Orthogonal reduction of sparse matrices to upper triangular form using Householder transformations

Alan George
Esmond Ng

Department of Computer Science
University of Waterloo
Waterloo, Ontario, CANADA

CS-84-05
February 1984

Orthogonal reduction of sparse matrices to upper triangular form using Householder transformations+

Alan George

Esmond Ng

Department of Computer Science
University of Waterloo
Waterloo, Ontario, CANADA

Research report CS-84-05
February 1984

ABSTRACT

In this paper we consider the problem of predicting the fill-in that occurs in the QR-decomposition of sparse matrices using Householder transformations. We show that a static data structure can be used throughout the numerical computation, and that the Householder transformations can be saved explicitly in a compact format.

1. Introduction

Let A be an $n \times n$ nonsingular matrix. In this paper we consider the problem of reducing A to upper triangular form using orthogonal transformations, where A is large and sparse. That is, we construct an $n \times n$ orthogonal matrix Q so that

$$A = QR \quad ,$$

where R is $n \times n$ and upper triangular. Since it is well known that computing such a decomposition is numerically stable, the QR-decomposition is useful in various numerical computations, such as the solution of nonsingular systems of linear equations. However, very few implementations of the QR-decomposition exist for A when it is large and sparse. This is apparently due to the general belief that the orthogonal matrix Q and the intermediate matrices may be dense even though A is sparse, and also due to the lack of efficient techniques for exploiting the sparsity of the orthogonal matrix and the intermediate matrices.

One such implementation is due to George and Heath [2]. They make use of the fact that the upper triangular matrix R is (mathematically) the Cholesky factor of the symmetric positive definite matrix A^TA (apart from possible sign differences in some rows). Thus, assuming A^TA and its Cholesky factor are sparse, one can easily determine the structure of R and set up a data structure for R using techniques developed for solving sparse symmetric positive definite systems [3]. Then R can be computed using the *static* data structure by applying Givens rotations to the

+ Research was supported in part by the Canadian Natural Sciences and Engineering Research Council under Grant A8111.

rows of A one at a time. The Givens rotations are not saved in their implementation. However in some applications, it is desirable or necessary to have the orthogonal matrix Q available. One such context is the solution of several problems that have the same coefficient matrix A but different right hand side vectors. In this paper, we show that if we compute the decomposition using Householder transformations, then the nonzeros in the transformations and in the intermediate matrices can be stored in a static data structure allocated for the Cholesky factors L and L^T of the matrix A^TA. The ideas presented in this paper are similar to those used in the implementation of Gaussian elimination for sparse matrices using partial pivoting [4].

An outline of this paper is as follows. In Section 2, we present the main results which show that the structures of the transformations and the intermediate matrices obtained in the orthogonal reduction of A are contained in the structures of the Cholesky factors of A^TA. The effect of permuting the columns of A in the orthogonal reduction is considered in Section 3. In Section 4, the basic technique of the paper is extended to handle rectangular matrices. Finally, some concluding remarks are provided in Section 5.

2. Basic results

Let A be an $n \times n$ nonsingular matrix. The following notation will be used throughout our discussion. The (i,j)-element of the matrix A is denoted by a_{ij}. The set of indices of the nonzeros in A is denoted by $Nonz(A)$; that is,

$$Nonz(A) = \{ (i,j) \mid a_{ij} \neq 0 \} \quad .$$

The matrix A is said to have a *zero-free diagonal* if all its diagonal elements are nonzero.

Lemma 2.1 [1]

Let A be an $n \times n$ nonsingular matrix. Then there exists a permutation matrix P such that PA has a zero-free diagonal. □

For convenience, we assume in the following discussion that the rows of A have been permuted so that A has a zero-free diagonal. The next result is useful in deriving the main results.

Lemma 2.2

Suppose A is $n \times n$ and has a zero-free diagonal, and let B be an $n \times p$ matrix. Then

$$Nonz(B) \subseteq Nonz(AB) \quad . \; \square$$

We will also assume that accidental *structural* cancellation does not occur; that is, we assume that $Nonz(A+B) = Nonz(A) \bigcup Nonz(B)$, for any $n \times n$ matrices A and B.

Now let $A_0 = A$ and partition A_0 into

$$A_0 = \begin{pmatrix} \alpha_1 & y_1^T \\ x_1 & B_1 \end{pmatrix} \quad ,$$

where B_1 is $(n-1)\times(n-1)$, and x_1 and y_1 are vectors of appropriate dimensions. By assumption, $\alpha_1 \neq 0$ and B_1 has a zero-free diagonal. Assume $x_1 \neq 0$ and consider annihilating the nonzeros of x_1 using a Householder transformation H_1. (If $x_1 = 0$, then $H_1 = I$.) One way of constructing the Householder matrix H_1 is as follows. Define an n-vector w_1 by

$$w_1 = \begin{pmatrix} \alpha_1 + \sigma_1 \\ x_1 \end{pmatrix} \quad ,$$

where $\sigma^2 = \alpha_1^2 + x_1^T x_1$. Let $\pi_1 = \frac{1}{2} w_1^T w_1$. Then it is easy to verify that

$$H_1 = I - \frac{1}{\pi_1} w_1 w_1^T$$

is orthogonal and

$$H_1 \begin{pmatrix} \alpha_1 \\ x_1 \end{pmatrix} = \begin{pmatrix} -\sigma_1 \\ 0 \end{pmatrix} .$$

There are other ways of constructing H_1 (see [6]) and they differ essentially in the way the vector $\begin{pmatrix} \alpha_1 \\ x_1 \end{pmatrix}$ is scaled. Thus we can assume that in general the Householder matrix H_1 has the form

$$H_1 = I - \frac{1}{\pi_1} w_1 w_1^T \quad ,$$

where

$$w_1 = \begin{pmatrix} \beta_1 \\ u_1 \end{pmatrix} \quad ,$$

for some appropriate π_1, β_1 and u_1, with $\beta_1 \neq 0$ and $Nonz(u_1) = Nonz(x_1)$. Note that by storing the nonzeros of u_1 (and β_1 and π_1), one can save H_1 in a compact format.

Consider applying H_1 to A_0. Let

$$H_1 A_0 = \begin{pmatrix} -\sigma_1 & z_1^T \\ 0 & A_1 \end{pmatrix} \quad ,$$

where

$$\begin{pmatrix} z_1^T \\ A_1 \end{pmatrix} = H_1 \begin{pmatrix} y_1^T \\ B_1 \end{pmatrix} = (I - \frac{1}{\pi_1} w_1 w_1^T) \begin{pmatrix} y_1^T \\ B_1 \end{pmatrix} = \begin{pmatrix} y_1^T \\ B_1 \end{pmatrix} - \frac{1}{\pi_1} \begin{pmatrix} \beta_1 \\ u_1 \end{pmatrix} \begin{pmatrix} \beta_1 & u_1^T \end{pmatrix} \begin{pmatrix} y_1^T \\ B_1 \end{pmatrix}$$

Thus,

$$z_1 = y_1 - \frac{1}{\pi_1} \beta_1 (\beta_1 y_1 + B_1^T u_1) \quad ,$$

and

$$A_1 = B_1 - \frac{1}{\pi_1} u_1 (\beta_1 y_1^T + u_1^T B_1) \quad .$$

Since $\beta_1 \neq 0$, and if exact structural cancellation does not occur,

$$Nonz(z_1) = Nonz(y_1) \cup Nonz(B_1^T u_1) \quad ,$$

and

$$Nonz(A_1) = Nonz(B_1) \cup Nonz(u_1 y_1^T) \cup Nonz(u_1 u_1^T B_1) \quad .$$

Furthermore, since $Nonz(u_1)=Nonz(x_1)$, we obtain the following which we state as a lemma for future reference.

Lemma 2.3

(1) $Nonz(z_1)=Nonz(y_1)\bigcup Nonz(B_1^Tx_1)$.

(2) $Nonz(A_1)=Nonz(B_1)\bigcup Nonz(x_1y_1^T)\bigcup Nonz(x_1x_1^TB_1)$. □

Corollary 2.4

A_1 has a zero-free diagonal. □

Note that similar results holds if Givens rotations are used to annihilate the nonzeros in x_1. The effect of applying a Givens rotation to eliminate a nonzero, say a_{k1}, is to replace the first and the $k-th$ rows of A by a *linear combination* of those two rows. Consequently, after annihilating a_{k1}, the structures of rows 1 and k of A are the *union* of the structures of the original rows. Thus, after *all* the nonzeros in x_1 have been annihilated, the structure of row 1 of A will be the union of the structures of the first row of A and of those rows such that they have a nonzero in column 1. That is, $Nonz(z_1)$ will be given by

$$Nonz(z_1) = Nonz(y_1) \bigcup Nonz(x_1^TB_1) \quad .$$

Using similar arguments, it is easy to see that, in the worst case, the structure of the remaining $(n-1)\times(n-1)$ matrix A_1 will be given by

$$Nonz(A_1) = Nonz(B_1) \bigcup Nonz(x_1(y_1^T + x_1^TB_1)) \quad .$$

Thus Lemma 2.3 holds even if Givens rotations are used. Now for each nonzero in x_1, there will be one Givens rotation. In order to save these Givens rotations in the space provided by the nonzeros in x_1, we need to represent each of them by a single number using the scheme proposed by Stewart [7].

We now show that the structures of u_1, z_1 and A_1 are related to the structures of the matrices obtained after applying one step of Cholesky decomposition to the symmetric positive definite matrix $A_0^TA_0$. Note that

$$A_0^TA_0 = \begin{pmatrix} \alpha_1 & x_1^T \\ y_1 & B_1^T \end{pmatrix}\begin{pmatrix} \alpha_1 & y_1^T \\ x_1 & B_1 \end{pmatrix} = \begin{pmatrix} \alpha_1^2 + x_1^Tx_1 & \alpha_1y_1^T + x_1^TB_1 \\ \alpha_1y_1 + B_1^Tx_1 & B_1^TB_1 + y_1y_1^T \end{pmatrix} = \begin{pmatrix} \tau_1 & v_1^T \\ v_1 & E_1 \end{pmatrix} \quad .$$

Applying one step of Cholesky decomposition to $A_0^TA_0$, we obtain

$$A_0^TA_0 = \begin{pmatrix} \tau_1^{1/2} & 0 \\ v_1/\tau_1^{1/2} & I \end{pmatrix}\begin{pmatrix} 1 & 0 \\ 0 & F_1 \end{pmatrix}\begin{pmatrix} \tau_1^{1/2} & v_1^T/\tau_1^{1/2} \\ 0 & I \end{pmatrix} \quad ,$$

where

$$F_1 = E_1 - \frac{1}{\tau_1}v_1v_1^T \quad .$$

The first observation is that

$$Nonz(v_1) = Nonz(y_1) \bigcup Nonz(B_1^Tx_1) \quad ,$$

assuming again exact structural cancellation does not occur and also because $\alpha_1 \neq 0$. Since B_1 has a zero-free diagonal, it follows from Lemma 2.2 that

$$Nonz(x_1) \subseteq Nonz(B_1^T x_1) \quad ,$$

and hence

$$Nonz(u_1) = Nonz(x_1) \subseteq Nonz(B_1^T x_1) \subseteq Nonz(y_1) \cup Nonz(B_1^T x_1) = Nonz(v_1) \quad .$$

Moreover, from Lemma 2.3,

$$Nonz(z_1) = Nonz(y_1) \cup Nonz(B_1^T x_1) = Nonz(v_1) \quad .$$

Consider the matrix F_1.

$$F_1 = E_1 - \frac{1}{\tau_1} v_1 v_1^T = (B_1^T B_1 + y_1 y_1^T) - \frac{1}{\tau_1}(\alpha_1 y_1 + B_1^T x_1)(\alpha_1 y_1^T + x_1^T B_1) \quad .$$

If exact structural cancellation does not occur, then

$$Nonz(F_1) = Nonz(B_1^T B_1) \cup Nonz(y_1 y_1^T)$$
$$\cup Nonz(B_1^T x_1 y_1^T) \cup Nonz(y_1 x_1^T B_1) \cup Nonz(B_1^T x_1 x_1^T B_1) \quad .$$

Recall from Lemma 2.3 that

$$Nonz(A_1) = Nonz(B_1) \cup Nonz(x_1 y_1^T) \cup Nonz(x_1 x_1^T B_1) \quad .$$

Since B_1 has a zero-free diagonal, it follows from Lemma 2.2 that

$$Nonz(A_1) \subseteq Nonz(B_1^T B_1) \cup Nonz(B_1^T x_1 y_1^T) \cup Nonz(B_1^T x_1 x_1^T B_1) \subseteq Nonz(F_1) \quad .$$

Thus we have proved the following result.

Theorem 2.5

Assume exact structural cancellation does not occur. Then

(1) $Nonz(u_1) \subseteq Nonz(v_1)$.

(2) $Nonz(z_1) \subseteq Nonz(v_1)$.

(3) $Nonz(A_1) \subseteq Nonz(F_1)$. □

That is, the structures of u_1, z_1 and A_1 which are obtained when x_1 is annihilated by an Householder transformation are *contained* in those of the matrices obtained after applying one step of Cholesky decomposition to $A_0^T A_0$. The fact that A_0 has a zero-free diagonal plays an important role here. Some of the results above may not hold if A_0 does not have a zero-free diagonal. For example, it is easy to construct an example in which $Nonz(B_1) \not\subseteq Nonz(B_1^T B_1)$, where B_1 does not have a zero-free diagonal.

Now partition A_1 into

$$A_1 = \begin{pmatrix} \alpha_2 & y_2^T \\ x_2 & B_2 \end{pmatrix} \quad ,$$

and assume $x_2 \neq 0$. Consider annihilating the nonzeros of x_2 using an Householder transformation H_2. Let

$$H_2 = I - \frac{1}{\pi_2} w_2 w_2^T \quad ,$$

where $w_2 = \begin{pmatrix} \beta_2 \\ u_2 \end{pmatrix}$ with $Nonz(u_2) = Nonz(x_2)$. As before, π_2, β_2 and u_2 are chosen so that

$$H_2 \begin{pmatrix} \alpha_2 \\ x_2 \end{pmatrix} = \begin{pmatrix} -\sigma_2 \\ 0 \end{pmatrix} \quad ,$$

where $\sigma_2^2 = \alpha_2^2 + x_2^T x_2$. Suppose

$$H_2 A_1 = \begin{pmatrix} -\sigma_2 & z_2^T \\ 0 & A_2 \end{pmatrix} \quad .$$

By Corollary 2.3, A_1 has a zero-free diagonal and hence Theorem 2.5 applies again. That is, the structures of u_2, z_2 and A_2 must be contained in the structures of the matrices obtained by applying one step of Cholesky decomposition to $A_1^T A_1$.

Apparently the results obtained so far do not provide us with a mechanism to implement the orthogonal reduction of sparse matrices efficiently using Householder transformations since we now have to consider the Cholesky decomposition of $A_1^T A_1$. However, the next result takes care of this problem.

Lemma 2.6

Assuming exact cancellation does not occur,

$$Nonz(A_1^T A_1) = Nonz(F_1) \quad .$$

Proof:

Recall that

$$A_1 = B_1 - \frac{1}{\pi_1} u_1 (\beta_1 y_1^T + u_1^T B_1) \quad .$$

It is then straightforward to verify that

$$A_1^T A_1 = B_1^T B_1 + \frac{\beta_1^2 u_1^T u_1}{\pi_1^2} y_1 y_1^T + \left(\frac{u_1^T u_1}{\pi_1^2} - \frac{2}{\pi_1}\right) B_1^T u_1 u_1^T B_1 + \left(\frac{\beta_1 u_1^T u_1}{\pi_1^2} - \frac{\beta_1}{\pi_1}\right)(B_1^T u_1 y_1^T + y_1 u_1^T B_1) \quad .$$

Thus, assuming exact structural cancellation does not occur and assuming $\beta_1 \neq 0$,

$$\begin{aligned} Nonz(A_1^T A_1) &= Nonz(B_1^T B_1) \cup Nonz(y_1 y_1^T) \cup Nonz(B_1^T u_1 y_1^T) \cup Nonz(y_1 u_1^T B_1) \\ &\qquad \cup Nonz(B_1^T u_1 u_1^T B_1) \\ &= Nonz(B_1^T B_1) \cup Nonz(y_1 y_1^T) \cup Nonz(B_1^T x_1 y_1^T) \cup Nonz(y_1 x_1^T B_1) \end{aligned}$$

$$\cup \; Nonz(B_1^T x_1 x_1^T B_1) \quad ,$$

since $Nonz(u_1)=Nonz(x_1)$. Hence

$$Nonz(A_1^T A_1) \;=\; Nonz(F_1) \quad . \;\square$$

Corollary 2.7

The Cholesky factors of $A_1^T A_1$ and F_1 have identical nonzero structures, assuming exact structural cancellation does not occur. □

Lemma 2.6 and Corollary 2.7 are important since they say that we do not have to worry about the Cholesky decomposition of $A_1^T A_1$. We only have to consider the Cholesky decomposition of F_1. That is, suppose

$$F_1 \;=\; \begin{pmatrix} \tau_2 & v_2^T \\ v_2 & E_2 \end{pmatrix} \;=\; \begin{pmatrix} \tau_2^{1/2} & 0 \\ v_2/\tau_2^{1/2} & I \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & F_2 \end{pmatrix} \begin{pmatrix} \tau_2^{1/2} & v_2/\tau_2^{1/2} \\ 0 & I \end{pmatrix} \quad .$$

Then,

$$Nonz(u_2) \;\subseteq\; Nonz(v_2) \quad ,$$

$$Nonz(z_2) \;\subseteq\; Nonz(v_2) \quad ,$$

and

$$Nonz(A_2) \;\subseteq\; Nonz(F_2) \quad .$$

By applying the arguments above recursively to A_2 and F_2, one can obtain a result which is a generalization of Theorem 2.5. Before stating the result, we introduce more notation.

Let A be an $n \times n$ matrix with a zero-free diagonal and let $A_0=A$. Consider the sequence of matrices

$$\{\, A_0, A_1, A_2, \cdots, A_{n-1} \,\} \quad ,$$

generated as follows. For $k=1,2,\cdots,n-1$, partition A_{k-1} into

$$A_{k-1} \;=\; \begin{pmatrix} \alpha_k & y_k^T \\ x_k & B_k \end{pmatrix} \quad .$$

Assume $x_k \neq 0$ and construct an Householder transformation H_k so that

$$H_k A_{k-1} \;=\; \begin{pmatrix} -\sigma_k & z_k^T \\ 0 & A_k \end{pmatrix} \quad ,$$

where $\sigma_k^2=\alpha_k^2+x_k^T x_k$. Assume

$$H_k \;=\; I_{n-k+1} - \frac{1}{\pi_k} w_k w_k^T \quad ,$$

where $w_k = \begin{pmatrix} \beta_k \\ u_k \end{pmatrix}$ with $Nonz(u_k)=Nonz(x_k)$. Here I_j denotes the identity matrix of order j. It is easy to see that

$$A = Q_1 Q_2 \cdots Q_{n-2} Q_{n-1} \begin{pmatrix} \alpha_1 & z_1^T & & \\ & \alpha_2 & z_2^T & \\ & & \alpha_3 & z_3^T \\ & & & \cdot \\ & & & & \cdot \\ & & & & & \cdot \end{pmatrix} = QR \quad ,$$

where

$$Q_k = \begin{pmatrix} I_{k-1} & O \\ O & H_k \end{pmatrix} \quad , \qquad \text{for} \quad k=1,2,\cdots,n-1 \quad ,$$

and

$$Q = Q_1 Q_2 \cdots Q_{n-1} \quad .$$

Also consider the sequence of matrices

$$\{ F_0, F_1, F_2, \cdots, F_{n-1} \} \quad ,$$

which is defined as follows. Let $F_0 = A^T A$. For $k=1,2,\cdots,n-1$, partition F_{k-1} into

$$F_{k-1} = \begin{pmatrix} \tau_k & v_k^T \\ v_k & E_k \end{pmatrix} \quad .$$

Applying one step of Cholesky decomposition to F_{k-1} yields

$$F_{k-1} = \begin{pmatrix} \tau_k^{1/2} & 0 \\ v_k/\tau_k^{1/2} & I_{n-k} \end{pmatrix} \begin{pmatrix} I_k & O \\ O & F_k \end{pmatrix} \begin{pmatrix} \tau_k^{1/2} & v_k/\tau_k^{1/2} \\ 0 & I_{n-k} \end{pmatrix} \quad .$$

If we define L_k by

$$L_k = \begin{pmatrix} I_{k-1} & 0 & 0 \\ 0 & \tau_k^{1/2} & 0 \\ 0 & v_k/\tau_k^{1/2} & I_{n-k} \end{pmatrix} \quad , \qquad k=1,2,\cdots,n-1 \quad ,$$

and L_n by

$$L_n = \begin{pmatrix} I_{n-1} & O \\ O & F_{n-1} \end{pmatrix} = \begin{pmatrix} I_{n-1} & O \\ O & \tau_n^{1/2} \end{pmatrix} \quad .$$

Then it is clear that

$$A^T A = F_0 = L_1 L_2 \cdots L_{n-1} L_n L_n^T L_{n-1}^T \cdots L_2^T L_1^T = LL^T \quad ,$$

where $L = L_1 L_2 \cdots L_{n-1} L_n$. Moreover, because of the way v_k and F_k are constructed, we have

$$Nonz(F_k) \subseteq \bigcup_{k=1}^{n} Nonz(L_k + L_k^T) = Nonz(L + L^T) \quad .$$

The following result is a generalization of Theorem 2.5. Its proof is similar to that of Theorem 2.5 and hence is omitted.

Theorem 2.8

Assume exact structural cancellation does not occur. Then for $k=1,2,\cdots,n-1$,

(1) A_k has a zero-free diagonal,

(2) $Nonz(u_k) \subseteq Nonz(v_k)$,

(3) $Nonz(z_k) \subseteq Nonz(v_k)$, and

(4) $Nonz(A_k) \subseteq Nonz(F_k) \subseteq Nonz(L+L^T)$. □

Theorem 2.8 has an important implication. If A is sparse, then it says that the structures of the vectors u_k (which are the major components in the construction of Q) and the upper triangular matrix R are *all contained* in the structure of the Cholesky factors of A^TA. The crucial point is that if A^TA and its Cholesky factor are sparse, then it is possible to determine the structure of the Cholesky factor L of A^TA from that of A^TA efficiently. The reader is referred to [3] for details. Knowing the structure of L, one can set up an efficient data structure that exploits the sparsity of L. Now Theorem 2.8 simply implies that one can compute the orthogonal decomposition using Householder transformations in that *static data structure.* No dynamic storage allocation is necessary. Furthermore, the orthogonal matrix Q (in factored form) can be retained. This may be useful in some situations, for example, when the QR-decomposition of A has to be used several times.

Of course, the success of the approach relies on the fact that A^TA and its Cholesky factor are sparse if A is sparse. There are examples in which this may not be true; the matrices A^TA and its Cholesky factor may be dense even if A is sparse. Fortunately, the latter situation arises usually because there are a relatively small number of dense rows in A. Even though identifying these rows is a difficult problem, there are schemes which can handle dense rows in an efficient manner. See [4,5].

3. Effect of permuting the columns

Let P_c be an $n \times n$ permutation matrix and denote the QR-decomposition of AP_c by

$$AP_c = \hat{Q}_1\hat{Q}_2 \cdots \hat{Q}_{n-2}\hat{Q}_{n-1}\hat{R} \quad ,$$

where $\hat{Q}_k$ is an appropriate Householder transformation and $\hat{R}$ is an $n \times n$ upper triangular matrix. Our results in the previous section indicates that the structures of $\hat{R}$ and the vectors used in constructing $\hat{Q}_k$ are contained in the structure of the Cholesky factor $\hat{L}$ of the symmetric positive definite matrix $(AP_c)^T(AP_c)=P_c^TA^TAP_c$. If A^TA is sparse, it is well known that the structure and the sparsity of $\hat{L}$ depend not only on the structure of A^TA, but also on the choice of the permutation matrix P_c. Thus it is desirable to choose P_c so that $\hat{L}$ is as sparse as possible. Unfortunately, the problem of finding such a permutation has shown to be an NP-complete problem [8]. On the other hand, there are many reliable heuristic algorithms for finding P_c that yields a reasonably sparse $\hat{L}$. Examples include the nested dissection algorithm and the minimum degree algorithm. See [3] for a detailed discussion of the ordering problem in sparse Cholesky decomposition.

Note that post-multiplying A by P_c may change the zero-nonzero pattern of A. In particular the matrix AP_c may no longer have a zero-free diagonal (assuming A originally has

$$A = \begin{pmatrix} \times & \times & \times & & \\ & \times & & & \times \\ \times & & \times & & \\ & & & \times & \\ & \times & & \times & \times \end{pmatrix} \qquad P_c = \begin{pmatrix} & & & & 1 \\ & & & 1 & \\ & & 1 & & \\ & 1 & & & \\ 1 & & & & \end{pmatrix} \qquad AP_c = \begin{pmatrix} & & \times & \times & \times \\ \times & & & \times & \\ & & \times & & \times \\ & \times & & & \\ \times & \times & & \times & \end{pmatrix}$$

Figure 3.1: An example illustrating the fact that AP_c may not have a zero-free diagonal even though A has one.

one). This is illustrated by an example in Figure 3.1. In order to preserve the zero-free diagonal, we can apply P_c to A symmetrically. That is, instead of looking at AP_c, we consider $P_c^T AP_c$. It is a simple exercise to verify that $P_c^T AP_c$ has a zero-free diagonal for the matrix A given in Figure 3.1. Also note that pre-multiplying AP_c by P_c has no effect on the structure of $\hat{L}$ since

$$(P_c^T AP_c)^T (P_c^T AP_c) = P_c^T A^T AP_c = \hat{L}\hat{L}^T \quad .$$

Another approach which solves this problem is to find a column permutation P_c first. Then we find a row permutation P_r to make sure that $P_r(AP_c)$ has a zero-free diagonal. The main observation here is that the Cholesky factor of $(P_r AP_c)^T (P_r AP_c)$ is mathematically the same as that of $(AP_c)^T (AP_c)$.

4. Generalization to rectangular matrices

In some situations, such as the solution of sparse linear least squares problems, it may be necessary to reduce a rectangular matrix to upper trapezoidal form. The approach we described in Sections 2 and 3 can be modified to handle these cases. Let A be an $m \times n$ sparse matrix with $m \geq n$. We assume that A has full column rank. Partition A into

$$A = \begin{pmatrix} B \\ C \end{pmatrix} \quad ,$$

where B is $n \times n$ and C is $(m-n) \times n$. For simplicity, we also assume B has a zero-free diagonal.

Denote the orthogonal decomposition of A by

$$A = Q_1 Q_2 \cdots Q_n \begin{pmatrix} R \\ O \end{pmatrix} \quad ,$$

where Q_k is an $m \times m$ Householder matrix and R is an $n \times n$ upper triangular matrix. Suppose

$$Q_k = \begin{pmatrix} I_{k-1} & O \\ O & H_k \end{pmatrix} \quad ,$$

with

$$H_k = I_{m-k+1} - \frac{1}{\pi_k} \begin{pmatrix} \beta_k \\ u_k \end{pmatrix} \begin{pmatrix} \beta_k & u_k^T \end{pmatrix} \quad .$$

Here u_k is an $(m-k)$-vector. Note that the decomposition is equivalent to performing the first n

steps in the orthogonal reduction of the $m \times m$ matrix $\overline{A}$:

$$\overline{A} = \begin{pmatrix} B & O \\ C & I \end{pmatrix} .$$

Consider the matrix $\overline{A}^T\overline{A}$.

$$\overline{A}^T\overline{A} = \begin{pmatrix} B^T & C^T \\ O & I \end{pmatrix}\begin{pmatrix} B & O \\ C & I \end{pmatrix} = \begin{pmatrix} B^TB + C^TC & C^T \\ C & I \end{pmatrix} = \begin{pmatrix} D & C^T \\ C & I \end{pmatrix} .$$

Applying the first n steps of the Cholesky decomposition to $\overline{A}^T\overline{A}$ yields

$$\overline{A}^T\overline{A} = \begin{pmatrix} D & C^T \\ C & I \end{pmatrix} = \begin{pmatrix} L & O \\ W & I \end{pmatrix}\begin{pmatrix} I & O \\ O & F \end{pmatrix}\begin{pmatrix} L^T & W^T \\ O & I \end{pmatrix} ,$$

where

$$LL^T = D = B^TB + C^TC ,$$

and

$$W = CL^{-1} .$$

Since $\overline{A}$ has a zero-free diagonal, the results in Section 2 apply. That is, the structure of u_k must be contained in the structure of the $k-th$ column of the matrix $\begin{pmatrix} L \\ W \end{pmatrix}$. Similarly, the structure of R must be contained in the structure of L^T. Thus, one way to implement the reduction of A is as follows.

(1) Determine the structure of $M = \overline{A}^T\overline{A}$.

(2) Perform the first n steps of symbolic Cholesky factorization to M, and determine the structures of L and $W = CL^{-1}$. Set up a data structure that exploits the sparsity of L^T and $\begin{pmatrix} L \\ W \end{pmatrix}$.

(3) Reduce the matrix A to upper trapezoidal form using Householder transformations, storing R and u_k's in the static data structure determined in Step 2.

Note that we only want to reduce $\begin{pmatrix} B \\ C \end{pmatrix}$ to upper trapezoidal form. Thus we do not want to worry about the last $(m-n)$ columns in $\overline{A}$. In other words, if we want to permute the columns of $\overline{A}$ so as to obtain a sparse Cholesky factorization, we should only permute the first n columns of $\overline{A}$.

5. Conclusion

We have shown in this paper that when a sparse matrix A is reduced to upper triangular form using Householder transformations, the structures of the transformations, the intermediate matrices and the final upper triangular matrix are contained in the structure of the Cholesky factors of A^TA. These results have an important practical implication. It is well known that the structure of the Cholesky factor of a sparse symmetric positive definite matrix B can be determined efficiently from the structure of B. Thus, by analyzing the structure of A^TA, we can determine the structure of the Cholesky factors L and L^T of A^TA, and can set up a data structure for L and L^T. Then we can perform the orthogonal reduction of A using this static data structure. This idea has been extended to handle the case in which A is rectangular.

Efficient implementation of the ideas described in this paper is currently under investigation.

6. References

[1] I.S. Duff, "Analysis of sparse systems", D.Phil. Thesis, Oxford University (1972).

[2] J.A. George and M.T. Heath, "Solution of sparse linear least squares problems using Givens rotations", *Linear Algebra and its Appl.,* **34** (1980), pp. 69-83.

[3] J.A. George and J.W.H. Liu, *Computer solution of large sparse positive definite systems,* Prentice-Hall Inc., Englewood Cliffs, N.J. (1981).

[4] J.A. George and E.G.Y. Ng, "An implementation of Gaussian elimination with partial pivoting for sparse systems", to appear in SIAM J. Sci. Stat. Comput. (1984).

[5] M.T. Heath, "Some extensions of an algorithm for sparse linear least squares problems", *SIAM J. Sci. Stat. Comput.,* **3** (1982), pp. 223-237.

[6] G.W. Stewart, *Introduction to matrix computations,* Academic Press, New York (1973).

[7] G.W. Stewart, "The economical storage of plane rotations", *Numer. Math.,* **25** (1976), pp. 137-138.

[8] M. Yannakakis, "Computing the minimum fill-in is NP-complete", *SIAM J. Alg. Disc. Meth.,* **2** (1981), pp. 77-79.