Numer. Math. Theor. Meth. Appl.
doi: 10.4208/nmtma.2015.w06si

Vol. **8**, No. 1, pp. 112-135
February 2015

Bootstrap Algebraic Multigrid: Status Report, Open Problems, and Outlook

Achi Brandt[1], James Brannick[2,*], Karsten Kahl[3] and Ira Livshits[4]

[1] *Department of Applied Mathematics & Computer Science, The Weizmann Institute of Science, 76100 Rehovot, Israel.*
[2] *Department of Mathematics, The Pennsylvania State University, University Park, PA 16802, USA.*
[3] *Department of Mathematics, University of Wuppertal, Wuppertal, Germany.*
[4] *Department of Mathematical Sciences, Ball State University, Muncie, IN 47306, USA.*

Received 2 November 2013; Accepted 19 June 2014

Abstract. This paper provides an overview of the main ideas driving the bootstrap algebraic multigrid methodology, including compatible relaxation and algebraic distances for defining effective coarsening strategies, the least squares method for computing accurate prolongation operators and the bootstrap cycles for computing the test vectors that are used in the least squares process. We review some recent research in the development, analysis and application of bootstrap algebraic multigrid and point to open problems in these areas. Results from our previous research as well as some new results for some model diffusion problems with highly oscillatory diffusion coefficient are presented to illustrate the basic components of the BAMG algorithm.

AMS subject classifications: 65F10, 65N22, 65N55

Key words: Bootstrap algebraic multigrid, compatible relaxation, algebraic distance, least squares interpolation

1. Introduction

Numerous scientific and engineering problems find their formulation in terms of (systems) of partial differential equations, which in turn require the solution of large-scale finite element, finite difference, or finite volume equations. Modern applications involve large-scale parallel processing of (linear) systems with millions or even billions of unknowns, for which multigrid (MG) methods often provide solvers that are optimal with respect to their computational complexity and, hence, their parallel scalability.

*Corresponding author. *Email addresses:* achi.brandt@weizmann.ac.il (A. Brandt), brannick@psu.edu (J. Brannick), kkahl@math.uni-wuppertal.de (K. Kahl), ilivshits@bsu.edu (I. Livshits)

Multigrid solvers for sparse systems of linear equations $Ax = b$ are all based on two complimentary processes: a local relaxation scheme (smoother) that eliminates certain components of the error by working on the fine level and a coarse-level correction that treats the remaining global error.

Generally, the design and analysis of these two MG components are based on the following smoothing property of relaxation. For any $0 < \rho < 1$, an error vector e is called ρ-smooth if all its *normalized* residuals[†] are smaller than $\rho\|e\|$. The basic observation [2] is that the convergence of a proper relaxation process[‡] slows down only when the current error is ρ-smooth with $\rho \ll 1$, the smaller the ρ the slower the convergence. This implies that when relaxation slows down, the error vector e can be approximated in a much lower-dimensional subspace. Very efficient *"geometric multigrid"* solvers have been developed for the case that the lower-dimensional subspace corresponds to functions on a well-structured grid (the *coarse* level), on which the smooth errors can be approximated by easy-to-derive equations, based for example on discretizing the same continuum operator that has given rise to the fine-level equations $Ax = b$ to define the coarse-level operator A_c. The coarse-level equations are solved using recursively a similar combination of relaxation sweeps and still-coarser-level approximations to the resulting smooth errors.

The basic two-grid method for solving $Ax = b$, from which a multigrid method is defined by recursion, involves a stationary linear iterative method applied to the fine-grid system, and a coarse-grid correction; Given an approximation $w \in \mathbb{C}^n$, compute an update $v \in \mathbb{C}^n$ by

1. Pre-smoothing: $y = w + M(f - Aw)$,
2. Correction: $v = y + PA_c^{-1}R(f - Ay)$.

Here, M is the approximate inverse of A that defines the multigrid smoother and $R : \mathbb{C}^n \mapsto \mathbb{C}^{n_c}$ and $P : \mathbb{C}^{n_c} \mapsto \mathbb{C}^n$ with $n_c < n$ are the restriction and interpolation operators that map information between the coarse grid of size n_c and the fine grid of size n.

To deal with more general situations, where the fine-level system may not be defined on a well-structured grid nor perhaps arise from any continuum problem, *"algebraic multigrid"* (AMG) methods were developed to derive the set of coarse-level *variables* and coarse-level *equations* directly from the given matrix A. The basic approach (developed in [9,10,34] and called today "classical AMG" or RS-AMG) involves the following two steps:

(1) Choosing the coarse-grid variables, e.g., as a subset C of the set of fine-grid variables, Ω, such that each variable in $F = \Omega \setminus C$ is "strongly connected" to variables in C.

[†]The vector of *normalized residuals* is $\tilde{A}e$, where $\tilde{A}$, the *normalized matrix*, is scaled so that the l_1 norm of each of its rows is 1. Correspondingly, we define the *normalized eigenvalues* of A to be the eigenvalues of $\tilde{A}$. The magnitude of any normalized eigenvalue is at most 1.

[‡]Kaczmarz relaxation can always be used (but better schemes are often available), supplemented when needed by local relaxation steps around exceptionally large residuals (as suggested in [1, §A.9]).