Algebraic Multigrid Theory: The Symmetric Case

Achi Brandt

Department of Applied Mathematics
The Weizmann Institute of Science
Rehovot 76100, Israel

Abstract. A rigorous two-level theory is developed for general symmetric matrices (and non-symmetric ones using Kaczmarz relaxation), without assuming any regularity, not even any grid structure of the unknowns. The theory applies to algebraic multigrid (AMG) processes, as well as to the usual (geometric) multigrid. It yields very realistic estimates and precise answers to basic algorithmic questions, such as: In what algebraic sense Gauss-Seidel (or Jacobi, Kaczmarz, line Gauss-Seidel, etc.) relaxation smoothes the error? When to use block relaxation? What algebraic relations must be satisfied by the coarse-to-fine interpolations? What is the algorithmic role of the geometric origin of the problem? The theory helps to rigorize local mode analyses and locally analyze cases where the latter is inapplicable.

The research reported herein is sponsored by the Air Force Weight Aeronautical Laboratories, Air Force Systems Command, United Air Force, under Grant AFOSR 84-0070. A preliminary version has appeared in April, 1983 [7].

Content

1. Introduction

The purpose of this work is to develop the theoretical background needed for *algebraic multigrid (AMG)* solvers (see [6, §13.1], [10], [11], [20] and [21]), and at the same time to develop precise *algebraic theory* also for usual (including geometric) multigrid solvers. The theory is more precisely quantitative than other approaches, excluding local mode analyses [8]. Its insights are particularly important for cases where local mode analysis cannot be applied, such as AMG or geometric multigrid for problems with everywhere strongly discontinuous coefficients, as in case of Lagrangian discretizations. The inability of mode analysis to give full insight to the multi-level solutions of such problems, and the pressing need for solving many such problems with fully efficient algorithms, that can only be guided by precise quantitative insights (with convergence estimates realistically dependent on the various algorithmic parameters), were indeed the main motivations for this work. Also, even when mode analysis is applicable, it usually needs to be supplemented by the algebraic theory at some particular regions (e.g., near structural singularities; see [8]). More generally, the present theory explores the most general extension of the multi-level approach (not only for symmetric matrices; see in particular Theorem 3.4 and the discussion that follows it, and also [6, §1.1]).

The term "algebraic" multigrid, as against "geometric" multigrid, refers mainly to the mode of constructing coarser levels (coarser grids) in multi-level solvers. In the usual, *geometric* approach, each coarser grid is a simple, well-ordered coarsening of the next finer grid; e.g., the lines of the former are obtained by taking every other line of the latter, whether these are lines of gridpoints or lines of grid cells (when discretization is made in terms of cells). Sometimes only lines in some directions are so coarsened, while in other directions all fine-grid lines are retained in the coarser level [6, §4.2.1]. At any rate, the coarse-level variables, and their relation to the fine-grid variables (through interpolation), are characterized by their geometrical location.

In the *algebraic* approach no organized grids are assumed. The coarse-level variables are just a set of variables selected so as to satisfy certain criteria based on the fine-level equations. The most basic criterion is typically that each fine-level variable should be "strongly connected" to some coarse-level variables (see Secs. 3.5, 4.3). The coarse-to-fine interpolations and the fine-to-coarse transfers are also mainly based on the algebraic connections (although "hybrid" schemes are envisioned in which some geometrical information is taken into account; see Sec. 5.2); in this respect AMG continues the line of development set in [1] and [13].

Each of these approaches has its advantages and domain of applicability. The geometric multigrid in many cases is more efficient in terms of storage and time,

because it need not store the (sparse) matrix entries (both on fine and on coarser grid) and it need not assemble the coarse grid equations. But in other cases this storage and work are still invested even though the geometric algorithm is used. For such cases and some others, AMG (or the mentioned hybrid schemes) has several advantages:

First, AMG can be used as a black-box solver. Since the algorithm is based on the given matrix, assuming only some general properties, it can be used for a wide range of problems, with no need to give special attention to boundary conditions, anisotropies, strong discontinuities in coefficients, etc., features which normally require advanced expertise. Moreover, since the AMG selection of the coarse grid is based on strong algebraic connections, extremely anisotropic equations, equations with strong discontinuities, quasi-elliptic equations (cf. [9, §3.4], [12], [14]) and many other cases where full h-ellipticity is irregularly lost, are all solved by the same AMG algorithm, employing simple (pointwise) Gauss-Seidel relaxation. This is particularly important for problems where several of these difficulties appear together. In some pathological cases (see [A1, §8]) even 5-point diffusion equations cannot be solved with the usual, geometric coarsening.

An obvious use of AMG is for problems where no organized grids are employed and hence the geometric multigrid is not even applicable. This includes finite elements discretization with arbitrary, irregular fine-grid triangulations, and large matrix equations which are not at all derived from continuous problems, such as the geodetic adjustment problem [18], multivariate interpolation of scattered data [17], large scale problems in economics (cf. references in [10]), circuitry queuing problems, network optimizations, image reconstruction equations, etc.

There may well be large-scale problems which cannot efficiently be solved by multi-level techniques. One of the purposes of AMG Research in general, and the present study in particular is to define classes of problems for which AMG makes sense. In particular we study the role of the *geometrical origin* of the problem (see Sec. 5.5).

Despite the irregularity of the treated problems, the present theory is very concrete and realistic in its estimates, avoiding general undetermined constants, and almost precise in studying the dependence of convergence rates on various factors. Concrete estimates are obtained for successive displacement relaxation schemes, such as Gauss-Seidel or Kaczmarz with arbitrary ordering, for simultaneous displacement schemes, such as Jacobi, and for block relaxation. This is obtained by introducing new error norms (or square norms, such as R and E in Sec. 3.1) which experience precise interplay under the corresponding relaxation schemes. Through examples it is shown that the obtained relations are essentially the best possible, including sometimes the best possible constants.

For purely algebraic multigrid algorithms the performance of two-level

schemes is derived by showing the "relaxability" (efficient reducibility by relaxation) of the error produced by coarse-grid corrections (Sec. 4.2). This general result does not depend on any regularity, in fact not even on any underlying geometric structure. But it is then shown (Sec. 4.7) that such purely algebraic algorithms have limited efficiency.

To obtain full efficiency geometric interpolation should be used. The main theoretical vehicle to analyze such situation is localization, briefly described in Sec. 5.3.

Various AMG concepts and tools are suggested and discussed by the present study. This includes the algebraic meaning of smoothing by relaxation (Sec. 3.5), the general "rule of block relaxation" and hence the AMG rule of interpolating along strong connections (Sec. 3.5), the possibility of using few interpolation points even when many points are strongly connected (Sec. 4.4), preliminary discussions of handling PDE systems (Sec. 5.1) and using geometric information in selecting coarse-grid variables and in constructing interpolation (Sec. 5.2), and partial relaxation passes (Sec. 5.4).

2. Classes of Symmetric Matrices

We deal with the solution of the linear system

$$\sum n_{j=1} a_{ij} X_j = b_i, \qquad (i = 1, \ldots, n) \tag{2.1}$$

or, in matrix notation

$$AX = b.$$

The matrix A can be fairly general, but for some of the processes described below to make sense we will see later that some properties of A are needed which are usually satisfied whenever the system of equations has a geometrical background; it may, but need not, arise from discretizing a continuous problem.

2.1 Local positive definiteness

In the present work we do not attempt to derive the most general theory. To make our study easier we start with the case that A is *symmetric*: $a_{ij} = a_{ji}$; or, in the complex case, *Hermitian*: $a_{ij} = \overline{a_{ji}}$. Except for Theorem 3.4, this will be assumed throughout. Also, we assume A to be *positive definite* (in case of simple point relaxation. More general cases may be considered by using distributive and weighted relaxation schemes, in which case it is $\tilde{A}$, defined in Sec. 3.3, that should be symmetric and positive definite).

We will use, more specifically, the assumption that A is "*locally*" (or "*sparsely*") positive definite, i.e., for any vector x

$$xTAx = \sum_k f_k(x)2, \tag{2.2}$$

where each $f_k(x)$ is a linear function involving only few of the components x_j. This class naturally involves least square problems with local (sparse) measurements and (variational) discretization of elliptic equations. In addition we note in the next paragraphs several important classes of locally positive definite matrices.

2.2 Positive type matrices

The simplest case of local positive definiteness is that of positive-type matrices, i.e., matrices A for which

$$a_{ij} \leq 0 \text{ for all } j \neq i, \text{ and } \sum_j a_{ij} \geq 0 \text{ for all } i. \tag{2.3}$$

It is easy to see that

$$xTAx = \tfrac{1}{2} \textstyle\sum_{i,j}(-a_{ij})(x_i - x_j)2 + \sum_i(\sum_j a_{ij})x_i 2, \tag{2.4}$$

which is of the form (2.2). Note from here that a sparsely positive definite matrix is itself not necessarily sparse.

2.3 Zero row sums

sometimes, without loss of generality, it is convenient to assume that a positive-type matrix A has zero row sums, i.e.,

$$\sum_j a_{ij} = 0 \text{ for all } i. \tag{2.5}$$

Indeed, if (2.5) is not satisfied, we can always add a *slack variable*, x_0 say, taking

$$a_{0i} = a_{i0} = -\sum n_{j=1} a_{ij}, \qquad (i = 0, 1, \ldots, n).$$

The new system clearly satisfies both (2.3) and (2.5). Note that with (2.5) the system is singular; the solution is determined at most up to an additive constant: If x is a solution, so is also $x - (c, c, \ldots, c)T$. Thus, by choosing $c = x_0$, we can always obtain a solution in which $x_0 = 0$ and hence the non-slack variables satisfy

the original system. Note also that, with the zero row sum assumption (2.5), relation (2.4) simplifies to

$$xTAx = \tfrac{1}{2}\sum_{ij}(-a_{ij})(x_i - x_j)2. \tag{2.6}$$

The introduction of a slack variable is not just a theoretical artifact. It turns out to be a useful algorithmic device as well. Including the slack variable and the corresponding equation in the Gauss-Seidel relaxation is equivalent, in terms of the original system, to a step of adding the same constant to all unknowns so as to satisfy the sum of all the equations. Such a step is indeed needed in some cases (e.g., Helmholz equation $-\Delta u + Cu = f$ with small $|C|$ and Neumann boundary conditions) in order to get fast convergence on the coarsest grid (cf. [23] and [1, p. 436]). On the finest grids this is usually not needed, so some work could be saved by avoiding there the introduction of the slack variable. (It can usually altogether be avoided by solving the coarsest system directly.)

2.4 Essentially positive type matrices

Quite often the matrix A is "almost" positive type, in the sense that $\sum_j a_{ij} \geq 0$ is still true for all i, but the requirement $a_{ij} \leq 0$ is violated for some $i \neq j$. If those positive a_{ij} are sufficiently small, they need not bother us. This, for example, is usually the situation in discretizing second-order elliptic equations with mixed terms ($\partial 2u/\partial x_1 \partial x_2$ etc.). For each $i \neq j$ with $a_{ij} > 0$, we can then usually find another variable, k say, such that both $-a_{ik} > \alpha a_{ij}$ and $-a_{jk} > \beta a_{ij}$, where $\alpha > 1$, $\beta > 1$ and $(\alpha - 1)(\beta - 1) = 1$. We can then replace $(x_i - x_j)2$ in (2.4) by

$$\alpha(x_i - x_k)2 + \beta(x_j - x_k)2 - \big(\overline{\alpha}x_i + \overline{\beta}x_j - (\overline{\alpha} + \overline{\beta})x_k\big)2 \tag{2.7}$$

where $\overline{\alpha} = (\alpha - 1)1/2$ and $\overline{\beta} = (\beta - 1)1/2$, so that the resulting quadratic will have the form (2.2). Whenever we can so repair (2.4) to get the form (2.2), we say that A is *essentially positive type*.

2.5 RB matrices and RB vectors

Another important class of matrices, which are used below mainly as examples for the sharpness of some estimates, is the class of RB (Red-Black) matrices. These are defined as matrices A for which there exist two disjoint subsets ("colors") S_R and S_B ("red" and "black") such that $S_R \cup S_B = \{1, 2, \ldots, n\}$ and such that $a_{ij}a_{ik} \geq 0$ if j and k are in the same subset and $a_{ij}a_{ik} \leq 0$ if j and k are in different subsets.

All positive-type matrices with Property A (see[24]), for example, are RB matrices, including all those arising from the common $(2d+1)$-point discretization of

second-order elliptic equations in d-dimensional space. In addition, the common 13-point discretization of the biharmonic operator, the staggered discretization of the Stokes or the compressible or incompressible Navier-Stokes equations (the principal and subprincipal terms; see [5] or [6]), and many other discretized problems, yield RB matrices, not necessarily symmetric and not necessarily of positive type.

With respect to a fixed RB matrix, a vector x will be called an *RB-constant vector* if $x_i = -x_j$ whenever $i \varepsilon S_R$ and $j \varepsilon S_B$.

3. Relaxation and Smoothing

3.1 Gauss-Seidel relaxation

In this section (except in Theorem 3.4) we assume A to be symmetric (or Hermitian) with positive diagonal: $a_{ii} > 0$ for $i = 1, \ldots, n$. Let x be some approximation to the solution X of (2.1). We denote the error by $e = x - X$ and the residual by $r = Ae$. Thus, $r_i = \sum_j a_{ij} x_j - b_i = \sum_j a_{ij} e_j$. We will use the square error norms

$$E = eTr \quad \text{and} \quad R = \sum_i |r_i|2/a_{ii}$$

and observe their changes in various processes. (The superscript T stands for transposed, and in case of complex values, for complex-conjugate transposed.)

Consider a Gauss-Seidel relaxation sweep. Let $e0$, $r0$, $E0$ and $R0$ denote the values of e, r, E and R, respectively, before the sweep, and $e1$, $r1$, $E1$ and $R1$ the corresponding values after the sweep. And let r_j* denote the "dynamic residual", i.e., the value of r_j just before the step of relaxing x_j. Here, "relaxing x_j" means the step of replacing x_j by a new value $x_j + \delta_j$, as a result of which each r_i is replaced by $r_i + a_{ij}\delta_j$, $(i = 1, \ldots, n)$. In Gauss-Seidel one takes

$$\delta_j = -r_j*/a_{jj} \tag{3.1}$$

so that, in particular, r_j is replaced by 0. Using symmetry (this is in fact the main point where symmetry is used), it is also easy to see that E is thereupon replaced by

$$E + \sum_i \overline{e}_i a_{ij} \delta_j = E + \overline{r}_j * \delta_j = E - \frac{|r_j*|2}{a_{jj}} = E - a_{jj}|\delta_j|2. \tag{3.2}$$

(The complex-conjugate bars are shown in case one is interested in complex functions, like those appearing in the corresponding mode analysis.) The Gauss-Seidel sweep consists of relaxing x_1, then x_2, then x_3, etc. up to x_n. Hence,

$$E1 = E0 - R*, \tag{3.3}$$

where

$$R* = \sum_j |r_j*|2/a_{jj} = \sum a_{jj}|\delta_j|2.$$

Thus, in a full sweep the value of E is decreased by a quantity $R*$, akin to $R0$ and to $R1$. Actually, we now prove that the decrease is at least comparable to $R0$ and to $R1$.

Theorem 3.1. The decrease of E in one Gauss-Seidel sweep is $E0 - E1 \geq \gamma_0 R0$, where $\gamma_0 = [(1+\gamma_-)(1+\gamma_+)]-1$ and

$$\gamma_- = \gamma_-(A) = \max_i \sum_{j<i} |a_{ij}|/a_{ii}, \qquad \gamma_+(A) = \max_i \sum_{j>i} |a_{ij}|/a_{ii}. \tag{3.4}$$

Proof. After relaxing x_i we have a zero residual at i, that is

$$r_i = r_i 0 + \sum_{j \leq i} a_{ij}\delta_j = 0. \tag{3.5}$$

Hence, by the Cauchy-Schwarz inequality and (3.4),

$$\begin{aligned}
\textstyle\sum_i |r_i 0|2/a_{ii} &= \textstyle\sum_i a_{ii}-1|\sum_{j\leq i} a_{ij}\delta_j|2 \\
&\leq \textstyle\sum_i a_{ii}-1\big(\sum_{j\leq i}|a_{ij}|\big)\big(\sum_{j\leq i}|a_{ij}||\delta_j|2\big) \\
&\leq (1+\gamma_-)\textstyle\sum_i\sum_{j\leq i}|a_{ij}||\delta_j|2 \\
&= (1+\gamma_-)\textstyle\sum_j|\delta_j|2\sum_{i\geq j}|a_{ij}| \\
&\leq (1+\gamma_-)(1+\gamma_+)\textstyle\sum_j a_{jj}|\delta_j|2
\end{aligned}$$

and the theorem follows from (3.3).

Corollary 3.1. If A is symmetric and positive definite than

$$E \geq \gamma_0 R.$$

This follows from Theorem 3.1, since $E1 \geq 0$. Numerical values of γ_0 and the sharpness of the estimate are discussed in the next section.

Theorem 3.2. The decrease of E in one Gauss-Seidel sweep is at least $\gamma_1 R1$, where $\gamma_1 = [\gamma_-(A)\gamma_+(A)]-1$.

Proof. After relaxing x_i we have $r_i = 0$, hence

$$r_i 1 = \sum_{j>i} a_{ij}\delta_j. \tag{3.6}$$

By Cauchy-Schwarz inequality

$$\begin{aligned}\textstyle\sum_i |r_i 1|2/a_{ii} &= \textstyle\sum_i a_{ii}{-}1|\sum_{j>i} a_{ij}\delta_j|2 \\ &\le \textstyle\sum_i a_{ii}{-}1\big(\sum_{j>i}|a_{ij}|\big)\big(\sum_{j>i}|a_{ij}||\delta_j|2\big) \\ &\le \gamma_+ \textstyle\sum_i \sum_{j>i}|a_{ij}||\delta_j|2 \\ &= \gamma_+ \textstyle\sum_j |\delta_j|2 \sum_{i<j}|a_{ij}|\end{aligned}$$

and the theorem follows from (3.4) and (3.3).

3.2 Sharpness of bounds and converse theorem

It will now be shown that the constants γ_0 and γ_1 in Theorems 3.1 and 3.2 are practically optimal: In the most important cases the stated bounds are actually attained, and in other cases they are reasonably approached. Moreover, in any case, constants as close to practical as one wishes can be derived. (The reader not interested in these details can skip this section.)

Consider first the case of an *infinite* RB matrix A (see Sec. 2.5) with

$$\gamma_- = \sum_{j<i} |a_{ij}|/a_{ii}, \qquad \gamma_+ = \sum_{j>i} |a_{ij}|/a_{ii} \tag{3.7}$$

for all i. Then, for an RB-constant vector of corrections δ, the inequalities in the above proofs all become equalities. Hence, γ_0 and γ_1 are the best possible constants in terms of γ_- and γ_+. Moreover, for a *finite* RB matrix A and an RB vector δ, if (3.7) is *approximately* satisfied, in the sense that for some small subset $S_\partial \subset \{1, 2, \ldots, n\}$ and for a small positive constant ε we have

$$\sum_{j<i} |a_{ij}|/a_{ii} \ge \gamma_- - \varepsilon, \qquad \sum_{j>i} |a_{ij}|/a_{ii} \ge \gamma_+ - \varepsilon \tag{3.8}$$

for all $i \notin S_\partial$, then the inequalities in the proofs are not far from equalities, hence the bounds γ_0 and γ_1 are almost attained. This is indeed the usual situation in geometrically ordered relaxation of finite difference approximations to differential equations on regular grids with meshsize h, where S_∂ is a subset of points near the boundary and $\varepsilon \to 0$ as $h \to 0$ if the coefficients of the differential equation are continuous. The bounds γ_0 and γ_1 are then attained in the limit $h \to 0$.

The above class of matrices include most of the regular symmetric matrices arising in differencing differential equations on regular grids. For matrices not in that class the bounds are not generally approached, but the practical values of γ_0 and γ_1 are not far from the stated ones, provided the definition of γ_+ and γ_- are

slightly modified by excluding some special values of i from the maximization in (3.4).

Moreover, for regular and large ($h \to 0$) grids, precise values of γ_0 and γ_1 can always be derived by local mode analysis (cf. [6]). Usually, values calculated by local mode analysis are rigorously true only for the infinite-domain (or the rectangular, etc.) problem. In determining γ_0 and γ_1, however, one can rigorously justify the mode analysis values, $\tilde{\gamma}_0$ and $\tilde{\gamma}_1$ say, for any domain, at least in the limit of vanishing meshsize ($h \to 0$) and with proper modification of the relaxation scheme at a small number of points. This is done in the following way: Given a small positive constant ε, before each relaxation sweep a (possibly large) number $\nu(\varepsilon)$ of relaxation sweeps is performed over all those gridpoints whose distance from the boundary is less than $k(\varepsilon)h$. With a proper selection of $\nu(\varepsilon)$, $k(\varepsilon)$ and $h_0(\varepsilon)$, we can prove that for any $0 < h \leq h_0(\varepsilon)$ Theorems 3.1 and 3.2 are satisfied with $\gamma_0 = \tilde{\gamma}_0 - \varepsilon$ and $\gamma_1 = \tilde{\gamma}_1 - \varepsilon$. The detailed proof is omitted since it is not used below.

For a symmetric constant-coefficient difference operator Lh with *geometrically consistent relaxation ordering* (i.e., if grid point $\underline{x}$ is relaxed after $\underline{x} - \underline{a}$, then it is relaxed before $\underline{x} + \underline{a}$), local mode analysis yields particularly simple results: $Lh \exp(i\underline{\theta x}/h)$ can be written as $[a + b(\underline{\theta}) + \overline{b}(\underline{\theta})] \exp(i\underline{\theta x}/h)$, where $a = a_{ii}$, and each relaxation sweep transforms the error component $\exp(i\underline{\theta x}/h)$ to $(a+b)^{-1}\overline{b} \exp(i\underline{\theta x}/h)$. Hence, it is easy to calculate that for each error component

$$\frac{E0 - E1}{R0} = |\frac{a}{a+b}|2, \qquad \frac{E0 - E1}{R1} = |\frac{a}{b}|2. \tag{3.9}$$

Since $b(0) = \overline{b}(0) = -a/2$, it follows that for smooth components ($\underline{\theta} \approx 0$) we have $E0 - E1 \approx 4R0 \approx 4R1$. Indeed, when relaxation is slow it is clear that $r0 \approx r1 \approx r*/2$.

There are, of course, particular error components for which $E0 - E1$ is considerably larger than $\gamma_0 R0$, and (other) components for which $E0 - E1$ is much larger than $\gamma_1 R1$. Still, the lower bounds of Theorems 3.1 and 3.2 are very realistic, because they are (almost) attained exactly for the components which are normally dominant at the stage for which each theorem is used. Theorem 3.1 is used (see for example Sec. 4.2) to estimate the drop in E at the first relaxation sweep after the coarse-grid correction. Due to the nature of interpolation, the dominant errors $e0$ at this stage are normally highly oscillatory (e.g., with a large RB component) which are indeed the type of errors for which the bound in Theorem 3.1 is attained or approached. Theorem 3.2, on the other hand, is used to estimate the smallness of R when relaxation slows down. At this stage the dominant errors are smooth (see Sec. 3.5), for which the bound of Theorem 3.2 is indeed attained (e.g. when A is positive-type), or reasonably approached.

For symmetric positive type matrices approximating continuous operators,

with geometrically consistent relaxation ordering, it is easy to see that $\gamma_-(A) \approx \gamma_+(A) \approx \frac{1}{2}$. Hence, $\gamma_0 \approx \frac{4}{9}$ and $\gamma_1 \approx 4$, values which are actually attained for the highest- and lowest-frequency errors, respectively. For this type of operators, as well as many others, the following theorem, a converse of Theorem 3.1, can be of interest.

Theorem 3.3. *If A is symmetric and $\gamma_-(A)\gamma_+(A) < 1$ then $E0 - E1 \leq [1 - (\gamma_-\gamma_+)1/2] - 2R0$.*

Proof. Denoting $D = \sum_i |\sum_{j<i} a_{ij}\delta_i|2/a_{ii}$ it is easy to see (cf. the proof of Theorem 3.2) that $D \leq \gamma_-\gamma_+ R*$. Hence

$$\begin{aligned} R* &= \textstyle\sum_i |r_i 0 + \sum_{j<i} a_{ij}\delta_i|2/a_{ii} \\ &\leq R0 + 2\textstyle\sum_i |r_i 0||\sum_{j<i} a_{ij}\delta_i|/a_{ii} + D \\ &\leq R0 + 2(R0D)1/2 + D \\ &\leq R0 + 2(\gamma_-\gamma_+ R0R*)1/2 + \gamma_-\gamma_+ R* \end{aligned}$$

which can be rewritten as

$$\big[\big(1 + (\gamma_-\gamma_+)1/2\big)(R*)1/2 + (R0)1/2\big]\big[\big(1 - (\gamma_-\gamma_+)1/2\big)(R*)1/2 - (R0)1/2\big] \leq 0$$

yielding, by (3.3), the theorem.

In particular, for symmetric positive type operators with constant coefficients and with consistent relaxation ordering we have $\gamma_+ = \gamma_- = \frac{1}{2}$, hence we get the interesting result

$$4R1 \leq E0 - E1 \leq 4R0. \tag{3.10}$$

When relaxation slows down then $R1 \approx R0$, hence we have quite a precise measure for the relation between that slowness and the smallness of the ratio R/E. Namely, the convergence factor is almost precisely $1 - 4R/E$. As we saw above (following (3.9)), this is in practice the slow convergence factor even when the operator is not of positive type.

3.3 Other point relaxation schemes

Theorems 3.1 and 3.2 can be extended to perhaps any reasonable point relaxation, including point Jacobi (Richardson) schemes, red-black Jacobi (even with couplings within each color), etc. In fact, if sufficiently small, but depending only on

$$\gamma_*(A) = \max_i \sum_j |a_{ij}|/a_{ii}, \tag{3.11}$$

under-relaxation parameter ω is used (as in most multigrid theories, but not in practice), then the proof is trivial, since $R0$, $R*$ and $R1$ are obviously comparable, and the effect of relaxing points simultaneously is $0\big(\omega\gamma_*(A)\big)$ compared to relaxing them successively.

In *point Jacobi*, for example, $\delta_i = -\omega r_i 0/a_{ii}$ is added to x_i, simultaneously for all i. Hence

$$E1 = E0 - 2\omega \sum_i |r_i 0|2/a_{ii} + \omega 2 \sum_{i,j} \frac{a_{ij}}{a_{ii}a_{jj}} \overline{r}_i 0 r_j 0$$

and by the Cauchy-Schwarz inequality and (3.11)

$$E1 \le E0 - \omega[2 - \omega\gamma_*(A)]R0. \tag{3.12}$$

Note that for positive type operators $\gamma_*(A) \le 2$, hence $E0 - E1 \ge \gamma_0 R0$, with $\gamma_0 = 2\omega(1-\omega)$. This value of γ_0 is again sharp; it is attained for infinite RB matrices satisfying (3.7) and closely approached for finite RB matrices.

Other types of point relaxation are the *distributive schemes* (see [9, §5.3], [4, §4.1] or [6, §3.4]) and *weighted schemes* (see [3, §3.3] or [6, §3.4]), and the combination of distributive and weighted relaxation. Such schemes should be used whenever $\gamma_*(A)$ is too large. Each such scheme is equivalent to a simple (non-distributive and not weighted) scheme for another matrix $\tilde{A} = BAC$, hence its analysis follows from applying the above theory to $\tilde{A}$ (if $\tilde{A}$ is symmetric). The distributed Gauss-Seidel (DGS) scheme, for example, is equivalent to a Gauss-Seidel relaxation of $\tilde{A} = AC$ (see [4, §4.1]).

Kaczmarz relaxation [16], in particular, is equivalent to Gauss-Seidel relaxation for $\tilde{A} = AAT$. Indeed, Kaczmarz relaxation for the system (2.1) is defined as follows. Denoting by a_i the i-th row of A, a *Kaczmarz step* corresponding to that row is defined as the replacement of x by $x - (r_i/a_i a_i T)a_i T$, thereby forcing r_i to zero. (This in fact is the way to force r_i to zero with least square changes in x.) A full *Kaczmarz sweep* is the employment of such a step for every row of A, in the natural order. Now, if (2.1) is rewritten as $\tilde{A}\tilde{X} = b$, where $\tilde{A} = AAT$ and $X = AT\tilde{X}$, it is easy to see that the Kaczmarz step corresponding to the i-th row is equivalent to replacing $\tilde{x}_i$ (the current approximation to $\tilde{X}_i$) by $\tilde{x}_i - r_i/a_i a_i T$, which is exactly the i-th step of the Gauss-Seidel relaxation for $\tilde{A}\tilde{X} = b$. Thus, to analyze the Kaczmarz relaxation, one simply applies Theorem 3.1 and 3.2 to $\tilde{A}$ and $\tilde{e}$ instead of A and e, where $e = AT\tilde{e}$. This immediately yields the following theorem.

Theorem 3.4. For any matrix A (not necessarily symmetric), let

$$\tilde{E} = \sum_i |e_i|2 \quad \text{and} \quad \tilde{R} = \sum_i \big(|r_i|2/\sum_\lambda |a_{i\lambda}|2\big) \tag{3.13}$$

and let $\tilde{E}0$ *and* $\tilde{R}0$ *be the initial values of these square norms and* $\tilde{E}1$ *and* $\tilde{R}1$ *their values after one Kaczmarz relaxation sweep. Then*

$$\tilde{E}1 \leq \tilde{E}0 - \tilde{\gamma}_0 \tilde{R}0 \tag{3.14}$$

and also

$$\tilde{E}1 \leq \tilde{E}0 - \tilde{\gamma}_1 \tilde{R}1 \tag{3.15}$$

where

$$\begin{aligned}
\tilde{\gamma}_0 &= [(1+\tilde{\gamma}_-)(1+\tilde{\gamma}_+)]-1 \\
\tilde{\gamma}_1 &= (\tilde{\gamma}_-\tilde{\gamma}_+)-1 \\
\tilde{\gamma}_- &= \max_i \big(\textstyle\sum_{j<i} |\sum_\lambda a_{i\lambda}\overline{a}_{j\lambda}|\big) / \sum_\lambda |a_{i\lambda}|2 \\
\tilde{\gamma}_+ &= \max_i \big(\textstyle\sum_{j>i} |\sum_\lambda a_{i\lambda}\overline{a}_{j\lambda}|\big) / \sum_\lambda |a_{i\lambda}|2
\end{aligned}$$

Observe that $\tilde{\gamma}_\pm$ are reasonably small in many cases where $\gamma_\pm$ of Theorem 3.1 are not. In fact, for any reasonable discretization of usual differential equations $\tilde{\gamma}_+$ and $\tilde{\gamma}_-$ are smaller than 3 or so. More generally, for systems not necessarily arising from differential equations, $\tilde{\gamma}_+$ and $\tilde{\gamma}_-$ can be regarded as measures of the overlap, or dependence, between the individual equations, hence $\tilde{\gamma}_0$ and $\tilde{\gamma}_1$ are measures of their independence.

Theorem 3.4 has in fact general implications. Since Kaczmarz relaxation can be applied to *any* (even general rectangular) matrix A, and it converges whenever a solution exists (see [22]), the theorem shows that fast convergence by a point relaxation is always obtainable as long as $\tilde{R}$ is comparable to $\tilde{E}$.

The converse is also true; namely, when $\tilde{R} \ll \tilde{E}$ then no point relaxation scheme can yield fast convergence. This is because any point relaxation scheme which is based on one residual r_i at a time must introduce into e changes which are at most comparable to the *normalized* residual $\overline{r}_i = r_i / \parallel a_i \parallel_2 = a_i e / \parallel a_i \parallel_2$ where $\parallel \cdot \parallel_2$ is the ℓ_2 norm (otherwise, for many errors the changes would be large compared with the errors themselves, bringing about fast divergence). Hence the changes introduced into $\parallel e \parallel_2 = \tilde{E}1/2$ in a complete sweep are at most comparable to $\parallel \overline{r} \parallel_2 = \tilde{R}1/2$.

Thus, very generally, *point relaxation must slow down when and only when the normalized residuals are small compared with the errors.* Note that for an arbitrary error vector most of the normalized residuals $\overline{r}_i$ would be *comparable* to e (unless a_i has a large number of dominant terms comparable to each other). Hence, point relaxation must slow down only for a special class of error components.

3.4 Block relaxation

In case of block relaxation, such as line relaxation or collective relaxation, the theory above still applies, but in block notation: Denoting by E_i the i-th block of e, by A_{ij} the corresponding blocks of A, a Gauss-Seidel step for the i-th block replaces E_i by $E_i - A_{ii}{-}1R_i*$, where $R_i = \sum_j A_{ij}E_j$ and R_i* is the value of R_i just *before* this step. Hence, analogously to (3.2), this step reduces the error energy E by

$$R_i*TA_{ii}{-}1R_i*. \tag{3.16}$$

The full *block Gauss-Seidel* sweep (making such a step successively for all the blocks) will therefore reduce E by

$$\hat{R}* = \sum_i R_i*TA_{ii}{-}1R_i*. \tag{3.17}$$

From here one can prove, analogously to the proofs of Theorems 3.1 and 3.2, the following theorem.

Theorem 3.5. *Let $\hat{R}0$ and $\hat{R}1$ respectively be the values of the block residual norm*

$$\hat{R} = \sum R_iTA_{ii}{-}1R_i$$

before and after a block-Gauss-Seidel relaxation sweep, for a symmetric matrix A. Then the decrease $E0 - E1$ of the error energy E due to that sweep is at least $\hat{\gamma}_0\hat{R}0$, and at least $\hat{\gamma}_1\hat{R}1$, where

$$\hat{\gamma}_0 = (\hat{\gamma}_+{\prime}\hat{\gamma}_-{\prime}){-}1, \qquad \hat{\gamma}_1 = (\hat{\gamma}_+\hat{\gamma}_-){-}1$$
$$\hat{\gamma}_- = \max_i(\zeta_i \textstyle\sum_{j<i} \alpha_{ij}), \qquad \hat{\gamma}_-{\prime} = \max_i(\zeta_i \textstyle\sum_{j\le i} \alpha_{ij})$$
$$\hat{\gamma}_+ = \max_i(\zeta_i \textstyle\sum_{j>i} \alpha_{ij}), \qquad \hat{\gamma}_+{\prime} = \max_i(\zeta_i \textstyle\sum_{j\ge i} \alpha_{ij})$$

ζ_j being the spectral radius of $A_{jj}{-}1$ and α_{ij} being the L_2 norm of A_{ij}.

This theorem, too, can be extended to other block relaxation schemes, such as block-Jacobi, zebra and distributed block relaxation schemes, the only change being the form of dependence of $\hat{\gamma}_0$ and $\hat{\gamma}_1$ on $\hat{\gamma}_\pm$ and $\hat{\gamma}_\pm{\prime}$. The derived constants are again sharp in various cases, while in other cases they are not far off, and optimal constants can be derived from local mode analyses, rigorously justifiable for sufficiently large ($h \to 0$) regular grids and geometric ordering of relaxation.

Concerning simultaneous relaxation, we would like to add here another result, which is an interesting generalization of formula (3.2) above and of Corollary 3.1, and which will be useful in future work.

Theorem 3.6. *If A is symmetric and positive definite, then relaxing simultaneously the I-th block reduces E at least by*

$$\gamma_0 I \sum_{i\varepsilon I} |r_i*|2/a_{ii}$$

where

$$\begin{aligned}\gamma_0 I &= [(1+\gamma_- I)(1+\gamma_+ I)]-1\\ \gamma_- I &= \max_{i\varepsilon I} \sum_{\substack{j<i\\ j\varepsilon I}} |a_{ij}|/a_{ii}\\ \gamma_+ I &= \max_{i\varepsilon I} \sum_{\substack{j>i\\ j\varepsilon I}} |a_{ij}|/a_{ii}\end{aligned}$$

The proof follows from (3.16) and from applying Corollary 3.1 to the matrix A_{II} (instead of A) and to the error $A_{ii}-1R_{i}*$ (instead of e).

3.5 Algebraic sense of smoothing

A basic premise of multigrid processing is that relaxation is very efficient in smoothing the error, which makes it possible to efficiently approximate relaxed errors on coarser grids. More precisely, it is asserted that when the convergence of relaxation becomes slow, the error must be smooth in some sense. It is instructive to see that this smoothing has a well defined algebraic meaning. (This discussion is heuristic. It is not used in the subsequent multigrid theory, except as a motivation.)

Indeed, by Theorem 3.2, when the convergence of Gauss-Seidel relaxation for a symmetric matrix A slows down, i.e., when $E0 - E1 \ll E0$, we must have

$$R \ll E. \tag{3.18}$$

This by itself can be regarded as an expression of smoothness of e, but we can give it even more intuitive interpretations. Consider for example the positive-type case (2.3). By the Cauchy-Schwarz inequality

$$E2 \leq R \sum a_{ii} e_i 2, \tag{3.19}$$

hence, the slow convergence condition (3.18) implies

$$E \ll \sum a_{ii} e_i 1, \tag{3.20}$$

or, by (2.4),

$$\sum_{i,j} (-a_{ij})(e_i - e_j)2 \ll \sum a_{ii} e_i 2. \tag{3.21}$$

Since $a_{ij} \leq 0$ for $i \neq j$, this means that $e_i - e_j$ is on the average small compared with $\| e \|$ whenever $|a_{ij}|$ is comparable with either $\max_{k\neq i} |a_{ik}|$ or $\max_{k\neq j} |a_{kj}|$, i.e., whenever i and j are "*strongly coupled*". We can thus say that the Gauss-Seidel relaxation (and similarly the other schemes) *smoothes the error along the strongest couplings* (and couplings comparable to the strongest).

Greater smoothing is of course obtained by block relaxation. When such relaxation slows down we must have, by Theorem 3.5,

$$\hat{R} \ll E. \tag{3.22}$$

Hence, applying now the Cauchy-Schwarz inequality to E in the form $E = \sum(A_{\ell\ell}1/2E_\ell)T(A_{\ell\ell}{-}1/2R_\ell)$, we obtain

$$\sum_{i,j}(-a_{ij})(e_i - e_j)2 \ll \sum E_\ell T A_{\ell\ell} E_\ell. \tag{3.23}$$

Hence, for $(e_i - e_j)$ to be small on the average, it is enough that a_{ij} is comparable to

$$\min(\max_{k\not\in\ell_i}|a_{ik}|, \max_{k\not\in\ell_j}|a_{jk}|),$$

where ℓ_i is the block containing i. Thus, *by including the strongest couplings in the blocks, block relaxation will smooth the error also along next-strongest couplings.*

We will indeed see later, more rigorously, that we obtain good multigrid rates if we see to it that our interpolation is along couplings which are about strongest (or next-strongest) in the above senses. From here follows the *rule of block relaxation* [6, §3.3]:

A locally strongly coupled block of unknowns, which is locally decoupled from (or weakly coupled with) the coarser-grid variables, should be relaxed simultaneously.

Although the smoothness may more easily be visualized through (3.21) or (3.23), it is the stronger conditions (3.18) or (3.22) which we will actually use in our theory below. In fact, from Example 4.2 below we can see that for algebraic multigrid the weaker conditions are not enough.

Condition (3.18) is indeed fundamentally stronger than (3.21), since it implies small changes in *derivatives* of e along strong couplings, while (3.21) implies only small changes in e itself.

For simplicity, we have discussed smoothing for positive type matrices only. The *generalization to locally positive definite matrices* (see Sec. 2.1) is straightforward. Instead of (2.4) one generally uses (2.2) in the left-hand side of (3.21), showing that when relaxation slows down $f_k(e)$ must on the average be small compared with $a_{ii}e_i2$. In case of essentially positive type matrices (Sec. 2.4) this still means smallness of certain error differences compared with the error itself (see (2.7)). More generally the smallness of $f_k(e)$ does not necessarily mean smallness of error differences, but can still be used for devising the coarse-to-fine interpolation. In various cases the linear forms f_k are not explicitly known. The sense

of "smoothness" (and hence interpolation) can then be based on *pre-relaxation* (i.e., on the local behavior of "smooth" vectors previously obtained by relaxing the homogeneous equation $AX = 0$) or on geometry (see Sec. 5.2).

The notion of "smoothness" should indeed be generalized, especially in the algebraic context, the most general concept being "approximability by a much lower dimensional subspace". For *any* matrix, even far from symmetry or far from positive definiteness, even indeed rectangular, a suitable relaxation, such as Kaczmarz, will converge fast as long as the normalized residuals are not small (compared with the error; see Theorem 3.4 and the discussion thereafter). Errors that must converge slowly for any relaxation scheme must therefore be approximable by the subspace spanned by eigenvectors of A with small eigenvalues. The meaning of this subspace (and hence how to operate it) depends on the kind of problem. If for example A has good h-ellipticity measure (see [6; §2.1]) then the subspace contains only smooth functions. If A is highly indefinite, as for example the discretization of standing wave operators, then the subspace is spanned by "rays", i.e. natural oscillations (unconstrained waves) modified by smooth multipliers. Etc.

4. Two-Level Algebraic Analysis

4.1 General description and notation

We have thus seen that when convergence by relaxation is slow, the error e must be smooth, in the sense that along strong algebraic couplings it has small variations (or sometimes in a more general sense). We can therefore hope to get good approximations to that error "on a coarser grid", i.e., by solving a much smaller system of equations. For that purpose we have to construct that "coarse level" system of equations and specify the "interpolation", i.e., the relation by which an approximation to e (the "fine level" error) is obtained from the coarse level solution. In geometrical setting there are several quite natural ways for selecting the coarse level and the interpolation, but in purely algebraic terms the methods should be different.

We will generally denote the coarse-level vector of unknowns by $ec = (e_1c, e_2c, \ldots, ec_{nc})T$ and assume a coarse-to-fine linear interpolation I_c given by

$$(I_c ec)_i = \sum_{k=1} ncw_{ik} ec_k, \qquad (1 \le i \le n). \tag{4.1}$$

Thus, $I_c ec$ is a fine-level vector, intended to approximate the fine-level error vector e.

We can usually (but not necessarily) think of ec as approximating a properly chosen subset of e, with each ec_k approximating a particular $e_{F(k)}$. The smoothness properties of e resulting from relaxation, as described above, should be used

in selecting this subset and in choosing the weights w_{ik}. In particular they imply that w_{ik} should be large only when i and $j = F(k)$ are *strongly connected*. By this we roughly mean that there exists a short chain $k_0, k_1, \ldots, k_\ell$ (usually $\ell = 1$ or 2) such that $k_0 = i$, $k_\ell = j$ and k_α is strongly coupled with $k_{\alpha-1}$, $(\alpha = 1, \ldots, \ell)$.

At this point we will not prescribe, however, any exact rules for choosing the weights. Instead, we will ask our theory below to furnish the general conditions that should be satisfied by the weights. As a result, the theory will apply to regular (geometric) multigrid as well.

Having computed ec, its interpolant $I_c ec$ is used to correct the fine-level approximation. In this section we will use superscripts $-$ and 0, respectively, to denote the fine-level parameters before and after this correction. In particular, $e-$, $r-$, $E-$ and $R-$ will respectively denote the values of the error vector e, the residual vector r and the square norms E and R *before* the coarse-level correction, while $e0$, $r0$, $E0$ and $R0$ will denote the corresponding values *after* that correction. Hence $e0 = e - -I_c ec$, $r\ell = Ae\ell$, $E\ell = \sum_i e_i \ell r_i \ell$ and $R\ell = \sum_i |r_i \ell| 2/|a_{ii}|$, $(\ell = -, 0)$.

What equations should be satisfied by ec? In regular (geometric) multigrid situations the cheapest way for deriving the coarse-level equations is usually to write coarse-grid approximations to the original differential equations (see more about this issue in [6, §11]). In algebraic multigrid this approach is not applicable, and a general prescription in the symmetric case (see [10]) is to require ec to satisfy

$$(I_c)TAI_c ec = (I_c)Tr - . \tag{4.2}$$

This is equivalent to choosing that ec for which $E0$ will be minimal. It therefore implies the following three properties of the coarse-level correction:

(i) $E0 \leq E-$.

(ii) $E0 \leq \tilde{E}0$, where $\tilde{E}0$ is the value of E after a correction by any other coarse-level vector $\tilde{e}c$.

(iii) $e0$ is a projection of $e-$; i.e., the linear transformation $e0 = Pe-$ satisfies $e0 = Pe0$. Hence we must have $(I_c)Tr0 = 0$, or, by (4.1),

$$\sum_j w_{jk} r_j 0 = 0 \text{ for all } k. \tag{4.3}$$

4.2 Relaxability after coarse level correction

In this section we show conditions under which the "energy" $E0$ of the error $e0$ produced after the coarse-level correction is efficiently reducible by subsequent

relaxation sweeps. Since the coarse-level correction itself reduces E, this is equivalent to showing the efficiency of the two-level cycle (composed of a coarse-level correction followed by relaxation sweeps). By "efficiency" here we mean that E is reduced to $E_{new} \leq CE$, where $C < 1$ and C does not depend on the size of the matrix A. Actually, C will depend only on "local" constants like $\gamma_*(A)$ (see (3.11)). We thus say that $e0$ is "*relaxable*" if a relaxation sweep will reduce $E0$ to $E1 \leq CE0$ with such C. From Secs. 3.1 and 3.3 it follows that $e0$ is "*point-relaxable*", i.e., relaxable by a point relaxation (Gauss-Seidel, under-relaxation Jacobi, etc.) if (and usually only if – see Sec. 3.2) one can ensure that $R0 \geq C_0 E0$. This relation, and the size of its constant, are governed by the choice of w_{ik}, as the following theorem shows.

Theorem 4.1. if there is a constant C_0 such that for any fine level vector e there exists a coarse level vector ec for which

$$C_0 \sum_j a_{jj}\big(e_j - \sum_k w_{jk} e_k c\big)2 \leq eTAe \tag{4.4}$$

then $R0 \geq C_o E0$.

In the above summation j extends over the fine-level and k over the coarse level, of course. This convention is also kept in the summations below.

Proof. Choosing an ec which satisfies (4.4) with $e = e0$ and using (4.3), one obtains

$$\begin{aligned} E0 &= \textstyle\sum_j e_j 0 r_j 0 - \sum_k e_k c \sum_j w_{jk} r_j 0 \\ &= \textstyle\sum_j r_j 0 \big(e_j 0 - \sum_k w_{jk} e_k c\big). \end{aligned}$$

Hence, by the Cauchy-Schwarz inequality and (4.4)

$$(E0)2 \leq R0 \sum_j a_{jj}\big(e_j 0 - \sum_k w_{jk} e_k c\big)2 \leq R0E0/C0$$

and hence $R0 \geq C_0 E0$.

Corollary 4.1. If I_c is a regular averaging, i.e.,

$$w_{jk} \geq 0 \text{ and } \sum_k w_{jk} = 1, \tag{4.5}$$

and if for any e there exists ec such that

$$C_0 \sum_{j,k} a_{jj} w_{jk} (e_j - e_k c)2 \leq eTAe \tag{4.6}$$

then $R0 \geq C_0 E0$.

This follows from Theorem (4.1) upon replacing e_j in (4.4) by $\sum_k w_{jk} e_j$ and applying the Cauchy-Schwarz inequality.

4.3 Conclusions for positive type matrices

The theorem and the corollary spell out various possible rules of interpolation that can guarantee relaxability. For simplicity we first discuss such rules for the positive-type (2.3) zero-row-sum (2.5) case. In this case interpolation is chosen to be a usual averaging (4.5), and the required inequality (4.6) can be written, by (2.6), as

$$\tfrac{1}{2} \sum_{j,i} (-a_{ji})(e_j - e_i)2 \geq C_0 \sum_{j,k} a_{jj} w_{jk} (e_j - e_k c)2. \tag{4.7}$$

The zero row sum property (2.5) is, of course, not essential. Even without that property it is enough to satisfy (4.7), as is evident from (2.4). But the addition of the slack variable makes it *easier* to satisfy (4.6) for a *larger* C_0. It makes it possible to exploit strong couplings to the slack.

A simple way to satisfy (4.7) is to think of ec as corresponding to a subset of e, with $e_k c$ corresponding to $e_{F(k)}$ say, and to let

$$w_{jk} = \begin{cases} 1 & \text{if} \quad j = F(k) \\ 0 & \text{if} \quad j = F(k\prime), k\prime \neq k \\ a_{j,F(k)} / \sum_{k\prime} a_{j,F(k\prime)} & \text{otherwise} \end{cases} \tag{4.8}$$

so that (4.7) is satisfied with

$$C_0 = \tfrac{1}{2} \min_j \big(\textstyle\sum_k |a_{j,F(k)}| / a_{jj} \big). \tag{4.9}$$

Hence, all we have to do to obtain good relaxability (large C_0) is to select the coarse level so that each fine-level variable j is strongly coupled with the coarse-level variables, i.e., so that the total of its coarse level couplings $\big(\sum_k |a_{j,F(k)}|\big)$ is a large fraction (C_0, at least) of its total couplings $\big(a_{jj} = \sum_i |a_{ji}|\big)$.

This approach was indeed the first one adopted in AMG [10]. It was later realized, however, that (4.7) can be satisfied, with about the same C_0, by an interpolation I_c that uses fewer coarse-level points k for any fine-level point j (see Sec. 4.4). This is important in order to keep the coarse-level problem (4.2) as sparse as possible.

Observe that to satisfy (4.7) it is not necessary that large weights w_{jk} correspond to direct strong couplings $a_{j,F(k)}$. It is enough that they correspond to strong *connection* from j to $F(k)$ (cf. Sec. 4.1). If, for example, $a_{j,F(k)} = 0$ but

there exists an ℓ such that both $a_{j\ell}$ and $a_{\ell,F(k)}$ are strong, then we can use large w_{jk}, because we can satisfy (4.7) through

$$\begin{aligned} w_{jk}(e_j - e_k c)2 &= w_{jk}(e_j - e_\ell - e_{F(k)})2 \\ &\leq \alpha w_{jk}(e_j - e_\ell)2 + \beta w_{jk}(e_\ell - e_{F(k)})2, \end{aligned} \tag{4.10}$$

for some $\alpha, \beta > 1$ such that $(\alpha - 1)(\beta - 1) = 1$. (The choice $\alpha = \beta = 2$ may be standard, but one can take α/β to be an increasing function of $a_{j\ell}/a_{\ell,F(k)}$). Thus, each fine-level point j must have strong connections, but not necessarily direct couplings, to the coarse level.

The interpolation from strongly connected (not necessarily directly coupled) points is the usual practice in *geometric* multigrid, hence the present theory applies to all such cases. It applies, in particular, to equations with strongly discontinuous coefficients, like those in [1].

4.4 Strong coupling overlaps

Consider now the case where each fine-level variable j is strongly coupled to *many*, say roughly M, other variables; i.e., $0(M)$ terms should be taken in the summation (4.9) to make $C_0 = 0(1)$. This is likely to be the situation, for example, on coarse levels produced by AMG procedure (4.2) (having not used in I_c the reduction in interpolation points being discussed hereafter). Does this mean then that we must interpolate to each point from $0(M)$ points in order to obtain relaxability? It does, if the problem is arbitrary, but it does not, if it has (implicit) geometrical background. In the latter case, if each point has $0(M)$ strong neighbors, then those strong neighborhoods must have overlaps of size $0(M)$; i.e., each point j must have (many) neighbors i such that i and j have $0(M)$ common strong neighbors. Choosing even just one of these neighbors i to be a coarse-level point k, and using (4.10) for every ℓ which has strong couplings to both $i = F(k)$ and j, we get $0(M)$ strong-term contributions to majoring the left-hand side of (4.7), which is what we needed in order to satisfy (4.7) with $C_0 = 0(1)$.

Thus, *even for problems with large neighborhoods of strong couplings, we can still get good relaxability by interpolating to any point j from just few neighbors, provided those neighbors' neighborhoods largely overlap the neighborhood of j.*

That such an overlap is *necessary* we can see from the following example.

Example 4.1. The fine grid is a rectangular lattice of M columns by N rows, so that each unknown i has a column number I_i and a row number J_i. The

coefficients of the matrix are defined by

$$a_{ij} = \begin{cases} N & \text{if} \quad j = i \\ -1 & \text{if} \quad j \neq i, I_j = I_i \\ -1 & \text{if} \quad I_j = I_i - 1 (\text{mod} M) \quad \text{and} \quad J_j = J_i \\ 0 & \text{otherwise.} \end{cases}$$

It is easy to find (e.g., by mode analysis) that a point relaxation smoothes the error (very efficiently) along columns, but not along rows. Hence the coarse grid may represent even just one point per column, with interpolation within each column the convergence will be fast. On the other hand the convergence will be slow (convergence factor $1 - 0(N-1)$) if we interpolate to each point from its left $(\text{mod} M)$ neighbor, with whom the point has poorly overlapping neighborhoods.

4.5 Conclusions in more general cases

The above conclusions for the positive-type case can be extended to general locally positive definite matrices (Sec. 2.1): By substituting $\sum_m f_m(e)2$ for $eTAe$ in (4.4), various interpolation strategies to guarantee relaxability, and various limitations, come into view. Basically, one can interpolate to a fine point j by taking that e_j which minimizes $\sum_m j f_m(e)2$, where $\sum_m j$ is a sum over all those $f_m(e)$ which include e_j and which otherwise include only coarse-level variables $e_{F(k)}$. As in Sec. 4.4, this can then be modified to sums which include fewer terms.Practical ways of finding such interpolation when the f_m forms are not explicitly given can be based on pre-relaxation (see Sec. 3.5) or on geometry (Sec. 5.2).

In case of *essentially positive type* matrices, for example, this approach shows that as long as the positive off-diagonal terms a_{ij} are small enough, the positive type procedures discussed above can remain unchanged. If, for example, for each $i \neq j$ with $a_{ij} > 0$ one can find an ℓ such that $-a_{i\ell} > 4a_{ij}$ and $-a_{\ell j} > 4a_{ij}$, with each ℓ assigned to only one such (i, j), then, by the argument of Sec. 2.5, condition (4.7) will retain its form, except that the sum $\sum_{j,i}$ would run only on the negative a_{ji} and some of these a_{ji} would be replaced by $a_{ji}/2$.

In more extreme cases of essentially positive type operators, however, the above approach shows that some modifications are required. Consider for example one of the most extreme cases of 7-point finite-difference approximation to second-order elliptic equations. This is the approximation to the (actually already parabolic) equation

$$U_{xx} - 2U_{xy} + U_{yy} = f,$$

based on the "wrong" seven points, yielding the difference stencil

$$\begin{pmatrix} 0 & 2 & -1 \\ 2 & -6 & 2 \\ -1 & 2 & 0 \end{pmatrix}. \tag{4.11}$$

In this case the positive form $E = eTAe$ (which can be written in a local form by the procedure of Sec. 2.4, or otherwise) does no longer have terms of the form $b_{ij}(e_j - e_i)2$. It is made solely of terms like

$$b_{\alpha,\beta}(2e_{\alpha,\beta} - e_{\alpha-1,\beta} - e_{\alpha,\beta+1})2$$

or

$$b\prime_{\alpha,\beta}(2e_{\alpha,\beta} - e_{\alpha+1,\beta} - e_{\alpha,\beta-1})2,$$

where (α, β) are the grid (column,row) numbers. This shows that in order to satisfy (4.4) one must use interpolation in which

$$w_{(\alpha,\beta)(\alpha+1,\beta)} = w_{(\alpha,\beta)(\alpha,\beta-1)} \quad \text{and} \quad w_{(\alpha,\beta)(\alpha-1,\beta)} = w_{(\alpha,\beta)(\alpha,\beta+1)}. \tag{4.12}$$

Indeed, using interpolation which does not always satisfy (4.12) yielded in this case bad results [10, Table 1].

4.6 Two-grid analysis with block relaxation

In case of block relaxation, relaxability after the coarse-level correction is obtained if $\hat{R}0 \geq C_0 E0$, where $\hat{R}$ is defined in Theorem 3.5 and $\hat{R}0$ is its value after the coarse-level correction. For this purpose we have the following generalization of Theorem 4.1.

Theorem 4.2. If for any fine-grid vector e there exists a coarse-level ec such that

$$C_0 \sum_J (e - I_c ec)_{J\hat{}} T A_{JJ} (e - I_c ec)_{J\hat{}} \leq eTAe \tag{4.13}$$

then $\hat{R}0 \geq C_0 E0$.

In (4.13) we used $(e - I_c ec)_{J\hat{}}$ to denote the J-th block of $e - I_c ec$. The proof is simply the block version of the proof of Theorem 4.1.

It is easy to see that this theorem implies relaxability for the usual anisotropic operators (such as $\varepsilon\partial_{xx} + \partial_{yy}$), provided the suitable block relaxation (such as y-line Gauss-Seidel relaxation) is employed. Theorem 4.1 shows, however, that block relaxation is *not* needed if interpolation has large weights only along strong connections (e.g., vertically, in case of $\varepsilon\partial_{xx} + \partial_{yy}$).

Another case of block relaxation, sometimes called *collective relaxation*, is used when each block contains a fixed number q of unknown and stand for the values of q different functions at the same grid point. It is then natural to coarsen the problem using the same block structure: The coarse-level vector ec is also made of q-blocks, its K-th block, denoted $\hat{e}_K c$, usually corresponding to one of the fine-level blocks, denoted $\hat{e}_{F(K)}$. The interpolation may then take the form of *block interpolation*

$$\hat{e}_J = \sum_K W_{JK} \hat{e}_K c, \tag{4.14}$$

where each W_{JK} is of course a $q \times q$ matrix. Theorem 4.2 can then be used to derive block interpolation strategies analogous to those in Secs. 4.3 and 4.4. (For a more general approach to treat systems with $q > 1$ unknown functions – see Sec. 5.1)

4.7 Limitations of purely algebraic interpolation

In purely algebraic multigrid algorithm no use is made of any geometric information, including the geometric location of unknowns. (The geometric origin of the problem may still play an implicit role, as in Sec. 4.4.) This may lead to severe losses in convergence rates, as the following examples show. We later discuss how to mend such situations.

Example 4.2. We start with the 3-point one-dimensional Poisson equation, in which case

$$a_{ii} = 2, \qquad a_{i,i+1} = a_{i,i-1} = -1, \qquad \text{otherwise} \quad a_{ij} = 0.$$

The coarse grid points correspond to every other fine grid point, i.e., $e_i c$ represents e_{2i}. The only difference from the usual (geometric) multigrid treatment of this problem is in defining the interpolation weights: We assume interpolation from the right neighbor only:

$$w_{ik} = \begin{cases} 1 & \text{if} \quad i = 2k \quad \text{or} \quad i = 2k-1 \\ 0 & \text{otherwise.} \end{cases} \tag{4.25}$$

Suppose a smooth error $e-$ is given, with (slowly varying) slope s. Its coarse-grid approximation ec, produced by (4.2), minimizes the L_2 norm of the slope of $e - -I_c ec$. But by (4.15) the slope of $I_c ec$ must vanish on half the intervals, so all that can be minimized is the L_2 norm of the slope of $e - -I_c ec$ on the rest of the intervals. The minimum of the latter is actually zero, produced by ec for which

$$e_k c - e_{k-1} c = (e_{-2k-1} - e_{-2k-2}) \approx \tfrac{1}{2}(e_{-2k} - e_{-2k-2})$$

(except at the last interval). Thus, the slope of ec is $s/2$. Hence the slope of $e0 = e - -I_c ec$ is on the average $s - s/2 = s/2$, but it is actually alternating between 0 and s. The subsequent relaxation sweeps will efficiently reduce these oscillations (thereby reducing E from $E0$ to $E0/2$); but, no matter how many sweeps are made (if the grid is sufficiently large), the (average) slope will remain $s/2$. Thus the complete two-grid cycle can have a convergence factor no better than .5, no matter how many relaxation sweeps it employs.

Incidentally, the error $e0$ produced in this example after the coarse-grid correction can serve as an example to an error e for which

$$E = \sum_{i,j}(-a_{ij})(e_i - e_j)2 \ll \sum_i a_{ii}e_i 2,$$

but which nevertheless cannot be reduced by a coarse-grid correction (cf. Sec. 3.5).

The fault in the example above is not exactly the fact that we interpolate from one point only. To stress this point we bring now a modification of this example.

Example 4.3. Take now the usual 5-point two-dimensional Poisson equation, with the coarse grid comprising of the red points (grid points (α, β) such that $\alpha + \beta$ is even). Assume first that the interpolation to each black (i.e., not red) point employs its west and north neighbors only. By an argument similar to the above it can be shown that, ignoring some effects near boundaries, the complete two-grid cycle can reduce the north-west-to-south-east slope of the error at most by a factor .5, no matter how many relaxation sweeps are utilized.

If, on the other hand, the two neighbors used in interpolation are the west and east (or north and south) ones, we get quite a usual cycle, where convergence factors can be as small as we please, depending on using enough relaxation sweeps (as can be verified by mode analysis [6], [9], [8], or by other theories).

Thus, geometry plays here a decisive role: Knowing in what direction to interpolate can dramatically change the convergence rates. This is even more so when multigrid V-cycles are used. Interpolations as in the above examples can be shown to give convergence factors no better than $1 - 21 - L$ per cycle, where L is the number of levels.

5. Concluding Remarks

Except for minor modifications, the above sections of this article have appeared before in a preliminary form [7]. The intention to extend this work has

so far kept it from regular publication. In the meantime results and insights described above have been used as keystones in the development of AMG (see [20]), and also in the applications of multigrid processes to non-PDE grid problems, and in some theoretical studies [19]. Theorem 3.4 above have served to broaden the multi-level outlook (see [6, §1.1]) and will be employed in a general justification for local mode analyses [8]. It is therefore at this point deemed useful to publish the paper even without the intended extensions, replacing the latter by the following remarks.

5.1 AMG for systems

The present author is grateful to a referee for pointing out a mistake (in Sec. 2.4 of [7], here omitted) related to "block positive type" matrices. It is related to the issue of how to develop AMG for matrices arising from discretizing PDE *systems*, i.e., q differential equations in q unknown functions, with $q > 1$. Indeed, since the appearance of [7], new approaches have been developed for such matrices.

It is usually assumed that the AMG solver is supplied with more information than just the matrix. It can be assumed, in particular, that the unknowns and equations come with *labels*, i.e., they are classified so that the solver knows which continuous function or PDE equation is approximated by each of them. Any AMG user can easily supply this classification along with his matrix. (There are also tricks how to disentangle this information by pre-processing, such as pre-relaxation of the homogeneous system with random initial approximations, which almost surely yields different local values for different functions.) A corresponding labelling can then be carried over to coarser levels, and the interpolation can be based on it, with each function being separately interpolated.

As for *relaxation* of PDE systems, a general approach for *geometric* multigrid solvers is described in [6, §3.7], based on the factors of the principal determinant of the system. This can also yield a prescription for the corresponding AMG relaxation. For example, it implies Kaczmarz relaxation in case of the Cauchy-Riemann system. In case of Stokes equations it implies Gauss-Seidel relaxation of the momentum equations and "compounded-Kaczmarz" relaxation of the continuity equation (where the velocities involved in the relaxed continuity equation are changed as in Kaczmarz, and in addition pressures involved in the momentum equations corresponding to these velocities are changed proportionately to their total weights in those equations, with each momentum equation being weighted proportionately to the change in the corresponding velocity). Etc.

5.2 Improved interpolations

The basic difficulty in all AMG algorithms is the construction of sufficiently

good coarse-to-fine interpolation formulae. When the order of interpolation fails to be as high as required by the usual multigrid rules (see [6, §4.3]), convergence rates will substantially deteriorate, as shown by Examples 4.2 and 4.3 above, and as confirmed by many numerical tests. For second-order PDE equations, as well as for many PDE systems of second order equations, the interpolation should be at least second order (i.e., linear interpolation). This is difficult to obtain without geometric information. Various rules added to more recent versions of AMG (the detailed rules in [20, §4.2]) tend, for some geometric reasons, to promote geometric-like coarsening and hence linear interpolation, and their high performance can be traced back to this fact. However, that linear interpolation is obtained very indirectly, and is therefore quite precarious: unlucky choice of point ordering, as well as various small perturbations to the algorithms, may spoil it, on all or at least some levels, causing marked decrease in efficiency.

Also, linearity is not enough. Even linear interpolation gives worse-than-usual convergence unless it is *convex*, i.e., unless each fine-grid point lies, at least approximately, in the convex hull of the coarse-grid points from which its corrections are interpolated. Moreover, higher order equations require higher-than-linear interpolations, which is still much harder to obtain.

A general cure to all these difficulties is to allow AMG the option of using additional geometric information, namely, the geometric location of each unknown. Most AMG users would have that information ready anyway. The algorithm can use it to directly construct linear (and even higher order) interpolations (keeping of course the rule of interpolating only along strong connections). It can also use it to obtain convex interpolation, provided a corresponding convexity rule is already employed at the stage of selecting the coarse-level unknowns. This convexity is highly desired for other reasons, too: it would ensure representation on the coarse grid for boundaries and for internal layers of strong discontinuities.

To reduce complexity of the constructed coarse-level problem, especially when the convexity rule is enforced, it is necessary to define strong connections not only in terms of direct strong couplings, but also in terms of *second order strength*, i.e., counting the total strength through all the length-two connecting chains ($\ell = 2$ in the notation of Sec. 4.1). Such second-order measure of connection strength is already used in some recent AMG codes.

5.3 Improved *E-R* theory. Localization blueprint

As shown in Sec. 4.7, the current theory essentially yields the best possible convergence estimates for purely algebraic algorithms. With the improved geometric interpolation discussed in Sec. 5.2, and in geometric multigrid algorithms in particular, improved convergence can be obtained, and improved theory should correspondingly be developed.

In case of uniform or piecewise uniform grids, local mode analysis yields concrete numerical convergence factors which are precisely attained for sufficiently small mesh-sizes, provided negligible amounts of local relaxation work, near structural singularities, are allowed to be added to the analyzed algorithm [8]. With a similar approach, the present E-R theory can be extended to the case of improved interpolation, and rigorously yield concrete convergence estimates. These estimates, although not as precise as the uniform-grid ones, can still be quite realistic, unlike the estimates with vague, undetermined constants obtained in most multigrid theoretical studies.

Whereas the relaxability-after-coarse-grid-correction theory of Sec. 4 is typically based on Theorem 3.1, a theory with improved convergence, implied by the improved interpolation, should typically be based on Theorem 3.2, that is, on the fact that when relaxation slows down, the final residual square-norm R is small compared with the final error energy E. The next step is now to show that, with suitable order of interpolation, the error energy is then always reduced by a coarse-grid correction to a new value, $\bar{E}$, which satisfies $\bar{E} \leq cR$, with a concrete and realistic c, hence deducing $\bar{E} \ll E$.

To obtain reasonably low constants c for general situations, possibly including various singularities (e.g., re-entrant corners), it is necessary to *localize* the theory. This means to first develop, for each non-singular neighborhood, *local convergence estimates*, i.e., estimates based on the assumption that the local conditions (local coefficients, local arrangement of finite elements, etc.) stretch out to the entire space, without boundaries; and then *justify the localization*, i.e., prove that for sufficiently small meshsizes the true convergence is as good as the worst local convergence, provided some extra local relaxation (costing negligible extra work) is applied near singularities.

To justify the localization in the E-R theory (which is different from the justification in [8]) two steps are needed. The first step is to show that, with the aid of local relaxation, R can always be made uniformly distributed, i.e., $R(N)$ can be made comparable at different neighborhoods N. This is easy to show by the theory of Sec. 3 above (typically Theorem 3.1, or 3.4, being here used again). Denoting by $\tilde{E}$ the value of the error energy after a coarse-grid correction with the function $\tilde{e}c$, where $\tilde{e}c$ is a coarse grid function derived from the old fine-grid error by injection, we have $\bar{E} \leq \tilde{E}$ (cf. (ii) in Sec. 4.1). The second step is then to show that, once R is uniformly distributed, for elliptic equations with sufficiently small meshsizes and away from singularities, the local value of $\tilde{E}$ is only determined by R.

The local convergence estimates themselves are quite straightforward to derive, at least in various interesting cases. For the 5-point Poisson equation on regular grids, for example, simple inequalities lead to the relation $\bar{E} \leq 4R$. (Mode analysis actually yield $\bar{E} \leq 2R$, but this kind of derivation is not extendible to

non-uniform grids and finite elements). This relation, together with Theorem 3.2, show that E is reduced at least by the factor $(\nu+1)-1$ by a two-level cycle with ν relaxation sweeps.

5.4 Partial relaxation

The theory sketched above suggests at least one important practical device that should be used by multigrid algorithms in general, and by AMG algorithms in particular, namely, *local*, or *partial relaxation* (cf. [3, §A.9]). Performance can very much be enhanced by extra relaxation passes restricted to the immediate neighborhood of singularities. For reentrant corners this is demonstrated in [2]. In AMG processing this can be done adaptively, by for example special passes over points exhibiting large normalized residuals (cf. Sec. 3.3).

5.5 Role of geometry

As discussed in Secs. 4.7 and 5.3, geometric information can be used to improve AMG performance. Beyond this, it should be emphasized that even purely algebraic multigrid algorithm depend for their success on the geometrical origin of the problem or on a similar structure of the matrix. It is due to geometric-like structure that the coarse-level matrices generated by AMG stay "local" and their sparsity is suitably checked, that is, the number of non-zero entries per row remains nicely bounded (see (3.10) in [10]).

In case the given (finest-level) matrix is itself dense (or a large fraction of its entries do not vanish) the sparsity on coarser levels is unimportant, and geometric-like structure is not needed. What counts in such matrices is only that they have γ constants (such as γ_0, γ_1, γ_*, $\tilde{\gamma}_0$, $\tilde{\gamma}_1$, $\hat{\gamma}_0$ or $\hat{\gamma}_1$ in Sec. 3) comparable to 1.

Acknowledgement. The author is grateful for the hospitality of the Institute for Computational Studies at Colorado State University, Fort Collins, Colorado, where the original version of this article [7] was written.

References

[1] R.E. Alcouffe, A. Brandt, J.E. Dendy, Jr., and J.W. Painter: The multi-grid methods for the diffusion equation with strongly discontinuous coefficients. *SIAM J. Sci. Stat. Comput.* **2** (1981), 430–454.

[2] D. Bai and A. Brandt: Local mesh refinement multilevel techniques. *SIAM*

J. Sci. Stat. Comp., to appear.

[3] A. Brandt: Multi-level adaptive solutions to boundary-value problems. *Math. Comp.* **31** (1977), 333–390. ICASE Report 76-27.

[4] A. Brandt: Numerical stability and fast solutions of boundary-value problems. *Boundary and Interior Layers – Computational and Asymptotic Methods* (J.J.H. Miller, ed.), Boole Press, Dublin 1980, 29–49.

[5] A. Brandt: Multigrid solutions to steady-state compressible Navier-Stokes equations. *Computing Methods in Applied Sciences and Engineering* **V** (R. Glowinski and J.L. Lions, eds.), North Holland 1982, 407–422.

[6] A. Brandt: *Multigrid Techniques: 1984 Guide With Applications to Fluid Dynamics*. Monograph, Weizmann Institute of Science, February, 1984. Available as GMD-Studie No. 85, from GMD-F1T, Postfach 1240, D-5205, St. Augustin 1, W. Germany. Sections 2 to 16 also appear (except for some recent additions) in [15], pp. 220–313. An expanded version will be published by North-Holland.

[7] A. Brandt: Algebraic multigrid theory: the symmetric case. *Preliminary Proceedings for International Multigrid Conference*, April 6-8, 1983, Copper Mountain, Colorado.

[8] A. Brandt: Rigorous local mode analysis. *Proc. 2nd European Conf. on Multigrid Methods*, to appear.

[9] A. Brandt and N. Dinar: Multigrid solutions to elliptic flow problems. *Numerical Methods for Partial Differential Equations* (S. Parter, ed.), Academic Press 1979, 53–147. ICASE Report 79-15.

[10] A. Brandt, S. McCormick and J. Ruge: Algebraic multigrid (AMG) for automatic multigrid solutions with application to geodetic computations. Report, Institute for Computational Studies, Fort Collins, Colorado, October 1982.

[11] A. Brandt, S. McCormick and J. Ruge: Algebraic multigrid (AMG) for sparse matrix equations. *Sparsing and Its Applications* (D.J. Evans, ed.), Cambridge Univ. Press, Cambridge 1985, 257–284.

[12] A. Brandt and S. Ta'asan: Multigrid solutions to quasi-elliptic schemes. *Progress and Supercomputing in Computational Fluid Dynamics* (E.M. Murman and S.S. Abarbanel, eds.), Birkhäuser, Boston 1985, 235–255.

[13] J.E. Dendy, Jr.: Black box multigrid. *J. Comp. Phys.* **48** (1982), 366–386.

[14] J.E. Dendy, Jr.: Multigrid semi-implicit hydrodynamics revisited. *Large Scale Scientific Computation* (S.V. Porter, ed.), Academic Press 1984.

[15] W. Hackbusch and U. Trottenberg (eds.): *Multigrid Methods*. Springer-Verlag, 1982.

[16] S. Kaczmarz: Angenaherte Auflosung von Systemen Linearer Gleichungen. *Bulletin de l'Académie Polonaise des Sciences et Lettres* **A** (1937), 355–357.

[17] J. Meinguet: Multivariate interpolation at arbitrary points made simple. *J. Applied Math. Phys.* (ZAMP) **30** (1979), 292–304.

[18] P. Meissle: Apriori prediction of roundoff error accumulation in the solution of a super-large geodetic normal equation system. NOAA Professional Paper 12, National Oceanic and Atmospheric Administration, Rockville, Maryland, 1980.

[19] J. Mandel: Étude algébrique d'une méthode multigrille pour quelques problemés de frontiére libre, *Compte Rendus*, t. 298, Série I, no. 18, 1984, 469–472.

[20] J. Ruge and K. Stüben: Efficient Solution of Finite Difference and Finite Element Equations by Algebraic Multigrid (AMG). *Proc. Bristol Multigrid Workshop* (D. Paddon, ed.).

[21] K. Stüben: Algebraic multigrid (AMG): Experiences and comparisons. *Appl. Math. Comp.* **13** (1983), 419–451.

[22] K. Tanabe: Projection methods for solving a singular system of linear equations and its applications. *Numerische Mathematik* **17** (1971), 203–214.

[23] F. de la Vallée Poussin: An accelerated relaxation algorithm for iterative solution of elliptic equations. *SIAM J. Numer. Anal.* **5** (1968), 340–351.

[24] D.M. Young: *Iterative Solution of Large Linear Systems*, Academic Press, New York 1971.