Lawrence Livermore Laboratory

On the Problem of Unstable Pivots in the Incomplete LU-Conjugate Gradient Method

David S. Kershaw

MASTER

June 19, 1978

Prepared for the Eighth Conference on Numerical Simulation of Plasmas,
Monterey, California, June 28-30, 1978.

This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

On the Problem of Unstable Pivots in the
Incomplete LU-Conjugate Gradient Method*

David S. Kershaw

University of California, Lawrence Livermore Laboratory
Livermore, CA 94550

ABSTRACT

Incomplete LU and incomplete-Cholesky conjugate gradient methods are becoming widely used in both laser and magnetic fusion research. In my original presentation of these methods [1], the problem of what to do if a pivot ($L_{ii}U_{ii}$) becomes very small or zero was raised and only partially answered by the suggestion that it be arbitrarily set to some non-zero value. In what follows it will be shown precisely how small the pivot can become before it must be fixed and precisely what value it should be set to in order to minimize the error in LU. Numerical examples will be given to show that not only does this prescription improve incomplete LU-conjugate gradient methods, but exact LU decomposition carried out with this prescription for handling small pivots and followed by a few linear or conjugate gradient iterations can be much faster than the permutations of rows and columns usually employed to circumvent small pivot problems.

*Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore Laboratory under contract No. W-7405-ENG-48.

NOTICE
This report was prepared as an account of work sponsored by the United States Government. Neither the United States nor the United States Department of Energy, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product or process disclosed, or represents that its use would not infringe privately owned rights.

DEPARTMENT OF ENERGY UNCLASSIFIED

124

In reference [1] the point was raised that in incomplete Cholesky decompositions pivots (L_{ij} in eq. (7b) of reference [1]) may arise which are ≤ 0 . Similarly, in incomplete LU decompositions pivots (U_{ij} in Appendix A of reference [1]) may arise which are $= 0$. In either case this leads to a breakdown of the algorithms unless something is done to "fix" these "bad" pivots. In reference [1] it was suggested that if in the course of an incomplete Cholesky decomposition one comes up with $L_{ij} \leq 0$, one should set

$$L_{ij} = \sum_{j=1}^{(i-1)} |L_{ij}| + \sum_{j=(i+1)}^N |L_{ij}|$$

thus assuring diagonal dominance of the i 'th row and column of L . This has worked quite well in practice. In [1] it was also suggested that if in the course of an incomplete LU decomposition one comes up with $U_{ij} = 0$, "simply set U_{ij} to a non zero value and go on with the algorithm." This is too vague to be very satisfactory and in what follows we shall derive a quantitative theory to determine:

1. Exactly how small must the pivot be before we alter it?

Obviously on a computer of given accuracy there is some pivot value which is small enough to make the algorithm go unstable but is not yet a hard zero.

2. If a pivot needs to be "fixed" what value should we set it to so as to get the best possible approximate inverse?

It shall be assumed in what follows that we are dealing with a floating point computer where the mantissa is stored with t binary digits so the accuracy with which any number can be stored is 1 part in 2^t .

1. BASIC THEORY

Consider first complete LU decomposition.

We first note the theorem of Wilkinson [2] which states that if one performs an LU decomposition of a matrix A with exact arithmetic except for an error which one makes in computing one element L_{ij} ($i > j$) then all elements of the error matrix $E = A - LU$ will be zero except E_{ij} . Similarly, if we make an error only in U_{ij} ($i \leq j$) then only E_{ij} will be nonzero. Thus errors do not propagate.

Now consider the effect of altering a pivot in LU decomposition. We take

$$U_{ii} = A_{ii} - \sum_{j < i} L_{ij} U_{ji} - \delta_i = \tilde{U}_{ii} - \delta_i,$$

where δ_i is our "fix" that alters the i 'th pivot away from its unstable value, and $\tilde{U}_{ii} = A_{ii} - \sum_{j < i} L_{ij} U_{ji}$, is the unfixed pivot. By Wilkinson's theorem the sole effect of this on the error matrix in exact arithmetic is to introduce one nonzero element,

$$E_{ji} = + \delta_i,$$

into $E = A - LU$. What is the effect on round off error in the rest of the decomposition? First, all the elements of L in the i 'th column become

$$\begin{aligned} L_{ji} &= (A_{ji} - \sum_{k < i} L_{jk} U_{ki}) / U_{ii} \\ &= S_{ji} / (\tilde{U}_{ii} - \delta_i), \quad (j > i) \end{aligned}$$

$$\text{where } S_{ji} = A_{ji} - \sum_{k < i} L_{jk} U_{ki}.$$

(Note that E_{ji} is independent of δ_i since $L_{ji}(\tilde{U}_{ii} - \delta_i)$ is a constant independent of the choice of δ_i). Now consider the round off error in the calculation of any subsequent L or U element which depends on L_{ji} . We have

$$\begin{aligned} L_{jn} U_{nn} &= A_{jn} - L_{ji} U_{in} - \sum_k L_{jk} U_{kn}, \\ (n < j), (k < n), (k \neq i), (j, n > i), \end{aligned}$$

and

$$\begin{aligned} U_{jn} &= A_{jn} - L_{ji} U_{in} - \sum_k L_{jk} U_{kn}, \\ (n \geq j), (k < j), (k \neq i), (j, n > i). \end{aligned}$$

Therefore the contribution from roundoff error to E_{jn} due to the L_{ji} term is given by

$$\begin{aligned} E_{jn} &= \pm L_{ji} U_{in} 2^{-t} \\ &= \pm S_{ji} U_{in} 2^{-t} / (\tilde{U}_{ij} - \delta_i) . \end{aligned}$$

To proceed further we need a definition of what constitutes the best LU decomposition. We define the best pivot fix, δ_i , to be the one that minimizes the maximum of all the δ_i dependent contributions to the error matrix E. These δ_i dependent error contributions are to E_{ij} , which has an error contribution δ_i and to all E_{jn} (with $j, n > i$) which have error contributions $S_{ji} U_{in} 2^{-t} / (\tilde{U}_{in} - \delta_i)$. If we set

$$\sigma_i = \text{MAX}_{j > i} |S_{ji}| ,$$

$$\text{and } \mu_i = \text{MAX}_{n > i} |U_{in}| ,$$

then the maximum error contribution will be

$$\text{MAX} \left(|\delta_i| , 2^{-t} \sigma_i \mu_i / |\tilde{U}_{ii} - \delta_i| \right) .$$

The value of δ_i which minimizes this is

$$\delta_i = - \text{sign}(\tilde{U}_{ii}) \left(\sqrt{(\tilde{U}_{ii}/2)^2 + 2^{-t} \sigma_i \mu_i} - (\tilde{U}_{ii}/2) \right) ,$$

or

$$U_{ii} = \tilde{U}_{ii} - \delta_i = (\tilde{U}_{ii}/2) + \text{sign}(\tilde{U}_{ii}) \sqrt{(\tilde{U}_{ii}/2)^2 + 2^{-t} \sigma_i \mu_i} .$$

For ease of computation we approximate this by

$$U_{ii} = \begin{cases} \tilde{U}_{ii}, & (\tilde{U}_{ii})^2 > 2^{-t} \sigma_i \mu_i \\ \text{sign}(\tilde{U}_{ii}) \sqrt{2^{-t} \sigma_i \mu_i}, & (\tilde{U}_{ii})^2 < 2^{-t} \sigma_i \mu_i. \end{cases}$$

Thus our prescription for LU decomposition is as follows:

1. As we compute the i 'th row of U we find

$$\mu_i = \max_{n > i} |U_{in}|.$$

2. As we compute the i 'th column of L (but before dividing by the i 'th pivot) we find

$$\sigma_i = \max_{j > i} |S_{ji}|.$$

3. If $(U_{ii})^2 > 2^{-t} \sigma_i \mu_i$ we leave the i 'th pivot (U_{ii}) unchanged.
If $(U_{ii})^2 < 2^{-t} \sigma_i \mu_i$ we set the i 'th pivot (the new U_{ii}) =
 $\text{sign}(U_{ii}) \sqrt{2^{-t} \sigma_i \mu_i}.$
4. We divide all the S_{ji} ($j > i$) by the new U_{ii} to get the i 'th column of L .

For incomplete LU decompositions the arguments and the resulting prescriptions are exactly the same except only elements within the sparsity pattern chosen are calculated, used and stored.

II. NUMERICAL RESULTS

We tested our new prescription on three matrices whose complete LU decompositions (without pivoting) are totally unstable.

Example 1

$A_{i,i+1} = -A_{i+1,i} = 1$ and all other elements = 0. If the dimension N is even the matrix is nonsingular while if N is odd it is singular. We took $N = 1000$. Since the determinant of every odd dimensional principal minor is zero, every other pivot is bad and LU decomposition (without pivoting) blows up immediately. On the CDC 7600 ($t = 48$ binary digit mantissa) we performed the LU decomposition using our modified pivot prescription. We then calculated

$$A - LU = E_1$$

and

$$I - (LU)^{-1}A = E_2.$$

The largest elements of E_1 and E_2 had absolute values of $\approx 2^{-24}$.

Example 2

$$A_{ij} = 1, i + j \leq N + 1,$$

$$\text{and } A_{ij} = 0, i + j > N + 1,$$

where N is the dimension and we used $N = 20$. A is nonsingular but all its principal minors except for A itself and the 1×1 principal minor are singular. Therefore LU decomposition is totally unstable but with our modified prescription we again obtain an LU decomposition which gives an E_1 and E_2 the largest elements of which have absolute values $\approx 2^{-24}$.

Example 3

$$A = B - \alpha I, \text{ where}$$

$$B_{ii} = 2, B_{i,i+1} = B_{i+1,i} = -1, \text{ and all other elements of } B \text{ were zero.}$$

The eigenvalues of B are

$$\lambda_j = 4 \sin^2 \left(\pi j / (2N + 2) \right); j = 1, 2, \dots, N,$$

where N is the dimension. We took $N = 1000$, and $\alpha = 4 \sin^2(\pi/5)$.

Then the i 'th principal minor of A is singular for $i = 5n - 1$; $n = 1, 2, 3, \dots, 200$, while A itself is not singular. Thus every 5th pivot is bad and regular LU decomposition blows up. Our prescription produced an LU decomposition such that the largest elements of E_1 and E_2 had absolute values $\approx 2^{-24}$.

Thus on all three test problems our modified prescription was able to take matrices whose regular LU decomposition (without pivoting) was completely unstable and produce an approximate LU decomposition such that $(LU)^{-1}A = I$ and $LU = A$ to 24 (binary digit) accuracy (on a 48 binary digit machine).

III. ADDITIONAL APPLICATIONS

Our experience with the three test problems suggests that, in addition to its intended application in incomplete factorization schemes, our pivot prescription might be a very viable alternative to pivoting in complete LU decomposition algorithms. Complete LU decomposition with our modified pivot prescription (and no permutations of rows or columns) followed by conjugate gradient if needed (i.e., if any pivots required modification) to further improve the answer has many advantages over the usual pivoting schemes such as:

1. In the sparse matrix case pivoting usually causes many more elements which were zero in the original matrix to fill in and become non zero in L and U than would have been the case if there had been no permutations of rows or columns. In our method we suffer from none of this additional fill in. Furthermore, various very effective minimal storage schemes (such as nested dissection) for Gauss elimination with sparse matrices have so far only been applicable to positive definite symmetric matrices because for more general matrices pivoting was required.

Our modified pivot prescription would make these methods applicable to any matrices with the appropriate sparsity patterns.

2. LU decomposition without pivoting vectorizes very well and Fong and Jordan [3] have shown that it can go at super vector speeds on the Cray 1 computer. However, permuting rows and columns does not vectorize well and so our modified pivot method, which already entails much less work than pivoting in scalar computing, will have an even greater advantage on vector machines.
3. Computational experience shows that in many production code applications of LU decomposition most of the matrices produced do not have "bad" pivots and so, with conventional methods, most of the time a lot of unnecessary permuting of rows and columns is being done when just as good results would have been obtained from LU decomposition without pivoting. It is only the occasional matrix with "bad" pivots that needs some special attention and our method does very little extra work except when a "bad" pivot shows up.

IV. A MORE EFFICIENT FORM FOR THE INCOMPLETE

LU - CONJUGATE GRADIENT ALGORITHM

The algorithm given in Ref. [1], Appendix A, eqs. (9'a) through (9'e), may be made computationally more efficient by the substitution $p_i(\text{old}) = (U^T U)^{-1} p_i(\text{new})$. This gives the more efficient form of the algorithm:

$$r_0 = y - Ax_0 \quad \text{and} \quad p_0 = A^T(LL^T)^{-1}r_0 ,$$

$$a_i = \frac{(r_i, (LL^T)^{-1}r_i)}{(p_i, (U^TU)^{-1}p_i)} , \quad (9'a)$$

$$x_{i+1} = x_i + a_i(U^TU)^{-1}p_i , \quad (9'b)$$

$$r_{i+1} = r_i - a_iA(U^TU)^{-1}p_i , \quad (9'c)$$

$$b_i = \frac{(r_{i+1}, (LL^T)^{-1}r_{i+1})}{(r_i, (LL^T)^{-1}r_i)} , \quad (9'd)$$

$$p_{i+1} = A^T(LL^T)^{-1}r_{i+1} + b_i p_i , \quad (9'e)$$

$$i = 0, 1, 2, \dots$$

REFERENCES

1. D. S. Kershaw, The Incomplete Cholesky-Conjugate Gradient Method for the Iterative Solution of Systems of Linear Equations, J. Comp. Physics 26 #1, January 1978, p. 43.
2. J. H. Wilkinson, The Algebraic Eigenvalue Problem, Clarendon Press, Oxford, 1965, Chapter 4, sec. 40.
3. K. Fong and T. L. Jordon, Some Linear Algebraic Algorithms and Their Performance on Cray 1, Los Alamos Scientific Laboratory Report, LA-6774, June 1977.

NOTICE

"This report was prepared as an account of work sponsored by the United States Government. Neither the United States nor the United States Department of Energy, nor any of their employees nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness or usefulness of any information, apparatus, product or process disclosed, or represents that its use would not infringe privately-owned rights."

Reference to a company or product names does not imply approval or recommendation of the product by the University of California or the U.S. Department of Energy to the exclusion of others that may be suitable.