

Predicting Fill for Sparse Orthogonal Factorization

THOMAS F. COLEMAN, ANDERS EDENBRANDT, AND JOHN R. GILBERT

Cornell University, Ithaca, New York

Abstract. In solving large sparse linear least squares problems $Ax \approx b$, several different numeric methods involve computing the same upper triangular factor R of A . It is of interest to be able to compute the nonzero structure of R , given only the structure of A . The solution to this problem comes from the theory of matchings in bipartite graphs. The structure of A is modeled with a bipartite graph, and it is shown how the rows and columns of A can be rearranged into a structure from which the structure of its upper triangular factor can be correctly computed. Also, a new method for solving sparse least squares problems, called block back-substitution, is presented. This method assures that no unnecessary space is allocated for fill, and that no unnecessary space is needed for intermediate fill.

Categories and Subject Descriptors: F.2.1 [Analysis of Algorithms and Problem Complexity]: Numerical Algorithms and Problems—*computations on matrices*; G.1.2 [Numerical Analysis]: Approximation—*least squares approximation*; G.1.3 [Numerical Analysis]: Numerical Linear Algebra—*linear systems; sparse and very large systems*; G.1.6 [Numerical Analysis]: Optimization—*least squares methods*; G.2.2 [Discrete Mathematics]: Graph Theory—*graph algorithms*

General Terms: Algorithms, Theory, Verification

Additional Key Words and Phrases: Bipartite graphs, block triangular form, fill, Givens rotations, least squares, matching, orthogonal factorization, sparse matrices

1. Introduction

Consider solving a system of linear equations

$$Ax \approx b,$$

where A is an $m \times n$, $m \geq n$, large, sparse matrix. In the following we always assume that A has full column rank. When A is rectangular, the system is overdetermined, and we seek the least squares solution—that is, the solution that minimizes

$$\|Ax - b\|_2.$$

There are several different methods for finding this solution [19]. We take a brief look at two of them, *orthogonalization* and the *normal equations* methods.

The work of T. F. Coleman was supported in part by the Applied Mathematical Sciences Research Program (KC-04-02) of the Office of Energy Research of the U.S. Department of Energy under contract DE-AC02-83ER13069. The work of A. Edenbrandt was supported in part by an IBM graduate student fellowship. The work of J. R. Gilbert was supported in part by National Science Foundation grant MCS 82-02948.

Authors' present addresses: T. F. Coleman and J. R. Gilbert, Department of Computer Science, Cornell University, 405 Upson Hall, Ithaca, NY 14853; A. Edenbrandt, Department of Math and Computer Science, Dartmouth College, Bradley Hall, Hanover, NH 03755.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1986 ACM 0004-5411/86/0700-0517 \$00.75




FIGURE 3





FIGURE 4

The second method we look at is the George–Heath algorithm [11]. These authors use the normal equations to predict the structure of R and then do the numeric computation by orthogonalization. The first (structural) phase of the algorithm can be described as follows:

- (1) Starting with the structure of A , compute the structure of $A^T A$.
- (2) From the structure of $A^T A$, compute the structure of L^T —using a symbolic Gaussian elimination algorithm.

George and Heath proved that the structure that results from the algorithm above will have room for all the nonzeros of R .

LEMMA 1.1 [11]. *The structure of R as predicted by symbolic factorization of $A^T A \supseteq$ the structure of R as predicted by the Local Givens Rule \supseteq the structure of R .*

COROLLARY 1.2. *If the structure of $R =$ the structure predicted by symbolic factorization of $A^T A$, then the structure of $R =$ the structure predicted by the Local Givens Rule.*

However, this method also may be too generous in allocating space for nonzeros in R , as is shown in the example in Figure 4. Here A is already in upper triangular form, so $R = A$. But, since A has one full row, $A^T A$ is full and R is predicted to be full. (By Lemma 1.1, symbolic factorization will also be too generous with the example in Figure 3.)

We have thus seen that both of the previously suggested methods for computing the structure of R may predict too much fill. It should be emphasized that for both of the methods described above the occurrence of this “bogus fill” is purely a consequence of the nonzero structure of the matrix, and does not depend on the actual numeric values of the nonzero entries.

Instead of coming up with a new rule for computing fill, we have asked the following questions: First, can we exhibit a class of matrix structures for which symbolic factorization of $A^T A$ gives the correct result? Second, can we reorder the columns and rows of A so that the Local Givens Rule is guaranteed to give the correct result when applied to this rearranged matrix? We answer these questions in the affirmative.

Since the nonzero structure of a matrix is a combinatoric rather than a numeric property, it is convenient to represent this structure by a graph—in our case a bipartite graph—and to model the operations performed on the matrix with operations performed on the vertices and edges of the graph. In Section 2 of this paper we give some basic graph-theoretic definitions, and introduce a structural concept called *the strong Hall property*. In Section 3 we analyze the computation of the structure of R using symbolic factorization of $A^T A$, and show that this gives the correct result if A has the strong Hall property. This fact can be exploited to provide an algorithm for computing the structure of R using the Local Givens Rule for an arbitrary matrix A (appropriately reordered), as is shown in Section 4. We also give an algorithm for solving the original system of equations without using the whole structure of R . Finally, Section 5 contains a summary of our results.

2. Definitions and Notation

2.1 BASIC GRAPH-THEORETIC NOTATION. An *undirected graph*, $G = (V, E)$, consists of a set V of *vertices* and a set E of *edges*. An *edge* (v, w) is an unordered pair of distinct vertices. If (v, w) is an edge, the vertices v and w are *adjacent*. The edge (v, w) is *incident* on vertices v and w , which are its *endpoints*. The *degree* of a vertex, written $\deg(v)$, is the number of edges incident on v . Given a vertex v , the set $\Gamma(v) = \{x : (v, x) \in E\}$ is called the *adjacency set* of v . Similarly, for a set of vertices, $X \subseteq V$, we say that $\Gamma(X) = \{y : (x, y) \in E \text{ for some } x \in X\}$ is the *adjacency set* of X .

A *matching* in a graph G is a set of edges, $F \subseteq E$, such that no two edges in F have an endpoint in common. A vertex that is the endpoint of some edge in F is *covered*; otherwise, it is *exposed*. If every vertex is covered by F , then F is a *perfect matching*.

A *path* between v and w in G is a sequence of distinct edges $(v_0, v_1), (v_1, v_2), \dots, (v_{k-1}, v_k)$, where $v = v_0$ and $w = v_k$ and all the vertices are distinct except possibly $v_0 = v_k$. If F is a matching in G , an *alternating path* is a path in which every other edge is in F .

A *subgraph* $G' = (V', E')$, of $G = (V, E)$ is a graph where $V' \subseteq V$ and $E' \subseteq E$. If $V' = V$, then G' is a *spanning subgraph*. The set of edges of G with both endpoints in V' is written as $E(V')$. If $E' = E(V')$, then G' is the subgraph of G *induced* by the vertex set V' .

For an $n \times n$ symmetric matrix A that has only nonzeros on its diagonal, we can use an undirected graph $G(A)$ to represent the zero/nonzero structure of A . The graph $G(A)$ has n vertices and (v_i, v_j) is an edge if and only if $A_{ij} \neq 0$. The changes in the structure of A during Gaussian elimination can be modeled as a process working on the graph: when column i is zeroed out in the matrix, we mark v_i "old" and add enough new edges to turn all of v_i 's unmarked neighbors into a complete graph. The final result of this process is called the *filled graph*, G^* , and we refer to this process as *symbolic Gaussian elimination*. The following characterization of fill, in terms of an undirected graph, is very useful.

LEMMA 2.1 [24]. *The edge (v, w) is an edge in the filled graph G^* if and only if there is a path in G from v to w going only through vertices marked earlier than both v and w .*

2.2 BIPARTITE GRAPHS AND THE HALL PROPERTY. An undirected graph whose vertex set can be divided into two disjoint sets V_1 and V_2 , such that every edge in E has one endpoint in V_1 and the other in V_2 , is called a *bipartite graph*. This graph is sometimes written $H = (V_1, V_2, E)$ to stress the partition of V .

Let $H = (V, W, E)$ be a bipartite graph with $|V| \geq |W|$. If, for every set of vertices $X \subseteq W$, X is adjacent to at least $|X|$ vertices in V , then we say that the bipartite graph H has the *Hall property*. If every proper subset X is adjacent to more than $|X|$ vertices, we say that H has the *strong Hall property*; or, for more precise definitions:

Definition 2.2. A bipartite graph $H = (V, W, E)$ with $|V| \geq |W|$ has the *Hall property* if

$$\forall X \subseteq W: |\Gamma(X)| \geq |X|.$$

Definition 2.3. A bipartite graph $H = (V, W, E)$ with $|V| \geq |W|$ has the *strong Hall property* if

$$\forall X \subset W, X \neq \emptyset: |\Gamma(X)| \geq |X|.$$

We list some basic and useful facts about matchings in bipartite graphs.

LEMMA 2.4 [17]. *Let $H = (V, W, E)$ be a bipartite graph. There is a matching in H of cardinality $|W|$ if and only if H has the Hall property.*

COROLLARY 2.5. *Let $H = (V, W, E)$ be a bipartite graph with $|V| = |W|$. Then H has a perfect matching if and only if H has the Hall property.*

A *directed graph* $D = (V, F)$ is a graph whose edges are ordered pairs of distinct vertices, $\langle v_i, v_j \rangle$ called *arcs*. A *directed path* from v to w in D is a sequence of distinct arcs $\langle v_0, v_1 \rangle, \langle v_1, v_2 \rangle, \dots, \langle v_{k-1}, v_k \rangle$, where $v = v_0$ and $w = v_k$ and all the vertices are distinct except possibly $v_0 = v_k$. If for every two vertices v and w in D there is a directed path from v to w , then D is *strongly connected*.

With any directed graph D of n vertices we can associate a bipartite graph H_D as follows: $H_D = (X, Y, E)$, with $|X| = |Y| = n$, and (x_i, y_j) is an edge in E if $\langle v_i, v_j \rangle$ is an arc in F . In addition, all the edges (x_i, y_i) for $1 \leq i \leq n$ are in E .

The following observation, due to Johnson et al. [22], is also proved in [21].

LEMMA 2.6 [22]. *D is strongly connected if and only if H_D has the strong Hall property.*

Note that this correspondence between the strong Hall property and strong connectedness is well defined only when the bipartite graph is "square," that is, when it has the same number of vertices in each part. Indeed, it is possible for a rectangular bipartite graph that has the strong Hall property not to be connected, as in Figure 5. However, every bipartite graph has a canonical decomposition into subgraphs with the strong Hall property. This topic is pursued further in Section 4.

2.3 REPRESENTING THE STRUCTURE OF A MATRIX WITH A BIPARTITE GRAPH. Let A be an $m \times n$ real matrix, with $m \geq n$, and full column rank. We represent

FIG. 5. A disconnected bipartite graph with strong Hall property.





FIG. 6. A product-bipartite graph where $W_1 = \{w_1, w_2, w_3, w_4\}$ and $W_2 = \{w_1, w_2, w_3, w_5\}$.

the zero/nonzero structure of A with a graph $H(A)$. $H = (V, W, E)$ is a bipartite graph with $V = \{v_1, \dots, v_m\}$, $W = \{w_1, \dots, w_n\}$, and $(v_i, w_j) \in E$ if and only if $A_{ij} \neq 0$. We refer to the vertices of W as *column vertices*, and to the vertices of V as *row vertices*. Note that since A has full column rank, H has a matching that covers all the column vertices.

Given any two sets $W_1, W_2 \subseteq W$ of column vertices, we can define a new bipartite graph $H' = (X, Y, S)$ where $|X| = |W_1|$, $|Y| = |W_2|$, and $(x_i, y_j) \in S$ if and only if $w_i \in W_1$ and $w_j \in W_2$ and $w_j \in \Gamma(\Gamma(w_i))$. We write this as $H' = \Phi(H, W_1, W_2)$. See Figure 6 for an example.

If U_1 is the matrix containing only the columns of A represented by W_1 , and U_2 is the matrix containing the columns of W_2 , then $H' = H(U_1^T U_2)$, provided that no two columns that have some nonzeros in the same rows are orthogonal. This is true, for example, if all the entries are nonnegative. In particular, under this assumption $\Phi(H(A), W, W) = H(A^T A)$.

By $N(A)$ we denote the set of $m \times n$ real matrices B such that $B_{ij} = 0$ if $A_{ij} = 0$.

For a matrix A , we define the *structural rank* of A to be the size of a maximum matching in $H(A)$. The *numeric rank* of A is the number of linearly independent columns of A . The structural rank of A is at least as large as its numeric rank.

To make it easier to remember all this notation, we summarize it in the following table.

Symbol	Denotes
$G = (V, E)$	Undirected graph
$H = (V, W, E)$	Bipartite graph
$D = (V, F)$	Directed graph
$\Gamma(X)$	Adjacency set of X
$\Phi(H, W_1, W_2)$	Product graph
$\langle H, r, p, q \rangle$	The special product graph $\Phi(H, W_1, W_2)$ where $W_1 = \{w_1, \dots, w_r\} \cup \{w_p\}$ and $W_2 = \{w_1, \dots, w_r\} \cup \{w_q\}$ (see Section 3.2)
A or M	Matrix
$M = \langle A, r, p, q \rangle$	Special submatrix of $A^T A$ (See Section 3.1)
$N(A)$	Matrices whose structure fits into A 's structure
$H(A)$	Bipartite graph that represents the structure of A

3. Correct Prediction of Fill

Let α_N be a numeric algorithm that takes an $m \times n$ matrix A as input and produces a matrix R as output; and let α_S be a corresponding structural algorithm that takes a structure H_A as input and outputs a structure H_R . We say that α_S *correctly models* α_N (or simply that α_S is *correct*) if the following two conditions hold. For all matrices A ,

- (1) $H(R) \subseteq H_R$, and
- (2) there is a matrix $A' \in N(A)$ such that $H(R') = H_R$.

We assume throughout that A has full column rank. Then $A^T A$ is positive definite, and Cholesky factorization can be computed using Gaussian elimination with pivots from the main diagonal. There are well-known and efficient algorithms for determining the structure of the matrix that results from Gaussian elimination [12]. However, any such algorithm produces the correct structure only under the assumption that no cancellation occurs during the Gaussian elimination.

It is worthwhile at this point to say what we mean by cancellation. Gaussian elimination (for our purposes) is an algorithm that zeros the subdiagonal elements of one column at a time, meanwhile replacing some of the elements in other columns to its right. Cancellation occurs when a nonzero is replaced by a zero outside the column being eliminated.

There are two kinds of cancellation. *Lucky cancellation* is cancellation that occurs only for certain values of the nonzeros in the given matrix. We shall disregard lucky cancellation because it cannot be predicted from the structure alone. *Essential cancellation* occurs no matter what the nonzero values are.

Brayton et al. [2] show that if the structure of a square matrix M with nonzero diagonal is given, then there is no essential cancellation when it is factored by Gaussian elimination. (In their terminology, Gaussian elimination is a "minimal algorithm" for such matrices. They point out that Gaussian elimination is not necessarily a minimal algorithm for matrices with zeros on the diagonal.) It follows that, for the class of positive definite matrices, symbolic factorization correctly models Gaussian elimination. Here, however, we begin not with the structure of $A^T A$, the matrix whose Cholesky factor we want, but with the structure of A . The




FIG. 7. A special submatrix.

entries of $A^T A$ are not independent: the structure of A may be such that cancellation will occur when we do Gaussian elimination on $A^T A$, regardless of the values of the nonzeros in A . This essential cancellation is the reason why symbolic factorization of $A^T A$ may predict too much fill. In the following we show that, provided A has the strong Hall property, no essential cancellation will occur when we perform Gaussian elimination on $A^T A$. From this we can conclude that symbolic factorization of $A^T A$ gives the correct result for matrices that have the strong Hall property.

Our way to the main theorem of this section is via three important lemmas. The first of these lemmas says that we need only consider certain special submatrices of $A^T A$. The second lemma talks about bipartite graphs, and shows how we can use the strong Hall property to make sure that each special submatrix has the property we desire. The third lemma tells us that it is enough to look at just one special submatrix at a time, rather than all of them simultaneously.

3.1 A SPECIAL SORT OF SUBMATRIX. The following lemma shows that if cancellation occurs at position (p, q) in $A^T A$ during Gaussian elimination, then this position is in the lower right-hand corner of a singular submatrix of $A^T A$ of a special form described below.

Consider the square matrix M consisting of the principal $r \times r$ submatrix of $A^T A$, plus the first r entries of row p and column q , plus entry (p, q) , where $1 \leq r < p \leq q \leq n$. See Figure 7. As a shorthand notation we write $M = \langle A, r, p, q \rangle$.

LEMMA 3.1. *If cancellation occurs in $A^T A$ at entry (p, q) during the r th stage of the Gaussian elimination, where row r is used to zero out entry (p, r) , then $M = \langle A, r, p, q \rangle$ has full structural rank and is singular.*

PROOF. To show that M has full structural rank, let $C = H(M)$ be the bipartite graph that represents the structure of M , and let $H = (X, Y, S) = H(A^T A)$ be the bipartite graph that represents the structure of $A^T A$. Clearly, C is a subgraph of H .

Since $A^T A$ is a symmetric matrix with a nonzero diagonal, $F = \{(x_1, y_1), \dots, (x_n, y_n)\}$ is a perfect matching in H . By Lemma 2.1 it then follows that position (p, q) in $A^T A$ fills in if and only if there is an alternating path in H from vertex x_p to vertex y_q going only through lower numbered vertices. Hence, if we have cancellation at position (p, q) , such a path exists, and, by interchanging the matching

and nonmatching edges of this path, we can construct a perfect matching in C . Any vertex x_i that does not lie on the alternating path will be matched with y_i .

To see that M is singular, note that performing Gaussian elimination on a square matrix does not change the value of its determinant; and since the resulting matrix is upper triangular and has one diagonal entry equal to zero, it must be singular. \square

3.2 A STRUCTURAL LEMMA. In this section we show that if A 's bipartite graph has the strong Hall property, then we can make any one of the special submatrices of $A^T A$ nonsingular by choosing appropriate values for the nonzeros of A .

As before, $H = (V, W, E)$ is a bipartite graph with the strong Hall property. Let $W_1 = \{w_1, \dots, w_r\} \cup \{w_p\}$ and $W_2 = \{w_1, \dots, w_r\} \cup \{w_q\}$. Let $C = \Phi(H, W_1, W_2)$. We also use the shorthand notation $C = \langle H, r, p, q \rangle$ for this graph.

Now, we are given the structure of a matrix, that is, a bipartite graph H . We want to show that provided that H has the strong Hall property, we can find a matrix A that fits into the given structure, and is such that for each subgraph $C = \langle H, r, p, q \rangle$ that has a perfect matching, the corresponding submatrix $M = \langle A, r, p, q \rangle$ is nonsingular. If we can show this, then by Lemma 3.1 we can say that there will be no essential cancellation during Gaussian elimination in $A^T A$.

We shall take on a somewhat easier task, though, namely, showing that for any single subgraph $C = \langle H, r, p, q \rangle$ that has a perfect matching, we can find a matrix A that fits into H and has $M = \langle A, r, p, q \rangle$ nonsingular.

Each nonzero term in the determinant of a square matrix corresponds to a different perfect matching in its bipartite graph. Hence, if C has exactly one perfect matching, then the determinant is nonzero—that is, M is nonsingular. On the other hand, if C has no perfect matching at all, this means that all terms in the determinant of M are zero, and the matrix is singular.

LEMMA 3.2. *Let $H = (V, W, E)$ be a bipartite graph, and let W_1 and W_2 be subsets of W such that $|W_1| = |W_2| = k$, and $|W_1 \cap W_2| \geq k - 1$. If H has the strong Hall property and if $C = \Phi(H, W_1, W_2)$ has a perfect matching, then there is a spanning subgraph \hat{H} of H such that $\hat{C} = \Phi(\hat{H}, W_1, W_2)$ has exactly one perfect matching.*

PROOF. Let $C = (S, T, U)$, with $S = \{s_1, \dots, s_k\}$ and $T = \{t_1, \dots, t_k\}$. It is convenient to distinguish between two cases:

Case 1. $|W_1 \cap W_2| = k$, that is, $W_1 = W_2$. Let $H_1 = (V, W_1, E)$. Since H_1 has the Hall property, we know there is a matching F_1 in H_1 that covers W_1 . Let $\hat{H} = (V, W, F_1)$. It is easy to see that any vertex s_i in \hat{C} is adjacent only to t_i . Therefore, \hat{C} has precisely one perfect matching, namely, $\{(s_1, t_1), \dots, (s_k, t_k)\}$.

Case 2. $|W_1 \cap W_2| = k - 1$, that is, there is a vertex $w_p \in W_1$ such that $w_p \notin W_2$, and a vertex $w_q \in W_2$ such that $w_q \notin W_1$. Without loss of generality, we may assume that $W_1 = \{w_1, \dots, w_{k-1}\} \cup \{w_p\}$, $W_2 = \{w_1, \dots, w_{k-1}\} \cup \{w_q\}$. For the proof of this second case, we need two more lemmas.

The following lemma is due to Ford and Fulkerson.

LEMMA 3.3 [9, Corollary 10.9]. *Let $H_1 = (V, W_1, E_1)$ and $H_2 = (V, W_2, E_2)$ be two bipartite graphs with one vertex set in common, and $|W_1| = |W_2| = k$. There is a set $P \subseteq V$ of k vertices, with a matching F_1 in H_1 covering P , and a matching F_2 in H_2 also covering P , if and only if*

$$\forall X \subseteq W_1 \forall Y \subseteq W_2: |X| + |Y| \leq k + |\Gamma(X) \cap \Gamma(Y)|. \quad (1)$$

LEMMA 3.4. *Let $H_1 = (V, W_1, E(V \cup W_1))$ and $H_2 = (V, W_2, E(V \cup W_2))$. If $C = \Phi(H, W_1, W_2)$ has a perfect matching, then there is a matching F_1 in H_1 and a matching F_2 in H_2 , both of cardinality k and both covering the same set $V' \subseteq V$, where $|V'| = k$.*

PROOF. By Lemma 3.3, it is enough to show that condition (1) holds for H_1 and H_2 .

Let $X \subseteq W_1$ and $Y \subseteq W_2$. If $|X| + |Y| \leq k$, then the inequality in (1) is immediate. Assume that $|X| + |Y| = k + \delta$, where $\delta \geq 1$. We must prove that $|\Gamma(X) \cap \Gamma(Y)| \geq \delta$. Let $Z = X \cap Y$. Clearly, $\Gamma(Z) \subseteq \Gamma(X) \cap \Gamma(Y)$.

Since X and Y are chosen from sets of size k that have $k - 1$ elements in common, we can conclude that $|Z| \geq \delta - 1$. We consider separately the cases $Z = \emptyset$ and $Z \neq \emptyset$.

If $Z = \emptyset$, then $\delta = 1$. Let $\Gamma_C(X)$ denote the adjacency set of X in the graph C . By assumption C has a perfect matching, which by Corollary 2.5 implies that $|\Gamma_C(X)| = |\Gamma(\Gamma(X)) \cap W_2| \geq |X|$. Since $\Gamma_C(X) \subseteq W_2$ and

$$|\Gamma_C(X)| + |Y| \geq |X| + |Y| = k + 1,$$

$\Gamma_C(X)$ and Y must have at least one vertex in common. That is, there must be at least one vertex $w \in Y$ such that $w \in \Gamma_C(X)$, and thus

$$|\Gamma(X) \cap \Gamma(Y)| \geq 1 = \delta.$$

Next, assume that $Z \neq \emptyset$. By the strong Hall property $|\Gamma(Z)| > |Z| \geq \delta - 1$. Then

$$|\Gamma(X) \cap \Gamma(Y)| \geq |\Gamma(Z)| \geq \delta.$$

We have thus shown that condition (1) is satisfied, and Lemma 3.4 now follows from Lemma 3.3. \square

We continue case 2 of the proof of Lemma 3.2.

Among the matchings F_1 in H_1 and F_2 in H_2 that satisfy the conclusion of Lemma 3.4, choose that pair F_1, F_2 that maximizes $|F_1 \cap F_2|$. Let $\hat{H} = (V, W, F_1 \cup F_2)$ be the graph with the vertices of the original graph H and the edges of the two matchings F_1 and F_2 . We show that $\hat{C} = \Phi(\hat{H}, W_1, W_2)$ has exactly one perfect matching. In fact, we see that \hat{C} consists of a path of odd length plus some isolated edges.

First note that every vertex in \hat{H} has degree at most 2, and that $\deg(w_p) = \deg(w_q) = 1$, since w_p and w_q are only covered by one matching each.

Let s_i be any vertex in \hat{C} . If $\deg(w_i) = 1$ and $i \neq p$ and $i \neq q$, then w_i was matched to the same vertex v_j in both F_1 and F_2 . This implies that $\deg(v_j) = 1$ in \hat{H} and that (s_i, t_i) is the only edge incident on s_i and also the only edge incident on t_i , in \hat{C} . In other words, s_i can be matched with exactly one other vertex, t_i .

If $\deg(w_i) = 1$ and $i = p$, so that s_i corresponds to w_p , then some edge $e_1 = (w_p, v_\alpha)$ is in F_1 . Some other edge $e_2 = (v_\alpha, w_\beta)$ must be in F_2 . Thus, s_i in \hat{C} is adjacent to exactly one vertex, namely, t_β . It follows that (s_i, t_β) must be a matching edge in \hat{C} .

If we delete e_1 and e_2 from \hat{H} , then we can repeat the argument above, beginning at w_β , to find another edge that must be in the matching on \hat{C} . This procedure will continue until $\deg(t_\beta) = 1$. Since the only other vertex of degree 1 that we have not dealt with is w_q , this will happen precisely when we reach w_q .

Finally, if $\deg(w_i) = 2$, and if w_i is not on the path between w_p and w_q , then w_i must lie on a cycle in \hat{H} . The cycle edges alternate between F_1 and F_2 . We can increase $|F_1 \cap F_2|$ by putting every other cycle edge into both F_1 and F_2 . This contradicts the assumption that the matchings were chosen to maximize $|F_1 \cap F_2|$.

This completes the proof that there is only one perfect matching in \hat{C} . \square

3.3 MEASURE THEORY. Lemma 3.2 in the previous section implies that if the graph $\langle H, r, p, q \rangle$ has a perfect matching, then the special submatrix $\langle A, r, p, q \rangle$ of $A^T A$ is made nonsingular by some assignment of values to A —that is, by some $A' \in N(A)$. In this section we show that there is a single assignment of values to A that makes all these special submatrices nonsingular at once.

LEMMA 3.5. *Let A be a matrix with t nonzeros, where in place of each nonzero we have put a variable, x_1, x_2, \dots, x_t . Let M_1, M_2, \dots, M_k be submatrices of $A^T A$. Suppose that, for each i with $1 \leq i \leq k$, there is an assignment of values $\mathbf{a}^{(i)} = (a_1^{(i)}, a_2^{(i)}, \dots, a_t^{(i)})$ to the nonzero entries of A such that M_i is nonsingular. Then there is a single assignment $\mathbf{a} = (a_1, \dots, a_t)$ such that every M_i , $1 \leq i \leq k$, is nonsingular.*

PROOF. Each $\det(M_i)$ is a nonzero polynomial $p_i(\mathbf{x})$ with integer coefficients. The zeros of such a polynomial form a set of measure zero [8]. Since there are only k such polynomials the set $S = \bigcup_i \{\mathbf{x} : p_i(\mathbf{x}) = 0\}$ is also of measure zero. Hence, there must exist a point $\mathbf{a} \notin S$; so $p_i(\mathbf{a}) \neq 0$ for all i . \square

3.4 MAIN THEOREM. With all the necessary lemmas in hand, we are ready for the main theorem of this section.

THEOREM 3.6. *Let A be an $m \times n$ matrix, $m \geq n$, with the strong Hall property, and an orthogonal factorization $A = QR$. The following algorithm will correctly compute the nonzero structure of the upper triangular factor R .*

- (1) Determine the structure of $A^T A$.
- (2) Apply symbolic Gaussian elimination to the structure of $A^T A$, yielding the structure of R .

PROOF. Recall the two conditions at the beginning of Section 3 that are required for a structural algorithm to be called correct.

The first of these conditions, that $H(R) \subseteq H_R$, is just Lemma 1.1, and was proved by George and Heath [11].

For the second condition, that there exist a matrix $A' \in N(A)$ such that $H(R') = H_{R'}$, recall that symbolic factorization of $A^T A$ gives an incorrect result when the symbolic Gaussian elimination does not produce the same structure as that from the actual Gaussian elimination. Furthermore, this happens when cancellation occurs during the actual Gaussian elimination. Hence, we want to show that there is a matrix $B \in N(A)$ in which no cancellation occurs during Gaussian elimination in $B^T B$. By Lemma 3.1, it is enough if we can show that there is a matrix $A' \in N(A)$ such that for each subgraph $C = \langle H, r, p, q \rangle$ that has a perfect matching, $M = \langle A', r, p, q \rangle$ is nonsingular.

Let $C = \langle H, r, p, q \rangle$, $1 \leq r < p \leq q \leq n$, be a subgraph of H that has a perfect matching. By Lemma 3.2, for every such C there is a subgraph H' of H , such that $C' = \langle H', r, p, q \rangle$ has precisely one perfect matching.




FIG. 8. The Dulmage-Mendelsohn reordering of a matrix.

Form A' by setting $A'_{ij} = 1$ if (w_i, w_j) is an edge in H' , and setting $A'_{ij} = 0$, otherwise. Clearly, $A' \in N(A)$, and $M' = \langle A', r, p, q \rangle$ is nonsingular—since $C = H(M')$, and hence its determinant has exactly one nonzero term.

Thus, for each subgraph $C = \langle H, r, p, q \rangle$ with a perfect matching, there is a matrix $A' \in N(A)$ such that $M' = \langle A', r, p, q \rangle$ is nonsingular. Finally, by Lemma 3.5, we conclude that there is one matrix, B , such that for all subgraphs $C = \langle H, r, p, q \rangle$ with a perfect matching, $M = \langle B, r, p, q \rangle$ is nonsingular. For this matrix B , no essential cancellation will occur during Gaussian elimination in $B^T B$, and the theorem follows. \square

4. Finding the Structure of R

Section 3 showed us that, provided A has the strong Hall property, we can determine the structure of R correctly with symbolic factorization of $A^T A$. But what if we want to use the Local Givens Rule? Or, most important, what can we do when A does not have the strong Hall property? We answer these questions below. First, note that the Local Givens Rule also works correctly on matrices with the strong Hall property:

THEOREM 4.1. *Let A be an $m \times n$ matrix, $m \geq n$, with the strong Hall property. If we compute the nonzero structure of R by doing symbolic Givens rotations using the Local Givens Rule, then the result is correct.*

PROOF. Follows directly from Theorem 3.6 and Corollary 1.2. \square

In the general case—when the structure of A does not necessarily have the strong Hall property—we can permute A into parts, each of which has the desired property. This canonical reordering scheme was first studied by Dulmage and Mendelsohn [5–7, 22].

The reordering we want is what is known as “block upper triangular form” [16]. It is a permutation of the columns and rows of A that leaves square submatrices on the diagonal, and a rectangular submatrix in the lower right-hand corner. See Figure 8.

As we pointed out in Section 2, the strong Hall property is closely related to the strong connectedness of a directed graph. In particular, if A is a square matrix, we

can do our reordering by finding the strongly connected components in the directed graph that corresponds to A [16, 25]:

- (1) Find a perfect matching in A and, with this matching fixed, construct the directed graph $D(A)$.
- (2) Find the strongly connected components of D in topological order.

Regardless of which perfect matching we choose in Step 1, Step 2 will give the same row partition and the same column partition.

If A is not square, we can proceed as follows. First, find a matching in A that covers all the columns. Then, the last, rectangular, component of A consists of every column and row that can be reached from an uncovered row via an alternating path. The remaining rows and columns are a square submatrix that can be reordered into strongly connected components in topological order, as above. The rectangular component and each of the square strongly connected components have the strong Hall property. The rows and columns in each component are independent of which maximum matching in A we use.

The time to find a maximum matching is $O(t\sqrt{m+n})$ [20, 23], where A has t nonzeros. The rectangular component can be found by depth-first search in $O(t)$ time. The strong components can also be found in $O(t)$ time [1, 25]. (In fact, the same depth-first search can be used to find the rectangular component and the strong components.) Thus, the worst-case time for the reordering is bounded by $O(t\sqrt{m+n})$. Duff's experiments with a heuristic maximum-matching algorithm [4] indicate that in practice a maximum matching can often be found in $O(t)$ time; in such cases, the entire reordering takes $O(t)$ time.

THEOREM 4.2. *Let A be an $m \times n$ rectangular matrix. If we use the algorithm above to permute A into A' , and use the Local Givens Rule on A' to find the structure of its factor R' , then the result is correct.*

PROOF. Let M be any of the square submatrices on the diagonal. If there are none, A itself has the strong Hall property, and the theorem follows from Theorem 4.1.

Let c_i, \dots, c_j be the columns that are in M , and let c_k be any other column in A' to the right of M . If c_k has only zeros in rows i, \dots, j , then the factoring of M does not affect the structure of c_k in these rows.

If c_k has at least one nonzero in rows i, \dots, j , consider the rectangular submatrix, U , made up from the columns c_i, \dots, c_{j-1}, c_k . Since column c_k has a nonzero in row k , U has the strong Hall property. Hence, by Theorem 4.1, the Local Givens Rule works on this matrix, and it follows that the Local Givens Rule works correctly on the whole of A' . \square

Although the result in the theorem above is interesting in its own right, it turns out that in practice it is probably better to solve the system by "block back-substitution"—and never compute the structure of R . This is done as follows (refer to Figure 9 for notation):

Algorithm 4.1

- (1) Reorder A as shown in Figure 9, with square diagonal blocks M_1, M_2, \dots, M_k and a rectangular diagonal block M_{k+1} .
- (2) Solve the rectangular least-squares system $M_{k+1}x_{k+1} \approx b_{k+1}$.
- (3) For $i = k, k-1, \dots, 1$ do:
Update the vector b , by setting $(b_1, \dots, b_i) \leftarrow (b_1, \dots, b_i) - U_{i+1}x_{i+1}$. Solve the square system $M_i x_i = b_i$.




FIG. 9. Block back-substitution.

This way of solving the system has several advantages. First, in step (2), we know that M_{k+1} has the strong Hall property, so by Theorem 4.1 we shall not use any unnecessary space—whether we use symbolic factorization of $A^T A$ or the Local Givens Rule. In case we use symbolic factorization of $A^T A$, note that the strong Hall property is preserved under column and row permutations. This means we can use a heuristic, like minimum degree, on the structure of $A^T A$ to find a column ordering that gives small fill, and still be guaranteed that we are not allocating any extra space for R .

Second, all the systems we solve in step (3) are square, which means that we can use any standard algorithm to solve these and to find orderings that give little fill. Of course, one way to solve the square system $M_i x_i = b_i$ is by a QR factorization of M_i [18]. If we do this, Theorem 3.6 says that symbolic factorization of $A^T A$ correctly predicts the structure of M_i 's triangular factor. Symbolic factorization of $A^T A$ can also be useful when solving $M_i x_i = b_i$ by Gaussian elimination with partial pivoting, as shown by George and Ng [13].

Third, by updating the vector b every time we have determined new values for some variables in x , we avoid having to store any fill at all in U , the part of A above the diagonal blocks.

5. Summary

We have described the nonzero structure of matrices and singled out a class of structures, namely, those with the strong Hall property. For matrices with this property, we have shown that symbolic factorization of $A^T A$ correctly predicts the structure of the upper triangular factor R . We have also shown that this condition is sufficient to say that the Local Givens Rule will correctly predict the structure of R . However, in neither case is it a necessary condition: there are matrices that do not have the strong Hall property and in which no bogus fill occurs.

We have given a way of reordering a matrix of arbitrary structure into blocks for which symbolic factorization of $A^T A$ correctly predicts fill. We have outlined an algorithm that uses this reordering to solve a least squares problem block by block, storing no unnecessary fill within each block and storing no fill at all between

blocks. The reordering also gives a way to permute columns so that the Local Givens Rule will correctly predict the structure of the triangular factor of the whole matrix.

Our analysis suggests that the fill from a Givens rotation cannot be characterized by a local rule—that is, a rule based on the structure of only a fixed number of rows—as opposed to Gaussian elimination. We have found it necessary to use both a local rule—the Local Givens Rule—and a concept that summarizes the structure of the whole matrix—the strong Hall property. On the other hand, the fill in Gaussian elimination can be given a precise *static* characterization: we can predict which positions will be filled in terms of paths in the initial structure, as in Lemma 2.1. This characterization has been a valuable theoretic tool. For this reason it might be interesting to find a similar static characterization for Givens rotations. George et al. explore this theme in [14].

REFERENCES

(Reference [3] is not cited in text.)

1. AHO, A. V., HOPCROFT, J. E., AND ULLMAN, J. D. *The Design and Analysis of Computer Algorithms*. Addison-Wesley, Reading, Mass., 1974.
2. BRAYTON, R. K., GUSTAVSON, F. G., AND WILLOUGHBY, R. A. Some results on sparse matrices. *Math. Comput.* 24 (1970), 937–954.
3. BUNCH, J. R., AND ROSE, D. J., Eds. *Sparse Matrix Computations*. Academic Press, Orlando, Fla., 1976.
4. DUFF, I.A. On algorithms for obtaining a maximum transversal. *ACM Trans. Math. Softw.* 7 (1981), 315–330.
5. DULMAGE, A. L., AND MENDELSON, N. S. Coverings of bipartite graphs. *Canad. J. Math.* 10 (1958), 517–534.
6. DULMAGE, A. L., AND MENDELSON, N. S. A structure theory of bipartite graphs of finite exterior dimension. *Trans. Roy. Soc. Can.* 53, III (1959), 1–13.
7. DULMAGE, A. L., AND MENDELSON, N. S. Two algorithms for bipartite graphs. *J. SIAM* 11 (1963), 183–194.
8. FEDERER, H. *Geometric Measure Theory*. Springer-Verlag, Berlin, 1969.
9. FORD, L. R., AND FULKERSON, D. R. *Flows in Networks*. Princeton University Press, Princeton, N.J., 1962.
10. GENTLEMAN, W. M. Row elimination for solving sparse linear systems and least squares problems. In *Proceedings of the Dundee Conference on Numerical Analysis* (1975), G. A. Watson, Ed., Lecture Notes in Mathematics, vol. 506. Springer-Verlag, New York, 1976, pp. 122–123.
11. GEORGE, A., AND HEATH, M. T. Solution of sparse linear least squares problems using Givens rotations. *Linear Alg. Appl.* 34 (1980), 69–83.
12. GEORGE, A., AND LIU, J.W. *Computer Solution of Large Sparse Positive Definite Systems*. Prentice-Hall, Englewood Cliffs, N.J., 1981.
13. GEORGE, A., AND NG, E. An implementation of Gaussian elimination with partial pivoting for sparse systems. *SIAM J. Sci. Stat. Comput.* 6 (1985), 390–409.
14. GEORGE, A., LIU, J., AND NG, E. Row ordering schemes for sparse Givens transformations. I: Bipartite graph model. *Linear Alg. Appl.* 61 (1984), 55–81.
15. GOLUB, G. H., AND BUSINGER, P. Linear least squares solutions by Householder transformations. *Numerische Mathematik*, 7, *Hand Book Series on Linear Algebra*. (1965), 269–276.
16. GUSTAVSON, F. Finding the block lower triangular form of a sparse matrix. In *Sparse Matrix Computations*, J. R. Bunch and D. J. Rose, Eds. Academic Press, Orlando, Fla., 1976, pp. 275–289.
17. HALL, P. On representatives of subsets. *J. Lond. Math. Soc.* 10 (1935), 26–30.
18. HEATH, M. T. Some extensions of an algorithm for sparse linear least squares problems. *SIAM J. Sci. Stat. Comput.* 3 (1982), 223–237.
19. HEATH, M. T. Numerical methods for large sparse linear least squares problems. *SIAM J. Sci. Stat. Comput.* 5 (1984), 497–513.
20. HOPCROFT, J. E., AND KARP, R. M. An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM J. Comput.* 2 (1973), 225–231.
21. HOWELL, T. D. Partitioning using PAQ. In *Sparse Matrix Computations*, J. R. Bunch and D. J. Rose, Eds. Academic Press, Orlando, Fla., 1976, pp. 23–37.

22. JOHNSON, D. M., DULMAGE, A. L., AND MENDELSON, N. S. Connectivity and reducibility of graphs. *Canad. J. Math.* 14 (1962), 529-539.
23. PAPADIMITRIOU, C. H., AND STEIGLITZ, K. *Combinatorial Optimization*. Prentice-Hall, Englewood Cliffs, N.J., 1982.
24. ROSE, D. J., TARJAN, R. E., AND LUEKER, G. S. Algorithmic aspects of vertex elimination on graphs. *SIAM J. Comput.* 5 (1976), 266-283.
25. TARJAN, R. Depth-first search and linear graph algorithms. *SIAM J. Comput.* 1 (1972), 146-160.

RECEIVED OCTOBER 1983; REVISED JULY 1985; ACCEPTED JULY 1985