Electronic Transactions on Numerical Analysis.
Volume 10, 2000, pp. 92-114.

ISSN 1068-9613.

ETNA
Kent State University
etna@mcs.kent.edu

BEHAVIOR OF PLANE RELAXATION METHODS AS MULTIGRID SMOOTHERS *

IGNACIO M. LLORENTE† AND N. DUANE MELSON‡

Abstract. This paper contains the first published numerical results and analysis of the behavior of alternating plane relaxation methods as multigrid smoothers for cell-centered grids. The results are very satisfactory: plane smoothers work well in general and their performance improves considerably for strong anisotropies in the right direction because they effectively become exact solvers. In fact, the convergence rate decreases (improves) linearly with increasing anisotropy strength. The methods compared are plane Jacobi with damping, plane Jacobi with partial damping, plane Gauss-Seidel, plane zebra Gauss-Seidel, and line Gauss-Seidel. Based on numerical experiments and local mode analysis, the smoothing factor and cost per cycle of the different methods in the presence of strong anisotropies for Dirichlet boundary conditions are compared. A four-color Gauss-Seidel method is found to have the best numerical and architectural properties of the methods considered in the present work. Although alternating direction plane relaxation schemes are simpler and more robust than other approaches, they are not currently used in industrial and production codes because they require the solution of a two-dimensional problem for each plane in each direction. We verify the theoretical predictions of Thole and Trottenberg that an exact solution of each plane is not necessary; in fact, a single two-dimensional multigrid cycle gives the same result as an exact solution, in much less execution time. As a result, alternating-plane smoothers are found to be highly efficient multigrid smoothers for anisotropic elliptic problems.

Key words. multigrid methods, anisotropic discrete operators, plane implicit methods, robust multigrid.

AMS subject classifications. 65M55.

1. Introduction and motivation. Multiblock structured grids are often used in fluid dynamic applications to capture complex geometries and facilitate parallel implementation without dealing with unstructured grids. Inside each block where high gradients in the solution are present, stretched grids are used to obtain improved discretization accuracy. As a result, the local discrete operator may contain strong anisotropies from both the coefficients of the equation and the highly stretched grid.

Standard multigrid techniques are efficient methods for solving many types of partial differential equations (pde's) due to their optimal complexity (work linearly proportional to the number of unknowns), optimal memory requirement, and good parallel efficiency and scalability in parallel implementations. Although highly efficient multigrid methods have been developed for a wide class of problems governed by pde's, they are underutilized in production and commercial codes. One reason for this is that the high level of efficiency is not maintained in anisotropic problems; i.e., the convergence rate of standard multigrid methods degenerates on problems that have anisotropic discrete operators. There is intensive ongoing research aimed at combining the high efficiency of multigrid with good robustness so that multigrid becomes more widely used in production and/or commercial codes.

Several methods have been proposed in the multigrid literature to deal with anisotropic operators. One popular approach is to use semi-coarsening, where the multigrid coarsening is not applied uniformly to all of the coordinate directions [18, 19, 21, 24]. By coarsening the grid only in the direction of the anisotropy, the coarse grid correction process can effectively

* Received May 25, 1999. Accepted for publication February 1, 2000. Recommended by S. McCormick. This research was supported by the National Aeronautics and Space Administration under NASA Contract No. NAS1-97046 while the first author was in residence at the Institute for Computer Applications in Science and Engineering (ICASE), NASA Langley Research Center, Hampton, VA 23681-2199. The first author was supported by the *Dirección General de Enseñanza Superior* of the Spanish Government.

†Departamento de Arquitectura de Computadores y Automática, Universidad Complutense, 28040 Madrid, Spain (llorente@dacya.ucm.es).

‡NASA Langley Research Center, Hampton, VA 23681-2199 (n.d.melson@larc.nasa.gov).

approximate and correct the error left by point relaxation.

Another approach for dealing with anisotropic problems is to develop and use multigrid block smoothers which can eliminate all high frequency errors in the presence of strong anisotropies [20, 23]. Other intermediate alternatives that combine implicit point, line or plane relaxation with partial and/or full coarsening have been presented in the multigrid literature [16, 5, 4, 28]. A comparison between the alternatives is difficult because there are many performance parameters involved that result in a great variety of numerical and architectural properties. We do not find one method always better than the others; rather, we find that each one can be optimum in certain situations. See ref. [12] for additional information about robust alternatives to deal with anisotropic discrete operators.

For rectangular global domains, global definitions of lines and/or planes exist that span the entire domain. Therefore, semicoarsening combined with block smoothing and full coarsening combined with alternating-plane smoothing can be applied to deal with anisotropies. Both alternatives present similar convergence rates for weak anisotropies. However, the alternating-plane smoothing process becomes an exact solver for strong anisotropies where the convergence factor tends to zero, as will be shown in this paper.

On the other hand, for the complex multiblock grids that are often used in computational fluid dynamics (CFD), the union of the subdomains does not result in a logically rectangular global domain. Therefore, global definitions of lines and/or planes do not exist and semicoarsening cannot be readily applied. Plane-implicit smoothers can still be applied but only within the current block.

The objective of this work is to study the behavior of traditional plane relaxation methods as smoothers for the multigrid solution using full coarsening of these discrete anisotropic operators. The model problem under study is the solution of the anisotropic elliptic model equation on a cell-centered grid, described in section 2, by a *full approximation scheme* (FAS) [2, 17, 26], described in section 3. This multigrid method is more involved than the simpler correction scheme, but can be applied to solve nonlinear equations.

Plane smoothers are considered in the multigrid literature to have poor numerical and parallel properties because of the expensive and parallel-inefficient solution of a large number of two-dimensional (2-D) problems. However, we demonstrate in section 4 that an exact solution of the planes is not needed and that one 2-D multigrid cycle is sufficient, which considerably reduces the execution time of a three-dimensional (3-D) smoothing sweep. This is similar to the result shown in [23] by Thole and Trottenberg for vertex-centered grids. Other authors have also used plane implicit methods as multigrid smoothers with approximated solutions for the planes; for example, a preconditioned GMRES method is used as the plane solver in [20] by Oosterlee.

On the other hand, the solution of each plane by a 2-D multigrid cycle involves the solution of a very large number of tridiagonal systems of equations (in a general case, band structured with constant bandwidth). This is not a problem in a sequential setting because very efficient band solvers exist. A single multigrid cycle could also be used as a 1-D solver. However, in a fine-grain parallel setting, the tridiagonal systems may be distributed across many processors, which leads to a high volume of interprocessor communication. The improvement of the parallel efficiency of tridiagonal solvers has been a focus of intensive research in the last few years, see for example [7, 6, 1, 15, 10].

Section 5 presents analytical formulae for the smoothing factors of some plane relaxation smoothers with periodic and Dirichlet boundary conditions. The analytical expressions are verified with several numerical experiments with Dirichlet boundary conditions and cell-centered grids. The formulae provide an accurate prediction of the numerical results. The dependence of the convergence rate on the strength of the anisotropy for the model problem on vertex-centered grids has been previously studied for the 2-D case, for example by Wesseling [26], and observed for the 3-D case with zebra Gauss-Seidel by Thole and Trottenberg [23].

Section 6 compares the plane smoothers, not considering the convergence rate per cycle but the convergence rate per work unit in order to include the work per cycle and the architectural properties of the smoothers in the study. The numerical results show that zebra Gauss-Seidel does not perform as well as expected; in fact, the standard (lexicographic order) Gauss-Seidel method obtains better convergence rates. This seems to contradict the results presented by Thole and Trottenberg in ref. [23] and by Yavneh in ref. [27] for vertex-centered grids. However, the deterioration of the smoothing factor of Gauss-Seidel with odd-even ordering on cell-centered grids has been reported previously by Gjesdal [8] for the 2-D case.

Grid stretching is commonly used in CFD grid generation to pack points in regions with large solution gradients while avoiding an excess of points in more benign regions. Stretching causes the strength of the anisotropies to be different in each cell. Section 7 studies the performance of alternating-direction plane methods as robust smoothers for these problems.

2. The model problem.

Consider the following anisotropic elliptic partial differential equation:

$$a\frac{\partial^2 u(x,y,z)}{\partial x^2} + b\frac{\partial^2 u(x,y,z)}{\partial y^2} + c\frac{\partial^2 u(x,y,z)}{\partial z^2} = f(x,y,z). \tag{2.1}$$

where a, b, and c are constants. This anisotropic Poisson-like equation is solved on a 3-D rectangular domain $\Omega \subset \Re^3$ with some suitable boundary conditions.

In cell-centered finite volume discretizations, the domain Ω is divided into cells and the conserved quantities are stored at the centers of these cells. Cell-centered grids have been widely used in CFD for the finite volume solution of the incompressible and compressible Navier-Stokes equations. We focus our attention to cell-centered grids because our goal is to study the behavior of the smoothers in a CFD setting.

The following difference equations (normalized by the z-direction coefficient), involving algebraic relationships between grid points, are obtained via a discretization of Eq. (2.1) on a uniform cell-centered computational grid by a finite volume technique:

$$\begin{aligned} &\epsilon_1(u_{i-1,j,k} - 2u_{i,j,k} + u_{i+1,j,k}) + \\ &\epsilon_2(u_{i,j-1,k} - 2u_{i,j,k} + u_{i,j+1,k}) + \\ &(u_{i,j,k-1} - 2u_{i,j,k} + u_{i,j,k+1}) = f'_{i,j,k} \end{aligned} \tag{2.2}$$

where $i = 1, ..., n_1, j = 1, ..., n_2, k = 1, ..., n_3$, $u_{i,j,k}$ and $f_{i,j,k}$ are the respective scaled discrete approximations to $u(x,y,z)$ and $f(x,y,z)$ on cell $(i-\frac{1}{2}, j-\frac{1}{2}, k-\frac{1}{2})$, and ϵ_1 and ϵ_2 are the strengths of the anisotropies determined by the coefficients in Eq. (2.1) and the mesh ratios hx, hy and hz used in the discretization.

The experimental convergence factors presented in the paper have been obtained for a homogeneous problem with right-hand side $f(x,y,z) = 0$, boundary condition $u(x,y,z) = 0$ and starting with a random initial guess.

The structured grids used in the present work allow a relatively easy sequential and parallel implementation using, for example, distribution strategies supported by the current versions of High Performance Fortran (HPF). Furthermore, parallel implementation and cache memory exploitation is possible due to the regular data structures in the structured grids.

If we consider the relative size of the resulting terms in the three coordinate directions, we see that there are several possible scenarios for a given cell:

ETNA
Kent State University
etna@mcs.kent.edu

- all three terms are relatively equal (isotropic with no directions dominant)
- one term is relatively larger than the other two terms (anisotropic with one direction dominant)
- two terms are relatively larger than the third term (anisotropic with two directions dominant)

3. Robust multigrid methods. The multigrid technique has many important advantages from the computational point of view. A well designed multigrid method has a computational complexity of $O(N)$, where N is the number of equations in the system, to achieve a numerical solution to the level of truncation error [2, 3, 17, 26]. Moreover, these multigrid methods offer very good parallel efficiencies and scalability on parallel computers [11, 13, 14].

In cell-centered coarsening, a coarse grid is typically obtained by taking unions of fine grid cells. Consequently, by applying cell-centered coarsening, a sequence Ω^l for $l = 1, 2, ..., L$ of increasingly coarser grids is obtained. We used volume weighted summation for the restriction operator R (steps 3 and 4). Trilinear interpolation in the computational space was used for the prolongation operator P (step 7). It is convenient to write (and program) a scheme in delta form where the scheme is written as a steady-state residual plus a delta or update at any point previously updated in the given iteration sweep. The advantage of treating a scheme in delta form is that it is easier to tell what values are 'new' and 'old' which makes the code easier to change and/or debug. The biggest disadvantage of the delta form is that it may be computationally less efficient.

The following iterative algorithm represents an FAS V(γ_1,γ_2)-cycle to solve the system $Lu = f$ where $\Omega^1 = \Omega$.

$u \leftarrow$ FAS(L, u, f):

step 1: Application of γ_1 sweeps of the smoothing method on $L^1u^1 = f^1$

RESTRICTION PART
for l = 2 to L
step 2: Computation of residual $r^{l-1} = f^{l-1} - L^{l-1}u^{l-1}$
step 3: Restriction of residual $r^l = R^{l-1}r^{l-1}$
step 4: Restriction of current approximation $u^l_{\text{old}} = R^{l-1}u^{l-1}$
step 5: Computation of right-hand side $f^l = r^l + L^l u^l_{\text{old}}$
step 6: Application of γ_1 sweeps of the smoothing method on $L^l u^l = f^l$

PROLONGATION PART
for l = L-1 to 1
step 7: Correction of current approximation $u^l = u^l - P^l(u^{l+1} - u^{l+1}_{\text{old}})$
step 8: Application of γ_2 sweeps of the smoothing method on $L^l u^l = f^l$

Plane-implicit smoothers (steps 6 and 8) require the solution of a large number of 2-D boundary value problems (Fig. 1). For example, an (x,y)-plane smoother requires the solution of n_3 problems ($K = 1, ..., n_3$) given by

$$\epsilon_1(u_{i-1,j,K} - 2u_{i,j,K} + u_{i+1,j,K})+$$

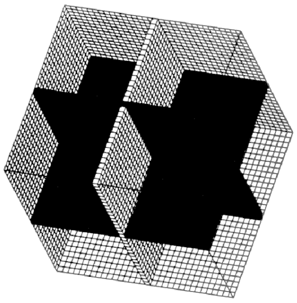

FIG. 1. *Alternating-plane implicit smoother.*

$$\epsilon_2(u_{i,j-1,K} - 2u_{i,j,K} + u_{i,j+1,K}) + (-2u_{i,j,K}) = f^{2-D}_{i,j,K}$$
$$i = 1, ..., n_1, j = 1, ..., n_2$$

where $f^{2-D}_{i,j,K}$ depends on the relaxation method. The resulting 2-D problems are more favorable because the systems may have more diagonal dominance than the original 3-D system. The 2-D problems can also be solved efficiently by using 2-D FAS cycles (section 4).

Notice that the grids visited in the 2-D coarsening are different from the grids used for the 3-D multigrid. Therefore, the grid metrics for the grid hierarchy to solve the planes does not correspond to the grid metrics for the 3-D grid hierarchy. To precompute all the metrics for the 2-D grid hierarchy would significantly increase the memory requirements of the multigrid code. In fact, the required memory for a 3-D cycle with a pointwise smoother grows in proportion to $\frac{8}{7}N$, but, to precalculate all the metrics of a planewise smoother, it would grow as $\frac{4}{3}N$, which is about 52% larger.

Due to the considerable improvement of convergence rate achieved by multigrid methods, the solution of pde's is moving from *time-critical applications* to *accuracy-critical applications* [14]. In these kinds of applications, memory usage is the limiting factor for solving larger problems. Savings in computing time are just used to solve even larger problems. Consequently, it is important to maintain the memory requirements of pointwise smoothers when implementing planewise smoothers. Therefore, the present scheme is coded using just one 2-D multigrid data structure and recomputing each 2-D system of equations every time a plane is visited. This implementation maintains the memory requirements of the original pointwise 3-D multigrid cycle, but increases the execution time. However, because the memory requirements are the same, the performance of the plane smoother can be more accurately compared with the performance of point and line smoothers. This implementation alternative also improves the data locality (temporal and spatial) of the algorithm because the same 2-D multigrid structure is used to solve each plane. The temporal locality is improved because the 2-D data structure fits in cache and the program uses the same data storage addresses when solving each of the planes. The spatial locality of the data is improved in this implementation because data are contiguously stored in memory, whereas the use of a global 2-D multigrid hierarchy to store all planes would present different memory access strides depending on the orientation of the plane.

4. The inner 2-D multigrid solver. Although plane relaxation is simpler and more robust than a semicoarsening approach, it has not been widely applied in practical situations because it requires the solution of a 2-D problem for each plane in each smoothing sweep. However, an exact solution for the 2-D problems is not needed. We show that an approximate solution is sufficient and can be obtained by applying just one 2-D multigrid cycle. This behavior has been also reported by Thole and Trottenberg [23] and by Oosterlee [20] for vertex-centered grids. Here we study this behaviour in cell-centered grids in a little depth.

Table 4.1 compares the experimental 3-D convergence rates for just one 2-D cycle in each plane using two cycling strategies (V(1,0) and V(1,1)) and for a 2-D "exact" solver (four 2-D V(2,1)-cycles). The results are very interesting. Using a single 2-D V(1,1)-cycle, we obtain essentially the same convergence factors as with the "exact" solver, but with significantly less computational work.

The amplification factor of a 2-D multigrid cycle can be approximated by the smoothing analysis of the relaxation scheme. Observe that we require a line relaxation smoother for the 2-D cycles because there may be an anisotropy in each plane; for example, to solve the (*x*,*y*)-planes we need a *y*-line smoother when $\epsilon_1 \leq \epsilon_2$ and an *x*-line smoother otherwise. (For robustness in the general variable coefficient case, where $\epsilon_1 \leq \epsilon_2$ in part of the plane and $\epsilon_1 \geq \epsilon_2$ in part of the plane, an alternating line scheme can be used.)

		Cycling strategies of the 2-D cycle		
ϵ_1	ϵ_2	V(1,0)	V(1,1)	Exact
1	1	0.45	0.34	0.34
1	10^2	0.27	0.25	0.25
1	10^4	1.8×10^{-2}	1.0×10^{-2}	1.0×10^{-2}
1	10^6	1.9×10^{-4}	1.0×10^{-4}	1.0×10^{-4}
1	10^8	1.9×10^{-6}	1.0×10^{-6}	1.0×10^{-6}

TABLE 4.1

Computational convergence factors, ρ_e, of one 3-D V(1,0)-cycle with (x,y)-plane Gauss-Seidel for different ϵ_2 and $\epsilon_1 = 1$ solving the planes with one 2-D V(1,0)-cycle, one 2-D V(1,1)-cycle, and an "exact" solver (four 2-D V(2,1)-cycles).

The behavior of the approximated smoothing method can be analyzed as a perturbation of the exact method. When γ_{2-D} multigrid cycles with an amplification factor of $\overline{\rho}_{2-D}$ are used to solve each plane, the smoothing factor, $\overline{\rho}_a$, of a plane relaxation smoother can be approximated by

$$\overline{\rho}_a \approx \overline{\rho} + (\overline{\rho}_{2-D})^{\gamma_{2-D}}, \tag{4.1}$$

where $\overline{\rho}$ is the smoothing factor of the plane smoother with an exact 2-D solver.

The smoothing factor of a 2-D line relaxation method is quite similar to its corresponding 3-D plane version ($\overline{\rho} \approx \overline{\rho}_{2-D}$)[12]. Consequently, if we use the same block relaxation method in 2-D and 3-D cycles, the 3-D smoothing factor with one 2-D V(1,0)-cycle is given by

$$\overline{\rho}_a \approx \overline{\rho} + (\overline{\rho}_{2-D}) \approx 2\overline{\rho},$$

and with one 2-D V(1,1)-cycle it could be approximated by

$$\overline{\rho}_a \approx \overline{\rho} + (\overline{\rho}_{2-D})^2 \approx \overline{\rho}.$$

Observe that the decrease of the convergence rate for strong anisotropies is also present in the 2-D smoother because ϵ_1 remains fixed and ϵ_2 increases, so the 2-D problem solved in each plane is also anisotropic. However, if both anisotropy values are increased together, the 3-D problem is anisotropic. The 2-D problem does not present an anisotropy, and so the convergence rate of the 2-D problem remains fixed, becoming the dominant term in (4.1):

$$\overline{\rho}_a \approx \overline{\rho} + (\overline{\rho}_{2-D})^{\gamma_{2-D}} \approx \overline{\rho}_{2-D}^{\gamma_{2-D}}.$$

		Cycling strategies of the 2-D cycle		
ϵ_1	ϵ_2	V(1,0)	V(1,1)	Exact
1	1	0.45	0.34	0.34
10^2	10^2	0.31	0.14	0.20
10^4	10^4	0.34	0.12	5.0×10^{-3}
10^6	10^6	0.34	0.12	5.1×10^{-5}
10^8	10^8	0.34	0.12	5.1×10^{-7}

TABLE 4.2

Computational convergence factors, ρ_e, of one 3-D V(1,0)-cycle with (x,y)-plane Gauss-Seidel for different ϵ_2 and ϵ_1 solving the planes with one 2-D V(1,0)-cycle, one 2-D V(1,1)-cycle, and an "exact" solver (four 2-D V(2,1)-cycles).

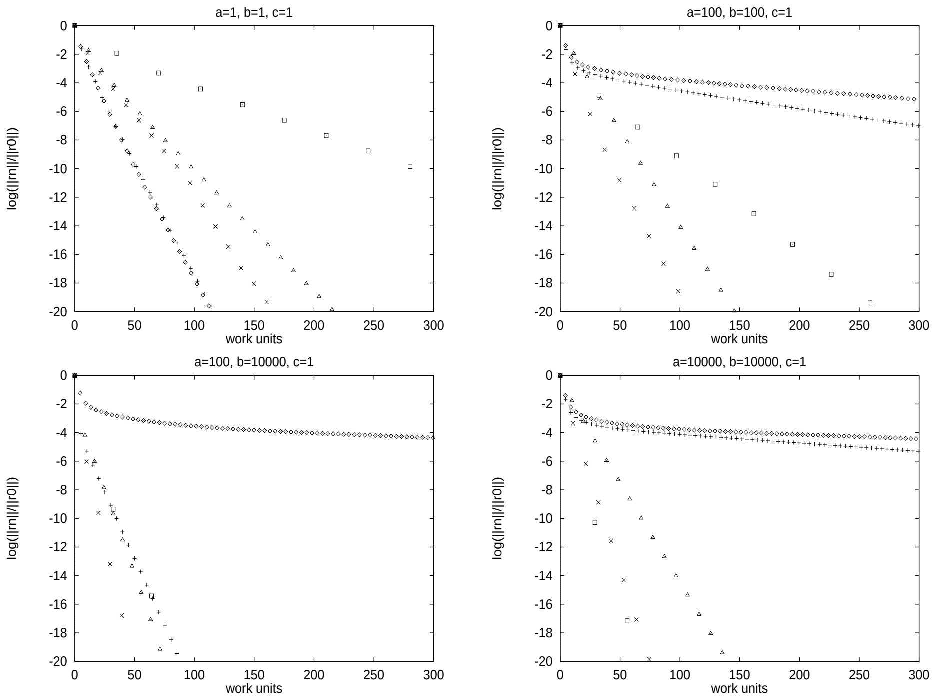

FIG. 2. *Residual versus work units for 3-D V(2,1)-cycles to solve four anisotropic equations on a 32×32×32 uniform grid with the following five smoothers: point-wise Gauss-Seidel (rhombus), y-line Gauss-Seidel (plus), (x,y)-plane Gauss-Seidel with "exact" solver (four 2-D V(2,1)-cycles) in each plane (square), (x,y)-plane Gauss-Seidel with one 2-D V(1,1)-cycle in each plane (little "x"), and (x,y)-plane Gauss-Seidel with one 2-D V(1,0)-cycle in each plane (triangle). The smoother used in the 2-D cycles is y-line Gauss-Seidel. Each symbol represents a 3-D cycle in order to compare the complexity of the cycles of different smoothers.*

This behavior is illustrated by Table 4.2. These results may suggest that the "exact" solver outperforms the approximate solver when both anisotropies are strong. However, we must take into account that one smoothing sweep with the "exact" solver is considerably more expensive than a smoothing sweep with an approximate solver, and so the overall efficiency can be better for the approximate method.

Fig. 2 shows the residual versus work units for 3-D V(2,1)-cycles to solve four anisotropic equations on a $32\times32\times32$ uniform grid with five smoothers:

- point-wise Gauss-Seidel (rhombus)
- y-line Gauss-Seidel (plus)

- (x,y)-plane Gauss-Seidel with "exact" solver in each plane (square)
- (x,y)-plane Gauss-Seidel with one 2-D V(1,1)-cycle in each plane (little "x")
- (x,y)-plane Gauss-Seidel with one 2-D V(1,0)-cycle in each plane (triangle)

These results allow a comparison of the performance of the plane approximate solvers with that of the plane "exact" solvers and with the behaviors of point and line smoothers. The smoother used in the 2-D cycles to solve the planes is y-line Gauss-Seidel. Each symbol is drawn at the completion of a 3-D multigrid cycle to compare the computational complexity of the cycles corresponding to different smoothers. Here a work unit is conservatively defined as the computer time consumed in a residual computation on the finest grid. The time needed to perform one pointwise iteration on the finest grid is about two work units.

As indicated in Fig. 2, the approximate plane solution version with one 2-D V(1,1)-cycle (little "x") is more efficient than the approximate version with one V(1,0)-cycle (triangle) and the "exact" version (square). Even when the 2-D problem solved in each plane is isotropic and the remaining direction is effectively decoupled ($a = b = 10000$, $c = 1$), it is not worth using the "exact" solver because each 3-D cycle consumes too much time. The behavior of the plane smoother for strong anisotropies is very good. The residual reduction per work unit or cycle increases as the anisotropy gets stronger. The two graphics at the bottom of the figure show similar residual reductions per work unit. In fact, when $\epsilon_2 = 10^4$, the solution is achieved to the level of truncation error in just two 3-D cycles (about 50 work units).

It is illustrative to study the behavior of point and line Gauss-Seidel for these cases. Both smoothers exhibit similar performance for the isotropic problem. The work per cycle is slightly greater in the line version; however, the asymptotic convergence rate of the line version is 1.42 times lower (better). The plane smoother, on the other hand, exhibits a less efficient behavior for the isotropic case. The convergence rate of the plane smoother is twice as slow as that of the point smoother. The plane smoother reduces the error more per multigrid cycle than the line smoother does, but its operation count is so much higher that it ends up being slower (for the isotropic case). The computer program used to generate the present results is not fully optimized. It was coded to deal with many different methods and situations, so the plane smoothers results may improve some with coding practice.

As was expected, the point smoother gives very poor convergence rates for aniso-tropic problems. However, the line smoother performance is very fast when $\epsilon_2 \gg \epsilon_1$, as will be demonstrated in section 5.5. This is because a single direction dominates the solution in this case.

5. Smoothing factors of plane relaxation methods. Fig. 2 shows an improvement of the smoothing properties of the plane relaxation methods with increasing anisotropy strength. After a brief description of the Fourier smoothing analysis, details of this behavior are given.

A simple and convenient tool to study the smoothing properties of a relaxation method is Fourier smoothing analysis. The results of this analysis give a measure of the quality of a numerical method. Fourier smoothing analysis does not consider the intergrid transfer process and the discrepancy between the coarse grid and fine grid discrete approximations of the operator, so the actual numerical performance can vary slightly from the predicted performance. To obtain a more accurate prediction, a two-level analysis, which takes into account the operations and discrepancies between levels, must be applied. Results of Fourier analysis can be found in the literature; for further examples, see refs. [22, 25, 26] for the 2-D case and ref. [27] for the 3-D case.

In spite of the fact that Fourier analysis gives the same results for vertex-centered and cell-centered grids, we have noticed differences between them with zebra plane Gauss-Seidel in our numerical experiments. We show how the lexicographic order performs significantly better than the zebra ordering. These differences have also been reported by Gjesdal [8] for

ETNA
Kent State University
etna@mcs.kent.edu

the 2-D case. They may be caused by the coarse grid correction and may be reflected in a two-level analysis.

The methodology and notations followed in this section are similar to those used by Wesseling [26]. Only the results of the analysis are presented; a more detailed study can be found in ref. [12]. The finest grid used in the numerical experiments of this section has 32 points in each direction, and all levels in the grid hierarchy are visited during a cycle. The asymptotic convergence rate monitored in the numerical tests is the asymptotic rate of reduction of the L_2 norm of the residual function per V(1,0)-cycle for the homogeneous problem.

5.1. The plane Jacobi method with damping. Assume $n_1 = n_2 = n_3 = n$ and $\epsilon_1 \leq \epsilon_2$, so the optimum value of the damping parameter, for minimizing the smoothing factor of the (*x*,*y*)-plane Jacobi method applied to Eq. (2.2) is given by

$$\omega = \frac{2 + 2\epsilon_1}{2 + 3\epsilon_1}. \tag{5.1}$$

With this optimum damping parameter, we have

$$\overline{\rho} = \frac{2 + \epsilon_1}{2 + 3\epsilon_1}. \tag{5.2}$$

If $\epsilon_1 < 1$, the optimum smoothing factor tends to 1; for example, $\overline{\rho} = .99$ for $\epsilon_1 = 10^{-2}$. ← Is this correct? Should it be "When $\epsilon_1 \to 0$, the smoothing... Therefore, in such cases (*x*,*y*)-plane relaxation is not a good smoother for our problem and we should use (*x*,*z*)-plane relaxation.

On the other hand, for $\epsilon_1 \geq 1$, the optimum value of ω depends slightly on ϵ_1; we have $\frac{4}{5} \geq \omega \geq \frac{2}{3}$ and $\frac{3}{5} \geq \overline{\rho} \geq \frac{1}{3}$.

For ω lower than the optimum, the smoothing factor is given by

$$\overline{\rho} = 1 - \frac{\omega\epsilon_1}{1 + \epsilon_1}, \tag{5.3}$$

and for ω greater than the optimum,

$$\overline{\rho} = |1 - 2\omega|. \tag{5.4}$$

We exclude from analysis the Fourier modes with $\theta_\alpha = 0$. To do this, we let $\varphi = \frac{2\pi}{n}$ and consider n large enough to approximate $(1 - \cos\varphi)$ by $\frac{2\pi^2}{n^2}$. We obtain the following optimum damping parameter (always considering $\epsilon_1 \leq \epsilon_2$):

$$\omega = 2\left(\frac{\epsilon_1 + \epsilon_2\frac{2\pi^2}{n^2} + \frac{2\pi^2}{n^2}}{1 + \epsilon_1 + \epsilon_2\frac{2\pi^2}{n^2}} + \frac{2 + \epsilon_1\frac{2\pi^2}{n^2} + \epsilon_2\frac{2\pi^2}{n^2}}{1 + \epsilon_1\frac{2\pi^2}{n^2} + \epsilon_2\frac{2\pi^2}{n^2}}\right)^{-1}, \tag{5.5}$$

and the corresponding smoothing factor:

$$\rho_D = \frac{\frac{\epsilon_1+\epsilon_2\frac{2\pi^2}{n^2}+\frac{2\pi^2}{n^2}}{1+\epsilon_1+\epsilon_2\frac{2\pi^2}{n^2}} - \frac{2+\epsilon_1\frac{2\pi^2}{n^2}+\epsilon_2\frac{2\pi^2}{n^2}}{1+\epsilon_1\frac{2\pi^2}{n^2}+\epsilon_2\frac{2\pi^2}{n^2}}}{\frac{\epsilon_1+\epsilon_2\frac{2\pi^2}{n^2}+\frac{2\pi^2}{n^2}}{1+\epsilon_1+\epsilon_2\frac{2\pi^2}{n^2}} + \frac{2+\epsilon_1\frac{2\pi^2}{n^2}+\epsilon_2\frac{2\pi^2}{n^2}}{1+\epsilon_1\frac{2\pi^2}{n^2}+\epsilon_2\frac{2\pi^2}{n^2}}}. \tag{5.6}$$

If $n \gg \epsilon_1, \epsilon_2$, the previous optimum damping and smoothing factors approach Eq. (5.1) and Eq. (5.2), respectively. Otherwise, the optimum damping parameter (5.5) approaches 1 and the optimum amplification factor (5.6) approaches 0 as $O(\frac{1}{\epsilon_1})$ if $n \ll \epsilon_1$, or as $O(\frac{1}{\epsilon_2})$ if

$n \ll \epsilon_2$. Therefore, the plane Jacobi method with $\omega = 1$ is a very good smoother for strong anisotropies with Dirichlet boundary conditions. Observe that this result totally contradicts the result obtained with periodic boundary conditions.

For damping parameters lower than the optimum, the smoothing factor is given by

$$\rho_D = 1 - \omega + \frac{1 - \frac{2\pi^2}{n^2}}{1 + \epsilon_1 + \epsilon_2 \frac{2\pi^2}{n^2}} \omega,$$

and for a damping greater than the optimum, the smoothing factor is given by

$$\rho_D = \left| 1 - \omega + \frac{1}{1 + \epsilon_1 \frac{2\pi^2}{n^2} + \epsilon_2 \frac{2\pi^2}{n^2}} \omega \right|.$$

$\epsilon_1 = \epsilon_2$						
		ω				
ϵ_1	ϵ_2	0.2	0.4	0.6	0.8	1.0
10^{-8}	10^{-8}	0.99	0.99	0.99	0.99	0.98
10^{-2}	10^{-2}	0.98	0.98	0.98	0.98	0.98
1	1	0.90	0.80	0.70	0.59	0.98
10^{2}	10^{2}	0.80	0.62	0.42	0.23	0.50
10^{4}	10^{4}	0.80	0.60	0.40	0.20	1.0×10^{-2}
10^{6}	10^{6}	0.80	0.60	0.40	0.20	1.0×10^{-4}
10^{8}	10^{8}	0.80	0.60	0.40	0.20	1.0×10^{-6}
$\epsilon_1 = 1$ and various ϵ_2						
		ω				
ϵ_1	ϵ_2	0.2	0.4	0.6	0.8	1.0
1	1	0.90	0.80	0.70	0.59	0.98
1	10^{2}	0.81	0.74	0.62	0.50	0.66
1	10^{4}	0.80	0.60	0.40	0.20	1.9×10^{-2}
1	10^{6}	0.80	0.60	0.40	0.20	2.0×10^{-4}
1	10^{8}	0.80	0.60	0.40	0.20	2.0×10^{-6}
$\epsilon_1 = 10^{-4}$ and various ϵ_2						
		ω				
ϵ_1	ϵ_2	0.2	0.4	0.6	0.8	1.0
10^{-4}	10^{-2}	0.99	0.99	0.99	0.99	0.98
10^{-4}	1	0.99	0.99	0.99	0.99	0.98
10^{-4}	10^{2}	0.93	0.87	0.80	0.73	0.66
10^{-4}	10^{4}	0.80	0.60	0.40	0.20	2.0×10^{-2}
10^{-4}	10^{6}	0.80	0.60	0.40	0.20	2.0×10^{-4}
10^{-4}	10^{8}	0.80	0.60	0.40	0.20	2.0×10^{-6}

TABLE 5.1
Computational convergence factors, ρ_e, of one 3-D V(1,0)-cycle with (x,y)-plane Jacobi with damping parameter ω.

Therefore:

- If $n \gg \epsilon_1$ and $n \gg \epsilon_2$, the smoothing factor approaches the periodic case. See, for example, in Table 5.1, how the results for $\epsilon_1 \leq 1$ and $\epsilon_2 \leq 1$ fully agree with results for Eq. (5.1) for the optimum damping parameter and with Eq. (5.3) and (5.4) for the smoothing factor.
- If $n \ll \epsilon_1$ and $n \ll \epsilon_2$, the smoothing factor approaches

$$\rho_D = 1 - \omega + O(\frac{1}{\epsilon_1 + \epsilon_2})\omega$$

ETNA
Kent State University
etna@mcs.kent.edu

as ϵ_1 and ϵ_2 become larger. For $\epsilon_1, \epsilon_2 \geq 10^2$, the top section of Table 5.1 presents some results that verify the above expression. Note that the smoothing factor falls linearly with the strength of the anisotropies when the damping parameter is equal to 1.

- If $n \gg \epsilon_1$ and $n \ll \epsilon_2$, the smoothing factor approaches

$$\rho_D = 1 - \omega + O(\frac{1}{\epsilon_2})\omega$$

as ϵ_2 becomes larger. The middle and lower sections of Table 5.1 give some numerical results that agree with the above expression. Observe that for low values of ϵ_1 (even though $\epsilon_1 \ll 1$) the method is a very good smoother if ϵ_2 is large enough.

We conclude that the periodic case can be considered as an asymptotic limit of the Dirichlet case when n approaches ∞ ($n \gg \epsilon$). However, for practical grid sizes, there is a huge difference between the cases. The Dirichlet case presents very good convergence rates with anisotropy values for which the periodic case does not converge. The multigrid algorithm reaches the solution accurate to the truncation error in just a few cycles for a strong value of just one of the anisotropies.

The behavior of the smoother with Dirichlet boundary conditions is attributed to the fact that, as the anisotropy grows, the method becomes an exact solver and so the error is reduced by a factor of $1 - \omega$. The optimum damping parameter depends strongly on ϵ in this case. One way to avoid this dependency and get good convergence rates for all ω is to apply the damping parameter only to the diagonal component of the method in the explicit direction, which is the plane Jacobi method with partial damping. In this way the method approaches an exact solver for strong anisotropies and all damping parameters.

5.2. The plane Jacobi method with partial damping. Assume $n_1 = n_2 = n_3 = n$ and $\epsilon_1 \leq \epsilon_2$, so the smoothing factor of the (*x*,*y*)-plane Jacobi method with partial damping applied to Eq. (2.2) is given by

$$\overline{\rho} = \max(1 - \omega, \frac{1}{1 + \epsilon_1 \omega}, |1 - 2\omega|). \tag{5.7}$$

If $\epsilon_1 < 1$, the optimum smoothing factor approaches 1; for example, $\overline{\rho} = .99$ for $\epsilon_1 = 10^{-2}$.

If $1 \leq \epsilon_1 \leq 3$, the optimum damping parameter can be obtained by equating the second and the third functions in Eq. (5.7). For example, for $\epsilon_1 = 1$, the optimum damping parameter is 0.78 and the optimum smoothing factor is 0.56. For ω lower than 0.78, the smoothing factor is given by

$$\overline{\rho} = \frac{1}{1 + \omega}, \tag{5.8}$$

and, for a damping parameter greater than 0.78, it is given by

$$\overline{\rho} = |1 - 2\omega|. \tag{5.9}$$

On the other hand, if $\epsilon_1 \geq 3$, the smoothing factor does not depend on the anisotropy and is given by

$$\overline{\rho} = \max(1 - \omega, |1 - 2\omega|),$$

and so the optimum smoothing parameter is $\frac{1}{3}$, corresponding to a damping of $\frac{2}{3}$. For ω lower than the optimum, the smoothing factor is given by

$$\overline{\rho} = 1 - \omega, \tag{5.10}$$

and, for a damping greater than the optimum, it is given by

$$\overline{\rho} = |1 - 2\omega|. \tag{5.11}$$

$\epsilon_1 = \epsilon_2$							
		ω					
ϵ_1	ϵ_2	0.2	0.4	0.6	0.8	0.9	1.0
10^{-8}	10^{-8}	0.99	0.99	0.99	0.99	0.99	0.98
10^{-2}	10^{-2}	0.99	0.99	0.99	0.99	0.99	0.98
1	1	0.81	0.70	0.65	0.55	0.77	0.98
3	3	0.80	0.59	0.41	0.57	0.76	0.95
10	10	0.76	0.56	0.35	0.53	0.71	0.88
10^2	10^2	0.66	0.42	0.27	0.33	0.42	0.5
10^4	10^4	4.6×10^{-2}	1.4×10^{-2}	7.8×10^{-3}	7.6×10^{-3}	9.0×10^{-3}	1.0×10^{-2}
10^6	10^6	5×10^{-4}	1.7×10^{-4}	8×10^{-5}	7.6×10^{-5}	9.1×10^{-5}	9.9×10^{-5}
10^8	10^8	5×10^{-6}	1.7×10^{-6}	8×10^{-7}	7.6×10^{-7}	9.1×10^{-7}	9.9×10^{-7}
$\epsilon_1 = 1$ and various ϵ_2							
		ω					
ϵ_1	ϵ_2	0.2	0.4	0.6	0.8	0.9	1.0
1	1	0.81	0.70	0.65	0.55	0.77	0.98
1	10	0.77	0.66	0.58	0.55	0.72	0.92
1	10^2	0.75	0.61	0.51	0.44	0.53	0.65
1	10^4	1.0×10^{-1}	4.7×10^{-2}	3.2×10^{-2}	2.4×10^{-2}	2.1×10^{-2}	1.9×10^{-2}
1	10^6	1.2×10^{-3}	5.3×10^{-4}	3.4×10^{-4}	2.5×10^{-4}	2.2×10^{-4}	2.0×10^{-4}
1	10^8	1.2×10^{-5}	5.3×10^{-6}	3.4×10^{-6}	2.5×10^{-6}	2.2×10^{-6}	2.0×10^{-6}
$\epsilon_1 = 10^{-4}$ and various ϵ_2							
		ω					
ϵ_1	ϵ_2	0.2	0.4	0.6	0.8	0.9	1.0
10^{-4}	10^{-2}	0.99	0.99	0.99	0.99	0.99	0.98
10^{-4}	1	0.99	0.99	0.99	0.99	0.99	0.98
10^{-4}	10	0.99	0.98	0.97	0.96	0.95	0.98
10^{-4}	10^2	0.91	0.84	0.77	0.72	0.69	0.66
10^{-4}	10^4	1.0×10^{-1}	4.7×10^{-2}	3.2×10^{-2}	2.4×10^{-2}	2.1×10^{-2}	1.9×10^{-2}
10^{-4}	10^6	1.2×10^{-3}	5.3×10^{-4}	3.4×10^{-4}	2.5×10^{-4}	2.2×10^{-4}	2.0×10^{-4}
10^{-4}	10^8	1.2×10^{-5}	5.3×10^{-6}	3.4×10^{-6}	2.5×10^{-6}	2.2×10^{-6}	2.0×10^{-6}

TABLE 5.2
Computational convergence factors, ρ_e, of one 3-D V(1,0)-cycle (x,y)-plane Jacobi with partial damping parameter ω.

We exclude from analysis the Fourier modes with $\theta_\alpha = 0$. Doing so, defining $\varphi = \frac{2\pi}{n}$ and considering that n is large enough, we obtain an expression for the smoothing factor that coincides with Eq. (5.7) if $n \gg \epsilon_1, \epsilon_2$.

Otherwise, if $n \ll \epsilon_1, \epsilon_2$, and always considering $\epsilon_1 \leq \epsilon_2$, the optimum damping factor remains $\frac{2}{3}$ with a corresponding optimum smoothing factor

$$\rho_D = \frac{\frac{1}{3}}{1 + \frac{2}{3}\frac{2\pi^2}{n^2}(\epsilon_1 + \epsilon_2)}.$$

Note that the previous optimum amplification factor of $O(\frac{1}{\epsilon_1+\epsilon_2})$ approaches 0. For damping parameters lower than the optimum, the smoothing parameter is given by

$$\rho_D = \frac{1-\omega}{1 + \omega\epsilon_1\frac{2\pi^2}{n^2} + \omega\epsilon_2\frac{2\pi^2}{n^2}},$$

ETNA
Kent State University
etna@mcs.kent.edu

and, for a damping greater than the optimum, the smoothing factor is given by

$$\rho_D = \left| \frac{1-2\omega}{1+\omega\epsilon_1\frac{2\pi^2}{n^2}+\omega\epsilon_2\frac{2\pi^2}{n^2}} \right|.$$

On the other hand, if ϵ_1 is small enough, the optimum smoothing parameter approaches 1 for increasing values of ϵ_2. For damping parameters lower than the optimum, the smoothing parameter is given by

$$\rho_D = \frac{1-\omega\frac{2\pi^2}{n^2}}{1+\omega\epsilon_1+\omega\epsilon_2\frac{2\pi^2}{n^2}},$$

and, for a damping greater than the optimum, the smoothing factor is given by

$$\rho_D = \left| \frac{1-2\omega}{1+\omega\epsilon_1\frac{2\pi^2}{n^2}+\omega\epsilon_2\frac{2\pi^2}{n^2}} \right|.$$

Therefore:

- If $n \gg \epsilon_1$ and $n \gg \epsilon_2$, the smoothing factor approaches the periodic case. See, for example, how the analytical expressions obtained for the periodic case (Eqs. (5.8) and (5.9) for $\epsilon_1 = 1$ and (5.10) and (5.11) for $\epsilon_1 = 3$) accurately agree with the numerical results presented in the top section of Table 5.2.
- If $n \ll \epsilon_1$ and $n \ll \epsilon_2$, the smoothing factor for a damping lower than the optimum, $\frac{2}{3}$, approaches

  $$\rho_D = \frac{1-\omega}{\omega O(\epsilon_1+\epsilon_2)},$$

  and, for a damping parameter greater than the optimum, it approaches

  $$\rho_D = \left| \frac{1-2\omega}{\omega O(\epsilon_1+\epsilon_2)} \right|$$

  as ϵ_1 and ϵ_2 become larger. The top section of Table 5.2, for $\epsilon_1, \epsilon_2 \geq 10^2$, contains some results that verify that the smoothing factor falls linearly with the anisotropy and is a function of ω.
- If $n \gg \epsilon_1$ and $n \ll \epsilon_2$, the smoothing factor for a damping parameter lower than the optimum approaches

  $$\rho_D = \frac{1}{\omega O(\epsilon_2)}$$

  and, for a damping greater than the optimum, it approaches

  $$\rho_D = \left| \frac{1-2\omega}{\omega O(\epsilon_2)} \right|$$

  as ϵ_2 becomes larger. The optimum damping parameter moves from 0.78 to 1. The middle and lower sections of Table 5.2 give some numerical results that verify the above expressions. This method is a very good smoother for very low values of ϵ_1 when ϵ_2 is large enough.

5.3. The plane Gauss-Seidel method. Assume $n_1 = n_2 = n_3 = n$ and $\epsilon_1 \leq \epsilon_2$. The smoothing factor of the (*x*,*y*)-plane Gauss-Seidel method is then given by

$$\overline{\rho} = \max(\frac{1}{1+2\epsilon_1}, \frac{1}{\sqrt{5}}).$$

Therefore, the smoothing factor is $\frac{1}{\sqrt{5}}$ for $\epsilon_1 \geq \frac{\sqrt{5}-1}{2} \approx 0.6$ and $\frac{1}{1+2\epsilon_1}$ for $\epsilon_1 < 0.6$. For $\epsilon_1 < 0.6$, the smoothing factor approaches 1 so (*x*,*y*)-plane relaxation is not a good smoother; we should use (*x*,*z*)-plane relaxation.

$\epsilon_1 = \epsilon_2$		
ϵ_1	ϵ_2	ρ_e
10^{-8}	10^{-8}	0.99
10^{-2}	10^{-2}	0.96
0.66	0.66	0.43
1	1	0.34
10^{2}	10^{2}	0.20
10^{4}	10^{4}	5.0×10^{-3}
10^{6}	10^{6}	5.1×10^{-5}
10^{8}	10^{8}	5.1×10^{-7}
$\epsilon_1 = 1$ and various ϵ_2		
ϵ_1	ϵ_2	ρ_e
1	1	0.34
1	10^{2}	0.25
1	10^{4}	1.0×10^{-2}
1	10^{6}	1.0×10^{-4}
1	10^{8}	1.0×10^{-6}
$\epsilon_1 = 10^{-4}$ and various ϵ_2		
ϵ_1	ϵ_2	ρ_e
10^{-4}	10^{-2}	0.99
10^{-4}	1	0.97
10^{-4}	10^{2}	0.50
10^{-4}	10^{4}	1.0×10^{-2}
10^{-4}	10^{6}	1.0×10^{-4}
10^{-4}	10^{8}	1.0×10^{-6}

TABLE 5.3
Computational convergence factors, ρ_e, of one 3-D V(1,0)-cycle with (x,y)-plane Gauss-Seidel.

Excluding from analysis the Fourier modes with $\theta_\alpha = 0$, we let $\varphi = \frac{2\pi}{n}$ and, for n large enough, we find an expression for the smoothing factor that coincides with the periodic case if $n \gg \epsilon_1, \epsilon_2$. For example, the top section of Table 5.3 shows that the experimental convergence rate when $\epsilon = 0.66$ is 0.43, which verifies the analytical prediction of the periodic case: $\frac{1}{\sqrt{5}}$.

Otherwise, the smoothing factor is

$$\rho_D = \left| \frac{1}{\sqrt{(2+2\epsilon_1\frac{2\pi^2}{n^2}+2\epsilon_2\frac{2\pi^2}{n^2})^2+1}} \right|$$

when $n \ll \epsilon_1, \epsilon_2$. Note that in this case the smoothing factor approaches 0 as $O(\frac{1}{\epsilon_1+\epsilon_2})$ for strong anisotropies. This dependence of the smoothing factor on the anisotropy is verified numerically by the results presented in Table 5.3.

However, for small values of ϵ_1, the smoothing factor tends to

$$\rho_D = \frac{1}{\sqrt{(1+\frac{2\pi^2}{n^2}+2\epsilon_1+2\epsilon_2\frac{2\pi^2}{n^2})^2+\frac{2\pi^2}{n^2}}}.$$

If ϵ_1 is small enough, ρ_D can be approximated by

$$\rho_D \approx \frac{1}{1+2\epsilon_2\frac{2\pi^2}{n^2}}, \tag{5.12}$$

and the smoothing factor decreases as $O(\frac{1}{\epsilon_2})$ for increasing ϵ_2 values. This behavior is also exhibited by the numerical experiments presented in the middle and lower sections of Table 5.3. Again we find that very good convergence rates can be achieved even though one anisotropy is lower than 1.

The numerical results show that it does not pay to use successive over-relaxation (SOR) ($\omega > 1$) or damped Gauss-Seidel ($\omega < 1$) as a smoother.

5.4. The plane zebra Gauss-Seidel method. The analytical study for this case is more involved because the Fourier modes are not invariant under this method; that is, the zebra ordering does not preserve the modes. However, the study can be performed considering that the operation of an iteration on a mode results in a combination of the mode and its harmonics [26]. In this section, we restrict ourselves to the presentation and discussion of the numerical results.

In reference [27], Yavneh presents results on zebra Gauss-Seidel in all combinations of block and point relaxation with full and partial coarsening for periodic boundary conditions. He indicates that the smoothing factor with r relaxation sweeps for the present case is given by

$$\overline{\rho} = \max\{(\frac{1}{1+\epsilon_1})^2, (\frac{2\mathrm{r}-1}{2\mathrm{r}})^2[\frac{1}{2(2\mathrm{r}-1)}]^{\frac{1}{\mathrm{r}}}\}. \tag{5.13}$$

	Cycling strategies of the 3-D cycle						
	V(1,0)	V(0,1)	V(2,0)	V(0,2)	V(1,1)	V(2,1)	V(1,2)
GS	0.34	0.34	0.14	0.14	0.13	0.08	0.08
ZGS	0.48	0.42	0.20	0.13	0.24	0.17	0.17
4cGS	0.28	0.35	0.14	0.14	0.12	0.10	0.10

TABLE 5.4

Computational convergence factors, ρ_e, of 3-D cycles for different V-cycle strategies with (x,y)-plane Gauss-Seidel (GS), (x,y)-plane zebra Gauss-Seidel (ZGS), and (x,y)-plane four-color Gauss-Seidel (4cGS) as smoothers.

The numerical results presented in Table 5.4 diverge considerably from those that we expected because the lexicographic ordering performs better than the zebra ordering, and the result for the isotropic case is larger than the 0.25 predicted by Eq. (5.13). Occurring in cell-centered grids in relation to vertex-centered grids, the translation of the grid points by a half space step does not affect the Fourier analysis, but it could affect the behavior of the zebra method in the coarse grid correction, as is shown in ref. [8] for the 2-D case. Based on the good results with four-color ordering reported in [8] for the 2-D case, we applied this ordering and obtained more robustness. (See Table 5.4.)

Table 5.5 shows the behavior of the four-color and zebra Gauss-Seidel methods used as smoothers. The methods have good behavior for strong values of the anisotropy where

$\epsilon_1 = \epsilon_2$			
ϵ_1	ϵ_2	ZGS	4cGS
10^{-8}	10^{-2}	0.99	0.99
10^{-2}	10^{-2}	0.96	0.96
1	1	0.48	0.28
10^{2}	10^{2}	0.15	0.13
10^{4}	10^{4}	7.2×10^{-4}	2.2×10^{-3}
10^{6}	10^{6}	7.3×10^{-7}	2.2×10^{-5}
10^{8}	10^{8}	9.3×10^{-10}	2.2×10^{-7}
$\epsilon_1 = 1$ and various ϵ_2			
ϵ_1	ϵ_2	ZGS	4cGS
1	1	0.48	0.28
1	10^{2}	0.22	0.16
1	10^{4}	2.0×10^{-3}	4.6×10^{-3}
1	10^{6}	2.0×10^{-6}	4.9×10^{-5}
1	10^{8}	2.0×10^{-9}	4.9×10^{-7}
$\epsilon_1 = 10^{-4}$ and various ϵ_2			
ϵ_1	ϵ_2	ZGS	4cGS
10^{-4}	10^{-2}	0.99	0.99
10^{-4}	1	0.97	0.97
10^{-4}	10^{2}	0.45	0.46
10^{-4}	10^{4}	2.0×10^{-3}	4.6×10^{-3}
10^{-4}	10^{6}	2.0×10^{-6}	4.9×10^{-5}
10^{-4}	10^{8}	2.0×10^{-9}	4.9×10^{-7}

TABLE 5.5
Computational convergence factors, ρ_e, of one 3-D V(1,0)-cycle with (x,y)-plane zebra Gauss-Seidel (ZGS), and (x,y)-plane four-color Gauss-Seidel (4cGS).

the convergence rate of $O(\frac{1}{\epsilon_1+\epsilon_2})$ decreases (improves) linearly with the strength of the anisotropy. The four-color ordering presents convergence rates similar to the lexicographic ordering and parallelizes easily, so that it is an attractive smoother.

5.5. The line Gauss-Seidel method. We also include line Gauss-Seidel in this analysis because of its good behavior observed for strong anisotropies in a single direction. Its performance improves considerably with Dirichlet boundary conditions when one of the anisotropies is stronger than the other (i.e., when one direction dominates). The application of Fourier analysis to study the smoothing properties of *y*-line Gauss-Seidel as applied to Eq. (2.2) is very cumbersome; instead, we present numerical results and some explicit formulae for the smoothing factor obtained by studying the behavior of the numerical results. The formulae are meant to explain the qualitative behavior of the method, and are not necessarily quantitatively accurate.

If $\epsilon_1 > 1$, the smoothing factor can be approximated by

$$\rho_D \approx \frac{\epsilon_1}{1+\epsilon_1+\epsilon_2\frac{2\pi^2}{n^2}},$$

which is equal to 0.5 for $\epsilon_1 = 1$ and approaches 1 for large values of ϵ_1; this behavior is observed in the results on the top section of Table 5.6. If $\epsilon_2 \gg \epsilon_1 > 1$, the previous expression can be approximated by

$$\rho_D \approx \frac{1}{1+\frac{\epsilon_2}{\epsilon_1}\frac{2\pi^2}{n^2}}, \tag{5.14}$$

ETNA
Kent State University
etna@mcs.kent.edu

$\epsilon_1 = \epsilon_2$			$\epsilon_1 = 1$ and various ϵ_2		
ϵ_1	ϵ_2	ρ_e	ϵ_1	ϵ_2	ρ_e
10^{-8}	10^{-8}	0.99	1	1	0.48
10^{-2}	10^{-2}	0.96	1	10^2	0.31
1	1	0.48	1	10^4	1.8×10^{-2}
10^2	10^2	0.95	1	10^6	1.9×10^{-4}
10^4	10^4	0.98	1	10^8	1.9×10^{-6}
10^6	10^6	0.99			
10^8	10^8	0.99			
$\epsilon_1 = 10^2$ and various ϵ_2			$\epsilon_1 = 10^4$ and various ϵ_2		
ϵ_1	ϵ_2	ρ_e	ϵ_1	ϵ_2	ρ_e
10^2	1	0.96	10^4	10^2	0.99
10^2	10^2	0.95	10^4	10^4	0.98
10^2	10^4	0.49	10^4	10^6	0.50
10^2	10^6	1.0×10^{-2}	10^4	10^8	1.0×10^{-2}
10^2	10^8	1.0×10^{-4}	10^4	10^{10}	1.0×10^{-4}
$\epsilon_1 = 10^{-4}$ and various ϵ_2			$\epsilon_1 = 10^{-2}$ and various ϵ_2		
ϵ_1	ϵ_2	ρ_e	ϵ_1	ϵ_2	ρ_e
10^{-4}	10^{-2}	0.99	10^{-2}	10^{-2}	0.96
10^{-4}	1	0.97	10^{-2}	1	0.95
10^{-4}	10^2	0.50	10^{-2}	10^2	0.50
10^{-4}	10^4	1.0×10^{-2}	10^{-2}	10^4	1.0×10^{-2}
10^{-4}	10^6	1.0×10^{-4}	10^{-2}	10^6	1.0×10^{-4}
10^{-4}	10^8	1.0×10^{-6}	10^{-2}	10^8	1.0×10^{-6}

TABLE 5.6
Computational convergence factors, ρ_e, of one 3-D V(1,0)-cycle with y-line Gauss-Seidel.

and the method presents good convergence rates. (See the second section of Table 5.6.) Eq. (5.14) shows that the smoothing factor depends on the ratio $\frac{\epsilon_2}{\epsilon_1}$. The results presented in the lower section of Table 5.6 also show this dependence.

On the other hand, if $\epsilon_1 < 1$, the smoothing factor can be approximated by

$$\rho_D \approx \frac{1}{1 + \epsilon_1 + \epsilon_2 \frac{2\pi^2}{n^2}},$$

which approaches 1 for decreasing values of ϵ_1. (See the top section of Table 5.6.) However, if $\epsilon_2 \gg n$, the smoothing factor can be approximated by

$$\rho_D \approx \frac{1}{1 + \epsilon_2 \frac{2\pi^2}{n^2}}.$$

In this case, the smoothing factor decreases linearly with ϵ_2 and not with $\frac{\epsilon_2}{\epsilon_1}$. (See the lower section of Table 5.6.)

The expressions for the periodic case are obtained by letting n go to ∞ in the previous expressions. As expected, the convergence rate improves linearly with ϵ_2 when $\epsilon_1 = 1$. However, another very important result is that, when $\epsilon_1 > 1$, the convergence rate grows with $\frac{\epsilon_2}{\epsilon_1}$ and, when $\epsilon_1 < 1$, the convergence rate grows with ϵ_2. Consequently, when one of the anisotropies is stronger than the other, only one term dominates, and the line smoother gives good convergence rates for practical mesh sizes.

6. Comparison of the plane smoothers. In the present section, we consider the convergence rate per work unit to compare the plane smoothers. Fig. 3 shows residual versus

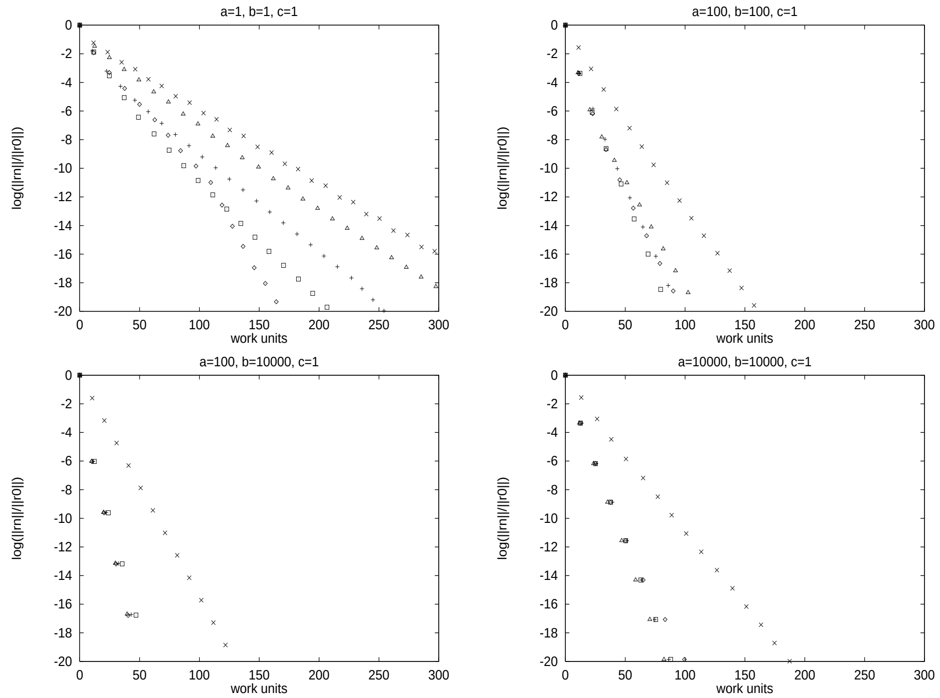

FIG. 3. *Residual versus work units for 3-D V(2,1)-cycles to solve four anisotropic equations on a 32×32×32 uniform grid with the following five (x,y)-plane smoothers: Gauss-Seidel (rhombus), zebra Gauss-Seidel (plus), four-color Gauss-Seidel (square), Jacobi with damping 0.7 (little "x"), and Jacobi with partial damping 0.7 (triangle). One 2-D V(1,1)-cycle is used to solve each plane. The smoother used in the 2-D cycles is the y-line version of the one used in the 3-D cycles.*

work units for 3-D V(2,1)-cycles to solve four anisotropic equations on a 32×32×32 uniform grid with the following five (x,y)-plane smoothers:

- Gauss-Seidel (rhombus)
- zebra Gauss-Seidel (plus)
- four-color Gauss-Seidel (square)
- Jacobi with damping 0.7 (little "x")
- Jacobi with partial damping 0.7 (triangle)

The best results are consistently obtained with the four-color plane Gauss-Seidel method. In general, the three Gauss-Seidel plane implicit methods and the Jacobi plane implicit method with partial damping give similar results with anisotropic equations. Jacobi with damping performs worse because its smoothing factor does not improve with the anisotropy. However, four-color Gauss-Seidel performs better in the isotropic case. Regarding the parallel implementation, zebra Gauss-Seidel, four-color Gauss-Seidel, and Jacobi methods are fully parallelizable; however, the Jacobi method is likely to give better parallel efficiencies because of its coarser granularity. On the other hand, for the Jacobi method the improvement of the convergence rate for strong anisotropies deteriorates slightly for increasing mesh sizes [12].

7. Isotropic coefficient equation on stretched grids. Grid stretching is commonly used in CFD. Each cell can have different aspect ratios and so the discretization of Eq. (2.1) is given by the following general discrete operator:

ETNA
Kent State University
etna@mcs.kent.edu

$$
\begin{aligned}
&\frac{2a}{hx_i}\left(\frac{u_{i-1,j,k}}{hx_i+hx_{i-1}}-\left(\frac{u_{i,j,k}}{hx_i+hx_{i+1}}+\frac{u_{i,j,k}}{hx_i+hx_{i-1}}\right)+\frac{u_{i+1,j,k}}{hx_i+hx_{i+1}}\right)+\\
&\frac{2b}{hy_j}\left(\frac{u_{i,j-1,k}}{hy_j+hy_{j-1}}-\left(\frac{u_{i,j,k}}{hy_j+hy_{j+1}}+\frac{1}{hy_j+hy_{j-1}}\right)+\frac{u_{i,j+1,k}}{hy_j+hy_{j+1}}\right)+\\
&\frac{2c}{hz_k}\left(\frac{u_{i,j,k-1}}{hz_k+hz_{k-1}}-\left(\frac{u_{i,j,k}}{hz_k+hz_{k+1}}+\frac{u_{i,j,k}}{hz_k+hz_{k-1}}\right)+\frac{u_{i,j,k+1}}{hz_k+hz_{k+1}}\right)=\\
&f_{i,j,k},\ i=1,...,n_1, j=1,...,n_2, k=1,...,n_3
\end{aligned}
\tag{7.1}
$$

where Dirichlet boundary conditions are imposed at the boundaries of the domain. Note that Eq. (7.1) includes the case with variable coefficients (different values of a, b, and c in different parts of the computational domain). Varying grid aspect ratios and values of the equation coefficients cause the strength of the anisotropies to be different in each cell.

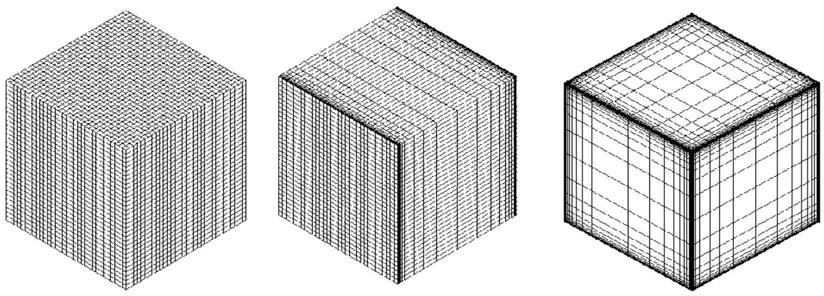

FIG. 4. *32×32×32 uniform grid, 32×32×32 stretched grid along x-direction (stretching ratio 1.5) and 32×32×32 stretched grid along all directions (stretching ratio 1.5).*

In the present work, the stretching of the grid in a given direction is determined by the stretching ratio (quotient between two consecutive space steps). See Fig. 4 for an example. Note that for exponential stretching in all directions, every cell can have different levels of anisotropy between the coefficients of the discrete operator. On the other hand, when there is exponential stretching in only one direction, two coefficients are similar and the third coefficient, corresponding to the stretching direction, changes in each cell.

Fig. 5 shows the residual versus work unit for 3-D V(1,0)-cycles to solve the isotropic equation on four 32×32×32 stretched grids with the following five smoothers:

- point-wise Gauss-Seidel (rhombus)
- alternating-line Gauss-Seidel (plus)
- alternating-plane Gauss-Seidel with exact solver in each plane (square)
- alternating-plane Gauss-Seidel with one 2-D V(1,1)-cycle in each plane (little "x")
- alternating-plane Gauss-Seidel with one 2-D V(1,0)-cycle in each plane (triangle)

The smoother used in the 2-D cycles to solve the planes is alternating-line Gauss-Seidel.

The approximate plane solver with one V(1,0) or V(1,1)-cycle performs better than the exact solver. Fig. 5 also shows an unexpected behavior of the alternating-line smoother. This optimal behavior is due to the use of stretching along the three directions that produces high discrepancies for the local values of the anisotropy in each cell, and therefore the local dominance of one of the directions. In fact, the numerical results obtained with grids stretched along just one direction show a poor behavior of the alternating-line smoother.

Fig. 6, on the other hand, shows residual versus work units for 3-D V(1,1)-cycles to solve

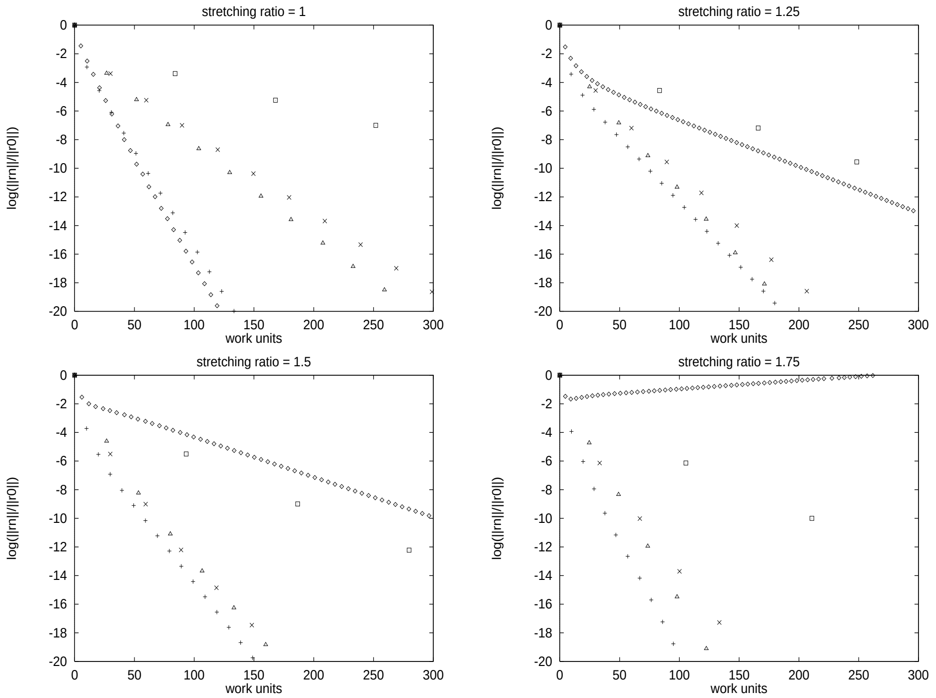

FIG. 5. *Residual versus work units for 3-D V(1,0)-cycles to solve the isotropic equation on four 32×32×32 stretched grids with the following five smoothers: point-wise Gauss-Seidel (rhombus), alternating-line Gauss-Seidel (plus), alternating-plane Gauss-Seidel with exact solver in each plane (square), alternating-plane Gauss-Seidel with one 2-D V(1,1)-cycle in each plane (little "x"), and alternating-plane Gauss-Seidel with one 2-D V(1,0)-cycle in each plane (trianle). The smoother used in the 2-D cycles is alternating-line Gauss-Seidel. Each symbol represents a 3-D cycle in order to compare the complexity with the cycles of different smoothers.*

the isotropic equation on four 32×32×32 stretched grids with the following five alternating-plane smoothers:

- Gauss-Seidel (rhombus)
- zebra Gauss-Seidel (plus)
- four-color Gauss-Seidel (square)
- Jacobi with damping 0.7 (little "x")
- Jacobi with partial damping 0.7 (triangle)

In this case, the best results are obtained with the lexicographic ordering.

8. Conclusions and future directions. We have shown numerically and analytically the smoothing factors of traditional plane relaxation methods with Dirichlet boundary conditions. For the multigrid solution of a discrete elliptic model equation on a cell-centered grid with strong anisotropies, the smoothing performance of the following relaxation methods have been investigated:

- plane Jacobi with damping
- plane Jacobi with partial damping
- plane Gauss-Seidel
- plane zebra Gauss-Seidel
- plane four-color Gauss-Seidel
- line Gauss-Seidel

We found that, as expected, as the anisotropy strength grows:

ETNA
Kent State University
etna@mcs.kent.edu

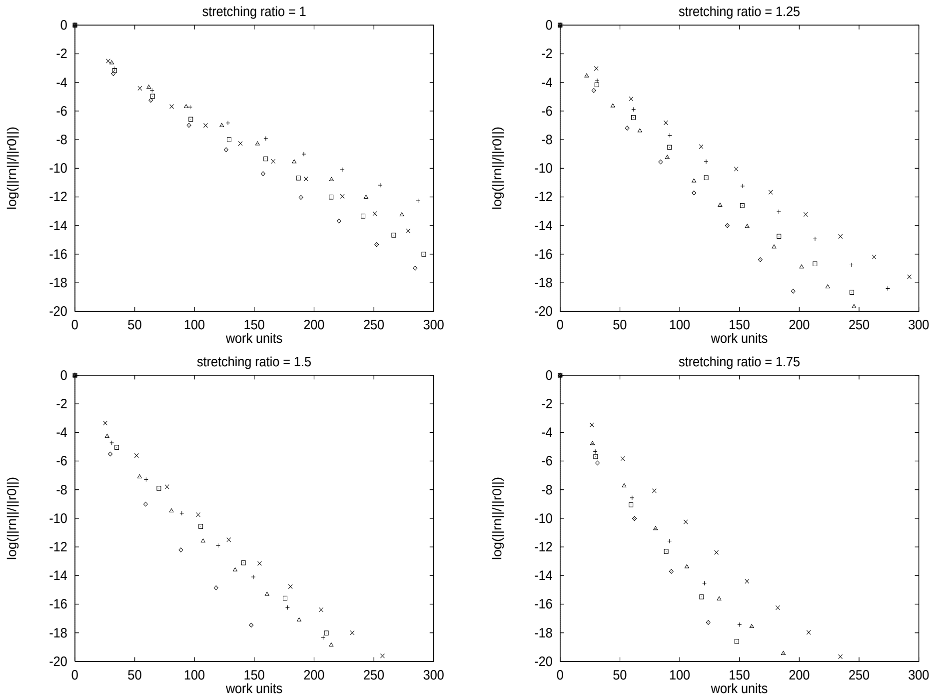

FIG. 6. *Residual versus work units for 3-D V(1,1)-cycles to solve the isotropic equation on four 32×32×32 stretched grids with the following five alternating-plane smoothers: Gauss-Seidel (rhombus), zebra Gauss-Seidel (plus), four-color Gauss-Seidel (square), Jacobi with damping 0.7 (little "x") and Jacobi with partial damping 0.7 (triangle). One 2-D V(1,1)-cycle is used to solve each plane. The smoother used in the 2-D cycles is the alternating-line version of the one used in the 3-D cycles.*

1. for periodic boundary conditions, the smoothing factor of a block smoother with full coarsening approaches the smoothing factor obtained with the point-wise version with full coarsening for the lower-dimensional problem defined by the coordinates that are not relaxed in the block, and
2. for Dirichlet boundary conditions, the smoothing factor of a block smoother with full coarsening approaches an exact solver.

Although line smoothers give very good results when one of the anisotropies is much larger than the other, they perform poorly when both anisotropies are similar, and hence cannot be considered for a robust method. All of the plane-implicit schemes studied, except the plane Jacobi with damping, show a linear decrease of smoothing factor with increasing anisotropy strength regardless of the relative strengths of the two anisotropies possible in 3-D problems considered. The good behavior of the plane smoothers deteriorates very slightly with an increase in the number of cells per side in the grid. Consequently, we believe that their excellent performance can be maintained for large, practical problems. Moreover, due to computer memory limitations, the number of cells per side is usually much lower in 3-D cases than in 1-D or 2-D cases.

The numerical results indicate that zebra Gauss-Seidel does not perform as well as expected on cell-centered grids. In fact, the lexicographic order Gauss-Seidel obtains better convergence rates, and the four-color plane Gauss-Seidel becomes an attractive smoother because of its good convergence rates and parallel properties. The plane Jacobi with partial damping is also a very good alternative; it performs worse in the isotropic case, but exhibits coarser granularity in a parallel setting.

The practical feasibility of plane relaxation as a multigrid smoother has been demonstrated. The solution of the 2-D boundary-value problem corresponding to each plane can be approximated with just one 2-D multigrid cycle. The same behavior can be expected if 1-D multigrid is applied to the solution of the tridiagonal systems of equations involved in the plane solution. As a result, alternating-plane smoothers are just two times slower than point-wise smoothers in the isotropic case and are robust multigrid smoothers that are orders-of-magnitude faster for anisotropic operators. Moreover, plane smoothers are easy to program, both on sequential and parallel computers.

Because we are interested in the applicability of plane smoothers with multiblock grids, we will continue to work on block-structured plane smoothers. In particular, we want to study the behavior of blocked plane smoothers and determine whether the results relating block size, overlap, and anisotropy strength obtained by Jones and Melson in ref. [9] hold in the 3-D case and in more complicated pde's and problem geometries. We believe that the smoothing performance will not suffer excessive deterioration with domain decomposition.

9. Acknowledgements. The authors would like to acknowledge Manuel Salas and James L. Thomas for their valuable comments on the CFD discipline and Fourier local mode analysis. James L. Thomas proposed the use of partial damping in the Jacobi method to improve its behavior for strong anisotropies. We also would like to thank David Sidilkover, Dimitri J. Mavriplis, Thor Gjesdal, Kees Oosterlee and Francisco Gaspar for their valuable comments and for providing copies of their work.

We would also like to thank Veer Vatsa for his many words of encouragement and his endless enthusiasm in the search for faster multigrid methods.

REFERENCES

[1] J. C. AGÜI AND J. JIMENEZ, *A binary tree implementation of a parallel distributed tridiagonal solver*, Parallel Comput., 21 (1995), pp. 233–241.
[2] A. BRANDT, *Multigrid techniques: 1984 guide with applications to fluid dynamics*, Tech. Rep. GMD-Studien 85, May 1984.
[3] W. L. BRIGGS, *A multigrid tutorial*, SIAM, 1988.
[4] J. E. DENDY JR., M. P. IDA, AND J. M. RUTLEDGE, *A semi-coarsening multigrid algorithm for SIMD machines*, SIAM J. Sci. Stat. Comput,, 13 (1992), pp. 1460–1469.
[5] J. E. DENDY, S. F. MCCORMICK, J. RUGE, T. RUSSELL, AND S. SCHAFFER, *Multigrid methods for three-dimensional petroleum reservoir simulation*, in Tenth SPE Symposium on Reservoir Simulation, february 1989.
[6] C. C. DOUGLAS, S. MALHOTRA, AND M. H. SCHULTZ, *Transpose free alternating direction smoothers for serial and parallel methods*. MGNET at http://www.mgnet.org/, 1997.
[7] D. ESPADAS, M. PRIETO, I. M. LLORENTE, AND F. TIRADO, *Parallel resolution of alternating-line processes by means of pipelined techniques*, in Proceedings of the 7th. Euromicro Workshop on Parallel and Distributed Systems, IEEE Computer Society Press, June 1999.
[8] T. GJESDAL, *Smoothing analysis of multicolor pattern schemes*, J. Comp. Appl. Math., 85 (1997), pp. 345–350.
[9] J. E. JONES AND N. D. MELSON, *A note on multi-block relaxation schemes for multigrid solvers*, Tech. Rep. 97-15, ICASE, 1997.
[10] A. KRECHEL, H. PLUM, AND K. STÜBEN, *Parallelization and vectorization aspects of the solution of tridiagonal linear systems*, Parallel Comput., 14 (1990), pp. 31–49.
[11] J. LINDEN, G. LONSDALE, H. RITZDORF, AND A. SCHÜLLER, *Scalability aspects of parallel multigrid*, Future Generation Computer Systems, 10 (1994), pp. 429–440.
[12] I. M. LLORENTE AND N. D. MELSON, *Robust multigrid smoothers for three dimensional elliptic equations with strong anisotropies*, Tech. Rep. 98-37, ICASE, 1998.
[13] I. M. LLORENTE AND F. TIRADO, *Relationships between efficiency and execution time of full multigrid methods on parallel computers*, IEEE Trans. on Parallel and Distributed Systems, 8 (1997), pp. 562–573.

[14] I. M. LLORENTE, F. TIRADO, AND L. VÁZQUEZ, *Some aspects about the scalability of scientific applications on parallel architectures*, Parallel Comput., 22 (1996), pp. 1169–1195.
[15] J. LÓPEZ, O. PLATAS, F. ARGÜELLO, AND E. L. ZAPATA, *Unified framework for the parallelization of divide and conquer based tridiagonal systems*, Parallel Comput., 23 (1997), pp. 667–686.
[16] D. J. MAVRIPLIS, *Multigrid strategies for viscous flow solvers on anisotropic unstructured meshes*, in Proceedings of the AIAA CFD Conference, June 1997.
[17] O. A. MCBRYAN, P. O. FREDERICKSON, J. LINDEN, A. SCHÜLLER, K. SOLCHENBACH, K. STÜBEN, C. THOLE, AND U. TROTTENBERG, *Multigrid methods on parallel computers-a survey of recent developments*, Impact of Computing in Science and Engineering, 3 (1991), pp. 1–75.
[18] W. MULDER, *A new multigrid approach to convection problems*, J. Comput. Phys., 83 (1989), pp. 303–323.
[19] N. H. NAIK AND J. V. ROSENDALE, *The improved robustness of multigrid elliptic solvers based on multiple semicoarsened grids*, SIAM J. Numer. Anal., 30 (1993), pp. 215–229.
[20] C. W. OOSTERLEE, *A GMRES-based plane smoother in multigrid to solve 3-D anisotropic fluid flow problems*, J. Comput. Phys., 130 (1997), pp. 41–53.
[21] A. OVERMAN AND J. V. ROSENDALE, *Mapping robust parallel multigrid algorithms to scalable memory architectures*, in Proceedings of the Sixth Copper Mountain Conference on Multigrid Methods, 1993.
[22] K. STÜBEN AND U. TROTTENBERG, *Multigrid methods: Fundamental algorithms, model problem analysis and applications*, in Multigrid methods, W. Hackbusch and U. Trottenberg, eds., vol. 960 of Lecture Notes in Mathematics, Springer-Verlag, 1982, pp. 1–176.
[23] C. THOLE AND U. TROTTENBERG, *Basic smoothing procedures for the multigrid treatment of elliptic 3d operators*, Appl. Math. Comput., 19 (1986), pp. 333–345.
[24] T. WASHIO AND K. OOSTERLEE, *An evaluation of parallel multigrid as a solver and a preconditioner for singularly perturbed problems, Part II: Flexible 2D and 3D multiple semicoarsening*, Tech. Rep. 1012, GMD, 1996 (to appear in SIAM J. Sci. Comput.).
[25] P. WESSELING, *A survey of smoothing analysis results*, in Multigrid Methods III, W. Hackbusch and U. Trottenberg, eds., Birkhäuser–Verlag, 1991, pp. 105–127.
[26] ———, *An Introduction to Multigrid Methods*, John Wiley & Sons, New York, 1992.
[27] I. YAVNEH, *Multigrid smoothing factors for red-black Gauss-Seidel relaxation applied to a class of elliptic operators*, SIAM J. Numer. Anal., 32 (1995), pp. 1126–1138.
[28] ———, *Multigrid smoothing factors of red-black Gauss-Seidel applied to a class of elliptic operators*, SIAM J. Numer. Anal., 32 (1995), pp. 1126–1138.