

A MINIMAL STORAGE IMPLEMENTATION
OF THE MINIMUM DEGREE ALGORITHM†

by

Alan George*

and

Joseph W.H. Liu**

Research Report CS-77-09

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

* Department of Computer Science
University of Waterloo
Waterloo, Ontario, Canada

** Computer Services Division
City of Mississauga
Mississauga, Ontario, Canada

† Research supported in part by Canadian
National Research Council grant A8111.

ABSTRACT

We describe an efficient implementation of the minimum degree algorithm, which experience has shown to be effective in finding low fill orderings for sparse positive definite systems. The algorithm is heuristic; at each step in the elimination the variable chosen to eliminate next is that which minimizes the fill suffered at that step. Thus, some representation of the partially factored matrix is required at each step of the ordering. Previous implementations have stored this representation in an explicit form, which requires a data structure which allows the matrix structure to change as the ordering proceeds. The implementation we describe in this paper works only with the graph of the original matrix, and all data structures used are fixed throughout the execution of the algorithm. In contrast to most previous implementations, the total storage needs of the algorithm are known before execution. Several effective techniques for speeding up the algorithm are described, and numerical experiments on some problems arising in finite element applications suggest that for these problems the execution time is $O(N)$, where N is the number of equations.

Introduction

Consider the symmetric positive definite system of linear equations

$$(1.1) \quad Ax = b,$$

where the N by N matrix A is sparse. If we solve (1.1) using Cholesky's method, A is first factored into LL^T , where L is lower triangular, and then we solve the triangular systems $Ly = b$ and $L^Tx = y$. Sparse matrices normally suffer some fill when they are factored, so $L + L^T$ is usually fuller than A . Since for any N by N permutation P , the matrix PAP^T is still symmetric and positive definite, we can still use Cholesky's method to solve the equivalent problem

$$(1.2) \quad (PAP^T)(Px) = Pb.$$

A judicious choice of P can drastically reduce fill, hence the interest in algorithms for finding such permutations.

A heuristic algorithm which experience has shown to be extremely effective in finding low-fill orderings is the so-called minimum degree algorithm [7]. It is a "local" algorithm which at each elimination step permutes the part of the matrix remaining to be factored so that a column (row) with the fewest nonzeros is in the pivot position. This implies that at each step of the algorithm we need a representation of the structure of the partially factored matrix. Previous implementations known to the authors [1, 3, 8] store some explicit representation of the partially factored matrix, which has two disadvantages. Since the structure changes as the elimination (or simulation thereof) proceeds, the data structure must be flexible enough to allow for such changes. Second, it is usually impossible to predict the maximum storage requirements for such implementations;

usually the storage requirement grows for a time and then tapers off near the conclusion of the ordering.

The implementation we describe in this paper operates only on the original matrix graph. The data structures remain fixed during the execution of the algorithm, and storage requirements consist of a small number of vectors of length N together with the storage needed for the graph.

An outline of the paper is as follows. In section 2 we introduce some graph theory notions which are required in describing the algorithm and its implementation. Section 3 contains a description of the basic algorithm, and section 4 contains some crucial refinements to the algorithm which appears to reduce its time complexity to $O(N)$ for finite element problems. Section 5 contains some programming details along with some numerical experiments, and section 6 contains our conclusions.

§2 Preliminaries

In this section we review some basic graph theory notions that are related to symmetric Gaussian elimination. We also consider graph theoretic ways of viewing the elimination process.

2.1 Some Graph Theory Terminology

A graph $G = (X, E)$ consists of a finite nonempty set X of nodes together with a prescribed edge set E of unordered pairs of distinct nodes. A graph $G' = (X', E')$ is a subgraph of $G = (X, E)$ if $X' \subset X$ and $E' \subset E$. For $Y \subset X$, $G(Y)$ refers to the subgraph $(Y, E(Y))$ of G , where $E(Y) = \{\{u, v\} \in E \mid u, v \in Y\}$.

Nodes x and y are said to be adjacent if $\{x, y\}$ is an edge in E .

For a subset Y of nodes, the adjacent set of Y is defined as

$$\text{Adj}(Y) = \{x \in X \setminus Y \mid \{x, y\} \in E \text{ for some } y \in Y\}.$$

If $Y = \{y\}$, we shall write $\text{Adj}(y)$ instead of the formally correct $\text{Adj}(\{y\})$.

The degree of a node x is the number of nodes adjacent to x , denoted by $|\text{Adj}(x)|$. Sometimes, we shall refer to $y \in \text{Adj}(x)$ as a neighbor of the node x .

A path of length ℓ is an ordered set of distinct nodes $(v_0, v_1, \dots, v_\ell)$ where $v_i \in \text{Adj}(v_{i-1})$, $1 \leq i \leq \ell$. A graph G is connected if there is a path connecting each pair of distinct nodes. If G is disconnected, it consists of two or more maximal connected subgraphs called components.

Let Y be a subset of the node set X . The component partitioning $C(Y)$ of Y is defined as:

$$C(Y) = \{Y' \subset Y \mid G(Y') \text{ is a connected component in the subgraph } G(Y)\}.$$

When $Y = X$, $C(X)$ is simply the set of component sets in the graph G .

A useful notion in the study of Gaussian elimination is the reachable set which we now define. Let S be a subset of X and $y \in X \setminus S$. The node y is said to be reachable from a node x through S if there exists a path (x, v_1, \dots, v_k, y) such that $v_i \in S$, for $1 \leq i \leq k$. Note that k can be zero, so that any adjacent node of y not in S is reachable from y through S .

The reachable set of y through S , denoted by $\text{Reach}(y, S)$, is then defined to be

$$\text{Reach}(y, S) = \{x \in X \setminus S \mid x \text{ is reachable from } y \text{ through } S\}.$$

We can extend this definition to subsets of X . Let $Y \subseteq X$ with $Y \cap S = \emptyset$.

The reachable set of Y through S is then

$$\text{Reach}(Y, S) = \{x \in X \setminus (S \cup Y) \mid x \text{ is reachable from some node } y \in Y \text{ through } S\}.$$

Note that $\text{Adj}(Y) \setminus S \subseteq \text{Reach}(Y, S)$. When $S = \emptyset$, it can be seen that

$$\text{Adj}(Y) = \text{Reach}(Y, \emptyset),$$

so that the reachable set concept may be regarded as a generalization of the adjacent set.

§2.2 The Elimination Process in terms of Elimination Graphs

In this section, we review the graph theory approach used by Parter [5] and Rose [7] to study the Gaussian elimination process. We first establish a correspondence between graphs and matrices.

Let A be an N by N symmetric matrix. The labelled undirected graph of A , denoted by $G^A = (X^A, E^A)$, is one for which X^A is labelled from 1 to N :

$$X^A = \{x_1, \dots, x_n\},$$

and $\{x_i, x_j\} \in E^A$ if and only if $A_{ij} \neq 0$. For any N by N permutation matrix P , the unlabelled graphs of A and PAP^T are the same, but the associated labellings differ.

Consider the symmetric factorization of the matrix A into LL^T . The elimination process applied to A can be interpreted as a sequence of graph transformations on G^A . Following Rose [7], we define an elimination graph as follows. Let $G = (X, E)$ be a graph and y be a node in G . The elimination graph of G by y , denoted by G_y , is the graph

$$(X \setminus \{y\}, E(X \setminus \{y\}) \cup \{\{u, v\} \mid u, v \in \text{Adj}(y)\}).$$

With this definition, the process of Gaussian elimination on a matrix A can be viewed as a sequence of elimination graphs

$$G_0, G_1, \dots, G_{N-1}$$

where $G_0 = G^A,$

and $G_i = (G_{i-1})_{x_i} = (X_i, E_i).$

Here the set of nodes $X_i = \{x_{i+1}, \dots, x_N\}$. The graph G_i precisely reflects

the structure of the matrix remaining to be factored after the i -th step of the Gaussian elimination. This interpretation provides insight into the elimination process, and it is useful in the study of the fill-in phenomenon.

We now relate properties of the triangular factor L of A with the elimination graph sequence. In the factor L , let

$$v_i = |\{L_{ji} : L_{ji} \neq 0, j > i\}| \quad \text{for } i = 1, \dots, N.$$

Lemma 2.1 [7] The quantity v_i is the degree of the node x_i in the elimination graph G_{i-1} . □

Since the number of off-diagonal nonzeros in the factor L is given by

$$\sum_{i=1}^N v_i,$$

and the number of multiplicative operations required for the factorization is

$$\frac{1}{2} \sum_{i=1}^N v_i(v_i+1),$$

the degree of the node x_i in G_{i-1} plays an important role in the storage and operation requirements of the elimination process. The minimum degree algorithm to be studied in section 3 is designed to reduce these requirements by a local minimization of the degrees.

§2.3 An Alternative View using Reachable Sets

In section 2.2, the quantities v_i in the factor L are related to node degrees in the sequence of elimination graphs. In this section, we provide a direct relation between these numbers v_i and the original graph G^A associated with the matrix A . This relationship has been established elsewhere but we include it here for completeness. The approach uses the notion of reachable sets introduced in section 2.1.

Let $G^A = (X^A, E^A)$ and $G^F = (X^F, E^F)$ where $F = L + L^T$. Here G^F is called the filled graph of G^A , and E^F consists of the edges E^A in G^A together with all the edges added during the factorization. Obviously $X^F = X^A$, and the edge sets E^A and E^F are related by the following lemma due to Parter [5].

Lemma 2.3 The unordered pair $\{x_i, x_j\} \in E^F$ if and only if $\{x_i, x_j\} \in E^A$, or $\{x_i, x_k\} \in E^F$ and $\{x_j, x_k\} \in E^F$ for some $k < \min\{i, j\}$. □

In order to relate the numbers v_i to the graph G^A , we characterize the edge set E^F using only G^A . The following results are quoted from [2]. Lemma 2.2 is equivalent to Lemma 4 in [8].

Lemma 2.2 Let $j > i$. The unordered pair $\{x_i, x_j\} \in E^F$ if and only if $x_j \in \text{Reach}(x_i, \{x_1, \dots, x_{i-1}\})$. □

Corollary 2.3 For $i = 1, \dots, N$,

$$v_i = |\text{Reach}(x_i, \{x_1, \dots, x_{i-1}\})|.$$
□

Not only does the reachable set concept characterize the qualities v_i , it also reflects the adjacency structure of the elimination graph G_i . Let G_0, G_1, \dots, G_{N-1} be the sequence of elimination graphs as defined by the nodes x_1, x_2, \dots, x_N . The next lemma follows from the definitions of elimination graphs and reachable sets, and it can be proved by induction.

Lemma 2.4 Let y be any node in the elimination graph $G_i = (X_i, E_i)$. The set of adjacent nodes of y in G_i is given by

$$\text{Reach}(y, \{x_1, \dots, x_i\}),$$

where the Reach operator is taken in the original graph. □

The above lemma is useful in the next section when the minimum degree algorithm is studied.

§3 The Minimum Degree Algorithm

§3.1 Description of the Algorithm Using Elimination Graphs

Following Rose [7], we describe the minimum degree algorithm using elimination graphs. Let $G_0 = (X, E)$ be an unlabelled graph.

Step 1 (Initialization) $i \leftarrow 1$

Step 2 (Minimum degree selection) In the elimination graph G_{i-1} , choose x_i to be a node such that

$$|\text{Adj}(x_i)| = \min_{y \in X_{i-1}} |\text{Adj}(y)|,$$

where $G_{i-1} = (X_{i-1}, E_{i-1})$.

Step 3 (Graph transformation) Form the new elimination graph

$$G_i = (G_{i-1})_{x_i}.$$

Step 4 (Loop or stop) $i \leftarrow i + 1$. If $i > |X|$, stop. Otherwise, go to step 2.

The above formulation of the algorithm involves the formation of the sequence of elimination graphs. An implementation of this description can be found in the Yale Sparse Matrix Package (Sherman [9]).

§3.2 Description of the Algorithm using Reachable Sets

In the description of the minimum degree algorithm in section 3.1, the sole purpose of step 3, the graph transformation, is to facilitate the selection of the next node from the new elimination graph. This step can be omitted if we can provide an alternate way to compute the degrees of the nodes in the elimination graph. Lemma 2.4 shows that the reachable set is the relevant concept to use.

The lemma relates the adjacency structure of the elimination graphs to that of the original graph. With this simple connection, we can restate the minimum degree algorithm in terms of reachable sets.

Step 1 (Initialization) $S \leftarrow \phi$.

$DEG(x) \leftarrow |Adj(x)|$, for $x \in X$.

Step 2 (Minimum degree selection) Pick a node $y \in X \setminus S$

where $DEG(y) = \min_{x \in X \setminus S} DEG(x)$.

Number the node y next and set $S \leftarrow S \cup \{y\}$.

Step 3 (Degree update)

$DEG(x) \leftarrow |Reach(x, S)|$ for $x \in X \setminus S$.

Step 4 (Loop or stop). If $S = X$, stop. Otherwise, go to step 2.

§3.3 Some Related Results on Reachable Sets

The observation in section 3.2 shows that the reachable sets deserve more detailed analysis. In this subsection, we establish some preliminary results in this direction.

Consider a graph $G = (X, E)$. Let S be a (possibly empty) subset of X and $C(S)$ be the component partitioning of S . (See section 2 for the definition).

Consider a node y in $X \setminus S$. Let

$$\{C_1, \dots, C_k\} = C(S)$$

be all the connected components in $G(S)$ with $y \in \text{Adj}(C_i)$, $1 \leq i \leq k$.

Note that the number k depends on the subset S and the node y . Define the neighborhood of y in S to be the subset

$$\text{Nbrhd}(y, S) = \bigcup_{i=1}^k C_i.$$

The following lemmas relate neighborhoods to reachable sets.

Lemma 3.1 $\text{Nbrhd}(y, S) = \bigcup_{T \subset S} \text{Reach}(y, T) \cap S.$

□

In other words, it can be expressed as:

$$\text{Nbrhd}(y, S) = \{s \in S \mid s \text{ is reachable from } y \text{ through a subset of } S\}.$$

Lemma 3.2 $\text{Reach}(y, S) = \text{Adj}(\text{Nbrhd}(y, S) \cup \{y\})$.

Proof Consider $u \in \text{Reach}(y, S)$. There exists a path (u, s_1, \dots, s_t, y) where $s_i \in S$, $1 \leq i \leq t$. If $t = 0$, $u \in \text{Adj}(y)$ so that $u \in \text{Adj}(\text{Nbrhd}(y, S) \cup \{y\})$. If $t \neq 0$, we have $s_1 \in \text{Nbrhd}(y, S)$ and thus $u \in \text{Adj}(\text{Nbrhd}(y, S))$.

On the other hand, let $v \in \text{Adj}(\text{Nbrhd}(y, S) \cup \{y\})$. Either $v \in \text{Adj}(y) \setminus S$ or $v \in \text{Adj}(s)$ for some $s \in \text{Nbrhd}(y, S)$. In both cases, v is reachable from y through S . \square

For convenience in later discussions, we introduce one more definition. The closure of y by S is defined by

$$\text{Closure}(y, S) = \text{Nbrhd}(y, S) \cup \{y\} \cup \text{Reach}(y, S).$$

It is clear that the closure is a disjoint union of

$$\text{Nbrhd}(y, S) \cup S$$

and $\text{Reach}(y, S) \cup \{y\} \subset X \setminus S.$

It is interesting to point out that "Reach", "Nbrhd", and "Closure" may be regarded as operators:

$$\text{Reach: } X \setminus S \rightarrow P(X \setminus S)$$

$$\text{Nbrhd: } X \setminus S \rightarrow P(S)$$

$$\text{Closure: } X \setminus S \rightarrow P(X),$$

where $P(\star)$ is the power set of the specified set.




Figure 1. A 9-node graph.

The definitions and results can be best illustrated by an example. Consider the graph in Figure 1. Let $S = \{s_1, s_2, s_3, s_4\}$ and $y \in S$. It can be seen that the corresponding value of k is two, with

$$C_1 = \{s_1\},$$

and

$$C_2 = \{s_2, s_4\}.$$

So, the neighborhood set is

$$\text{Nbrhd}(y, S) = \{s_1, s_2, s_4\}.$$

By lemma 3.2, we have

$$\begin{aligned} \text{Reach}(y, S) &= \text{Adj}(\{s_1, s_2, s_4, y\}) \\ &= \{a, b, c\}. \end{aligned}$$

The closure set is then

$$\text{Closure}(y, S) = \{s_1, s_2, s_4, y, a, b, c\}.$$

The following results contain observations that are crucial to our implementation of the minimum degree algorithm to be discussed later in the paper.

Lemma 3.3 Let $x \in X \setminus S$. If

$$\text{Adj}(x) \subset \text{Closure}(y, S)$$

then $\text{Nbrhd}(x, S) \subset \text{Nbrhd}(y, S)$.

Proof Consider any $s \in \text{Nbrhd}(x, S)$. There exists a path x, s_1, \dots, s_t, s , where $s_i \in S$ for $1 \leq i \leq t$. If $t = 0$, $s \in \text{Adj}(x) \cap S \subset \text{Closure}(y, S) \cap S = \text{Nbrhd}(y, S)$. If $t \neq 0$, s_1 will belong to $\text{Nbrhd}(y, S)$ so that $s \in \text{Nbrhd}(y, S)$. \square

Theorem 3.4 Let $x \in X \setminus S$. If

$$\text{Adj}(x) \subset \text{Closure}(y, S)$$

then $\text{Reach}(x, S) \subset \text{Reach}(y, S) \cup \{y\}$.

Proof Consider a node $w \in \text{Reach}(x, S)$ and $w \neq y$. There exists a path x, s_1, \dots, s_t, w , where $s_i \in S$ for $1 \leq i \leq t$.

If $t = 0$, $w \in \text{Adj}(x) \cap (X \setminus S) \subset \text{Closure}(y, S) \cap (X \setminus S)$, so that $w \in \text{Reach}(y, S)$. Otherwise, if $t \neq 0$, $s_1 \in \text{Adj}(x) \cap S \subset \text{Closure}(y, S) \cap S = \text{Nbrhd}(y, S)$. So a path in S can be traced from w to y ; in other words $w \in \text{Reach}(y, S)$. \square

Corollary 3.5 Let x be as in theorem 3.4. Then $|\text{Reach}(x, S)| \leq |\text{Reach}(y, S)|$.

Proof If $y \notin \text{Reach}(x, S)$, by theorem 3.4

$$\text{Reach}(x, S) \subset \text{Reach}(y, S),$$

so that the result follows. On the other hand, assume $y \in \text{Reach}(x, S)$.

Then $x \in \text{Reach}(y, S)$ and we have $\text{Reach}(x, S) \cup \{x\} \subset \text{Reach}(y, S) \cup \{y\}$.

$$\begin{aligned}\text{Thus, } |\text{Reach}(x,S)| &= |\text{Reach}(x,S) \cup \{x\}| - 1 \\ &\leq |\text{Reach}(y,S) \cup \{y\}| - 1 \\ &= |\text{Reach}(y,S)|.\end{aligned}$$

□

It is instructive to illustrate theorem 3.4 with an example. In the example of figure 1, we have noted that

and
$$\begin{aligned}\text{Nbrhd}(y,S) &= \{s_1, s_2, s_4\} \\ \text{Reach}(y,S) &= \{a, b, c\}.\end{aligned}$$

The nodes a, b, d satisfy the condition in theorem 3.4 since

$$\begin{aligned}\text{Adj}(a) &= \{s_4\} \\ \text{Adj}(b) &= \{y, c\} \\ \text{Adj}(d) &= \{c\}.\end{aligned}$$

However, the node c fails.

The definitions introduced in this subsection can be extended to nonempty subsets. For any nonempty subset $Y \subset X \setminus S$, the readers are left to formalize the notions $\text{Nbrhd}(Y,S)$ and $\text{Closure}(Y,S)$.

§4 Refinements of the Algorithm using Reachable Sets

In this section, we consider novel features in our implementation of the minimum degree algorithm as described in section 3.2. It should be emphasized that we only operate on the adjacency structure of the original graph. This has the obvious advantage of keeping the original graph structure intact. In addition, no complicated data structure is necessary to allow for graph transformations.

§4.1 Minimum Degree Selection

For each execution of step 2 of the algorithm described in section 3.2, a node of minimum degree is selected and numbered. In what follows, we show that it may be possible to number a set of nodes at one time and that the amount of extra work involved is small. This basic observation has already been made and utilized elsewhere [3,9]. However, its previous application used information which is not available to us directly, since we do not have an explicit representation of the partially eliminated matrix. Our objective here is to establish conditions whereby we can carry out a "mass elimination", in terms of information provided by the REACH operator.

$$\text{Let } y \in X \setminus S \text{ satisfying } |\text{Reach}(y, S)| = \min_{z \notin S} |\text{Reach}(z, S)|.$$

Theorem 4.1 Let $x \in X \setminus S$ and

$$(4.1) \quad \text{Adj}(x) \subset \text{Closure}(y, S).$$

$$\text{Then} \quad |\text{Reach}(x, S)| = |\text{Reach}(y, S)|.$$

Proof From the choice of the node y , we have

$$|\text{Reach}(y, S)| \leq |\text{Reach}(x, S)|.$$

The result then follows from corollary 3.5. □

Among those nodes satisfying the adjacency condition in theorem 4.1, we restrict ourselves only to those in $\text{Reach}(y, S) \cup \{y\}$. To this end, we define the set

$$(4.2) \quad Y = \{x \in \{\text{Reach}(y, S) \cup \{y\} \mid \text{Adj}(x) \subset \text{Closure}(y, S)\}\}.$$

We now establish some properties of the set Y .

Lemma 4.2 Let $x \in Y$. Then

$$\text{Reach}(x, S) \cup \{x\} = \text{Reach}(y, S) \cup \{y\}.$$

Proof Assume $x \neq y$. Then $x \in Y \setminus \{y\} \subset \text{Reach}(y, S)$. Together with theorem 3.4, we have

$$\text{Reach}(x, S) \cup \{x\} \subset \text{Reach}(y, S) \cup \{y\}.$$

By theorem 4.1, the two sets must be the same. □

We can generalize the result of lemma 4.2 to an arbitrary subset of Y . Let Y' be a non-empty subset of Y .

Lemma 4.3 $\text{Reach}(Y', S) \cup Y' = \text{Reach}(y, S) \cup \{y\}$.

Proof By definition, $Y' \subset Y \subset \text{Reach}(y, S) \cup \{y\}$. Note also that $\text{Reach}(Y', S) = (\bigcup_{x \in Y'} \text{Reach}(x, S)) \setminus Y'$, so that by theorem 3.4 we have

$$\text{Reach}(Y', S) \subset \bigcup_{x \in Y'} \text{Reach}(x, S) \subset \text{Reach}(y, S) \cup \{y\}.$$

On the other hand, clearly $y \in \text{Reach}(Y', S) \cup Y'$. Let $v \in \text{Reach}(y, S)$ and $v \notin Y'$. Pick any $x \in Y'$; by lemma 4.2, $v \in \text{Reach}(x, S)$ so that v is reachable from some node in Y' through S . □

The next lemma follows directly from definition and from lemma 4.3

Lemma 4.4 $\text{Nbrhd}(Y, S) = \text{Nbrhd}(y, S),$

$$\text{Reach}(Y, S) \cup Y = \text{Reach}(y, S) \cup \{y\},$$

and $\text{Closure}(Y, S) = \text{Closure}(y, S).$ □

Lemma 4.5 For any subset $Y' \subset Y$, if $z \notin S \cup \text{Reach}(Y, S) \cup Y$ then $\text{Reach}(z, S \cup Y') = \text{Reach}(z, S)$.

Proof Since $z \notin \text{Reach}(Y, S)$, we have $Y \cap \text{Reach}(z, S) = \emptyset$.

It then follows that $\text{Reach}(z, S) = \text{Reach}(z, S \cup Y')$. □

Lemma 4.6 For any subset $Y' \subset Y$, if $x \in Y \setminus Y'$, then $\text{Reach}(x, S \cup Y') \cup Y' = \text{Reach}(x, S)$.

Proof By lemma 4.2,

$$Y' \subset Y \subset \text{Reach}(y, S) \cup \{y\} = \text{Reach}(x, S) \cup \{x\}.$$

But $x \notin Y'$, so that $Y' \subset \text{Reach}(x, S)$. Furthermore if $v \in \text{Reach}(x, S \cup Y')$, by the definition of Y , the node v must be reachable from x through S . Thus

$$\text{Reach}(x, S \cup Y') \cup Y' \subset \text{Reach}(x, S).$$

The other inclusion is immediate. □

Corollary 4.7 If $x \in Y \setminus Y'$, then $|\text{Reach}(x, S \cup Y')| = |\text{Reach}(x, S)| - |Y'|$. □

Theorem 4.8 For any subset $Y' \subset Y$, if $x \in Y \setminus Y'$, $|\text{Reach}(x, S \cup Y')| \leq |\text{Reach}(z, S \cup Y')|$ for all $z \notin S \cup Y'$.

Proof Let $x \in Y \setminus Y'$ and $z \notin S \cup Y'$. By theorem 4.1,

$$|\text{Reach}(x, S)| \leq |\text{Reach}(z, S)|.$$

Then, by lemma 4.4 and corollary 4.6, $|\text{Reach}(x, S \cup Y')| \leq |\text{Reach}(z, S \cup Y')|$. □

Theorem 4.8 has the following important implication: if y is the node selected in the minimum degree algorithm, the whole set Y can be numbered together.

§4.2 Degree Update

Having selected and numbered the subset Y as defined in the previous section, we need to update the sizes of the new reachable sets in preparation for the next node selection step. For our discussion, we let S and Y be as before and define $\bar{S} = S \cup Y$. The problem is to determine $|\text{Reach}(u, \bar{S})|$ for all $u \in X \setminus S$.

Theorem 4.9 Let $u \in X \setminus \bar{S}$. Then

$$\text{Reach}(u, \bar{S}) = \begin{cases} (\text{Reach}(u, S) \cup \text{Reach}(Y, S)) \setminus (Y \cup \{u\}), & \text{if } u \in \text{Reach}(Y, S) \\ \text{Reach}(u, S) & \text{otherwise.} \end{cases}$$

Proof If $u \notin \text{Reach}(Y, S)$, it follows from lemma 4.4 that $\text{Reach}(u, \bar{S}) = \text{Reach}(u, S)$. On the other hand, when $u \in \text{Reach}(Y, S)$, the result follows from the definition of $\text{Reach}(u, \bar{S})$ and $\text{Reach}(Y, S)$. \square

Theorem 4.9 says that only the nodes in $\text{Reach}(Y, S)$ need to have their degrees updated. In addition, for $u \in \text{Reach}(Y, S)$, its new degree is given by

$$|\text{Reach}(u, \bar{S})| = |(\text{Reach}(u, S) \cup \text{Reach}(Y, S)) \setminus (Y \cup \{u\})|.$$

Since the expression in theorem 4.7 can be written as the disjoint set union

$$[\text{Reach}(u, S) \setminus (Y \cup \text{Reach}(Y, S))] \cup [\text{Reach}(Y, S) \setminus \{u\}],$$

we have

$$|\text{Reach}(u, \bar{S})| = |\text{Reach}(u, S) \setminus (Y \cup \text{Reach}(Y, S))| + |\text{Reach}(Y, S) \setminus \{u\}|.$$

But the size of $\text{Reach}(Y, S)$ is known, so the problem is reduced to the determination of

$$(4.3) \quad |\text{Reach}(u, S) \setminus (Y \cup \text{Reach}(Y, S))|$$

for every $u \in \text{Reach}(Y, S)$.

As we do not intend to keep the set $\text{Reach}(u, S)$ but rather the number $|\text{Reach}(u, S)|$, the set given by (4.3) has to be generated in order to find its size. In what follows, we investigate the possibility of updating the degrees of a set of nodes in $\text{Reach}(Y, S)$. Let $u \in \text{Reach}(Y, S)$. The following lemma is obvious.

Lemma 4.10 $\text{Nbrhd}(Y, S) \cup Y \in \mathcal{C}(\text{Nbrhd}(u, \bar{S}))$. □

Lemma 4.11 If $\text{Adj}(u) \subset \text{Nbrhd}(u, \bar{S}) \cup \text{Adj}(\text{Nbrhd}(u, \bar{S}))$, then

$$|\mathcal{C}(\text{Nbrhd}(u, \bar{S}))| > 1.$$

Proof Assume for contradiction that $|\mathcal{C}(\text{Nbrhd}(u, \bar{S}))| = 1$. By Lemma 4.10, $\text{Nbrhd}(Y, S) \cup Y = \text{Nbrhd}(u, \bar{S})$. But $\text{Adj}(u) \subset \text{Nbrhd}(u, \bar{S}) \cup \text{Adj}(\text{Nbrhd}(u, \bar{S}))$

$$\begin{aligned} &= \text{Nbrhd}(Y, S) \cup Y \cup \text{Adj}(\text{Nbrhd}(Y, S) \cup Y) \\ &= \text{Nbrhd}(Y, S) \cup Y \cup \text{Reach}(Y, S) \\ &= \text{Closure}(Y, S) \\ &= \text{Closure}(y, S). \end{aligned}$$

By the definition (4.2) of the set Y , the node u should belong to the set Y . □

Lemma 4.12 If $\text{Adj}(u) \subset \text{Nbrhd}(u, \bar{S}) \cup \text{Adj}(\text{Nbrhd}(u, \bar{S}))$, then

$$\text{Adj}(\text{Nbrhd}(u, \bar{S})) = \text{Adj}(\text{Nbrhd}(u, \bar{S}) \cup \{u\}) \cup \{u\}. \quad \square$$

Now, let $u \in \text{Reach}(Y, S)$ satisfy the conditions

$$\begin{aligned} (4.4) \quad &\text{Adj}(u) \subset \text{Nbrhd}(u, \bar{S}) \cup \text{Adj}(\text{Nbrhd}(u, \bar{S})) \\ &\text{and } |\mathcal{C}(\text{Nbrhd}(u, \bar{S}))| = 2. \end{aligned}$$

Then the following theorem holds.

Theorem 4.13 If $v \in \text{Reach}(Y, S)$ with

$$\begin{aligned} &\text{Adj}(v) \subset \text{Closure}(u, \bar{S}) \\ &\text{and } \text{Adj}(v) \cap (\text{Nbrhd}(u, \bar{S}) \setminus (\text{Nbrhd}(Y, S) \cup Y)) \neq \emptyset, \\ &\text{then } |\text{Reach}(v, \bar{S})| = |\text{Reach}(u, \bar{S})|. \end{aligned}$$

Proof By lemma 3.3,

$$\text{Nbrhd}(v, \bar{S}) \subset \text{Nbrhd}(u, \bar{S}).$$

Together with $|\mathcal{C}(\text{Nbrhd}(u, \bar{S}))| = 2$ and $\text{Adj}(v) \cap (\text{Nbrhd}(u, \bar{S}) \cup \text{Nbrhd}(Y, S) \cup Y) \neq \emptyset$, we have $\text{Nbrhd}(v, \bar{S}) = \text{Nbrhd}(u, \bar{S})$. Therefore, by lemma 4.12,

$$\begin{aligned} \text{Reach}(v, \bar{S}) \cup \{v\} &= \text{Adj}(\text{Nbrhd}(v, \bar{S}) \cup \{v\}) \cup \{v\} \\ &\supset \text{Adj}(\text{Nbrhd}(v, \bar{S})) \cup \{v\} \\ &= \text{Adj}(\text{Nbrhd}(u, \bar{S})) \cup \{v\} \\ &= \text{Reach}(u, \bar{S}) \cup \{u\}. \end{aligned}$$

By corollary 3.5, the two sets must be the same. □

Thus, when we update the degrees of the nodes in $\text{Reach}(Y, S)$, it may be possible to update the degrees of a set of nodes $u \in \text{Reach}(Y, S)$, with only one application of the REACH operator.

§4.3 Determination of Reachable Sets

The reach set operator

$$\text{Reach}: X \setminus S \rightarrow \mathcal{P}(X \setminus S)$$

is used repeatedly in our implementation of the minimum degree algorithm.

In the minimum degree selection step, it is used to determine the set Y of nodes to be eliminated as given by (4.2). In the degree update step, the new degrees are obtained by finding the size of the set (4.3) through reachable sets. Thus, an efficient method for determining reachable sets is extremely important in the overall performance of the algorithm.

Let $y \in X \setminus S$. To find $\text{Reach}(y, S)$ as $\text{Adj}(\text{Nbrhd}(y, S) \cup \{y\})$, it is apparent that all the neighbors of the nodes in $\text{Nbrhd}(y, S) \cup \{y\}$ have to be inspected for membership. It is natural to ask whether it is possible to generate $\text{Reach}(y, S)$ with less effort.

In this connection, we look for a subset V of S so that for $x \in X \setminus S$, the reachable set $\text{Reach}(x, S)$ is the same as $\text{Reach}(x, S \setminus V)$, where

the latter operator is taken in the subgraph $G(X \setminus V) = (X \setminus V, E(X \setminus V))$. It is clear that this subset would satisfy the following conditions:

$$\text{Nbrhd}(x, S \setminus V) \subset \text{Nbrhd}(x, S)$$

and $\text{Reach}(x, S \setminus V) \setminus V = \text{Reach}(x, S)$.

By marking the subset V as inactive, we need only to consider the subgraph $G(X \setminus V)$, and this can imply a drastic reduction in effort to generate $\text{Reach}(x, S)$.

Consider the example in figure 4.1. The subset S contains 23 nodes marked with s . By making 7 of them inactive (those nodes marked in dark), the reachable sets are preserved in the subgraph. Furthermore, to find $\text{Reach}(x, S)$, 16 nodes have to be inspected; whereas, in the subgraph, only 10 is necessary.

Consider an elimination step in the algorithm. Let S be the set of nodes already eliminated, and Y be the set (4.2) of nodes with minimum degree to be eliminated. If we use \bar{S} to denote the new set of eliminated nodes $S \cup Y$, then

$$\text{Nbrhd}(Y, S) \cup Y \in C(\bar{S}),$$

and $\text{Reach}(Y, S) = \text{Adj}(\text{Nbrhd}(Y, S) \cup Y)$.

All nodes in $\text{Nbrhd}(Y, S) \cup Y$ except a subset R can be marked inactive so long as the subset R defines a connected subgraph $G(R)$ and satisfies

$$(4.4) \quad \text{Reach}(Y, S) \subset \text{Adj}(R).$$

The best choice of R would be one with the smallest size satisfying the above conditions. For practical reasons, we look for one that can be **determined easily**.




Figure 4.1 An example to illustrate the effect
of de-activating eliminated nodes.

Obviously, the subset

$$(4.5) \quad \text{Adj}(\text{Reach}(Y,S)) \cap (\text{Nbrhd}(Y,S) \cup S)$$

satisfies (4.4), but it may not be connected. We simply take R to be a connected subgraph of $\text{Nbrhd}(Y,S) \cup Y$ containing the set (4.5).

Thus, the way to "de-activate" eliminated nodes is completely specified if we can provide an algorithm for the following problem. We shall describe it in a general setting. Let $G = (X,E)$ be a connected graph and Z be a (possibly disconnected) subset of X . The problem is to determine a small subset \bar{Z} such that

$$Z \subset \bar{Z} \subset X$$

and such that the subgraph $G(\bar{Z})$ is connected. The following algorithm serves the purpose though it may not produce a smallest \bar{Z} .

- Step 1 (Initialization) Initialize $\bar{Z} \leftarrow Z$ and find a component $C \in \mathcal{C}(Z)$.
- Step 2 (Test for termination) If $C = \bar{Z}$, stop.
- Step 3 (Reach for another component) Determine a node $z \in Z \setminus C$, which is closest to some node c in C . Let c, x_1, \dots, x_t, z be a shortest path with $x_i \in X \setminus Z$, $1 \leq i \leq t$.
- Step 4 (Expansion of component) Find the component $C' \in \mathcal{C}(Z)$ with $z \in C'$. Put
- $$\bar{Z} \leftarrow \bar{Z} \cup \{x_1, \dots, x_t\}$$
- and
- $$C \leftarrow C \cup \{x_1, \dots, x_t\} \cup C'.$$
- Step 5 (Loop) Go to step 2.

§4.4 An Overview of the Refined Algorithm

Refinements have been given to improve the overall performance of the minimum degree algorithm as described in section 3.2. In this section, we combine them and provide an overall picture of the entire algorithm.

- Step 1 (Initialization) Let $G = (X, E)$ and initialize $S \leftarrow \phi$,
 $DEG(x) \leftarrow |Adj(x)|$, for $x \in X$.
- Step 2 (Minimum degree selection) Pick a node $y \in X \setminus S$ where
 $DEG(y) = \min_{x \in X \setminus S} DEG(x)$.
- Step 3 (Mass minimum degree elimination) Determine the sets $Reach(y, S)$ and $Closure(y, S)$. Then find
 $Y = \{x \in Reach(y, S) \cup \{y\} \mid Adj(x) \subset Closure(y, S)\}$, and number the nodes in Y next.
- Step 4 (Mass degree update) For $u \in Reach(Y, S)$, determine
 $Reach(u, S) \setminus (Reach(Y, S) \cup Y)$ and put
 $DEG(u) \leftarrow |Reach(u, S) \setminus (Reach(Y, S) \cup Y)| + |Reach(Y, S) \setminus \{u\}|$.
 Check for mass degree update conditions and perform the update accordingly, if possible.
- Step 5 (Deactivate some eliminated nodes) Find a subset V of
 $Nbrhd(Y, S) \cup Y$ such that
 $G((Nbrhd(Y, S) \cup Y) \setminus V)$ is connected
 and $Reach(Y, S) \subset Adj((Nbrhd(Y, S) \cup Y) \setminus V)$.
 Then put
 $(X, E) \leftarrow (X \setminus V, E(X \setminus V))$
- and $S \leftarrow (S \cup Y) \setminus V$.
- Step 6 (Loop or stop) If $S = X$, stop. Otherwise, go to step 2.


In the next sections, we shall consider some implementation details and provide some experimental results on the performance of this scheme.

§5 Implementation Details and Numerical Experiments

In this section we describe the important aspects of our implementations (MD), and report some numerical experiments showing its performance. As a basis for comparison, we have used the appropriate routines from the Yale Sparse Matrix Package (YSMP) [9], which is fairly widely distributed and enjoys a good reputation.

In order to avoid doing any searching for a node of minimum degree, the degrees of the nodes are stored using the three arrays ORG, NEXT, and PREV, each of length N , as depicted in Figure 5.1. Nodes having the same degree are stored in a doubly linked degree list; the beginning of the list containing nodes of degree i is stored in $ORG(i)$, and the arrays NEXT and PREV contain the usual forward and backward pointers of a doubly linked list. If node q is the first in the degree list, then $PREV(q)$ is $-\deg(q)$, and if node q is the last node in the degree list, then $NEXT(q) = 0$. Note that when the degree of a node is changed, it can be deleted from the old degree list and inserted in the new list in a fixed number of operations, independent of N . By maintaining a pointer to the first nonzero entry in ORG, we can easily find a node of minimum degree [6].

In the set of subroutines which implement our algorithm, the set S is maintained using a three-state array SMASK, where $SMASK(i) > 0 \Leftrightarrow$ node $i \in S$, $SMASK(i) = 0 \Leftrightarrow$ node $i \notin S$, and $SMASK(i) < 0$ if node i has been deactivated using the algorithm of section 4.3. In addition, four other arrays of length N are used, one of which is, like SMASK, only required to store these different states of a node. The graph itself is stored as a sequence of adjacency lists in the array pair (XADJ, ADJNCY), as shown in Figure 5.1. Since our implementation does not exploit the fact that two of



Graph G

1	0	4	-2
2	1	11	6
3	6	9	-4
4	3	0	1
5	8	0	8
6	7	2	-3
7	0	0	-6
8	0	5	-5
9	0	0	3
10	0	0	11
11	0	10	2

ORG

NEXT

PREV




Figure 5.1 Representation of the adjacency and degree structure of the graph using the arrays XADJ, ADJNCY, ORG, NEXT, and PREV.

the vectors of length need only represent three states, our total array storage requirements for the ordering program is $9N + 2|E|$ integer locations.

In order to obtain some evidence bearing on the asymptotic execution time of our implementation, the program was applied to the graded L mesh (graph) shown in Figure 5.2, subdivided by increasing subdivision factors s , yielding s^2 as many triangles as in the original mesh. The YSMP subroutine SORDER was also run on the same problems to obtain comparative results. The results are summarized in Table 5.1. Execution times are in seconds on an IBM 360/75 computer. The programs are written in Fortran, compiled using the optimizing version of the H-level compiler.

The entries in the storage column for the YSMP program were obtained by inserting a statement in SORDER to monitor the maximum storage used in some working storage arrays. The user must estimate this number, and provide at least this much storage to allow the program to execute. This is a disadvantage shared by many minimum degree algorithm implementations. We regard the fact that our implementation uses a modest **fixed** amount of storage as a major advantage. The quality of the orderings produced by the two programs was almost identical; i.e., the amount of fill suffered by the corresponding ordered matrices when factored was nearly identical.

Note that the asymptotic execution time of MD appears to be of a lower order than the YSMP program. Indeed, the numerical results in Table 5.1 suggest quite strongly that the execution time of our program is $O(N)$. Other similar experiments with finite element discretizations of two and three dimensional regions suggest that our program runs in $O(N)$ time for this class of problems.




Figure 5.2 Graded L mesh

EXECUTION TIME						
MD			YSMP		STORAGE	
N	TIME	$\frac{\text{TIME}}{N}$	TIME	$\frac{\text{TIME}}{N}$	MD	YSMP
265	1.20	4.52	.90	3.21	4139	6667
406	1.97	4.85	1.50	3.65	6371	10428
577	2.95	5.11	2.31	3.85	9083	15559
778	3.89	5.00	3.39	4.27	12275	23044
1009	5.14	5.10	4.60	4.45	15947	30211
1270	6.30	4.96	6.28	4.95	20099	40220
1561	7.98	5.11	8.14	5.21	24731	49407
1882	9.87	5.24	10.65	5.66	29843	61488
2233	11.98	5.37	12.68	5.68	35435	76427
2614	13.63	5.21	-	-	41507	-
3025	15.99	5.29	-	-	48059	-
3466	18.50	5.34	-	-	55091	-
		($\times 10^{-3}$)		($\times 10^{-3}$)		

Table 5.1 Comparisons of execution times and storage requirements for the MD and YSMP programs.

In order to demonstrate the effectiveness of (a) the deactivation technique described in section 4.3, and (b) the "mass degree update" technique (Theorem 4.13), we solved the problems reported on above, with one or both of these features removed from our program. The results are summarized in Table 5.2, and illustrate their value for large N.

N	NODE DEACTIVATION AND MASS DEGREE UPDATE	NO NODE DEACTIVATION BUT MASS DEGREE UPDATE	NODE DEACTIVATION BUT <u>NO</u> MASS DEGREE UPDATE	NO NODE DEACTIVATION AND NO MASS DEGREE UPDATE
265	1.20	1.30	1.49	1.73
406	1.97	2.16	2.81	3.61
577	2.95	3.28	4.15	5.82
778	3.89	4.30	5.90	8.98
1009	5.14	5.82	7.87	13.57
1270	6.30	7.07	10.65	18.24
1561	7.98	9.04	13.92	24.98
1882	9.87	11.78	17.66	32.39
2233	11.98	14.49	21.74	41.81

Table 5.2 Results showing the effect of node deactivation and mass degree update.

§6 Conclusions

We have described an implementation of the minimum degree algorithm which requires only $O(N)$ storage in addition to that required for the original graph. The storage required by the program is known before execution, and is independent of the amount of fill that is suffered by the correspondingly reordered matrix, when it is factored. The results of Table 5.2 show that the success of our approach depends upon two important techniques; "mass degree update", and "node deactivation". The latter of these techniques has important application in other areas of sparse matrix computation, particularly in the area of storage allocation and related problems [4]. Finally, the results in Table 5.1, along with other experiments, suggest that the execution time of our implementation, for problems arising in finite element applications is $O(N)$. Unfortunately, it is not difficult to construct mesh-like problems where the running time of our algorithm is superlinear.

§7 References

- [1] I.S. Duff, A.M. Erisman and J.K. Reid, "On George's nested dissection algorithm", SIAM J. Numer. Anal., 5 (1976), pp. 686-695.
- [2] Alan George and Joseph W.H. Liu, "An algorithm for automatic nested dissection and its application to general finite element problems", Proc. Sixth Conference on Numerical Mathematics and Computing, Winnipeg, Manitoba, Sept. 30 - Oct. 2, (1976).
- [3] Alan George and Dávid R. McIntyre, "On the application of the minimum degree algorithm to finite element systems", SIAM J. Numer. Anal., to appear.
- [4] Alan George and Joseph W.H. Liu, "On finding diagonal block envelopes of triangular factors of partitioned matrices", Dept. of Computer Science Technical Report, CS-77-10, 1977.
- [5] S.V. Parter, "The use of linear graphs in Gauss elimination". SIAM Rev., 3 (1961), pp. 364-369.
- [6] J.K. Reid, Private communication.
- [7] D.J. Rose, "A graph theoretic study of the numerical solution of sparse positive definite systems", in Graph Theory and Computing, R.C. Read, editor, Academic Press, (1972).
- [8] D. Rose, R. Tarjan and G. Luecker, "Algorithmic aspects of vertex elimination on directed graphs", SIAM J. Computing, 5 (1975), pp. 266-283.
- [9] Andrew H. Sherman, "Yale Sparse Matrix Package Users Guide", Lawrence Livermore Laboratory Report, UCID-30114, August 1975.