On the existence problem of incomplete factorisation methods, Lapack Working Note 144, UT-CS-99-435

Victor Eijkhout*

1 December 1999

Abstract

Incomplete factorisation methods can suffer from breakdown in that they may give zero or negative pivots where an exact factorisation would show only positive pivots. This breakdown effectively prevents the factorisation from being used in iterative methods such as Conjugate Gradients. We give an overview of strategies that have been proposed to prevent this breakdown, and we touch briefly on various related issues in incomplete factorisations.

1 Introduction

For the efficient solution of sparse linear systems $Au = b$ by an iterative method, the choice of a proper preconditioner is crucial. A preconditioner is a matrix M that approximates A, but for which solving the system $Mu = b$ is computationally cheap. In addition, M itself should be easily constructable.

Since the original coefficient matrix A is sparse, people have sought to construct sparse factorisations $M = LU \approx A$. The exact LU factorisation of A is not sparse, so M is constructed by a so-called incomplete factorisation, where the update

$$a_{ij} \leftarrow a_{ij} - a_{ik}{a_{kk}}^{-1}a_{kj} \tag{1}$$

is executed subject to some decision process.

Ideally, the only question concerning incomplete factorisations would be their accuracy. For instance, typically the condition number $\kappa(A) \sim h^{-2}$ where h is the mesh width, and one would hope that $\kappa(M^{-1}A)$ is smaller, preferably of

*This work was supported in part or in whole by Center for Parallel Computation, subcontract #292-3-54397, and Lawrence Livermore National Laboratory, subcontract #B503913, and Los Alamos National Laboratory, subcontract #D0252-0019-2G, and Los Alamos National Laboratory Copmputer Science Institute (LASCI) through LANL contract number 03891-99-23, as part of the prime contract (W-7405-ENG-36) between the Department of Energy and the Regents of the University of California.

a lower order than h^{-2}. In practice, however, it is already hard to guarantee the existence of the factorisation. To begin with, if $i = j$ in equation (1), the updated value of a_{ii} can be zero leading to breakdown in the i-th elimination step. A negative value of a_{ii} is a problem too, since, if A is positive definite, we want M to be so, and consequently all pivots of the incomplete factorisation to be positive.

Incomplete factorisation methods are well-defined for M-matrices, but for any other type of matrix, even symmetric positive definite ones, they can suffer from breakdown of some form or other. A number of remedies have been proposed, but all suffer from certain disadvantages. We will give in sections 2 and 3 an overview of various factorisation algorithms, and related issues in the theory of incomplete factorisations.

We do not give any numerical tests in this report since it is mostly concerned with the theoretical aspects of the methods. For comparative tests of the various methods we refer the reader to [18].

2 Overview of earlier factorisation algorithms

In this section we will give an overview of several representative methods for computation of a factorisation[1]

$$M = (D_M + L_M){D_M}^{-1}(D_M + U_M) \tag{2}$$

of a matrix

$$A = D_A + L_A + U_A. \tag{3}$$

The choice $D_M = D_A$, $L_M = L_A$, $U_M = U_A$ gives the SSOR method (section 2.2); $D_M \neq D_A$, $L_M = L_A$, $U_M = U_A$ is called ILU-D (section 2.3); the general case $D_M \neq D_A$, $L_M \neq L_A$, $U_M \neq U_A$ describes all other incomplete LU methods.

2.1 Classifying factorisation methods

There are several criteria with which to distinguish incomplete factorisation methods. We will give a few in this subsection; they are not orthogonal, nor do they apply to all of the methods presented below.

2.1.1 Algorithmic criteria

As a first criterium we consider the dropping strategy for fill-in elements. This strategy can be one of the following:

positional There is a set $S \subset \mathbf{N}^2$ of matrix positions in which we don't allow fill-in, no matter its numerical value. This set can be determined prior to starting the algorithm, or it can be constructed adaptively. A common

[1] There is more than one mathematically equivalent way to write a factorisation; we choose the form (2) to bring out the symmetry, even though it is not computationally optimal.

choice is for S to comprise all zero positions of A. This method has the advantage that storage requirements of the incomplete factorisation are predictable.

numerical We apply some test to fill-in values, and anything falling under the threshold is dropped; elements of sufficient magnitude are accepted. This method has indeterminate storage requirements (see [4] for a discussion on how to limit the storage by modifying the threshold parameter), but it is likely to be more accurate than the positional dropping strategy.

2.1.2 Pivot repair

A further algorithmic issue to consider is how the method deals with zero or negative pivots, should they occur.

The existence question of incomplete LU factorisations was fully solved in [34] for the case of M-matrices[2]. This paper and subsequent generalisations such as [1] established that for M-matrices fill-in can be totally or partially ignored, while the M-matrix property is preserved for the remaining submatrix. As a result, all pivots are guaranteed to be positive and no repair strategy is needed.

For other matrices than M-matrices, even for symmetric positive definite ones, an incomplete factorisation can fail in the sense that pivots can become zero or negative. Kershaw [30] gave the following example:

$$\begin{pmatrix} 3 & -2 & 0 & 2 \\ -2 & 3 & -2 & 0 \\ 0 & -2 & 3 & -2 \\ 2 & 0 & -2 & 3 \end{pmatrix}$$

which is positive definite, but will have a negative fourth pivot in an ILU(0) factorisation.

There are the two choices to preventing breakdown: one can adopt an ad-hoc strategy to copy with non-positive pivots when they arise, or the factorisation method can be designed in such a way that no breakdown will even occur. We will see examples of both strategies.

2.1.3 Order reduction: modified methods

Next, incomplete factorisations can be distinguished by the fact that they try to preserve spectral properties of the coefficient matrix in the preconditioner matrix. While the original matrix satisfies $\kappa(A) = O(h^{-2})$, simply dropping elements will lead to $\kappa(M) = O(1)$, and consequently $\kappa(M^{-1}A)$ will be of the same order as $\kappa(A)$. However, the constant of proportionality may be substantially lowered.

[2] There are several equivalent definitions of an M-matrix. For our purpose, the most convenient one is that A is an M-matrix if it is positive definite and has nonpositive off-diagonal elements. It follows that an M-matrix has positive diagonal elements.

In order to estimate the condition of the preconditioned system, one can use the estimate based on the error matrix $R = M - A$ (see [6])

$$\kappa(M^{-1}A) \leq \omega^*[1 + \|A^{-1}\|\|R\|]$$

where A and C are symmetric positive definite, and

$$\omega^* = \min\{\omega : (2 - \omega^{-1})D_M - D_A \text{ is SPD}\};$$

(see equations (3) and (2) for the definition of D_A and D_M.) The value of ω^* is typically $O(1)$: for the ILU factorisation of the central difference Laplace problem it is $\sqrt{2}$.

The so-called *modified incomplete factorisation* methods aim at lowering the order of the precoditioned system. In their simplest form they replace the conditional execution of (1) by

$$\text{if fill is ignored in } (i,j) \text{ position,} \quad a_{ii} \leftarrow a_{ii} - a_{ik}{a_{kk}}^{-1}a_{kj} \tag{4}$$

This is commonly referred to as 'moving fill to the diagonal', and it can lead to a condition number of lower order; see section 2.6.2 for more details. Another interpretation of modified methods is that they force $Av = Mv$ for some positive vector v.

2.1.4 Ordering

A final criterium to distinguish incomplete factorisation algorithms is the ordering strategy of the unknowns.

Symmetric permutations PAP^t of the coefficient matrix could be used to put off, or perhaps avoid altogether, problems with zero or negative pivots. In [47] the authors explore the idea that an incomplete factorisation exists if the ordering is such that the factorisation is exact. From this they derive a sufficient condition on the sparsity patterns of the matrix and the preconditioner.

Incomplete factorisation methods are sensitive to (symmetric) permutations of the coefficient matrix. Thus, orderings that increase the parallelism of the preconditioner solve may incur a larger number of iterations than the same factorisation algorithm applied to the matrix under the natural ordering. This was observed in [16] and analysed in [15, 14, 19].

In [13] the ordering was chosen, while keeping the maximum fill level fixed, in such a way as to minimise the size of fill elements.

In the following subsections, we will give an overview of various incomplete and modified incomplete factorisation algorithms, remarking on their existence properties and practical behaviour. We will ignore the issue of the influence of the ordering on the existence of the factorisation, and concentrate solely on the modification strategy.

2.2 SSOR

The Symmetric Successive Over-Relaxation preconditioner is defined as the product $(D_A + L_A){D_A}^{-1}(D_A + U_A)$, where D_A, L_A, and U_A are the diagonal and strict upper and lower triangular parts of the coefficient matrix A. Clearly, constructing this factorisation carries zero cost. At most, one wants to compute and store ${D_A}^{-1}$ explicitly. This factorisation also has the pleasant property that the question of well-definedness of the factorisation is trivially satisfied[3]. Therefore, some people advocate using variants of $SSOR$ as preconditioner for complicated problems [32].

On the downside, this preconditioner will be less effective than incomplete factorisation methods when the latter exist. Introduction of a relaxation parameter can alleviate this [2, Ch. 1], but the calculation of the optimal relaxation parameter is usually nontrivial.

2.3 Types of ILU factorisations: ILU-D, ILU(0), ILU(r), ILU(k)

There are various types of ILU methods. In the simplest type, not only is all fill-in ignored, but only diagonal nonzero elements of the matrix are ever altered:

$$\text{if } i = j\text{:} \qquad a_{ij} \leftarrow a_{ij} - a_{ik}{a_{kk}}^{-1}a_{kj} \tag{5}$$

This method has the important practical property that the preconditioner largely re-uses the matrix elements. Practical implications are that only one vector's worth of elements needs to be stored, and that an efficient implementation of the solution step is possible [22], which combines it with the matrix product, giving a substantial reduction of the operation count. The name ILU-D for this factorisation was coined in [40].

Slightly more elaborately, the $ILU(0)$ factorisation ignores all fill-in, but allows modification of off-diagonal nonzeros:

$$\text{if } a_{ij} \neq 0\text{:} \qquad a_{ij} \leftarrow a_{ij} - a_{ik}{a_{kk}}^{-1}a_{kj} \tag{6}$$

The storage needed for this factorisation equals that of the original coefficient matrix.

The ILU-D and ILU(0) methods are identical for certain matrices, one common case being that of the central difference operator on a regular grid. More general, they are identical if there are no triangles in the matrix graph.

In the general case, fill-in can be accepted or discarded in any position. If a threshold parameter r is used, for instance as in

$$\text{if } |a_{ij}| \geq r a_{ii}\text{:} \qquad a_{ij} \leftarrow a_{ij} - a_{ik}{a_{kk}}^{-1}a_{kj}, \tag{7}$$

we call this a 'numerical dropping strategy', and term the method ILU(r).

If the notion of level is applied to fill-in, the method is called ILU(k):

$$\text{initially:} \qquad \forall_{i,j}\colon \ell_{ij} = 0; \tag{8}$$

[3] We will call a factorisation well-defined if, when applied to a (symmetric) positive definite matrix, it yields a (symmetric) positive definite preconditioner. We will not consider the indefinite case in this paper.

$$\begin{array}{rl} \text{if } \ell_{ip} \le k \text{ and } \ell_{pj} \le k: & a_{ij} \leftarrow a_{ij} - a_{ip}{a_{pp}}^{-1}a_{pj} \\ \text{if fill-in and } \ell_{ij} = 0: & \ell_{ij} \leftarrow 1 + \min\{\ell_{ip}, \ell_{pj}\} \\ \text{otherwise:} & \text{ignore fill} \end{array} \tag{9}$$

2.4 Kershaw's method

The breakdown of the incomplete factorisation method can be remedied, as was suggested by Kershaw [30], by substituting an arbitrary positive value for zero or negative pivots. While this trivially guarantees the existence of the factorisation, it is likely to lead to a large condition number for $\kappa(M^{-1}A)$.

Choosing the repair value is not trivial: too small a value will give unstable recurrences during the solution process [23, 24], so, even while the preconditioner is SPD, the iterative method may diverge. Overestimates of the optimal repair value will make the preconditioner too diagonally dominant, in effect turning it into a Jacobi method.

Kershaw's choice for the repair pivot is $m_{kk} = \sum_{j<k} |m_{kj}| + \sum_{j>k} |m_{jk}|$. A similar choice $m_{kk} = \sum_{j<k} |m_{kj}| + \sum_{j>k} |m_{kj}|$. was used by van der Vorst [43] motivated by considerations of stability; see section 3.5.

2.5 Manteuffel's method

Since the main problem with incomplete factorisation is pivots becoming negative, it makes sense that adding a sufficient number to the diagonal of the matrix will give a well-defined factorisation.

Trivially, adding enough to make the matrix diagonally dominant is a sufficient condition, but this is likely to lead to an ill-conditioned system. Manteuffel [33] proposed to make several attempts at finding a small enough value of α such that $A + \alpha I$ has a well-defined factorisation M, and which does not give a too large condition number $\kappa(M^{-1}A)$.

2.6 Modified Incomplete LU

The idea of moving fill-in to the diagonal (equation 4) has been around in various forms for a long time. Already in [17] it was shown that this, when combined with small perturbations (see below), could lower the condition number of the preconditioned system to $O(h^{-1})$.

2.6.1 Existence of MILU factorisations

All the conditioning theory for modified methods holds only for symmetric matrices, and the methods are not guaranteed to be well-defined for other matrices than M-matrices. Example: the following is a symmetric positive definite matrix for which $MILU$ is not spd [4] for some sparsity patterns:

$$A = \begin{pmatrix} 1 & -1 & \epsilon & -\epsilon \\ -1 & (1+\epsilon/2) & 0 & 0 \\ \epsilon & 0 & 1 & 0 \\ -\epsilon & 0 & 0 & 1 \end{pmatrix}$$

This matrix is positive definite for $0 < \epsilon < -1+\sqrt{2}$. After one elimination step, the Schur complement is

$$A_2 = \begin{pmatrix} \epsilon/2 & \epsilon & -\epsilon \\ \epsilon & 1-\epsilon^2 & \epsilon^2 \\ -\epsilon & \epsilon^2 & 1-\epsilon^2 \end{pmatrix}.$$

Clearly, simply moving the $(1,3)$ position of the Schur complement to the diagonal leads to breakdown of the algorithm.

In fact, even for M-matrices the existence theory is based on finding a vector $v > 0$ such that $Mv = Av > 0$. Moving fill-in to the diagonal corresponds to letting v be the vector $e = (1, 1, \ldots)^t$, and typically we only have $Ae \geq 0$. In a finite element or finite difference context this means that the rows corresponding to interior nodes of the domain have zero rowsums. In such rows the factorisation can break down. Thus, the MILU factorisation can break down for symmetric M-matrices; see [20, 36]. The global idea is that modified methods preserve row sums. Therefore, if a row with zero rowsum is an 'endpoint' of the factorisation (there are no nodes with a higher number connected to it), it will have a zero pivot.

As an example of this, the matrix

$$A = \begin{pmatrix} 2 & -1 & -1 \\ -1 & 1 & 0 \\ -1 & 0 & 1+\epsilon \end{pmatrix}$$

is an M-matrix for any $\epsilon > 0$. However, a modified incomplete factorisation will give a zero pivot after the first elimination step. (Using $v = e$ in the definition of the modified factorisation, we find from the Gershgorin circle theorems that in the incomplete factorisation pivots can not become negative.)

As a heuristic statement, one can say that under the natural ordering MILU methods do not break down on M-matrices, provided Neuman boundaries of the domain are not ordered last. When no breakdown occurs, the modification entails adding a negative semi-definite matrix to a principal 2×2 minor of the coefficient matrix. Since this matrix, call it F, is chosen such that $Fv = 0$, with v the positive vector for which $Av > 0$, this leaves the minor, and consequently the factorisation, an M-matrix and therefore positive definite.

2.6.2 Accuracy of MILU factorisations

Often the fill is multiplied by a parameter less than 1 before being moved to the diagonal. This is refered to as a 'relaxed modified incomplete factorisation';

see [7, 11, 20, 44]. Also, some methods perturb the diagonal by adding elements of order h^2 to the diagonal; see [2, 8]. Such (perturbed) modified incomplete factorisation algorithms can be proved to give $\kappa(M^{-1}A) = O(h^{-1})$; see [26, 36, 10]. (There is of course an equivalence of sorts between the relaxed and perturbed modified methods.)

An intuitive way of explaining the order reduction is to observe that the rest matrix becomes a zero-rowsum difference matrix, giving

$$(Ru)_i = O(h), \quad \forall_i$$

for vectors that are $C_0^1(\Omega)$. Gustafsson [26] showed the relation of this condition to a sufficient condition for order reduction.

Most studies of modified incomplete factorisations implicitly assume a natural ordering of the unknowns. Discussions of the ordering issue in this context can be found in, e.g., [9, 28].

2.7 Gustafsson's method of modified elements

In effect the problem with incomplete factorisations of symmetric non-M-matrices lies in the positive off-diagonal elements. Therefore, Gustafsson [27] proposed to eliminate these elements by moving them to the diagonal prior to the factorisation. (This method was also explored in [3].)

Since this can be considered a modification on the element matrix level, specifically by adding positive semi-definite matrices, we find that the so modified matrix A_m has a relative condition to the original matrix independent of the matrix size:

$$u^t A u \leq u^t A_m u \leq \alpha u^t A u$$

where $\alpha > 1$ is the maximum relative condition over all element matrices. The crucial fact here is that α is independent of the matrix size.

A modified incomplete factorisation M of A_m will then give

$$\kappa(M^{-1}A) \leq \alpha\kappa(M^{-1}A_m) = O(h^{-1}).$$

The philosophical problem with this method is that it can wind up preconditioning the wrong operator. For instance, the biharmonic operator has a stencil

$$\begin{array}{ccccc} & & 1 & & \\ & 2 & -8 & 2 & \\ 1 & -8 & 20 & -8 & 1 \\ & 2 & -8 & 2 & \\ & & 1 & & \end{array}$$

which after elimination of the positive coefficients becomes

$$\begin{array}{ccc} & -8 & \\ -8 & 32 & -8\,, \\ & -8 & \end{array}$$

that is, a multiple of the Laplacian, a different operator altogether.

Although Gustafsson formulated this method on the element matrix level, it can be applied algebraically to a fully formed matrix. The only difficulty arises along the boundary of the domain, where certain element matrix coefficients have gone into forming the right hand side, and hence can not be retrieved from the matrix. We then have to apply a heuristic, for instance forcing the corresponding rows to have zero rowsums. As an example, performing algebraic removal on the above biharmonic stencil, the points at distance $2h$ from the boundary will have rowsums -1.

2.8 Jennings and Malik's partial elimination

Moving fill-in to the diagonal in modified incomplete factorisations of M-matrices corresponds to adding negative semi-definite matrices to the original matrix. In the M-matrix case this does not make the matrix indefinite, but in general this is not true, potentially leading to indefinite matrices and negative pivots.

Jennings and Malik [29] proposed to add the absolute size of the fill elements to the diagonal. This corresponds to adding *positive* definite or semi-definite matrices. As this only increases the diagonal, it is easy to see that the factorisation never becomes undefined. We note that in the case where a_{ij} and a_{ji} are of opposite sign this strategy is theoretically defensible: for such 'sign anti-symmetric' matrices it reduces to MILU. In general however, it guarantees nothing beyond the mere existence of the factorisation; in particular, for the model case of M-matrices it does not reduce to any other method.

This method was also proposed by Robert [41].

2.9 Eijkhout's 'weighted modification' method

Eijkhout's method aims to conserve the following matrix properties when making an incomplete factorisation:

1. Any symmetry of the original matrix is to be preserved. Note that this is not trivial in the case of threshold dropping, as can be easily seen from equation (7).

2. The factorisation should be well-defined in the sense that, given a matrix with positive diagonal, all pivots should be positive. Note that this is a stronger condition than the traditional one that pivots should be positive for matrices that are positive definite. This makes the method potentially suitable for certain indefinite systems.

3. The spectrum is not to be disturbed too much, in particular, the dropping strategy should reduce to some variant of MILU for M-matrices.

4. Applying the factorisation to an M-matrix should yield an M-matrix.

These aims are accomplished by eliminating the off-diagonal fill elements in (i, j) and (j, i) elements together by added weighted combinations of them to the (i, i) and (j, j) diagonal elements:

$$a_{ii} \leftarrow a_{ii} - a_{ij}\sqrt{a_{ii}/a_{jj}}/2, \qquad a_{jj} \leftarrow a_{jj} - a_{ji}\sqrt{a_{jj}/a_{ii}}/2.$$

This basic idea is augmented by some heuristics and dynamic conditions to skip the modification in certain cases.

3 Other topics in incomplete factorisations

3.1 H-matrices

Most M-matrix theory can easily be generalised to H-matrices. (A matrix is an H-matrix if its 'comparison matrix' is an M-matrix; basic theory of H-matrices can be found in Ostrowski [37] and Neumaier [35].) For the scalar case this generalization was done by Varga et al [45], and for the block case by Polman [39].

3.2 Ordering of unknowns

Symmetric permutations PAP^t of the coefficient matrix could be used to put off, or perhaps avoid altogether, problems with zero or negative pivots. In [47] the authors explore the idea that an incomplete factorisation exists if the ordering is such that the factorisation is exact. From this they derive a sufficient condition on the sparsity patterns of the matrix and the preconditioner.

However, in incomplete factorisations orderings are usually chosen for other reasons. As we indicated in sections 2.6.1 and 2.6.2, certain orderings would not lead to a well-defined factorisation, or would not give the desired accuracy.

Considerations of parallelism or vectorisability may also dictate the ordering of the unknowns, at least to an extent. Duff and Meurant [16] reported tests on the influence of ordering strategies on the convergence speed of the conjugate gradients method. These results were explained in the model case by Doi [15] and more general by Eijkhout [19]. In [13] the ordering was chosen, while keeping the maximum fill level fixed, in such a way as to minimise the size of the fill elements.

3.3 Preconditioning from equivalent operators

Often, the existence problems of incomplete factorisation methods arise from the fact that the discretisation uses higher-order finite difference or finite element schemes, or from the asymmetry of the problem. One could obviate these problems by basing the factorisation on a matrix or an operator related to the original one. Given a matrix A, a related matrix A_0, and a preconditioner, accuracy estimates then follow from equations such as

$$\kappa(M^{-1}A) \leq \kappa(M^{-1}A_0) \cdot \kappa({A_0}^{-1}A),$$

where often analytic results hold for $\kappa({A_0}^{-1}A)$.

For instance, one could base the preconditioner on the symmetric part $(A + A^t)/2$ of the coefficient matrix. This idea was explored in, e.g., [25, 46].

If the actual PDE operator is known, the preconditioner M can be based on a lower order finite element discretisation A_0. In this case we have $\kappa({A_0}^{-1}A) = O(1)$ [2]. The advantage of this strategy lies in that the lower order discretisation may very well given an M-matrix, so the existence problem of the preconditioner is solved. For an example of this approach, see[21].

3.4 Off-diagonal modification

Axelsson and Munksgaard [4] made, like Jennings and Malik (section 2.8), the observation that modification decreases the diagonal elements, thereby possibly making the factorisation indefinite, was also made. Unlike Jennings and Malik, they proposed not altering the fill element to be moved, but rather propose to add it to any positive off-diagonal element. This method has no theoretical guarantees, and the authors therefore suggest that adding a positive number to the diagonal (as happens in the methods of Kershaw, section 2.4, and Manteuffel, section 2.5) can be used as a final remedy.

3.5 Stability of the preconditioner solve

If the preconditioner is well-defined and gives a small enough rest matrix, there is still a further complication that may lead to slowly or not converging iterative methods. By considering the solution of the triangular systems as a recurrence relation, we see that build-up of round-off error is possible if the roots of the characteristic polynomial are greater than 1 in absolute value.

For example, solving $Lu = v$ for a matrix L with constant diagonals $L = (\ell_n, 0, \ldots, 0, \ell_1, d)$ corresponds to the recurrence $du_i + \ell_1 u_{i-1} + \ell_n u_{i-n} = v_i$, with a characteristic polynomial $dx^n + \ell_1 x^{n-1} + \ell_n$. The stability of this was analysed by Elman [23, 24], and found to be equivalent to the factors being diagonally dominant.

For a short proof, consider the recurrence

$$a_0 x_i + \sum_{j=1}^{n} a_j x_{i-j} = f_i$$

with characteristic solutions $x_i = \lambda^i$ where λ is a solution of

$$a_0 \lambda^i + \sum_{j=1}^{n} a_j \lambda^{i-j} = 0.$$

Now suppose that the matrix is diagonally non-strictly dominant, that is, $a_0 \geq \sum |a_j|$. The assumption $|\lambda| > 1$ gives by

$$1 = \left| \sum \frac{a_j}{a_0} \lambda^{-j} \right| \leq \sum \left| \frac{a_j}{a_0} \lambda^{-j} \right| < \sum \left| \frac{a_j}{a_0} \right|$$

a contradiction, therefore the characteristic roots satisfy $|\lambda| \leq 1$.

Elman's stability condition was also suggested by van der Vorst [43]. It should be noted that this analysis mostly pertains to problems with slowly varying or constant diagonals.

3.6 Block factorisations

In this report we will not consider block factorisation methods. These treat the matrix in terms of larger subblocks than the scalar entries, for instance deriving blocks from lines of grid points in the physical domain of the PDE. The scalar factorisation problem often appears in these methods, since they may require an approximation to the inverse of the pivot blocks. The interested reader is referred to [1, 6, 5, 10, 12, 31, 38, 42].

4 Conclusion

While ideally an incomplete factorisation should be judged solely on its accuracy properties, in practice there is first the hurdle of guaranteeing its existence. Various strategies have been proposed to ascertain this Some of them, such as the methods of Gustafsson and of Kershaw, become effectively no-ops in the cases where the classical algorithms are well-defined. Others, such as the methods of Eijkhout and of Jennings and Malik, are always applied. In either case, the repair strategy can have more or less theoretical justification, with consequent implications for the approximation accuracy of the preconditioner.

In this paper we have given an overview of the existing methods and several attendant issues. Since no set of tests can ever be comprehensive we have not included any numerical results in this survey; the reader is referred to [18] for numerical results on the methods discussed here.

References

[1] O. Axelsson. A general incomplete block-matrix factorization method. *Lin. Alg. Appl.*, 74:179–190, 1986.

[2] O. Axelsson and A.V. Barker. *Finite element solution of boundary value problems. Theory and computation.* Academic Press, Orlando, Fl., 1984.

[3] O. Axelsson and L. Yu. Kolotilina. Diagonally compensated reduction and related preconditioning methods. *Num. Lin. Alg. Appl.*, 1, 1995.

[4] O. Axelsson and N. Munksgaard. Analysis of incomplete factorizations with fixed storage allocation. In D. Evans, editor, *Preconditioning Methods – Theory and Applications*, pages 265–293. Gordon and Breach, New York, 1983.

[5] O. Axelsson and B. Polman. On approximate factorization methods for block-matrices suitable for vector and parallel processors. *Lin. Alg. Appl.*, 77:3–26, 1986.

[6] Owe Axelsson and Victor Eijkhout. Vectorizable preconditioners for elliptic difference equations in three space dimensions. *J. Comp. Appl. Math.*, 27:299–321, 1989.

[7] Owe Axelsson and Gunhild Lindskog. On the rate of convergence of the preconditioned conjugate gradient method. *Numer. Math.*, 48:499–523, 1986.

[8] R. Beauwens. On Axelsson's perturbations. *Lin. Alg. Appl.*, 68:221–242, 1985.

[9] R. Beauwens. Approximate factorizations with S/P consistently ordered M-factors. *BIT*, 29:658–681, 1989.

[10] Robert Beauwens and Mustafa Ben Bouzid. Existence and conditioning properties of sparse approximate block factorizations. *SIAM Numer. Anal.*, 25:941–956, 1988.

[11] Tony F. Chan. Fourier analysis of relaxed incomplete factorization preconditioners. *SIAM J. Sci. Stat. Comput.*, 12:668–680, 1991.

[12] P. Concus, G.H. Golub, and G. Meurant. Block preconditioning for the conjugate gradient method. *SIAM J. Sci. Stat. Comput.*, 6:220–252, 1985.

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

[14] S. Doi and A. Lichnewsky. An ordering theory of incomplete LU factorizations on regular grids. Technical report, INRIA, 1990.

[15] Shun Doi. On parallelism and convergence of incomplete LU factorizations. *App. Num. Math.*, 7:417–436, 1991.

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

[17] T. Dupont, R. Kendall, and H. Rachford. An approximate factorization procedure for solving self-adjoint elliptic difference equations. *SIAM J. Numer. Anal.*, 5:559–573, 1968.

[18] Victor Eijkhout. The 'weighted modification' incomplete factorisation method. Technical Report Lapack working note 145, UT-CS-99-436, Computer Science Department, University of Tennessee.

[19] Victor Eijkhout. Analysis of parallel incomplete point factorizations. *Lin. Alg. Appl.*, 154–156:723–740, 1991.

[20] Victor Eijkhout. Beware of unperturbed modified incomplete point factorizations. In Robert Beauwens and Pieter de Groen, editors, *Proceedings of the IMACS International Symposium on Iterative Methods in Linear Algebra, Brussels Belgium*, pages 583–591, 1992.

[21] Victor Eijkhout and Svetozar Margenov. Preconditioned iterative solution of spline finite element systems. *Mathematica Balkanica*, 4:350–367, 1990.

[22] Stanley C. Eisenstat. Efficient implementation of a class of preconditioned conjugate gradient methods. *SIAM J. Sci. Stat. Comput.*, 2:1–4, 1981.

[23] Howard C. Elman. A stability analysis of incomplete LU factorization. *Math. Comp.*, 47:191–217, 1986.

[24] Howard C. Elman. Relaxed and stabilized incomplete factorizations for non-self-adjoint linear systems. *BIT*, 29:890–915, 1989.

[25] Howard C. Elman and Martin H. Schultz. Preconditioning by fast direct methods for non self-adjoint nonseparable elliptic equations. *SIAM J. Numer. Anal.*, 23:44–57, 1986.

[26] I. Gustafsson. A class of first-order factorization methods. *BIT*, 18:142–156, 1978.

[27] Ivar Gustafsson. An incomplete factorization preconditioning method based on modification of element matrices. *BIT*, 36:86–100, 1996.

[28] Louis A. Hageman and David M. Young. *Applied Iterative Methods*. Academic Press, New York, 1981.

[29] A. Jennings and G.M. Malik. Partial elimination. *J. Inst. Maths Applics*, 20:307–316, 1977.

[30] D.S. Kershaw. The incomplete cholesky-conjugate gradient method for the iterative solution of systems of linear equations. *J. Comp. Phys.*, 26:43–65, 1978.

[31] L. Yu. Kolotilina and A. Yu. Yeremin. On a family of two-level preconditionings of the incomlete block factorization type. *Sov. J. Numer. Anal. Math. Modelling*, pages 293–320, 1986.

[32] L. Yu. Kolotilina and A. Yu. Yeremin. Block SSOR preconditionings for high order 3d FE systems. *BIT*, 29:805–823, 1989.

[33] T.A. Manteuffel. An incomplete factorization technique for positive definite linear systems. *Math. Comp.*, 34:473–497, 1980.

[34] J.A. Meijerink and H.A. van der Vorst. An iterative solution method for linear systems of which the coefficient matrix is a symmetric m-matrix. *Math Comp*, 31:148–162, 1977.

[35] A. Neumaier. On the comparison of H-matrices with M-matrices. *Lin. Alg. Appl.*, 83:135–141, 1986.

[36] Y. Notay. Solving positive (semi)definite linear systems by preconditioned iterative methods. In O. Axelsson and L.Yu. Kolotilina, editors, *Preconditioned Conjugate Gradient Methods*, pages 105–125, Nijmegen, 1989. Lecture Notes in Mathematics, vol. 1457.

[37] A.M. Ostrowski. über die Determinanten mit überwiegenden Hauptdiagonale. *Comment. Math. Helv.*, 10:69–96, 1937.

[38] Seymour V. Parter. Multi-line iterative methods for elliptic difference equations and fundamental frequencies. *Numer. Math.*, 3:305–319, 1961.

[39] Ben Polman. Incomplete blockwise factorizations of (block) H-matrices. *Lin. Alg. Appl.*, 90:119–132, 1987.

[40] Claude Pommerell. *Solution of Large Unsymmetric Systems of Linear Equations*, volume 17 of *Series in Micro-electronics, volume 17*. Hartung-Gorre Verlag, Konstanz, 1992.

[41] Yves Robert. Regular incomplete factorizations of real positive definite matrices. *Lin. Alg. Appl.*, 48:105–117, 1982.

[42] K. Takahishi, J. Fagan, and M.S. Chen. Formation of a sparse bus impedance matrix and its application to short circuit study. In *8th PICA conference, Minneapolis*, pages 63–69, 1973.

[43] Henk van der Vorst. Iterative slution methods for certain sparse linear systems with a non-symmetric matrix arising from PDE-problems. *J. Comput. Phys.*, 44:1–19, 1981.

[44] Henk van der Vorst. High performance preconditioning. *SIAM J. Sci. Stat. Comput.*, 10:1174–1185, 1989.

[45] R.S. Varga, E.B. Saff, and V. Mehrmann. Incomplete factorizations of matrices and connections with H-matrices. *SIAM J. Numer. Anal.*, 17:787–793, 1980.

[46] P.S. Vassilevski. Preconditioning nonsymmetric and indefinite finite element matrices. *J. Numer. Alg. Appl.*, 1:59–76, 1992.

[47] X. Wang, K. A. Gallivan, and R. Bramley. A necessary and sufficient symbolic condition for the existence of incomplete Cholesky factorization. Technical Report 440, Indiana University–Bloomington, Bloomington, IN 47405, 1995. Submitted to IMACS.