A NOTE ON THE COLUMN ELIMINATION TREE

JOHN R. GILBERT[†] AND LAURA GRIGORI[‡]

Abstract. This short communication considers the LU factorization with partial pivoting and shows that an *all-at-once* result is possible for the structure prediction of the column dependencies in L and U. Specifically, we prove that for every square strong Hall matrix A there exists a permutation P such that every edge of its column elimination tree corresponds to a symbolic nonzero in the upper triangular factor U. In the symbolic sense, this resolves a conjecture of Gilbert and Ng [6].

Key words. column elimination tree, sparse partial pivoting, structure prediction, lower bounds

AMS subject classifications. 65F50, 65F05, 68R10

1. Introduction. Sparsity in matrix computations offers opportunities to save memory space by storing only non-zero elements, shorten execution times by eliminating computations on zeros and exploit parallelism exposed by independent non-zero structures. Exploiting these opportunities often relies on a symbolic computation phase that predicts as accurately as possible which elements will have or can have nonzero values during the numerical computation itself, based only on the nonzero structure of the input matrix.

In LU factorization with partial pivoting, a square matrix A is factored as $PA = LU$, where P is a permutation matrix that depends on the values of the nonzeros of A and cannot be predicted only from the nonzero structure of A. Two structure prediction questions have been studied for this problem. The first is to predict bounds on the nonzero structure of the factors L and U. The second is to predict which columns of L and U depend directly or indirectly on which earlier columns. We restrict our attention to the class of matrices that satisfy an irreducibility condition called the *strong Hall* property.

George and Ng [4] developed upper bounds on the nonzero structure of L and U by employing a *row merge graph*. Gilbert and Ng [6] showed that this upper bound is as tight as possible in what they called "the exact sense." This means that, given the nonzero structure of a strong Hall matrix, for every edge in the row merge graph there is a choice of values for the nonzeros of A and a pivoting permutation P such that the corresponding element of L or U is nonzero. This is a *one-at-a-time* result [6]: any single position in the predicted structure can be made nonzero, but it may be the case that no single choice of nonzero values makes all the predicted elements nonzero at once.

The *column elimination tree* is a tree whose vertices are the columns of A, and whose edges correspond to potential dependencies between columns (a complete definition is below.) Gilbert and Ng [6] showed that if k is the parent of j in the column elimination tree of a strong Hall matrix A, there exists a choice of nonzero values of A that will make column j update column k during factorization with partial pivoting—that is, a choice of nonzero values for A that will make $u_{jk} \neq 0$. This is again a one-at-a-time result.

[†]Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, CA 94304 (gilbert@parc.xerox.com). The work of this author was performed during a visit to CERFACS, Toulouse, France.

[‡]INRIA, Domaine de Voluceau Rocquencourt - B.P. 105, 78153 Le Chesnay Cedex - France, currently visiting Lawrence Berkeley National Laboratory (LGrigori@lbl.gov). The research of this author was conducted as a Ph.D. student at INRIA Lorraine - Université Henri Poincaré, France.

A stronger statement would be an *all-at-once* result, showing that *all* the predicted positions can be made nonzero for the same input values. Unlike the case of sparse QR factorization, no tight all-at-once prediction is possible for the structure of L and U. The purpose of this short communication is to show that if we consider only the edges of the column elimination tree, an all-at-once result is possible in the symbolic sense. We prove that for every square strong Hall matrix A, there exists a permutation P such that every edge of the column elimination tree corresponds to a symbolic nonzero in the upper triangular factor U of A with partial pivoting. This resolves a variant of a conjecture of Gilbert and Ng [6].

Our result is *symbolic* in the sense that we assume that addition or subtraction of nonzeros always yields a nonzero result. Gilbert and Ng [6] also consider what they call *exact* results; we discuss this further in the conclusion.

A motivation for the current result is its impact on solvers that use the column elimination tree to model factorization in parallel. In solvers like the one described by Gilbert [5] and in the shared memory version of SuperLU [3], the tasks are scheduled dynamically on processors by using the precedence given by the column elimination tree. Our result shows that, in fact, for every strong Hall nonzero structure there is a matrix for which every dependency in the column elimination tree is a real constraint on the order of computation of the columns of the factor.

The next section presents background results and notation used in the paper. Section 3 introduces new results on the structure of the matrix during elimination. These results help to prove the all-at-once structure prediction of the column elimination tree. Section 4 concludes the paper.

2. Background. Let $A = (a_{rc})$ be a square, possibly unsymmetric, sparse $n \times n$ matrix which is to be factored as $PA = LU$ using partial pivoting.

In the following we introduce the commonly used tree and graph structures, the strong Hall property, a previously published theorem and lemma that will subsequently be used in our proofs. Most of our notation is similar or identical to that of Gilbert and Ng [6].

The *column intersection graph* $G_\cap(A)$ is undirected and has n vertices (one for each column) and an edge (i, j) if there is an r such that $a_{ri} \neq 0$ and $a_{rj} \neq 0$. This graph is equal to the graph of $A^T A$, unless there is numerical cancellation; in general $G(A^T A) \subseteq G_\cap(A)$.

The *directed graph* $G(A)$ has n vertices and an edge (i, j) for each nonzero element a_{ij}. The *bipartite graph* $H(A)$ has $2n$ vertices (one for each row and one for each column) and an edge (r', c) whenever a_{rc} is nonzero. In the bipartite graph, we use primes on the names of row vertices. For any graph G and vertex v, we write $Adj(v, G)$ to represent the set of vertices w such that (v, w) is an edge of G.

The *elimination tree* structure (*etree*) was first introduced for the Cholesky factorization of symmetric positive definite (SPD) matrices [8]. If L is the Cholesky factor of the SPD matrix A, then this tree has n vertices, and k is the parent of j if and only if $k = \min\{r > j : l_{rj} \neq 0\}$. Later the elimination tree was adapted to the LU factorization with partial pivoting [5]; the *column elimination tree* is the elimination tree of the column intersection graph $G_\cap(A)$, or equivalently the elimination tree of $A^T A$ if there is no numerical cancellation when computing or factoring $A^T A$.

A *strong Hall graph* is a bipartite graph with m rows and n columns that has the strong Hall property [2, 6]: every set of k column vertices is adjacent to at least $k+1$ row vertices, for all $1 \leq k < n$. A square matrix has the strong Hall property if and only if it is a *fully indecomposable matrix*, that is, there are no two permutations P

and Q such that PAQ is block triangular.

Before introducing the necessary theorem and lemma, let us elaborate on an additional definition, that of a sequence of bipartite graphs which model the structure of L and U during the elimination. Let $H_0 = H(A)$ be the bipartite graph of A. Suppose a_{rc} is nonzero and is chosen as pivot at step 1. The *deficiency* of the edge (r', c) of H_0 is defined as the set of edges

$$\{(i', j) : c \in Adj(i', H_0), j \in Adj(r', H_0), \text{ and } j \notin Adj(i', H_0)\}$$

It corresponds to the zero elements of A that become nonzero when a_{rc} is used as a pivot in Gaussian elimination.

Knowing the sequence of pivoting elements $(r'_1, c_1), (r'_2, c_2), \ldots, (r'_{n-1}, c_{n-1})$, we can construct a sequence of bipartite graphs $H_0, H_1, \ldots, H_n$, where H_i describes the structure of the $(n-i)\times(n-i)$ Schur complement remaining after step i. The bipartite graph H_i of the $(n-i) \times (n-i)$ submatrix that remains after eliminating (r'_i, c_i) is obtained as follows: delete from H_{i-1} vertices r'_i and c_i and all edges incident to them, then add the edges in the deficiency of (r'_i, c_i). The *bipartite filled graph* $H^+(A)$ is the bipartite graph containing all the edges of all H_i.

If the diagonal elements of A are nonzero, and the pivots are chosen in the order $(1', 1)$, $(2', 2), \ldots, (n', n)$, then we write $G^+(A)$ for the *filled graph* of A, which is obtained from $H^+(A)$ by merging each row vertex v' with its corresponding column vertex v. The *filled column intersection graph* $G^+_\cap(A)$ is the filled graph of the column intersection graph of A, that is, $G^+(G_\cap(A))$. If H_0 is the bipartite graph of A, then $G_\cap(H_0)$ is equivalent to $G_\cap(A)$ ($G^+_\cap(H_0)$ is equivalent to $G^+_\cap(A)$).

With these definitions at hand we now mention two results on which ours is based.

THEOREM 2.1 (Gilbert and Ng [6]). *Let H_0 be a bipartite graph and let (r', c) be an edge of H_0. Let H_1 be the bipartite graph resulting from the elimination of edge (r', c). If H_0 has the strong Hall property, then H_1 also has the strong Hall property.*

For the following lemma (called the *fill path* lemma), a *path* is a sequence of edges $P = [(v_0, v_1), (v_1, v_2), \ldots, (v_{p-1}, v_p)] = [v_0, v_1, \ldots, v_p]$ in which all the vertices are distinct. The length of this path P is p.

LEMMA 2.2 (Rose, Tarjan and Lueker [7]). *Let G be a directed or undirected graph whose vertices are the integers 1 through n, and let G^+ be its filled graph. Then (x, y) is an edge of G^+ if and only if there is a path in G from x to y whose intermediate vertices are all smaller than* $\min(x, y)$.

We conclude this section by presenting several previous results, outlining the role of the different graphs, introduced here, in the structure prediction of L and U. If the matrix A can be factored without row or column interchanges, then $G(L + U)$ is equal to $G^+(A)$ unless numerical cancellation occurs.

If pivoting is necessary during the Gaussian elimination, then only upper bounds on the structures of L and U can be predicted. The filled column intersection graph of A represents such an upper bound: $G(U) \subseteq G^+_\cap(A)$, and a slightly different representation of L is also a subgraph of $G^+_\cap(A)$. Thus the graph $G^+_\cap(A)$ contains an edge for each element of L and U that can possibly be nonzero during the numerical computation. If the matrix A has the strong Hall property, then the filled column intersection graph is a tight exact bound for the nonzero structure of U [6].

3. Structure prediction and the column elimination tree.

3.1. An example. The elimination tree plays an important role in the parallel sparse Cholesky factorization of symmetric positive definite matrices. This tree describes all the dependencies between column computations, and it represents the task scheduling model of almost all parallel sparse Cholesky solvers.

In the LU factorization with partial pivoting, the column elimination tree predicts all potential dependencies between columns, and hence it can be used as a task scheduling model in the unsymmetric case. For example, in the shared memory version of SuperLU [3] this tree helps identifying two levels of parallelism in the LU factorization with partial pivoting. As described in [3], a first level of parallelism exploits the property that computations in disjoint subtrees are independent thus leading to assigning disjoint subtrees to different processors; a second level of parallelism sequences in a pipelining manner the computation of dependent columns in a subtree. This level is especially useful in the superior part of the tree, where there are more idle processors than disjoint subtrees.

The nonzero structure of U cannot be, in general, exactly determined prior to the numerical factorization, and thus the column elimination tree can overestimate the real column dependencies. Consider for example the strong Hall matrix A in figure 3.1 with its bipartite graph H_0, the filled column intersection graph $G_\cap^+(H_0)$ and its column elimination tree.

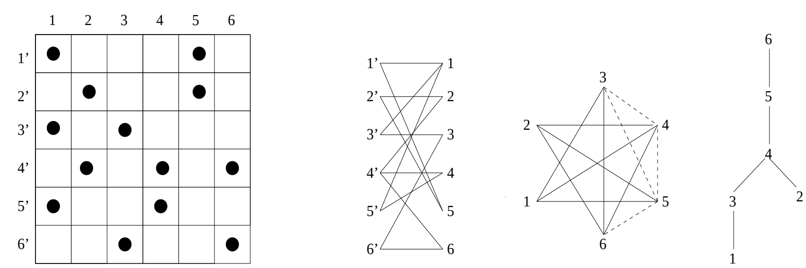

FIG. 3.1. *Matrix example A, the bipartite graph H_0, the filled column intersection graph $G_\cap^+(H_0)$ and its column elimination tree. The dotted lines in the filled column intersection graph $G_\cap^+(H_0)$ represent fill-in.*

Suppose that at the first elimination step the diagonal element is used as pivot. This means that the element u_{13} is zero and there is no dependency between the computations of columns 1 and 3. In other words, the dependency between the nodes 1 and 3 in the column elimination tree of A corresponds to an overestimation of the real dependencies.

Let us now analyze the later stages of the elimination. Consider the matrix $P_1 A$ in figure 3.2 (P_1 describes the first elimination step), the bipartite graph H_1 resulting from the elimination of edge $(1', 1)$, followed by its filled column intersection graph $G_\cap^+(H_1)$ and the corresponding column elimination tree.

We note that while the edge $(3, 4)$ is present in the filled column intersection graph of H_0, it does not belong to the filled column intersection graph of H_1, and thus the structures of these two graphs are different. Hence, in general the graph $G_\cap^+(H_1)$ cannot be simply obtained by deleting the vertex 1 and its incident edges from the graph $G_\cap^+(H_0)$.

As a consequence, the structure of the column elimination trees related to the elimination graphs H_i can change from one elimination step to another. In our example, after the first step of elimination there is no potential dependency between the computations of columns 3 and 4, and thus the edge $(3, 4)$ is not present in the column elimination tree of H_1, and is replaced by the edge $(3, 5)$.

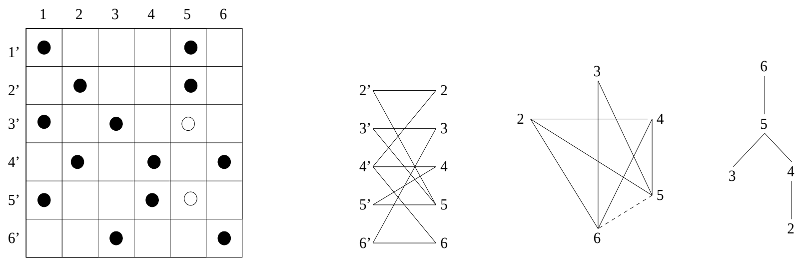

FIG. 3.2. *Matrix $P_1 A$ (including the deficiency of $(1', 1)$ represented by o), the bipartite graph H_1, the filled column intersection graph $G^+_\cap(H_1)$ and its column elimination tree.*

This simple example shows that the column elimination tree may overestimate the dependencies between columns. However, it has been shown that for a strong Hall matrix, this is the tightest information we can obtain before the numerical factorization of A. In other words, for each edge of the column elimination tree, there exists a choice of numerical values of A such that this edge corresponds to a real column dependency.

3.2. Main result. In this section we prove the main result of the paper, which is that every strong Hall nonzero pattern admits a pivoting permutation for which every edge of the column elimination tree corresponds to a symbolic column dependency. We first prove a lemma saying that there is a choice of pivot element such that the first elimination step creates the correct dependency (corresponding to an edge in the column elimination tree) for the first column and also does not change the structure of the filled column intersection graph. The lemma essentially says that (with this pivoting order) we never learn anything more about the uncomputed rows of U than we knew from $G^+_\cap$ at the beginning. We then prove the main theorem by induction.

We make two observations about symbolic elimination. First, the fact that $i+1$ is the least-valued vertex in H_i implies that there is no fill-in edge having $i+1$ as an endpoint. This gives us

$$Adj(i+1, G^+_\cap(H_i)) = Adj(i+1, G_\cap(H_i)) \tag{3.1}$$

$$= \{v : i+1 \in Adj(t', H_i) \text{ and } v \in Adj(t', H_i)\}. \tag{3.2}$$

Second, the fill path lemma implies that the vertices in the set $\{i+1\} \cup Adj(i+1, G^+_\cap(H_i))$ form a complete subgraph.

The next lemma shows that when we pivot on an element that is not the only element in its column, enough fill is added to preserve the structure of the filled column intersection graph. For each vertex i, we denote its parent in the column elimination tree by parent$[i]$; by definition this is $\min\{j > i : j \in Adj(i, G^+_\cap(H_0))\}$.

LEMMA 3.1. *Let H_0 be the structure of a square matrix A with at least two nonzero elements in column 1. Let P_1 be the permutation matrix that interchanges row r' with row 1 such that the edge $(1, \text{parent}[1])$ of the elimination tree of $G_\cap(A)$ corresponds to a nonzero in the upper triangular factor U. If H_1 is the bipartite graph resulting from the elimination of edge $(r', 1)$, then the filled column intersection graph of H_1 is obtained from the filled column intersection graph of H_0 just by deleting vertex 1 and its incident edges. That is,*

$$G^+_\cap(H_1) = G^+_\cap(H_0) - \{1\} \tag{3.3}$$

Proof. We will prove that the deficiency set of $(r', 1)$ introduces all the edges and preserves all the paths that can disappear by the deletion of row r' and column 1, while constructing the graph H_1. We will also show that this deficiency set does not introduce new edges or new fill paths in $G_\cap^+(H_1)$ compared to $G_\cap^+(H_0)$.

Let us analyze what happens when adding the deficiency of $(r', 1)$. Let $S = Adj(1, H_0)$ be the set of row indices of nonzeros in column 1. From the lemma statement, recall that column 1 is adjacent to at least two row vertices, so that r' is not the only element in set S.

By using the definition of the deficiency of $(r', 1)$, for each $t' \in S$ such that $t' \neq r'$ and for each edge $v \neq 1$ adjacent to r' in H_0, we see that v belongs to $Adj(t', H_1)$. For each two vertices $v_1, v_2 \in Adj(t', H_1)$, by using the definition of the column intersection graph, we see that (v_1, v_2) is an edge of $G_\cap(H_1)$.

Let us make an analysis depending on the origin of vertices v_1, v_2. First, if v_1, v_2 are adjacent to r' in H_0, (that is $v_1, v_2 \in Adj(r', H_0)$) the fact that (v_1, v_2) is an edge of $G_\cap(H_1)$ proves that the deletion of the row r' does not change the structure of $G_\cap(H_1)$ compared to the structure of $G_\cap(H_0)$.

Second, if $v_1 \in Adj(r', H_0)$ and $v_2 \in Adj(t', H_0)$, then v_1 and v_2 are both adjacent to 1 in the column intersection graph of H_0, so $v_1, v_2 \in Adj(1, G_\cap(H_0))$. By using the observation at the beginning of this section, we see that (v_1, v_2) belongs to $G_\cap^+(H_0)$. This proves that the deficiency set does not introduce new edges in $G_\cap^+(H_1)$.

Using this analysis of edges introduced in H_1, we can easily check that

$$Adj(1, G_\cap(H_0)) - \{v\} \subseteq Adj(v, G_\cap(H_1)), \forall v \in Adj(r', H_0), v \neq 1. \tag{3.4}$$

Suppose that $[x_1, \ldots, x_r]$, $r > 2$ is a fill path in $G_\cap^+(H_0)$ and has 1 as an intermediate vertex. This means that $x_k < \min\{x_1, x_r\}$ for all $k = 2, \ldots, r-1$, and the edge (x_1, x_r) belongs to $G_\cap^+(H_0)$. Suppose that $x_k = 1, k > 1, k < r$. By using relation (3.4), we see that $x_{k-1}, x_{k+1} \in Adj(\text{parent}[1], G_\cap(H_1))$. If $x_{k-1} = \text{parent}[1]$ or $x_{k+1} = \text{parent}[1]$, it is evident that the fill path is preserved, since we can suppress 1 from the path while preserving adjacency in the path. Otherwise, vertex 1 can be replaced by parent[1] in the path $[x_1, \ldots, x_{k-1}, \text{parent}[1], x_{k+1}, \ldots, x_r]$. By using the definition of the column etree, we see that $x_{k-1}, x_{k+1} \geq \text{parent}[1]$, and this shows that $[x_1, \ldots, x_r]$ is a fill path in $G_\cap^+(H_1)$. This proves that all the fill paths are preserved in $G_\cap^+(H_1)$ and no new fill path is introduced. □

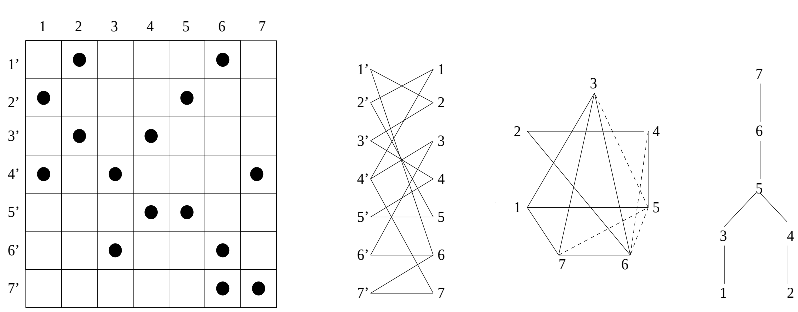

FIG. 3.3. *Matrix example A, the bipartite graph H_0, the filled column intersection graph $G_\cap^+(H_0)$ and its column elimination tree. The dotted lines in the filled column intersection graph $G_\cap^+(H_0)$ represent fill-in.*

Figure 3.3 shows a matrix example A, its bipartite graph H_0, followed by the filled column intersection graph $G_\cap^+(H_0)$ with its column elimination tree. Figure 3.4 presents the permuted matrix P_1A, the bipartite graph H_1 with its filled column intersection graph $G_\cap^+(H_1)$ and the corresponding column elimination tree.

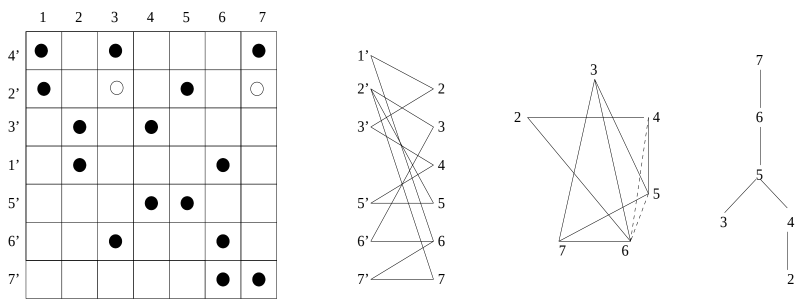

FIG. 3.4. *Matrix P_1A (including the deficiency of $(4',1)$ represented by o), the bipartite graph H_1, the filled column intersection graph $G_\cap^+(H_1)$ and its column elimination tree.*

Consider the elimination of edge $(4',1)$ in matrix example A, figure 3.3. The vertices 3 and 7 are adjacent to $4'$ in the bipartite graph H_0. Deleting the row $4'$ has as consequence that the edge $(3,7)$ dissapears from $G_\cap^+(H_1)$. By adding the deficiency of $(4',1)$, the edge $(3,7)$ is introduced in $G_\cap^+(H_1)$ due to row vertex $2'$. Now consider the vertex 3 adjacent to vertex $4'$ and the vertex 5 adjacent to vertex $2'$ in the bipartite graph H_0. By the permutation of row $4'$ with row $1'$, the edge $(3,5)$ is introduced in the filled column intersection graph $G_\cap^+(H_1)$. However, we remark that $(3,5)$ was already present in the filled column intersection graph $G_\cap^+(H_0)$. Finally, we consider the fill path [5 1 3 6] in $G_\cap^+(H_0)$ which is preserved in $G_\cap^+(H_1)$ in a compact form [5 3 6].

The next theorem is the main result of this paper. It proves the conjecture of Gilbert and Ng [6] in the symbolic sense, that is, if we assume that zeros are introduced only by explicit elimination and not by cancellation.

THEOREM 3.2. *Let A be an unsymmetric square sparse matrix having the strong Hall property. There is a permutation P such that every edge of the elimination tree of $G_\cap(A)$ corresponds to a nonzero in the upper triangular factor U in symbolic sense, when the factorization $PA = LU$ is computed.*

Proof. We will prove this by induction. Let $H_0 = H(A)$ be the bipartite graph of A.

Initial phase. We show that there exists a permutation P_1 such that the element $u_{1,\text{parent}[1]}$ is nonzero.

Using relation (3.1), we see that $(1, \text{parent}[1])$ belongs to $G_\cap(H_0)$. There exists a row vertex r_1' such that $(r_1', 1)$ and $(r_1', \text{parent}[1])$ are edges of H_0. We choose r_1' as pivot, and P_1 describes this permutation. Row 1 is interchanged with row r_1'; therefore the element $u_{1,\text{parent}[1]}$ is nonzero.

Induction phase $(m-1 \to m)$. We suppose that there is a sequence of permutations $P_{m-1}, \ldots, P_1$ such that for all $k = 1 \ldots m-1$, $u_{k,parent[k]}$ is nonzero. We show that there is a permutation P_m such that the element $u_{m,parent[m]}$ is nonzero.

According to Theorem 2.1, at each elimination step k, the bipartite graph H_k is strong Hall because H_{k-1} is. In particular, this means that at each elimination step we have at least two elements as choices to pivot on (column vertex k is adjacent to at least two row vertices in the graph H_{k-1}).

Therefore, Lemma 3.1 applies, and says that at each elimination step k, the structure of the filled column intersection graph is preserved

$$G_\cap^+(H_k) = G_\cap^+(H_{k-1}) - \{k\} \quad 1 \leq k < m. \tag{3.5}$$

This relation shows that the induction hypothesis has as direct consequence that the structure of the filled column intersection graph was preserved until this step m of elimination. We can deduce that $(m, \text{parent}[m])$ belongs to $G_\cap^+(H_{m-1})$. Even more, relation (3.1) says that this edge belongs to $G_\cap(H_{m-1})$.

Thus, there is some vertex r'_m such that (r'_m, m) and $(r'_m, \text{parent}[m])$ are edges of H_{m-1}. We choose r'_m as pivot, and let P_m describe this permutation. The permutation of the row m with the row r'_m will make the element $u_{m,\text{parent}[m]}$ be nonzero.

Let $P = P_{n-1}, \ldots, P_1$ be the permutation matrix that includes the $n-1$ row interchanges. We have proved that every edge of the column elimination tree corresponds to a symbolic nonzero in the upper triangular factor U, when the factorization $PA = LU$ is computed with partial pivoting. □

4. Concluding remarks. The main result of this paper is Theorem 3.2, which gives an all-at-once structure prediction result, under the assumption that the matrix A is strong Hall. In the proof, we showed that if at each elimination step k the element $u_{k,\text{parent}[k]}$ is nonzero, the structure of the filled column intersection graph is preserved during the elimination. One way to interpret this result is that (for a strong Hall matrix) there is a pivot sequence for which the only information about the structure of U exposed by each elimination step is the single newly computed row. In other words, the elimination does not give progressively more partial information about the uncomputed rows of U than was available from $G_\cap^+$ at the beginning.

We remark that, in the proof of Theorem 3.2, the strong Hall property was used in only one place for each elimination step. We used the strong Hall property to conclude that at each step (except the last), there is always a choice of at least two elements to pivot on. One could ask whether the strong Hall property is necessary as well as sufficient for this.

Our result is symbolic, in the sense that we assume that during Gaussian elimination the result of adding or subtracting two nonzeros is never zero. A stronger result would be what Gilbert and Ng [6] called *exact*, which would assume only that the nonzero values in A were algebraically independent from each other; in other words, it would assume that any computed zeros were due to combinatorial properties of the nonzero structure rather than to coincidence in choice of values. We do not know whether the exact version of our main theorem holds or not, though we conjecture that it does. An exact version holds, for example, for the class of strong Hall matrices with exactly two nonzeros in every row and every column, because every elimination step creates exactly one new nonzero and that nonzero is algebraically independent of the other remaining nonzeros.

We conclude by mentioning an open problem: what is the case for non strong Hall matrices, either for the elimination tree or for the structures of L and U? In this case, it is known that $G_\cap^+(A)$ may not be a tight bound for U. Is there a tight bound on U? If so, does it share the property that there is no new information revealed during the elimination except the structure of the current row of U?

Acknowledgments. The authors thank the anonymous reviewers for their helpful comments and suggestions to improve the presentation of the paper.

REFERENCES

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

[2] R. A. Brualdi and B. A. Shader, *Strong Hall matrices*, SIAM J. Matrix Anal. Appl., 15 (1994), pp. 359–365.

[3] J. W. Demmel, J. R. Gilbert, and X. S. Li, *A Parallel Supernodal Approach to Sparse Partial Pivoting*, SIAM J. Matrix Anal. Appl., 20 (1999), pp. 720–755.

[4] A. George and E. Ng, *Symbolic factorization for sparse gaussian elimination with partial pivoting*, SIAM J. Sci. Stat. Comput., 8 (1987), pp. 877–898.

[5] J. R. Gilbert, *An efficient parallel sparse partial pivoting algorithm*, Tech. Report 88/45052-1, Christian Michelsen Institute, 1988.

[6] J. R. Gilbert and E. G. Ng, *Predicting Structure in Nonsymmetric Sparse Matrix Factorizations*, in Graph Theory and Sparse Matrix Computation, A. George, J. R. Gilbert, and J. W. H. Liu, eds., Springer Verlag, 1994, pp. 107–139.

[7] D. J. Rose and R. E. Tarjan, *Algorithmic aspects of vertex elimination on directed graphs*, SIAM J. Appl. Math., 34 (1978), pp. 176–197.

[8] R. Schreiber, *A new implementation of sparse Gaussian elimination*, ACM Trans. Math. Software, 8 (1982), pp. 256–276.