Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 24, pp. 153-167, October 2012

ACCURATE AND EFFICIENT LDU DECOMPOSITIONS OF DIAGONALLY DOMINANT M-MATRICES*

A. BARRERAS† AND J. M. PEÑA†

Key words. *LDU* decomposition, rank revealing decomposition, accuracy, M-matrices, diagonal dominance.

AMS subject classifications. 65F05, 65F35, 65F15, 15A18

Abstract. An efficient method for the computation to high relative accuracy of the LDU decomposition of an $n \times n$ row diagonally dominant M–matrix is presented, assuming that the off–diagonal entries and row sums are given. This method costs an additional $\mathcal{O}(n^2)$ elementary operations over the cost of Gaussian elimination, and leads to a lower triangular, column diagonally dominant matrix and an upper triangular, row diagonally dominant matrix. Comparisons with other methods in the literature are commented and illustrated.

1. Introduction.

Recent advances in Numerical Linear Algebra have shown that certain classes of matrices allow computation of certain matrix functions to high relative accuracy, independently of the size of the classical condition number. Some of these classes of matrices are defined by special sign or other structure and require knowledge of some natural parameters to high relative accuracy. In most of those cases, accurate spectral computation (eigenvalues, singular values) is assured once we have an accurate matrix factorization with a suitable pivoting. For instance, the bidiagonal decomposition in the case of totally nonnegative matrices (see also [5], [10]) or an LDU factorization after a symmetric pivoting in the case of diagonally dominant matrices (cf. [4], [13], [14]).

Let us focus now on the problem considered in this paper. An algorithm published in [2] computes with high relative accuracy the LDU factorization of an $n \times n$ row diagonally dominant M–matrix, if the off–diagonal entries and the row sums are given. The trick is to modify Gaussian elimination to compute the off–diagonal entries and the row sums of each Schur complement without performing subtractions. In addition, symmetric complete pivoting was used in [4] in order to obtain well conditioned L and U factors (U is even row diagonally dominant). This factorization is a special case of a

*Received by the editors on December 22, 2011. Accepted for publication on August 1, 2012. Handling Editor: Carlos Fonseca.

†Dept. Applied Mathematics/IUMA, University of Zaragoza, Spain; *e-mails*: albarrer@unizar.es, jmpena@unizar.es Research Partially Supported by the Spanish Research Grant MTM2009-07315, by Gobierno de Aragón and Fondo Social Europeo.

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 24, pp. 153-167, October 2012

rank revealing decomposition, as will be recalled in Section 2. Demmel et al. showed in [3], with the corresponding algorithm, that the singular value decomposition can be computed accurately and efficiently for matrices that admit accurate rank revealing decompositions. To implement symmetric complete pivoting, the algorithm in [4] computes all the diagonal entries and all Schur complements and this increases the cost in $\mathcal{O}(n^3)$ flops with respect to standard Gaussian elimination. In [13] another pivoting strategy was used, also with a subtraction-free implementation and a similar computational cost, but leading to both triangular matrices L and U column and row diagonally dominant, respectively. In Section 2 we recall nice bounds for the condition number of such matrices L and U (see also [11] or [12]).

This paper, in Section 3, also presents, for a diagonally dominant M–matrix A, a factorization $PAP^T = LDU$, where L is a unit lower triangular, column diagonally dominant matrix and U is a unit upper triangular, row diagonally dominant matrix. For this purpose, we develop an accurate algorithm that requires $\mathcal{O}(n^2)$ elementary operations beyond the cost of Gaussian elimination. To achieve this, the main idea is to update $c^{(i)}$, the sum of each column in the process of Gaussian elimination, and to avoid subtractions, decompose $c^{(i)}$ into $c^{(i)} = h^{(i)} - s^{(i)}$, where $h^{(i)}$ and $s^{(i)}$ can be updated without subtractions. Let us recall that, although in [11] it was shown that the pivoting strategy used in [13] could be implemented with an additional cost of $\mathcal{O}(n^2)$ elementary operations beyond the cost of standard Gaussian elimination, the implementation of [11] was not subtraction-free, in contrast to those of [4], [13] and the present paper.

In Remark 3.2 it is also shown that our method is valid for diagonally dominant matrices having certain sign patterns: with off–diagonal entries of the same sign or satisfying a chessboard pattern. Numerical examples at the end of Section 3 show that the lower triangular matrices obtained with our method can be much better conditioned than those obtained with symmetric complete pivoting. Section 4 contains some concluding remarks. Finally, let us mention that the problem of computing an accurate LDU decomposition of diagonally dominant matrices has been solved by Ye in [14] (see also [6]).

2. Basic concepts, methods and notations. Let us start by introducing some classes of matrices used in this paper. A real matrix with nonpositive off–diagonal elements is called a *Z–matrix*. We say that a matrix $A = (a_{ij})_{1\le i,j\le n}$ is *row diagonally dominant* if, for each $i = 1, \dots, n$, $|a_{ii}| \ge \sum_{j\ne i} |a_{ij}|$ (these matrices are called weakly diagonally dominant in [4]). If A^T is row diagonally dominant, then we say that A is column diagonally dominant. Given a matrix $A = (a_{ij})_{1\le i,j\le n}$, its *comparison matrix* $\mathcal{M}(A) = (m_{ij})_{1\le i,j\le n}$ is the Z–matrix defined by $m_{ii} := |a_{ii}|$ and $m_{ij} := -|a_{ij}|$ if $i \ne j$, $1 \le i, j \le n$. Let us recall that if a Z–matrix A can be expressed as $A = sI - B$, with $B \ge 0$ and $s \ge \rho(B)$ (where $\rho(B)$ is the spectral radius of B),

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 24, pp. 153-167, October 2012

then it is called an *M–matrix*. Let us also recall that a Z–matrix A is a nonsingular M–matrix if and only if A^{-1} is nonnegative. Observe that a Z–matrix row diagonally dominant with nonnegative diagonal entries is necessarily an M–matrix.

Given an algorithm using only additions of numbers of the same sign, multiplications and divisions, and assuming that each initial real datum is known to high relative accuracy, then it is well–known that the output of that algorithm can be computed to high relative accuracy (cf. [3, p. 52]). That is, the only forbidden operation is true subtraction, due to possible cancellation in leading digits. In this paper we will use the word *accurately* to mean *to high relative accuracy*. A *rank revealing decomposition* of a matrix A is defined in [3] as a decomposition $A = XDY^T$, where X, Y are well conditioned and D is a diagonal matrix. In [3] Demmel et al. showed, with the corresponding algorithm, that the singular value decomposition can be computed accurately and efficiently for matrices that admit accurate rank revealing decompositions.

Let us recall that an idea that has played a crucial role in some recent works on accurate computations has been the need to reparametrize matrices belonging to some special classes. For instance, in the class of totally nonnegative matrices the parameters are (see [5] and [10]) the multipliers of an elimination process called Neville elimination (see [8]). In the class of M–matrices, the natural parameters that allow us to derive accurate and efficient algorithms are the off–diagonal entries and the row sums (or the column sums): see [1], [2] and [4], where the class of M–matrices row diagonally dominant was considered. For instance, in the field of digital electrical circuits, the column sums are given by the quotient between the conductance and capacitance of each node (see [1]).

As usual, an LDU factorization of a square matrix $A = LDU$ means that L is a lower triangular matrix with unit diagonal (unit lower triangular), D is a diagonal matrix and U is an upper triangular matrix with unit diagonal (unit upper triangular). Given $k \in \{1, 2, \ldots, n\}$, let α, β be two increasing sequences of k positive integers less than or equal to n. Then we denote by $A[\alpha|\beta]$ the $k \times k$ submatrix of A containing rows numbered by α and columns numbered by β. For principal submatrices, we use the notation $A[\alpha] := A[\alpha|\alpha]$. Gaussian elimination with a given pivoting strategy, for nonsingular matrices $A = (a_{ij})_{1\le i,j\le n}$, consists of a succession of at most $n-1$ major steps resulting in a sequence of matrices as follows:

$$A = A^{(1)} \longrightarrow \tilde{A}^{(1)} \longrightarrow A^{(2)} \longrightarrow \tilde{A}^{(2)} \longrightarrow \cdots \longrightarrow A^{(n)} = \tilde{A}^{(n)} = DU, \tag{2.1}$$

where $A^{(t)} = (a_{ij}^{(t)})_{1\le i,j\le n}$ has zeros below its main diagonal in the first $t-1$ columns and DU is upper triangular with the pivots on its main diagonal. The matrix $\tilde{A}^{(t)} = (\tilde{a}_{ij}^{(t)})_{1\le i,j\le n}$ is obtained from the matrix $A^{(t)}$ by reordering the rows and/or columns $t, t+1, \ldots, n$ of $A^{(t)}$ according to the given pivoting strategy and satisfying $\tilde{a}_{tt}^{(t)} \neq 0$.

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 24, pp. 153-167, October 2012

To obtain $A^{(t+1)}$ from $\tilde{A}^{(t)}$ we produce zeros in column t below the *pivot element* $\tilde{a}_{tt}^{(t)}$ by subtracting multiples of row t from the rows beneath it. If the matrix A is singular, in this paper we allow the resulting matrices in (2.1) to have $\tilde{a}_{tt}^{(t)} = 0$, but (as we shall see later) in this case its corresponding column and row are null,

$$A^{(t)}[t, \ldots, n|t] = 0, \quad A^{(t)}[t|t, \ldots, n] = 0, \tag{2.2}$$

and we continue the elimination process with $A^{(t+1)}[t+1, \ldots n] = A^{(t)}[t+1, \ldots n]$.

We say that we carry out a *symmetric pivoting strategy* when we perform the same row and column exchanges, that is, $PAP^T = LDU$, where P is the associated permutation matrix. Let us present several symmetric pivoting strategies for row diagonally dominant matrices that either have been used in other papers or will be used in this paper. Since row diagonal dominance is inherited by Schur complements in the Gaussian elimination, Gaussian elimination with symmetric pivoting preserves it, that is, all matrices $A^{(t)}$ of (2.1) are row diagonally dominant (and, in particular, DU and so U). Therefore, it is sufficient to describe the choice of the first pivot $\tilde{a}_{11} = a_{kk}$. On the one hand, the symmetric pivoting that selects the maximum entry on the diagonal for the pivot will be equivalent to complete pivoting and was used in [4]. It leads to U row diagonally dominant, and so well conditioned, and to L, which is usually well conditioned as well. On the other hand, since A is row diagonally dominant, we have

$$\sum_{i=1}^{n} |a_{ii}| \geq \sum_{i=1}^{n} \sum_{j=1, j\neq i}^{n} |a_{ij}|,$$

and there exists k such that column k is diagonally dominant, that is,

$$|a_{kk}| \geq \sum_{i=1, i\neq k}^{n} |a_{ik}|. \tag{2.3}$$

The symmetric pivoting strategy that chooses the first pivot $\tilde{a}_{11} = a_{kk}$ was called in [14] *column diagonal dominance pivoting*. In [13] the first pivot $\tilde{a}_{11} = a_{kk}$ was chosen so that it gives the most diagonal dominance in (2.3) (i.e., the largest difference between the absolute value of a diagonal entry and the sum of the absolute values of the off–diagonal entries of the corresponding row), and this strategy is a particular case of column diagonal dominance pivoting. In this paper we shall use a strategy that we call *weak column diagonal dominance pivoting*: it is a symmetric pivoting strategy that chooses the first pivot $\tilde{a}_{11} = a_{kk}$ satisfying (2.3), and without the necessity of being nonzero. If $\tilde{a}_{11} = 0$, then its row and column diagonal dominance implies that its row and column are null, and we continue the elimination process with $A^{(2)}[2, \ldots, n] = A[2, \ldots, n]$ (as we had announced for the t-th pivot in (2.2)). In order to uniquely determine this strategy, we can choose the first index k satisfying (2.3).

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 24, pp. 153-167, October 2012

Column diagonal dominance pivoting and weak column diagonal dominance pivoting lead to U row diagonally dominant and to L column diagonally dominant. Then both triangular matrices are always well conditioned. In fact, since L is unit lower triangular column diagonally dominant, we know by [13, Proposition 2.1] and [13, Remark 2.2] that

$$\kappa_\infty(L) = \|L\|_\infty \, \|L^{-1}\|_\infty \leq n^2 \quad \text{and} \quad \kappa_1(L) = \|L\|_1 \, \|L^{-1}\|_1 \leq 2n. \tag{2.4}$$

Analogously, with U unit upper triangular and row diagonally dominant, we have

$$\kappa_\infty(U) \leq 2n \quad \text{and} \quad \kappa_1(U) \leq n^2.$$

In contrast, symmetric complete pivoting leads to L that is usually well conditioned, but it is not necessarily column diagonally dominant. Finally, let $e := (1, \dots, 1)^T$ and let

$$r := Ae \tag{2.5}$$

be the vector of row sums.

It is well known (cf. [2]) that we can carry out the Gaussian elimination of a diagonally dominant M–matrix with high relative accuracy because there is no subtraction involved throughout the process. Summarizing the process of [2, Algorithm 1], it starts with (2.5) and at each step of the Gaussian elimination it is only necessary to update the vector r. Diagonal entries of the matrix are not computed at each step (except the pivot) and so, the computational cost is of order $\mathcal{O}(n^2)$ beyond the cost of Gaussian elimination. We can also conclude that it is possible to compute the inverse of a nonsingular diagonally dominant M–matrix, A, with high relative accuracy, by the following procedure. We obtain the LDU factorization of A accurately. Then, it is well known (cf. [9, Section 13.2]) that we can compute the inverse of L and U without subtraction in the process. Thus, we can compute $A^{-1} = U^{-1}D^{-1}L^{-1}$ with high relative accuracy.

If A is a row and column diagonally dominant M–matrix, then no pivoting strategy is necessary to compute an accurate LDU factorization with L and U column and row diagonally dominant, respectively, because L also inherits through Gaussian elimination the column diagonal dominance from A. In fact, Gaussian elimination can be applied without row or column exchanges and so, for each $t = 1, \dots, n-1$, $A^{(t)} = \tilde{A}^{(t)}$ (see (2.1)) and all matrices $A^{(t)}[t, \dots, n]$ are row and column diagonally dominant. In conclusion, given the off–diagonal elements of a row and column diagonally dominant M–matrix $A = (a_{ij})_{1 \leq i,j \leq n}$ and the vector r of row sums (see (2.5)), we can calculate with high relative accuracy the LDU decomposition of A, where L is column diagonally dominant and U is row diagonally dominant. Moreover, this computation is subtraction-free and so can be performed with high relative accuracy.

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 24, pp. 153-167, October 2012

3. Accurate and efficient LDU decomposition of row diagonally dominant M–matrices. In this section we provide an accurate and efficient method for obtaining the LDU factorization (with L column diagonally dominant and U row diagonally dominant) of a row diagonally dominant M–matrix provided its off–diagonal entries and its row sums. Using A^T instead of A, we have also an accurate method for obtaining the LDU factorization of a column diagonally dominant M–matrix provided its off–diagonal entries and its column sums. The comparison with the computational cost of the methods presented in [4] and in [13, Section 4] can be seen in Remark 3.1. This method produces a matrix U with a similar conditioning as in those papers because it is also row diagonally dominant and a matrix L that can be better conditioned than that of [4] (as the matrices of (3.28) shows) and satisfies bounds (2.4) because it is column diagonally dominant, as commented previously. We start by presenting our algorithm to compute the LDU decomposition of a row diagonally dominant M–matrix.

Algorithm 1

Input: $A = [a_{ij}]$ $(i \neq j)$ and $r = [r_i] \geq 0$

For $i = 1 : n$

$p_i = \sum_{j=1, j\neq i}^{n} a_{ij}$

$a_{ii} = r_i - p_i$

$s_i = \sum_{j=1, j\neq i}^{n} a_{ji}$

$h_i = a_{ii}$

End For

Choose an interchange permutation P_1 such that $A = P_1 A P_1^T$ satisfies $h_1 \geq -s_1$, where $h = P_1 h$, $s = P_1 s$

Initialize: $P = P_1$; $L = I$; $D = \text{diag}(d_i)_{i=1}^{n} = \text{diag}(h_1, 0 \ldots, 0)$; $r = P_1 r$

For $k = 1 : (n-1)$

If $d_k = 0$

For $i = (k+1) : n$

$l_{ik} = 0$

$n_{ki} = 0$

End For

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 24, pp. 153-167, October 2012

Else

For $i = (k+1) : n$

$l_{ik} = a_{ik}/a_{kk}$

$n_{ki} = a_{ki}/a_{kk}$

$r_i = r_i - l_{ik} r_k$

$h_i = h_i - n_{ki} h_k$

$s_i = s_i - n_{ki} s_k$

For $j = (k+1) : n$

If $i \neq j$

$a_{ij} = a_{ij} - l_{ik} a_{kj}$

End If

End For

End For

End If

Choose interchange permutation P_2 such that $A = P_2 A P_2^T$ satisfies $h_{k+1} \geq -s_{k+1}$, where $h = P_2 h$, $s = P_2 s$

$P = P_2 P$; $L = P_2 L P_2$; $r = P_2 r$

$p_{k+1} = \sum_{j=k+2}^{n} a_{k+1,j}$

$a_{k+1,k+1} = r_{k+1} - p_{k+1}$

$d_{k+1} = a_{k+1,k+1}$

End For

In **output**, the algorithm produces the factorization $PAP^T = LDU$ (nontrivial entries of U are stored in $N = (n_{ij})_{1 \leq i < j \leq n}$).

The following result proves the nice properties of the previous algorithm.

THEOREM 3.1. *Given the off–diagonal elements of a row diagonally dominant M–matrix $A = (a_{ij})_{1 \leq i,j \leq n}$ and the vector r of row sums (see (2.5)), we can compute, by Algorithm 1, with high relative accuracy the LDU decomposition of PAP^T, where P the permutation matrix associated to a weak column diagonal dominance pivoting*

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 24, pp. 153-167, October 2012

strategy applied when performing Gaussian elimination of A and such that L is column diagonally dominant and U is row diagonally dominant. Moreover, this computation is subtraction-free and can be performed with a computational cost that exceeds that of the Gaussian elimination by at most $(7n^2-11n+6)/2$ additions, $n(n-1)$ multiplications, $n(n-1)/2$ quotients and $n(n-1)/2$ comparisons.

Proof. If we consider the linear system (2.5) $Ae = r$ (where $e := (1,\dots,1)^T$), simultaneously to the sequence of matrices (2.1), we can obtain the corresponding sequence of vectors in $\mathbf{R}^n$

$$r = r^{(1)} \longrightarrow \tilde{r}^{(1)} \longrightarrow r^{(2)} \longrightarrow \tilde{r}^{(2)} \longrightarrow \cdots \longrightarrow r^{(n)}, \tag{3.1}$$

giving each $r^{(k)}$ and $\tilde{r}^{(k)}$ the row sums of the corresponding matrices $A^{(k)}$ and $\tilde{A}^{(k)}$, respectively, because e is the solution of $Ax = r$ and also of the equivalent systems $A^{(k)}x = r^{(k)}$ and $\tilde{A}^{(k)}x = \tilde{r}^{(k)}$. Now, for each $k = 1,\dots,n-1$ such that $\tilde{a}^{(k)}_{kk} \neq 0$ and for each $j = k+1,\dots,n$, we have

$$r^{(k+1)}_j = \tilde{r}^{(k)}_j - \frac{\tilde{a}^{(k)}_{jk}}{\tilde{a}^{(k)}_{kk}}\tilde{r}^{(k)}_k. \tag{3.2}$$

We can also consider the sums $p^{(k)}_i$ of the off–diagonal entries of each row i of $A^{(k)}$: for each k

$$p^{(k)}_i = \sum_{j=1, j\neq i}^{n} a^{(k)}_{ij}, \quad i = 1,\dots,n, \tag{3.3}$$

and form the vector $p^{(k)} = (p^{(k)}_1,\dots,p^{(k)}_n)^T$.

Let us introduce the corresponding sequence of vectors in $\mathbf{R}^n$ (analogous to (3.1)):

$$c^{(1)} \longrightarrow \tilde{c}^{(1)} \longrightarrow c^{(2)} \longrightarrow \tilde{c}^{(2)} \longrightarrow \cdots \longrightarrow c^{(n)} = \tilde{c}^{(n)}, \tag{3.4}$$

giving the last $n-k+1$ components of $c^{(k)}$ and $\tilde{c}^{(k)}$ the column sums of the corresponding matrices $A^{(k)}[k,\dots,n]$ and $\tilde{A}^{(k)}[k,\dots,n]$, respectively (see (2.1), where the matrices $A^{(k)}$ and $\tilde{A}^{(k)}$ have zeros below the first $k-1$ diagonal entries). In contrast,

$$c^{(k)}_j := \tilde{c}^{(k-1)}_j, \quad j < k. \tag{3.5}$$

If $\tilde{a}^{(k)}_{kk} = 0$, then $c^{(k+1)}_j = \tilde{c}^{(k)}_j$ for all j. Let us see a simple relation of the last $n-k$ components $c^{(k+1)}_j$ $(j = k+1,\dots,n)$ of the vector $c^{(k+1)}$ and the corresponding components of $\tilde{c}^{(k)}$ $(k = 1,\dots,n-1)$ when $\tilde{a}^{(k)}_{kk} \neq 0$. In fact,

$$c^{(k+1)}_j = \sum_{i=k+1}^{n} a^{(k+1)}_{ij} = \sum_{i=k+1}^{n}\left(\tilde{a}^{(k)}_{ij} - \frac{\tilde{a}^{(k)}_{ik}}{\tilde{a}^{(k)}_{kk}}\tilde{a}^{(k)}_{kj}\right) = \sum_{i=k+1}^{n}\tilde{a}^{(k)}_{ij} - \tilde{a}^{(k)}_{kj}\left(\sum_{i=k+1}^{n}\frac{\tilde{a}^{(k)}_{ik}}{\tilde{a}^{(k)}_{kk}}\right). \tag{3.6}$$

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 24, pp. 153-167, October 2012

Taking into account that

$$\sum_{i=k+1}^{n} \frac{\tilde{a}_{ik}^{(k)}}{\tilde{a}_{kk}^{(k)}} = \frac{\tilde{c}_k^{(k)} - \tilde{a}_{kk}^{(k)}}{\tilde{a}_{kk}^{(k)}},$$

we deduce from (3.6) that

$$c_j^{(k+1)} = \sum_{i=k}^{n} \tilde{a}_{ij}^{(k)} - \frac{\tilde{a}_{kj}^{(k)}}{\tilde{a}_{kk}^{(k)}} \tilde{c}_k^{(k)} \tag{3.7}$$

and so

$$c_j^{(k+1)} = \tilde{c}_j^{(k)} - \frac{\tilde{a}_{kj}^{(k)}}{\tilde{a}_{kk}^{(k)}} \tilde{c}_k^{(k)}, \quad j = k+1, \ldots, n. \tag{3.8}$$

Finally, we can consider the column sums of the off–diagonal elements of $A = A^{(1)}$:

$$s_j^{(1)} = \sum_{i=1, i \neq j}^{n} a_{ij}^{(1)}, \quad j = 1, \ldots, n. \tag{3.9}$$

The initial matrix $A = A^{(1)}$ is a row diagonally dominant M–matrix. These properties are inherited by performing a row permutation and the same column permutation, and it is well known that they are also inherited by Schur complements (see [7] for the property of being an M–matrix). So, since our method applies a symmetric pivoting strategy, these properties are inherited by all matrices $A^{(t)}[t, \ldots, n]$ and $\tilde{A}^{(t)}[t, \ldots, n]$ for $t = 1, \ldots, n$. The row diagonal dominance, together with the nonnegativity of the diagonal entries and the nonpositivity of the off–diagonal entries, implies that all vectors of row sums $r^{(t)}, \tilde{r}^{(t)}$ (see (3.1)) satisfy

$$r^{(t)}, \tilde{r}^{(t)} \geq 0, \quad t = 1, \ldots, n. \tag{3.10}$$

By (2.3), our pivoting strategy chooses an index $j_t \in \{t, \ldots, n\}$ satisfying the following property:

$$a_{j_t j_t}^{(t)} \geq \sum_{i=t, i \neq j_t}^{n} |a_{i j_t}^{(t)}|. \tag{3.11}$$

Taking into account the signs of the entries of $A^{(t)}[t, \ldots, n]$, and that the last $n - k + 1$ components $c_j^{(t)}$ of $c^{(t)}$ give the column sums of $A^{(t)}[t, \ldots, n]$, (3.11) is equivalent to

$$c_{j_t}^{(t)} \geq 0. \tag{3.12}$$

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 24, pp. 153-167, October 2012

Observe that if $(\tilde{a}_{tt}^{(t)} =) a_{j_t j_t}^{(t)} = 0$, then (3.11) implies that the column $A^{(t)}[t,\ldots,n|j_t]$ is null, and the row $A^{(t)}[j_t|t,\ldots,n]$ is also null by the row diagonal dominance of $A^{(t)}[t,\ldots,n]$ (as required in (2.2)). The *multiplier* l_{it} and the quotient n_{ti} $(i > t)$ are defined by

$$l_{it} := \frac{\tilde{a}_{it}^{(t)}}{\tilde{a}_{tt}^{(t)}}, \quad n_{ti} := \frac{\tilde{a}_{ti}^{(t)}}{\tilde{a}_{tt}^{(t)}} \tag{3.13}$$

if $0 \neq \tilde{a}_{tt}^{(t)}$, and by $l_{it} := 0$ and $n_{ti} := 0$ if $0 = \tilde{a}_{tt}^{(t)}$. Taking into account the signs of the entries of $\tilde{A}^{(t)}[t,\ldots,n]$, we conclude that all the elements defined in (3.13) are nonpositive. From (3.5) and (3.8) (when $\tilde{a}_{tt}^{(t)} \neq 0$), we derive

$$c_j^{(t+1)} = \tilde{c}_j^{(t)} - n_{tj}\tilde{c}_t^{(t)}, \quad j = t+1,\ldots,n, \tag{3.14}$$

and, from (3.2),

$$r_j^{(t+1)} = \tilde{r}_j^{(t)} - l_{jt}\tilde{r}_t^{(t)}. \tag{3.15}$$

Let us now define two new sequences of vectors which will be used for our choice of the pivot. Since

$$c_j^{(1)} = a_{jj}^{(1)} + s_j^{(1)}, \quad j = 1,\ldots,n,$$

denoting $h_j^{(1)} := a_{jj}^{(1)}$, $h^{(1)} = (h_1^{(1)},\ldots,h_n^{(1)})^T$ and $s^{(1)} = (s_1^{(1)},\ldots,s_n^{(1)})^T$, we can write

$$c^{(1)} = h^{(1)} + s^{(1)}. \tag{3.16}$$

Then we can generate two sequences of vectors

$$h^{(1)} \longrightarrow \tilde{h}^{(1)} \longrightarrow h^{(2)} \longrightarrow \tilde{h}^{(2)} \longrightarrow \cdots \longrightarrow h^{(n)} = \tilde{h}^{(n)} \tag{3.17}$$

and

$$s^{(1)} \longrightarrow \tilde{s}^{(1)} \longrightarrow s^{(2)} \longrightarrow \tilde{s}^{(2)} \longrightarrow \cdots \longrightarrow s^{(n)} = \tilde{s}^{(n)}, \tag{3.18}$$

in the same way as in (3.4), that is, $\tilde{h}^{(t)}$ and $\tilde{s}^{(t)}$ are obtained from $h^{(t)}$ and $s^{(t)}$, respectively, with the same permutation of indices as used to obtain $\tilde{c}^{(k)}$ from $c^{(k)}$, and $h^{(t+1)}$ and $s^{(t+1)}$ are obtained from $\tilde{h}^{(t)}$ and $\tilde{s}^{(t)}$ analogously to (3.14). In particular:

$$h_j^{(t+1)} = \tilde{h}_j^{(t)} - n_{tj}\tilde{h}_t^{(t)}, \quad j = t+1,\ldots,n \tag{3.19}$$

and

$$s_j^{(t+1)} = \tilde{s}_j^{(t)} - n_{tj}\tilde{s}_t^{(t)}, \quad j = t+1,\ldots,n. \tag{3.20}$$

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 24, pp. 153-167, October 2012

Clearly,

$$\tilde{c}^{(t)} = \tilde{h}^{(t)} + \tilde{s}^{(t)}, \quad c^{(t)} = h^{(t)} + s^{(t)}, \qquad t = 1, \ldots, n. \tag{3.21}$$

Taking into account that the multipliers are nonpositive and that $h^{(1)} \geq 0$ and $s^{(1)} \leq 0$, we conclude that

$$\tilde{h}^{(t)}, h^{(t)} \geq 0, \quad \tilde{s}^{(t)}, s^{(t)} \leq 0, \qquad t = 1, \ldots, n. \tag{3.22}$$

We have seen above that our symmetric pivoting strategy will choose an index $j_t \in \{t, \ldots, n\}$ satisfying (3.12), which is equivalent, by (3.21) and (3.22), to

$$h_{j_t}^{(t)} \geq |s_{j_t}^{(t)}|. \tag{3.23}$$

Finally, let us summarize the main steps of our method and its computational cost additional to that of Gaussian elimination. Assuming we know the vector $r = r^{(1)}$ of row sums of A and the off–diagonal entries of A, we shall show that each step is subtraction-free and so can be performed accurately. The initial steps, recalled in the following paragraph, only have to be carried out once with the initial matrix, in contrast to the remaining steps, which should be applied at each major step of Gaussian elimination.

We first calculate the vector $p^{(1)} = (p_1^{(1)}, \ldots, p_n^{(1)})^T$ with the sums of the off–diagonal entries of the rows of $A = A^{(1)}$ by means of (3.3) for $k = 1$, which can be performed accurately because all off–diagonal elements of A are nonpositive. It requires $n(n-2)$ additions. Then we calculate the diagonal elements of $A = A^{(1)}$: $a_{jj}^{(1)} = r_j^{(1)} - p_j^{(1)}$, $j = 1, \ldots, n$. This can be calculated accurately because $r^{(1)}$ is nonnegative and $p^{(1)}$ is nonpositive, and it requires n additions. We also calculate the vector $s^{(1)} = (s_1^{(1)}, \ldots, s_n^{(1)})^T$ with the sums of the off–diagonal entries of the columns of $A = A^{(1)}$ by means of (3.9), which can be performed accurately because all off–diagonal elements of A are nonpositive. It requires again $n(n-2)$ additions.

Then we choose a pivot index j_1 from (3.23) for $t = 1$, which uses the diagonal elements $h_j^{(1)} = a_{jj}^{(1)}$ of A and the sums $s_j^{(1)}$ of the off–diagonal entries of the columns. This requires $n-1$ comparisons. If the first pivot $\tilde{a}_{11}^{(1)} = 0$, then $\tilde{A}^{(1)} = A^{(2)}$. If $\tilde{a}_{11}^{(1)} \neq 0$, then we calculate (accurately) the multipliers l_{i1} (of Gaussian elimination) and quotients n_{1i} $(i > 1)$ by (3.13) with $t = 1$, and, if $\tilde{a}_{11}^{(1)} = 0$, then $l_{i1} := 0$ and $n_{1i} := 0$. Using this procedure to calculate the elements n_{ki} (at step k of Gaussian elimination) requires $n-k$ quotients. We now calculate the last $n-(k-1)$ components of the vector $r^{(k)}$ with the row sums of $A^{(k)}[k, \ldots, n]$ which (by (3.2)) are given by

$$r_j^{(k)} = \tilde{r}_j^{(k-1)} - l_{j,k-1}\tilde{r}_{k-1}^{(k-1)}, \quad j = k, \ldots, n.$$

This calculation is subtraction-free because $\tilde{r}^{(k-1)}$ is nonnegative and the multipliers are nonpositive. It requires the same cost as that of updating diagonal elements

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 24, pp. 153-167, October 2012

in Gaussian elimination. Then we calculate the last $n-(k-1)$ components of the vector $h^{(k)}$ by (3.19). It can be calculated accurately because $\tilde{h}^{(k-1)}$ is nonnegative and the multipliers are nonpositive. This computation requires $n-k$ additions and multiplications. The calculation of the last $n-(k-1)$ components of the vector $s^{(k)}$ by (3.20) can be performed accurately because $\tilde{s}^{(k-1)}$ is nonpositive and the multipliers are nonpositive. This computation requires $n-k$ additions and multiplications. Now we calculate the off–diagonal elements of $A^{(k)}[k,\dots,n]$, as in Gaussian elimination: $a_{ij}^{(k)} = \tilde{a}_{ij}^{(k-1)} - l_{i,k-1}\tilde{a}_{k-1,j}^{(k-1)}$ (observe that $\tilde{a}_{ij}^{(k-1)}$ and $-l_{i,k-1}\tilde{a}_{k-1,j}^{(k-1)}$ are both nonpositive). We chooose the k-th pivot index j_k from (3.23) for $t=k$, which uses the elements $h_j^{(k)}$ of A and the elements $s_j^{(k)}$. Choosing the k-th pivot requires $n-k$ comparisons. We calculate the k-th pivot $\tilde{a}_{kk}^{(k)} = a_{j_kj_k}^{(k)} = r_{j_k}^{(k)} - p_{j_k}^{(k)}$. This can be calculated accurately because $r_{j_k}^{(k)}$ is nonnegative and $p_{j_k}^{(k)}$ is nonpositive, which, in turn, can be calculated accurately as a sum of nonpositive numbers: see (3.3) for $i=j_k$. So, computing the k-th pivot requires $n-(k+1)$ additions.

As we have recalled above, all matrices $A^{(t)}[t,\dots,n]$ are again row diagonally dominant M–matrices. So, we can iterate the procedure described in the previous paragraph with the corresponding sequence of matrices

$$A^{(2)}[2,\dots,n],\dots,A^{(n-1)}[n-1,n],$$

until obtaining L and DU. From this last matrix we can compute accurately D and U. We increase the computational cost of Gaussian elimination without row or column exchanges with $(7n^2-11n+6)/2$ additions, $n(n-1)$ multiplications, $n(n-1)/2$ quotients and $n(n-1)/2$ comparisons. □

Remark 3.1. The method of the previous theorem has less computational cost than those of [4] (symmetric complete pivoting) and [13, Section 4] because it requires $\mathcal{O}(n^2)$ (instead of $\mathcal{O}(n^3)$) elementary operations beyond the cost of Gaussian elimination. The reason for the lower computational cost comes from the fact that the method of Theorem 3.1 does not require, for each $t>1$, the calculation of all diagonal elements $a_{jj}^{(t)}$ $(j \geq t)$ of the matrices $A^{(t)}[t,\dots,n]$ in order to choose the pivot $\tilde{a}_{tt}^{(t)}$. However, in the case of symmetric complete pivoting, Ye suggested in [14, p. 2202], that we can use the diagonal entries as computed by standard Gaussian elimination to determine the pivot and permutation and then compute the pivot $a_{tt}^{(t)}$. With this procedure, symmetric complete pivoting also requires $\mathcal{O}(n^2)$ elementary operations beyond the cost of Gaussian elimination, although the possible pivots are not then computed accurately for the choice.

Remark 3.2. Theorem 3.1 can be applied to any row diagonally dominant matrix $A=(a_{ij})_{1\leq i,j\leq n}$ satisfying

$$\operatorname{sign}(a_{ij}) \leq 0, \quad j \neq i, \qquad \operatorname{sign}(a_{ii}) \geq 0, \qquad i=1,\dots,n, \tag{3.24}$$

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 24, pp. 153-167, October 2012

given its off–diagonal entries and its vector r of row sums (see (2.5)) and so the method of [3] allows us to calculate accurately all its singular values. Let us observe that we can also apply the method of Theorem 3.1 (and so the method of [3] allows us to calculate accurately all its singular values) to any row diagonally dominant matrix A satisfying any of the following sign patterns:

$$\text{sign}(a_{ij}) = (-1)^{i+j+1}, \quad j \neq i, \qquad \text{sign}(a_{ii}) \geq 0, \qquad i = 1, \ldots, n, \tag{3.25}$$

$$\text{sign}(a_{ij}) \geq 0, \quad j \neq i, \qquad \text{sign}(a_{ii}) \leq 0, \qquad i = 1, \ldots, n. \tag{3.26}$$

$$\text{sign}(a_{ij}) = (-1)^{i+j}, \quad j \neq i, \qquad \text{sign}(a_{ij}) \leq 0, \qquad i = 1, \ldots, n, \tag{3.27}$$

assuming that we know its off–diagonal entries and the vector of row sums of its comparison matrix $\mathcal{M}(A)$. In fact, let us define the diagonal $n \times n$ matrix $J = \text{diag}(1, -1, \ldots, (-1)^{n-1})$ and observe that $J^{-1} = J$ and that, if A satisfies (3.25), then the matrix $J^{-1}AJ = JAJ = \mathcal{M}(A)$ satisfies (3.24), has the same singular values as A and we can calculate them with the method of [3] after obtaining the accurate LDU factorization of $\mathcal{M}(A)$ by the method of Theorem 3.1. Analogously, if A satisfies either (3.26) or (3.27), then we apply the procedure of Theorem 3.1 to $-A$ or to $J(-A)J$, respectively. Diagonally dominant matrices with arbitrary sign patterns were considered in [6] and [14], as commented in the introduction.

We have carried out numerical experiments with randomly generated diagonally dominant M–matrices. Although for (symmetric) complete pivoting there are no satisfactory bounds on the condition number of L as those of (2.4), we can conclude from the numerical experiments that the matrices L obtained using complete pivoting strategy in Gausssian elimination (method of [4]) are well conditioned in general. However, let us introduce a family of examples of matrices that shows that the condition number of the computed L, $\kappa_\infty(L)$, can be much worse using complete pivoting than with weak column diagonal dominance pivoting. Let us consider the $n \times n$ matrix ($n \geq 2$)

$$A_n = \begin{pmatrix} n-1 & -(n-1) & 0 & \cdots & \cdots & 0 \\ 0 & n & -1 & \cdots & -1 & -2 \\ \vdots & -(n-1) & n-1 & 0 & \cdots & 0 \\ \vdots & \vdots & 0 & \ddots & & \vdots \\ \vdots & \vdots & \vdots & & n-1 & 0 \\ 0 & -(n-1) & 0 & \cdots & 0 & n-1 \end{pmatrix}. \tag{3.28}$$

These matrices are diagonally dominant M–matrices. We have computed the LDU factorization of A_n by using Algorithm 1 and the method of [4] (with complete pivoing)

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 24, pp. 153-167, October 2012

and we shall compare the condition number of the matrices L_1 (Algorithm 1) and L_2 (complete pivoting).

It can be checked that in the $P_1 A_n P_1 = L_1 D_1 U_1$ decomposition obtained from the Algorithm 1, we have

$$L_1 = \begin{pmatrix} 1 & & & & & \\ 0 & 1 & & & & \\ \vdots & 0 & \ddots & & & \\ \vdots & \vdots & \ddots & 1 & & \\ \vdots & 0 & \cdots & 0 & 1 & \\ 0 & -\frac{1}{n-1} & \cdots & -\frac{1}{n-1} & -\frac{2}{n-1} & 1 \end{pmatrix}.$$

Thus, we can derive that $\|L_1\|_\infty = \|L_1^{-1}\|_\infty = 2$ for $n \geq 5$. The equivalent decomposition for complete pivoting gives the following matrix

$$L_2 = \begin{pmatrix} 1 & & & & & \\ -\frac{n-1}{n} & 1 & & & & \\ \vdots & 0 & 1 & & & \\ \vdots & \vdots & -\frac{1}{n-1} & 1 & & \\ \vdots & \vdots & \vdots & \ddots & \ddots & \\ -\frac{n-1}{n} & 0 & -\frac{1}{n-1} & \cdots & -\frac{1}{3} & 1 \end{pmatrix},$$

wich has $\|L_2\|_\infty = \frac{2n-1}{n} + \sum_{i=3}^{n-1} \frac{1}{i}$ and $\|L_2^{-1}\|_\infty = \frac{2n-1}{3}$.

Then we conclude that the condition number of L_1 is $\kappa_\infty(L_1) = 4$ for all $n \geq 5$ but the condition number of L_2 can be arbitrarily large as n increases.

We present in Table 3.1 the condition numbers of the matrices L_1 (second column) and L_2 (third column) of sizes $n = 10, 20, 30, 40$ and 50. In Table 3.2 we have the relative errors of the factors of the LDU decomposition of our example matrix A_n (with sizes $n = 10, 20, 30, 40$ and 50) computed with the method of [4], that is, with complete pivoting. The exact matrices are denoted by L_2, D_2, U_2 and the computed matrices are denoted by $\hat{L}_2, \hat{D}_2, \hat{U}_2$. Let us observe that there exists a small relative error for these matrices computed by the algorithm of [4]. We have also carried out analogous computations using Algorithm 1 and the computed factors L_1, D_1 and U_1 are exactly computed, that is, the relative error is 0.

4. Concluding remarks. Finally, we summarize some conclusions of this paper. It presents an efficient and subtraction–free implementation of a weak column diagonal dominance pivoting strategy of row diagonally dominant M–matrices for

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 24, pp. 153-167, October 2012

n	$\kappa_\infty(L_1)$	$\kappa_\infty(L_2)$
10	4	20.4501
20	4	51.9706
30	4	87.0903
40	4	124.5183
50	4	163.6538

TABLE 3.1
Condition numbers.

n	$\|\hat{L}_2 - L_2\|_2/\|L_2\|_2$	$\|\hat{D}_2 - D_2\|_2/\|D_2\|_2$	$\|\hat{U}_2 - U_2\|_2/\|U_2\|_2$
10	1.8922e-017	1.7764e-016	7.6823e-017
20	3.5440e-017	3.5527e-016	1.2123e-016
30	3.9756e-017	5.9212e-016	1.7154e-016
40	4.3490e-017	8.8818e-016	2.1188e-016
50	4.8376e-017	8.5265e-016	2.3833e-016

TABLE 3.2
Relative errors with complete pivoting.

computing their LDU decompositions, assuming that their off–diagonal entries and row sums are given. These decompositions can be performed with high relative accuracy and, in addition, the lower triangular factor L is column diagonally dominant and the upper triangular factor U is row diagonally dominant. So, both triangular factors are well conditioned and the lower triangular factor L can be considerably better conditioned than the lower triangular factor obtained with complete pivoting, as an example shows. For an $n \times n$ matrix, the computational cost of our method exceeds that of Gaussian elimination by at most $\mathcal{O}(n^2)$ elementary operations. The LDU factorization presented in this paper is a special case of rank revealing decomposition and so, it can be used for the accurate computation of the singular values through the algorithm presented in [3].

REFERENCES

[1] A.S. Alfa, J. Xue and Q. Ye. Entrywise perturbation theory for diagonally dominant M-matrices with applications. *Numer. Math.*, 90:401–414, 1999.

[2] A.S. Alfa, J. Xue and Q. Ye. Accurate computation of the smallest eigenvalue of a diagonally dominant M-matrix. *Math. Comp.*, 71:217–236, 2001.

[3] J. Demmel, M. Gu, S. Eisenstat, I. Slapnicar, K. Veselic and Z. Drmac. Computing the singular value decomposition with high relative accuracy. *Linear Algebra Appl.*, 299:21–80, 1999.

[4] J. Demmel and P. Koev. Accurate SVDs of weakly diagonally dominant M-matrices. *Numer. Math.*, 98:99–104, 2004.

Electronic Journal of Linear Algebra ISSN 1081-3810
A publication of the International Linear Algebra Society
Volume 24, pp. 153-167, October 2012

[5] J. Demmel and P. Koev. The accurate and efficient solution of a totally positive generalized Vandermonde linear system. *SIAM J. Matrix Anal. Appl.*, 27:142–152, 2005.

[6] F.M. Dopico and P. Koev. Perturbation theory for the LDU factorization and accurate computations for diagonally dominant matrices. *Numer. Math.*, 119:337–371, 2001.

[7] K. Fan. Note on M–matrices. *Quart. J. Math. Oxford Ser. (2)*, 11:43–49, 1961.

[8] M. Gasca and J.M. Peña. Total positivity and Neville elimination. *Linear Algebra Appl.*, 165:25–44, 1992.

[9] N.J. Higham. *Accuracy and Stability of Numerical Algorithms, second ed.*. SIAM, Philadelphia, PA, 2002.

[10] P. Koev. Accurate eigenvalues and SVDs of totally nonnegative matrices. *SIAM J. Matrix Anal. Appl.*, 27:1–23, 2005.

[11] J.M. Peña. Pivoting strategies leading to diagonal dominance by rows. *Numer. Math.*, 81:293–304, 1998.

[12] J.M. Peña. Scaled pivots and scaled partial pivoting strategies. *SIAM J. Numer. Anal.*, 41:1022–1031, 2003.

[13] J.M. Peña. LDU decompositions with L and U well conditioned. *Electronic Transactions of Numerical Analysis*, 18:198–208, 2004.

[14] Q. Ye. Computing singular values of diagonally dominant matrices to high relative accuracy. *Math. Comp.*. 77:2195–2230, 2008.