

ornl

OAK RIDGE
NATIONAL
LABORATORY

MARTIN MARIETTA

MARTIN MARIETTA ENERGY SYSTEMS LIBRARIES

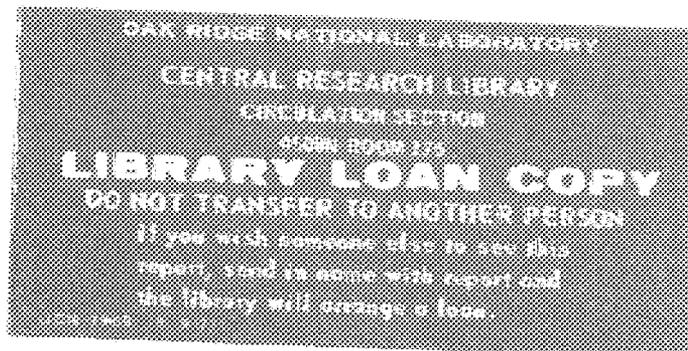


3 4456 0262425 7

ORNL/TM-10452

The Evolution of the Minimum Degree Ordering Algorithm

Alan George
Joseph W. H. Liu



OPERATED BY
MARTIN MARIETTA ENERGY SYSTEMS, INC.
FOR THE UNITED STATES
DEPARTMENT OF ENERGY

Printed in the United States of America. Available from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Road, Springfield, Virginia 22161
NTIS price codes—Printed Copy: A03; Microfiche A01

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Engineering Physics and Mathematics Division
Mathematical Sciences Section

**THE EVOLUTION OF THE
MINIMUM DEGREE ORDERING ALGORITHM**

Alan George †

Joseph W. H. Liu ††

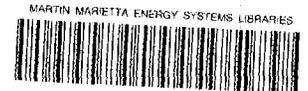
† Also a member of the
Departments of Computer Science and Mathematics
The University of Tennessee
Knoxville, Tennessee 37996-1301

†† Department of Computer Science
York University
North York, Ontario
Canada M3J 1P3

Date Published: May 1987

Research was supported by in part by the Natural Sciences and Engineering Research Council of Canada under grants A8111 and A5509, by the Applied Mathematical Sciences Research Program of the Office of Energy Research, U.S. Department of Energy under contract DE-AC05-84OR21400 with Martin Marietta Energy Systems Inc., by the U.S. Air Force Office of Scientific Research under contract AFOSR-ISSA-85-00013, and by the Science Alliance, a state supported program at the University of Tennessee.

Prepared by the
Oak Ridge National Laboratory
Oak Ridge, Tennessee 37831
operated by
Martin Marietta Energy Systems, Inc.
for the
U.S. DEPARTMENT OF ENERGY
under Contract No. DE-AC05-84OR21400



3 4456 0262425 7

CONTENTS

Abstract	v
1. Introduction	1
2. The Basic Minimum Degree Ordering	2
3. Development of Improvements to the Minimum Degree Algorithm	4
4. Timing Results of Various Improvement Techniques	10
5. On Tie-Breaking Strategies	12
6. Odds and Ends	14
7. Future Research Directions	16
References	17

The Evolution of the Minimum Degree Ordering Algorithm

Alan George †

Mathematical Sciences Section
Engineering Physics and Mathematics Division

and

Department of Computer Science
The University of Tennessee
Knoxville, Tennessee 37996-1301

Joseph W.H. Liu

Department of Computer Science
York University
North York, Ontario
Canada M3J 1P3

ABSTRACT

Over the past fifteen years, the implementation of the minimum degree algorithm has received much study, and many important enhancements have been made to it. In this article, we describe these various enhancements, trace their historical development, and provide some experiments showing how very effective they are in improving the execution time of the algorithm. We also present a shortcoming that exists in all of the widely used implementations of the algorithm, namely, that the quality of the ordering provided by the implementations is surprisingly sensitive to the initial ordering. For example, changing the input ordering can lead to an increase (or decrease) of as much as a factor of three in the cost of the subsequent numerical factorization. This sensitivity is caused by the lack of an effective tie-breaking strategy, and our experiments illustrate the importance of developing such a strategy.

1. Introduction

Consider the n by n symmetric positive definite system of equations

$$A x = b ,$$

where n is large and A is sparse. When A is factored using Cholesky's method, it normally suffers some fill. Since PAP^T is also symmetric and positive definite for any permutation matrix P , we can instead solve the reordered system

$$(P A P^T) (P x) = P b .$$

The choice of P can have a dramatic effect on the amount of fill that occurs during the factorization. Thus, it is standard practice to reorder the rows and columns of the matrix before performing the factorization.

The overall solution of a sparse positive definite system of equations is typically divided into four distinct independent phases:

- (a) Find an appropriate ordering P for A .
- (b) Set up a data structure for L , the Cholesky factor of PAP^T .
- (c) Numerically factor PAP^T into LL^T .
- (d) Solve $(LL^T)(Px) = Pb$.

Note that steps (a) and (b) depend only on the structure of A , and are independent of its numerical values.

The problem of finding a best ordering for A in the sense of minimizing the fill is computationally intractable: an NP-complete problem [5]. We are therefore obliged to rely on heuristic algorithms. One of the most effective of these is the minimum degree algorithm. This algorithm is a symmetric analog of an algorithm proposed by Markowitz in 1957 for reordering equations arising in linear programming applications [2]. Loosely speaking, the Markowitz algorithm begins with the given matrix, and at each step of Gaussian elimination, row and column permutations are performed so as to minimize the product of the number of off-diagonal nonzeros in the pivot row and pivot column. Thus, one minimizes the amount of arithmetic that must be performed at each step of Gaussian elimination. This will also tend to minimize the amount of fill that occurs. Of course such a local minimization strategy will not in general provide a global minimum for either the arithmetic requirements or fill. Nevertheless, the strategy has proved to be very effective in reducing arithmetic and fill. The cost of employing the strategy is almost always far outweighed by the savings in execution time that accrue. In practice, in order to preserve numerical stability, care must be taken to avoid using pivot elements that are too small in magnitude. Thus, there is a trade-off between limiting fill and preserving numerical accuracy. A widely used modern implementation of Markowitz's basic strategy, along with many improvements, is Duff's MA28 code [1].

Tinney and Walker employed a similar strategy in solving the large sparse systems arising in the analysis of power systems [4]. These problems have symmetric structure, and do not require interchanges for numerical stability provided that the pivots are taken from the diagonal of A . Thus, Tinney and Walker employed a symmetric version of the Markowitz strategy. That is, whenever rows were interchanged, the corresponding columns were interchanged as well, thus preserving the symmetric structure of the matrix. Subsequently, Rose [3] developed a graph theoretic model of the algorithm, and for reasons that will be clear in later parts of this paper, he renamed Tinney and Walker's symmetric analog of Markowitz's original strategy as the minimum degree algorithm.

In terms of the matrix A , the minimum degree algorithm can be described as follows. Generally, it works only with the structure of A , and simulates in some manner the n steps of symmetric Gaussian elimination. At each step, a row and corresponding column interchange is applied to the part of the matrix remaining to be factored so that the number of nonzeros in the pivot row and column is minimized. (Note that since the structure of the matrix is symmetric, the number of nonzeros in the pivot row and pivot column is the same.) After n steps, the entire factorization has been simulated, and the order in which the pivot rows and columns were chosen is the ordering. An important observation which we will focus on later is that ties usually occur in the choice of the pivot row and column, which implies that the ordering obtained will depend on how these ties are resolved.

Rose's 1970 doctoral thesis work [24] did much to promote interest in the minimum degree algorithm, and its efficient implementation has received a great deal of attention over the past fifteen years. Researchers at Harwell, Waterloo and Yale in particular have devoted much effort in this direction [4, 5, 7, 9, 12, 13, 16, 21]. A number of effective enhancements have been proposed, and are now standard techniques in most state-of-the-art implementations. One of our objectives in this article is to trace these developments, and to provide some experimental results illustrating the extent to which the technology has advanced. A second objective of our work is to describe and highlight a significant shortcoming that exists in all of the widely used implementations of the algorithm. Specifically, the quality of the ordering provided by these implementations is quite sensitive to the initial ordering of A . This is a reflection of the fact that the quality of the ordering depends quite strongly on the way in which ties are resolved, and all of the implementations break ties more or less arbitrarily. Finding an effective tie-breaking strategy is one of the many interesting problems related to the algorithm that still remain unsolved.

An outline of the paper is as follows. In section 2 we describe the minimum degree algorithm in graph theory terms and provide a small example. In section 3 we trace the development of the various enhancements that have been proposed to reduce the execution time of the algorithm and section 4 contains some numerical experiments showing the relative effectiveness of the enhancements. In section 5 we provide some experiments which illustrate the importance of tie-breaking. Section 6 mentions briefly some implementation issues, and describes how the implementation might be adapted to special structures arising in certain applications, such as the finite element method and least squares problems. Section 8 contains some remarks on future research directions.

2. The Basic Minimum Degree Ordering

2.1. Description of the Minimum Degree Algorithm

We assume that the reader is familiar with basic graph theory notions, and the correspondence between undirected graphs and the structure of symmetric matrices. See, for example, Chapter 3 of [14]. Let G be an undirected graph and v a node of G . Throughout this paper, we shall use the notation $Adj_G(v)$ to refer to the set of nodes adjacent to v in G . The degree of the node v in G will be denoted by $degree_G(v)$, which is simply $|Adj_G(v)|$.

The basic minimum degree ordering can be best described in terms of elimination graphs. Following Rose [25], we use the notation G_v to represent the elimination graph obtained after the elimination of the node v from the graph G . The graph G_v can be obtained by deleting the node v and its incident edges from G and then adding edges to make the nodes adjacent to v into a clique.

The basic algorithm is as follows:

Algorithm 2.1: (Basic minimum degree ordering)
begin

$G :=$ given symmetric graph ;
 while $G \neq \emptyset$ *do*
 begin
 select a node y of minimum degree in G and order y next ;
 $G := G_y$;
 end ;

end .

The above simple formulation of the basic algorithm captures the main idea of the ordering strategy. Implicitly assumed in this description is a scheme to represent the elimination graphs and to transform them. This allows the selection of a minimum degree node in the new elimination graph. Based on the transformation rule, we note that if a node v is not adjacent to y in G ,

$$Adj_{G_y}(v) = Adj_G(v).$$

However, if $v \in Adj_G(y)$, then we have

$$Adj_{G_y}(v) = (Adj_G(y) \cup Adj_G(v)) - \{v, y\}.$$

Therefore, the degree of a node may change after the elimination graph transformation due to the deletion of edges incident to y and the possible addition of new edges joining nodes adjacent to y . In other words, $degree_G(v)$ may be quite different from $degree_{G_y}(v)$. The following are simple observations on the relation between $degree_{G_y}(v)$ and $degree_G(v)$.

Observation 2.1: If v is not adjacent to y , then $degree_{G_y}(v) = degree_G(v)$.

Observation 2.2: If v is adjacent to y , then

$$degree_{G_y}(v) \geq degree_G(v) - 1 \geq degree_G(y) - 1.$$

Algorithm 2.1 is not completely specified, since at the node selection stage, there may be several nodes of minimum degree. These "ties" must be resolved in some manner, usually arbitrarily. Figure 2.1 gives two minimum degree orderings by applying this algorithm to the 5-by-5 regular grid model problem (9-point difference), where the ties were resolved in different ways. Different orderings can be obtained due to the freedom in the selection of minimum degree nodes. The example in Figure 2.1 is designed to illustrate this point.

4	8	11	7	3	4	7	21	12	3
20	21	22	16	15	11	16	22	15	8
10	14	23	13	9	20	19	23	18	17
6	18	24	25	5	6	14	24	13	9
2	17	12	19	1	2	10	25	5	1

Figure 2.1: Two minimum degree orderings on a 5-by-5 grid.

Often, the quality of the resulting minimum degree ordering depends crucially on the "correct" node selection from the set of minimum degree nodes. We use the term tie-breaking strategy to refer to any scheme that provides a means of choosing a node from such a set. Some tie-breaking strategies will be discussed in later sections.

2.2. Quality of the Minimum Degree Ordering

It is easy to see that a minimum degree ordering on a tree structure will be one that is a perfect ordering (with no fill), and hence a minimum fill ordering. But for general graphs, a minimum degree ordering may not be a minimum fill ordering. In [25], Rose provides an example to show this. Indeed, his example is a graph associated with a perfect elimination matrix; by definition, such matrices can always be ordered so that they suffer no fill. However, a minimum degree ordering of this example will create some fill from elimination. Recently, Hempel [19] gives a graph example where a minimum degree can generate fill that is more than a constant factor greater than that given by a minimum fill ordering.

If we consider the k -by- k regular grid problem, a nested dissection ordering of it [10] is known to be optimal. It can also be regarded as a minimum degree ordering with a special way of selecting nodes of minimum degree. In other words, a special type of minimum degree ordering on the grid problem can be optimal. Unfortunately, not all minimum degree orderings for the grid are optimal. More will be said about this later in this paper when we consider tie-breaking strategies.

3. Development of Improvements to the Minimum Degree Algorithm

It is apparent from the description of Algorithm 3.1 that the transformation of the elimination graph

$$G := G_y$$

is central to the implementation of this algorithm. In general, this elimination step creates fill among nodes adjacent to y . It is therefore necessary to have an efficient representation of the resulting elimination graph to accommodate additional fill (either implicitly or explicitly). Moreover, the degrees of nodes adjacent to y may change, so that a re-calculation of their degrees is required in preparation for the next node selection step. It has been recognized by researchers that this "degree update" is the most time consuming step of the entire algorithm.

In this section, we discuss various methods that have been developed to improve the performance of the basic minimum degree algorithm. They are presented in the sequence in which they appeared in the literature, and in each instance, we try to attribute the technique to those who first suggested them.

3.1. Mass Elimination

In their study of the minimum degree algorithm on finite element problems, George and McIntyre [16] observe that in the elimination of a node y of minimum degree, often there is a subset of nodes adjacent to y that can be eliminated immediately after y . This happens especially often in later stages of the elimination. The next theorem contains the theoretical basis of their observation.

Theorem 3.1: If y is selected as the minimum degree node in the graph G , then the subset

$$Y = \{ z \in Adj_G(y) \mid degree_G(z) = degree_G(y) - 1 \}$$

can be selected next (in any order) in the minimum degree algorithm.

Proof: We first show that any node $z \in Y$ is a node of minimum degree in G_y . For any node $v \neq y$, it is clear that

$$\text{degree}_G(y) \leq \text{degree}_G(v).$$

Furthermore, by definition of the set Y and observation 2.2, we have

$$\begin{aligned} \text{degree}_{G_y}(z) &= \text{degree}_G(y) - 1 \\ &\leq \text{degree}_G(v) - 1 \leq \text{degree}_{G_y}(v). \end{aligned}$$

The theorem then follows by a repeated application of the same argument and the observation that for any $v \in Y - \{z\}$,

$$\text{degree}_{(G_y)_z}(v) = \text{degree}_{G_y}(z) - 1.$$

Theorem 3.1 allows us to avoid some graph transformations and degree update steps, since it provides a set of nodes of minimum degree that may be selected next. Instead of having to perform the transformation $G := G_y$ and the degree update, we can eliminate nodes in $Y \cup \{y\}$ simultaneously. This implies that the elimination graph transformation and degree update need only be performed once for the whole set instead of $|Y \cup \{y\}|$ times. This will also save node selection time. We describe how to identify Y in section 3.2.

To illustrate mass elimination, consider the actual application of the minimum degree algorithm to the 5-by-5 grid. Figure 3.1 shows an intermediate stage after 14 nodes have been eliminated during the ordering (compared with the left ordering in Figure 2.1). Nodes marked with "*" have already been eliminated. It is easy to verify that the node $y = 15$ has minimum degree of 8 and the node $z = 16$ satisfy the condition in Theorem 3.1. Therefore, $\{y, z\}$ can be eliminated together.

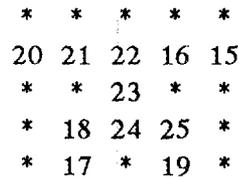


Figure 3.1: An intermediate stage of the minimum degree ordering.

3.2. Indistinguishable Nodes

The condition in Theorem 3.1 for mass elimination has a simple equivalent property. As before, let y be a node of minimum degree in G and $z \in \text{Adj}_G(y)$. By the nature of the elimination graph transformation,

$$\text{Adj}_{G_y}(z) = (\text{Adj}_G(y) \cup \text{Adj}_G(z)) - \{y, z\}.$$

Theorem 3.2: Let $z \in \text{Adj}_G(y)$. Then, $\text{degree}_{G_y}(z) = \text{degree}_G(y) - 1$ if and only if $\text{Adj}_G(z) \cup \{z\} = \text{Adj}_G(y) \cup \{y\}$.

Proof: "if" part: If $Adj_G(z) \cup \{z\} = Adj_G(y) \cup \{y\}$, then

$$\begin{aligned} degree_{G_y}(z) &= |(Adj_G(y) \cup Adj_G(z)) - \{y, z\}| \\ &= |Adj_G(y) - \{z\}| = degree_G(y) - 1. \end{aligned}$$

"only if" part: Let $degree_{G_y}(z) = degree_G(y) - 1$. We first show that $Adj_G(z) \cup \{z\} \subseteq Adj_G(y) \cup \{y\}$. Assume for contradiction that there exists a node $x \in Adj_G(z) - Adj_G(y)$. Then we have

$$\begin{aligned} Adj_{G_y}(z) &= (Adj_G(y) \cup Adj_G(z)) - \{y, z\} \\ &\supseteq (Adj_G(y) - \{z\}) \cup \{x\}. \end{aligned}$$

This implies that $degree_{G_y}(z) > degree_G(y) - 1$, and hence proves the claim that

$$Adj_G(z) \cup \{z\} \subseteq Adj_G(y) \cup \{y\}.$$

However, y is selected to be a node of minimum degree so that the two sets must be the same.

Together with Theorem 3.1, we note that the equivalent property in Theorem 3.2 can be used to identify nodes for mass elimination. We now show that it can also help in reducing the number of degree updates. Formally, we say that two nodes u and v are indistinguishable in G [14] if

$$Adj_G(u) \cup \{u\} = Adj_G(v) \cup \{v\}.$$

Sometimes the set $Adj_G(u) \cup \{u\}$ is referred to as the neighborhood set of u (see for example the book by Golubic [18]). So, two nodes with identical neighborhood sets are indistinguishable. As shown in [14], this property is preserved under elimination graph transformation.

Theorem 3.3: If two nodes are indistinguishable in G , they remain indistinguishable in G_y .

It is obvious that if two nodes are indistinguishable, their degrees must be the same. Theorem 3.3 says that if two nodes become indistinguishable at some stage of elimination, they will have identical neighborhood sets (and hence the same degree) thereafter. Furthermore, Theorems 3.1 and 3.2 imply that they can be eliminated together whenever one of them is picked for elimination.

Therefore, as far as the minimum degree algorithm is concerned, indistinguishable nodes can be merged together and treated as one. In this way, we need only to consider one representative from each group of indistinguishable nodes. In the literature, these representative nodes have been referred to as supernodes, supervariables [6], and prototype nodes [7].

The advantage of using indistinguishable nodes should be clear. We need to perform the degree update only on the representative nodes. This reduces the operating size of the current elimination graph in terms of both nodes and edges.

To illustrate this notion, consider the stage as given in Figure 3.1. The remaining 11 nodes can be divided into 5 groups of indistinguishable nodes, namely:

- {15, 16}
- {17, 18}
- {19, 25}
- {20, 21}
- {22, 23, 24}.

In effect, we are now working on a graph with 5 (super) nodes rather than with 11 nodes.

3.3. Representation of Elimination Graphs

As noted earlier, a compact and efficient representation of the sequence of elimination graphs is crucial to the overall performance of the minimum degree ordering algorithm. The elimination graph transformation " $G := G_y$ " involves both deletion and addition of edges. In other words, the number of edges in the elimination graph may either increase or decrease depending on how the nodes adjacent to y are connected. The representational scheme must be able to accommodate such changes.

In 1972, Speelpenning [27] uses a generalized element approach to consider the elimination process on mesh problems arising from the finite element method. Using this view, we can interpret each elimination graph as a collection of cliques or generalized elements. Indeed, the original graph with $|E|$ edges can be regarded as one with $|E|$ cliques, each consisting of two nodes (or equivalently an edge).

This not only gives a conceptually different view of the elimination process, but also provides a compact scheme to represent the elimination graphs. The advantage of this representation in terms of storage is based on the following observations.

Let $\{K_1, K_2, \dots, K_q\}$ be the set of cliques for the current graph G . Assume that y is a node of minimum degree selected for elimination. Let $\{K_{s_1}, \dots, K_{s_t}\}$ be the subset of cliques to which y belongs. Then the elimination graph transformation amounts to:

- 1) the removal of the cliques K_{s_1}, \dots, K_{s_t} from $\{K_1, K_2, \dots, K_q\}$;
- 2) the addition of the new clique:

$$K = (K_{s_1} \cup \dots \cup K_{s_t}) - \{y\}$$

into the clique set.

The next theorem is obvious and it is the basis of the advantage of this generalized element representation scheme.

Theorem 3.4: $|K| < \sum_{i=1}^t |K_{s_i}|.$

This theorem says that in using the generalized element approach for the representation of the sequence of elimination graphs, the amount of storage required will never exceed the amount needed to represent the original graph. The storage requirement to carry out the minimum degree algorithm is therefore known beforehand in spite of the rather dynamic nature of the elimination process. An economical representation of the structures has an added advantage of reducing processor time in their manipulations.

The use of the generalized element storage scheme appears in the Harwell code MA27 [5], the Waterloo Sparse Matrix Package (SPARSPAK) [15], and the Yale Sparse Matrix Package (YSMP) [7]. This representation technique has been referred to as a generalized element model, superelements [6, 8], and the quotient graph model [12].

Apparently, the developers of YSMP were the first to adopt this approach. Both the Harwell code and YSMP use linked lists to implement the set of generalized elements. In the SPARSPAK implementation, linked lists of subarrays are used to take advantage of the initial data representation of the structure of the given matrix.

We again use the example in Figure 3.1 to illustrate this generalized element approach. At this point, the elimination graph has four cliques:

$$\begin{aligned} &\{15, 16, 19, 22, 23, 24, 25\} \\ &\{15, 16, 20, 21, 22\} \\ &\{17, 18, 20, 21, 22, 23, 24\} \\ &\{17, 18, 19, 24, 25\} \end{aligned}$$

Upon the elimination of the nodes $\{15, 16\}$ (by mass elimination), it is easy to see that there will be three cliques left and they are:

$$\begin{aligned} &\{19, 20, 21, 22, 23, 24, 25\} \\ &\{17, 18, 20, 21, 22, 23, 24\} \\ &\{17, 18, 19, 24, 25\} \end{aligned}$$

Notice that the amount of storage required is less than before.

3.4. Incomplete Degree Update

The notion of indistinguishable nodes in section 3.2 can be regarded as an extension of the condition used for mass elimination. The use of indistinguishable nodes helps not only at elimination but also at the time of degree update. By merging indistinguishable nodes together, we need only to recompute the degrees of the representatives. Here, we consider yet a further generalization, which was first used in the YSMP implementation [7]. This technique speeds up the minimum degree algorithm by avoiding the degree computation for nodes that are known not to be minimum degree.

Following [7], given two nodes u and v in the graph G , the node v is said to be outmatched by u if

$$Adj_G(u) \cup \{u\} \subseteq Adj_G(v) \cup \{v\}.$$

A direct consequence of this condition is that $degree_G(u) \leq degree_G(v)$. More importantly, this property is preserved under elimination graph transformation [21].

Theorem 3.5: If the node v is outmatched by u in G , it is also outmatched by u in the graph G_y .

Corollary 3.6: If a node v becomes outmatched by u in the elimination process, the node u can be eliminated before v in the minimum degree ordering algorithm.

An important consequence of Corollary 3.6 is that if v becomes outmatched by u at some stage during the elimination, it is not necessary to update the degree of v until the node u has been eliminated. In the case of the indistinguishable property, degree update is saved since some degrees can be deduced from the representatives. Here, using the outmatched property, redundant degree updates are skipped for those nodes which will not participate in the next round of minimum degree selection.

Again using the example in Figure 3.1, we note that the nodes in $\{22,23,24\}$ are all outmatched by the remaining nodes in $\{15,16,17,18,19,20,21,25\}$. In other words, it is unnecessary to recompute the degrees of nodes 22, 23, and 24 until all others have been eliminated. This is indeed a very powerful technique in improving the performance of the ordering algorithm.

3.5. Element Absorption

The technique of element absorption was first proposed by Duff and Reid in [6]. This idea can be most easily explained in terms of the generalized element approach described in section 3.3. Recall that this approach provides a compact representation of elimination graphs by storing cliques rather than edges. Let $\{K_1, K_2, \dots, K_q\}$ be the set of cliques for the current graph G .

Theorem 3.7: If $K_s \subseteq K_t$ for some s and t , then the graph G can be represented by the clique set:

$$\{K_1, K_2, \dots, K_q\} - K_s.$$

In Theorem 3.7, the clique K_s can be viewed as being absorbed by the clique K_t . The technique of element absorption is to remove any detected clique redundancy in the representation. Reducing the number of cliques will speed up the ordering process since less overhead in manipulating the set of cliques is involved.

A good example to illustrate element absorption is the star graph and its variants. Consider a graph with 8 nodes and its structure is given by the following set of four cliques:

$$\begin{aligned} &\{1, 5\} \\ &\{2, 5, 6\} \\ &\{3, 5, 6, 7\} \\ &\{4, 5, 6, 7, 8\}. \end{aligned}$$

After the elimination of the nodes 1, 2 and 3 (in that order by the minimum degree algorithm), the technique of element absorption reduces the number of cliques to only one: namely, $\{4, 5, 6, 7, 8\}$. The new cliques formed $\{5\}$, $\{5, 6\}$, and $\{5, 6, 7\}$ after the elimination of nodes 1, 2 and 3 respectively can all be absorbed into the clique $\{4, 5, 6, 7, 8\}$. For larger versions of this example, a great reduction in degree update time can be achieved.

3.6. Multiple Elimination

The incomplete degree update described in section 3.4 can also be regarded as a technique using delayed degree update. The enhancement of multiple elimination proposed by Liu [21] extends this idea of delaying the degree update. Instead of performing a degree update step after each minimum degree node selection and elimination graph transformation, the technique of multiple elimination postpones the degree update step to a later stage.

The basis for multiple elimination is quite simple. It makes use of the observation that in the elimination of the node y from the graph G , the structure associated with nodes not in $Adj_G(y)$ remains unchanged. The idea is to suspend the degree update for nodes in $Adj_G(y)$ and select a node with the same degree as y in the remaining subgraph $G - (Adj_G(y) \cup \{y\})$. This process is repeated until there are no nodes in the remaining subgraph with degree $degree_G(y)$. A degree update step is then performed.

In essence, before each degree update step, an independent set of nodes with minimum degree is selected. The fact that the set is independent allows the delay of degree update until the entire set is determined and eliminated. As noted in [21], this modification does not always result in a genuine minimum degree ordering, but it very rarely produces an ordering that is inferior to that provided by the "true" minimum degree algorithm. Indeed, the quality of the resulting ordering is maintained while the execution time of the ordering is reduced.

To understand multiple elimination, we consider the 5-by-5 grid in Figure 3.2, where 8 nodes have been eliminated. As before, nodes marked with "*" have already been eliminated. At this stage, the minimum degree is 5. The set {9, 10, 11, 12} forms an independent set with degree 5. (Note that {15, 17, 19, 20} and {12, 15, 20} are two other different independent sets with degree 5). By eliminating them all together before a degree update helps to reduce the update time. We need to update the degree of nodes in

{13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 24, 25}

only once at the elimination of {9, 10, 11, 12}. On the other hand, using the conventional eliminate-update approach, the nodes in {15, 16, 17, 18, 19, 20, 21, 25} would have to be updated twice during the course of eliminating the nodes 9, 10, 11, 12 in four separate steps.

```

*   *  11  *   *
20 21 22 16 15
10 14 23 13  9
*  18 24 25  *
*  17 12 19  *
```

Figure 3.2: An intermediate stage of the minimum degree ordering.

3.7. External Degree

In [21], Liu suggests the use of external degree instead of true degree in the minimum degree algorithm. In the conventional scheme, the degree used is the number of adjacent nodes in the current elimination graph. By the external degree of a node, we mean the number of nodes adjacent to it that are not indistinguishable from itself.

The motivation comes from the underlying reason for the success of the minimum degree ordering in terms of fill reduction. Eliminating a node of minimum degree implies the formation of the smallest possible clique due to elimination. Since we are using the technique of mass elimination (section 3.1), the size of the resulting clique after mass elimination is the same as the external degree of nodes eliminated by the mass elimination step. For example, in Figure 3.1, since nodes 15 and 16 are indistinguishable, the external degree of node 15 is 7 (and its true degree is 8). Experimental results in [21] show that using external degree rather than true degree yields a reduction in the number of nonzeros in the factor matrix of 3 to 7 percent for the grid model problem.

4. Timing Results of Various Improvement Techniques

So far in the literature, there has been no formal analysis of the time complexity of the minimum degree ordering algorithm. Such an analysis appears to be very difficult, and will be further complicated by the various enhancements to the basic scheme that have been described in the previous section. As partial compensation for the lack of a theoretical analysis, in this section we provide some timing statistics for the various improvement techniques. Our objective is to illustrate to the reader the very substantial gains in efficiency that these enhancements provide. We use the minimum degree ordering algorithm developed by Liu [21], which is now incorporated in SPARSPAK [15]. In order to illustrate the effectiveness of the various techniques, we mask off each improvement strategy in the basic code and obtain timing results.

In all the experiments, the internal data structure to represent elimination graphs is the quotient graph scheme described in [12]. Fixing the data structure allows us to gauge the importance of each improvement technique on the algorithm. Using a different technique for representing the elimination graphs, such as generalized elements or linked lists, would undoubtedly change the absolute times reported. However, we do not think the relative ranking of the various techniques, in terms of their effectiveness in reducing execution time, would be changed significantly.

We choose as our test problem the 180-by-180 regular grid model problem corresponding to the use of a nine-point difference operator. For this matrix example, the number of unknowns is 32,400 and the number of off-diagonal nonzeros in the original matrix is 128,522. We first note that for this particular example, the quality of the resulting minimum degree ordering using external degree is better than that with true degree in terms of fill reduction. The numbers of off-diagonal nonzeros in the Cholesky factors are given by:

External Degree	1,180,771
True Degree	1,273,517

To make a uniform comparison, we used external degree for all the experiments reported hereafter. In Table 4.1, we have tabulated the number of off-diagonal nonzeros in the Cholesky factors and the amount of time in CPU seconds on a SUN 3/50 for each version of the minimum degree algorithm. In [2, page 5] and [3, page 105], Duff provides some timing statistics of various minimum degree ordering implementations between the period 1970 to 1981. Our timing results here are given in correspondence with each enhancement technique and can be regarded as supplements to Duff's statistics.

Version	Minimum Degree Algorithm	Off-diagonal Factor Nonz	Ordering Time
Md#1	Final minimum degree	1,180,771	43.90
Md#2	Md#1 without multiple elimination	1,374,837	57.38
Md#3	Md#2 without element absorption	1,374,837	56.00
Md#4	Md#3 without incomplete deg update	1,374,837	83.26
Md#5	Md#4 without indistinguishable nodes	1,307,969	183.26
Md#6	Md#5 without mass elimination	1,307,969	2289.44

Table 4.1: Statistics on various versions of minimum degree algorithm.

It should be emphasized that the effectiveness of each technique in reducing the ordering time is problem dependent. Here, we have simply used the regular grid as an example for illustration. The version "Md#1" is the current version of the minimum degree ordering routine in SPARSPAK. "Md#6" represents the basic scheme with no enhancement other than the use of clique representation (section 3.3). That is, it is almost the same as the original algorithm as described by Tinney and Walker [28] or by Rose [25], except that external degree is used. We did not optimize the codes for versions "Md#2-Md#6", but rather simply commented out the part associated with the enhancement in version "Md#1". We do not believe that such optimization would change the numbers in any significant way.

From Table 4.1, there is an overall reduction of ordering time by a factor of more than 50. Indeed, the saving has not even included the possible reduction due to the use of some generalized element representation.

Another aspect worth mentioning is the number of off-diagonal nonzeros in the Cholesky factors. For this example, the final version "MD#1" produces an ordering with 10% fewer nonzeros than that provided by those with some of the enhancements removed. This suggests that Liu's multiple elimination technique has an important effect on the quality of the ordering produced by the algorithm, in addition to reducing its the execution time. We shall return to this point in the next section.

A final point to note is that the reduction in the number of nonzeros in the factor matrix of 10% will lead to a much larger relative reduction in the numerical factorization time. We provide some experiments in the next section to illustrate this fact.

5. On Tie-Breaking Strategies

5.1. An Example: Regular Grid

Let A be a matrix associated with the k -by- k regular grid. It is known that the best ordering on A will yield a Cholesky factor with at least $O(k^2 \log_2 k)$ nonzeros [20]. Furthermore, it is shown by George [10] that a nested dissection ordering on the grid is one such optimal ordering. Nested dissection can be viewed as a minimum degree ordering with a perfect tie-breaking strategy for the grid.

On the other hand, it has been a long outstanding question as to how bad a minimum degree ordering can be for the grid problem. Or, to phrase the question differently, can the minimum degree algorithm produce an ordering for A for which the resulting Cholesky factor has more than $O(k^2 \log_2 k)$ nonzeros (in the order of magnitude sense)? This problem can be interpreted as a search for an imperfect tie-breaking strategy for the minimum degree algorithm on the grid. Recently, Berman and Schnitger have devised one such tie-breaking scheme [26], yielding an ordering for which the Cholesky factor is shown to have $O(k^{2 \log_3 4})$ nonzeros.

For the regular grid, there is therefore a wide class of minimum degree orderings, ranging from the optimal nested dissection ordering to the ordering by Berman and Schnitger. This shows the significance of choosing an effective tie-breaking strategy for the minimum degree algorithm. In the next subsection, we further illustrate its significance by considering a number of existing implementations of the minimum degree ordering.

5.2. Current Implementations: Random Tie-breaking

Perhaps the most well known implementations of the minimum degree algorithm are those found in the Harwell MA27 code [5], the Waterloo SPARSPAK package [15], and the Yale sparse matrix package (YSMP) [7]. They incorporate most of the enhancements described in section 3. It is not our objective here to compare the relative execution speed of these three implementations. Instead, we want to illustrate the practical importance of having an effective tie-breaking strategy.

In some sense, Liu's multiple elimination technique [21] provides a limited form of tie-breaking for the minimum degree algorithm. It forces the algorithm to search for a maximal set of independent minimum degree nodes before a degree update is performed. Other than this, none of the three implementations (Harwell, SPARSPAK, YSMP) has incorporated any tie-breaking strategy. The selection of the next minimum degree node from the candidate set is, in effect, random, since the initial ordering essentially determines the way ties are resolved.

To illustrate this, we again use the 180-by-180 grid problem. A random permutation is applied to rearrange the rows and columns before passing the matrix to the minimum degree ordering routine. In other words, the algorithm is ordering a different form of the same matrix problem. This is repeated ten times and we have tabulated in Table 5.1 the best and the worst orderings for all three implementations. For a basis for comparison, we have also included the case when the matrix problem is presented in a row-by-row initial ordering. In the table, "Factor Nonz" is the number of off-diagonal nonzeros in the resulting Cholesky factor, while "Factor Opns." is the number of multiplicative operations to perform the factorization (in millions of operations).

Minimum Degree Code	Row-by-row			Ten Random Initial Orderings					
	Initial Ordering			Best Min Deg Ordering			Worst Min Deg Ordering		
	Order Time	Factor Nonz	Factor Opns.	Order Time	Factor Nonz	Factor Opns.	Order Time	Factor Nonz	Factor Opns.
Harwell	115.2	1,605,697	132.0M	116.2	1,731,733	144.6M	116.2	1,821,235	174.9M
Sparspak	43.9	1,180,771	62.2M	58.5	1,567,324	118.0M	58.8	1,651,998	139.7M
YSMP	49.0	1,241,967	71.9M	52.9	1,710,895	139.5M	53.3	1,868,070	191.3M

Table 5.1: Best and worst minimum degree orderings from ten runs on the 180-by-180 grid with different initial orderings.

Results in Table 5.1 demonstrate that a random form of tie-breaking (as basically used by the three implementations) is clearly not enough to ensure a good quality ordering from a practical standpoint. Indeed, among the three YSMP minimum degree orderings tabulated, there is a difference of about 50% in terms of factor matrix nonzeros and over 160% in terms of factorization operation counts. The difference is even more dramatic among all the orderings in Table 5.1. The worst one requires more than three times as many arithmetic operations as that required by the best ordering. Note that they are all variants of the same basic minimum degree algorithm on the same matrix problem. This illustrates the practical importance of an effective tie-breaking strategy for the minimum degree algorithm.

It should be emphasized that the ordering time used by the minimum degree algorithm is negligible when compared to the actual numerical factorization time. Even a small percentage reduction in factorization operations would easily offset a relatively large increase in ordering time that might result from adding sophisticated tie-breaking strategies to the algorithm. To illustrate this, we actually performed the numerical factorization and triangular solution of the 180-by-180 grid problem using the best ordering from Table 5.1 (SPARSPAK minimum degree with row-by-row initial ordering). On a SUN 3/50 the numerical factorization requires 1847.08 CPU seconds (about 62.2 million arithmetic operations) and the forward and backward solution requires 78.52 seconds. The ordering time of 43.9 seconds is therefore an insignificant fraction of the overall cost.

For the sake of comparison, we also solved the same linear system using the worst minimum degree ordering from SPARSPAK. The numerical factorization required 4529.24 CPU seconds for the 139.7 million operations. This strongly suggests that it is important to devise effective tie-breaking strategies for reducing fill (and operations). The potential pay-off in the subsequent numerical phase can be tremendous.

It may be worthwhile to point out one major difference among these three implementations. Both MA27 and YSMP use linked lists in the representation of the sequence of elimination graphs for the generalized element model. By Theorem 4.4, the amount of storage required will never exceed the storage requirement for the linked list representation of the original given graph. On the other hand, SPARSPAK uses a quotient graph storage scheme as described in [12], whereby the ordering is performed "in-place" within the given adjacency structure representation. Higher overhead in manipulating the structure is required in exchange for a reduction of approximately 50% in storage requirements.

6.2. Adaptation to Some Applications

The minimum degree algorithm (with the enhancements) is a general purpose ordering scheme that can be applied to any symmetric matrix structure (or undirected graph). It is interesting to note that in some situations, it may be advantageous to tailor the general minimum degree algorithm to specific applications. In this section, we shall mention two such applications.

Consider computing the least squares solution of the large m by n sparse system

$$A x \approx b,$$

where $m \geq n$ and A has full rank. It is well known that such a problem can be solved via the orthogonal decomposition of the matrix A into $Q \begin{pmatrix} R \\ 0 \end{pmatrix}$, where Q is an orthogonal matrix and R is upper triangular. The solution to the least square problem above is known to be equivalent mathematically to solving the normal equations

$$A^T A x = A^T b.$$

Based on this connection, George and Heath [11] provide an overall scheme to solve the sparse system $A x \approx b$. In particular, they suggest solving the equivalent system:

$$(AP)(P^T x) \approx b,$$

where P is a minimum degree ordering of the symmetric matrix $A^T A$. This is motivated by the observation that the factor R is mathematically the same as the Cholesky factor of the matrix $A^T A$ so that using a minimum degree ordering of $A^T A$ will result in a sparse factor R .

To use the minimum degree algorithm in this setting, one can form the structure of the matrix $M = A^T A$ and then apply the ordering algorithm to M . A natural question is whether one can adapt the algorithm to determine a minimum degree ordering for M directly from the structure of A . We first note that the structure of each row of A corresponds to a clique in the graph associated with M . Therefore we can use the generalized element approach to represent M as a sequence of cliques consisting of the rows of A . Since this is an intermediate data structure used by the minimum degree algorithm, one can easily adapt the algorithm by allowing a set of cliques as the initial graph representation. This will save the extra step of forming the graph structure of the matrix $M = A^T A$ explicitly. Adapting the algorithm so that it efficiently finds an ordering of M , given A , is also important in the context of solving sparse indefinite systems using Gaussian elimination with partial pivoting [17].

A different application is in the numerical solution of sparse linear systems arising from the finite element method. The structure of the matrix from such systems is governed by the underlying finite element mesh. The nodes associated with each element form a clique.

In practice, the matrix structure is often provided in terms of the elements; that is, it is given in the form of node membership in each element. Of course, the adjacency structure of the matrix can be generated based on this element-node membership information and then passed to the minimum degree ordering. An alternative is to perform the ordering directly on the element-node information without forming the adjacency structure.

It is interesting to note that this is exactly the same situation as in the case of the sparse least squares problem. Consider the structure of the matrix B with each row corresponding to an element of the finite element mesh and each column to a node. The structure of each row of B is given by the node membership of the associated element. Then, it can be verified that the structure of the finite element matrix is the same as the structure of $B^T B$. Therefore, a minimum degree ordering of the finite element matrix can be obtained by performing the ordering on B in the same manner as the sparse least squares problem.

This approach has another important advantage. In the use of higher order elements, the representation of the finite element matrix structure implicitly by B usually requires significantly less storage than the explicit adjacency structure representation. For example, consider the k -by- k grid where each square/element is associated with 9 variables/nodes (with one interior node, one at each corner, and one along each side). Using the adjacency structure representation, we need storage amounting to about $46k^2$ items. On the other hand, the storage requirement for the structure of B is only $16k^2$.

7. Future Research Directions

The various enhancements to the basic minimum degree ordering algorithm over the past fifteen years have made the algorithm a truly practical approximate solution to the NP-complete minimum fill ordering problem [29]. The ordering time has been drastically reduced, as our example with an over 50-fold reduction in ordering time in Table 4.1 illustrates. This makes the ordering time a small fraction of the overall solution time for a given large sparse problem.

In the authors' opinion, the next significant advance of this algorithm will probably be the development of effective and practical tie-breaking strategies in the selection of minimum degree nodes. Given the possible variation in the quality of the orderings as exemplified by experimental results in Table 5.1, significant reduction in factorization storage requirement can be attained by having a good tie-breaking strategy. More importantly, reducing the number of nonzeros in the Cholesky factor implies an even more substantial reduction in operation counts (and hence factorization time). Indeed, the potential gain makes it apparently very worthwhile to invest more time in the ordering phase to determine a better quality ordering.

In section 6, we have offered a "partial" solution to the tie-breaking problem based on preordering. The reverse Cuthill-McKee profile ordering is used to pre-order a matrix before the minimum degree algorithm is applied. Some success in "stabilizing" the overall ordering phase is reported in section 5.

Another promising avenue for tie-breaking strategies is based on the notion of independent sets. As pointed out in section 5, Liu's multiple elimination technique [21] provides a form of tie-breaking by selecting a maximal independent set of nodes of minimum degree. An apparently better strategy would be to select at each stage a maximum independent set of minimum degree nodes. In view of the NP-completeness of this maximum independent set problem, we must instead rely on practical approximate solutions to find a "nearly" maximum independent set from the set of minimum degree nodes. Some of these ideas are currently under investigation by the authors.

Another related problem of practical interest is to generate a fill-reducing ordering appropriate for parallel elimination. Some recent progress in this area can be found in [22]. In terms of the minimum degree algorithm, we are interested in finding a suitable tie-breaking strategy so that the resulting minimum degree ordering is appropriate for parallel elimination. nh

We have tried the following strategy. At a given stage of the ordering, let S be the set of nodes already eliminated. If there are more than one node in the uneliminated set with the current minimum degree, we select a node y with the smallest connected component containing y in the subgraph of $S \cup \{y\}$. Intuitively, we try to enlarge the smaller components of the subgraph defined by S within the guideline of the minimum degree algorithm. Experiments on this (and other slight variants of this) tie-breaking strategy were tried. However, in terms of parallelism, the resulting orderings are not significantly different from those reported in [22].

References

1. I. S. DUFF, "MA28 - A set of FORTRAN subroutines for sparse unsymmetric linear equations," Technical Report AERE R-8730, Harwell, 1977.
2. I. S. DUFF, "A sparse future," in *Sparse Matrices and their Uses*, ed. I.S. Duff, pp. 1-9, Academic Press, 1981.
3. I. S. DUFF, "Research directions in sparse matrix computations," in *Studies in Numerical Analysis, Volume 24*, ed. G.H. Golub, pp. 83-139, Math Association of America, 1984.
4. I. S. DUFF AND J. K. REID, "A comparison of sparsity orderings for obtaining a pivotal sequence in Gaussian elimination," *JIMA*, vol. 14, pp. 281-291, 1974.
5. I. S. DUFF AND J. K. REID, "MA27 - A set of FORTRAN subroutines for solving sparse symmetric sets of linear equations," Report AERE R 10533, Harwell, 1982.
6. I. S. DUFF AND J. K. REID, "The multifrontal solution of indefinite sparse symmetric linear systems," *TOMS*, vol. 9, pp. 302-325, 1983.
7. S. C. EISENSTAT, M. C. GURSKY, M. H. SCHULTZ, AND A. H. SHERMAN, "The Yale Sparse Matrix Package, I. The symmetric code," *Internat. J. Numer. Meth. Engrg.*, vol. 18, pp. 1145-1151, 1982.
8. S. C. EISENSTAT, M. H. SCHULTZ, AND A. H. SHERMAN, "Applications of an element model for Gaussian elimination," in *Sparse Matrix Computations*, ed. J. R. Bunch and D. J. Rose, pp. 85-96, Academic Press, 1976.
9. S. C. EISENSTAT, M. H. SCHULTZ, AND A. H. SHERMAN, "Algorithms and data structures for sparse symmetric Gaussian elimination," *SISSC*, vol. 2, pp. 225-237, 1981.
10. J. A. GEORGE, "Nested dissection of a regular finite element mesh," *SINUM*, vol. 10, pp. 345-363, 1973.
11. J. A. GEORGE AND M. T. HEATH, "Solution of sparse linear least squares problems using Givens rotations," *LAA*, vol. 34, pp. 69-83, 1980.
12. J. A. GEORGE AND J. W. H. LIU, "A fast implementation of the minimum degree algorithm using quotient graphs," *TOMS*, vol. 6, pp. 337-358, 1980.
13. J. A. GEORGE AND J. W. H. LIU, "A minimal storage implementation of the minimum degree algorithm," *SINUM*, vol. 17, pp. 282-299, 1980.
14. J. A. GEORGE AND J. W. H. LIU, *Computer Solution of Large Sparse Positive Definite Systems*, Prentice-Hall Inc., Englewood Cliffs, N.J., 1981.

15. J. A. GEORGE, J. W. H. LIU, AND E. NG, "User guide for SPARSPAK, Waterloo Sparse Linear Equations Package," Res. Rep. CS-78-30 (revised 1980), DCS, University of Waterloo, Waterloo, Ont. Canada, 1980.
16. J. A. GEORGE AND D. MCINTYRE, "On the application of the minimum degree algorithm to finite element systems," *SINUM*, vol. 15, pp. 90-112, 1978.
17. J. A. GEORGE AND E. NG, "Symbolic factorization for sparse Gaussian elimination with partial pivoting," *SISSC*, 1987. (to appear)
18. M. C. GOLUMBIC, *Algorithmic Graph Theory and Perfect Graphs*, Academic Press, 1980.
19. C. HEMPEL, *Minimum degree can miss by more than a constant factor*, DCS, Cornell University, 1985. Unpublished manuscript
20. A. J. HOFFMAN, M. S. MARTIN, AND D. J. ROSE, "Complexity bounds for regular finite difference and finite element grids," *SIAM J. Numer. Anal.*, vol. 10, pp. 364-369, 1973.
21. J. W. H. LIU, "Modification of the minimum degree algorithm by multiple elimination," *TOMS*, vol. 11, pp. 141-153, 1985.
22. J. W. H. LIU, "Reordering sparse matrices for parallel elimination," Tech. Report CS-87-01, DCS, York University, 1987.
23. H. M. MARKOWITZ, "The elimination form of the inverse and its application to linear programming," *Management Sci.*, vol. 3, pp. 255-269, 1957.
24. D. J. ROSE, "Symmetric elimination on sparse positive definite systems and potential flow network problem," Ph.D Thesis, Harvard University, 1970.
25. D. J. ROSE, "A graph-theoretic study of the numerical solution of sparse positive definite systems of linear equations," in *Graph Theory and Computing*, ed. R. Read, pp. 183-217, Academic Press, 1972.
26. P. BERMAN AND G. SCHNITGER, *On the performance of the minimum degree heuristic for Gaussian elimination*, Dept. of Computer Science, The Pennsylvania State University, 1987. Unpublished manuscript
27. B. SPEELPENNING, "The generalized element method," Report UIUCDCS-R-78-946, DCS, Univ. of Illinois at Urbana-Champaign, 1978.
28. W. F. TINNEY AND J. W. WALKER, "Direct solutions of sparse network equations by optimally ordered triangular factorization," *Proc. IEEE*, vol. 55, pp. 1801-1809, 1967.
29. M. YANNAKAKIS, "Computing the minimum fill-in is NP-complete," *SIAM J. Alg. & Disc. Meth.*, vol. 2, pp. 77-79, 1981.

INTERNAL DISTRIBUTION

- | | | | |
|--------|-------------------|--------|---|
| 1. | E. Chu | 25-29. | R. C. Ward |
| 2. | J. B. Drake | 30. | M. A. Williams |
| 3. | E. L. Frome | 31. | D. G. Wilson |
| 4. | G. A. Geist | 32. | A. Zucker |
| 5-9. | J. A. George | 33. | P. W. Dickson (Consultant) |
| 10. | L. J. Gray | 34. | G. H. Golub (Consultant) |
| 11-12. | R. F. Harbison | 35. | R. M. Haralick (Consultant) |
| 13. | M. T. Heath | 36. | D. Steiner (Consultant) |
| 14-18. | J. K. Ingersoll | 37. | Central Research Library |
| 19. | F. C. Maienschein | 38. | K-25 Plant Library |
| 20. | T. J. Mitchell | 39. | ORNL Patent Office |
| 21. | E. G. Ng | 40. | Y-12 Technical Library/
Document Reference Station |
| 22. | G. Ostrouchov | 41. | Laboratory Records - RC |
| 23. | C. H. Romine | 42-43. | Laboratory Records Department |
| 24. | S. Thompson | | |

EXTERNAL DISTRIBUTION

44. Dr. Donald M. Austin, Office of Scientific Computing, Office of Energy Research, ER-7, Germantown Building, U.S. Department of Energy, Washington, DC 20545
45. Dr. Robert G. Babb, Department of Computer Science and Engineering, Oregon Graduate Center, 19600 N.W. Walker Road, Beaverton, OR 97006
46. Dr. Jesse L. Barlow, Department of Computer Science, Pennsylvania State University, University Park, PA 16802
47. Prof. Ake Bjorck, Department of Mathematics, Linkoping University, Linkoping 58183, Sweden
48. Dr. James C. Browne, Department of Computer Sciences, University of Texas, Austin, TX 78712
49. Dr. Bill L. Buzbee, C-3, Applications Support & Research, Los Alamos National Laboratory, P.O. Box 1663, Los Alamos, NM 87545
50. Dr. Donald A. Calahan, Department of Electrical and Computer Engineering, University of Michigan, Ann Arbor, MI 48109
51. Dr. Tony Chan, Department of Computer Science, Yale University, P.O. Box 2158 Yale Station, New Haven, CT 06520
52. Dr. Jagdish Chandra, Army Research Office, P.O. Box 12211, Research Triangle Park, North Carolina 27709
53. Dr. Paul Concus, Mathematics and Computing, Lawrence Berkeley Laboratory, Berkeley, CA 94720
54. Dr. Jane K. Cullum, IBM T. J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598

55. Dr. George Cybenko, Department of Computer Science, Tufts University, Medford, MA 02155
56. Dr. George J. Davis, Department of Mathematics, Georgia State University, Atlanta, GA 30303
57. Dr. Jack J. Dongarra, Mathematics and Computer Science Division, Argonne National Laboratory, 9700 South Cass Avenue, Argonne, IL 60439
58. Dr. Stanley Eisenstat, Department of Computer Science, Yale University, P.O. Box 2158 Yale Station, New Haven, CT 06520
59. Dr. Howard C. Elman, Computer Science Department, University of Maryland, College Park, MD 20742
60. Dr. Albert M. Erisman, Boeing Computer Services, 565 Andover Park West, Tukwila, WA 98188
61. Dr. Geoffrey C. Fox, Booth Computing Center 158-79, California Institute of Technology, Pasadena, CA 91125
62. Dr. Paul O. Frederickson, Computing Division, Los Alamos National Laboratory, Los Alamos, NM 87545
63. Dr. Fred N. Fritsch, L-300, Mathematics and Statistics Division, Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, CA 94550
64. Dr. Robert E. Funderlic, Department of Computer Science, North Carolina State University, Raleigh, NC 27650
65. Dr. Dennis B. Gannon, Computer Science Department, Indiana University, Bloomington, IN 47405
66. Dr. David M. Gay, Bell Laboratories, 600 Mountain Avenue, Murray Hill, NJ 07974
67. Dr. C. William Gear, Computer Science Department, University of Illinois, Urbana, Illinois 61801
68. Dr. W. Morven Gentleman, Division of Electrical Engineering, National Research Council, Building M-50, Room 344, Montreal Road, Ottawa, Ontario, Canada k1A 0R8
69. Prof. Gene H. Golub, Department of Computer Science, Stanford University, Stanford, CA 94305
70. Dr. Joseph F. Grcar, Division 8331, Sandia National Laboratories, Livermore, CA 94550
71. Dr. Don E. Heller, Physics and Computer Science Department, Shell Development Co., P.O. Box 481, Houston, TX 77001
72. Dr. Robert E. Huddleston, Computation Department, Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, CA 94550
73. Dr. Ilse Ipsen, Department of Computer Science, Yale University, P.O. Box 2158 Yale Station, New Haven, CT 06520
74. Dr. Harry Jordan, Department of Electrical and Computer Engineering, University of Colorado, Boulder, CO 80309

75. Dr. Linda Kaufman, Bell Laboratories, 600 Mountain Avenue, Murray Hill, NJ 07974
76. Dr. Robert J. Kee, Applied Mathematics Division 8331, Sandia National Laboratories, Livermore, CA 94550
77. Ms. Virginia Klema, Statistics Center, E40-131, MIT, Cambridge, MA 02139
78. Dr. Richard Lau, Office of Naval Research, 1030 E. Green Street, Pasadena, CA 91101
79. Dr. Alan J. Laub, Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA 93106
80. Dr. Robert L. Launer, Army Research Office, P.O. Box 12211, Research Triangle Park, North Carolina 27709
81. Prof. Peter D. Lax, Director, Courant Institute of Mathematical Sciences, New York University, 251 Mercer Street, New York, NY 10012
82. Dr. Michael R. Leuze, Computer Science Department, Box 1679 Station B, Vanderbilt University, Nashville, TN 37235
83. Dr. Joseph Liu, Department of Computer Science, York University, 4700 Keele Street, Downsview, Ontario, Canada M3J 1P3
84. Dr. Franklin Luk, Electrical Engineering Department, Cornell University, Ithaca, NY 14853
85. Dr. Thomas A. Manteuffel, Computing Division, Los Alamos National Laboratory, Los Alamos, NM 87545
86. Dr. Paul C. Messina, Applied Mathematics Division, Argonne National Laboratory, Argonne, IL 60439
87. Dr. Cleve Moler, Intel Scientific Computers, 15201 N.W. Greenbrier Parkway, Beaverton, OR 97006
88. Dr. Dianne P. O'Leary, Computer Science Department, University of Maryland, College Park, MD 20742
89. Maj. C. E. Oliver, Office of the Chief Scientist, Air Force Weapons Laboratory, Kirtland Air Force Base, Albuquerque, NM 87115
90. Dr. James M. Ortega, Department of Applied Mathematics, University of Virginia, Charlottesville, VA 22903
91. Prof. Chris Paige, Basser Department of Computer Science, Madsen Building F09, University of Sydney, N.S.W., Sydney, Australia 2006
92. Dr. John F. Palmer, NCUBE Corporation, 915 E. LaVieve Lane, Tempe, AZ 85284
93. Prof. Beresford N. Parlett, Department of Mathematics, University of California, Berkeley, CA 94720
94. Prof. Merrell Patrick, Department of Computer Science, Duke University, Durham, NC 27706
95. Dr. Robert J. Plemmons, Departments of Mathematics and Computer Science, North Carolina State University, Raleigh, NC 27650

96. Dr. John K. Reid, CSS Division, Building 8.9, AERE Harwell, Didcot, Oxon, England OX11 0RA
 97. Dr. John R. Rice, Computer Science Department, Purdue University, West Lafayette, IN 47907
 98. Dr. Garry Rodrigue, Numerical Mathematics Group, Lawrence Livermore Laboratory, Livermore, CA 94550
 99. Dr. Donald J. Rose, Department of Computer Science, Duke University, Durham, NC 27706
 100. Dr. Ahmed H. Sameh, Computer Science Department, University of Illinois, Urbana, IL 61801
 101. Dr. Michael Saunders, Systems Optimization Laboratory, Operations Research Department, Stanford University, Stanford, CA 94305
 102. Dr. Robert Schreiber, Department of Computer Science, Rensselaer Polytechnic Institute, Troy, NY 12180
 103. Dr. Martin H. Schultz, Department of Computer Science, Yale University, P.O. Box 2158 Yale Station, New Haven, CT 06520
 104. Dr. David S. Scott, Intel Scientific Computers, 15201 N.W. Greenbrier Parkway, Beaverton, OR 97006
 105. Dr. Lawrence F. Shampine, Numerical Mathematics Division 5642, Sandia National Laboratories, P.O. Box 5800, Albuquerque, NM 87115
 106. Dr. Danny C. Sorensen, Mathematics and Computer Science Division, Argonne National Laboratory, 9700 South Cass Avenue, Argonne, IL 60439
 107. Prof. G. W. Stewart, Computer Science Department, University of Maryland, College Park, MD 20742
 108. Capt. John P. Thomas, Air Force Office of Scientific Research, Building 410, Bolling Air Force Base, Washington, DC 20332
 109. Prof. Charles Van Loan, Department of Computer Science, Cornell University, Ithaca, NY 14853
 110. Dr. Robert G. Voigt, ICASE, MS 132-C, NASA Langley Research Center, Hampton, VA 23665
 111. Dr. Andrew B. White, Computing Division, Los Alamos National Laboratory, Los Alamos, NM 87545
 112. Mr. Patrick H. Worley, Computer Science Department, Stanford University, Stanford, CA 94305
 113. Dr. Arthur Wouk, Army Research Office, P.O. Box 12211, Research Triangle Park, North Carolina 27709
 114. Dr. Margaret Wright, Systems Optimization Laboratory, Operations Research Department, Stanford University, Stanford, CA 94305
 115. Office of Assistant Manager for Energy Research and Development, Department of Energy, Oak Ridge Operations Office, Oak Ridge, TN 37830
- 116-146. Technical Information Center