Orderings for Incomplete Factorization Preconditioning of Nonsymmetric Linear Systems

Michele Benzi
Scientific Computing Group,
Los Alamos National Laboratory,
New Mexico 87545, USA
benzi@lanl.gov

Daniel B. Szyld
Department of Mathematics,
Temple University, Philadelphia,
Pennsylvania 19122-6094, USA
szyld@math.temple.edu

Arno C.N. van Duin
Department of Computer Science,
Leiden University, Leiden,
The Netherlands
arno@cs.leidenuniv.nl

Abstract *Numerical experiments are presented whereby the effect of reorderings on the convergence of preconditioned Krylov subspace methods for the solution of nonsymmetric linear systems is shown. The preconditioners used in this study are different variants of incomplete factorizations. It is shown that reorderings for direct methods, such as Reverse Cuthill-McKee and Minimum Degree, can be very beneficial. The benefit can be seen in the reduction of the number of iterations and also in measuring the appropriate deviation of the preconditioned operator from the identity.*

1 INTRODUCTION

In this paper, we study experimentally how different reorderings affect the convergence of Krylov subspace methods for nonsymmetric systems of linear equations when incomplete LU factorizations are used as preconditioners. Given a sparse linear system of equations $Av = b$, where v and b are n-dimensional vectors, we consider symmetric permutations of the matrix A, i.e., of the form $P^T AP$, and then solve the equivalent system $P^T APw = P^T b$, with $v = Pw$, by way of some preconditioned iterative method. We consider several iterative methods for nonsymmetric systems, including GMRES,Bi-CGSTAB,and transpose free QMR (TFQMR); for a description of these, as well as for a description of incomplete factorizations, see, e.g., [1].

In this study, we mainly concentrate on orderings originally devised for matrix factorizations, i.e., those used to reduce fill-in in the factors. We want to call attention to the fact that a permutation of the variables (and equations) using reordering methods designed for direct solvers can have an important positive effect on the robustness and performance of preconditioned Krylov subspace methods when applied to nonsymmetric linear systems. This is especially the case if the matrices are far from symmetric. This observation, although not completely new, is not fully appreciated, to the point that some authors have concluded that direct solver reorderings should not be used with preconditioned iterative methods. It is hoped that the present study will contribute to a reassessment of direct solver reorderings in the context of incomplete factorization preconditioning. Furthermore, we hope that the evidence of our experiments can set the stage for more widespread use of these reorderings.

In Sect. 2 we survey some of the contributions on the effect of permutations on the convergence of preconditioned Krylov subspace methods. In Sect. 3 we present our numerical experiments and comment on those results. Finally, in Sect. 4 we present our conclusions. This paper is an abbreviated version of [3]

2 OVERVIEW OF THE LITERATURE

The influence of reorderings on the convergence of preconditioned iterative methods has been considered by a number of authors. Several of these papers are concerned with symmetric problems only [4, 8, 12, 14, 21, 22, 23, 27]. Duff and Meurant [12] have performed a very detailed study of the effects of reorderings for preconditioned conjugate gradients, i.e., in the symmetric positive definite case. Based on their extensive experiments, they concluded that direct solver reorderings the number of iterations required for convergence is usually about the same, and sometimes considerably higher, than with the natural (lexicographic) ordering. Throughout the paper we abuse the notation and use A to denote both the original matrix and the permuted one. Similarly, $\bar{L}$ and $\bar{U}$ refer to the incomplete factors of A or those of $P^T AP$, depending on the context.

As it can be seen in the experiments in Sect. 3, when the matrices are nearly symmetric, our observations are in agreement with those of Duff and Meurant. On the other hand, for highly nonsymmetric matrices, i.e., when the nonsymmetric part is large, our conclusions are that permutations that are usually regarded as ineffective for the (nearly) symmetric case turn out to be very beneficial, often improving the robustness and performance of the preconditioned iteration dramatically. This is an important observation, since problems with a large departure from symmetry are often very difficult to solve with the "natural" ordering. We will see that for problems that are strongly nonsymmetric and/or far from being diagonally dominant, the norm of the residual matrix R alone is usually not a reliable indicator of the quality of the corresponding preconditioner.

A more revealing measure of the quality of the preconditioner can be obtained by considering the Frobenius norm of the deviation of the preconditioned matrix from the identity, i.e., $\|I - A(\bar{L}\bar{U})^{-1}\|_F$; see, e.g. [7]. This quantity is equal to $\|R(\bar{L}\bar{U})^{-1}\|_F$. Even if R is small in norm, it could happen that $(\bar{L}\bar{U})^{-1}$ has very large entries, resulting in a large deviation of the preconditioned matrix from the identity. As a result, the preconditioned iteration fails to converge. A notable example of this phenomenon occurs when convection-dominated convection-diffusion equations are discretized with centered finite differences. When the lexicographic ordering is used, the incomplete factors resulting from a no-fill ILU factorization tend to be very ill-conditioned, having very large entries in their inverses, even if the coefficient matrix itself is well-conditioned. Allowing more fill-in in the factors (e.g., using ILU(1) rather than ILU(0)) will solve the problem in some cases, but not always.

This kind of instability of ILU factorizations was first noticed by van der Vorst [28], and analyzed in detail by Elman [15]. We will see that in some cases, reordering the coefficient matrix before performing the incomplete factorization can have the effect of producing stable triangular factors, and hence more effective preconditioners.

The effects of permutations on preconditioned Krylov subspace methods for nonsymmetric problems have been considered in [5, 8, 9, 10, 11, 13, 20, 26]. Some authors have concluded that the reorderings designed for sparse direct solvers are not recommended for use with preconditioned iterative methods; see, e.g., [8], [20], [26].

Of course, there have been cases in which such preorderings have been tried for nonsymmetric problems. Simon [26] used Quotient Minimum Degree and Nested Dissection [17] in conjunction with an ILU preconditioner for some oil simulation problems, and found essentially no improvement over the original ordering. Similar conclusions were reached by Langtangen [20], who applied a Minimum Degree reordering (MD) with ILU(0) preconditioning of matrices arising from a Petrov-Galerkin formulation for convection-diffusion equations. Dutto [13], in the context of a specific application (solving the compressible Navier-Stokes equations with finite elements on unstructured grids), was possibly the first to observe that MD and other direct solver reorderings can have a positive effect on the convergence of GMRES with ILU(0) preconditioning. This is mostly consistent with some of our own experiments reported here. For further review of the literature, see [3].

We also mention the Minimum Discarded Fill (MDF) algorithm (see [8]) which takes into account the numerical values of the entries of A. This method can be very effective, but it is often too expensive to be practical, except for rather simple problems.

In addition to reorderings which were originally designed for direct methods, one can use the permutation produced by the algorithm TPABLO [6], which also uses the magnitude of the entries of the matrix. This algorithm produces a permuted matrix with dense diagonal blocks, while the entries outside the blocks on the diagonal have magnitude below a prescribed threshold. Hence, like MDF, the TPABLO reordering is based both on graph information and on the numerical values. The motivation for TPABLO was to produce good block diagonal preconditioners, and also blocks for the treatment of certain Markov chain problems. It turns out that this reordering is also useful for point incomplete factorizations, where it is often better than the natural ordering [2]. However, for most cases considered in this paper, TPABLO is not as good as the reorderings designed for direct methods. This reordering might prove useful in the context of *block* incomplete factorizations, but this topic is outside the scope of the present paper.

3 NUMERICAL EXPERIMENTS

In this section we show, by means of numerical experiments, that direct solver reorderings can be very beneficial when solving *difficult* nonsymmetric linear systems (obviously, at best, little gains can be expected from reordering problems which are easily solved with the original ordering). The results reported here are a representative selection from a large number of experiments with nonsymmetric matrices arising from the numerical solution of partial differential equations. In the first subsection we focus on an important class of problems, while in the second we investigate reasons why reordering improves the performance of the preconditioners. Additional experiments on a selection of nonsymmetric matrices from various sources can be found in [3].

Convection-Diffusion Equations

A source of linear systems which are challenging for iterative methods is the following partial differential equation in the open unit square $\Omega = (0,1) \times (0,1)$

$$-\varepsilon \Delta u + \frac{\partial e^{xy} u}{\partial x} + \frac{\partial e^{-xy} u}{\partial y} = g \qquad (1)$$

with homogeneous Dirichlet boundary conditions. Equation (1) has been repeatedly used as a model problem in the literature; see, e.g., [24]. The problem is discretized using centered differences for both the second order and first order derivatives with grid size $h = 1/33$, leading to a block tridiagonal linear system of order 1024 with 4992 nonzero coefficients. Experiments were also performed with finer grids and with a three-dimensional analogue of the same PDE, and the results obtained were similar to those reported in Table 1 below. For the right-hand side we used a random vector. Results similar to those in Table 1 were obtained with other choices of b. We used $v_0 = 0$ as initial guess and we stopped the iterations when the 2-norm of the (unpreconditioned) residual $b - Av_k$ had been reduced to less than 10^{-9}, or when a maximum number of iterations was reached. The parameter $\varepsilon > 0$ controls the difficulty of the problem —the smaller ε is, the harder it is to solve the discrete problem by iterative methods. For our experiments, we generated linear systems of increasing difficulty, corresponding to $\varepsilon^{-1} = 100, 200, \ldots, 1000$. The coefficient matrix A becomes more nonsymmetric as ε decreases.

We report the results of experiments with the following accelerators: Bi-CGSTAB, TFQMR, and GMRES with restart parameter $m = 20$. The preconditioners used were standard incomplete factorizations based on levels of fill (ILU(0) and ILU(1)) and Saad's dual threshold ILUT; see [25]. For ILUT, we used two different sets of parameters, $(10^{-2}, 5)$ and $(10^{-3}, 10)$. The latter results in a very powerful but expensive preconditioner, containing up to five times the number of nonzeros in A. Right preconditioning was used in all cases.

In Table 1 we present the number of iterations for the different orderings and the three Krylov subspace methods used in this study. In these tables, and in the ones that follow, n/o stands for natural (or original) ordering, CM for Cuthill-McKee, RC for Reverse Cuthill-McKee, MD for Multiple Minimum Degree, and TP for TPABLO; for a description of the direct solver reorderings see, e.g., [17]. The symbol † indicates that convergence was not achieved in 250 iterations for Bi-CGSTAB and TFQMR (two matrix-vector products and applications of the preconditioner per iteration), and in 500 iterations for GMRES(20) (one matrix-vector product and preconditioner application per iteration). In each case the number in bold indicates the run which required the least amount of work. This is not always the same as the run which

Table 1: Number of iterations for different orderings, different preconditioners

		Bi-CGSTAB					GMRES(20)					TFQMR				
preconditioner	ε^{-1}	n/o	CM	RC	MD	TP	n/o	CM	RC	MD	TP	n/o	CM	RC	MD	TP
ILU(0)	100	**11**	**11**	**11**	30	14	**16**	**16**	**16**	112	25	**10**	**10**	**10**	30	15
	400	50	51	**46**	59	53	**100**	**100**	**100**	139	108	**49**	51	54	60	50
	600	158	191	127	**81**	105	313	316	315	**178**	211	200	213	231	**83**	89
	800	†	†	†	**108**	132	†	†	†	**230**	395	†	†	†	**112**	143
	1000	†	†	†	**136**	†	†	†	†	**305**	†	†	†	†	**140**	220
ILU(1)	100	**6**	**6**	**6**	13	9	**10**	**10**	**10**	20	16	**6**	**6**	**6**	13	10
	400	24	**12**	13	34	56	39	**20**	**20**	64	61	21	**12**	13	32	38
	600	52	**19**	**19**	58	†	80	**30**	**30**	137	490	36	**17**	19	50	117
	800	120	**28**	**28**	83	†	153	**44**	**44**	172	295	60	26	**24**	71	102
	1000	†	43	**38**	109	†	331	63	**62**	230	†	98	32	**31**	84	†
ILUT(10^{-2}, 5)	100	4	3	**3**	6	5	6	5	**5**	12	8	4	3	**3**	7	6
	400	10	**3**	4	36	7	15	**6**	7	68	13	10	4	**4**	38	8
	600	14	5	**5**	168	10	23	9	**9**	†	18	14	6	**6**	†	12
	800	23	7	**6**	†	24	37	11	**10**	†	52	21	7	**7**	†	29
	1000	†	9	**7**	†	†	†	16	**12**	†	†	†	10	**8**	†	†
ILUT(10^{-3}, 10)	100	2	**2**	2	5	4	4	**3**	3	7	6	2	**2**	2	5	4
	400	5	2	**2**	8	5	9	4	**4**	14	9	5	**2**	3	8	5
	600	8	3	**3**	11	6	14	5	**5**	18	11	9	3	**3**	11	6
	800	13	4	**4**	12	7	19	6	**6**	23	13	13	4	**4**	13	7
	1000	17	4	**4**	15	9	29	8	**7**	28	15	17	4	**4**	17	10

required the least number of iterations because, in general, different reorderings result in preconditioners with a different number of nonzeros (except for ILU(0), where the number of nonzeros is always equal to the number of nonzero entries in A). For this particular class of matrices, the amount of fill-in in the incomplete factors is comparable for the natural ordering, Cuthill-McKee, Reverse Cuthill-McKee and TPABLO, whereas MD produces much less fill-in.

We now comment on these results. We notice that for moderately nonsymmetric problems (smaller values of ε^{-1}), the alternative permutations offer little or no advantage over the natural ordering. The Cuthill-McKee and Reverse Cuthill-McKee reorderings are essentially equivalent to the natural ordering. Incidentally, we observe that, in the case of ILU(0), the different number of iterations with the natural ordering and the Cuthill-McKee reordering is an artifact of finite precision computation, since it is known (see, e.g., [29], Sections 5.5–5.6) that for five-point stencils and no-fill factorizations, the Cuthill-McKee reorderings are equivalent to the natural ordering, in the sense that the incomplete factors of the permuted matrix are just the permuted incomplete factors of the original matrix. Hence the ILU(0) preconditioners with the natural ordering and Cuthill-McKee reorderings are mathematically equivalent. The fact that these discrepancies are sometimes rather large, as in the case of $\varepsilon^{-1} = 600$, can be explained with the observation that the ILU(0) factors are badly conditioned for small ε, and this causes large differences in the finite precision results.

For larger values of ε, TPABLO gives somewhat worse results than the natural ordering, and MD causes a serious degradation of the convergence rate, especially with ILU(0). Note that the differences in behavior are not as pronounced with the ILUT preconditioners, which use a drop tolerance. These results are in good agreement with the conclusions drawn in [12] for symmetric systems.

As the coefficient matrix becomes increasingly nonsymmetric, however, things change. While the number of iterations increases for all reorderings, the rate of increase is not the same for all reorderings, suggesting that some orderings are less sensitive than others to the degree of nonsymmetry. With ILU(1), the natural ordering and especially TPABLO perform poorly, MD is only slightly better, but the Cuthill-McKee reorderings are both very good. With ILUT(10^{-2}, 5) MD is very bad, the natural ordering and TPABLO are roughly equivalent, and the Cuthill-McKee orderings are both excellent. With ILUT(10^{-3}, 10), which gives very good (but expensive) approximations of A, all orderings produce effective preconditioners, but the performance is particularly good with the Cuthill-McKee reorderings. Notice that Reverse Cuthill-McKee is only slightly better than Cuthill-McKee.

For the more powerful preconditioners (ILU(1) and ILUT) the Cuthill-McKee reorderings give the best results for *all* values of ε. Since these reorderings were never worse than the natural ordering, and in most cases considerably better, we see no reason why Cuthill-McKee reorderings should not be routinely used for this class of problems. If ILU(0) is the only option available, then MD may be used to have greater robustness for difficult problems. The TPABLO reordering, although better than the natural ordering in many cases, was not competitive with the direct solver reorderings, and is of course more expensive to compute. For this reason, we don't consider it further in our experiments.

As far as the relative performance of the three Krylov subspace solvers is concerned, we observe that they are more or less equivalent for this particular class of problems, with GMRES(20) requiring slightly fewer matrix-vector products than the other solvers in many cases.

This set of problems was generated using second order, centered difference approximations for both the second and first partial derivatives in (1). It is well known that for large values of h/ε, this discretization can become unstable. Alternative discretizations, such as those which use upwinding for the first order terms, do not suffer from this problem and give rise to matrices with very nice properties from the point of view of iterative solutions, such as diagonal dominance. However, this may not be true for nonuniform

Table 2: Number of iterations for different orderings, several preconditioners

preconditioner	P	Bi-CGSTAB n/o	Bi-CGSTAB CM	Bi-CGSTAB RC	Bi-CGSTAB MD	GMRES(20) n/o	GMRES(20) CM	GMRES(20) RC	GMRES(20) MD	TFQMR n/o	TFQMR CM	TFQMR RC	TFQMR MD
ILU(1)	500	236	†	**196**	164	333	†	**252**	494	145	†	**99**	155
	1000	†	†	†	†	†	†	†	†	†	†	†	†
ILUT(10^{-2}, 5)	500	49	**15**	26	†	89	**29**	53	†	44	**17**	28	†
	1000	†	**28**	35	†	†	**51**	71	†	†	**29**	34	†
ILUT(10^{-3}, 10)	500	28	**11**	22	23	63	**21**	42	49	30	**12**	22	26
	1000	†	**16**	27	34	†	**30**	67	74	†	**17**	27	37

grids (see [18]) and moreover, such approximations are only first order accurate and in many cases are unable to resolve fine features of the solution, such as boundary layers. In this case, as suggested in [19], a uniform coarse grid could be used to determine the region where the boundary layer is located (this corresponds to "wiggles" in an otherwise smooth solution). This would require solving linear systems such as those considered in the previous set of experiments. Subsequently, a local mesh refinement can be performed in the region containing the boundary layer. This solves the instability problem, and the approximation is still second order accurate (except at a few points on the interface between the coarse and the fine grids), but the resulting linear system, like the one corresponding to the uniform grid, is very challenging for iterative methods if convection is strong. Again, a simple reordering of the coefficient matrix can improve the situation dramatically. To illustrate this, we take the following example from Elman [16]. Consider the following partial differential equation in $\Omega = (0,1) \times (0,1)$

$$-\Delta u - 2P\frac{\partial u}{\partial x} + 2P\frac{\partial u}{\partial y} = g \qquad (2)$$

where the right-hand side g and the boundary conditions are determined by the solution

$$u(x,y) = \frac{e^{2P(1-x)} - 1}{e^{2P} - 1} + \frac{e^{2Py} - 1}{e^{2P} - 1}.$$

This function is nearly identically zero in Ω except for boundary layers of width $\mathcal{O}(\delta)$ near $x = 0$ and $y = 1$, where $\delta = 1/2P$. A uniform coarse grid was used in the region where the solution is smooth, and a uniform fine grid was superimposed to the regions containing the boundary layers, so as to produce a stable and accurate approximation; see [16] for details.

We performed experiments with $P = 500$ and $P = 1000$. These values are considerably larger than those used in [16]. The resulting matrices are of order 5041 and 7921, with 24921 and 39249 nonzeros, respectively. When ILU(0) preconditioning was used, no iterative solver converged in the maximum allowed number of iterations, independent of the reordering used. However, with Minimum Degree the three solvers appeared to be slowly converging, whereas with the other reorderings the iteration either diverged or stagnated. The results for ILU(1) and ILUT preconditioning and various orderings are reported in Table 2. We note that the natural ordering and the Cuthill-McKee reorderings produced the same or comparable fill-in for all preconditioners, whereas Multiple Minimum Degree resulted in higher fill-in with ILU(1) and considerably less fill-in with the ILUT preconditioners, with respect to the other reorderings.

From these results, we observe that Reverse Cuthill-McKee improves the quality of ILU(1) when $P = 500$. On the other hand, Cuthill-McKee behaves very poorly, and there is no advantage in using Multiple Minimum Degree. The results are very different with the ILUT preconditioners, for which Cuthill-McKee produces by far the best results. Reverse Cuthill-McKee is better than the natural ordering but not as good as Cuthill-McKee. Multiple Minimum Degree is bad with ILUT(10^{-2}, 5) but it performs well with ILUT(10^{-3}, 10), although it is not as effective as Cuthill-McKee. Notice that for $P = 1000$, all preconditioners fail when the natural ordering is used. For this particular example Cuthill-McKee dramatically improves the performance of ILUT preconditioners.

Further Analysis of the Results

The results of the experiments presented show that a simple reordering of the coefficient matrix can bring about a dramatic improvement in the quality of incomplete factorization preconditioners. In particular, we saw problems where all preconditioned iterative solvers failed with the natural ordering, and all converged rapidly after a symmetric permutation of the coefficient matrix. In this subsection, we investigate the reasons behind this.

An incomplete factorization preconditioner can fail or behave poorly for several reasons. A common cause of failure is instability of the incomplete factorization, which is caused by numerically zero pivots or exceedingly small ones. The result of this type of instability is that the incomplete factorization is very inaccurate, that is, the norm of the residual matrix $R = A - \bar{L}\bar{U}$ is large. This is a very real possibility for matrices that do not have some form of diagonal dominance, and for highly unstructured problems. Of course, an inaccurate factorization can also occur in the absence of small pivots, when many large fill-ins are dropped from the incomplete factors. Another kind of instability, which can take place whether or not small pivots occur, is severe ill-conditioning of the triangular factors, which reflects the instability of the long recurrences involved in the forward and backward solves when the preconditioning is applied [15], [28]. In this situation, $\|R\|_F$ need not be very large, but $\|I - A(\bar{L}\bar{U})^{-1}\|_F = \|R(\bar{L}\bar{U})^{-1}\|_F$ will be. Again, this is a common situation when the coefficient matrix is far from being diagonally dominant. Of course, both types of instabilities can simultaneously occur for a given problem: see [7] for an extensive experimental study of the causes of failure of incomplete factorizations. In [7], however, only scant attention is devoted to the effect of reorderings.

In order to gain some insight about the effect of reorderings, we computed the Frobenius norms of R and $R(\bar{L}\bar{U})^{-1}$ for each test matrix, reordering, and preconditioner. Those for the matrices arising from the discretization of problem (1) are reported in Table 3, where $N1 = \|A - \bar{L}\bar{U}\|_F$ and $N2 = \|I - A(\bar{L}\bar{U})^{-1}\|_F$.

We also monitored the size of the pivots in the course of the incomplete factorizations, and we did not find any very small pivots. Hence, failure or poor behavior of an incomplete factorization could be due either to the fact that significantly large fill-ins have been dropped, or to unstable triangular solves, or both.

The results in Table 3 give a clear explanation of the convergence behavior of iterative methods with ILU(0) preconditioning reported in Table 1. For the three equivalent permutations (natural ordering, Cuthill-McKee and Reverse Cuthill-McKee) the degradation and eventual failure in the convergence as ε^{-1} increases is not due to inaccuracy of the incomplete factorizations, but to instability of the triangular solves. As ε^{-1} increases, the condition number of the no-fill incomplete factors grows rapidly, the preconditioned matrix $A(\bar{L}\bar{U})^{-1}$ becomes more and more ill-conditioned, and the number of iterations increases. Furthermore, for large enough ε^{-1}, Elman [15] observed that the symmetric part of $A(\bar{L}\bar{U})^{-1}$ becomes indefinite, and this in turn can cause failure of the Krylov subspace accelerators. Inspection of the last column in Table 3 reveals that Minimum Degree has the effect of stabilizing the ILU(0) triangular factors. The preconditioner remains well-conditioned even for large values of ε^{-1}, and all three Krylov subspace methods converge. The fact that the number of iterations still increases with increasing ε^{-1} appears to be due to the fact that the ILU(0) factorization becomes less accurate. We observe that TPABLO (not shown in Table 3) also has the effect of damping the instability of the ILU(0) factors, but is somewhat less effective than MD. It is also worthwhile observing that for moderately nonsymmetric problems ($\varepsilon^{-1} \leq 400$) the number of iterations is almost directly related to $N1 = \|A - \bar{L}\bar{U}\|_F = \|R\|_F$, whereas for $\varepsilon^{-1} \geq 500$ this norm is not a good indicator of the effectiveness of the preconditioner, and $N2 = \|I - A(\bar{L}\bar{U})^{-1}\|_F$ becomes a more reliable indicator. Notice that MD always results in less accurate ILU(0) factors (larger $N1$) than the other reorderings.

The results in Table 3 show that ILU(1) preconditioning is more robust and effective than ILU(0) for this class of problems, and the Frobenius norms presented in Table 3 show that this is due to the fact that ILU(1) does not suffer from the kind of instability that plagues ILU(0). Because of this, $N1 = \|A - \bar{L}\bar{U}\|_F$ is now a very accurate indicator of the performance of the preconditioner for all values of ε^{-1}. The Cuthill-McKee reorderings outperform the natural ordering because they result in more accurate ILU(1) factors, and MD is inferior to the others because the incomplete factorization is less accurate.

Similar results apply to the ILUT preconditioners However, there are two phenomena that occur with ILUT$(10^{-2}, 5)$ and not with ILUT$(10^{-3}, 10)$, that deserve to be mentioned. For $\varepsilon^{-1} \geq 900$, the ILUT$(10^{-2}, 5)$ preconditioner with natural ordering fails rather dramatically. An inspection of the corresponding entries of Table 3 shows that this is due to the simultaneous occurrence of inaccuracy in the factorization (large $\|R\|_F$—too many large fill-ins have been dropped) *and* instability in the triangular factors (as revealed by a much larger value of $\|R(\bar{L}\bar{U})^{-1}\|_F$). This is an illustration of the fact that for strongly nonsymmetric problems, increasing fill-in in the incomplete factors does not necessarily result in an improved preconditioner, unless the factorization approaches a complete one; see also [7]. Note that the number of nonzeros in the factors is considerably higher for ILUT$(10^{-2}, 5)$ than for ILU(1). The other interesting phenomenon is that with ILUT$(10^{-2}, 5)$, the incomplete factors obtained with Minimum Degree are not only less accurate than those for the other reorderings, but also unstable when ε^{-1} becomes large. This is the opposite of what happens for ILU(0), and confirms that the relative performance of a given reordering is different for different preconditioning strategies, as already observed in [12]. When ILUT$(10^{-3}, 10)$ is used, the accuracy of the incomplete factorization approaches that of a direct solve, and none of the orderings suffers from instability (note that the original coefficient matrix A is fairly well-conditioned for all values of ε^{-1}). The norm of $R = A - \bar{L}\bar{U}$ becomes again a very reliable indicator of the performance of the preconditioners corresponding to the different permutations. The Cuthill-McKee reorderings give better results than the other orderings because they make the incomplete factorization more accurate. For fixed values of the ILUT parameters, the amount of fill-in in the incomplete factors is only slightly less with the Cuthill-McKee reorderings than with the natural ordering, whereas the number of nonzeros in the *complete* LU factors is much less. Hence, when compared to the natural ordering, the Cuthill-McKee reorderings allow one to compute a more accurate incomplete factorization for roughly the same arithmetic and storage costs for this class of problems.

The Frobenius norms $\|R\|_F$ and $\|R(\bar{L}\bar{U})^{-1}\|_F$ were computed also for the two matrices arising from problem (2). Our analysis of these can be found in [3]

4 CONCLUSIONS

We have provided evidence that reorderings originally designed for use with sparse direct solvers can significantly improve the performance of iterative methods preconditioned with incomplete LU factorizations. While not entirely new, an examination of the literature reveals that this fact is not widely known. Moreover, some of the conclusions of this study are somewhat at variance with those reached by other authors.

The benefit of reordering the coefficient matrix depends in part on how far the matrix is from being symmetric, as well as on the type of incomplete factorization preconditioner used. In our experiments we found that when the coefficient matrix is nearly symmetric, very little is gained from reordering it. This is consistent with the observations for symmetric matrices in [12]. On the other hand, if the matrix is strongly nonsymmetric, large reductions of the number of iterations can be obtained by (symmetrically) reordering the matrix.

A somewhat surprising result of our experiments is that the "natural," or "original" ordering of the test matrices used in this study is almost never the best from the point of view of incomplete factorization preconditioning and is usually the worst. It should be stressed that in most cases for which Reverse Cuthill-

Table 3: Norms of $A - \bar{L}\bar{U}$ and $I - A(\bar{L}\bar{U})^{-1}$ for different orderings, several preconditioners

		n/o		CM		RC		MD	
preconditioner	ε^{-1}	$N1$	$N2$	$N1$	$N2$	$N1$	$N2$	$N1$	$N2$
ILU(0)	100	1.46e-01	4.81e+01	1.49e-01	4.30e+01	1.46e-01	4.44e+01	9.29e-01	2.50e+01
	400	3.53e-01	2.68e+05	3.57e-01	2.73e+05	3.53e-01	2.74e+05	2.41e+00	9.97e+01
	600	3.86e-01	8.45e+05	3.92e-01	1.01e+07	3.86e-01	8.61e+05	3.48e+00	1.50e+02
	800	4.06e-01	3.26e+06	4.13e-01	4.86e+08	4.06e-01	2.79e+06	4.53e+00	2.00e+02
	1000	4.22e-01	2.69e+07	4.26e-01	6.25e+09	4.22e-01	2.26e+07	5.57e+00	2.50e+02
ILU(1)	100	4.60e-02	1.70e+00	4.90e-02	2.13e+00	4.59e-02	2.09e+00	4.21e-01	1.30e+01
	400	2.09e-01	1.44e+01	1.30e-01	8.12e+00	1.06e-01	9.42e+00	8.49e-01	3.23e+01
	600	2.43e-01	2.08e+01	1.69e-01	1.15e+01	1.18e-01	1.27e+01	1.17e+00	4.71e+01
	800	2.64e-01	2.63e+01	2.02e-01	1.49e+01	1.26e-01	1.56e+01	1.49e+00	6.21e+01
	1000	2.77e-01	3.14e+01	2.30e-01	1.85e+01	1.33e-01	1.83e+01	1.82e+00	7.71e+01
ILUT(10^{-2}, 5)	100	9.20e-03	5.65e-01	4.96e-03	2.06e-01	2.57e-03	1.07e-01	1.39e-01	4.64e+00
	400	7.75e-02	9.22e+00	1.91e-02	1.31e+00	1.29e-02	8.31e-01	7.09e-01	3.45e+01
	600	1.18e-01	1.45e+01	3.82e-02	2.78e+00	2.04e-02	1.62e+00	1.45e+00	2.58e+03
	800	1.78e-01	2.79e+01	6.21e-02	4.89e+00	2.83e-02	2.54e+00	7.39e+01	5.81e+06
	1000	6.47e+28	3.98e+43	9.82e-02	8.33e+00	3.54e-02	3.52e+00	2.28e+01	3.71e+06
ILUT(10^{-3}, 10)	100	1.48e-03	8.56e-02	5.09e-04	2.05e-02	3.77e-04	1.63e-02	4.76e-02	1.62e+00
	400	3.14e-02	2.49e+00	3.41e-03	2.36e-01	2.06e-03	1.39e-01	2.55e-01	9.00e+00
	600	6.16e-02	5.42e+00	1.02e-02	7.86e-01	4.26e-03	3.41e-01	4.46e-01	1.55e+01
	800	9.85e-02	1.08e+01	1.90e-02	1.55e+00	7.88e-03	6.98e-01	6.45e-01	2.29e+01
	1000	1.48e-01	1.98e+01	2.96e-02	2.54e+00	1.18e-02	1.13e+00	8.39e-01	3.08e+01

McKee was not best, it still gave good results (that is, it was not found to be much worse than the best ordering). Hence, overall, Reverse Cuthill-McKee appears to be superior to the other orderings in the context of incomplete factorization preconditioning. As revealed by a direct inspection of the residual matrices $A-\bar{L}\bar{U}$, in most cases this was simply due to the fact that this reordering produced more accurate incomplete factorizations than those obtained with the natural ordering (with a comparable amount of fill-in in the factors). In some cases the improvement was due to the fact that the reordering resulted in a stabilization of the incomplete triangular factors. However, none of the orderings considered in this paper were found to be completely immune from potential instabilities in the corresponding triangular solves. In most cases where instabilities occurred, the problem disappeared by allowing more fill-in in the incomplete factors, but not always. Indeed, there were a few cases where increasing the amount of fill-in made things worse, in the sense that the instability of the factors increased; see also [7].

In general, Cuthill-McKee cannot be recommended. While it performed very well for some problems, its behavior is rather erratic. Reverse Cuthill-McKee should be preferred. Of course, Reverse Cuthill-McKee is not trouble-free. The quality of the corresponding preconditioner will be affected, in general, by the choice of the initial node, and by the ordering of nodes within level sets. In particular, it is not clear how different tie-breaking strategies will affect the incomplete factorization. It is also possible that some orderings within level sets will produce a poor pivot sequence and a breakdown of the incomplete factorization process.

Nevertheless, based on the results of our experiments, we conclude that much can be gained from reordering strongly nonsymmetric matrices before performing an incomplete factorization, and not much should be lost, particularly when the Reverse Cuthill-McKee reordering is used. For convection-diffusion problems on rectangular grids, ILU(1) or ILUT preconditioning combined with Reverse Cuthill-McKee is recommended, whereas the lexicographic ordering behaves rather poorly and should be avoided. For matrices that do not have a “natural” ordering, such as those arising from unstructured meshes, we recommend Reverse Cuthill-McKee as the original ordering.

Acknowledgements. The first author was supported in part by DOE through grant W-7405-ENG-36 with Los Alamos National Lab. The second author was supported by NSF grant DMS-9625865.

REFERENCES

[1] R. Barrett, M. Berry, T.F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H.A. van der Vorst. *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods.* SIAM, Philadelphia, 1994.

[2] M. Benzi, H. Choi, and D.B. Szyld. Threshold ordering for preconditioning nonsymmetric problems. In G. Golub et al., editors, *Sci. Comput., Proc. Workshop, 10–12 March 1997, Hong Kong*, pages 159–165, Springer Verlag

[3] M. Benzi, D.B. Szyld, and A.C.N. van Duin. Orderings for Incomplete Factorization Preconditionings of Nonsymmetric Problems. Research Report 97-91, Department of Mathematics, Temple University, September 1997

[4] T.F. Chan, C.-C. Jay Kuo, and C. Tong. Parallel elliptic preconditioners: Fourier analysis and performance on the Connection Machine. *Comput. Phys. Comm.*, 53:237–252, 1989.

[5] M.P. Chernesky. On preconditioned Krylov subspace methods for discrete convection-diffusion problems. *Numer. Methods Partial Differential Equations*, 13:321–330, 1997.

[6] H. Choi and D.B. Szyld. Application of threshold partitioning of sparse matrices to Markov chains. In *Proc. IEEE Internat. Computer Performance and Dependability Symposium, IPDS'96, Urbana-Champaign, Illinois, September 4–6, 1996*, pages 158–165, IEEE Computer Society Press, Los Alamitos, California, 1996.

[7] E. Chow and Y. Saad. Experimental study of ILU preconditioners for indefinite matrices. To appear in *J. Comput. Appl. Math.*.

[8] E.F. D'Azevedo, P.A. Forsyth, and W. Tang. Ordering methods for preconditioned conjugate gradient methods applied to unstructured grid problems. *SIAM J. Matrix Anal. Appl.*, 13:944–961, 1992.

[9] M.A. DeLong and J.M. Ortega. SOR as a preconditioner. *Appl. Numer. Math.*, 18:431–440, 1995.

[10] S. Doi. On parallelism and convergence of incomplete LU factorizations. *Appl. Numer. Math.*, 7:417–436, 1991.

[11] S. Doi and A. Lichnewsky. A graph-theory approach for analyzing the effects of ordering on ILU preconditioning. INRIA, Rapport de Recherche No. 1452, Domain de Voluceau Rocquencourt, France, June 1991.

[12] I.S. Duff and G.A. Meurant. The effect of ordering on preconditioned conjugate gradients. *BIT*, 29:635–657, 1989.

[13] L.C. Dutto. The effect of ordering on preconditioned GMRES algorithm, for solving the compressible Navier-Stokes equations. *Internat. J. Numer. Methods Engrg*, 36:457–497, 1993.

[14] V. Eijkhout. Analysis of parallel incomplete point factorizations. *Linear Algebra Appl.*, 154-156:723–740, 1991.
[15] H.C. Elman. A stability analysis of incomplete LU factorizations. *Math. Comput.*, 47:191–217, 1986.
[16] H.C. Elman. Relaxed and stabilized incomplete factorizations for non-self-adjoint linear systems. *BIT*, 29:890–915, 1989.
[17] A. George and J.W. Liu. *Computer Solution of Large Sparse Positive Definite Systems.* Prentice-Hall, Englewood Cliffs, New Jersey, 1981.
[18] G. Golub, D. Silvester and A. Wathen. Diagonal dominance and positive definiteness of upwind approximations for advection diffusion problems. In D.F. Griffiths and G.A. Watson, eds, *Numerical Analysis. A.R. Mitchell 75th Birthday Volume*, 125–131, World Scientific Publishing Co., River Edge, NJ, 1996.
[19] P.M. Gresho and R.L. Lee. Don't suppress the wiggles—they're telling you something! *Comput. & Fluids*, 9:223–253, 1981.
[20] H.P. Langtangen. Conjugate gradient methods and ILU preconditioning of non-symmetric matrix systems with arbitrary sparsity patterns. *Internat. J. Numer. Methods Fluids*, 9:213–233, 1989.
[21] Y. Notay. Ordering methods for approximate factorization preconditioning. Report IT/IF14-11, Université Libre de Bruxelles, Belgium, January 1993.
[22] J.M. Ortega. Orderings for conjugate gradient preconditioners. *SIAM J. Optim.*, 1:565–582, 1991.
[23] E.L. Poole and J.M. Ortega. Multicolor ICCG methods for vector computers. *SIAM J. Numer. Anal.*, 24:1394–1418, 1987.
[24] Y. Saad. Preconditioning techniques for nonsymmetric and indefinite linear systems. *J. Comput. Appl. Math.*, 24:89–105, 1988.
[25] Y. Saad. ILUT: A dual threshold incomplete LU factorization. *Numer. Linear Algebra Appl.*, 1:387–402, 1994.
[26] H.D. Simon. Incomplete LU preconditioners for conjugate-gradient-type iterative methods. Paper SPE 13533. In *Proc. 1985 SPE Reservoir Simulation Symposium*, Richardson, Texas, 1985, pages 387–396, Society of Petroleum Engineers.
[27] S.A. Stotland and J.M. Ortega. Orderings for parallel conjugate gradient preconditioners. *SIAM J. Sci. Comput.*, 18:854–868, 1997.
[28] H.A. van der Vorst. Iterative solution methods for certain sparse linear systems with a non-symmetric matrix arising from PDE-problems. *J. Comput. Phys.*, 44:1–19, 1981.
[29] D.M. Young. *Iterative Solution of Large Linear Systems.* Academic Press, New York, 1971.