short and long range error components. Corresponding solvers are much more complicated than standard solvers. However, depending on the actual grid size and the concrete application itself, the resulting computational gain may be enormous.

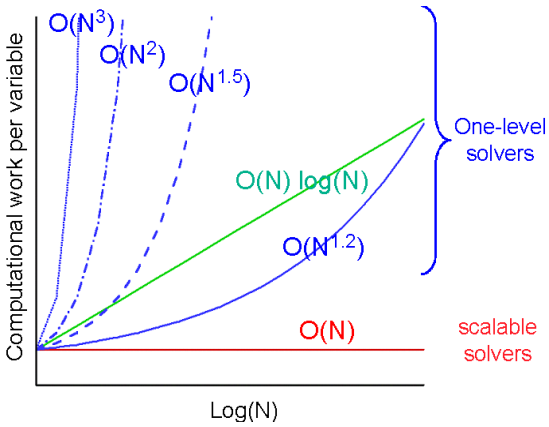

Figure 1. Computational work as a function of the number of variables, N

The most important progress in the numerical solution of discretized partial differential equations (PDEs) during the past three decades, was due to the multigrid principle. Any method based on this principle operates not just on the given discretization grid but rather on a hierarchy of grids, defined a priori by coarsening the given grid in a geometrically natural way. Unfortunately, integrating such a "geometric" multigrid method into an existing software package is very difficult. At the very least, it requires re-writing most of the code. Moreover, practically relevant grid models are often so complex that the explicit construction of a "natural" hierarchy of grids is very complicated if at all possible. In practice, this often limits the usefulness of geometric multigrid.

Consequently, there is an increasing demand for algebraically oriented "plug-in" solvers which are still scalable but do not explicitly exploit geometrical properties of the given problem. Over 15 years ago, the first research code based on the idea of *algebraic multigrid*, became available [AMG1R5, Ruge&Stüben 1986]. Its main purpose was to demonstrate that numerical scalability can, in principle, be achieved without exploiting geometric properties, at least for certain classes of scalar elliptic PDEs. Recently, AMG1R5 is being used for increasingly complex large-scale applications. In particular, it has successfully been used to

demonstrate the strong potential of algebraic multigrid in speeding up ground water flow simulations (see [Mehl&Hill 2001]). However, these applications have also revealed some serious drawbacks of the original algorithm, in particular its unacceptably high memory requirements. During the last years, a much more advanced AMG algorithm has been developed, realized in the software package SAMG, which overcomes most of the drawbacks of AMG1R5, see [Stüben&Clees 2003, Stüben 2001]. The purpose of this paper is to outline the method and present some typical results.

ALGEBRAIC MULTIGRID

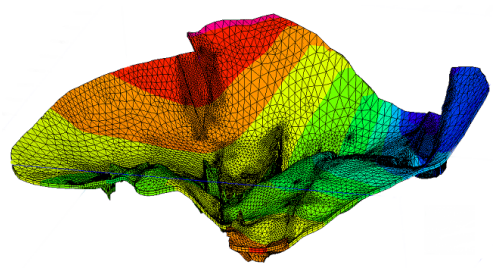

Figure 2. Modelling a reservoir by Finite Elements (courtesy of Wasy GmbH)

The idea of AMG is to generalize the two main principles of geometric multigrid - error smoothing by relaxation and coarse-grid correction - directly to certain classes of sparse matrix equations. This makes a corresponding solver particularly attractive as a black-box in connection with existing simulation software since only the matrix problem, $Au=f$, has to be provided. Ideally, no further information about the origin of the problem such as discretization, shape of the domain or the type of the discretization grid needs to be known. It does not matter whether a problem is 2D or 3D, whether the discretization grid is structured or unstructured and whether coefficients are smooth or discontinuous.

Rather than operating on a hierarchy of *grids*, AMG operates on a hierarchy of *increasingly smaller matrix equations*, constructed fully automatically, based on *algebraic information* (explicitly or implicitly contained in the discretization matrix). However, as for geometric multigrid, AMG is not a fixed method but rather provides a methodology. The details of how the individual algorithmical components are finally constructed, strongly influence the resulting efficiency in terms of speed of convergence, memory requirement as well as robustness.

The previously mentioned SAMG is actually not just a single code to solve particular matrix equations but rather provides a complete algebraic multi-level framework. For many years, its development has been driven by industrial requirements and its range of applicability is continuously being extended. Currently SAMG is being used in such diverse application fields as general purpose CFD, multiphase flow in porous media (both ground water modeling and oil reservoir simulation), structural mechanics, semiconductor process- and device-simulation, circuit simulation, casting and molding. Compared to its academic forerunner AMG1R5, SAMG is much simpler to use. It is considerably more robust and efficient, and – most importantly – it has an efficient memory management and new algorithmical components which allow to reduce memory requirements to a minimum.

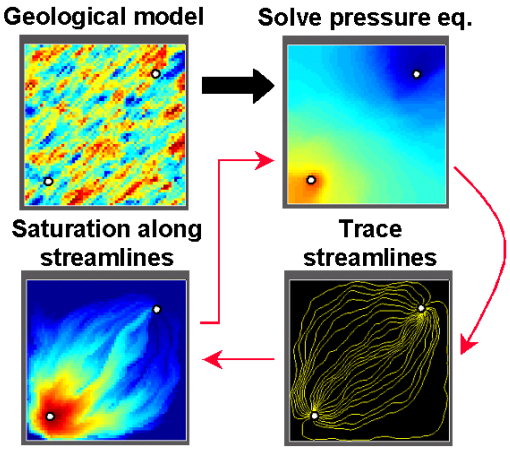

Figure 3. Streamline approach in oil reservoir simulation

Some of the above applications lead to highly complex coupled systems of PDEs for which the search for optimal AMG components still requires some basic research. For other cases, SAMG is very mature. For instance, SAMG is highly efficient in solving matrix problems of the type which occur if CFD applications are tackled by solving sequences of *scalar* PDEs instead of directly addressing the fully coupled system. Popular approaches are the *segregated solution approach* to solve the Navier-Stokes equations or the *streamline approach* in oil reservoir simulation. Typically, the resulting problems to be solved by SAMG then correspond to diffusion equations for the pressure or related quantities (such as the head in ground water modeling), or diffusion-convection-reaction equations. Simulation software such as MODFLOW requires highly efficient software modules for the numerical solution of such types of equations.

BENCHMARKS FOR SELECTED MATRIX PROBLEMS

For a demonstration of SAMG's efficiency, we compare the elapsed computational time of SAMG with that of a standard one-level solver (ILU-preconditioned CG or BI-CGSTAB, depending on whether or not the underlying problem is symmetric). We consider the solution of eight individual matrix problems, each one taken from a full simulation run in *ground water modelling*, *oil reservoir simulation* or *external flow simulation*. The following cases are considered:

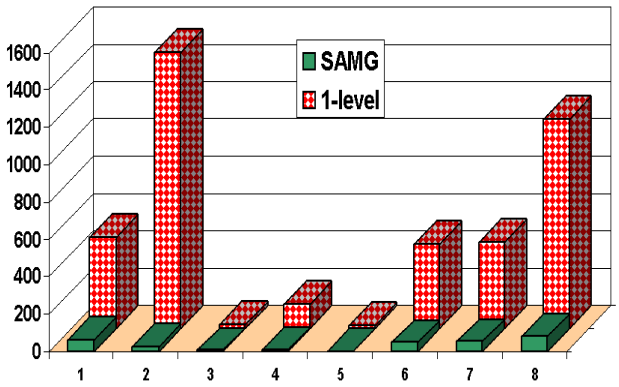

Figure 4. CPU times: SAMG versus standard one-level solver (residual reduction by 9 digits)

Cases 1-5: MODFLOW simulations as described in Table 1 of [Mehl&Hill 2001]. The size of these cases varies between 75,000 and 1,728,000 variables.
Cases 6-7: Streamline-based oil reservoir simulations by the commercial codes 3DSL (StreamSim Technologies) and FrontSim (GeoQuest), respectively. Both cases consist of about 1,200,000 grid cells; the first one is symmetric, the second one non-symmetric.
Case 8: Solution of a single pressure equation at one particular time step of a full CFD simulation of the flow around a car by the commercial code STAR-CD from Computational Dynamics (2.3 million cells, see [Stüben&Clees 2003]).

For problems of the type and size considered here, the observed speedup through the use of SAMG is up to two orders of magnitude. Note that, generally, the concrete speedup depends on various details of the application such as the size and regularity of the grid, the type of the underlying discretization, the discontinuities and anisotropies of the coefficients, etc. In any case, however, due to the scalability properties of SAMG, the computational gain will grow further with increasing grid size.

The following figures demonstrate the typical convergence behavior of SAMG. The rapid and uniform convergence behavior of SAMG is directly related to its hierarchical way of operation, a consequence of which is that AMG globally reduces errors much more effectively than any one-level method.

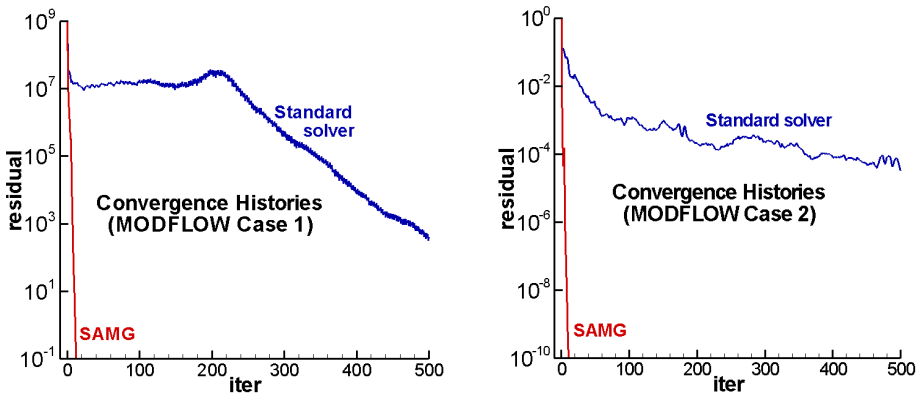

Figure 5. Comparison of convergence histories: SAMG versus standard one-level solver

As already mentioned, one of the major drawbacks of AMG1R5 is its memory requirement which, depending on the application, may easily be one order of magnitude higher than that of a standard one-level solver. Memory requirement typically becomes unacceptable in 3D situations using unstructured meshes and/or if coefficients vary in a strongly anisotropic or discontinuous way (see [Mehl&Hill 2001]). SAMG, on the other hand, has been designed to require, on the average, not more than twice the memory occupied by the given matrix *A*. Although a reduction of memory requirements in algebraic multigrid typically is at the expense of a

slower convergence, some precautions have been taken in SAMG which make it even (often considerably) faster than AMG1R5 in terms of elapsed computational time. This has been achieved by combining efficient techniques based on aggressive coarsening and special interpolation operators (see [Stüben, 2001]). To demonstrate SAMG's advantages in terms of memory requirement, we consider the *operator complexity*,

$$c_A = \sum_{\ell=1}^{L} |A^{(\ell)}| / |A^{(1)}| > 1.0 \ ,$$

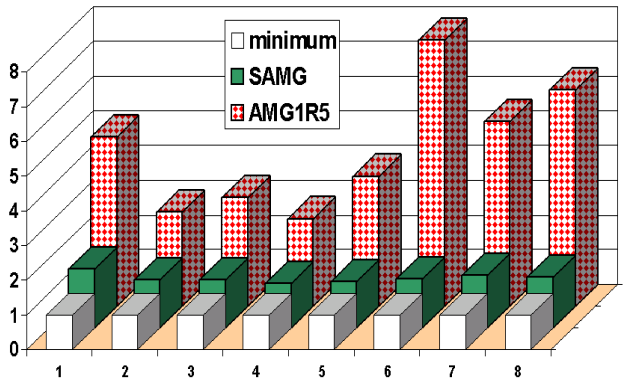

Figure 6. Comparison of operator complexity: SAMG versus AMG1R5

where $|A^{(\ell)}|$ denotes the number of non-zero entries contained in the ℓ-th level matrix (the 1st level being the finest one). That is, c_A defines the memory overhead caused by the need to store a sequence of coarser matrices in addition to just the given matrix. Although the true memory requirement is still somewhat higher (additional workspace is needed, for instance, to store the interpolation matrices), c_A is closely related to the overall memory requirement. Figure 6 compares the operator complexity of SAMG with that of AMG1R5 for the same eight cases as before. While AMG1R5's complexity values vary between 2.5 and 8, the average value for SAMG is only c_A~1.5 or even less.

APPLICATIONS IN GROUND WATER MODELING

In the course of research and development, many new technologies have effectively demonstrated remarkable improvements over established technologies in controlled testing environments, only to fall well short of expectations during implementation in 'real-world' scenarios. The purpose of this study was to test the SAMG Solver on several real-world groundwater flow models and compare the performance and results against those obtained using the PCG Solver, developed by the USGS and provided with MODFLOW-2000 [Hill, 1997], and the WHS Solver provided with Visual MODFLOW [Waterloo Hydrogeologic, 2003].

A total of six groundwater flow models were selected to run the comparisons for this study. All of these models are 'project models' previously prepared by Waterloo Hydrogeologic's Consulting Division using Visual MODFLOW, and were run using the USGS MODFLOW-2000 program.

- Model 1 contains 44908 grid cells with dimensions of 109 rows, 103 columns and 4 layers.
- Model 2 contains 110166 grid cells with dimensions of 122 rows, 129 columns and 7 layers.
- Model 3 contains 210574 grid cells with dimensions of 178 rows, 169 columns and 7 layers.
- Model 4 contains 138276 grid cells with dimensions of 167 rows, 138 columns and 6 layers.
- Model 5 contains 85680 grid cells with dimensions of 120 rows, 102 columns and 7 layers.
- Model 6 contains 59616 grid cells with dimensions of 138 rows, 144 columns and 3 layers.

Model 1, Model 2, Model 3 and Model 4 are regional scale groundwater flow models prepared for the Province of Ontario Groundwater Studies program. Model 5 is a regional wellhead protection model originally developed by the USGS and later extended and refined by Waterloo Hydrogeologic. Model 6 is regional scale model of a proposed mine site in Chile for a water supply feasibility study.

For the purposes of this study, Waterloo Hydrogeologic implemented support for the SAMG Solver in Visual MODFLOW. All of the models for this study were run using Visual MODFLOW v.3.2 (Beta) on a Pentium III 500MHz processor with 256 Mb RAM. The performance of the solvers was measured by comparing the time required for the solution to converge to a reasonable solution. Since the SAMG Solver uses a different type of closure criteria than the PCG Solver and the WHS Solver, it was decided to use a mass balance

error of 0.5% as the criteria to indicate a 'reasonable solution'. In order to achieve this, the closure criteria for each solver was adjusted on a trial and error basis for each model, so as to achieve the fastest convergence while maintaining a mass balance error of less than 0.5%. A visual inspection of the head distribution contour map was also performed to ensure the reasonableness of the solution.

A comparison of the solution times of each solver for the six models is presented in the following figure. For all but one of the models tested, the SAMG Solver demonstrated a significant performance advantage over both the PCG Solver and the WHS Solver. In the cases where the solvers all converged to a solution, the SAMG Solver was faster than the WHS Solver by a factor of between 3.7 and 9.7, and faster than the PCG Solver by a factor of between 2.4 and 11.3.

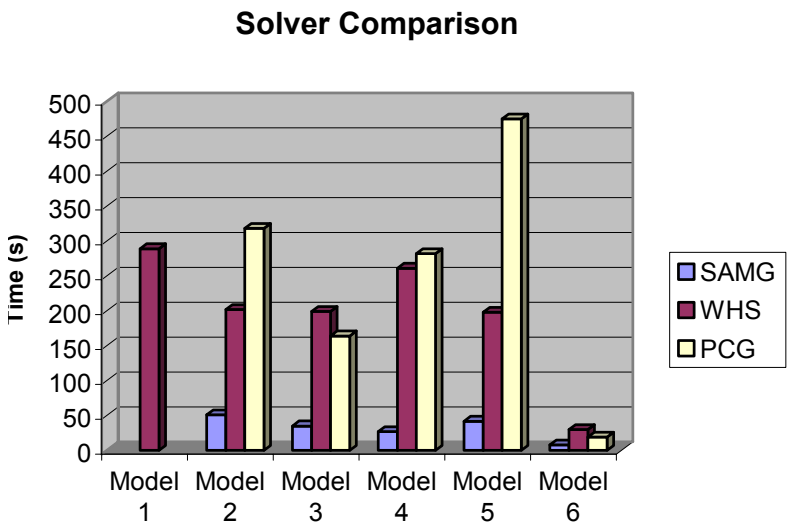

Figure 7. Comparison of Solution Convergence Times for the SAMG, WHS and PCG Solvers.

In the case of Model 1, we were not able to get the model to converge to a reasonable solution in a reasonable amount of time using the PCG Solver. Although the SAMG Solver did converge to a solution with a mass balance less than 0.5%, a visual inspection of the solution indicated it was not a reasonable solution to model. Due to time constraints we were not able to further investigate the source of the solution instability.

Although the models involved in this study contained a wide range of grid dimensioning, there was not a noticeable correlation between the model grid size (total number of grid cells) and the performance improvements realized by the SAMG solver. However, this can likely be attributed to the fact that the scaling benefits of the SAMG Solver are not fully realized until you are dealing with models containing one million or more grid cells.

With modelers continuing to build larger models with finer grid discretization, the time and cost savings of using the SAMG Solver will allow more time for creating better model conceptualization, improving the calibration, performing sensitivity analysis, and interpreting the results. This results in the development of better, more reliable models.

REFERENCES

Hill, M.C., 1990. Preconditioned Conjugate Gradient 2 (PCG2), A Computer Program for Solving Groundwater Flow Equations, Water Resources Investigation Report 90-4048, U.S. Geological Survey.

Mehl, S.W., Hill, M.C., 2001. MODFLOW-2000, the U.S. Geological Survey Modular Ground-Water Model – User Guide to the Link-AMG (LMG) Package for Solving Matrix Equations Using an Algebraic Multigrid Solver, U.S. Geological Survey, Denver (CO).

Ruge, J.W., Stüben, K., 1986. Algebraic Multigrid (AMG), in "Multigrid Methods" (S. McCormick, ed.), Frontiers in Applied Mathematics, Vol 5, SIAM, Philadelphia.

Stüben, K., 2001. A Review of Algebraic Multigrid, Journal of Computational and Applied Mathematics (JCAM) 128, pp. 281-309.

Stüben, K., Clees, T., 2003. SAMG User's Manual, Release 21c. Available through samg@scai.fhg.de.

Waterloo Hydrogeologic, Inc., 2003. Visual MODFLOW User's Manual, Waterloo Hydrogeologic, Inc.