Predicting the Structure of Sparse Orthogonal Factors

Alex Pothen*

Department of Computer Science
University of Waterloo
Waterloo, Ontario, N2L 3G1, Canada
*(*apothen@narnia.uwaterloo.ca,
na.pothen@na-net.ornl.gov*).*

Submitted by Richard A. Brualdi

ABSTRACT

The problem of correctly predicting the structures of the orthogonal factors Q and R from the structure of a matrix A with full column rank is considered in this paper. Recently Hare, Johnson, Olesky, and van den Driessche have described a method to predict these structures, and they have shown that corresponding to any specified nonzero element in the predicted structures of Q or R, there exists a matrix with the given structure whose factor has a nonzero in that position. In this paper this method is shown to satisfy a stronger property: there exist matrices with the structure of A whose factors have exactly the predicted structures. These results use matching theory, the Dulmage-Mendelsohn decomposition of bipartite graphs, and techniques from algebra. The proof technique shows that if values are assigned randomly to the nonzeros in A, then with high probability the elements predicted to be nonzero in the factors have nonzero values. It is shown that this stronger requirement cannot be satisfied for orthogonal factorization with column pivoting. In addition, efficient algorithms for computing the structures of the factors are designed, and the relationship between the structure of Q and

*This author was supported by NSF grant CCR-9024954 and by U. S. Department of Energy grant DE-FG02-91ER25095 at the Pennsylvania State University and by the Canadian Natural Sciences and Engineering Research Council under grant OGP0008111 at the University of Waterloo.

LINEAR ALGEBRA AND ITS APPLICATIONS xxx:1–xx (1992)1

655 Avenue of the Americas, New York, NY 10010 0024-3795/92/$5.00

the Householder array is described.

---

1. INTRODUCTION

Given the structure of an $m \times n$ real or complex matrix A, where $m \geq n$ and A has full column rank, we consider the problem of correctly predicting the structure of its orthogonal factors Q and R. (Here Q is $m \times n$ and R is $n \times n$. The full rank assumption is necessary for the factors to be unique.) Algorithms for structure prediction give valuable insight into the nature of sparse factorizations, and enable us to set up data structures for the factors so that the numerical factorization can be computed in time proportional to the number of floating point operations.

Coleman, Edenbrandt, and Gilbert [4] proved that if A has a combinatorial property called the *strong Hall property (SHP)*, then two algorithms called the George-Heath algorithm [10] and the Local Givens Rule correctly predict the structure of R. More recently Hare, Johnson, Olesky, and van den Driessche [16] have shown how the structures of Q and R may be predicted when A does not have SHP. Let $\mathcal{A}$ denote the set of matrices with full column rank whose structures are contained in that of A, i.e, the set of full rank matrices B such that $b_{ij} \neq 0 \rightarrow a_{ij} \neq 0$. The structures predicted by their method satisfy two requirements:
(1.1) The predicted structures are large enough to contain the structures of the factors of any matrix in $\mathcal{A}$, and
(1.2) Corresponding to any specified nonzero element (i, j) in the predicted structure of Q or R, there exists a matrix $A_{ij} \in \mathcal{A}$ whose factor has a nonzero in that position.

A natural question that arises is if there exists a single matrix $A' \in \mathcal{A}$ whose factors have exactly the predicted structures. This leads, instead of (1.2) to the requirement:
(1.2)$'$ There exists a matrix $A' \in \mathcal{A}$ whose factors Q' and R' are simultaneously nonzero in every nonzero element in the predicted structures.

It is this stronger requirement that has been considered in previous work on predicting the structure of R in the orthogonal factorization of strong Hall matrices [4]. Additionally, it is known that the predicted structures in Cholesky factorization [23], unsymmetric Gaussian elimination of a matrix with nonzero diagonal [24], the solution of a linear system when the coefficient matrix has a nonzero diagonal [13], and eigendecompositions of certain matrices [13] satisfy $(1.2)'$. It is also known that this requirement cannot be satisfied for unsymmetric Gaussian elimination when the matrix has zero elements on the diagonal [1] and for Gaussian elimination with pivoting.

In this paper we prove that the structures predicted by the methodology of Hare et al. satisfy the stronger requirement $(1.2)'$. We also show that this requirement cannot be satisfied for orthogonal factorization with column pivoting. In addition, we describe efficient algorithms for predicting the structure of orthogonal factors.

To prove that the predicted data structures satisfy $(1.2)'$, we will make use of the concepts of a Hall set and an auxiliary graph introduced by these authors, and then employ matching theory, the Dulmage-Mendelsohn decomposition of bipartite graphs, and some results from algebra. Hall sets can be computed efficiently from the Dulmage-Mendelsohn decomposition, and the efficient structure prediction algorithms we describe later in this paper make use of this decomposition as well. Hence the use of this decomposition in this context is quite natural. The techniques used by Hare et al. to characterize nonzero elements in the structures cannot be extended to establish these results since they may assign different values to a particular element in A to show that two elements of Q are nonzero.

These results are of theoretical interest since they bring structure prediction for orthogonal factorization on a par with known results for structure prediction for other factorizations. In addition, they have practical implications as well. If a predicted structure satisfied the requirement (1.2) but not $(1.2)'$, then after the numerical factorization is computed, it may be worthwhile to 'compress' the data structures by removing those elements which are actually zero while predicted to be nonzero to reduce storage requirements and to avoid arithmetic on zero elements. (Such schemes have been considered in [5, 11, 12].) Our proof technique shows that if numerical values are assigned randomly to the nonzeros in A, then with high probability the predicted nonzero elements in the factors are actually nonzero. An important conclusion we can draw is that when the numerical values in A are reasonably random, the use of a post-processing phase to remove zeros from the data structures for Q and R will not be worth the trouble.

The rest of this paper is organized as follows. In section 2 we briefly describe the Dulmage-Mendelsohn decomposition of bipartite graphs, and discuss in more detail the work of Hare et al. We characterize maximum Hall sets and the structure of alternating paths in an auxiliary graph B_j in section 3. We make use of some algebraic results and the characterizations in the previous section to show how the structures of the orthogonal factors may be predicted in section 4. In section 5 we describe efficient algorithms to compute the structures, relate the structure of the Householder array to that of the orthogonal factor, and provide an example to illustrate that the factors obtained from orthogonal factorization with column pivoting cannot satisfy the requirement $(1.2)'$. Section 6 discusses the significance of these results to computing sparse orthogonal factorization.

Notation. Throughout this paper, A will denote an $m \times n$ matrix with full column rank, where $m \geq n$. We represent the structure of A by means of the bipartite graph $H = H(A) = (\mathcal{R}, \mathcal{C}, \mathcal{E})$, where $\mathcal{R} = \{r_1, \ldots, r_m\}$ is the set of *row* vertices, $\mathcal{C} = \{c_1, \ldots, c_n\}$ is the set of *column* vertices, and an edge $(r_i, c_j) \in \mathcal{E}$ if and only if $a_{ij} \neq 0$. For convenience, we will assume without loss of generality that the rows are numbered such that A has a nonzero diagonal. For $R_1 \subseteq \mathcal{R}$ and $C_1 \subseteq \mathcal{C}$, the *induced subgraph* $H_1 = (R_1, C_1, E_1)$ is the subgraph of H whose vertex sets are R_1 and C_1, and whose edge set $E_1 \subseteq \mathcal{E}$ contains those edges of H with one endpoint in R_1 and the other in C_1. We shall often write it as the subgraph (R_1, C_1). *The subgraph of H induced by a column subset C_1* is the subgraph whose edge set consists of all edges with one endpoint in C_1, and whose row set consists of all rows which are endpoints of such edges.

2. BACKGROUND

In this section we review the Dulmage-Mendelsohn decomposition and the work of Hare, Johnson, Olesky, and van den Driessche [16].

2.1. The Dulmage-Mendelsohn decomposition

The block triangular form (btf) of A induced by the Dulmage-Mendelsohn (D-M) decomposition of the bipartite graph $H(A)$ has been described by Dulmage, Johnson and Mendelsohn [7, 8, 9, 17], and by Brualdi [2, 3]. Recent descriptions of this decomposition in terms of bipartite matching theory may be found in

[20, 22], with proofs included in [20]. Since this discussion of the D-M decomposition will be brief, the reader unfamiliar with this decomposition will find it helpful to consult [22].

If the rows and columns of a matrix A with full column rank are permuted appropriately, then the D-M decomposition leads to the btf

$$A = \begin{pmatrix} A_s & X \\ 0 & A_v \end{pmatrix},$$

where A_s is a square submatrix, A_v is an overdetermined matrix, and 'X' denotes a possibly nonzero submatrix of apt dimensions. (The submatrix A_s has a block upper triangular structure, and A_v is block diagonal.) The D-M decomposition is conveniently described with respect to a maximum matching in the bipartite graph $H(A)$. The terminology and results on matchings in graphs used here may be found in Lovász and Plummer [18].

A bipartite graph $H(A)$ with a matching is shown in Fig. 2.1, where the matched edges $\{(r_i, c_i) : i = 1, \ldots, 7\}$ are drawn as 'horizontal' edges. A *walk* is a sequence of vertices v_0, v_1, …, v_n such that (v_i, v_{i+1}) is an edge for $i = 0$, …, $n-1$. Vertices or edges may be repeated in a walk. An *alternating walk* in the graph is a walk with alternate edges in M. An *alternating tour* is an alternating walk whose endpoints are the same. An *alternating path* is an alternating walk with no repeated vertices. Following Gilbert [14], depending on the direction in which the matching edges are traversed, we distinguish between two kinds of alternating paths: In an r-alternating path, the matched edges are traversed from a column to a row, and in a c-alternating path, they are traversed from a row to a column. In Fig. 2.1, the path r_8, c_7, r_7, c_1, r_1 is an r-alternating path from r_8 to r_1 (it is a c-alternating path from r_1 to r_8); the path r_2, c_2, r_3, c_3 is c-alternating from r_2 to c_3.

A *maximum matching* is a matching of maximum cardinality. A matching is *column-perfect* if every column vertex is matched; it is *row-perfect* if every row vertex is matched. A matching is *perfect* if it is column-perfect and row-perfect. The matching in Fig. 2.1 is a maximum matching since it is column-perfect.

The D-M decomposition is described with respect to a maximum matching in the graph $H(A)$, but since it is a canonical decomposition of the matrix, any other maximum matching would lead to the same column and row sets in the decomposition.

Let SR denote the rows and SC the columns of A_s, and VR denote the rows and VC the columns of A_v. We call the sub-

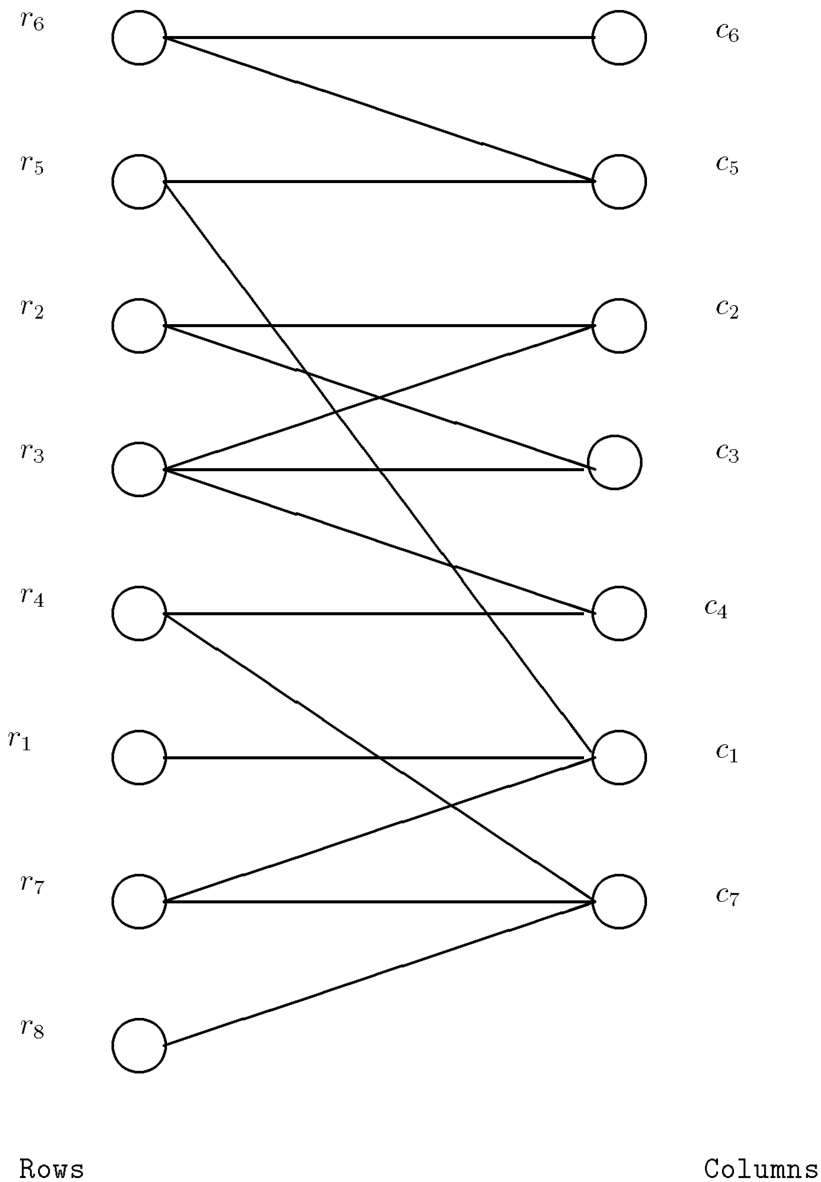

Figure 2.1: A bipartite graph H, a maximum matching, and its Dulmage-Mendelsohn decomposition.

graph of $H(A)$ induced by (SR, SC) the *square subgraph* H_s, and that induced by (VR, VC), the *overdetermined subgraph* H_v. The set VR can be characterized as the set of rows reached by r-alternating paths from unmatched rows, and VC is the set of columns thus reached. Note that all unmatched rows are included in VR since they can be reached by r-alternating paths of length zero, and that all columns in VC are matched to rows in VR. All remaining rows are perfectly matched to all remaining columns, and we call these sets SR and SC, respectively.

The overdetermined subgraph H_v may have more than one connected component. The overdetermined submatrix H_v has a block diagonal structure, corresponding to the connected components of H_v. We list the diagonal blocks of A_v as V_1, V_2, …, V_q, and denote the row set of V_i by U_i and its column set by D_i.

In the bipartite graph $H(A)$ shown in Fig. 2.1, the square subgraph H_s has columns $SC = \{c_6, c_5, c_2, c_3, c_4\}$ and its row set SR is the set of rows matched to these columns. The overdetermined subgraph H_v has columns $VC = \{c_1, c_7\}$ and rows $VR = \{r_1, r_7, r_8\}$, and has only one connected component.

The square subgraph H_s has a finer decomposition which leads to a block upper triangular form for the submatrix A_s. Define two columns in SC to be *related* if there is an alternating tour joining them. This is an equivalence relation, and let the classes of this relation be the column sets C_1, C_2, …, C_p. Let R_i denote the row set matched to C_i. It is possible to renumber the sets $\{R_i\}$ (and $\{C_i\}$) such that if $i > j$, then no edge joins a vertex in R_i to a column in C_j. (This renumbering may not be unique.) Henceforth we assume that these row sets and column sets have been renumbered to satisfy this property. In Fig. 2.1, $C_1 = \{c_6\}$; $C_2 = \{c_5\}$; $C_3 = \{c_2, c_3\}$; and $C_4 = \{c_4\}$.

Permuting the rows and columns in the above order leads to the block upper triangular form of A_s. The diagonal blocks of this form are square submatrices induced by the row set R_i and the column set C_i. We number the block diagonal submatrices T_1, T_2, …, T_p, and each submatrix T_i is irreducible.

A bipartite graph $H(A)$ with m rows and n columns $(m \geq n)$ has the *Hall property* (HP) if every set of k columns $(1 \leq k \leq n)$ is adjacent to at least k rows. It has the *strong Hall property* (SHP) if every set of k columns $(1 \leq k < m)$ is adjacent to at least $k + 1$ rows. Thus when $m > n$, every set of $k \leq n$ columns satisfies the adjacency requirement, and when $m = n$, every set of $k < n$ columns satisfies it. Notice the asymmetry in the definitions of HP and SHP for square bipartite graphs.

Philip Hall proved that the graph $H(A)$ has a column-perfect matching if and only if it has the HP. If the corresponding matrix A has full column rank, then it has a square nonsingular submatrix of order n. Hence there is at least one nonzero term in the alternating sum expansion of the determinant of the submatrix, from which we can conclude that A has a column-perfect matching, and hence the HP.

Each connected component of the overdetermined subgraph (the subgraph induced by each U_i and D_i) has SHP. Also, the square subgraph corresponding to each diagonal block T_i in H_s has SHP. Henceforth we call the diagonal blocks the *strong Hall components* of the respective subgraphs.

An easy consequence of the existence of an M-alternating tour joining any two vertices in a strong Hall component of a square subgraph T is that there is a c-alternating path from any vertex v to any other vertex w with respect to *any* perfect matching M in T. The next result, due to Gilbert, characterizes a strong Hall component of the overdetermined subgraph.

Lemma 2.1 ([14]) *Let V be a strong Hall component of the overdetermined subgraph in the D-M decomposition of a bipartite graph H, and let a vertex v and a column c belonging to V be specified. Then there exists a column-perfect matching M (which depends on v and c) in V such that there is a c-alternating path from v to c.* ∎

2.2. Previous work

For $1 \leq j \leq n$, let $\underline{a}_j$ denote the j-th column of A. We will find it necessary in this paper to consider A_j, the submatrix of A consisting of the first j columns. We let $J = \{c_1, \ldots, c_j\}$ be the set of the first j column vertices of H, and represent the structure of A_j by the bipartite subgraph $H_j = H(A_j)$, the subgraph of $H(A)$ induced by J.

Hare et al. [16] introduced the following two concepts. The first of these is the *Hall set*, a set of column vertices which is adjacent in H to exactly as many rows. A *maximum Hall set* S_j is a Hall set of largest cardinality in J. (The set S_0 is defined to be the empty set.) Let s_j denote the set of rows adjacent to S_j; then from the definition of a Hall set, these two sets are equal in size. The second concept is that of an auxiliary bipartite graph $B_j = (R', C', E')$, the subgraph of H_j from which the column set S_{j-1} and the row set s_{j-1} have been excluded. (This is the subgraph of H_j induced by the columns in $J \setminus S_{j-1}$.) Define p_j (P_j) to be the the set of row (column) vertices of B_j that belong to the same connected component as c_j; let u_j (U_j) be the remaining row (column) vertices of B_j; and let t_j denote the set of row vertices of H that are not adjacent to any column in J. Then after appropriate row and column permutations, A_j has the structure

$$A_j = \begin{array}{c} \\ t_j \\ s_{j-1} \\ u_j \\ p_j \end{array} \begin{array}{c} \begin{array}{ccc} S_{j-1} & U_j & P_j \end{array} \\ \left(\begin{array}{ccc} 0 & 0 & 0 \\ A_S & A_{SU} & A_{SP} \\ 0 & A_U & 0 \\ 0 & 0 & A_P \end{array} \right) \end{array}. \tag{2.1}$$

Here the submatrices are zero in the first row since rows in t_j are not adjacent to any column in J. The zero submatrices in the first column follow from the definition of a Hall set. Finally, since the vertex sets $p_j \cup P_j$ and $u_j \cup U_j$ belong to different connected components of B_j, the other zero submatrices in the third and fourth rows follow.

Hare et al. predict the structure of $\underline{q}_j$, the j-th column of Q, by means of the auxiliary graph B_j.

THEOREM 2.2 ([16]) *Let the bipartite graph $H(A)$ represent the structure of an $m \times n$ matrix A with full column rank, where $m \geq n$. Let Q denote the $m \times n$ orthogonal factor of a matrix $A_{ij} \in \mathcal{A}$. For $1 \leq j \leq n$ and $1 \leq i \leq m$,*

1. *if $r_i \in \mathcal{R} \setminus p_j$, then $q_{ij} = 0$.*
2. *if $r_i \in p_j$, then there exist values for the nonzeros in A_{ij} such that $q_{ij} \neq 0$.* ■

The authors also proved that values can be assigned to a matrix $A_{kl} \in \mathcal{A}$ such that any nonzero element r_{kl} in the structural product $Q^T A = R$ is nonzero. To prove the second part of the above theorem, the authors construct a matrix A_{ij} with values $\{\pm 1, \epsilon\}$ (here ϵ is a small positive value) for the nonzeros in A_{ij}, and show by a direct computation that $q_{ij} \neq 0$. They did not address the question if there was a single assignment of values to the nonzeros in A that simultaneously makes every such element of Q and R nonzero. Furthermore, since they assign specific numerical values from the set $\{\pm 1, \epsilon\}$, a common nonzero might be given different values in two different submatrices, and thus this technique cannot be extended to prove such a result.

3. A CLOSER LOOK AT THE AUXILIARY GRAPH

In this section, we characterize Hall sets, maximum Hall sets, the auxiliary graph B_j, and the subsets p_j and P_j by means of the D-M decomposition.

3.1. Hall sets

In characterizing Hall sets by means of the D-M decomposition, we will find the concept of a predecessor of a column set C_i useful. We assume that the column sets of the square subgraph H_s have been renumbered as described in section 2. A column set C_i *precedes* a set C_k ($i < k$) if and only if there is an r-alternating path from some column $c \in C_i$ to some column $d \in C_k$ in H_s. (Henceforth we will say that there is an r-alternating path from C_i to C_k.) The set of *predecessors* of C_k includes all the column sets which precede C_k (this set does not include C_k itself). The *least predecessor* of C_k is its lowest-numbered predecessor.

A Hall set is *simple* if it not the union of two or more Hall sets. We now characterize the simple Hall sets of $H(A)$ by means of its D-M decomposition.

LEMMA 3.1. *Let C_i be the column set of a square strong Hall component T_i in the D-M decomposition of a bipartite graph H. The columns in C_i and its set of predecessors together form a simple Hall set.*

Proof. First, we consider the case when C_i has no predecessor. From the renumbering of these sets in the D-M decomposition, no edge can join a column in C_i to a row in some R_k, where $k > i$. Since C_i has no predecessor, there is no edge from a column in C_i to a row in some R_h, where $h < i$. Thus the columns in C_i are adjacent only to the rows in R_i, and since these two sets are perfectly matched, they constitute a Hall set. Further, since the induced subgraph (R_i, C_i) has SHP, any proper subset S of columns in C_i is adjacent to more than $|S|$ rows. Hence no proper subset of columns in C_i is a Hall set, and it follows that these columns form a simple Hall set.

Now we consider the case when C_i has one or more predecessors. By the definition of a predecessor, there exists an r-alternating path from some row set R_g to C_i, with $g < i$. Choose a predecessor C_h, matched to the row set R_h, such that some row $r \in R_h$ is adjacent to a column in C_i. The columns in C_i are adjacent to more than $|C_i|$ rows. Since the induced subgraph (R_i, C_i) has SHP, the induced subgraph $(R_i \cup \{r\}, C_i)$ has SHP. (Note that the former is a square subgraph, and that the latter is an overdetermined subgraph, and hence our definitions of SHP in the two cases differ slightly.) Thus C_i by itself cannot be a simple Hall set.

However, the columns in C_i and its set of predecessors are adjacent only to the row sets perfectly matched to them, and thus form a Hall set. As in the first case, a proper subset of columns in C_i cannot be a Hall set. If S is a set including all columns in C_i together with some proper subset of the columns in its predecessors, then since each predecessor is a strong Hall component of the square subgraph H_s, S is adjacent to more than $|S|$ rows. Thus S cannot form a Hall set. Hence columns in C_i and its predecessors together constitute a simple Hall set. ∎

The sets $\{c_6\}$, $\{c_5, c_6\}$, $\{c_2, c_3\}$, and $\{c_4, c_2, c_3\}$ are the simple Hall sets in Fig. 2.1. We proceed to characterize a maximum Hall set S_j by means of the D-M decomposition.

LEMMA 3.2. *S_j consists of all column sets C_i such that C_i and its predecessors have all their columns numbered less than or equal to c_j.*

Proof. Since the subgraph induced by a column set C_h and the row set R_h has SHP, a nonempty proper subset of C_h cannot be a Hall set. Thus if C_h has one or more columns greater than c_j, its remaining columns cannot be in S_j.

If the column set C_i has predecessors, from Lemma 3.1, the columns in C_i and its set of predecessors together form a Hall set. If all these columns are numbered less than or equal to c_j, then this Hall set belongs to S_j.

By the characterization of simple Hall sets in Lemma 3.1, two incomparable simple Hall sets cannot have any column vertices in common. (It is possible for a simple Hall set to be contained in another, as the example in Fig. 2.1 shows.) Also, the union of vertex-disjoint simple Hall sets is a Hall set. Thus S_j is obtained by the union of all column sets C_i such that C_i and its predecessors have all columns less than or equal to c_j. ∎

In Fig. 2.1, $S_2 = S_1 = S_0 = \emptyset$; $S_3 = \{c_2, c_3\}$; $S_4 = \{c_2, c_3, c_4\}$; $S_5 = S_4$; $S_6 = \{c_2, c_3, c_4, c_5, c_6\}$; and $S_7 = S_6$.

3.2. Paths in the auxiliary graph B_j

Let s_j denote the set of rows adjacent in $H(A)$ to columns in S_j. For $j = 1, \ldots, n$, recall that the bipartite graph B_j is the subgraph of $H(A_j)$ obtained by excluding the columns in S_{j-1} and the rows in s_{j-1}. There is a pretty characterization of the structure of B_j in terms of its D-M decomposition. We use unprimed entities to refer to the graph $H(A)$ and primed entities to refer to B_j.

THEOREM 3.3. *The D-M decomposition of the graph B_j has one of the following mutually exclusive structures:*

1. *c_j belongs to the overdetermined subgraph H'_v; the square subgraph $H_s{}'$ is empty.*
2. *c_j belongs to the square subgraph H'_s; then $c_j \in C_1{}'$, where $C_1{}'$ is the least predecessor of all other square strong Hall components.* ■

The proof of this theorem is by a lengthy case analysis which obtains the D-M decomposition of B_j in terms of that of $H(A)$, and makes use of Lemmas 3.1 and 3.2. The proof is omitted here but may be found in [21].

An example of the structure of B_j when the square subgraph is present may be seen from Fig. 2.1. In the Figure, when $j = 6$, since $S_5 = \{c_2, c_3, c_4\}$, the D-M decomposition of the graph B_j is $C'_1 = \{c_6\}$, $C'_2 = \{c_5\}$, $VC' = \{c_1\}$, and $VR' = \{r_1, r_7\}$.

We now make use of the structural characterization of B_j in Theorem 3.3 to prove the main result of this section.

THEOREM 3.4. *Given a vertex $v \in p_j \cup P_j$, there exists a column-perfect matching M (which may depend on v) in the auxiliary graph B_j such that there is a c-alternating path from v to c_j.*

Proof. Since $v \in p_j \cup P_j$, there is a path in B_j from v to c_j. What the Theorem asserts is that we can choose the path to be c-alternating from v to c_j, relative to some column-perfect matching that depends on v.

From Theorem 3.3, the graph B_j has two possible structures. If c_j belongs to the overdetermined subgraph, then the square subgraph is empty, and B_j has the SHP. Hence by Lemma 2.1, it is possible to construct a column-perfect matching M in B_j such that there is a c-alternating path from v to c_j.

If c_j belongs to the square subgraph, then it belongs to C'_1. There are now two cases to consider.

The first case is when v belongs to the square subgraph of B_j. If the vertex v is a row, let R'_k denote the row set it belongs to, and let C'_k be the column set matched in any perfect matching of the square subgraph of B_j to R'_k. If v is a column vertex, let C'_k be the column set that it belongs to. Then from Theorem 3.3, C'_1 is a predecessor of C'_k.

Let M be any column perfect matching of B_j. Let $r_j \in R'_1$ be the row matched to c_j. If v is a row, let $c \in C'_k$ be the column matched to v, and otherwise, let c denote the column v. By the definition of a predecessor, there is an r-alternating path from c_j to v in B_j. By traversing this path in the reverse direction, we find the desired c-alternating path from v to c_j.

Finally, consider the case when the vertex v belongs to the overdetermined subgraph of B_j. If v is a column, let D'_i denote the column set of a connected overdetermined strong Hall component that it belongs to. If v is a row, let U'_i denote the row set of a connected overdetermined strong Hall component that it belongs to, and let D'_i be the column set of this component. Since $v \in p_j \cup P_j$, there is a path (not necessarily c-alternating) from v to c_j in B_j. Hence there exists a column $c_l \in D'_i$ which is adjacent to some row $r_k \in R'_s$ such that the edge (c_l, r_k) lies on the above path from v to c_j. From Lemma 2.1, there is a column-perfect matching N_1 of the overdetermined subgraph such that there is a c-alternating path from v to c_l. The last edge of this path is a matched edge. From the column c_l, we take the edge (c_l, r_k) as an unmatched edge, and then continue as in the preceding paragraph to find a c-alternating path (with respect to any perfect matching N_2 of the square subgraph) from r_k to c_j. We let $M = N_1 \cup N_2$, and obtain a c-alternating path from v to c_j by concatenating the path from v to c_l, the edge (c_l, r_k), and the path from r_k to c_j. ■

By definition, every vertex on a c-alternating path from r to c_j is matched to another vertex on the path. Since the path is defined with respect to a column-perfect matching M, a column of B_j not on the path continues to be matched in M to a row of B_j. Further, since the column set S_{j-1} is perfectly matched to the row set s_{j-1}, and both these sets are outside B_j, M can be extended to a column perfect matching of $H(A_j)$. This fact will enable us in the next section to construct a matrix of full column rank such that the j-th column of its orthogonal factor has nonzero elements in the row set p_j.

4. STRUCTURE PREDICTION

In this section, we use the c-alternating path characterization of the set $p_j \cup P_j$ and some algebraic techniques to characterize the structures of the orthogonal factors. In addition to the bipartite graph of a matrix, we will work with two other classes of graphs: the adjacency graph of a symmetric matrix, and a product bipartite graph computed from two bipartite graphs.

Let $\overline{A}$ be a symmetric matrix of order k with a nonzero diagonal. We will find it useful to consider the adjacency graph $\overline{G} = G(\overline{A}) = (V, \overline{E})$ of $\overline{A}$ in predicting the structure of the factor Q. The *vector structure* of a k-vector $\underline{b}$ is $structure(\underline{b}) = \{i : b_i \neq 0\}$. We interpret this set as a subset of vertices in the adjacency graph $\overline{G}$. For ease of notation, we will say that a vertex v is in $\underline{b}$ to indicate that it belongs to $structure(\underline{b})$. The *closure* of $\underline{b}$ with respect to $\overline{G}$, $closure(\underline{b})$, is the set of vertices of $\overline{G}$ which are reachable by undirected paths from vertices in $\underline{b}$.

We will make use of the following result due to Gilbert in characterizing the nonzero structure of Q.

THEOREM 4.1 ([13]) *Consider the symmetric system $\overline{A}\underline{x} = \underline{b}$, where the nonzeros in $\overline{A}$ and $\underline{b}$ are specified, and $\overline{A}$ has a nonzero diagonal. Then there exist symmetric values for the nonzeros in $\overline{A}$ such that $structure(\underline{x}) = closure(\underline{b})$.* ■

We need to clarify what we mean by the phrase 'there exist values ...' in the statement of the Theorem. To do so, we require some algebra. A finite set $\hat{x}_1, \ldots, \hat{x}_t$ of complex numbers is *algebraically independent* over the rational field $\mathbf{Q}$ if $\hat{x}_1, \ldots, \hat{x}_t$ is not a root of any nonzero polynomial with integer coefficients in the t variables $x_1, \ldots, x_t$. If we assign algebraically independent values to the nonzeros of $\overline{A}$, then the result of the Theorem holds.

We now show that it is possible to assign values to the nonzeros in the overdetermined matrix A to make the element $q_{ij} \neq 0$

for every $r_i \in p_j$. Since the j-th column of Q depends only on the first j columns of A, we indicate only how nonzeros in the submatrix A_j should be assigned values. In the proof, we make use of the fact that each distinct perfect matching of a square matrix contributes a term to the determinant. Thus if a matrix has a unique perfect matching, then any assignment of nonzero values to the elements corresponding to the edges in the perfect matching will make the matrix nonsingular.

THEOREM 4.2. *There exists a single assignment of values to the nonzeros in A_j to make $q_{ij} \neq 0$ for every $r_i \in p_j$.*

Proof. Consider the structure of A_j shown in (2.1), and recall that $H(A_j)$ represents the structure of a matrix with full column rank. Hence the subgraph of $H(A_j)$ induced by the columns in $S_{j-1} \cup U_j$ has the Hall property. Thus we can find a column-perfect matching in this induced subgraph, and assign algebraically independent values to the nonzeros corresponding to the matched edges and the value zero to the unmatched edges. With this assignment of values, the submatrix of A_j induced by $S_{j-1} \cup U_j$ has full column rank. By Theorem 2.2, $q_{ij} = 0$ for every row $r_i \in \mathcal{R} \setminus p_j$. Since the nonzero values in $\underline{q}_j$ are determined only by the columns in P_j and rows in p_j, we need consider only how the submatrix A_P induced by the sets (p_j, P_j) should be assigned values.

Let $\underline{q}$ ($\underline{a}$) denote the restriction of $\underline{q}_j$ ($\underline{a}_j$) to the rows in p_j, and let $|P_j| \equiv K$. We order the columns in A_P in their natural ordering, and thus $\underline{a}$ is the last column in A_P. The bipartite graph $H(A_P)$ corresponding to A_P is a subgraph of B_j induced by the connected component whose row set is p_j and column set is P_j.

Since the vector $\underline{q}$ belongs to the linear space spanned by the columns of P_j, there exists a K-vector $\underline{y}$ such that

$$A_P \, \underline{y} = \underline{q}. \tag{4.1}$$

Further, the vector $\underline{q}$ is orthogonal to all the columns of P_j except $\underline{a}$. Thus

$$A_P^T \, \underline{q} = \underline{e}_K. \tag{4.2}$$

Combining these two equations, we obtain the symmetric system

$$A_P^T \, A_P \, \underline{y} = \underline{e}_K. \tag{4.3}$$

Our strategy will be to first predict the structure of $\underline{y}$ from (4.3), and then to obtain the structure of $\underline{q}$ from (4.1).

Replace each nonzero in A_P by a variable x_l. We will show how to assign values to the variables in $\underline{x}$ to make $q_{ij} \neq 0$ for every $r_i \in p_j$.

Let c be any column in P_j. By Theorem 3.4, there is a c-alternating path from c to c_j with respect to some column-perfect matching M in the bipartite graph $H(A_P)$. Let $r \in p_j$ be the row matched in M to c. Choose a subgraph $\hat{B}$ of $H(A_P)$ to consist of the edges on the c-alternating path from r to c_j, and the other matched edges in $H(A_P)$. Since columns on the c-alternating path are matched to rows on the path, $\hat{B}$ has a unique perfect matching. Let $\hat{A}$ be the submatrix of A_P with nonzeros corresponding to edges in $\hat{B}$. The submatrix $\hat{A}^T\hat{A}$ has an unique nonzero diagonal because of the unique column-perfect matching in $\hat{B}$. The adjacency graph $G(\hat{A}^T\hat{A})$ is the column-intersection graph of $\hat{A}$, i.e., its vertices are the columns of $\hat{A}$, and it has an edge (c_k, c_l) if the columns c_k and c_l have nonzeros in a common row of $\hat{A}$. Thus the c-alternating path from c to c_j in B_j induces an undirected path between c and c_j in $G(\hat{A}^T\hat{A})$. Hence the set $closure(\underline{e}_K)$ with respect to $G(\hat{A}^T\hat{A})$ contains all column vertices on the path from c to c_j.

Because of the nonzero diagonal in the matrix $\hat{A}^T\hat{A}$, we can assign values to $\underline{x}$ to make $\det(\hat{A}^T\hat{A})$ nonzero. Then by Theorem 4.1, $structure(\underline{y}) = closure(\underline{e}_K)$ includes the column c. We can repeat this argument for each column $c \in P_j$, to show that the component of $\underline{y}$ corresponding to column c is nonzero. Corresponding to each column c, we have identified a submatrix $\hat{A}^T\hat{A}$ with a nonzero determinant. Each determinant $\det(\hat{A}^T\hat{A})$ is a polynomial with integer coefficients in $\underline{x}$, and hence its roots lie in a set of measure zero. Since each determinant vanishes on a set of measure zero, the union of these K sets has measure zero. Thus we can assign values to $\underline{x}$ such that none of the determinants vanish, and then $structure(\underline{y}) = P_j$.

We now relate the structure of $\underline{q}_j$ to that of $\underline{y}$ by means of the transformation $A_P\underline{y} = \underline{q}$. By Cramer's rule, each element y_l

is the ratio of two determinants,

$$y_l = \det(A_P^T A_P|_l^K)/\det(A_P^T A_P),$$

where the matrix in the numerator is obtained by replacing column l of $A_P^T A_P$ by the right-hand-side vector $\underline{e}_K$. Since each determinant is a polynomial with integer coefficients in $\underline{x}$, the component y_l is a rational function in $\underline{x}$.

Since there is a path from $r_i \in p_j$ to c_j in the auxiliary graph B_j, $\underline{\alpha}_i{}^T$, the row of A_P corresponding to the i-th row of A, has at least one nonzero. Now $q_{ij} = \underline{\alpha}_i{}^T\, \underline{y}$ implies that q_{ij} has at least one nonzero term since $\underline{y}$ is full. Furthermore, each q_{ij} is a rational function of $\underline{x}$. Since a rational function vanishes on a set of measure zero, we can choose values for $\underline{x}$ such that $q_{ij} \neq 0$ simultaneously for every $r_i \in p_j$. ■

By the above Theorem, the adjacency list of a column vertex c_j in the graph H_Q is the row set p_j. We can represent by a bipartite graph $H_Q = (\mathcal{R}, \mathcal{C}, E_Q)$, the structure obtained by repeatedly applying Theorem 4.2 for every column $j = 1, \ldots, n$. Clearly, by construction, the structure of the orthogonal factor of a matrix in $\mathcal{A}$ is then contained in H_Q.

We can predict the structure of the triangular factor R by forming the structural product $Q^T A$. To represent the structure of R by means of a bipartite graph H_R, we describe the concept of a product bipartite graph. Let $H_1 = (\mathcal{R}, \mathcal{C}, E_1)$ and $H_2 = (\mathcal{R}, \mathcal{C}, E_2)$ be two bipartite graphs with common row and column sets. We number $\mathcal{R} = \{r_1, \ldots, r_m\}$ and $\mathcal{C} = \{c_1, \ldots, c_n\}$. The (upper triangular) product bipartite graph $\Phi = \Phi(H_1, H_2) = (\mathcal{C}, \mathcal{C}, E_\Phi)$ has its row and column vertices both numbered from 1 to n, and for $i \leq j$, has an edge (i, j) joining vertices i and j if and only if (r_k, c_i) is an edge in H_1 and (r_k, c_j) is an edge in H_2, for some $1 \leq k \leq m$. The bipartite graph $H_R = \Phi(H_Q, H(A))$ then represents the predicted structure of R.

We are now in a position to prove the major result in this paper.

THEOREM 4.3. *Let A be an $m \times n$ matrix with full column rank, where $m \geq n$, and let H_Q and H_R denote the structures of the orthogonal factors predicted, as described above, from the bipartite graph $H(A)$. There exists a matrix $A' \in \mathcal{A}$ with factors Q' and R' such that $H(Q') = H_Q$, and $H(R') = H_R$.*

Proof. Let each nonzero in A be assigned a variable x_l. We will first prove that the diagonal elements and zero elements of R are predicted correctly in H_R, and then prove that H_Q and H_R are simultaneously tight for some matrix $A' \in \mathcal{A}$.

The j-th column of Q, $\underline{q}_j$, belongs to the linear space spanned by columns in P_j, and is orthogonal to all these columns except $\underline{a}_j$. Hence $r_{jj} = \underline{q}_j^T \underline{a}_j$ is nonzero from the assumption of full column rank in A. We need to show that there is an edge (j, j) in H_R corresponding to this diagonal element. In the bipartite graph H_Q, the column vertex c_j is adjacent to all rows in p_j, by Theorem 4.2. In the bipartite graph $H(A)$, c_j is adjacent to some row in p_j, since by definition of the row set p_j, there is a path in B_j from every row in p_j to c_j with intermediate vertices belonging only to p_j and P_j. Now by its definition, the product bipartite graph H_R contains the edge (j, j).

Now consider a fixed element r_{ij}, where $1 \leq i \leq n-1$, and $j > i$.

If the columns $\underline{q}_i$ and $\underline{a}_j$ do not have a nonzero element in a common row, then these two columns are structurally orthogonal, and $r_{ij} = \underline{q}_i^T \underline{a}_j$ is zero. By Theorem 4.2, p_i is the the set of row vertices that c_i is adjacent to in H_Q, and by assumption, in $H(A)$ the vertex c_j is not adjacent to any vertex in p_i. By the definition of H_R, then it does not contain the edge (i, j), and thus the zero elements in H_R are predicted correctly.

Now consider the situation when $\underline{q}_i$ and $\underline{a}_j$ have a nonzero element in a common row. In this case, the edge (i, j) is present in H_R, and we need to show that values can be assigned to A to make the element r_{ij} nonzero. We proved in Theorem 4.2 that the element q_{ij} has the form $\underline{\alpha}_i^T \underline{y}$, where $\underline{\alpha}_i^T$ is the row of A_P which corresponds to the i-th row of A, and $\underline{y}$ is a vector whose components are rational functions in $\underline{x}$. Thus $r_{ij} = \underline{q}_i^T \underline{a}_j$ is also a rational function in $\underline{x}$, and the set

$$Z_{ij} = \{\underline{x} : \underline{q}_i^T \underline{a}_j = 0\}$$

has measure zero. Thus it is possible to choose values for $\underline{x}$ such that r_{ij} is nonzero.

We now show that it is possible to assign values to $\underline{x}$ to make the structures of the factors Q and R exactly equal to the predicted structures H_Q and H_R. Associate with each edge (r_i, c_j) of H_Q the sets

$$X_{jk} = \{\underline{x} : \det(\hat{A}^T \hat{A}) = 0\}, \qquad \text{and} \qquad Y_{ij} = \{\underline{x} : \underline{\alpha}_i{}^T \underline{y} = 0\},$$

where the index k ranges over every $c_k \in P_j$. For reasons given before, each set X_{jk}, Y_{ij}, and Z_{ij} (from the preceding paragraph) has measure zero. The union of all these sets corresponding to every edge in in H_Q and H_R, being a finite union, also has measure zero. It is thus possible to assign a set of values $\underline{x}'$ outside these sets to obtain a matrix $A' \in \mathcal{A}$ whose factors satisfy $H(Q') = H_Q$ and $H(R') = H_R$. ∎

5. ALGORITHMS AND PIVOTING

In this section we describe efficient algorithms for predicting the structures of the factors Q and R, discuss the structure of the Householder array, and consider structure prediction for orthogonal factorization with pivoting.

5.1. Algorithms for structure prediction

We assume that the D-M decomposition of $H(A)$ has been computed by means of a maximum matching. This step requires $\mathcal{O}(n^{1/2}\tau(A))$ time and $\mathcal{O}(\tau(A))$ space, where $\tau(A)$ is the number of edges in $H(A)$ ([6]).

The bipartite graph H_Q can be computed by identifying the adjacency lists of the column vertices c_j, in order from $j = 1$, ..., n. The adjacency list of c_j in H_Q is p_j, the set of rows which belong to the same connected component of B_j as c_j. The set p_j can be computed by an appropriate search of the graph $H(A)$, without forming B_j as follows. We search the adjacency lists of vertices in $H(A)$, starting from the vertex c_j, and continuing the search from each as yet unvisited row and column vertex reached. We can exclude rows belonging to s_{j-1} from the search, since such rows do not belong to B_j; similarly, we exclude columns numbered greater than c_j, since such columns also do not belong to B_j. By the definition of a Hall set, columns in S_{j-1} are adjacent only to rows in s_{j-1}, and thus these columns will not be reached by the search since rows in s_{j-1} are excluded.

The search from c_j can be implemented in $\mathcal{O}(\tau(A))$ time and space. Thus the structure of H_Q can be computed in $\mathcal{O}(n\tau(A))$ time using space $\mathcal{O}(\max\{\tau(A), \tau(Q)\})$, where $\tau(Q)$ is the number of edges in H_Q. This algorithm is an improvement on an $\mathcal{O}(mn^2(h+1))$-time algorithm described by Hare et al. [16], where h is the number of distinct, nonempty maximum Hall sets S_j. Note that $h = \mathcal{O}(n)$.

Now we turn to the computation of H_R. Since $\underline{r}_i^T = \underline{q}_i^T A$, the structure of the i-th row of R can be predicted from the structures of A and the i-th column of Q. This is an important advantage when only the structure of R is required, since then H_Q need not be stored.

Recall that for $i \leq j$, there is an edge (i,j) in H_R when (r_k, c_i) is an edge in H_Q and (r_k, c_j) is an edge in $H(A)$, for some $1 \leq k \leq m$. The adjacency list of c_i in H_Q is given by the set p_i. We can thus compute the structure of the i-th row of R by forming the union

$$(\cup_{r \in p_i} adj(r)) \cap \{c_i, \ldots, c_n\}.$$

This set can also be computed in $\mathcal{O}(\tau(A))$ time, and thus H_R can be computed in $\mathcal{O}(n\tau(A))$ time.

Since the time complexity of our structure prediction algorithms is the same as the complexity of symbolic factorization in sparse Cholesky factorization, these algorithms can be used practically for setting up data structures for orthogonal factors.

5.2. *The Householder array*

The following remarks concern a data structure that has been considered by George, Liu, and Ng [11]. When the orthogonalization is computed by means of Householder transformations, the orthogonal factor is not explicitly computed, but is implicitly stored in terms of the Householder vectors. Define an $m \times n$ Householder array $\mathcal{H}$ whose columns are the Householder vectors; this is a lower trapezoidal matrix. After these results were mentioned without proofs in [21], Ng and Peyton [19] have shown that when either A is a strong Hall matrix with a nonzero diagonal, or A is a Hall matrix with columns and rows numbered consistent with its D-M decomposition (i.e., columns in the square subgraph numbered before columns in the overdetermined subgraph; within the square subgraph, columns in predecessors numbered before columns in a strong Hall component;

rows numbered such that r_i is the row matched to column c_i for $i = 1, \ldots, n$), then the structure of the j-th column of $\mathcal{H}$ is obtained from the adjacency set of c_j in the bipartite graph H_Q by omitting the superdiagonal rows, i.e., $p_j \cap \{r_j, r_{j+1}, \ldots, r_m\}$. Hence in these cases, the lower trapezoidal structure of Q and the Householder array are identical. It can also be seen from examples that when A is a Hall matrix with columns in some arbitrary ordering, then the structure of the j-th column of the Householder array may not be contained in the structure of the j-th column of Q.

The undirected adjacency graph of the triangular factor R of a strong Hall matrix A is a chordal graph with the column ordering of A corresponding to a perfect elimination ordering, since it has the same structure as the transposed Cholesky factor of $A^T A$. Then George, Liu, and Ng show that the row structure of $\mathcal{H}$ can be obtained in terms of an appropriately defined path in the elimination tree of R. Unfortunately, when A is a Hall matrix and not strong Hall, then the adjacency graph of R is no longer a chordal graph with vertices in a perfect elimination ordering. Thus there is no elimination tree corresponding to R.

Now consider the directed graph $D(R)$ with vertices numbered from 1 to n, and for $i < j$, an edge (i, j) if $r_{ij} \neq 0$. We could form the transitive reduction of $D(R)$ and then ask if a similar path characterization may be obtained for the row structure of $\mathcal{H}$. The answer turns out to be no again.

5.3. Orthogonal factorization with pivoting

We show by means of an example that when column pivoting is incorporated into sparse orthogonal factorization, there cannot exist structures H_Q or H_R which satisfy the requirements (1.1) and (1.2)$'$. Let

$$A = \begin{pmatrix} \times & & & \\ \times & \times & & \\ & \times & \times & \\ & & \times & \times \end{pmatrix}.$$

Depending on the numerical values of the nonzero elements, when A is factored using column pivoting, three among the possible structures for Q are

$$\begin{pmatrix} & & & \times \\ & & \times & \\ & \times & & \\ \times & & & \end{pmatrix}, \begin{pmatrix} \times & \times & \times & \times \\ \times & \times & \times & \times \\ & \times & \times & \times \\ & & \times & \times \end{pmatrix},$$

and

$$\begin{pmatrix} & & \times & \times \\ \times & & \times & \times \\ \times & & \times & \times \\ & \times & & \end{pmatrix}.$$

Since the structure of the first column of Q is the structure of the column of A which is chosen to be factored first, it cannot be full. However, the smallest structure that contains the three possible structures shown above for Q is a full matrix. Similarly, it can be shown that the smallest structure that contains all possible structures of R is a full upper triangular matrix. But since the first column of Q contains at most two nonzeros in consecutive rows of A, in the structural product $Q^T A = R$, the first row of R cannot be full.

Hence for orthogonal factorization with column pivoting, we will have to be satisfied with the weaker requirements (1.1) and (1.2).

6. CONCLUSIONS

The results in this paper have important implications for computing the orthogonal factorization of sparse matrices.

For well-conditioned matrices, these results stress the importance of first computing the block triangular form of the given matrix, and then factoring its strong Hall components rather than the given matrix. Important advantages then accrue from the perspective of designing data structures to represent the structures of the factor matrices. The adjacency graph of the triangular factor of a strong Hall component is a chordal graph with vertices ordered in a perfect elimination ordering, and thus elimination trees and clique trees may be used to represent its structure. The structure of the Householder array (which implicitly represents the orthogonal matrix) can then be compactly represented in terms of paths in the elimination tree. On the other hand, if the matrix A is not strong Hall, then the adjacency graph of its triangular factor is not a chordal graph with vertices in a perfect elimination ordering. Hence there is no elimination tree or clique tree representation, and no corresponding compact representation for its Householder array.

For rank-deficient and ill-conditioned matrices, it no longer suffices to factor only the strong Hall components. The techniques described here are potentially useful in predicting the structures of the factors within the context of orthogonal factorization with column pivoting and rank-revealing orthogonal factorization. The Dulmage-Mendelsohn decomposition can be used to guide the selection of the pivot column. Such an algorithm for orthogonal factorization with pivoting would be similar in spirit to algorithms for sparse unsymmetric Gaussian elimination with pivoting in which combinatorial structure prediction and numerical computations are interleaved.

The above discussion points out the importance of the block triangular form of a sparse matrix in computing its orthogonal factorization. An algorithm for computing this form via the Dulmage-Mendelsohn decomposition has been implemented in [22]. The block triangular form has also been employed in sparse Matlab [15] to solve unsymmetric systems of linear equations.

REFERENCES

1 R. K. Brayton, F. G. Gustavson, and R. A. Willoughby. Some results on sparse matrices. *Math. Comput.*, 24:937–954, 1970.

2 R. A. Brualdi. Term rank of the direct product of matrices. *Can. J. Math.*, 18:126–138, 1966.

3 R. A. Brualdi and H. J. Ryser. *Combinatorial Matrix Theory*. Cambridge University Press, 1991.

4 T. F. Coleman, A. Edenbrandt, and J. R. Gilbert. Predicting fill for sparse orthogonal factorization. *J. Assoc. Comput. Mach.*, 33:517–532, 1986.

5 I. S. Duff. MA28—a set of FORTRAN subroutines for sparse unsymmetric linear equations. Technical Report AERE R-8730, Harwell, England, 1977.

6 I. S. Duff and T. Wiberg. Implementations of $O(n^{1/2}\tau)$ assignment algorithms. *ACM Trans. on Math. Software*, 4:267–287, 1988.

7 A. L. Dulmage and N. S. Mendelsohn. Coverings of bipartite graphs. *Can. J. Math.*, 10:517–534, 1958.
8 A. L. Dulmage and N. S. Mendelsohn. A structure theory of bipartite graphs of finite exterior dimension. *Trans. Roy. Soc. Can. Sec III*, 53:1–13, 1959.
9 A. L. Dulmage and N. S. Mendelsohn. Two algorithms for bipartite graphs. *J. Soc. Ind. Appl. Math.*, 11:183–194, 1963.
10 J. A. George and M. T. Heath. Solution of sparse linear least-squares problems using Givens rotations. *Linear Algebra and its Appl.*, 34:69–83, 1980.
11 J. A. George, J. W. H. Liu, and E. G. Y. Ng. A data structure for sparse QR and LU factorizations. *SIAM J. Sci. Stat. Comput.*, 9:100–121, 1988.
12 J. A. George and E. G. Y. Ng. Symbolic factorization for sparse Gaussian elimination with partial pivoting. *SIAM J. Sci. Stat. Comput.*, 8:877–898, 1987.
13 J. R. Gilbert. Predicting structure in sparse matrix computations. Technical Report TR-86-750, Computer Science, Cornell University, May 1986.
14 J. R. Gilbert. An efficient parallel sparse partial pivoting algorithm. Technical Report 88/45052-1, Center for Computer Science, Chr. Michelsen Institute, Bergen, Norway, Aug. 1988.
15 J. R. Gilbert, C. B. Moler, and R. S. Schreiber. Sparse matrices in MATLAB: design and implementation. Technical Report CSL-91-4, Xerox Palo Alto Research Center, Palo Alto, CA, July 1991.
16 D. R. Hare, C. R. Johnson, D. D. Olesky, and P. van den Driessche. Sparsity analysis of the QR factorization. Technical Report DMS-536-IR, Mathematics and Statistics, University of Victoria, Victoria, B. C., Canada, V8W 3P4, Nov. 1991. To appear in SIAM J. Matrix. Anal. Appl. (This is a revision of a Mar. 1990 preprint by Johnson, Olesky, and van den Driessche.)
17 D. M. Johnson, A. L. Dulmage, and N. S. Mendelsohn. Connectivity and reducibility of graphs. *Can. J. Math.*, 14:529–539, 1962.
18 L. Lovász and M. D. Plummer. *Matching Theory.* North Holland, Amsterdam, 1986.
19 E. G. Y. Ng and B. W. Peyton. A tight and explicit representation of Q in sparse QR factorization. Technical Report ORNL/TM-12059, Oak Ridge National Laboratory,

Oak Ridge, TN, May 1992.

20 A. Pothen. *Sparse Null Bases and Marriage Theorems.* PhD thesis, Cornell University, Ithaca, New York, 1984.

21 A. Pothen. Predicting the structure of sparse orthogonal factors. Technical Report CS-91-27, Computer Science, The Pennsylvania State University, Sep. 1991.

22 A. Pothen and C. J. Fan. Computing the block triangular form of a sparse matrix. *ACM Trans. on Math. Software*, 16:303–324, Dec. 1990.

23 D. J. Rose. A graph-theoretic study of the numerical solution of sparse positive definite systems of linear equations. In R. C. Read, editor, *Graph Theory and Computing*, pages 183–217. Academic Press, 1972.

24 D. J. Rose and R. E. Tarjan. Algorithmic aspects of vertex elimination on directed graphs. *SIAM J. Appl. Math.*, 34:176–197, 1978.