

BLOCK STATIONARY METHODS FOR NONSYMMETRIC CYCLICALLY REDUCED SYSTEMS ARISING FROM THREE-DIMENSIONAL ELLIPTIC EQUATIONS*

CHEN GREIF[†] AND JAMES VARAH[‡]

Abstract. We consider a three-dimensional convection-diffusion model problem and examine systems of equations arising from performing one step of cyclic reduction on an equally spaced mesh, discretized using the seven-point operator. We present two ordering strategies and analyze block splittings of the resulting matrices. If the matrices are consistently ordered relative to a given partitioning, Young's analysis for the block Gauss–Seidel and block SOR methods can be applied. We compare partitionings for which this property holds with ones where the matrices do not have Property A yet still give rise to an efficient solution process. Bounds on convergence rates are derived and the work involved in solving the systems is estimated.

Key words. cyclic reduction, stationary methods, three-dimensional problems, convection-diffusion

AMS subject classifications. 65F10, 65N22

PII. S0895479897317715

1. Introduction. Consider the three-dimensional (3D) convection-diffusion equation with constant coefficients

$$(1.1) \quad -(u_{xx} + u_{yy} + u_{zz}) + \sigma u_x + \tau u_y + \mu u_z = p(x, y, z)$$

on the unit cube $\Omega = [0, 1] \times [0, 1] \times [0, 1]$, subject to Dirichlet-type boundary conditions. We focus on applying seven-point finite difference discretizations, for example centered differences to the diffusive terms, and centered differences or first-order upwind approximations to the convective terms. Let us define n and h so that n^3 is the number of unknowns and $h = 1/(n + 1)$ is the mesh size, and let F denote the corresponding difference operator, after scaling by h^2 , so that for a gridpoint $u_{i,j,k}$ not next to the boundary we have

$$(1.2) \quad \begin{aligned} F u_{i,j,k} = & a u_{i,j,k} + b u_{i,j-1,k} + c u_{i-1,j,k} \\ & + d u_{i+1,j,k} + e u_{i,j+1,k} + f u_{i,j,k-1} + g u_{i,j,k+1}. \end{aligned}$$

If we denote the mesh Reynolds numbers by

$$(1.3) \quad \beta = \frac{\sigma h}{2}, \quad \gamma = \frac{\tau h}{2}, \quad \delta = \frac{\mu h}{2},$$

then the values of the components of the computational molecule are given by

$$(1.4) \quad \begin{aligned} a = 6, \quad b = -1 - \gamma, \quad c = -1 - \beta, \quad d = -1 + \beta, \\ e = -1 + \gamma, \quad f = -1 - \delta, \quad g = -1 + \delta \end{aligned}$$

*Received by the editors March 3, 1997; accepted for publication (in revised form) by Z. Strakoš February 27, 1998; published electronically July 9, 1999.

<http://www.siam.org/journals/simax/20-4/31771.html>

[†]SCCM Program, Gates Building, Stanford University, Stanford, CA 94305 (greif@sccm.stanford.edu).

[‡]Department of Computer Science, University of British Columbia, Vancouver, BC, V6T 1Z4, Canada (varah@cs.ubc.ca).




FIG. 1.1. Sparsity patterns of the matrices corresponding to two possible orderings of the unknowns

if centered difference approximations of the first derivatives are used and by

$$(1.5) \quad \begin{aligned} a &= 6 + 2(\beta + \gamma + \delta), & b &= -1 - 2\gamma, & c &= -1 - 2\beta, \\ d &= -1, & e &= -1, & f &= -1 - 2\delta, & g &= -1 \end{aligned}$$

if backward first-order accurate schemes are used.

The sparsity pattern of the underlying matrix depends on the ordering of the unknowns. In Fig. 1.1 the sparsity patterns associated with two possible ordering strategies are illustrated. The natural lexicographic ordering in (a) is one where the unknowns are numbered rowwise and then planewise. The red/black ordering in (b) means we color the gridpoints using two colors, in a checkerboard fashion, and then number all the points that correspond to one of the colors first.

As is evident from Fig. 1.1(b), if we split the matrix into four blocks of the same size, we can see that the two diagonal blocks are diagonal matrices. This means that the matrix has Property A [24]. A cheap and relatively simple process of elimination of all the points that correspond to one color (say, red) leads to a smaller system of equations, whose associated matrix is the Schur complement of the original matrix, and is still fairly sparse. This procedure amounts to performing one step of cyclic reduction. Notice that in general both the original and the reduced matrices are nonsymmetric.

The cyclic reduction step can be repeated until a small system of equations is obtained, which can then be solved directly. This procedure is called *complete* cyclic reduction. It has been studied in several papers, mainly for symmetric systems arising from two-dimensional (2D) self-adjoint elliptic problems. A general overview of the algorithm and a list of references can be found in [12]. Early papers that present and analyze the algorithm are those of Hockney [17], Buneman [2], and Buzbee, Golub, and Nielson [4]. Buzbee et al. [3] use cyclic reduction for solving the Poisson equation on irregular regions; Concus and Golub [6] discuss 2D nonseparable cases. Application of cyclic reduction to matrices with arbitrary dimensions is done by Sweet [20], [21]. Detyna [7] presents a fast $O(n^2)$ algorithm and discusses its stability and efficiency.

One step of cyclic reduction for symmetric positive definite systems is analyzed by Hageman and Varga in [16] and later by Hageman, Luk, and Young [15], where it is shown that the reduced solver generally converges faster than the unreduced solver.

In [1], Axelsson and Gustafsson use cyclic reduction in conjunction with the conjugate gradient method. Elman and Golub have conducted an extensive investigation for 2D elliptic non-self-adjoint problems [8], [9], [10] and have shown that one step of cyclic reduction leads to systems with several valuable properties, such as symmetrizability for a large set of the underlying PDE coefficients, which is effectively used to derive bounds on the convergence rates of iterative solvers, and fast convergence.

Preliminary analysis for the non-self-adjoint 3D model problem (1.1) has been done by the authors in [14], where one step of 3D cyclic reduction has been described in detail, and a block Jacobi solver has been analyzed, which is based on a certain block splitting (referred to as 1D splitting throughout this paper), in conjunction with what we called a two-plane ordering strategy.

The computational molecule of the reduced operator consists of 19 points, located on 5 parallel planes. Let R denote the reduced difference operator, after scaling by ah^2 . Then for an interior gridpoint, $u_{i,j,k}$, we have

$$\begin{aligned}
 (1.6) \quad Ru_{i,j,k} = & (a^2 - 2be - 2cd - 2fg) u_{i,j,k} - f^2 u_{i,j,k-2} - 2ef u_{i,j+1,k-1} \\
 & - 2cf u_{i-1,j,k-1} - 2df u_{i+1,j,k-1} - 2bf u_{i,j-1,k-1} - e^2 u_{i,j+2,k} \\
 & - 2de u_{i+1,j+1,k} - c^2 u_{i-2,j,k} - d^2 u_{i+2,j,k} - 2bc u_{i-1,j-1,k} \\
 & - b^2 u_{i,j-2,k} - 2eg u_{i,j+1,k+1} - 2cg u_{i-1,j,k+1} - 2ce u_{i-1,j+1,k} \\
 & - 2bd u_{i+1,j-1,k} - 2dg u_{i+1,j,k+1} - 2bg u_{i,j-1,k+1} - g^2 u_{i,j,k+2} .
 \end{aligned}$$

The following results hold for any ordering strategy. See [14] for the proofs, which have been obtained by using the techniques of Elman and Golub [8], [9].

THEOREM 1.1. *The reduced matrix can be symmetrized by a real diagonal similarity transformation if and only if the products $bcde$, $befg$, and $cdfg$ are positive.*

THEOREM 1.2. *If be , cd , $fg > 0$, then both the reduced matrix and the symmetrized reduced matrix are diagonally dominant M -matrices.*

In this paper our purpose is to extend the analysis initiated in [14] and examine block stationary methods as solvers for the reduced system. In section 2 we present the ordering strategies that are examined. In section 3 two block splittings are presented, and bounds on convergence rates are derived. In section 4 we analyze the reduced system in the context of consistently ordered matrices. In section 5 the amount of computational work involved in solving the linear systems is estimated, a comparison of the reduced system with the unreduced system is conducted, and some numerical results which validate our analysis and illustrate the fast convergence of the reduced system are given. Finally, in section 6 we conclude.

2. Orderings for the reduced system. We consider two ordering strategies for the reduced grid. The two-plane ordering has been described in detail in [14]. It corresponds to ordering the unknowns by gathering blocks of $2n$ gridpoints from two horizontal lines and two adjacent planes. This ordering strategy is depicted in Fig. 2.1(a). In the figure, the numbers are the indices of the gridpoints, which are to be expressed below by ℓ in (2.1) and (2.15).

The connection between the index of a gridpoint, ℓ , and its coordinate values $(i, j, k) = (\frac{x}{h}, \frac{y}{h}, \frac{z}{h})$ is given below. The term fix is borrowed from MATLAB and means rounding to the nearest integer toward zero.

$$(2.1a) \quad i = \text{fix}\{[(\ell - 1) \bmod (2n)]/2\} + 1,$$

$$(2.1b) \quad j = \begin{cases} 2 \cdot [\text{fix}(\frac{\ell-1}{n^2}) + 1], & \ell \bmod 4 = 0 \text{ or } 1, \\ 2 \cdot \text{fix}(\frac{\ell-1}{n^2}) + 1, & \ell \bmod 4 = 2 \text{ or } 3, \end{cases}$$




FIG. 2.2. Sparsity patterns of the reduced matrices associated with the two ordering strategies (the matrices correspond to $6 \times 6 \times 6$ grids). Each square corresponds to an $n^2 \times n^2$ block (n^2 gridpoints form two coupled planes in the reduced grid).

$$(2.11) \quad S_{j,j+1}^{(0)} = \begin{cases} \text{tri}(-2cg, -2bg \cdot E_{10} - 2eg \cdot E_{01}, -2dg), & j \text{ odd,} \\ \text{tri}(-2cg, -2bg \cdot E_{01} - 2eg \cdot E_{10}, -2dg), & j \text{ even,} \end{cases}$$

$$(2.12) \quad S_{j,j+1}^{(1)} = \begin{cases} \text{diag}(-2eg \cdot E_{10}), & j \text{ odd,} \\ \text{diag}(-2eg \cdot E_{01}), & j \text{ even.} \end{cases}$$

Finally, the matrices $S_{j,j-2}$ and $S_{j,j+2}$ are diagonal:

$$(2.13) \quad S_{j,j-2} = \text{diag}(-f^2), \quad j = 3, \dots, n;$$

$$(2.14) \quad S_{j,j+2} = \text{diag}(-g^2), \quad j = 1, \dots, n-2.$$

The connection between the gridpoint's index and its coordinate values is given by

$$(2.15a) \quad i = [(\ell - 1) \bmod n] + 1,$$

$$(2.15b) \quad k = \text{fix} \left(\frac{\ell - 1}{n^2/2} \right) + 1,$$

$$(2.15c) \quad j = \begin{cases} 2 \cdot (\text{fix}((\text{fix}(\ell - (k - 1) \cdot (n^2/2)) - 1)/n)) + 1, & k \text{ odd,} \\ 2 \cdot (\text{fix}((\text{fix}(\ell - (k - 1) \cdot (n^2/2)) - 1)/n)) + 2, & k \text{ even.} \end{cases}$$

The sparsity patterns of the matrices corresponding to two-plane ordering and two-line ordering are depicted in Fig. 2.2.




FIG. 3.1. Sparsity patterns of the block diagonal matrices associated with the block Jacobi splitting, for the two suggested block splittings, using the two-plane ordering strategy.

3. Block splittings and bounds on convergence rate. For the two ordering strategies presented in section 2 the matrices can be expressed as block tridiagonal, of the form

$$(3.1) \quad S = \text{tri}[S_{j,j-1}, S_{j,j}, S_{j,j+1}].$$

S is an $(n^3/2) \times (n^3/2)$ matrix. In the case of two-plane ordering, each block $S_{i,j}$ is of size $n^2 \times n^2$ and is block tridiagonal with respect to $2n \times 2n$ blocks. In the case of two-line ordering, each block $S_{i,j}$ is of size $(n^2/2) \times (n^2/2)$ and is block tridiagonal with respect to $n \times n$ blocks.

In solving the reduced system using a stationary method, various splittings are possible. We consider two obvious ones, based on dimension. We use the term 1D splitting for a splitting which is based on partitioning the matrix into $O(n)$ blocks ($2n \times 2n$ blocks for the two-plane ordering and $n \times n$ blocks for the two-line ordering). A 2D splitting is one which is based on partitioning the matrix into $O(n^2)$ blocks ($n^2 \times n^2$ blocks for the two-plane ordering and $(n^2/2) \times (n^2/2)$ blocks for the two-line ordering). Notice that the 1D splitting for both ordering strategies is essentially associated with blocks of gridpoints that are x -oriented. However, the 2D splitting for the two-line ordering corresponds to x - y oriented planes of gridpoints, whereas for the two-plane ordering it corresponds to x - z oriented planes of gridpoints. (These observations can be deduced by referring to Fig. 2.1.) Different orientations can be obtained by simply reordering the unknowns so that the roles of x , y , and z are interchanged.

The sparsity patterns of the block diagonal parts of the splittings associated with the block Jacobi scheme are depicted in Fig. 3.1.

We now compare the orderings. We have the following useful result.

THEOREM 3.1. *If $be, cd, fg > 0$, then for the 1D splitting the spectral radius of the Jacobi iteration matrix associated with two-plane ordering is smaller than the spectral radius of the iteration matrix associated with the two-line ordering.*

Proof. By Theorem 1.2 the matrices are M -matrices. Each ordering strategy produces a matrix which is merely a symmetric permutation of a matrix associated with the other ordering. Suppose $S_1 = M_1 - N_1$ is a 1D splitting of the two-plane ordering matrix and $S_2 = M_2 - N_2$ is a 1D splitting for the two-line order-

ing. There exists a permutation matrix P such that $P^T S_2 P = S_1$. Consider the splitting $P^T S_2 P = P^T M_2 P - P^T N_2 P$. It is straightforward to show by examining the matrix entries that $P^T N_2 P \geq N_1$. The latter are both nonnegative matrices; therefore by [23, Thm. 3.15] it follows that $0 < \rho(M_1^{-1} N_1) < \rho(P^T M_2^{-1} N_2 P) = \rho(M_2^{-1} N_2) < 1$. \square

The same result applies to 2D splitting, provided that the orientation of the planes of gridpoints is identical for both ordering strategies. The proof for this is identical to the proof of Theorem 3.1.

The results indicated in Theorem 3.1 can be observed in Fig. 3.2. It is interesting to observe that the superiority of the two-plane ordering carries over to the case $be, cd, fg < 0$, which corresponds to the region of mesh Reynolds numbers larger than 1 (for which the PDE is considered convection-dominated). We remark, however, that the amount of computational work per each iteration is somewhat higher for the system which corresponds to two-plane ordering. In Fig. 3.2 a few cross sections of mesh Reynolds numbers are examined. For example, graph (a) corresponds to flow with the same velocity in x, y , and z directions. Graph (b) corresponds to flow only in x and y directions, and no convection in z direction, and so on. (See (1.3) for definitions of β, γ , and δ .)

We now derive bounds on convergence rates. Below we shall attach the subscripts 1 and 2 to matrices associated with the 1D splitting and 2D splitting, respectively. Since two-plane ordering gives rise to a more-efficient solution procedure than two-line ordering, we focus on it.

Denote the two splittings for the block Jacobi scheme by $S = D_1 - C_1 = D_2 - C_2$. In [14] we have shown that if be, cd , and fg have the same sign then a real diagonal nonsingular symmetrizer can be found, and thus (since the symmetrizer is diagonal) the sparsity patterns of the original nonsymmetric matrix and the symmetrized matrix are identical. Let us attach the hat sign to a matrix to denote application of the similarity transformation that symmetrizes it. That is, for a given matrix X and a diagonal symmetrizer Q , $Q^{-1} X Q$ is to be denoted by \hat{X} .

The matrices $\hat{D}_1^{-1} \hat{C}_1$ and $\hat{D}_2^{-1} \hat{C}_2$ are similar to the original iteration matrices $D_1^{-1} C_1$ and $D_2^{-1} C_2$, respectively, and thus have the same spectral radii. Following Elman and Golub's strategy [8], [9], the symmetric matrix can be handled more easily as far as computing the spectral radius is concerned, since we can use the following:

$$(3.2) \quad \rho(D_i^{-1} C_i) = \rho(\hat{D}_i^{-1} \hat{C}_i) \leq \|\hat{D}_i^{-1}\|_2 \|\hat{C}_i\|_2 = \frac{\rho(\hat{C}_i)}{\lambda_{\min}(\hat{D}_i)}, \quad i = 1, 2.$$

The results presented below are for the case $be, cd, fg > 0$, using two-plane ordering. These conditions are equivalent to $|\beta|, |\gamma|, |\delta| < 1$ if centered differences are used to discretize the convective terms. No restriction on the magnitude of the mesh Reynolds numbers is imposed if upwind differences are used. For these values tight bounds for the spectral radius of the iteration matrix can be obtained.

For \hat{D}_1 the minimal eigenvalue has been found in [14, Thm. 3.8], and the relevant part of this theorem is quoted below.

PROPOSITION 3.2. *The minimal eigenvalue of \hat{D}_1 is*

$$(3.3) \quad \eta = a^2 - 2be - 2fg - 2\sqrt{befg} - 4(\sqrt{bcde} + \sqrt{cdfg}) \cdot \cos(\pi h) - 4cd \cos^2(\pi h).$$

A lower bound for $\hat{D}_2 - \hat{D}_1$ is given by the following proposition.

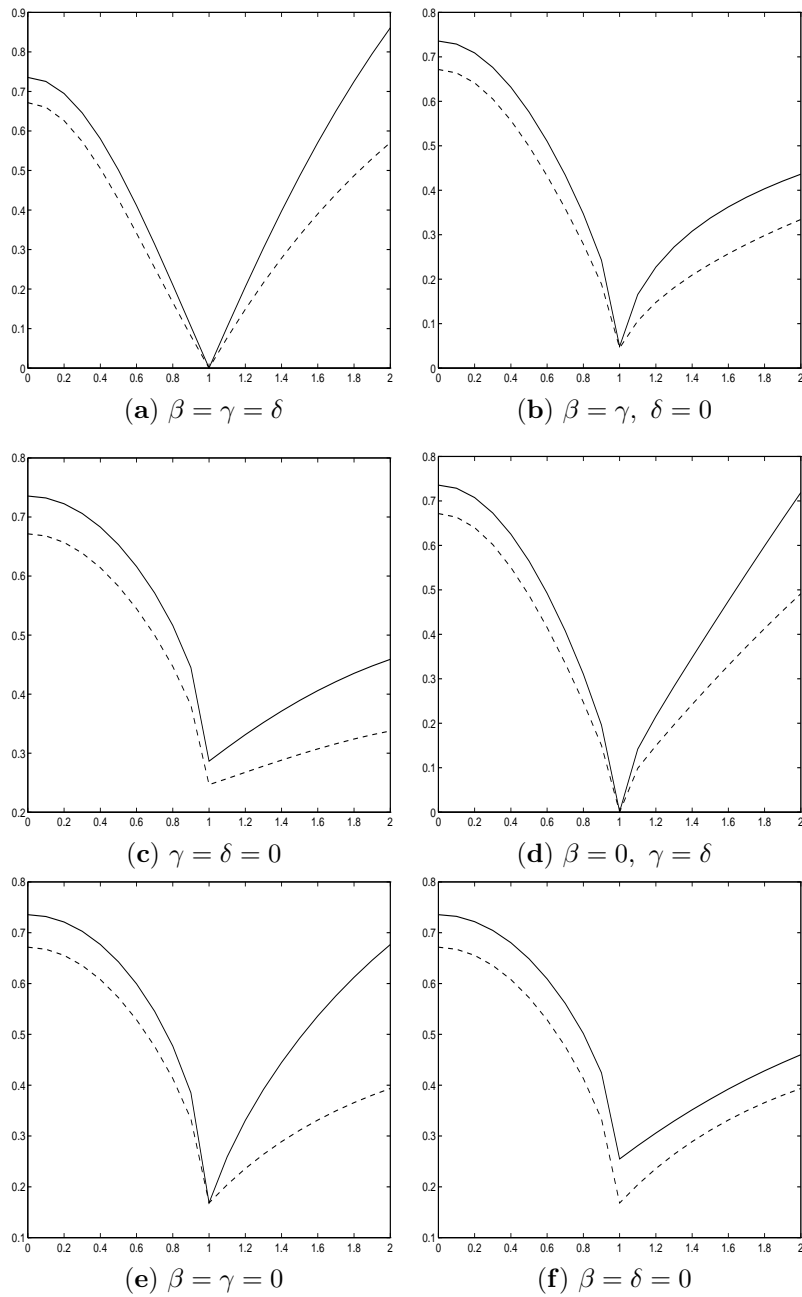


FIG. 3.2. Spectral radii of iteration matrices versus mesh Reynolds numbers for the block Jacobi scheme, using 1D splitting and centered difference discretization. The broken lines correspond to two-plane ordering. The solid lines correspond to two-line ordering.

PROPOSITION 3.3. The minimal eigenvalue of $\hat{D}_2 - \hat{D}_1$ is bounded from below by $-\xi$, where

$$(3.4) \quad \xi = 2fg \cos\left(\frac{\pi}{\frac{n}{2} + 1}\right) + \sqrt{4bfg + 16 \cdot cdfg \cdot \cos^2(\pi h) + 16\sqrt{bcde} \cdot fg \cdot \cos(\pi h)}.$$

TABLE 3.1

Comparison between the computed spectral radius and the bound, for the 2D splitting, with $\beta = \gamma = \delta = 0.5$.

Scheme	Upwind			Centered		
	ρ	bound	ratio	ρ	bound	ratio
4	0.265	0.430	1.62	0.203	0.309	1.52
6	0.411	0.530	1.29	0.297	0.368	1.24
8	0.499	0.583	1.17	0.350	0.398	1.14
10	0.553	0.615	1.11	0.381	0.415	1.09
12	0.588	0.634	1.08	0.400	0.425	1.06
14	0.611	0.647	1.06	0.413	0.432	1.05

The proof for this part follows from [14, Lem. 3.10], where it is shown that the spectral radius of $\hat{D}_2 - \hat{D}_1$ is bounded by ξ .

Combining Propositions 3.2 and 3.3, and applying Rayleigh quotients to the matrices \hat{D}_1 and $\hat{D}_2 - \hat{D}_1$, we obtain the following lemma.

LEMMA 3.4. *The minimal eigenvalue of \hat{D}_2 is bounded from below by $\eta - \xi$, where η and ξ are the expressions given in (3.3) and (3.4).*

The bound for \hat{C}_2 can be obtained by combining [14, Lems. 3.11–3.13], as follows.

LEMMA 3.5. *The spectral radius of the matrix \hat{C}_2 is bounded by*

$$(3.5) \quad \phi = 4\sqrt{befg} + 4\sqrt{bcde} \cdot \cos \frac{\pi}{n+1} + 2be \cos \left(\frac{\pi}{\frac{n}{2} + 1} \right).$$

Finally, Lemmas 3.4 and 3.5 lead to the following theorem.

THEOREM 3.6. *The spectral radii of the iteration matrices $D_1^{-1}C_1$ and $D_2^{-1}C_2$ are bounded by $\frac{\phi + \xi}{\eta}$ and $\frac{\phi}{\eta - \xi}$, respectively, where η , ξ , and ϕ are defined in (3.3), (3.4), and (3.5), respectively.*

COROLLARY 3.7. *If $be, cd, fg > 0$ then the block Jacobi iteration converges for both the 1D and 2D splittings.*

Proof. For this we can use Varga’s result on M -matrices [23, Thm. 3.13]. Alternatively, Taylor expansions of the bounds given in Theorem 3.6 are given by

$$(3.6) \quad \rho(D_1^{-1}C_1) \leq \frac{\phi + \xi}{\eta} = 1 - \left(\frac{10}{9}\pi^2 + \frac{1}{6}\mu^2 + \frac{1}{6}\tau^2 + \frac{1}{6}\sigma^2 \right) h^2 + o(h^2)$$

and

$$(3.7) \quad \rho(D_2^{-1}C_2) \leq \frac{\phi}{\eta - \xi} = 1 - \left(2\pi^2 + \frac{3}{10}\mu^2 + \frac{3}{10}\tau^2 + \frac{3}{10}\sigma^2 \right) h^2 + o(h^2)$$

and thus are smaller than 1. \square

In Table 3.1 we give some indication on the quality of the bound for the 2D splitting. Results with a similar level of accuracy have been obtained and presented in [14] for the 1D splitting. As can be observed, the bounds are tight and become tighter as n increases, which suggests that they are asymptotic to the spectral radii.

We now discuss other stationary methods, namely, Gauss–Seidel and SOR. Relative to a given partitioning, if the reduced matrix is consistently ordered, then it is straightforward to apply Young’s analysis, and the bounds in Theorem 3.6 can be used for estimating the rate of convergence of the Gauss–Seidel and SOR schemes. The reader is referred to [19, Defs. 4.3 and 4.4] for definitions of Property A and consistent ordering. As stated in [19], a matrix that is consistently ordered has Property

A; conversely, a matrix with Property A can be permuted so that it is consistently ordered. We mentioned in the introduction that the matrix of the unreduced system has Property A. For the reduced system, we have the following observations.

PROPOSITION 3.8. *The reduced matrix associated with two-line ordering, S_L , does not have Property A relative to 1D or 2D partitionings.*

Proof. Let $S_{i,j}$ denote the (i, j) th $n \times n$ block of S_L , and let Q be an $(n^2/2) \times (n^2/2)$ matrix, whose entries satisfy $q_{i,j} = 1$ if $S_{i,j} \neq 0$ and $q_{i,j} = 0$ otherwise. Let T be an $n \times n$ matrix, such that $t_{i,j} = 1$ if the (i, j) th $(n^2/2) \times (n^2/2)$ block submatrix of S is nonzero, and $t_{i,j} = 0$ otherwise. Clearly, T is a pentadiagonal matrix and thus does not have Property A. Since Q can be referred to as a partitioning of T into $(n/2) \times (n/2)$ blocks, it also does not have Property A. \square

PROPOSITION 3.9. *The reduced matrix associated with two-plane ordering, S_P , does not have Property A relative to 1D partitioning.*

Proof. Let $S_{i,j}$ denote the (i, j) th $2n \times 2n$ block of S_P , and let Q be an $(n^2/4) \times (n^2/4)$ matrix, whose entries satisfy $q_{i,j} = 1$ if $S_{i,j} \neq 0$ and $q_{i,j} = 0$ otherwise. It is straightforward to see that the nonzero pattern of Q is identical to that of the matrix associated with using a nine-point operator for a 2D grid. Since the latter does not have Property A relative to partitioning into 1×1 matrices, the result follows. \square

On the other hand, we have the following proposition.

PROPOSITION 3.10. *The reduced matrix associated with two-plane ordering, S_P , has Property A and, moreover, is consistently ordered relative to 2D partitioning.*

Proof. The matrix is block tridiagonal relative to this partitioning [24]. \square

For the SOR scheme we have the following result, which is completely analogous to Elman and Golub’s result for the 2D problem [9, Thm. 4].

THEOREM 3.11. *Let \mathcal{L}_ω denote the block SOR operator associated with 2D splitting and using two-plane ordering. If either $be, cd, fg > 0$ or < 0 , then the choice $\omega^* = \frac{2}{1 + \sqrt{1 - \rho^2(D_2^{-1}C_2)}}$ minimizes $\rho(\mathcal{L}_\omega)$ with respect to ω , and $\rho(\mathcal{L}_{\omega^*}) = \omega^* - 1$.*

The proof of this theorem is essentially identical to the proof of Elman and Golub in [9, Thm. 4] and follows from Young [24, Chap. 14, Sects. 5.2 and 14.3]. The algebraic details on how to pick the signs of the diagonal symmetrizer so that the symmetrized block diagonal part of the splitting is a diagonally dominant M -matrix are omitted. That $\rho(D_2^{-1}C_2) < 1$ is known by Corollary 3.7. The reduced matrix is consistently ordered by Proposition 3.10.

A way to approximately determine an optimal relaxation parameter for the case $be, cd, fg > 0$ is to replace $\rho(D_2^{-1}C_2)$ by the bound for it (given in Theorem 3.6) in the expression for ω^* in Theorem 3.11. If the bound for the block Jacobi scheme is tight, then the estimate of ω^* is fairly accurate.

PROPOSITION 3.12. *Suppose $be, cd, fg > 0$. For the system associated with 2D splitting and for h sufficiently small, the choice*

$$(3.8) \quad \tilde{\omega}^* = \frac{2(\eta - \xi)}{\eta - \xi + \sqrt{(\eta - \xi)^2 - \phi^2}}$$

approximately minimizes $\rho(\mathcal{L}_\omega)$. The spectral radius of the iteration matrix is approximately $\tilde{\omega}^ - 1$.*

The Taylor expansion of the estimate for the optimal relaxation parameter is given by

$$(3.9) \quad \tilde{\omega}^* = 2 - \frac{2}{5} \sqrt{100\pi^2 + 15(\tau^2 + \mu^2 + \sigma^2)} \cdot h + O(h^2).$$




FIG. 4.1. The sparsity pattern of the matrix C_d .

From (3.9) it follows that the estimated asymptotic rate of convergence of the block SOR scheme is approximately the second term in (3.9) (with the negative sign removed) and is thus $O(h)$.

4. Near-Property A for 1D splitting of the two-plane matrix. Although the matrix associated with two-plane ordering does not have Property A relative to the 1D partitioning, some interesting observations can be made: As before, let $\{S_{i,j}\}$ denote the $n^2 \times n^2$ blocks of the reduced matrix. Each block $S_{i,j}$ is a block tridiagonal matrix relative to $2n \times 2n$ blocks. We attach superscripts to mark how far a block diagonal is from the main block diagonal, and we define

$$(4.1) \quad C_d = - \left(S_{j,j+1}^{(-1)} + S_{j,j+1}^{(1)} + S_{j,j-1}^{(-1)} + S_{j,j-1}^{(1)} \right).$$

See [14] for specification of the entries of these matrices. As in section 3 (with a slight change in notation), let $S_P = D - C$ be the 1D splitting of the matrix, and define \tilde{C} so that $S_P = D - (\tilde{C} + C_d)$. The matrix $D - \tilde{C}$ has Property A, but S_P does not. Let us examine the matrix that prevents S_P from having Property A, namely, C_d . It is an extremely sparse matrix, and the magnitude of the nonzero values in this matrix is bounded by 2 if $be, cd, fg > 0$. The nonzero pattern of C_d is depicted in Fig. 4.1.

We wish to estimate how far the reduced matrix S_P is from having block Property A, relative to the 1D partitioning. Let us denote the upper part and the lower part of C_d by U_d and L_d , respectively, and let \tilde{U} and \tilde{L} be the upper part and lower part of \tilde{C} , respectively. Then the spectral radius of the block Gauss-Seidel matrix satisfies

$$(4.2) \quad \rho_{GS} = \rho((D - \tilde{L} - L_d)^{-1}(\tilde{U} + U_d)) \leq \frac{\|(D - \tilde{L})^{-1}\tilde{U}\|_2 + \|(D - \tilde{L})^{-1}U_d\|_2}{1 - \|(D - \tilde{L})^{-1}L_d\|_2}.$$

$\|(D - \tilde{L})^{-1}\tilde{U}\|_2$ is significantly larger than the other norms in the above inequality, which means that the spectral radius of the Gauss-Seidel iteration matrix associated with the two-plane ordering can be estimated by replacing the two-plane matrix by $D - \tilde{C}$, which does have Property A and thus is easier to analyze. Alternatively, the following observation has been obtained by numerical experiments:

$$(4.3) \quad \rho(D^{-1}C) \approx \rho(D^{-1}\tilde{C}) + \rho(D^{-1}C_d).$$

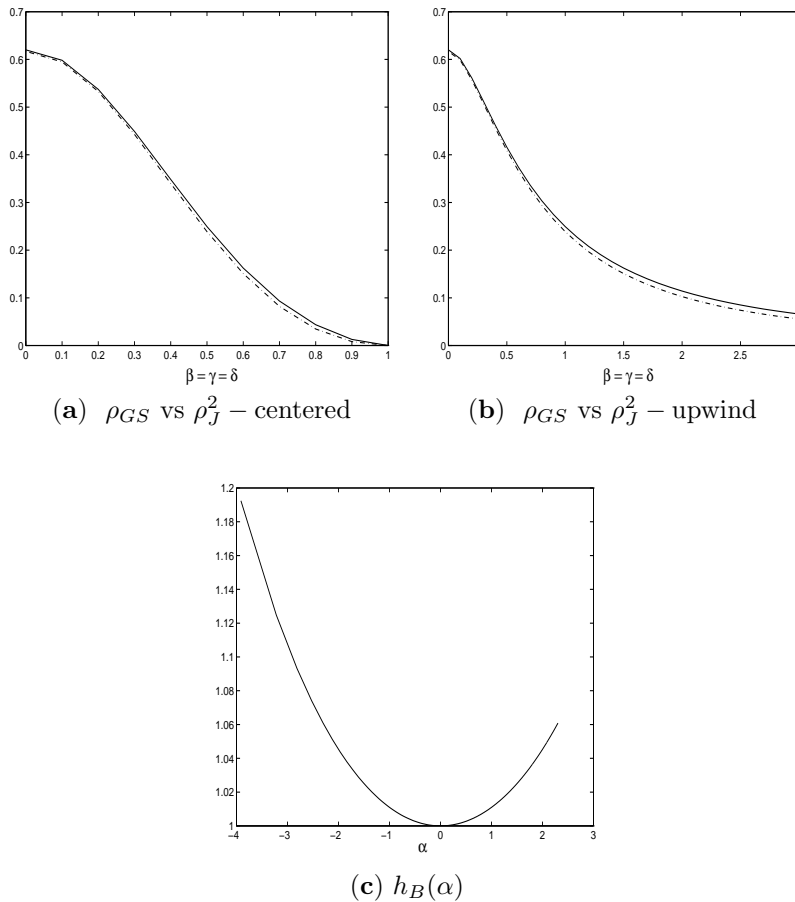


FIG. 4.2. “Near Property A” for the 1D splitting.

Young’s analysis can be applied directly to both $D - \tilde{C}$ and $D - C_d$ (both have Property A), and thus an approximate relationship between the eigenvalues of the block Jacobi iteration matrix and the eigenvalues of the block Gauss–Seidel iteration matrix can be obtained.

For $be, cd, fg > 0$ we have observed that the spectral radius of the block Jacobi iteration matrix satisfies

$$(4.4) \quad \rho_J^2 \approx \rho_{GS}.$$

The first two graphs in Fig. 4.2 illustrate this phenomenon numerically. The broken lines in graphs (a) and (b) correspond to the square of the spectral radius of the iteration matrix associated with block Jacobi, for a 256×256 matrix. The solid lines correspond to the spectral radius of the block Gauss–Seidel iteration matrix. As can be seen, the curves are almost indistinguishable. This phenomenon becomes more dramatic as the systems become larger.

Some analysis can be done using Varga’s work on extensions of the theory of p -cyclic matrices [22], [23, Sect. 4.4]. (In this paper we are concerned only with $p = 2$.) Recall [23, Def. 4.2], which defines a set S of matrices as follows. The square matrix $B \in S$ if B satisfies the following properties:

1. $B \geq 0$ with zero diagonal entries.
2. B is irreducible and convergent, i.e., $0 < \rho(B) < 1$.
3. B is symmetric.

If $be, cd, fg > 0$, the reduced matrix S_P is a diagonally dominant M -matrix which can be symmetrized, and $\hat{D}^{1/2}$ is well defined. Define $\tilde{S} = \hat{D}^{-1/2}\hat{C}\hat{D}^{-1/2} = I - \hat{D}^{-1/2}\hat{C}\hat{D}^{-1/2}$.

Applying block Jacobi to the original reduced system is analogous to applying point Jacobi to \tilde{S} , in the sense that the spectra of the iteration matrices associated with both systems are identical. The iteration matrix associated with \tilde{S} is $B = \hat{D}^{-1/2}\hat{C}\hat{D}^{-1/2}$. Showing that the matrix B belongs to the set S defined above is easy and is omitted. Let L be the lower part of B . Define $M_B(\theta) = \theta L + L^T, \theta \geq 0$, and $m_B(\theta) = \rho(M_B(\theta))$. Let

$$(4.5) \quad h_B(\ln \theta) = \frac{m_B(\theta)}{\rho(B)\theta^{1/2}}, \quad \theta > 0.$$

Then we have [23, Thm. 4.7] (with a slight modification so as to match the terminology used in this paper), as follows.

THEOREM 4.1. *Let $B \in S$. Then, $h_B(\alpha) \equiv 1$ if and only if B is consistently ordered.*

In some sense $h_B(\ln \theta)$ measures the departure of the matrix B from having block Property A. For matrices that are not consistently ordered, the following result applies [23, Thm. 4.6].

THEOREM 4.2. *If $B \in S$, then either $h_B(\alpha) \equiv 1$ for all real α , or $h_B(\alpha)$ is strictly increasing for $\alpha \geq 0$. Moreover, for any $\alpha \neq 0$,*

$$(4.6) \quad 1 \leq h_B(\alpha) \leq \cosh\left(\frac{\alpha}{2}\right).$$

Figure 4.2(c) demonstrates how close the function h_B is to 1 for the reduced matrix when 1D partitioning is used and provides another way to illustrate the near-Property A of the matrix. In the figure, the function h_B is computed for a symmetrized block Jacobi 256×256 matrix, where $\beta = \gamma = \delta = 0.5$.

We can now analyze the Gauss–Seidel and SOR schemes. Recall [23, Thm. 4.8] (slightly modified), as follows.

THEOREM 4.3. *Let $\mathcal{L}_{B,\omega}$ denote the SOR iteration matrix. If $B \in S$ then the Gauss–Seidel iteration matrix, which corresponds to the case $\omega = 1$, satisfies*

$$(4.7) \quad \rho^2(B) \leq \rho(\mathcal{L}_{B,1}) < \frac{\rho(B)}{2 - \rho(B)},$$

with equality possible only if B is consistently ordered.

This is a sharpened form of the Stein–Rosenberg theorem [23]. Applying this theorem to our reduced matrix, we have the following theorem.

THEOREM 4.4. *If the bound for the block Jacobi iteration matrix tends to the actual spectral radius as $h \rightarrow 0$, then the spectral radius of the block Gauss–Seidel iteration matrix coincides with the square of the bound for the spectral radius of the block Jacobi iteration matrix up to $O(h^2)$ terms.*

Proof. Since the iteration matrix B has the same spectral radius as $D^{-1}C$, where $S_P = D - C$, we can use the bound for the 1D iteration matrix, which was presented in Theorem 3.6. For simplicity of notation, denote this bound by Φ . Clearly, since




FIG. 4.3. Spectral radius of the SOR iteration matrix versus the relaxation parameter. The uppermost curve corresponds to 1D splitting for the unreduced system, and then we have, in order, 2D splitting for the unreduced system, 1D splitting for the reduced system, and 2D splitting for the reduced system.

$$0 < \rho(B) \leq \Phi,$$

$$(4.8) \quad \frac{\rho(B)}{2 - \rho(B)} \leq \frac{\Phi}{2 - \Phi}.$$

Since Φ has a Taylor expansion of the form $1 - ch^2 + O(h^2)$, by (3.6), it follows that $\frac{\Phi}{2 - \Phi}$ and Φ^2 have the same Taylor expansion up to $O(h^2)$ terms, of the form $1 - 2ch^2 + O(h^3)$. Indeed, in terms of the PDE coefficients,

$$(4.9) \quad \frac{\Phi}{2 - \Phi} = 1 - \left(\frac{20}{9}\pi^2 + \frac{1}{3}\sigma^2 + \frac{1}{3}\tau^2 + \frac{1}{3}\mu^2 \right) h^2 + O(h^3),$$

and the same for Φ^2 . It has been shown that the bound for $\rho(B)$ is extremely tight as $h \rightarrow 0$, and so we can replace the spectral radii by the bounds for the spectral radii in Theorem 4.3 to obtain the desired result. \square

The actual meaning of this result is that for systems of equations that are large enough, the matrix nearly has Property A relative to 1D partitioning, at least as far as the convergence properties of the block Gauss–Seidel scheme are concerned. Since the solution process for small mesh Reynolds numbers is more efficient for the 1D splitting, compared to the 2D splitting, as we shall see in section 5, it was our aim to overcome the difficulty of not being able to apply Young’s analysis directly.

For the block SOR scheme, the upper bound for the spectral radius is given in [23, Thm. 4.9] as $\sqrt{\omega^* - 1}$ and is not tight. However, it is numerically evident that the bound for the Jacobi iteration matrix can be effectively used to estimate the optimal SOR parameter. In Fig. 4.3 we can observe that the behavior for the 1D splitting is qualitatively identical to the behavior of two-cyclic consistently ordered matrices. Here we present results for centered difference discretization of the problem with $\beta = \gamma = \delta = 0.5$. The reduced matrix is 256×256 . In the figure we also present the behavior of the SOR iteration matrix of the unreduced system.

5. Computational work and numerical experiments. Having done some analysis, in this section we examine which of the 1D and 2D solvers is more efficient overall and show that the reduced system is superior to the unreduced system.

5.1. Aspects of computational work. If $be, cd, fg > 0$, then by [23, Thm. 3.15] or by (3.6) and (3.7), it is evident that the spectral radius of the iteration matrix associated with the 2D splitting is smaller than that of the 1D iteration matrix. However, inverting D_1 involves less computational work than inverting D_2 . We now compare these two solution procedures.

We begin with the block Jacobi scheme. Asymptotically, there is a fixed ratio of 1.8 between the rate of convergence of the two splittings (see (3.6) and (3.7)). In rough terms, this number characterizes the ratio between number of iterations until convergence for the two solvers.

As far as the computational work per iteration is concerned, if $D_1 = L_1U_1$ and $D_2 = L_2U_2$ are the LU decompositions of the matrices of the systems that are to be solved in each iteration, we can assume that the number of operations per iteration is approximately the number of nonzeros in $L_i + U_i$ plus the number of nonzeros in the other part of the splitting. In order to avoid costly fill-in using Gaussian elimination for D_2 (whose band is sparse), we use instead a technique of inner-outer iterations.

Let k_1 and k_2 denote the number of iterations for the schemes associated with the 1D splitting and the 2D splitting, respectively. Let us also define cost functions as follows: $c_1(n)$ and $c_2(n)$ represent the overall number of floating point operations for each of the solvers, and $c_{in}(n)$ represents the cost of the inner solve. Then

$$(5.1a) \quad c_1(n) \approx [nz(L_1 + U_1) + nz(S - D_1)] \cdot k_1 = [10n^3 - 19n^2 + 4n] \cdot k_1,$$

$$(5.1b) \quad c_2(n) \approx [c_{in}(n) + nz(S - D_2)] \cdot k_2 = [c_{in}(n) + 3n^3 - 8n^2 + 4n] \cdot k_2.$$

The term $nz(X)$ stands for the number of nonzeros of a matrix X , and S stands for the reduced matrix.

PROPOSITION 5.1. *For n large enough, the scheme associated with the 2D splitting is cheaper than the one associated with the 1D splitting only if $c_{in}(n) < 15n^3$.*

Proof. If n is large enough we can use the relation $\frac{k_1}{k_2} = 1.8$ and refer only to the leading power of n in the expressions for $c_1(n)$ and $c_2(n)$. So doing, it follows that

$$(5.2) \quad \frac{c_1(n)}{c_2(n)} \approx \frac{18n^3}{c_{in}(n) + 3n^3},$$

and the result stated in the proposition readily follows. \square

What is left now is to examine the amount of work involved in solving the inner system of equations. A natural choice of a splitting for this system is $D_2 = D_1 - (D_1 - D_2)$. It is straightforward to show the following by Propositions 3.2 and 3.3.

PROPOSITION 5.2. *If block Jacobi based on the splitting $D_2 = D_1 - (D_1 - D_2)$ is used, then the spectral radius of the inner iteration matrix, namely, $I - D_1^{-1}D_2$, is bounded by $\frac{\xi}{\eta}$, where η and ξ are defined in (3.3) and (3.4).*

For considering methods that are faster than block Jacobi for the inner system, we have the following useful result.

PROPOSITION 5.3. *The inner matrix is block consistently ordered relative to 1D partitioning.*

Proof. The inner matrix is block tridiagonal relative to this partitioning. \square

We are now ready to prove the main result of this subsection.

PROPOSITION 5.4. *If $be, cd, fg > 0$, then if 1D splitting is used in solving the inner system, the cost of solving it is higher than $15n^3$ floating point operations, for block Jacobi as well as block Gauss–Seidel and block SOR, and thus, for n large enough and the methods considered in this paper, the 1D solver is faster than the 2D solver.*

Proof. The Taylor expansion of the bound in Proposition 5.2 is

$$(5.3) \quad \frac{\xi}{\eta} = \frac{4}{9} - \left(\frac{43}{81}\pi^2 + \frac{65}{648}\mu^2 + \frac{29}{648}\tau^2 + \frac{25}{324}\sigma^2 \right) h^2 + o(h^2).$$

For h small enough, we can simply examine the leading term: The bound is approximately $\frac{4}{9}$ if block Jacobi is used, and since by Proposition 5.3 the matrix is consistently ordered, Young's analysis shows that the spectral radius is approximately $\frac{16}{81}$ if block Gauss-Seidel is used and approximately 0.055 if block SOR with the optimal relaxation parameter is used. For both of these schemes each iteration costs about $7n^3$ floating point operations. Since reducing the initial error by a factor of 10^m takes roughly $-m/\log_{10} \rho$ iterations, where ρ is the spectral radius of the associated iteration matrix, it follows that even for the block SOR scheme with the optimal relaxation parameter, which is the fastest scheme considered here, after two iterations the error is reduced only by a factor of approximately $10^{2.5}$, which is obviously far from satisfactory. Thus the iteration count is larger than 2, and the cost of inner solve is larger than $15n^3$ floating point operations. \square

We remark that an inexact inner solve can also be considered (see, for example, Elman and Golub's paper on inexact Uzawa algorithms [11]), but this is beyond the scope of this work.

It is our conclusion that the solver associated with 1D splitting is more efficient than the one associated with the 2D splitting if upwind differences are used or if centered differences with mesh Reynolds numbers smaller than 1 in magnitude are used.

5.2. Comparison with the unreduced system. One step of cyclic reduction results in a more complicated difference operator compared to the original, unreduced system, and a grid which is more difficult to handle as far as ordering of the unknowns is concerned. Moreover, the unreduced matrix is block consistently ordered relative to *both* 1D and 2D splittings (we refer to the straightforward one-line and one-plane partitionings as the basis for 1D and 2D splittings in case of the unreduced system) and thus Young's analysis can be easily applied. One could ask, therefore, what the advantages of using cyclic reduction are. In this subsection we illustrate the superiority of the reduced system over the unreduced system.

We start with the block Jacobi scheme. For the unreduced system we shall refer to natural lexicographic ordering of the unknowns, so that the lines of gridpoints are x -oriented and the planes are x - y oriented. We start with quoting the following result, given in [14, Sects. 2 and 4].

LEMMA 5.5. *The spectral radius of the block Jacobi scheme associated with the 1D splitting for the unreduced system is*

$$(5.4) \quad \frac{2\sqrt{be} \cdot \cos(\pi h) + 2\sqrt{fg} \cdot \cos(\pi h)}{a + 2\sqrt{cd} \cdot \cos(\pi h)}.$$

The Taylor expansion of (5.4) about $h = 0$ is given by

$$(5.5) \quad 1 - \left(\frac{3}{4}\pi^2 + \frac{1}{16}\sigma^2 + \frac{1}{16}\tau^2 + \frac{1}{16}\mu^2 \right) h^2 + o(h^2).$$

In [14] we have shown that the spectrum of the iteration matrix of the unreduced system can be found by a sequence of diagonalizations and permutations that form

a similarity transformation of the matrix into a matrix whose associated iteration matrix is easy to analyze, as far as its spectrum is concerned. The reader is referred to the proof of [14, Thm. 2.1] for full details. For the 2D splitting a similar procedure can be applied. The technique we have used is similar to the one presented in [14] and the algebraic details are omitted.

LEMMA 5.6. *The spectral radius of the block Jacobi iteration matrix associated with 2D splitting is given by*

$$(5.6) \quad \frac{2\sqrt{fg} \cdot \cos(\pi h)}{a + 2\sqrt{cd} \cdot \cos(\pi n h) + 2\sqrt{be} \cdot \cos(\pi n h)},$$

and its Taylor expansion about $h = 0$ is

$$(5.7) \quad 1 - \left(\frac{3}{2}\pi^2 + \frac{1}{8}\sigma^2 + \frac{1}{8}\tau^2 + \frac{1}{8}\mu^2 \right) h^2 + o(h^2).$$

The same type of analysis that has been done in the previous section, comparing the 1D splitting to the 2D splitting for the reduced system, is possible for the unreduced system. Below we sketch the main details: Suppose inner-outer iterations are used in solving the scheme associated with the 2D splitting. Denote, again, this splitting for the inner system as $D_2 = D_1 - (D_1 - D_2)$ (D_1 and D_2 are now different than the ones defined in section 3). Then we have the following proposition.

PROPOSITION 5.7. *Consider the unreduced system. Suppose $be, cd, fg > 0$, n is sufficiently large, and 1D splitting is used in solving the inner system. Then for the stationary methods considered in this paper, the 1D solver is faster than the 2D solver.*

Proof. The ratio between the asymptotic rate of convergence between the 1D solver and the 2D solver is 2. The number of nonzeros of the whole matrix is approximately $7n^3$, the number of nonzeros of D_1 is approximately $3n^3$, and the number of nonzeros of D_2 is approximately $5n^3$. Since the spectral radii for the two splittings are available, we can find the spectral radius for the iteration matrix of the inner system. Its Taylor expansion is given by $\frac{1}{2} - \left(\frac{3}{8}\pi^2 + \frac{1}{16}\sigma^2 + \frac{1}{32}\tau^2 \right) h^2 + o(h^2)$. Defining cost functions analogous to the ones defined in section 3 for the reduced system, and using the same line of argument, we have

$$(5.8) \quad \frac{c_1(n)}{c_2(n)} \approx \frac{14n^3}{c_{in}(n) + 2n^3},$$

and from this it follows that only if $c_{in}(n) < 12n^3$ the 2D solver is more efficient. However, as in Proposition 5.4, this means at most two iterations of the inner solve can be performed, which is not enough for the required accuracy. \square

Since the 1D splitting for both the reduced and the unreduced systems gives rise to a more efficient solve, we compare these two systems, focusing on this splitting. See also [14, Sect. 4]. The LU decomposition for the solution of the system in each iteration is done once and for all (see [12] for operation count) and its cost is negligible in comparison with the amount of work done in the iterative process.

Each iteration in the reduced system costs about $10n^3$ floating point operations, whereas each iterate for the unreduced system costs approximately $7n^3$ floating point operations per iteration. Hence, the amount of computational work per iteration is cheaper for the unreduced system by a factor of about $10/7$. However, using the asymptotic formulas (3.6) and (5.5), it is evident that the number of iterations required for the unreduced system is larger than that required for the reduced system, and in




FIG. 5.1. Comparison between the spectral radii of the Gauss-Seidel iteration matrices of the reduced and unreduced systems. The uppermost curve corresponds to 1D splitting for the unreduced system, and then we have, in order, 2D splitting for the unreduced system, 1D splitting for the reduced system, and 2D splitting for the reduced system.

the worst case, the ratio between the work required for solving the reduced system versus the unreduced system is roughly $(10/7) \cdot (27/40)$, which is $27/28$ and is still smaller than 1, thus the reduced solver is more efficient. If the convective terms are nonzero, then this ratio becomes smaller, and in practice we have observed substantial savings, as is illustrated in the test problem discussed in section 5.3.

Moving from comparing the block Jacobi scheme for both the reduced and the unreduced systems to comparing Gauss-Seidel and SOR is straightforward if Young's analysis can be used. In section 4 we showed that even though the reduced matrix is not consistently ordered relative to 1D partitioning, it is *nearly* consistently ordered. In general, convergence analysis for the Jacobi scheme does not always indicate the behavior of the Gauss-Seidel and SOR schemes. Nevertheless, for two-cyclic consistently ordered matrices (or matrices that are nearly so) the strong connections between the spectra of the Jacobi iteration matrix and the Gauss-Seidel and SOR iteration matrices [24] allow us to conclude that once the superiority of the reduced system over the unreduced system has been shown for Jacobi, this superiority is carried over to the other stationary schemes. Indeed, our numerical experiments verify this observation, as is illustrated in section 5.3.

In Fig. 5.1 the superiority of the reduced system over the unreduced system for the Gauss-Seidel scheme is illustrated numerically. The graphs were created for a small 512-point grid. It is interesting to notice that the reduced 1D Gauss-Seidel iteration matrix is well behaved (i.e., its spectral radius is significantly smaller than 1), even for the convection-dominated case, when centered differences are used. Convergence does not occur when the block Jacobi scheme with the same values of mesh Reynolds numbers is used. We have no bounds on convergence rates for this range of mesh Reynolds numbers and thus cannot explain this phenomenon analytically.

The superiority of the reduced system is evident also for the SOR scheme (see Fig. 4.3). Notice that for the SOR scheme it is difficult to determine the optimal relaxation parameter when be , cd , and fg are negative.

We end this subsection with a remark regarding the case of convection-dominated equations. Our convergence analysis does not cover the case of mesh Reynolds num-

TABLE 5.1

Comparison between iteration counts for the reduced and unreduced system, for different values of mesh Reynolds numbers. *N/C* marks no convergence after 2,000 iterations. *GS* = Gauss-Seidel.

System		Reduced				Unreduced			
PDE coeff. ($\sigma = \tau = \mu$)		10	20	100	1000	10	20	100	1000
Jacobi	centered	393	173	53	N/C	1030	444	N/C	N/C
GS	centered	188	77	14	322	492	198	N/C	N/C
SOR	centered	36	25	-	-	61	38	-	-
Jacobi	upwind	455	239	75	43	1194	620	179	89
GS	upwind	219	111	27	10	574	287	63	16
SOR	upwind	39	27	18	9	66	45	24	11

bers that are greater than 1 in magnitude in conjunction with centered difference discretization. Since the numerical solution might be oscillatory when a centered difference scheme is used [18], analysis for this case is of less interest. Nevertheless, Fourier analysis based on Chan and Elman's technique [5], which shows that when one of the mesh Reynolds numbers tends to ∞ the scheme still converges, is presented in [13].

5.3. Test problem. Consider (1.1), where the right-hand side is such that the solution for the continuous problem is $u(x, y, z) = \sin(\pi x) \cdot \sin(\pi y) \cdot \sin(\pi z)$ and the domain is the unit cube. The Dirichlet boundary conditions in this case are zero. The performance of the solvers for this specific problem well represents the performance for other test problems that we have examined.

We have taken the zero vector as our initial guess and have used $\|r_i\|_2/\|r_0\|_2 < 10^{-10}$ as a stopping criterion (here r_i denotes the residual at the i th iterate). The program stopped if the stopping criterion was not satisfied after 2,000 iterations. Our numerical experiments were executed on an SGI Origin 2000, which has four parallel 195 MHz processors, 512 MB RAM, and 4MB cache. The program was written in MATLAB 5.

In the experiments that are presented, the 1D solver is used. In Table 5.1, the grid is of size $32 \times 32 \times 32$. The matrix of the underlying system of equations is of size $32,768 \times 32,768$. In the table, iteration counts for the Jacobi scheme and the Gauss-Seidel scheme are presented for four values of the PDE coefficients and for two discretization schemes.

The PDE coefficients referred to in Table 5.1 are specified in (1.1). For the values of these coefficients in the table, namely 10, 20, 100, and 1,000, the corresponding values of the mesh Reynolds numbers are 0.1515, 0.3030, 1.515, and 15.15. Notice that the last two are larger than 1, and so for these values we have no analytical way of knowing the optimal relaxation parameter and the experiments for these values were not performed.

The following observations can be made.

1. Overall the reduced solver is substantially faster than the unreduced solver. There are cases where the reduced solver converges whereas the unreduced solver does not. We remark that in all cases that were examined, the CPU time for the reduced solver was less (much less in most cases) than the CPU time for the unreduced system.
2. For $\sigma = 50$ convergence is faster than for $\sigma = 10$. This illustrates a phenomenon, which is supported by the analysis and holds also for the two-dimensional case [9], that for small-enough mesh Reynolds numbers, the

“more nonsymmetric” systems converge faster than the “close to symmetric” ones (close in the sense of PDE coefficients close to zero).

3. The upwind difference scheme converges more slowly than the centered difference scheme when the mesh Reynolds numbers are small in magnitude, but convergence is extremely fast for large mesh Reynolds numbers. This applies to both the reduced and the unreduced systems and follows from the fact that as the PDE coefficients grow larger, the underlying matrix is more diagonally dominant when upwind schemes are used.

6. Concluding remarks. We have presented ordering strategies for a cyclically reduced matrix arising from discretizing a 3D model problem with constant coefficients. We have derived bounds on convergence rates for block stationary schemes associated with what we called 1D splitting or 2D splitting. We have compared the amount of work involved in solving the system with the suggested splittings. In general, the 1D splitting gives rise to more-efficient solvers. Since the matrices associated with this splitting are not consistently ordered, we have analyzed their departure from block Property A and have shown that, in fact, these matrices are nearly block consistently ordered. We have shown, both analytically and numerically, that one step of cyclic reduction results in a system which is easier to solve, compared to the original, unreduced system.

Acknowledgments. We would like to thank the referees for their helpful comments, which substantially improved this manuscript.

REFERENCES

- [1] O. AXELSSON AND I. GUSTAFSSON, *On the use of preconditioned conjugate gradient methods for red-black order five-point difference schemes*, J. Comput. Phys., 35 (1980), pp. 284–289.
- [2] O. BUNEMAN, *A compact non-iterative Poisson solver*, Report 294, Stanford University Institute for Plasma Research, Stanford, CA, 1969.
- [3] B. L. BUZBEE, F. W. DORR, J. A. GEORGE, AND G. H. GOLUB, *The direct solution of the discrete Poisson equation on irregular regions*, SIAM J. Numer. Anal., 8 (1971), pp. 722–736.
- [4] B. L. BUZBEE, G. H. GOLUB, AND C. W. NIELSON, *On direct methods for solving Poisson’s equations*, SIAM J. Numer. Anal., 7 (1970), pp. 627–656.
- [5] T. F. CHAN AND H. C. ELMAN, *Fourier analysis of iterative methods for elliptic problems*, SIAM Rev., 31 (1989), pp. 20–49.
- [6] P. CONCUS AND G. H. GOLUB, *Use of fast direct methods for the efficient numerical solution of nonseparable elliptic equations*, SIAM J. Numer. Anal., 10 (1973), pp. 1103–1120.
- [7] E. DETYNA, *Point cyclic reductions for elliptic boundary-value problems I: The constant coefficient case*, J. Comput. Phys., 33 (1979), pp. 204–216.
- [8] H. C. ELMAN AND G. H. GOLUB, *Iterative methods for cyclically reduced non-self-adjoint linear systems*, Math. Comp., 54 (1990), pp. 671–700.
- [9] H. C. ELMAN AND G. H. GOLUB, *Iterative methods for cyclically reduced non-self-adjoint linear systems II*, Math. Comp., 56 (1991), pp. 215–242.
- [10] H. C. ELMAN AND G. H. GOLUB, *Line iterative methods for cyclically reduced discrete convection-diffusion problems*, SIAM J. Sci. Stat. Comput., 13 (1992), pp. 339–363.
- [11] H. C. ELMAN AND G. H. GOLUB, *Inexact and preconditioned Uzawa algorithms for saddle point problems*, SIAM J. Numer. Anal., 31 (1994), pp. 1645–1661.
- [12] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 3rd ed., Johns Hopkins University Press, Baltimore, MD, 1996.
- [13] C. GREIF, *Analysis of Cyclic Reduction for the Numerical Solution of Three-Dimensional Convection-Diffusion Equations*, Ph.D. thesis, University of British Columbia, Vancouver, BC, 1998.

- [14] C. GREIF AND J. M. VARAH, *Iterative solution of cyclically reduced systems arising from discretization of the three-dimensional convection diffusion equation*, SIAM J. Sci. Comput., (1998), pp. 1918–1940.
- [15] L. A. HAGEMAN, F. T. LUK, AND D. M. YOUNG, *On the equivalence of certain iterative acceleration methods*, SIAM J. Numer. Anal., 17 (1980), pp. 852–873.
- [16] L. A. HAGEMAN AND R. S. VARGA, *Block iterative methods for cyclically reduced matrix equations*, Numer. Math., 6 (1964), pp. 106–119.
- [17] R. W. HOCKNEY, *A fast direct solution of Poisson's equation using Fourier analysis*, J. Assoc. Comput. Mach., (1965), pp. 95–113.
- [18] K. W. MORTON, *Numerical Solution of Convection-Diffusion Problems*, 1st ed., Chapman and Hall, London, 1996.
- [19] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, PWS Publishing Company, Boston, 1996.
- [20] R. A. SWEET, *A generalized cyclic reduction algorithm*, SIAM J. Numer. Anal., 11 (1974), pp. 506–520.
- [21] R. A. SWEET, *A cyclic reduction algorithm for solving block tridiagonal systems of arbitrary dimension*, SIAM J. Numer. Anal., 14 (1977), pp. 706–720.
- [22] R. S. VARGA, *p-Cyclic matrices: A generalization of the Young-Frankel successive over-relaxation scheme*, Pacific J. Math., 9 (1959), pp. 617–628.
- [23] R. S. VARGA, *Matrix Iterative Analysis*, Prentice-Hall, Englewood Cliffs, NJ, 1962.
- [24] D. M. YOUNG, *Iterative Solution of Large Linear Systems*, Academic Press, New York, 1971.