A Box Shaped Cyclically Reduced Operator*

Chen Greif† L. Robert Hocking‡

Abstract

A new procedure of cyclic reduction is proposed, whereby instead of performing a step of elimination on the original cartesian mesh using a two-color ordering and a standard 5-point or 7-point operator, we perform the decoupling step on the reduced mesh associated with one color, using non-standard operators that are better aligned with that mesh. This yields a cartesian mesh and box shaped 9-point (in 2D) or 27-point (in 3D) operators that are easy to deal with. Convergence analysis for multi-line and multi-plane orderings is carried out. Numerical experiments demonstrate the merits of the approach taken.

1 Introduction

The technique of cyclic reduction has been studied and analyzed in several settings in the last few decades. Early work demonstrates the merits of applying this procedure to the discrete Poisson equation associated with standard finite differences on a uniform mesh; see [4] for a review of history and applications, and references therein. The nonzero pattern of the matrix in this case allows for efficiently eliminating half of the unknowns, while preserving the block structure. The procedure can be applied repeatedly until a small system that can be easily solved is obtained. Recovering the solution for the unknowns that have been eliminated throughout the process is straightforward, and the overall computational cost is attractively low.

In the nonsymmetric case, for example in the case of the discrete convection-diffusion equation (on which we focus in this work), some of the attractive features of the Laplacian are lost and a recursive cyclic reduction approach may be numerically unstable. In the early 1990s Elman & Golub offered a thorough analysis of the spectral properties and convergence behavior of linear systems arising from a procedure of *one step of cyclic reduction* [1, 2, 3]. Using red-black ordering and eliminating all the ones corresponding to one of the two

*This work was supported in part by the Natural Sciences and Engineering Research Council of Canada.

†Department of Computer Science, The University of British Columbia, Vancouver, BC, V6T 1Z4, CANADA, greif@cs.ubc.ca.

‡Department of Computer Science, The University of British Columbia, Vancouver, BC, V6T 1Z4, CANADA, rob.l.hocking@gmail.com.

colors, results in a linear system associated with a diamond shaped 9-point computational molecule, in contrast to the original unreduced operator, which has the well known plus shaped 5-point stencil.

The Schur complement (reduced) matrix obtained after the elimination is only half the dimension of the original (unreduced) matrix. Thus, the cost of performing matrix-vector products on the reduced system is similar to, or in fact marginally lower than the cost of matrix-vector products on the unreduced system. As a result, when comparing the iterative solution procedure for the original system to that for the cyclically reduced system, the overall performance of solvers depends almost exclusively on the spectral structure of the reduced vs. the unreduced operators, and not on the cost of a single iteration. The analysis and numerical experiments in [1, 2, 3] show that iterative solvers for the reduced system converge faster, and hence it pays off to perform one step of cyclic reduction in the 2D case.

In the late 1990s, Greif & Varah [7, 6, 8] showed that gains can be made for the three-dimensional case as well. However, in 3D the original 7-point operator is replaced by a 19-point one, applied to half of the unknowns, and hence matrix-vector products for the reduced system are more expensive, in contrast to the 2D case. Despite that, the improvement in the spectral structure and convergence rates lead to the conclusion that in 3D it still pays off to perform a step of cyclic reduction.

Despite their attractive numerical properties and the computational savings, cyclically reduced operators of the form just described have not been widely used. One possible reason for this is the nature of the computational molecule and the mesh, which present a computational challenge, for example in handling boundary conditions. Applying one step of cyclic reduction to a cartesian mesh, using a standard 5-point operator (in 2D) or 7-point operator (in 3D), yields a non-cartesian mesh with 'holes'. The reduced stencil includes gridpoints not immediately next to each other.

In this paper we propose to perform a step of decoupling on the reduced mesh. We show that one step of cyclic reduction on a non-cartesian mesh can be done by using a non-standard operator aligned with that mesh, and yields a (further) reduced mesh that is cartesian again, associated with a standard box shaped operator. It is this last observation which forms the basis for the procedure we propose.

The basic idea is quite simple and amounts to adopting a bottom-top approach, in the sense that the reduced grid is handled before the unreduced grid is taken care of. Instead of carrying out the step of cyclic reduction on the cartesian unreduced mesh with a standard operator, we first decouple the gridpoints using two colors, but postpone any step of elimination for later. We define a two-color ordering on the reduced grid, and deal with the gridpoints that belong to one of the colors by using non-standard operators that are naturally aligned with the grid. In 3D this is more complicated and involves a second decoupling step, but the idea is similar. The result is a grid that now contains only $1/2^d$ of the original gridpoints (with $d = 2, 3$ being the dimension), but this grid is again cartesian, as was the original grid, and the operator associated with

it is box shaped: it looks like the standard 9-point operator in 2D, and the standard 27-point operator in 3D. Boundary conditions can now be straightforwardly handled. Implementation is therefore straightforward, and one benefit from the fact that the mesh is 1/4 and and 1/8 of the original one, in 2D and 3D respectively, is that matrix-vector products are cheaper.

We use a block ordering strategy that exploits the structure of the stencil. Specifically, we consider k-line and k-plane orderings in 2D and 3D respectively, where k is a small integer: 2 or 3. We derive exact analytic expressions for the spectral radius of the block Jacobi iteration matrix with lexicographic ordering (which is equivalent to 1-line or 1-plane ordering), and tight bounds for a 2-line (2-plane) ordering in 2D (3D). We show that the latter orderings are superior to lexicographic ordering. Our experiments suggest that a 3-line ordering is even more effective, although we are not able to obtain tight bounds in this case. Since the matrices in question are consistently ordered with respect to the block partitioning we consider, the analytic bounds carry over to block Gauss-Seidel and block SOR. In our numerical results we also consider ILU preconditioned GMRES iterations.

The remainder of this paper is organized as follows. In Section 2 we introduce the model convection-diffusion problem and derive our new operator. In Section 3 we discuss the notion of block grids. In Section 4 we apply multi-line and multi-plane orderings for the 2D and 3D cases respectively, and offer a convergence analysis for block Jacobi. In Section 5 we provide numerical comparisons between our operators, the "traditional" cyclically reduced operators, and the standard unreduced 5-point and 7-point operators. Finally, in Section 6 we draw some conclusions.

2 The Proposed Procedure of Cyclic Reduction

Consider the two- or three-dimensional convection-diffusion model problem with constant coefficients:

$$-\Delta u + \vec{w} \cdot \nabla u = f. \tag{1}$$

The domain $\Omega = (0,1)^d$ is the unit square ($d = 2$) or unit cube ($d = 3$), subject to Dirichlet type boundary conditions. When the row vector $\vec{w}$ is constant, we denote its components by (σ, τ) for 2D and (σ, τ, μ) for 3D.

We discretize the problem using a uniform rectangular grid G in Ω. Our method requires the number of grid points to be odd; we thus set $G = \{h, 2h, ..., (2n+1)h\}^d$ where $h = \frac{1}{2n+2}$. The mesh Reynolds numbers are defined as

$$\gamma = \frac{\sigma h}{2}, \qquad \delta = \frac{\tau h}{2}, \qquad \eta = \frac{\mu h}{2}.$$

2.1 The Reduction Step in 2D

We will be using a combination of two different discrete operators, as follows. Denote the standard plus shaped second order 5-point discretization of (1) in

2D, after scaling by h^2, by

$$F^+u_{i,j} = 4u_{i,j}+(-1-\delta)u_{i,j-1}+(-1-\gamma)u_{i-1,j}+(-1+\gamma)u_{i+1,j}+(-1+\delta)u_{i,j+1}.$$

Denote by $F^\times$ an $\times$ shaped second order 5-point discretization of (1):

$$F^\times u_{i,j} = au_{i,j} + bu_{i+1,j+1} + cu_{i-1,j+1} + du_{i-1,j-1} + eu_{i+1,j-1},$$

after scaling by $2h^2$. We have in this case

$$a = 4,\ b = -1+\gamma+\delta,\ c = -1-\gamma+\delta,\ d = -1-\gamma-\delta,\ e = -1+\gamma-\delta. \quad (2)$$

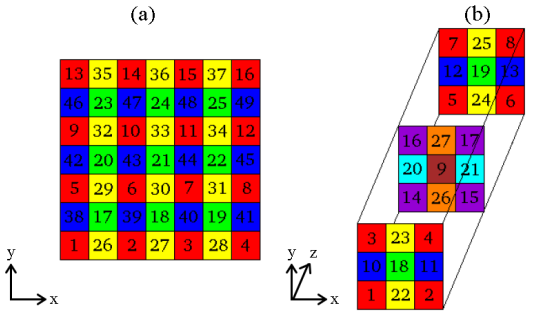

Figure 1: (a) Four color ordering applied to a 7×7 grid. (b) Eight color ordering applied to a $3 \times 3 \times 3$ grid.

Suppose we apply a four-color ordering to the grid G, as in Figure 2.1(a). Red points (i, j) have both i and j odd; green is for i, j even; blue is for i odd and j even; and yellow signifies i even and j odd.

We then discretize (1) by applying F^+ at the yellow and blue points, and $F^\times$ at the red and green points.

The resulting $(2n + 1)^2 \times (2n + 1)^2$ linear system can be written in block form as

$$\begin{pmatrix} D_1 & B_1 & 0 & 0 \\ B_2 & D_2 & 0 & 0 \\ B_3 & B_4 & D_3 & 0 \\ B_5 & B_6 & 0 & D_4 \end{pmatrix} \begin{pmatrix} u^{red} \\ u^{green} \\ u^{blue} \\ u^{yellow} \end{pmatrix} = \begin{pmatrix} 2h^2 f^{red} \\ 2h^2 f^{green} \\ h^2 f^{blue} \\ h^2 f^{yellow} \end{pmatrix},$$

where the matrices D_1, D_2, D_3, and D_4 are diagonal. We apply a block elimination procedure to obtain a reduced (Schur complement) system of size $n^2 \times n^2$, involving only the green points:

$$(D_2 - B_2D_1^{-1}B_1)u^{green} = 2h^2(f^{green} - B_2D_1^{-1}f^{red}). \quad (3)$$

The remaining three quarters of the unknowns can be recovered by computing the following three equations, where the inversions are trivial because $D_i,\ i =$

$1, 3, 4$, are diagonal:

$$\begin{array}{lcl} u^{red} & = & D_1^{-1}(2h^2 f^{red} - B_1 u^{green}), \\ u^{blue} & = & D_3^{-1}(h^2 f^{blue} - B_3 u^{red} - B_4 u^{green}), \\ u^{yellow} & = & D_4^{-1}(h^2 f^{yellow} - B_5 u^{red} - B_6 u^{green}). \end{array}$$

The $n^2 \times n^2$ reduced matrix in (3) is block tridiagonal with tridiagonal blocks:

$$\begin{array}{lcl} D_2 - B_2 D_1^{-1} B_1 & = & \mathrm{tri}_{n\times n}[\mathrm{tri}_{n\times n}[-d^2, -2de, -e^2], \\ & & \quad \mathrm{tri}_{n\times n}[-2cd, a^2 - 2bd - 2ce, -2be], \\ & & \quad \mathrm{tri}_{n\times n}[-c^2, -2bc, -b^2]]. \end{array} \tag{4}$$

Notice that the 4×4 block system that we started with was reducible. Indeed, the use of four colors above has been done mainly for the purpose of illustration. In practice, it is sufficient to think of the blue and yellow points as points of one color (black, if traditional red-black ordering is considered): dealing with them is postponed for later, and we start off with the red and green points. In terms of traditional two-color orderings, these latter points were originally all red, but we have re-colored them, defining the green color and then eliminating the remaining red points. The operator for the green points looks like a standard 9-point operator, as opposed to the diamond shaped 9-point typically cyclically reduced operator. We observe one of its properties as a discrete differential operator, as follows.

Proposition 1 *The difference equation for the operator for the green points is equivalent to standard finite difference discretization of the differential equation*

$$-\left[\left(1 + \frac{\sigma^2 h^2}{4}\right) u_{xx} + \left(1 + \frac{\tau^2 h^2}{4}\right) u_{yy}\right] + \sigma u_x + \tau u_y = f + O(h^2).$$

Proof. After expanding each $u_{i+\Delta i, j+\Delta j}$ term using a Taylor series about $u_{i,j}$, the left hand side of (3) reduces, after dividing by $16h^2$, to

$$\begin{array}{rl} \sigma u_x + \tau u_y - (1 + \frac{\sigma^2 h^2}{4}) u_{xx} - (1 + \frac{\tau^2 h^2}{4}) u_{yy} - \frac{\sigma\tau h^2}{2} u_{xy} & \\ + \frac{2\sigma h^2}{3} u_{xxx} + \frac{2\tau h^2}{3} u_{yyy} + \tau h^2 u_{xxy} + \sigma h^2 u_{xyy} & \\ - \frac{h^2}{3} u_{xxxx} - \frac{h^2}{3} u_{yyyy} - h^2 u_{xxyy} + o(h^2). & \end{array}$$

Similarly, the right hand side reduces to

$$f - \frac{\sigma h^2}{4} f_x - \frac{\tau h^2}{4} f_y + \frac{h^2}{4} \Delta f + o(h^2).$$

This computation was done using MATLAB's symbolic toolbox. □

Proposition 1 shows that the new cyclically reduced operator is in fact a second order operator for the 2D convection-diffusion equation. The additional

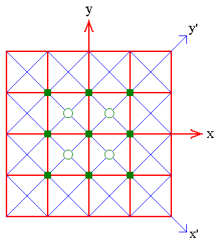

Figure 2: The standard 2D cyclically reduced operator defined on a mesh with spacing $\sqrt{2}h$ becomes a compact rectangular operator when viewed on a grid rotated 45 degrees with mesh spacing h.

$O(h^2)$ terms on the left hand side represent the introduction of artificial dissipation, and this suggests the operator may be better behaved for high mesh Reynolds numbers, compared to the original operator. In this regard, this operator satisfies properties similar to the ones for the "traditional" cyclically reduced operators.

We now make the point that in fact the new operator is nothing but a 'tilted' version of the typical cyclically reduced operator for a related equation; we note that this is *not* the case for 3D.

Proposition 2 *Suppose the new cyclic reduction process is applied to the 2D convection-diffusion problem, yielding a box shaped 9 point operator. Then the operator is identical to the standard cyclically reduced operator applied to a grid rotated 45 degrees clockwise with mesh spacing* $\sqrt{2}h$.

The proof follows from transforming to the rotated grid, where the convection-diffusion equation becomes

$$-\Delta u + \frac{1}{\sqrt{2}}(\sigma - \tau, \sigma + \tau) \cdot \nabla u = f,$$

and discretizing with mesh spacing $\sqrt{2}h$ using the standard 5-point operator; see Figure 2 for an illustration, where the original and transformed coordinates are denoted, respectively, by (x, y) and (x', y').

2.2 The Reduction Step in 3D

In three dimensions, our elimination procedure requires the use of the standard second order 7-point operator $\hat{F}$, as well as three nonstandard second order operators $\tilde{F}$, $\grave{F}$, and F. After scaling by h^2, $2h^2$, $2h^2$ and $4h^2$ respectively,

these operators are

$$\begin{aligned}
\hat{F}u_{i,j,k} &= 6u_{i,j,k} + (-1+\gamma)u_{i+1,j,k} + (-1-\gamma)u_{i-1,j,k} + (-1+\delta)u_{i,j+1,k} \\
&\quad +(-1-\delta)u_{i,j-1,k} + (-1+\eta)u_{i,j,k+1} + (-1-\eta)u_{i,j,k-1} \ , \\
\tilde{F}u_{i,j,k} &= 8u_{i,j,k} + (-1+\gamma+\delta)u_{i+1,j+1,k} + (-1-\gamma+\delta)u_{i-1,j+1,k} \\
&\quad +(-1-\gamma-\delta)u_{i-1,j-1,k} + (-1+\gamma-\delta)u_{i+1,j-1,k} \\
&\quad +2(-1+\eta)u_{i,j,k+1} + 2(-1-\eta)u_{i,j,k-1} \ , \\
\grave{F}u_{i,j,k} &= 8u_{i,j,k} + (-1-\gamma+\eta)u_{i-1,j,k+1} + (-1+\gamma+\eta)u_{i+1,j,k+1} \\
&\quad +(-1-\gamma-\eta)u_{i-1,j,k-1} + (-1+\gamma-\eta)u_{i+1,j,k-1} \\
&\quad +2(-1+\delta)u_{i,j+1,k} + 2(-1-\delta)u_{i,j-1,k} \ ,
\end{aligned}$$

and

$$\begin{aligned}
Fu_{i,j,k} &= au_{i,j,k} + bu_{i+1,j+1,k+1} + cu_{i-1,j+1,k+1} + du_{i-1,j-1,k+1} + eu_{i+1,j-1,k+1} \\
&\quad +pu_{i+1,j+1,k-1} + qu_{i-1,j+1,k-1} + ru_{i-1,j-1,k-1} + su_{i+1,j-1,k-1} \ .
\end{aligned}$$

Here $a = 8,\ b = -1+\gamma+\delta+\eta,\ c = -1-\gamma+\delta+\eta,\ d = -1-\gamma-\delta+\eta,\ e = -1+\gamma-\delta+\eta,\ p = -1+\gamma+\delta-\eta,\ q = -1-\gamma+\delta-\eta,\ r = -1-\gamma-\delta-\eta,\ s = -1+\gamma-\delta-\eta.$ See Table 1 for a classification of grid point colors in terms of i, j, and k.

We order the mesh G with an eight-color ordering as in figure 2.1(b). Equation (1) is then discretized by applying F to red and brown points, $\hat{F}$ to cyan and yellow points, $\grave{F}$ to orange and blue points, and $\tilde{F}$ to green and purple points.

For each color $c \in C \equiv$ {red, brown, orange, blue, purple, green, cyan, yellow}, we obtain a system

$$D_c u^c + \sum_{x \in C, x \neq c} B_{c,x} u^x = \alpha h^2 f^c,$$

where the matrix D_c is diagonal and $\alpha = 1, 2$, or 4, depending on which operator was applied to gridpoints of color c. As can be seen in Table 1, unknowns of a given color only depend on unknowns of a few other colors, so most of the $B_{c,x}$'s are zero.

In fact, as soon as u is known for all the brown points it can be found for every other color by inverting a diagonal matrix and performing a few matrix vector products – simply apply

$$u^c = D_c^{-1} \left(\alpha h^2 f^c - \sum_{x \in C, x \neq c} B_{c,x} u^x \right)$$

to the red, green, purple, blue, orange, yellow, and cyan points in that order.

A block elimination procedure completely analogous to the 2D case gives an $n^3 \times n^3$ reduced system involving only the brown unknowns. The reduced matrix

$$A = \mathrm{tri}_{n \times n}[A_1, A_2, A_3] \tag{5}$$

Color	i	j	k	Depends On
red	odd	odd	odd	red, brown
brown	even	even	even	brown, red
green	even	even	odd	green, red, brown
purple	odd	odd	even	purple, red, brown
blue	odd	even	odd	blue, red, brown
orange	even	odd	even	orange, red, brown
yellow	even	odd	odd	yellow, red, green, orange
cyan	odd	even	even	cyan, blue, purple, brown

Table 1: Classification of grid colors in terms of i, j, and k. Dependence of grid unknowns of a given color on grid unknowns of other colors, after discretizing (1) as above.

is block triagonal, with each block itself block tridiagonal with tridiagonal blocks. We have

$$\begin{aligned} A_1 &= \mathrm{tri}_{n\times n}[\mathrm{tri}_{n\times n}[-r^2,-2rs,-s^2],\\ &\qquad \mathrm{tri}_{n\times n}[-2qr,-2(qs+pr),-2ps],\\ &\qquad \mathrm{tri}_{n\times n}[-q^2,-2qp,-p^2]]\ ;\\ \\ A_2 &= \mathrm{tri}_{n\times n}[\mathrm{tri}_{n\times n}[-2dr,-2(re+ds),-2es],\\ &\qquad \mathrm{tri}_{n\times n}[-2(cr+qd),a^2-2br-2cs-2dp-2eq,-2(pe+bs)],\\ &\qquad \mathrm{tri}_{n\times n}[-2cq,-2(cp+bq),-2bp]]\ ;\\ \\ A_3 &= \mathrm{tri}_{n\times n}[\mathrm{tri}_{n\times n}[-d^2,-2de,-e^2],\\ &\qquad \mathrm{tri}_{n\times n}[-2cd,-2(bd+ce),-2be],\\ &\qquad \mathrm{tri}_{n\times n}[-c^2,-2bc,-b^2]]\ . \end{aligned}$$

It is tedious but straightforward to show that as long as $\|\frac{h}{2}\vec{w}\|_1 < 1$, the reduced matrices (4) and (5) are strictly diagonally dominant irreducible M-matrices.

A Taylor expansion similar to the one in 2D shows that the 3D operator has similar properties and can be interpreted as a discretization of the convection-diffusion equation with some artificial viscosity. More precisely, the difference equation for the reduced system for the brown points is equivalent to standard finite difference discretization of the differential equation

$$-\left[\left(1+\frac{\sigma^2h^2}{4}\right)u_{xx}+\left(1+\frac{\tau^2h^2}{4}\right)u_{yy}+\left(1+\frac{\mu^2h^2}{4}\right)u_{zz}\right]+\sigma u_x+\tau u_y+\mu u_z = f+O(h^2).$$

The proof is almost identical to the 2D case, Proposition 1, and is omitted.

On the other hand, it is not the case that there is a clear connection between this 3D operator and the typical cyclically reduced 3D operator. In this regard there is a fundamental difference between the 2D case and the 3D case.

In Section 4 we will show that in the case that $\vec{w}$ is aligned with the x-axis, both systems are symmetrizeable by a diagonal symmetrizer, and are positive definite.

3 Block Grid Orderings

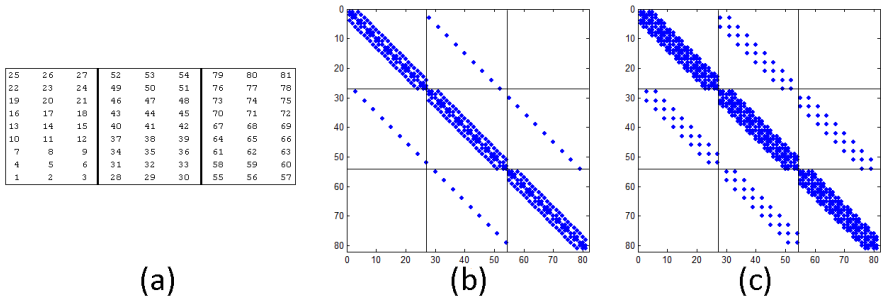

Figure 3: (a) 3-line ordering applied to a 9×9 grid. The grid blocks are separated by lines. (b) The induced block structure of the 5-point operator matrix. (c) The induced block structure of the 9-point operator matrix.

In some cases it may be beneficial to use grid orderings that generate dense diagonal blocks; see, e.g. [9]. For this purpose, multi-line and multi-plane block orderings have been developed. See [5] for an analysis of 2-line and 2-plane orderings applied to the standard centered differences discretizations of (1).

We explore the effect of k-line and k-plane orderings on the convergence of iterative solvers applied to 4 and 5. For convenience, let us define k-line and k-plane orderings within the context of general block grid orderings.

An ordering on a grid G is a function $i_G : G \to \{1, 2, ..., |G|\}$. A block ordering on a grid G is defined by a set of grid blocks $\{B\}$ partitioning G, an ordering of the grid blocks, and an ordering of the grid elements within each block. The block ordering is given by

$$i_G(p) = i_{B_I}(p) + \sum_{j=1}^{I-1} |B_j|,$$

where B_I is the grid block p belongs to, and $i_{B_I}(p)$ is the order of p with respect to B_I.

Unless stated otherwise, all block orderings considered in this paper will use lexicographic ordering for both the grid blocks, and the elements within each block. Note, however, that the resulting ordering is not lexicographic; see Fig. 3 and 4.

Suppose we work on an $n \times n$ grid G_{2D}, with n an integer multiple of k. The k-line ordering is defined by partitioning G_{2D} into $\frac{n}{k}$ vertical blocks containing k lines of n grid points each, as in figure 3(a) for $k = 3$, $n = 9$.

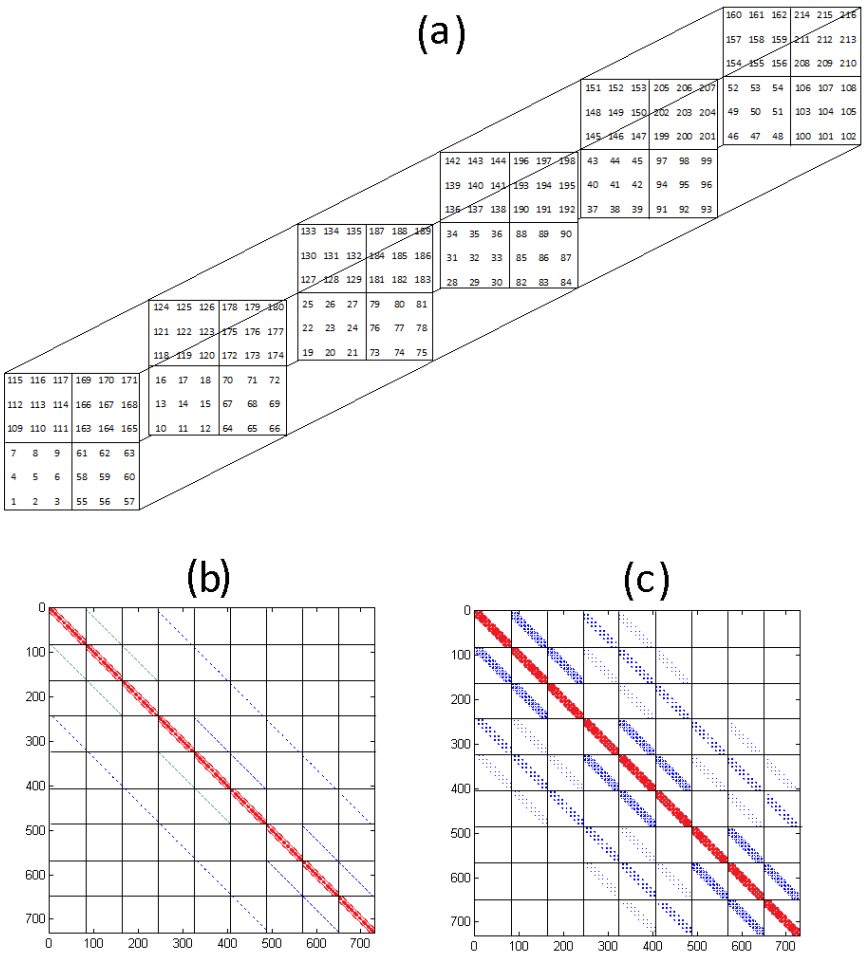

Figure 4: (a) 3-plane ordering applied to a $6 \times 6 \times 6$ grid. (b) The block matrix A_{3L} of the 5-point operator on a $9 \times 9 \times 9$ grid. (c) The block matrix A_{3L} of the 9-point operator on a $9 \times 9 \times 9$ grid.

Suppose now we work on an $n \times n \times n$ grid G_{3D} with n again an integer multiple of k. The k-plane ordering is defined by partitioning G_{3D} into $\frac{n^2}{k^2}$ blocks containing nk^2 grid points each, as in figure 4(a) for $k = 3$, $n = 6$.

Let A denote the matrix obtained by discretizing (1) on a grid G with a block ordering. A has a natural block structure that it inherits from the block ordering of G. For each ordered pair of grid blocks (B_i, B_j), there is a matrix block A_{ij} containing all the dependencies of members of B_i on members of B_j.

We denote by A_{kL} and A_{kP} the matrices that correspond to k-line and k-plane orderings. For k-line ordering, there are $\frac{n}{k}$ grid blocks of nk grid points each. A_{kL} can therefore be referred to as a block matrix with $\frac{n}{k} \times \frac{n}{k}$ blocks, each of size $nk \times nk$. For the k-plane ordering, there are $\frac{n^2}{k^2}$ grid blocks containing nk^2 points each. Therefore A_{kP} is a $\frac{n^2}{k^2} \times \frac{n^2}{k^2}$ block matrix with blocks of size $nk^2 \times nk^2$.

The sparsity pattern of the matrix depends, of course, also on the discretization used. In Fig. 3(b) and 3(c) we show the matrices corresponding to 3-line ordering, namely A_{3L}, for 5-point and 9-point discretizations. In Fig. 4(b) and 4(c) we show the matrices for 3-plane ordering, A_{3P}, using 7-point and 27-point stencils respectively.

4 Ordering Strategies and Bounds on Convergence Rates

Our analysis makes extensive use of Kronecker products and their properties, and of known results of spectra of tridiagonal Toeplitz matrices. The following two elementary results are used so often in our analysis that they are worth stating explicitly.

Proposition 3 *The eigenvalues of the tridiagonal matrix* $\mathrm{tri_{n\times n}[a, b, c]}$ *are given by*

$$\lambda_i = b + 2\sqrt{ac}\cos\left(\frac{i\pi}{n+1}\right), \qquad i = 1, 2, ..., n.$$

Proposition 4 *Suppose A and B are square matrices of sizes $n \times n$ and $m \times m$ respectively, with respective eigenpairs $\{\lambda_i, \vec{x}_i\}$, $i = 1, 2, ..., n$ and $\{\mu_j, \vec{y}_j\}$, $j = 1, 2, ...m$. Then $A \otimes B$ is $nm \times nm$ with eigenpairs $\{\lambda_i \mu_j,\ \vec{x} \otimes \vec{y}\}$, $i = 1, 2, ..., n,\ j = 1, 2, ..., m$.*

We will also be using the following technique (used effectively for analysis of the standard cyclically reduced operator). Let M and N be symmetric matrices, with M positive definite. Then

$$\rho(M^{-1}N) \leq \|M^{-1}\|_2 \|N\|_2 = \frac{\rho(N)}{\lambda_{min}(M)}. \tag{6}$$

The general case seems difficult to analyze; in our analysis we will consider the simpler case where convection is aligned with the x-axis, namely $\tau = \eta = 0$.

4.1 2D Case: k-Line Ordering

Consider the block Jacobi splitting $A_{kL} = M_{kL} - N_{kL}$, with M_{kL} consisting of all the blocks $(A_{kL})_{ii}$ discussed in section 3.

As will be evident from the analysis, it is useful to consider the general setting of a uniform $n_x \times n_y$ grid, to establish the eigenvalues of A_{kL}; we will then resort back to $n_x = n_y = n$.

The matrix A_{1L} is block tridiagonal and has dimensions $n_x \times n_x$. Each block is itself tridiagonal, with dimensions $n_y \times n_y$. Since $\tau = 0$, we have $b = e$ and $c = d$, and thus the matrix of (4) reduces to

$$\begin{aligned} A_{1L} &= \mathrm{tri}_{n_x \times n_x}[\mathrm{tri}_{n_y \times n_y}[-c^2, -2bc, -b^2], \\ &\qquad \mathrm{tri}_{n_y \times n_y}[-2c^2, a^2 - 4bc, -2b^2], \\ &\qquad \mathrm{tri}_{n_y \times n_y}[-c^2, -2bc, -b^2]] \\ &= a^2 I_{n_x n_y} + \mathrm{tri}_{n_x \times n_x}[1, 2, 1] \otimes \mathrm{tri}_{n_y \times n_y}[-c^2, -2bc, -b^2]. \end{aligned}$$

Proposition 5 *The eigenvalues of the matrix A_{kL} are given by*

$$\lambda_{i,j}(A_{kL}) = a^2 - 4bc\left[1 + \cos\left(\frac{i\pi}{n_y + 1}\right)\right]\left[1 + \cos\left(\frac{j\pi}{n_x + 1}\right)\right], \qquad i = 1, 2, ..., n_y,\ j = 1, 2, ..., n_x.$$

Proof. Using Propositions 3 and 4, the eigenvalues for A_{1L} are easily calculated, and since all the A_{kL} are permutations of each other, the result follows. □

The eigenvalues of M_{kL} can now be found as a simple corollary; we will set $n_x = n_y = n$ from this point on. Since the matrices M_{kL} are *not* permutations of each other, there is an explicit dependence on k.

Corollary 1 *The eigenvalues of M_{kL} are given by*

$$\lambda_{i,j}(M_{kL}) = a^2 - 4bc\left[1 + \cos\left(\frac{i\pi}{k + 1}\right)\right](1 + \cos(2\pi jh)), \qquad i = 1, 2, ..., k, \quad j = 1, 2, ..., n,$$

each with multiplicity $\frac{n}{k}$. The minimal eigenvalue of M_{kL} is given by

$$\lambda_{\min}(M_{kL}) = a^2 - 4bc\left[1 + \cos\left(\frac{\pi}{k + 1}\right)\right](1 + \cos(2\pi h)).$$

Also, M_{kL} is symmetrizable, and the symmetrized matrix is positive definite.

Proof. M_{kL} is block diagonal with $\frac{n}{k}$ identical blocks of size kn. The structure of each block $(M_{kL})_{ii}$ is that of the matrix A_{1L} discretized on a $k \times n$ grid.

Substituting $m = k$ and $h = \frac{1}{2n+2}$ into proposition 5, there are $\frac{n}{k}$ blocks and the minimal eigenvalue is attained when $i = j = 1$. From (2) it follows that $\lambda_{\min}$ is positive. Since M_{kL} is a Kronecker product of tridiagonal Toeplitz matrices, it is symmetrizable by a diagonal similarly transformation. □

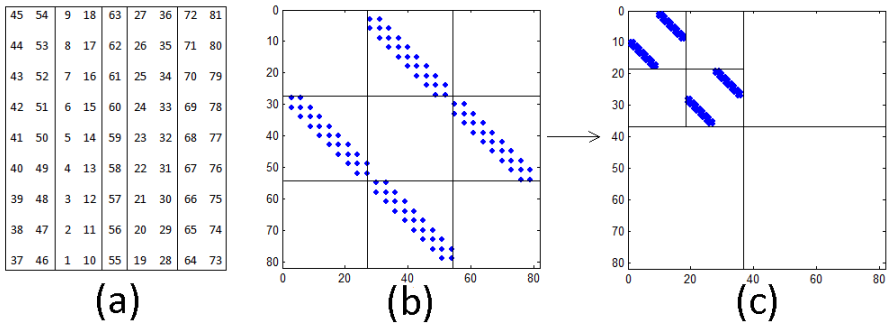

Figure 5: When unknowns are reordered as in (a), the matrix N_{kL} (b) takes on block diagonal form (c).

Proposition 6 *The eigenvalues of N_{kL} are as follows. When $k = 1$,*

$$\lambda_{i,j}(N_{1L}) = 4bc(1 + \cos(2\pi ih))\cos(2\pi jh), \qquad i, j = 1, 2, ..., n.$$

When $1 < k \leq n$, N_{kL} has $2n(\frac{n}{k} - 1)$ eigenvalues of the form

$$\lambda_i(N_{kL}) = \pm 2bc(1 + \cos(2\pi ih)), \qquad i = 1, 2, ..., n,$$

each of multiplicity $\frac{n}{k} - 1$. The remaining $n^2 - 2n(\frac{n}{k} - 1)$ eigenvalues are zero.

Proof. For $k = 1$ we have

$$N_{1L} = \mathrm{tri}_{n\times n}[1, 0, 1] \otimes \mathrm{tri}_{n\times n}[c^2, 2bc, b^2], \tag{7}$$

and the eigenvalues follow from Propositions 3 and 4.

Let $\partial_R B_i$ and $\partial_L B_i$ denote the boundary of B_i on the right and left respectively. We partition G_{2D} into $\frac{n}{k}$ new grid blocks $\tilde{B}_i$ defined by

$$\tilde{B}_i = \partial_R B_i \bigcup \partial_L B_{i+1}, \qquad i = 1, 2, ..., \frac{n}{k} - 1,$$

with $\tilde{B}_{\frac{n}{k}} = G_{2D} - \bigcup_{i=1}^{\frac{n}{k}-1} \tilde{B}_i$. The result is shown in figure 5, on a 9×9 grid which originally had 3-line ordering.

The matrix $\tilde{N}_{kL}$ has $\frac{n}{k} - 1$ diagonal blocks, $\mathrm{tri}_{2\times 2}[c^2, 0, b^2] \otimes \mathrm{tri}_{n\times n}[1, 2, 1]$. The expression for the nonzero eigenvalues follows from Propositions 3 and 4. The final diagonal block of $\tilde{N}_{kL}$ consists entirely of zeros and is of size $n^2 - 2n(\frac{n}{k} - 1)$, showing the existence and multiplicity of zero eigenvalues, as claimed. □

Corollary 2 *The spectral radii of N_{kL} are given by*

$$\rho(N_{1L}) = 4bc(1+\cos(2\pi h))\cos(2\pi h)\ ; \quad \rho(N_{kL}) = 2bc(1+\cos(2\pi h)),\ 1 < k < n.$$

We note that the case $k = n$ is a triviality, since $N_{nL} = 0$. We are now in a position to make a statement regarding the convergence of block Jacobi.

Theorem 1 *The spectral radius of the block Jacobi iteration matrix obeys the bound*

$$\rho_{kL} \leq \tilde{\rho}_{kL} = \frac{2bc(1+\cos(2\pi h))}{a^2 - 4bc(1+\cos\left(\frac{\pi}{k+1}\right))(1+\cos(2\pi h))}[1+\delta_{1k}(2\cos(2\pi h)-1)].$$

In the case $k = 1$, the bound is attained, that is $\rho_{1L} = \tilde{\rho}_{1L}$.

Proof. By Corollary 1, M_{kL} can be symmetrized, and its symmetrized version is also positive definite. It is easy to show that N_{kL} can also be symmetrized, and hence (6) can be used. Thus, by Corollaries 1, 2 and Eq. (6), the stated bound holds.

For $k = 1$ we can obtain an *exact* expression, as follows. We have $M_{1L} = a^2 I_{n^2} - 2I_n \otimes \mathrm{tri}_{n\times n}[c^2, 2bc, b^2]$ and N_{1L} given in Kronecker product form by (7). If $\vec{x}$ and $\vec{y}$ are the dominant eigenvectors of $\mathrm{tri}_{n\times n}[1,0,1]$ and $\mathrm{tri}_{n\times n}[c^2, 2bc, b^2]$, by Proposition 4, $\vec{x}\otimes\vec{y}$ is both the dominant eigenvector of N_{1L} and the eigenvector corresponding to the minimal eigenvalue of M_{1L}. Therefore the bound is attained. □

Table 2 shows experimentally the tightness of our bound. In the case $k = 2$ the bound becomes arbitrarily tight as $h \to 0$. For $k > 2$, however, the bound is not tight.

n	ρ_{2L}	$\tilde{\rho}_{2L}$	ρ_{3L}	$\tilde{\rho}_{3L}$
6	0.339	0.383	0.302	0.455
12	0.400	0.415	0.345	0.501
18	0.415	0.422	0.356	0.511
24	0.421	0.425	0.360	0.515
30	0.423	0.426	0.362	0.517

Table 2: Spectral radii ρ_{2L} and ρ_{3L} of the 2-line and 3-line iteration matrices, as well as the bounds $\tilde{\rho}_{2L}$ and $\tilde{\rho}_{3L}$. For the 2-line ordering the bound grows tight for large n, but not for the 3-line ordering. $\gamma = 0.5$ and $\delta = 0$.

We expect the convergence of block Jacobi to improve as k gets larger, since a greater proportion of the matrix A_{kL} ends up in M_{kL}. Since our bound $\rho_{\tilde{kL}}$ is not tight for $k > 2$, we cannot predict analytically how big the improvement is. We have nevertheless observed this experimentally.

The following proposition shows that doubling k always improves the convergence of block Jacobi.

Proposition 7 *For any given $k \geq 1$, $\rho_{2kL} < \rho_{kL}$.*

Proof. It is straightforward to show that $N_{kL} \geq N_{2kL} \geq 0$, when N_{kL} and N_{2kL} are permuted by reordering grid unknowns lexicographically. Since

the matrix A_{kL} is an irreducibly diagonally dominant M-matrix, $A_{kL}^{-1} > 0$ [10, pg 91]. The claimed result follows from [10, pg 97, Theorem 3.32]. □

We can use the results of Theorem 1 to determine exactly (approximately) the optimal parameter for block SOR with the 1-line (2-line) ordering. Since A_{kL} is block tridiagonal, it has block property A, and hence the analysis of Young [11, Chap. 14, Sections 5.2 and 14.3] applies.

Comparison of Convergence Rates (2D)

Let us compare the convergence of block Jacobi applied to our reduced system, the traditional reduced system, and the unreduced system. When referring to quantities of interest we use the superscripts $\square$, $\diamond$, and $+$ respectively, motivated by the shapes of the corresponding computational molecules.

A bound on the spectral radius of the standard 2D cyclically reduced system with a diagonal ordering strategy is given by [1]

$$\rho_{1L}^{\diamond} \leq \tilde{\rho}_{1L}^{\diamond} = \frac{(\sqrt{1-\gamma^2}+\sqrt{1-\delta^2})^2}{8-(\sqrt{1-\gamma^2}+\sqrt{1-\delta^2})^2+2\sqrt{(1-\gamma^2)(1-\delta^2)}(1-\cos(\pi h))}, \tag{8}$$

valid for $|\gamma|, |\delta| < 1$. The diagonal ordering is equivalent to a 1-line ordering if the grid is viewed at a 45 degree angle. It was observed experimentally that $\tilde{\rho}_{1L}^{\diamond} \to \rho_{1L}^{\diamond}$ as $h \to 0$ with (γ, δ) fixed.

Setting $\delta = 0$, we expand (8) as well as the results of Theorem 1 in a Taylor series centered at $h = 0$.

$$\tilde{\rho}_{1L}^{\diamond} = 1 - \left(-\frac{\pi^2}{4} + \frac{\sigma^2}{4}\right) h^2 + o(h^2);$$

$$\rho_{1L}^{\square} = 1 - \left(4\pi^2 + \frac{\sigma^2}{2}\right) h^2 + o(h^2);$$

$$\rho_{2L}^{\square} \leq \tilde{\rho}_{2L}^{\square} = 1 - (4\pi^2 + \sigma^2) h^2 + o(h^2).$$

As observed in Table 2, $\tilde{\rho}_{2L}^{\square} \to \rho_{2L}^{\square}$ as $h \to 0$ with γ fixed.

For block Jacobi applied to the standard 5-point centered differences operator, we have

$$\rho_{1L}^{+} = 1 - \left(\pi^2 + \frac{\sigma^2}{8}\right) h^2 + o(h^2)$$

and

$$\rho_{2L}^{+} \leq \tilde{\rho}_{2L}^{+} = 1 - \left(\pi^2 + \frac{\sigma^2}{4}\right) h^2 + o(h^2),$$

with the latter bound observed experimentally to grow arbitrarily tight as $h \to 0$ with γ fixed [5].

The asymptotic convergence rate is $R_\infty = -\log(\rho)$. Noting that $\log(1 + ch^2) = ch^2 + o(h^2)$, we have for $\sigma^2 \gg \pi^2$ and $h \ll 1$ the following hierarchy of relative asymptotic convergence rates:

$$(R_\infty)_{2L}^{\square} \approx 2(R_\infty)_{1L}^{\square} \approx 4(R_\infty)_{1L}^{\diamond} \approx 4(R_\infty)_{2L}^{+} \approx 8(R_\infty)_{1L}^{+}.$$

4.2 3D Case: k-Plane Ordering

Consider the block Jacobi splitting $A_{kP} = M_{kP} - N_{kP}$ with M_{kP} consisting of the blocks $(A_{kP})_{ii}$ of the k-plane matrix discussed in section 3. As previously stated, we assume that $\vec{w}$ is aligned with the x-axis. Thus, we have $\delta = \eta = 0$, which means $b = e = p = s$ and $c = d = q = r$. As we did for the 2D case, it is useful to consider the general setting of a uniform $n_x \times n_y \times n_z$ grid to establish the eigenvalues of A_{kP}; we will then resort back to $n_x = n_y = n_z$.

A_{1P} is a block tridiagonal matrix of size $n_y \times n_y$. Each block is itself a block tridiagonal matrix of size $n_x \times n_x$ with respect to $n_z \times n_z$ tridiagonal blocks.

Assuming that the ordering the gridpoints goes in a z-x-y fashion without loss of generality, we have

$$A_{1P} = \mathrm{tri}_{n_y \times n_y}[B, C, B],$$

where

$$\begin{aligned} B &= \mathrm{tri}_{n_x \times n_x}[\mathrm{tri}_{n_z \times n_z}[-c^2, -2c^2, -c^2], \\ &\qquad \mathrm{tri}_{n_z \times n_z}[-2bc, -4bc, -2bc], \\ &\qquad \mathrm{tri}_{n_z \times n_z}[-b^2, -2b^2, -b^2]] \ ; \\ \\ C &= \mathrm{tri}_{n_x \times n_x}[\mathrm{tri}_{n_z \times n_z}[-2c^2, -4c^2, -2c^2], \\ &\qquad \mathrm{tri}_{n_z \times n_z}[-4bc, a^2 - 8bc, -4bc], \\ &\qquad \mathrm{tri}_{n_z \times n_z}[-2b^2, -4b^2, -2b^2]] \ . \end{aligned}$$

Using Kronecker products, we have

$$A_{1P} = a^2 I_{n_x n_y n_z} - \mathrm{tri}_{n_y \times n_y}[1, 2, 1] \otimes \mathrm{tri}_{n_x \times n_x}[c^2, 2bc, b^2] \otimes \mathrm{tri}_{n_z \times n_z}[1, 2, 1]. \quad (9)$$

Proposition 8 *The eigenvalues of the matrix A_{kP} are*

$$\lambda_{i,j,k} = a^2 - 8bc\left[1 + \cos\left(\frac{i\pi}{n_x + 1}\right)\right]\left[1 + \cos\left(\frac{j\pi}{n_y + 1}\right)\right]\left[1 + \cos\left(\frac{\ell\pi}{n_z + 1}\right)\right],$$
$$i = 1, 2, ..., n_x, \qquad j = 1, 2, ..., n_y, \qquad \ell = 1, 2, ..., n_z.$$

Proof. This follows from (9) and Propositions 3 and 4. □

Once again, the eigenvalues of M_{kP} follow as a simple corollary. From this point on we assume $n_x = n_y = n_z = n$.

Corollary 3 *The eigenvalues of the matrix M_{kP} are given by*

$$\lambda_{i,j,k}(M_{kP}) = a^2 - 8bc\left[1 + \cos\left(\frac{i\pi}{k + 1}\right)\right]\left[1 + \cos\left(\frac{j\pi}{k + 1}\right)\right](1 + \cos(2\pi\ell h)),$$
$$i = 1, 2, ..., k, \qquad j = 1, 2, ..., k, \qquad \ell = 1, 2, ..., n,$$

each with multiplicity $\frac{n^2}{k^2}$.

The minimal eigenvalue of M_{kP} is

$$\lambda_{min}(M_{kP}) = a^2 - 8bc\left[1 + \cos\left(\frac{\pi}{k + 1}\right)\right]^2 (1 + \cos(2\pi h)).$$

Finally, M_{kP} is symmetrizable, and the symmetrized matrix is positive definite.

Proof. M_{kP} is block diagonal, with each block equal to the matrix A_{kP} discretized on a $k \times k \times n$ grid. The result follows from applying Proposition 8 to a $k \times k \times n$ subgrid. Since M_{kP} is a Kronecker product of tridiagonal Toeplitz matrices, it is symmetrizable by a diagonal similarly transformation. By inspecting the values of a, b and c, it readily follows that $\lambda_{\min}$ is positive. □

Lemma 1 *The eigenvalues of N_{1P} are given by*

$$\lambda_{ij\ell}(N_{1P}) = 8bc\left[(1+\cos(2\pi ih))\,(1+\cos(2\pi jh)) - 1\right][1+\cos(2\pi\ell h)] \qquad i,j,\ell = 1,2,...,n.$$

The spectral radius of N_{1P} is thus given by

$$\rho(N_{1P}) = 8bc\cos(2\pi h)\,[1+\cos(2\pi h)]\,[2+\cos(2\pi h)]\,.$$

When $k > 1$, a bound on $\rho(N_{kP})$ is given by

$$\rho(N_{kP}) \le 2bc\left[1+4\left(1+\cos\left(\frac{\pi}{k+1}\right)\right)\right][1+\cos(2\pi h)]\,.$$

Proof. We have

$$N_{1P} = (\mathrm{tri}_{n\times n}[1,2,1] \otimes \mathrm{tri}_{n\times n}[c^2, 2bc, b^2] - 4bcI_{n^2}) \otimes \mathrm{tri}_{n\times n}[1,2,1].$$

The claim on its eigenvalues now follows from Propositions 3 and 4. The spectral radius is attained at $i = j = \ell = 1$.

When $k > 1$ we use the following splitting:

$$N_{kP} = (N_{kP})_1 + (N_{kP})_2 + (N_{kP})_3,$$

where the eigenvalues of the three matrices on the right can be computed exactly. The sparsity patterns of these matrices are depicted in Figure 4.2.

It is useful to introduce some notation regarding the faces and edges of each grid block B_i. We denote by $\partial_{x^+}B_i$ and $\partial_{x^-}B_i$ the faces on the positive x and negative x sides of B_i. The faces in the y direction are denoted in the same fashion. The edge where the positive x and y faces meet is denoted by $\partial_{x^+y^+}B_i$, with obvious modifications for the other edges.

There exist three reorderings of G_{3D}, each bringing one of the pieces of N_{kP} to block diagonal form.

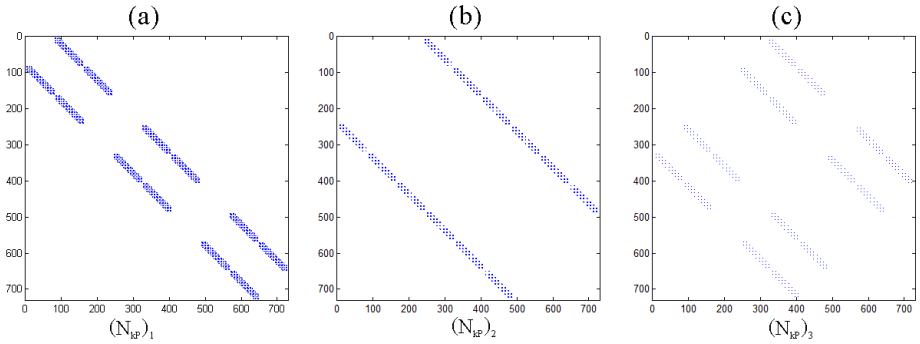

Figure 6: Partition of N_{kP} into (a) $(N_{kP})_1$, (b) $(N_{kP})_2$, and (c) $(N_{kP})_3$. The grid is $9 \times 9 \times 9$ with $k = 3$.

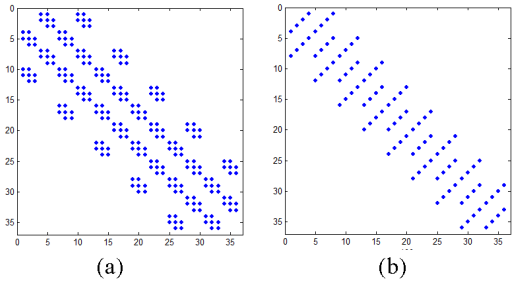

Figure 7: (a) Nonzero blocks of permuted $(N_{kP})_1$ and $(N_{kP})_2$. (b) Nonzero blocks of permuted $(N_{kP})_3$. Grid is $9 \times 9 \times 9$, $k = 3$.

We redefine the grid blocks as

$$\tilde{B}_i = \partial_{x^+} B_i \bigcup \partial_{x^-} B_{i+1} \qquad i = 1, 2, ..., \frac{n}{k}\left(\frac{n}{k} - 1\right)$$

and

$$\tilde{B}_{\frac{n}{k}(\frac{n}{k}-1)+1} = G_{3D} - \bigcup_{i=1}^{\frac{n}{k}(\frac{n}{k}-1)} \tilde{B}_i.$$

The permuted matrix $(\tilde{N}_{kP})_1$ is block diagonal with $\frac{n}{k}(\frac{n}{k} - 1)$ blocks of size $2kn \times 2kn$, and one block of zeroes of size $(n^3 - 2n^2(\frac{n}{k} - 1)) \times (n^3 - 2n^2(\frac{n}{k} - 1))$. Each nonzero block is the matrix

$$\text{tri}_{n\times n}[1, 2, 1] \otimes \text{tri}_{2\times 2}[-c^2, 0, -b^2] \otimes \text{tri}_{k\times k}[1, 2, 1],$$

whose sparcity pattern is shown in 7(a). By Propositions 3 and 4, the eigenvalues of $(N_{kP})_1$ are

$$\pm 4bc\left(1 + \cos\left(\frac{i\pi}{k+1}\right)\right)(1 + \cos(2\pi ih)), \qquad i = 1, 2, ..., k, \qquad j = 1, 2, ..., n,$$

each of multiplicity $\frac{n}{k}(\frac{n}{k}-1)$ and 0 of multiplicity $(n^3-2n^2(\frac{n}{k}-1))$. Clearly,

$$\rho((N_{kP})_1) = 4bc\left(1+\cos\left(\frac{\pi}{k+1}\right)\right)(1+\cos(2\pi h)). \tag{10}$$

By an almost identical process (using $\tilde{B}_i = \partial_{y^+}B_i \bigcup \partial_{y^-}B_{i+1}$), we have

$$\rho((N_{kP})_2) = \rho((N_{kP})_1). \tag{11}$$

We now define grid blocks

$$\grave{B}_i = \partial_{x^+y^+}B_i \bigcup \partial_{x^-y^+}B_{i+1} \bigcup \partial_{x^+y^-}B_{i+\frac{n}{k}} \bigcup \partial_{x^-y^-}B_{i+\frac{n}{k}+1} i = 1,2,...,\left(\frac{n}{k}-1\right)^2,$$

and

$$\grave{B}_{\left(\frac{n}{k}-1\right)^2+1} = G_{3D} - \bigcup_{i=1}^{\left(\frac{n}{k}-1\right)^2} \grave{B}_i.$$

The block diagonal matrix $(\grave{N}_{kP})_3$ has $(\frac{n}{k}-1)^2$ diagonal blocks of size $4n\times 4n$, followed by a zero block of size $(n^3-4n(\frac{n}{k}-1)^2)\times(n^3-4n(\frac{n}{k}-1)^2)$. Each nonzero block is the matrix

$$\mathrm{tri}_{n\times n}[1,2,1]\otimes \mathrm{tri}_{2\times 2}[1,0,1]\otimes \mathrm{tri}_{2\times 2}[-c^2,0,-b^2],$$

whose sparcity pattern is shown in 7(b). By Propositions 3 and 4, the eigenvalues of $(N_{kP})_3$ are

$$\pm 2bc(1+\cos(2\pi ih)), \qquad i=1,2,...,n,$$

each of multiplicity $(\frac{n}{k}-1)^2$ and 0 of multiplicity $n^3-4n(\frac{n}{k}-1)^2)$. Clearly,

$$\rho((N_{kP})_3) = 2bc(1+\cos(2\pi h)). \tag{12}$$

The bound on $\rho(N_{kP})$ now follows from (10), (11), (12), and the observation that $\rho(N_{kP}) \le \rho((N_{kP})_1) + \rho((N_{kP})_2) + \rho((N_{kP})_3)$. □

Theorem 2 *The spectral radius of the block Jacobi iteration matrix obeys the bound*

$$\begin{aligned}\rho_{kP} \le \tilde{\rho}_{kP} \quad = \quad & \frac{2bc\left[1+4\left(1+\cos\left(\frac{\pi}{k+1}\right)\right)\right](1+\cos(2\pi h))}{a^2-8bc\left[1+\cos\left(\frac{\pi}{k+1}\right)\right]^2(1+\cos(2\pi h))} \\ & \cdot\left[1+\delta_{1k}\left(\frac{4}{5}\cos(2\pi h)(2+\cos(2\pi h))-1\right)\right].\end{aligned}$$

Furthermore, the bound is exact for $k=1$, that is $\tilde{\rho}_{1P} = \rho_{1P}$.

Proof. By Corollary 3, M_{kP} is symmetrizeable and positive definite. It is straightforward to show N_{kP} is symmetrizeable too. Therefore (6) may be used, and the bound follows from Lemma 1 and Corollary 3.

For $k = 1$ we have

$$M_{1P} = a^2 I_{n^3} - 4bcI_{n^2} \otimes \mathrm{tri}_{n\times n}[1, 2, 1]$$

and

$$N_{1P} = (\mathrm{tri}_{n\times n}[1, 2, 1] \otimes \mathrm{tri}_{n\times n}[c^2, 2bc, b^2] - 4bcI_{n^2}) \otimes \mathrm{tri}_{n\times n}[1, 2, 1].$$

If $\vec{x}$ and $\vec{y}$ denote the dominant eigenvectors of $\mathrm{tri}_{n\times n}[1, 2, 1]$ and $\mathrm{tri}_{n\times n}[b^2, 2ab, a^2]$, then from Proposition 4 the vector $\vec{x} \otimes \vec{y} \otimes \vec{x}$ is both the dominant eigenvector of N_{1P} and the eigenvector corresponding to the minimal eigenvalue of M_{1P}. The bound is therefore attained in this case. □

Table 3 shows experimentally the tightness of our bound. In the case $k = 2$, we see that the bound becomes arbitrarily tight as $h \to 0$. For $k = 3$ (and in fact for $k > 2$ in general) the bound is not tight, but note that the computed spectral radius for $k = 3$ is smaller compared to the one for $k = 2$.

n	ρ_{2P}	$\tilde{\rho}_{2P}$	ρ_{3P}	$\tilde{\rho}_{3P}$
6	0.430	0.521	0.372	0.726
12	0.524	0.554	0.454	0.784
18	0.547	0.561	0.475	0.797
24	0.556	0.564	0.483	0.802
30	0.556	0.565	0.487	0.805
36	0.562	0.566	0.489	0.806
42	0.564	0.566	0.490	0.807

Table 3: Spectral radii ρ_{2P} and ρ_{3P} of the 2-plane and 3-plane iteration matrices, as well as the bounds $\tilde{\rho}_{2P}$ and $\tilde{\rho}_{3P}$. For the 2-plane ordering the bound grows tight for large n, but not for the 3-line ordering. We have taken $\gamma = 0.5$ and $\delta = 0 = \eta = 0$.

Comparison of Convergence Rates (3D)

A bound on the spectral radius of the block Jacobi iteration matrix obtained by applying a 2-plane ordering to the traditional 19-point 3D cyclically reduced operator is given by [8]

$$\rho_{2P}^{19pt} \leq \tilde{\rho}_{2P}^{19pt} = 1 - \left(\frac{10}{9}\pi^2 + \frac{1}{6}\mu^2 + \frac{1}{6}\tau^2 + \frac{1}{6}\sigma^2\right) h^2 + o(h^2).$$

Expanding the results of theorem 2 in a Taylor series gives

$$\rho_{1P}^{27pt} = 1 - (4\pi^2 + \frac{\sigma^2}{3})h^2 + o(h^2)$$

and

$$\rho_{2P}^{27pt} \leq \tilde{\rho}_{2P}^{27pt} = 1 - \frac{4}{7}(4\pi^2 + \sigma^2)h^2 + o(h^2).$$

For the standard 7-point centered differences operator, we have [5]

$$\rho_{1P}^{7pt} = 1 - (\frac{3}{4}\pi^2 + \frac{1}{16}\sigma^2 + \frac{1}{16}\tau^2 + \frac{1}{16}\mu^2)h^2 + o(h^2)$$

and

$$\rho_{2P}^{7pt} \leq \tilde{\rho}_{2P}^{7pt} = 1 - (\frac{1}{2}\pi^2 + \frac{1}{8}\sigma^2 + \frac{7}{64}\tau^2 + \frac{7}{64}\mu^2)h^2 + o(h^2).$$

We have found experimentally that all the above bounds grow arbitrarily tight as $h \to 0$ with mesh Reynolds numbers held constant. Noting $-\log(1 + ch^2) = ch^2 + o(h^2)$, we construct Table 4 showing the expected relative asymptotic convergence rates, valid for $h \ll 1$ and $\sigma^2 \gg \pi^2$.

	7-pt $1P$	7-pt $2P$	19-pt $2P$	27-pt $1P$	27-pt $2P$
7-pt $1P$	1				
7-pt $2P$	2	1			
19-pt $2P$	2.67	1.34	1		
27-pt $1P$	5.35	2.67	2	1	
27-pt $2P$	9.14	4.57	3.42	1.71	1

Table 4: Relative asymptotic convergence rates for different combinations of discretization scheme and grid ordering. The number in row i and column j is the ratio of the asymptotic convergence rate of method i to method j, using Taylor expansions.

Computational Cost. For the operators we consider, it is straightforward to estimate the computational cost per iteration. Due to the process of reduction, our cyclically reduced operators entail a lower computational cost of a single iteration, compared to the other operators. For example, in the 2D case the new operator is 9-point, as is the traditional cyclically reduced operator, but we work on a grid with only 1/4 of the unknowns rather than 1/2. In the 3D case our operator is 27-point (as opposed to 19-point for the traditional cyclically reduced operator) but we solve for only 1/8 of the unknowns rather than 1/2. Solving for the unknowns that were eliminated in the process entails a negligible cost, as it involves a very small number of diagonal system solves.

5 Numerical Experiments

In this section we present an experimental examination of our approach. We show results for block Jacobi, for which we have carried out a detailed analysis, and then briefly explore the performance of modern Krylov solvers. Specifically, we apply GMRES, preconditioned with incomplete LU. All experiments were

done using MATLAB. In both 2D and 3D we have looked at three linear systems arising from discretizing (1), corresponding to the standard second order centered difference discretization, the reduced system arising from standard cyclic reduction, and the reduced system arising from our new approach.

5.1 Test Problem 1: Constant Coefficients

Table 5 shows the number of iterations required for the convergence of block Jacobi. Then, we show the results of applying GMRES, preconditioned with ILUTP to each system in turn. A variety of values of σ, τ and μ were tried, for several different grid sizes. The results of a few of these experiments are shown in Table 6.

In the tables we use $+$, $\diamond$ and $\square$ respectively to denote the 5-point operator, standard cyclically reduced operator, and the new operator, as per the shapes of the corresponding computational molecules.

	n=128	n=256	n=512
$+$, 1L	877	3294	12254
$+$, 2L	443	1652	6131
$\diamond$, 1L	450	1668	6187
$\diamond$, 2L	230	839	3098
$\square$, 1L	241	887	3301
$\square$, 2L	125	449	1655

(a) 2D

	$n = 17$	$n = 33$	$n = 65$
7-pt, 1P	75	287	1098
7-pt, 2P	42	148	554
19-pt, 1P	46	170	643
19-pt, 2P	33	112	416
27-pt, 1P	15	56	208
27-pt, 2P	11	35	125

(b) 3D

Table 5: Number of block Jacobi iterations it takes to reduce the 2-norm of the relative residual by a factor of 10^{-4} for constant coefficients. The initial guess is the zero vector. The right hand side was constructed so that the solution was a vector of all 1s. Results are shown for $\sigma = 60$ in 2D and $\sigma = 30$ in 3D, and $\tau = \mu = 0$.

The results validate our analysis. The reduced system is solved within approximately half the number of iterations of the standard cyclically reduced

1-line	n=129	n=257	n=513
+	20	46	97
◇	11	26	55
□	7	15	34
2-line	n=129	n=257	n=513
+	18	36	76
◇	12	27	63
□	6	14	31

(a) 2D

1-plane	n=17	n=33	n=65
7-point	5	9	21
19-point	4	8	16
27-point	2	4	8
2-plane	n=17	n=33	n=65
7-point	5	10	21
19-point	4	7	15
27-point	2	4	8

(b) 3D

Table 6: Constant coefficient problem: number of iterations for GMRES preconditioned with ILU(0.01) to bring the norm of the relative residual down to 10^{-4}. $\sigma = 60$, $\tau = \mu = 0$, with $h = \frac{1}{n+1}$.

1-line	n=129	n=257	n=513
+	40	65	115
◇	18	30	57
□	14	23	42
2-line	n=129	n=257	n=513
+	28	45	85
◇	20	37	63
□	15	25	43

(a) 2D

1-plane	n=17	n=33	n=65
7-point	7	12	23
19-point	5	9	14
27-point	3	6	11
2-plane	n=17	n=33	n=65
7-point	7	13	24
19-point	5	8	14
27-point	3	6	11

(b) 3D

Table 7: Variable coefficient problem: number of iterations for GMRES preconditioned with ILU(0.01) to bring the norm of the relative residual down by a factor of 10^{-4}.

system. For block stationary schemes, the 2-line/plane orderings are provably superior in terms of speed of convergence, and this is manifested also in the experiments. For GMRES the benefit of using these orderings in place of standard lexicographic ordering is much less obvious. The two cyclically reduced operators typically perform similarly to lexicographic ordering, and in fact in some cases require more iterations for convergence.

5.2 Test Problem 2: Variable Coefficients

We now modify problem (1) so that the vector $\vec{w}$ is a function of the spatial coordinates. Table 7 shows the results for a 2D circular flow problem, $\vec{w} = 20\left(\frac{1}{2} - y, \frac{1}{2} + x\right)$, and for a 3D sink problem, $\vec{w} = 20\left(\frac{1}{2} - x, \frac{1}{2} - y, \frac{1}{2} - z\right)$.

The results lead to very similar conclusions as the constant coefficient case.

6 Conclusions

We have presented a new cyclically reduced operator and have provided an analysis of its spectral properties, for the discrete convection-diffusion equation. The derivation has been motivated by the following key point: instead of using a standard unreduced operator that leads to a non-standard cyclically reduced operator, as has been previously done in the literature, we start off with a non-standard operator, and after reduction the resulting operator has a simple computational molecule. As such, it is much easier to implement, while it maintains the same good spectral properties.

References

[1] Howard C. Elman and Gene H. Golub. Iterative methods for cyclically reduced nonselfadjoint linear systems. *Math. Comp.*, 54(190):671–700, 1990.

[2] Howard C. Elman and Gene H. Golub. Iterative methods for cyclically reduced nonselfadjoint linear systems. II. *Math. Comp.*, 56(193):215–242, 1991.

[3] Howard C. Elman and Gene H. Golub. Line iterative methods for cyclically reduced discrete convection-diffusion problems. *SIAM J. Sci. Statist. Comput.*, 13(1):339–363, 1992.

[4] Walter Gander and Gene H. Golub. Cyclic reduction—history and applications. In *Scientific computing (Hong Kong, 1997)*, pages 73–85. Springer, Singapore, 1997.

[5] Gene H. Golub, Chen Greif, and James M. Varah. Block orderings for tensor-product grids in two and three dimensions. *Numer. Algorithms*, 30(2):93–111, 2002.

[6] Chen Greif. Reduced systems for three-dimensional elliptic equations with variable coefficients. *SIAM J. Matrix Anal. Appl.*, 21(1):29–44, 1999.

[7] Chen Greif and James Varah. Iterative solution of cyclically reduced systems arising from discretization of the three-dimensional convection-diffusion equation. *SIAM J. Sci. Comput.*, 19(6):1918–1940, 1998.

[8] Chen Greif and James Varah. Block stationary methods for nonsymmetric cyclically reduced systems arising from three-dimensional elliptic equations. *SIAM J. Matrix Anal. Appl.*, 20(4):1038–1059, 1999.

[9] James O'Neil and Daniel B. Szyld. A block ordering method for sparse matrices. *SIAM J. Sci. Statist. Comput.*, 11(5):811–823, 1990.

[10] Richard S. Varga. *Matrix iterative analysis.* Prentice-Hall Inc., Englewood Cliffs, N.J., 1962.

[11] David M. Young. *Iterative solution of large linear systems.* Academic Press, New York, 1971.