Title: Results on the Effect of Orderings on SSOR and ILU
Preconditioning

RECEIVED

MAY 0 3 1999

U S I

Author(s): Wayne D. Jourbert
Emanuel H. Knill

Submitted to: External Distribution as Requested

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

MASTER

Los Alamos
NATIONAL LABORATORY

DISCLAIMER

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

RESULTS ON THE EFFECT OF ORDERINGS ON SSOR AND PRECONDITIONINGS *

RECEIVED

MAY 0 3 1999

OSTI

W. JOUBERT [†] AND E. KNILL [†]

Abstract. It is known that for SSOR and ILU preconditionings for solving systems of linear equations, orderings can have an enormous impact on robustness, convergence rate and parallelism. Unfortunately, it has been observed that there is an inverse relation between the convergence rate and the parallelism of typical orderings used in practice. This paper presents some numerical experiments with simple matrices to illustrate this behavior as well as a new theoretical result which sheds some light on this phenomenon and also gives an upper bound on the convergence rate of a number of preconditioners in popular use.

Key words. linear systems, iterative methods, preconditioning, incomplete factorizations, incomplete Cholesky preconditioning, SSOR preconditioning, parallel computation

AMS subject classifications. 65F10, 65F15

1. Introduction. The solution of sparse linear systems of the form

$$Au = b$$

is vital to many computational modeling and simulation processes of interest. Furthermore, it is recognized that parallel computers are necessary to solve very large problems of importance, necessitating the use of effective parallel linear solver algorithms.

Well-known preconditioners such as incomplete Cholesky and ILU preconditioning which lie at the heart of many linear solvers have been notoriously difficult to parallelize. ILU preconditioning with standard natural or reverse Cuthill-McKee orderings are robust and lead to rapid convergence of the iterative method in many cases, but they are difficult to parallelize. On the other hand, orderings such as red-black or multicolor orderings, though more parallelizable, in many cases lead to slower convergence. This inverse relation between parallelism and convergence rate has been observed, for example, by [Duff/Meurant].

The purpose of this paper is to examine the validity of this claim in more detail by the use of numerical experiments. We will also present a result for an important prototypical

model problem showing that a number of preconditionings, including multicolor IC and SSOR, which have high parallelism, can only improve upon the unpreconditioned case by a constant factor for very large problem sizes. Such convergence behavior is qualitatively worse than, say, MILU preconditioning with natural ordering, for which the improvement over the unpreconditioned case grows as the problem size grows.

The arrangement of this paper is as follows. In Section 2 we give numerical experiments to examine the relationship between convergence and parallelism of preconditioned iterative methods with orderings. Then in Section 3 we give the new result on the asymptotic performance of various preconditioners.

2. Numerical experiments with small matrices. In principle, it would be desirable to extend the study of [Duff/Meurant] to test all possible orderings of unknowns for simple model problems. Unfortunately, the number of orderings of unknowns grows exponentially with the number of unknowns for certain simple model problems of interest such as those of [Duff/Meurant]; thus, such a calculation is not computationally feasible. Nonetheless, looking at the performance of preconditionings under various orderings for very small cases, for which it is computationally feasible to examine all orderings, may give some clues as to what factors are key in controlling the performance of these methods.

For these experiments we will consider matrices derived from the 2-D Laplace equation,

$$-u_{xx} - u_{yy} = 0$$

on a rectangular domain, with homogeneous Dirichlet boundary conditions. The problem is discretized with standard central finite differences with $(n_x + 1) \times (n_y + 1)$ grid points, leading after elimination of the boundary unknowns to a matrix with five nonzero diagonals (before reordering), with 4's on the main diagonal and -1's as the other nonzero entries of the matrix.

For simplicity of analysis, we use the SSOR preconditioner,

$$Q = [\omega/(2 - \omega)][(1/\omega)D + L]D^{-1}[(1/\omega)D + U],$$

with relaxation parameter $\omega = 1$ (symmetric Gauss-Seidel). Here, $A = D + L + U$ is the decomposition of A into main diagonal, strictly lower triangular and strictly upper triangular matrices, and the preconditioned system to be solved is

$$Q^{-1}Au = Q^{-1}b.$$

The SSOR and SGS preconditioners have an identical structure to the no-fill incomplete Cholesky ordering for this case, and the SSOR and M/IC methods involve similar convergence and parallelism issues.

These experiments are performed on a Sun workstation with 64-bit IEEE floating point arithmetic. For each preconditioner, the preconditioned matrix is formed as a dense matrix, and the condition number is calculated from the extremal·eigenvalues as computed by LAPACK routines. For a given choice of $n_x \times n_y$, we compute all possible orderings of the unknowns, and for each ordering, apply the ordering to A, precondition, and then compute the resulting condition number.

The naive approach to enumerating all possible cases is to generate all $[(n_x n_y)!]$ permutations of the numbered grid points. However, many redundancies exist in such an approach, in the sense that, as is known [Young], two different orderings can lead to an identical preconditioner. Furthermore, other symmetries exist due to the nature of the model problem (regular grid, constant coefficients).

To enumerate the possible orderings without redundancies of preconditioner, we use the following approach. For any pair of grid points in the grid which are horizontally or vertically adjacent, one might draw a small arrow from one point to the other to indicate the relative priority of which point is ordered or updated before the other. This approach gives rise to $2^{n_x(n_y-1)+n_y(n_x-1)}$ different cases, since each arrow has a choice of two directions. This enumeration has exponential order, considerably better than the factorial order of the naive approach.

From these cases, it is necessary to eliminate each case that has a "cycle," i.e., a path through the grid that starts and ends at the same point. After this is done, it is easy to see that an ordering can be associated with each case: for the given case, recursively a point is found with all out-arrows and that point is numbered and eliminated. Furthermore, this scheme covers all possible cases according to [Young]. Note that this scheme confirms that the number of orderings producing distinct preconditioners is, at most, exponential in the number of grid points.

Since the problem is regular and has constant coefficients, redundancies due to reflections in x and y may be removed. If $n_x = n_y$, then redundancies from exchanging x and y may also be removed.

The resulting set of orderings still has some redundancies which may be removed, associ-

ated with the fact that if the ordering is reversed, the condition number of the preconditioned matrix is unchanged.

These calculations are summarized in Table 1 for a set of small grids. The final column of the table is generated empirically by counting the number of unique condition numbers which are generated by the code.

n_x	n_y	$n_x \cdot n_y$ # points	$(n_x n_y)!$ # perms	$2^{n_x(n_y-1)+n_y(n_x-1)}$ # update pat	# update pat no cycles	# update pat no cycles/refl	# cond
1	1	1	1	1	1	1	1
1	2	2	2	2	2	1	1
1	3	3	6	4	4	3	2
1	4	4	24	8	8	4	3
1	5	5	120	16	16	10	6
1	6	6	720	32	32	16	10
1	7	7	5040	64	64	36	19
1	8	8	40320	128	128	64	36
1	9	9	362880	256	256	136	69
1	10	10	3628800	512	512	256	135
1	11	11	39916800	1024	1024	528	261
1	12	12	479001600	2048	2048	1024	527
1	13	13	6.227020e09	4096	4096	2080	1030
1	14	14	8.717829e10	8192	8192	4096	2053
1	15	15	1.307674e12	16384	16384	8256	3992
1	16	16	2.092279e13	32768	32768	16384	7706
1	17	17	3.556874e14	65536	65536	32896	—
1	18	18	6.402373e15	131072	131072	65536	—
1	19	19	1.216451e17	262144	262144	131328	—
1	20	20	2.432902e18	524288	524288	262144	—
2	1	2	2	2	2	1	1
2	2	4	24	16	14	3	3
2	3	6	720	128	98	28	17
2	4	8	40320	1024	686	175	101
2	5	10	3628800	8192	4802	1225	639
2	6	12	479001600	65536	33614	8428	4308
2	7	14	8.717829e10	524288	235298	58996	—
2	8	16	2.092279e13	4194304	1647086	411943	—
3	1	3	6	4	4	3	2
3	2	6	720	128	98	28	17
3	3	9	362880	4096	2398	345	168
3	4	12	479001600	131072	58670	14839	7409
3	5	15	1.307674e12	4194304	1435414	360933	—
4	1	4	24	8	8	4	3
4	2	8	40320	1024	686	175	101
4	3	12	479001600	131072	58670	14839	7409
4	4	16	2.092279e12	16777216	5015972	627829	—

Table 1. Number of distinct orderings for small problems.

Figure 1. Path length vs. condition, grids 1x2, 1x3, 1x4.

Figure 2. Path length vs. condition, 1x5 grid.

Figure 3. Path length vs. condition, 1x6 grid.

Figure 4. Path length vs. condition, 1x7 grid.

Figure 5. Path length vs. condition, 1x8 grid.

Figure 6. Path length vs. condition, 1x9 grid.

Figure 7. Path length vs. condition, 1x10 grid.

Figure 8. Path length vs. condition, 1x11 grid.

Figure 9. Path length vs. condition, 1x12 grid.

Figure 10. Path length vs. condition, 1x13 grid.



Figure 11. Path length vs. condition, 1x14 grid.

Figure 12. Path length vs. condition, 1x15 grid.

Figure 13. Path length vs. condition, 1x16 grid.

Figure 14. Path length vs. condition, grids 2x1, 2x2, 2x3.
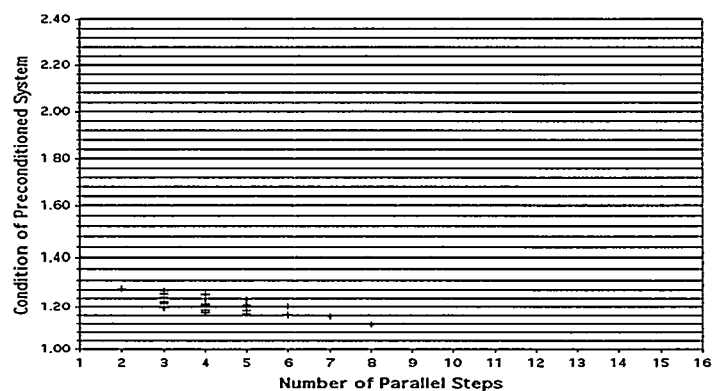
Figure 15. Path length vs. condition, 2x4 grid.

Figure 16. Path length vs. condition, 2x5 grid.

Figure 17. Path length vs. condition, 2x6 grid.

Figure 18. Path length vs. condition, grids 3x1, 3x2, 3x3.

Figure 19. Path length vs. condition, 3x4 grid.

Figure 20. Path length vs. condition, grids 4x1, 4x2, 4x3.

Figure 21. Correlation of condition with squared Frobenius norm, 3x3 grid.

Figures 1 through 20 give results of these numerical experiments for a variety of small

matrices. For each case, the number of parallel steps ("wavefronts," "independent sets" or "colors") induced by the ordering is given on the horizontal axis, and the condition number of the preconditioned system is given on the vertical axis.

We may draw several conclusions from these experiments:

1. As would be expected, larger problem sizes lead to generally larger condition numbers as well as more orderings and potentially more parallel steps.

2. In every case, the worst performer in terms of condition number is the red-black ordering (and equivalent orderings), which has higher condition number than any other ordering.

3. In every case, the best performer is the natural lexicographical ordering (which in this case is equivalent to reverse Cuthill McKee). In fact, this ordering is even slightly better than orderings with more parallel steps, such as the spiral or snake orderings.

4. The general trend of [Duff/Meurant] is confirmed, that, generally, the more parallel orderings give larger condition numbers than the less parallel orderings, though there is much variability around the general trendline.

5. It is possible to have slightly more than two parallel steps and have much better performance than the red-black ordering. On the other hand, there are three-color orderings which perform nearly as poorly as two-color orderings.

6. Some orderings with the largest number of parallel steps perform nearly as well as the natural ordering. However, in some cases such orderings perform poorly. Therefore, a large number of parallel steps is not necessarily a guarantee of very fast convergence.

7. The "lower envelope" on data values for low numbers of parallel steps seems to form a decreasing lower bound on the condition number as the number of parallel steps increases. This suggests the result of Section 3.

The convergence behavior of these methods is often analyzed in terms of the remainder matrix $R = A - Q$. Note that the model problem has constant main diagonal; thus, we may scale A to have $D = I$ without changing the performance of the methods. In this case, $Q = (I + L)(I + U)$, and thus $R = -LU$.

Though it is unclear whether a tight relationship holds between the size of R and the condition of $Q^{-1}A$, it has been observed that the squared Frobenius norm $||R||_F^2$ and the

condition κ are related. In fact, Figure 21 gives a scatter plot relating these two values for the case of the 3x3 grid, and the correlation is almost perfectly linear. In turn, large contributions to $||R||_F^2$ are made by "incompatible nodes" [Doi/Lichnewsky] or "non-naturally ordered nodes" [Eijkhout], which gives some indication of how the orderings impact the convergence rate.

3. Theoretical limitations on the performance of sparse preconditioners. The empirical results of the previous section suggest that it may be possible to put a lower bound on the condition number of the preconditioned system based on the number of parallel steps implied by the ordering. The theorem presented in this section does just this. It is based on the observation that if an ordering has a small number of parallel steps, this in fact puts a limit on the sparsity of Q, i.e., the maximal number of nonzeros per row. By showing that the condition number of a preconditioned system for which Q's nonzeros per row is bounded independent of problem size is bounded below, we are able to show the desired result, and at the same time show a convergence rate for other preconditioners such as certain sparse approximate inverse methods.

Let us begin with definitions. Let $\lambda_{min}(M)$ and $\lambda_{max}(M)$ be the extremal eigenvalues of any matrix M with real spectrum. Let the quantity $\kappa(M) = \lambda_{max}(M)/\lambda_{min}(M)$ be the spectral condition number of M.

The following result applies to d-dimensional regular Laplace equations such as those described in the previous section.

Theorem. Let d, $n \geq 1$. Let $A_{n,1}$ be the square sparse matrix of dimension n defined by $e_i^* A_{n,1} e_i = 1$ for $1 \leq i \leq n$ and $e_i^* A_{n,1} e_{i+1} = e_{i+1}^* A_{n,1} e_i = -1/2$ for $1 \leq i \leq n-1$, with all other elements of $A_{n,1}$ equal to zero. For fixed d let

$$A_{n,d} = \frac{1}{d} \sum_{i=1}^{d} [\otimes_{j=1}^{i-1} I_n] \otimes A_{n,1} \otimes [\otimes_{j=i+1}^{d} I_n].$$

Let $M_{n,d}$ be any symmetric matrix such that each row of $M_{n,d}$ has no more than b entries, for some $b \geq 1$. Then there exists a constant $C > 0$ independent of n such that for all $n \geq 1$, $\kappa(M_{n,d} A_{n,d}) \geq Cn^2$.

Proof: Since the eigenvalues of $M_{n,d} A_{n,d}$ and $A_{n,d}^{1/2} M_{n,d} A_{n,d}^{1/2}$ are the same, we may equivalently consider $\kappa(A_{n,d}^{1/2} M_{n,d} A_{n,d}^{1/2})$. By constant scaling, it is enough to consider $M_{n,d}$ such that $\lambda_{max}(A_{n,d}^{1/2} M_{n,d} A_{n,d}^{1/2}) = 1$. Thus it is sufficient to show that for such $M_{n,d}$,

$$\lambda_{min}(A_{n,d}^{1/2} M_{n,d} A_{n,d}^{1/2}) \leq [1/C]/n^2.$$

First we will show that for such $M_{n,d}$, $\|M_{n,d}\|$ is bounded independent of n.

Let us examine the eigenvalue decomposition of $A_{n,d}$. Let $V_{n,1} = \{\sin(\pi k\ell/(n+1))\}_{k,\ell=1}^n$ and $D_{n,1} = [\sqrt{(n+1)/2}]I_n$. Let $\hat{V}_{n,1} = V_{n,1}D_{n,1}^{-1}$. Let $\hat{V}_{n,d} = \otimes_{i=1}^d \hat{V}_{n,1}$. Note $A_{n,d}\hat{V}_{n,d} = \hat{V}_{n,d}\Lambda_{n,d}$, where

$$\Lambda_{n,d} = \frac{1}{d}\sum_{i=1}^d [\otimes_{j=1}^{i-1} I_n] \otimes \Lambda_{n,1} \otimes [\otimes_{j=i+1}^d I_n]$$

and $\Lambda_{n,1} = \text{diag}\{(1 - \cos(k\pi/(n+1)))/2\}_{k=1}^n$. Also, $\hat{V}_{n,d}^* \hat{V}_{n,d} = I$. Thus,

$$\Lambda_{n,d}^{1/2} \hat{V}_{n,d}^* M_{n,d} \hat{V}_{n,d} \Lambda_{n,d}^{1/2} = \hat{V}_{n,d}^* A_{n,d}^{1/2} M_{n,d} A_{n,d}^{1/2} \hat{V}_{n,d}.$$

Note that $\Lambda_{n,d}$ is a diagonal matrix whose entries are $[1/2] - [1/(2d)]\sum_{i=1}^d \cos(k_i\pi/(n+1))$, for $1 \leq k_i \leq n$ and $1 \leq i \leq d$.

Let $c = (1/2)(4b)^{-2/d}$. Note $0 < c < 1$. Let $[\otimes_{i=1}^d I_n] = F_{n,d} + G_{n,d}$, where $F_{n,d} = \otimes_{i=1}^d F_{n,1}$ and

$$F_{n,1} = \begin{bmatrix} I_{cn} & 0 \\ 0 & 0 \end{bmatrix},$$

where I_{cn} is the identity matrix of size $\lfloor cn \rfloor$. Note $0 \leq \lfloor cn \rfloor \leq cn \leq \lfloor cn \rfloor + 1 \leq n$, and the rank of $F_{n,d}$ is $\lfloor cn \rfloor^d$. Then

$$G_{n,d}^* \hat{V}_{n,d}^* M_{n,d} \hat{V}_{n,d} G_{n,d} = G_{n,d}^* \Lambda_{n,d}^{-1/2} \hat{V}_{n,d}^* A_{n,d}^{1/2} M_{n,d} A_{n,d}^{1/2} \hat{V}_{n,d} \Lambda_{n,d}^{-1/2} G_{n,d}. \qquad (1)$$

Note also

$$\|\Lambda_{n,d}^{-1/2} G_{n,d}\|^2 = \frac{2}{(1 - \cos((\lfloor cn \rfloor + 1)\pi/(n+1)))}$$

$$\leq \frac{2}{(1 - \cos(cn\pi/(n+1)))} \leq \frac{2}{(1 - \cos(c\pi/2))} = C_1$$

is bounded independent of n, thus the 2-norm of the right hand side of (1) is bounded by C_1, independent of n.

Now

$$\hat{V}_{n,d} G_{n,d}^* \hat{V}_{n,d}^* M_{n,d} \hat{V}_{n,d} G_{n,d} \hat{V}_{n,d}^* = \hat{V}_{n,d}(I_{nd} - F_{n,d})\hat{V}_{n,d}^* M_{n,d} \hat{V}_{n,d}(I_{nd} - F_{n,d})\hat{V}_{n,d}^*$$

$$= M_{n,d} - [\hat{V}_{n,d} F_{n,d} \hat{V}_{n,d}^*]M_{n,d} - M_{n,d}[\hat{V}_{n,d} F_{n,d} \hat{V}_{n,d}^*]$$

$$+[\hat{V}_{n,d}F_{n,d}\hat{V}^*_{n,d}]M_{n,d}[\hat{V}_{n,d}F_{n,d}\hat{V}^*_{n,d}]. \qquad (2)$$

Note that $\hat{V}_{n,d}$ has entries of magnitude bounded by $(\sqrt{2/(n+1)})^d$. From this it follows that

$$\|e^*_i\hat{V}_{n,d}F_{n,d}\| \le \sqrt{\lfloor cn\rfloor^d(2/(n+1))^d} \le (\sqrt{2c})^d,$$

for any i. Similarly, $\|F_{n,d}\hat{V}^*_{n,d}e_i\| \le (\sqrt{2c})^d$. Since $\hat{V}_{n,d}$ unitary, we have $\|e^*_i\hat{V}_{n,d}F_{n,d}\hat{V}^*_{n,d}\| \le (\sqrt{2c})^d$, and furthermore $\|\hat{V}_{n,d}F_{n,d}\hat{V}^*_{n,d}e_i\| \le (\sqrt{2c})^d$.

Thus, from (2),

$$|e^*_iM_{n,d}e_j| \le (2(\sqrt{2c})^d + (2c)^d)\|M_n\| + C_1 \le 3(\sqrt{2c})^d\|M_n\| + C_1$$

for any i, j. Thus there are symmetric matrices $\hat{M}_{n,d}$ and $\check{M}_{n,d}$ of the same sparsity pattern as $M_{n,d}$ whose entries are bounded in magnitude by 1 such that

$$M_{n,d} = 3(\sqrt{2c})^d\|M_{n,d}\|\hat{M}_{n,d} + C_1\check{M}_{n,d}.$$

But by Gershgorin's theorem, $\|\hat{M}_{n,d}\|$ and $\|\check{M}_{n,d}\|$ are each bounded by b. Thus

$$\|M_{n,d}\| \le 3(\sqrt{2c})^db\|M_{n,d}\| + C_1b.$$

This implies

$$\|M_{n,d}\|(1 - 3(\sqrt{2c})^db) \le C_1b.$$

By definition of c, this implies

$$\|M_{n,d}\| \le 4C_1b,$$

so $\|M_{n,d}\|$ is bounded.

Let $v_n = \hat{V}_{n,d}[\otimes^d_{i=1}e_1]$, the (normalized) eigenvector of $A_{n,d}$ corresponding to the smallest eigenvalue. Then

$$\lambda_{min}(A^{1/2}_{n,d}M_{n,d}A^{1/2}_{n,d}) \le v^*_nA^{1/2}_{n,d}M_{n,d}A^{1/2}_{n,d}v_n$$

$$= \frac{1-\cos(\pi/(n+1))}{2}v^*_nM_{n,d}v_n \le \frac{1-\cos(\pi/(n+1))}{2}\|M_{n,d}\|.$$

By Taylor's theorem,

$$\left|\frac{1-\cos(\pi/(n+1))}{2} - \frac{\pi^2}{4(n+1)^2}\right| \le \frac{\pi^4}{48(n+1)^4}.$$

Thus for $n \geq 1$,

$$\lambda_{min}(A_{n,d}^{1/2} M_{n,d} A_{n,d}^{1/2}) \leq \frac{\pi^2}{2(n+1)^2} \|M_{n,d}\| \leq \frac{\pi^2}{8n^2} \|M_{n,d}\|.$$

Since $\lambda_{max}(A_{n,d}^{1/2} M_{n,d} A_{n,d}^{1/2}) = 1$ and since $\|M_{n,d}\|$ is bounded independent of n, the result is shown, with

$$C = \frac{1}{\pi^2 b}(1 - \cos(\pi/(4(4b)^{2/d}))).$$

∎

Note that in fact the following bound holds, for $c = (1/2)(4b)^{-2/d}$:

$$\kappa(M_{n,d} A_{n,d}) \geq \frac{1 - (2(\sqrt{2c})^d + (2c)^d)b}{b} \frac{1 - \cos(cn\pi/(n+1))}{1 - \cos(\pi/(n+1))}$$

$$\geq \frac{1}{\pi^2 b}(1 - \cos(\pi/(4(4b)^{2/d})))n^2.$$

Table 2 shows some representative values of the lower bound of the condition number using this tighter bound. It is assumed here that the number of elements of $A_{n,d}$, i.e. n^d, is 10^9. Note that the condition numbers for the unpreconditioned systems for $d = 1$, 2 and 3, given by $(1 + \cos(\pi/(n+1)))/(1 - \cos(\pi/(n+1)))$, are $4.0528e+17$, $4.0531e+08$ and $4.0610e+05$, respectively, suggesting that the bound is weak.

Tables 3 through 5 give *upper* bounds for the condition number for the preconditioned system using the optimal preconditioner for a given sparsity pattern. The values are based on Chebyshev polynomial preconditioning of the original matrix. These figures are based on the fact that the associated spectral radius for degree-k polynomial preconditioning is given by $1/\cosh((k+1)\log((\sqrt{\kappa}-1)/(\sqrt{\kappa}+1)))$, where κ is the condition number of the original matrix. Again, the difference between the lower and upper bounds is fairly substantial, indicating that the bounds are weak, though the asymptotic behavior of the bounds indicate that such preconditioning can only improve the condition number of the system by a constant amount irrespective of the problem size.

b	bound, $d = 1$	bound, $d = 2$	bound, $d = 3$
1	4.269030439932792e+14	6.748544386655379e+06	1.667595146195060e+04
2	1.430439665228187e+13	9.125899360897805e+05	3.615384676085781e+03
3	1.925644271241498e+12	2.768998273522413e+05	1.442617057436141e+03
4	4.619345477052554e+11	1.181606700624436e+05	7.470629594715116e+02
5	1.523435542232577e+11	6.090618404239972e+04	4.472971144458266e+02
6	6.148513730152788e+10	3.540282128498064e+04	2.937963276729456e+02
7	2.853345519941855e+10	2.236437040612905e+04	2.057759094330227e+02
8	1.466832767451235e+10	1.501735487411991e+04	1.510904184372026e+02
9	8.154236430678563e+09	1.056621528182532e+04	1.150201295009818e+02
10	4.821776956584233e+09	7.713852431773913e+03	9.009815904167844e+01
11	2.997389479188783e+09	5.802330105678729e+03	7.222877142738187e+01
12	1.941852206587198e+09	4.473628423024728e+03	5.902222420273417e+01
13	1.302438301834628e+09	3.521525491870981e+03	4.901254221151281e+01
14	8.997768525540059e+08	2.821515277786319e+03	4.126234588520582e+01
15	6.376457374211799e+08	2.295393593880919e+03	3.515094870732632e+01
16	4.620233028072453e+08	1.892352481355578e+03	3.025492079562255e+01
17	3.413660263764376e+08	1.578406337283923e+03	2.627785885359187e+01
18	2.566144242327273e+08	1.330236426188525e+03	2.300755825865897e+01
19	1.959004344586539e+08	1.131480641219106e+03	2.028905490699708e+01
20	1.516342154863592e+08	9.704278090094737e+02	1.800717759966682e+01
21	1.188449579287516e+08	8.385455934108473e+02	1.607497414330090e+01
22	9.420675040977293e+07	7.295177085026331e+02	1.442586326643019e+01
23	7.545084936784838e+07	6.386004785643117e+02	1.300820699191690e+01
24	6.100226725721111e+07	5.621843546788020e+02	1.178148956451273e+01
25	4.974999976031717e+07	4.974897717065060e+02	1.071358361418229e+01
26	4.089877660773254e+07	4.423527603434260e+02	9.778765199218023e+00
27	3.387161736121851e+07	3.950715821205910e+02	8.956253044489959e+00
28	2.824468496991656e+07	3.542955218269618e+02	8.229120156699228e+00
29	2.370304033552378e+07	3.189432358839651e+02	7.583473584101552e+00
30	2.001015406729655e+07	2.881421052333836e+02	7.007829699144931e+00

Table 2. Lower bound on condition number.

d	b	bound
1	1	4.05284735379920e+17
1	3	1.01321183844980e+17
1	5	4.50316372644356e+15
1	7	2.53302959612450e+15
1	9	1.62113894151968e+15
1	11	1.12579093161089e+15
1	13	8.27111704856980e+15
1	15	6.33257399031125e+15
1	17	5.00351525160395e+15
1	19	4.05284735379920e+15
1	21	3.34946062297455e+15
1	23	2.81447732902722e+15
1	25	2.39813452887527e+15
1	27	2.06777926214245e+15
1	29	1.80126549057742e+15
1	31	1.58314349757781e+15
1	33	1.40236932657412e+15
1	35	1.25087881290099e+15
1	37	1.12267239717429e+15
1	39	1.01321183844980e+15
1	41	9.19013005396645e+14

Table 3. Upper bound on condition number.

d	b	bound
2	1	4.05310363247520e+08
2	5	1.01327592341674e+08
2	13	4.50344857177101e+07
2	25	2.53318985465594e+07
2	41	1.62124153061923e+07
2	61	1.12586219459801e+07
2	85	8.27164079201994e+06
2	113	6.33297513785550e+06
2	145	5.00383234611387e+06
2	181	4.05310432678792e+06
2	221	3.34967311412197e+06
2	265	2.81465598721645e+06
2	313	2.39828685651837e+06
2	365	2.06791069781018e+06
2	421	1.80138007170369e+06
2	481	1.58324428415013e+06
2	545	1.40245868091822e+06
2	613	1.25095858666467e+06
2	685	1.12274406310617e+06
2	761	1.01327658188175e+06
2	841	9.19071791570317e+05

Table 4. Upper bound on condition number.

d	b	bound
3	1	4.060950426525602e+05
3	7	1.01524260664692e+05
3	25	4.51222640004881e+04
3	63	2.53815651686360e+04
3	129	1.62444417103714e+04
3	231	1.12810660056218e+04
3	377	8.28830700103547e+03
3	575	6.34589130201644e+03
3	833	5.01417749465436e+03
3	1159	4.06161044298305e+03
3	1561	3.35681855421417e+03
3	2047	2.82076652356833e+03
3	2625	2.40359199612413e+03
3	3303	2.07257678042185e+03
3	4089	1.80553059688894e+03
3	4991	1.58697286489927e+03
3	6017	1.40583756528113e+03
3	7175	1.25404442352200e+03
3	8473	1.12558189486907e+03
3	9919	1.01590267229735e+03
3	11521	9.21515664198286e+02

Table 5. Upper bound on condition number.

The following result allows us to leverage the previous results for application to a larger class of matrices of interest, in particular, matrices which are spectrally equivalent to those of the previous theorem.

Theorem. Let d, $n \geq 1$. Let $A_{n,d}$ and $M_{n,d}$ be defined as in the previous theorem. Let $B_{n,d}$ be any symmetric positive definite matrix of dimension n^d such that $\kappa(B_{n,d}^{-1}A_{n,d})$ is bounded independent of n. Then there exists a constant $C' > 0$ independent of n such that for all $n \geq 1$, $\kappa(M_{n,d}B_{n,d}) \geq C'n^2$.

Proof: Note that for any symmetric nonsingular A, B and M,

$$\lambda_{max}(MA) = \max_v \frac{v^*Av}{v^*M^{-1}v} \leq \max_v \frac{v^*Av}{v^*Bv} \max_v \frac{v^*Bv}{v^*M^{-1}v} = \lambda_{max}(B^{-1}A)\lambda_{max}(MB),$$

$$\lambda_{min}(MA) = \min_v \frac{v^*Av}{v^*M^{-1}v} \geq \min_v \frac{v^*Av}{v^*Bv} \min_v \frac{v^*Bv}{v^*M^{-1}v} = \lambda_{min}(B^{-1}A)\lambda_{min}(MB),$$

thus $\kappa(MA) \leq \kappa(B^{-1}A)\kappa(MB)$. Therefore, $\kappa(M_{n,d}B_{n,d}) \geq (C/\kappa(B_{n,d}^{-1}A_{n,d}))n^2$, giving the result. ■

This result applies directly to discretized differential operators that are spectrally equiv-

alent to the diffusion operator in d dimensions. Consider the boundary value problem

$$-\sum_{i=1}^{d} \frac{d}{dx_i} c_i(x) \frac{d}{dx_i} u(x) = f(x), \quad x \in \Omega = [0,1]^d,$$

$$u(x) = g(x), \quad x \in \partial\Omega.$$

Suppose this is discretized by central $(2d+1)$-point finite differencing. Then setting $c_i(x) = 1$ gives rise to the matrix described in the first theorem. Furthermore, under appropriate conditions on c_i, the resulting matrix $B_{n,d}$ satisfies the conditions of the previous theorem. The result for such matrices is that any set of sparse preconditioners $M_{n,d}$ satisfying the given sparsity properties must necessarily give condition number growth order n^2, whose growth rate is no better than that for the unpreconditioned system.

We will now prove a sequence of results to apply the above theorems.

Corollary.(*Sparse Approximate Inverses.*) Let $d, n \geq 1$. Let $B_{n,d}$ be defined as previously. Let $M_{n,d}$ be a matrix whose sparsity pattern $\Sigma(M_{n,d})$ satisfies $\Sigma(M_{n,d}) = \cup_{i=0}^{k}\Sigma(A_{n,d}^i)$, for a fixed value of $k \geq 0$. Then there exists a constant $C' > 0$ independent of n such that for all $n \geq 1$, $\kappa(M_{n,d}B_{n,d}) \geq C'n^2$.

Proof: Let $b_{d,k}$ be the maximal number of nonzeros per row of $M_{n,d}$. Note $b_{d,k} \leq (2k+1)^d$. The number $b_{d,k}$ denotes the number of nodes in the graph corresponding to $\Sigma(A_{n,d})$ that can be traversed in no more than k steps. This satisfies the recurrence $b_{d,k} = b_{d-1,k} + 2\sum_{i=0}^{k-1} b_{d-1,i}$. In particular, $b_{1,k} = 2k+1$, $b_{2,k} = 2k(k+1)+1$, and $b_{3,k} = 2k+1+2k(k+1)(2k+1)/3$. ∎

Corollary.(*Polynomial Preconditioning.*) Let $d, n \geq 1$. Let $B_{n,d}$ be defined as previously. Let $M_{n,d}$ be a polynomial of degree no greater than k in $B_{n,d}$, for a fixed value of $k \geq 0$. Then there exists a constant $C' > 0$ independent of n such that for all $n \geq 1$, $\kappa(M_{n,d}B_{n,d}) \geq C'n^2$.

Proof: The number of nonzeros per row has bound equal to that for the previous corollary. ∎

Corollary.(*Factorized Sparse Approximate Inverses.*) Let $d, n \geq 1$. Let $B_{n,d}$ be defined as previously. Let $L_{n,d}$ be a lower triangular matrix whose sparsity pattern is any subset of $\cup_{i=0}^{k}\Sigma(A_{n,d}^i)$, for a fixed value of $k \geq 0$. For any diagonal matrix $D_{n,d}$ let $M_{n,d} = L_{n,d}D_{n,d}L_{n,d}^*$. Then there exists a constant $C' > 0$ independent of n such that for all $n \geq 1$, $\kappa(M_{n,d}B_{n,d}) \geq C'n^2$.

Proof: The number of nonzeros per row of $M_{n,d}$ is bounded by $[(2k+1)^d]^2$. A tighter bound is given by $2k+1$ ($d=1$) (sharp), $(k^2+k+1)^2$ ($d=2$), and $(1+k+2k^2+k(k-1)(2k-1)/3)^2$ ($d=3$). ∎

Corollary.(*Fixed-Size Overlapping Subgrid Preconditioning.*) Let d, $n \geq 1$. Let $B_{n,d}$ be defined as previously. Let $M_{n,d}$ be a sum of block diagonal matrices each of which is an arbitrary symmetric matrix on a subgrid of the grid and zero elsewhere. The subgrids may overlap. Let b bound the rank of any of these subgrid matrices, bounded independent of n. Then there exists a constant $C' > 0$ independent of n such that for all $n \geq 1$, $\kappa(M_{n,d}B_{n,d}) \geq C'n^2$.

Corollary.(*Banded Preconditioning.*) Let d, $n \geq 1$. Let $B_{n,d}$ be defined as previously. Let $M_{n,d}$ be a banded matrix of bandwidth bounded by b, where b is bounded independent of n. Then there exists a constant $C' > 0$ independent of n such that for all $n \geq 1$, $\kappa(M_{n,d}B_{n,d}) \geq C'n^2$.

Corollary.(*IC/MIC/SSOR.*) Let d, $n \geq 1$. Let $B_{n,d}$ be defined as previously. Suppose a 2-sided permutation is applied to $B_{n,d}$, and $M_{n,d}$ represents either IC(0), MIC(0) or SSOR preconditioning for $B_{n,d}$. Let ℓ_n be the number of "wavefronts," or parallel steps required to apply a forward or backward sweep of $M_{n,d}$, where ℓ_n is bounded over n. Then there exists a constant $C' > 0$ independent of n such that for all $n \geq 1$, $\kappa(M_{n,d}B_{n,d}) \geq C'n^2$.

Proof: Let $M_{n,d} = (I - L^*)^{-1}D(I - L)^{-1}$, for D diagonal and L strictly lower triangular. Note ℓ_n is the lowest integer for which $L^{\ell_n} = 0$. To bound the number of nonzeros per row of $M_{n,d}$, first note that $(I - L)^{-1} = \sum_{i=0}^{\ell_n-1} L^i$, and L^i has at most $(2d)^i$ nonzeros. Thus, $M_{n,d}$ has at most $(2d)^{2\ell_n}$ nonzeros. A second bound may be obtained by noting that applying $(I - L^*)^{-1}$ to a vector can connect a point to neighbors within an enclosing cube of size $2\ell_n + 1$ points per edge; thus, the number of nonzeros per row of $M_{n,d}$ is bounded by $(2\ell_n + 1)^{2d}$.

Note for red/black ordering ($\ell_n = 2$), these tighter bounds hold: 5 ($d=1$), 13 ($d=2$), and 25 ($d=3$). ∎

Corollary.(*Multicolor IC/MIC/SSOR.*) Let d, $n \geq 1$. Let $B_{n,d}$ be defined as previously. Suppose a 2-sided permutation is applied to $B_{n,d}$, and $M_{n,d}$ represents either IC(0), MIC(0) or SSOR preconditioning for $B_{n,d}$. Suppose the ordering represents a coloring of the graph

into k independent sets. Then there exists a constant $C' > 0$ independent of n such that for all $n \geq 1$, $\kappa(M_{n,d}B_{n,d}) \geq C'n^2$.

Proof: A graph colored with k colors requires k parallel steps to perform a forward or backward sweep. ∎

Corollary. *(M/IC(k).)* Let d, $n \geq 1$. Let $B_{n,d}$ be defined as previously. Suppose a 2-sided permutation is applied to $B_{n,d}$, and $M_{n,d}$ represents either IC(k) or MIC(k) preconditioning for $B_{n,d}$. Let ℓ_n be the number of "wavefronts," or parallel steps required to apply a forward or backward sweep of $M_{n,d}$, where ℓ_n is bounded over n. Then there exists a constant $C' > 0$ independent of n such that for all $n \geq 1$, $\kappa(M_{n,d}B_{n,d}) \geq C'n^2$.

Proof: A forward or backward sweep connects a point to neighbors as far away as $\ell_n 2^k$. Thus the number of nonzeros per row of $M_{n,d}$ is bounded by $(2\ell_n 2^k + 1)^{2d}$. A second bound on the number of nonzeros per row, as counted by connections, is $\left(\sum_{i=0}^{k}(2d)^{2^i}\right)^{2\ell_n} \leq (2d)^{2\ell_n(2^k+1)}$. ∎

We have shown that for a number of well-known preconditioners, preconditioning can only improve the iteration count by a constant factor independent of the grid size over the unpreconditioned case. One can compare this, for example, with MIC preconditioning on the 2-D Laplace equation matrix with natural ordering, for which the condition is of order n, rather than order n^2. This has to do with the fact that in the latter case, the matrix $M = Q^{-1}$, which is intended to approximate the dense matrix A^{-1}, is a dense matrix, due to the fact that the number of parallel steps grows with the problem size, and thus is not subject to the limitations imposed by the above results for sparse preconditioners.

These results address some questions raised in the paper [Greenbaum/Rodrigue] regarding optimal preconditionings of a given sparsity pattern. In particular, it is shown that for preconditioners with fixed sparsity pattern, the condition of the preconditioned system must grow by order n^2, regardless of what specific type of preconditioner is used.

For practical problems, one wants the iteration count for the model problem to grow as n^α, with $\alpha \geq 0$ as small as possible. These results show that certain entire approaches to preconditioning cannot by themselves improve over $\alpha = 2$—that is, the methods do not fail to attain $\alpha < 2$ due to some minor deficiency in how the coefficients are chosen, but it is due to the intrinsic nature of the approach.

Of course, the constant multiplier in front of n^2 may be quite small for these methods,

and for small to intermediate problem regimes such methods may in fact be the most effective. Such methods may also have robustness properties that make them more able to converge than than other "faster" methods which cannot be made to work for larger problems. Furthermore, it is unclear whether for all modeling and simulation problems of interest there exist methods for which $\alpha < 2$, e.g., for unstructured problems with complex physics, for which the growth rate of n is measured in some appropriate way, e.g., based on a uniform grid refinement scheme. Thus, the methods described here are by no means obsolete.

One may ask similar questions for other orderings for SSOR and ILU preconditionings. Let $M_{n,d}^{-1} = Q_{n,d} = L_{n,d}U_{n,d}$ for some sparse lower and upper triangular matrices $L_{n,d}$ and $U_{n,d}$, and assume that $L_{n,d}$ and $U_{n,d}$ have a bounded number of nonzeros per row independent of the problem size, and let $d \geq 2$ be the dimension of the model problem. Then there exists α defined by

$$\alpha = \sup\{\alpha' : \exists c : \kappa(M_{n,d}A_{n,d}) \geq cn^{\alpha'} \forall M_{n,d}\},$$

where $M_{n,d}$ is defined as above. For $d = 2$ it is known that $\alpha \leq 1$, though in general α is not known, and it would be worthwhile to know what its value is.

4. Conclusions. This paper has examined the effect of preconditioner orderings for a set of small test problems and has presented a new theoretical result on the limitations of certain orderings. Further research may shed more light on how orderings impact convergence and may suggest improved algorithmic approaches to solving large linear systems of interest in parallel.

REFERENCES

[1]   Doi, Shun and Alain Lichnewsky, "A Graph-Theory Approach for Analyzing the Effects of Ordering on ILU Preconditioning," INRIA report 1452, June 1991.

[2]   Duff, I. S. and G. A. Meurant, "The Effect of Ordering on Preconditioned Conjugate Gradient," *BIT*, 29:635–657, 1989.

[3] Eijkhout, Victor, "Analysis of parallel incomplete point factorizations," *Linear Algebra Appl.*, 154/156 (1991), 723–740.

[4] Greenbaum, A., and G. H. Rodrigue, "Optimal Preconditioners of a Given Sparsity Pattern," *BIT* 29 (1989), pp. 610-634.

[5] Ortega, James M., "Orderings for Conjugate Gradient Preconditionings," *SIAM J. Optimization*, vol. 1, no. 4, pp. 565–582, November 1991.

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