

Comparative Analysis of the Cuthill-McKee and
the Reverse Cuthill-McKee Ordering Algorithms
For Sparse Matrices

Wai-Hung Liu^{*}

Andrew H. Sherman⁺

Research Report #28

^{*} Department of Applied Analysis and Computer Science, University of Waterloo, Waterloo, Ontario, Canada. This author's work supported in part by University of Waterloo Graduate Bursary funds.

⁺ Department of Computer Science, Yale University, New Haven, Connecticut 06520. This author's work supported in part by the Office of Naval Research under grant N0014-67-0097-0016.

Abstract

In this paper we examine the popular Cuthill-McKee algorithm for ordering the unknowns and equations in sparse, positive definite linear systems of equations. For a given linear system $Ax = b$, this algorithm is designed to produce a permutation matrix P such that PAP^T has a small bandwidth. If we wish to exploit zeros in the band of A which occur before the first nonzero in each row and column, it has been experimentally observed that reversing the ordering produced by the Cuthill-McKee algorithm is often very much better than the original ordering in terms of the amount of storage and work required to factor A . We prove that for band elimination methods, the two orderings are equivalent and that, surprisingly, the reverse ordering is always at least as good as the original one when envelope elimination techniques are used. We give a condition on the matrix A under which the reverse ordering is strictly better than the original one, and we include several numerical experiments and analyses of practical examples to illustrate our results.

1. Introduction

Consider the N by N sparse linear algebraic system

$$A\bar{x} = \bar{b} \quad (1.1)$$

where A is symmetric and positive definite. To solve (1.1) we factor A into the product LDL^T , where L is unit lower triangular and D is a positive diagonal matrix. We then solve $L\bar{z} = \bar{b}$, $D\bar{y} = \bar{z}$, and $L^T\bar{x} = \bar{y}$. Direct methods such as this fall into two basic classes: the general sparse methods, which ignore all the zeros in L , and the band or envelope methods, which ignore only those zeros outside a particular region of L . In this paper we consider mainly methods in the latter group, particularly with respect to the effect on such methods of various ordering strategies for the variables and equations of the system (1.1).

Let P be any permutation matrix. The permuted linear system

$$PAP^T(P\bar{x}) = P\bar{b} \quad (1.2)$$

remains sparse, symmetric, and positive definite. It is well known that PAP^T possesses a triangular factorization LDL^T and that the factorization is numerically stable [8], so that we may solve (1.2) rather than (1.1). The possible advantage of solving the permuted system is that the solution might require a smaller number of arithmetic operations and/or storage locations. For the methods with which we deal in this paper, this means that the band or envelope of PAP^T may be smaller than the corresponding structure of A .

A popular ordering strategy for band solution of (1.1) is the Cuthill-McKee algorithm [2]. It tries to find a permutation matrix P for which PAP^T has

a small bandwidth. The reverse of the Cuthill-McKee ordering has been observed to produce a permutation matrix P for which PAP^T has a small envelope [3,1]. In this paper we give a detailed analysis of these two ordering strategies for band and envelope methods. In particular, we will show that the Reverse Cuthill-McKee algorithm (RCM) always produces an ordering at least as good as that produced with the Cuthill-McKee algorithm (CM), and we will give estimates of the improvement with RCM for several classes of problems arising in the application of the finite element method.

2. Definitions

Let A be a given N by N symmetric or lower triangular matrix. For the i -th row of A , $i = 1, 2, \dots, N$, we define

$$f_i(A) \equiv \min \{j : A_{ij} \neq 0\}, \quad (2.1)$$

$$\beta_i(A) \equiv i - f_i(A), \quad (2.2)$$

$$\omega_i(A) \equiv |\{k : k > i \text{ and } \exists \ell \leq i \text{ such that } A_{k\ell} \neq 0\}|; \quad (2.3)$$

that is, $f_i(A)$ is the column subscript of the first non-zero element of the i -th row of A , $\beta_i(A)$ is the "bandwidth" of the i -th row of A , and $\omega_i(A)$ is the "frontwidth" of the i -th row of A (see Figure 2.1). Following Cuthill and McKee [2], we define the bandwidth of A by

$$b(A) \equiv \max \{|i-j| : A_{ij} \neq 0\} \quad (2.4)$$

and note that

$$b(A) = \max \{\beta_i(A)\}. \quad (2.5)$$

The band, envelope, and transpose envelope structures of A are now defined by

$$B(A) \equiv \{A_{ij} : j \leq i \text{ and } |i-j| \leq b(A)\}, \quad (2.6)$$

$$\text{Env}(A) \equiv \{A_{ij} : f_i(A) \leq j \leq i\}. \quad (2.7)$$

$$\text{Tenv}(A) \equiv \{A_{ij} : j \leq i \text{ and } \exists k \geq i \text{ such that } A_{kj} \neq 0\}. \quad (2.8)$$

Clearly

$$\text{Tenv}(A) = \text{Env}(A^{\#}) \quad (2.9)$$

where $A^{\#}$ is the transpose of A about its minor diagonal.

We may easily express the size of the envelope of A in terms of the

$\beta_i(A)$ or the $\omega_i(A)$:

$$|\text{Env}(A)| = N + \sum_{i=1}^N \beta_i(A) = N + \sum_{i=1}^N \omega_i(A). \quad (2.10)$$

Assuming that A has the symmetric factorization LDL^T , we define

$$\text{Fill}(A) \equiv \{A_{ij} : A_{ij} = \emptyset \text{ and } L_{ij} \neq \emptyset\}. \quad (2.11)$$

Since $\text{Fill}(A) \subseteq \text{Env}(A)$ (see [5]), the number of locations required to store all of the nonzero entries of L is always less than or equal to $|\text{Env}(A)|$. Methods of the class considered in this paper do not exploit zeros within the envelope (resp. band) of A , so that the storage for L will be exactly $|\text{Env}(A)|$ (resp. $|B(A)|$).

The following lemma on general sparse symmetric decompositions which exploit all of the zeros in L is due to George [4].

Lemma 2.1: If the symmetric factorization of matrix A into the product LDL^T requires $\theta(A)$ multiplicative operations, then

$$\theta(A) \leq \frac{1}{2} \sum_{i=1}^N \omega_i(A) [\omega_i(A) + 3] \quad (2.12)$$

with equality exactly when L has a full envelope.

In order to facilitate what follows, we introduce the undirected graph $G(A) = (X(A), E(A))$ associated with the symmetric matrix A . Here $X(A) = \{x_1, x_2, \dots, x_N\}$ is the set of vertices corresponding to and labelled as the rows of A , and $E(A)$ is the set of edges, where $\{x_i, x_j\} \in E(A)$ if and only if $A_{ij} \neq \emptyset$. The adjacency and the degree of a vertex $x_i \in X(A)$ are then defined by

$$\text{Deg}(x_i) \equiv |\text{Adj}(x_i)| \equiv |\{x_j : \{x_i, x_j\} \in E(A)\}|. \quad (2.13)$$

For any permutation matrix P , the graphs $G(A)$ and $G(PAP^T)$ are structurally identical, but the node labels in $G(PAP^T)$ have been permuted according to P .

We now restate several of our definitions in terms of the graph $G(A)$:

$$f_i(A) \equiv \min \{j : x_j \in \text{Adj}(x_i)\}; \quad (2.14)$$

$$\text{Env}(A) \equiv \{\{x_i, x_j\} : f_i(A) \leq j \leq i\}; \quad (2.15)$$

$$\text{Tenv}(A) \equiv \{\{x_i, x_j\} : j \leq i \text{ and } \exists k \geq i \text{ such that } x_k \in \text{Adj}(x_j)\}; \quad (2.16)$$

$$\text{Fill}(A) \equiv \{\{x_i, x_j\} : \{x_i, x_j\} \notin E(A) \text{ and } \{x_i, x_j\} \in E(L)\}. \quad (2.17)$$

3. The Cuthill-McKee Algorithm

The ordering algorithm proposed by Cuthill and McKee [2] is most easily described in terms of labelling the graph structure $G(A)$ associated with the N by N matrix A . Although our results and the Cuthill-McKee algorithm apply to graphs with any number of connected components, we shall henceforth assume that $G(A)$ is connected, or, equivalently, that the matrix A is irreducible. In the following algorithm, Q , R , and S are first-in-first-out data structures known as queues. The operation of adding an element to a queue (denoted by $Q \leftarrow Q, x_i$) is always performed by placing the element at the end or back of the queue. The operation of removing the first element of the queue (denoted by $x_i \leftarrow \text{Top}(Q)$) assigns the first or oldest element of the queue to x_i and deletes that element from the queue. Whenever a set of vertices is assigned to a queue, the vertices are assumed to be ordered by increasing vertex degree in $G(A)$. The empty set is denoted by Λ .

Step 0: $i \leftarrow 1$; $Q \leftarrow \Lambda$; $R \leftarrow X(A)$;

$x_1 \leftarrow \text{Top}(R)$; $S \leftarrow \text{Adj}(x_1) \cap R$; $R \leftarrow R - S$;

Step 1: *while* $S \neq \Lambda$ *do* ($i \leftarrow i + 1$; $x_i \leftarrow \text{Top}(S)$; $Q \leftarrow Q, x_i$);

if $R = \Lambda$ *then stop*;

$z \leftarrow \text{Top}(Q)$; $S \leftarrow \text{Adj}(z) \cap R$; $R \leftarrow R - S$; *go to* Step 1;

As noted in [2], the Cuthill-McKee (CM) ordering algorithm corresponds to the generation of a spanning tree for the graph $G(A)$ in a breadth-first or level-by-level fashion. Such a spanning tree is formed by adding each new vertex as far to the top and left as possible.

Lemma 3.1: Let A_c be the matrix A ordered by the Cuthill-McKee algorithm.

Then A_c satisfies the monotone envelope property that if $i \leq j$, then $f_i \leq f_j$.

Furthermore, for $i > 1$, $f_i < i$.

Proof: Let $G(A_c)$ be the labelled graph associated with A_c , and let

$X = \{x_1, x_2, \dots, x_N\}$ be the ordered vertex set of $G(A_c)$. Then the first part of the lemma is equivalent to the condition that if x_i precedes x_j in X , then

$$I \equiv \min \{\ell : x_\ell \in \text{Adj}(x_i)\} \leq J \equiv \min \{\ell : x_\ell \in \text{Adj}(x_j)\}.$$

But this follows directly from the breadth-first nature of the Cuthill-McKee algorithm, for if $J < I$, then the algorithm would have labelled x_j before x_i , contradicting the assumption that x_i precedes x_j .

The second part of the lemma follows immediately from the first, for if there were a row with $f_i = i$, then the monotone envelope property would imply that A_c is reducible, contradicting the assumption made earlier.

The following theorem is a restatement, in our notation, of a result of George and Liu [5]. By Lemma 3.1, both it and its corollary apply to A_c .

Theorem 3.2: Let $A = LDL^T$ with $f_i(A) < i$ for $i > 1$. Then $\text{Env}(L)$ is full.

Corollary 3.3: Let $A = LDL^T$ with $f_i(A) < i$ for $i > 1$. Then the number of nonzero entries in L is exactly $|\text{Env}(A)|$, and the forward and backward solving can be done in $2 \cdot |\text{Env}(A)| - N$ multiplications.

We now turn to the problem of estimating the number of multiplications required for the symmetric factorization of an envelope-oriented N by N matrix A . A is assumed to satisfy the monotone envelope property and the hypothesis that $f_i(A) < i$ for $i > 1$. Therefore, by Lemma 3.1, the results apply directly to the matrix A_c .

The factorization process can be defined by the following equations for $i = 1, 2, \dots, N$.

$$a'_{ij} = a_{ij} - \sum_{k=\max(f_i, f_j)}^{j-1} a'_{ik} l_{jk}, \quad j = f_i, f_i+1, \dots, i-1 \quad (3.1a)$$

$$l_{ij} = a'_{ij} / d_{jj}, \quad j = f_i, f_i+1, \dots, i-1 \quad (3.1b)$$

$$d_{ii} = a_{ii} - \sum_{k=f_i}^{i-1} a'_{ik} l_{ik}. \quad (3.1c)$$

It should be clear that (3.1) is simply a variant form of the conventional defining equations for the Choleski decomposition for symmetric matrices [7]. Here in (3.1), zeros to the left of the first nonzero in each row of A are exploited. Since A is assumed to be a monotone envelope matrix, (3.1a) can be rewritten in a simpler form as:

$$a'_{ij} = a_{ij} - \sum_{k=f_i}^{j-1} a'_{ik} l_{jk}, \quad j = f_i, f_i+1, \dots, i-1. \quad (3.1a')$$

Recalling that $\beta_i = i - f_i$, we have the following theorem.

Theorem 3.4: Let A be a symmetric monotone envelope matrix satisfying $f_i(A) < i$ for $i > 1$. Then the symmetric factorization of A requires $\theta(A) = \frac{1}{2} \sum_{i=2}^N \beta_i (\beta_i + 3)$ multiplications.

Proof: From (3.1a') and Theorem 3.2 we observe that the a'_{ij} are all nonzero at positions within the envelope of A. Thus the products $a'_{ik} l_{jk}$ and $a'_{ik} l_{ik}$ appearing in (3.1a') and (3.1c), respectively, involve nonzero operands. In performing the i -th step defined by (3.1a'), (3.1b), and (3.1c), it is clear that:

a) for $j = f_i, f_i+1, \dots, i-1$, the computation of a'_{ij} requires $j - f_i$

multiplications. Thus the total number of multiplications for (3.1a') is:

$$\sum_{j=f_i}^{i-1} j - f_i = \frac{1}{2} (i - f_i)(i - f_i - 1);$$

- b) to compute ℓ_{ij} for $j = f_i, f_i + 1, \dots, i - 1$ requires $i - f_i$ multiplications;
 c) to compute d_{ii} requires $i - f_i$ multiplications.

Hence for the i -th step of the factorization, the total number of multiplications required is

$$\frac{1}{2} (i - f_i)(i - f_i + 3) = \frac{1}{2} \beta_i (\beta_i + 3).$$

Summing on i , then, the entire factorization process requires

$$\theta(A) = \frac{1}{2} \sum_{i=2}^N \beta_i (\beta_i + 3)$$

multiplications, since $\beta_1 = \emptyset$.

It should be noted that for a general sparse matrix, the operation count derived in Theorem 3.4 is only an upper bound for its sparse symmetric factorization. Only for irreducible monotone envelope matrices whose triangular factors have a full envelope is the bound achieved exactly. As an immediate corollary of this observation, we have the well-known result for band matrices [8].

Corollary 3.5: The number of multiplications required for the symmetric factorization of a band oriented matrix with bandwidth m is bounded by $\frac{1}{2} m(m + 3)N$.

Proof: For all i , $\beta_i \leq m$.

4. The Reverse Cuthill-McKee Ordering

In his study of envelope methods, George [3] considered the Reverse Cuthill-McKee algorithm (RCM) which renumbers the CM ordering in the reverse way. Surprisingly, this simple modification often yields an ordering superior to the original ordering in terms of efficiency, although the bandwidth remains unchanged. Experimental evidence of the superior performance of the reversed ordering for matrices arising from the finite element method has been reported in the thesis of George [3] and in the survey paper by Cuthill [1]. In this section we shall prove that the reverse scheme is always at least as good, as far as storage and operation counts are concerned.

Much of what we will develop in the remainder of this paper may be most easily seen by examining several examples. Therefore, we now introduce three graphs, and their associated systems of equations (1.1), which will be used to illustrate our results. The simplest of these is the star graph with N vertices. For any ordering of the vertices of such a graph, we obtain an N by N system (1.1) in which $A_{ij} \neq \emptyset$ if and only if the nodes x_i and x_j are connected in the graph. Figure 4.1 illustrates a star graph with seven nodes ordered by both the CM and RCM algorithms.

Our other two examples arise in the use of finite difference or finite element methods for the solution of partial differential equations. Consider the mesh graph obtained by subdividing the unit square into $N = n^2$ small square elements of side $\frac{1}{n}$ (see Figure 4.2). We may derive an N by N system (1.1) from any ordering of the mesh points of the n by n square mesh by using one of several techniques. A five-point finite difference system is obtained by

considering each mesh point to be adjacent to those vertices to which it is joined by a line in the mesh. Then the matrix A associated with such a system satisfies the condition that $A_{ij} \neq 0$ if and only if the nodes x_i and x_j are connected by a line in the mesh.

Alternatively, we may derive a nine-point finite element system by considering any mesh point x_i to be adjacent to all those vertices which border on small square elements containing x_i . The matrix A associated with the system will then be such that $A_{ij} \neq 0$ if and only if the nodes x_i and x_j are both in the same small square element of the mesh. Figures 4.2 and 4.3 respectively illustrate the CM and RCM orderings of the five- and nine-point square meshes.

As before, we denote by A_c the matrix A ordered by the Cuthill-McKee algorithm. Let A_r be the matrix obtained by reversing the ordering. It is easy to see that

$$A_r = A_c^\# = \tilde{P}A_c\tilde{P}^T \quad (4.1)$$

where

$$\tilde{P} = \begin{bmatrix} & & & & 1 \\ & & & 1 & \\ & & \dots & & \\ & 1 & & & \\ 1 & & & & \end{bmatrix}$$

The following lemma is immediate.

Lemma 4.1: $\text{Env}(A_r) = \text{Tenv}(A_c)$.

It is helpful to consider the labelled symmetric graphs

$G(A_c) = (X(A_c), E(A_c))$ and $G(A_r) = (X(A_r), E(A_r))$ associated with the matrices A_c and A_r respectively. Let $X(A_c) = \{x_1, x_2, \dots, x_N\}$ and $X(A_r) = \{y_1, y_2, \dots, y_N\}$, where x_i is the i -th node in the Cuthill-McKee ordering and y_j is the j -th node in the reverse ordering. As may be seen from Figures 4.1-4.3, $G(A_c)$ and $G(A_r)$ are identical structurally, and x_i and y_{N-i+1} represent the same node in the underlying unlabelled graph structure. We now have the following theorem and its corollaries.

Theorem 4.2: Let A_c be the matrix A ordered by the CM algorithm. Then

$$\text{Tenv}(A_c) \subseteq \text{Env}(A_c).$$

Proof: Assume that for some $i \geq j$, $(A_c)_{ij} \in \text{Tenv}(A_c)$ and $(A_c)_{ij} \notin \text{Env}(A_c)$.

$$(A_c)_{ij} \in \text{Tenv}(A_c) \text{ implies that } \exists k \geq i \text{ such that } x_k \in \text{Adj}(x_j), \quad (4.2)$$

and

$$(A_c)_{ij} \notin \text{Env}(A_c) \text{ implies that for all } \ell \leq j, x_i \notin \text{Adj}(x_\ell). \quad (4.3)$$

But (4.2) and (4.3) together imply that $f_k(A_c) < f_i(A_c)$, contradicting Lemma 3.1. This proves the theorem.

Corollary 4.3: For all i , $\omega_{N-i+1}(A_r) \leq \beta_i(A_c)$.

Proof: If for some i , $\omega_{N-i+1}(A_r) > \beta_i(A_c)$, then there would be some element $(A_c)_{ij}$ such that $(A_c)_{ij} \in \text{Tenv}(A_c) = \text{Env}(A_r)$ and $(A_c)_{ij} \notin \text{Env}(A_c)$ both held. This would contradict Theorem 4.2.

Corollary 4.4: $|\text{Env}(A_r)| \leq |\text{Env}(A_c)|$.

Proof: Theorem 4.2 and Lemma 4.1 together imply that

$$\text{Env}(A_r) \subseteq \text{Env}(A_c).$$

The corollary follows immediately from this.

Corollary 4.5: $|\text{Fill}(A_r)| \leq |\text{Fill}(A_c)|$.

Proof: Let $n(A)$ be the number of nonzeros in the lower triangle of the matrix A . By Lemma 3.1 and Theorem 3.2, the symmetric factorization of A_c into the product LDL^T yields an L having a full envelope, so that

$$|\text{Fill}(A_c)| = |\text{Env}(A_c)| - n(A).$$

On the other hand, we have that

$$|\text{Fill}(A_r)| \leq |\text{Env}(A_r)| - n(A).$$

With Corollary 4.4, this proves the result.

Corollary 4.6: $\theta(A_r) \leq \theta(A_c)$.

Proof: On combining Lemmas 2.1 and 3.1, Theorem 3.4, and Corollary 4.3,

we have:

$$\begin{aligned} \theta(A_r) &= \frac{1}{2} \sum_{k=1}^N \omega_k(A_r) [\omega_k(A_r) + 3] \\ &\leq \frac{1}{2} \sum_{k=1}^N \beta_{N-k+1}(A_c) [\beta_{N-k+1}(A_c) + 3] \\ &= \frac{1}{2} \sum_{i=1}^N \beta_i(A_c) [\beta_i(A_c) + 3] \\ &= \theta(A_c). \end{aligned}$$

5. Comparisons of CM and RCM

In the previous section we saw that the RCM ordering was always at least as good as the CM ordering for solving (1.1). In this section we give a theorem concerning the conditions under which the RCM ordering is actually strictly better than the CM ordering. We also present several examples and numerical experiments which illustrate the results of this paper. Finally, we give an analysis of certain special cases of the finite element method in order to demonstrate the savings which are actually achievable in practical problems.

For a graph $G(A_c)$ associated with the CM ordering, we will say that $G(A_c)$ has Property \tilde{P} if there exist vertices x_i , x_j , and x_k with $i < j < k$ such that $x_k \in \text{Adj}(x_i)$ and for all $l \geq k$, $x_l \notin \text{Adj}(x_j)$. Then we have the following theorem.

Theorem 5.1: $|\text{Env}(A_r)| < |\text{Env}(A_c)|$ if and only if $G(A_c)$ has Property \tilde{P} .

Proof: $G(A_c)$ has Property \tilde{P} if and only if there exist $i < j < k$ such that $(A_c)_{ki} \in \text{Env}(A_c)$, $(A_c)_{kj} \in \text{Env}(A_c)$, and $(A_c)_{kj} \notin \text{Tenv}(A_c)$. But this is equivalent to $|\text{Tenv}(A_c)| < |\text{Env}(A_c)|$, and the result follows by Lemma 4.1.

By examining Figure 4.2, we see that the five-point finite difference mesh cannot satisfy Property \tilde{P} , so that we have the following corollary.

Corollary 5.2: For a rectangular five-point finite difference mesh,

$$|\text{Env}(A_r)| = |\text{Env}(A_c)|.$$

In contrast, we can let $i = 14$, $j = 21$, and $k = 23$ in the CM ordering of the nine-point finite element mesh shown in Figure 4.3. Then we see that the mesh graph has Property \tilde{P} , so that we have:

Corollary 5.3: For a rectangular nine-point finite element mesh,

$$|\text{Env}(A_r)| < |\text{Env}(A_c)|.$$

For both meshes, similar results hold for the storage requirements and operation counts for the solution of the associated system of equations.

Example 1

For our first example we use the unlabelled star graph with N nodes. The best possible CM ordering is to start with one of the "planet" vertices, and the resulting CM and RCM orderings for $N = 7$ are shown in Figure 4.1. It is clear that for the star graph with N nodes,

$$\frac{|\text{Env}(A_r)|}{|\text{Env}(A_c)|} = \frac{2N - 1}{\frac{N(N-1)}{2} + 2} \approx \frac{4}{N}$$

and

$$\frac{\theta(A_r)}{\theta(A_c)} \approx \frac{3}{N}$$

This example shows that the RCM ordering can be significantly better than the CM ordering.

Example 2

We now analyze some examples that arise in the application of the finite element method. As before, we consider the mesh obtained by subdividing a unit square into n^2 small square elements of side $\frac{1}{n}$, and the nine-point finite element system associated with it. We number the nodes starting at the lower left hand corner using the CM algorithm. The case for $n = 4$ is given in Figure 4.3.

By inductive analysis for $n \geq 2$:

$$\begin{aligned} \text{(i)} \quad |\text{Env}(A_c)| &= 1 + \sum_{k=2}^n (4k^2 - 3k - 1) \\ &= \frac{4}{3} n^3 + \frac{1}{2} n^2 - \frac{11}{6} n + 1 \end{aligned}$$

$$(ii) \frac{|\text{Env}(A_r)|}{|\text{Env}(A_c)|} \approx 1$$

$$(iii) \theta(A_c) = 1 + \sum_{k=2}^n (4k^3 - 2k^2 - 5k + 2) \\ = n^4 + \frac{4}{3}n^3 - \frac{5}{2}n^2 - \frac{5}{6}n + 2$$

$$(iv) \frac{\theta(A_r)}{\theta(A_c)} \approx 1$$

The savings for different values of n are given in Table 5.1. It is of interest to note that for this example, the relative savings with RCM decrease with increasing n and that the well known natural, or row-by-row mesh ordering is better than either CM or RCM.

Example 3

The savings in Example 2 are not particularly impressive. We now consider the same model problem with triangular elements (that is, each small square in Figure 4.3 is subdivided into two right triangles), and its corresponding mesh system. The basic mesh is shown in Figure 5.1.

With quadratic elements, each triangle has three vertex unknowns and three edge unknowns. The corresponding CM ordering is shown in Figure 5.2 for $n = 2$. The results given in Table 5.2 show that substantial savings in computation and storage can be achieved if the RCM rather than the CM ordering is used.

In the case of cubic elements having ten nodal unknowns, the savings are even more dramatic. The CM ordering for $n = 2$ is shown in Figure 5.3. The storage and work requirements for CM and RCM are given in Table 5.3.

Finally, we consider elements with three vertex unknowns and a variable number p_\emptyset of interior unknowns. The case for $p_\emptyset = 1$ is shown in Figure 5.4. The results of experiments for $p_\emptyset = \emptyset$ and $p_\emptyset = 1$ are shown in Tables 5.4 and 5.5.

The examples in this section have shown that dramatic savings may be achieved by using RCM rather than CM. We will now give a more detailed analysis of the last example.

The basic finite element mesh to be analyzed is shown in Figure 5.1, and as above, we consider generalizations of the basic mesh with exactly p_\emptyset interior nodes in each triangular element. These meshes may be seen to be simplifications of the first two meshes presented in Example 3, so that the results which are obtained now will be closely related to the experimental results given there. In fact, our investigations have led to the conclusion that the difference between the CM and RCM orderings is only increased by the addition of either interior or edge nodes to the regular right triangular mesh of Figure 5.1.

As mentioned in Section 3, the CM algorithm generates a breadth-first spanning tree for the mesh graph. A sample spanning tree for the mesh in Figure 5.4 is shown in Figure 5.5. The analysis below will be based upon the characteristics of the generated spanning tree and the relationship between CM and RCM which is formalized in the following lemma.

Lemma 5.4: Consider a spanning tree generated by the CM algorithm, numbered top-to-bottom, left-to-right (see Figure 5.5). For any node x_i on level l of the tree:

$$\beta_i(A_c) = i - f_i(A_c) = (\text{the number of nodes to the left of } x_i \text{ on level } \ell) + (\text{the number of nodes to the right of the father of } x_i \text{ on level } \ell-1) + 1.$$

If there are N nodes in the tree, x_i will be relabelled x_{N-i+1} in the RCM ordering. Then, if x_i has a son on level $\ell+1$:

$$\beta_{N-i+1}(A_r) = N - i + 1 - f_{N-i+1}(A_r) = (\text{the number of nodes to the right of } x_i \text{ on level } \ell) + (\text{the number of nodes to the left of the rightmost son of } x_i \text{ on level } \ell+1) + 1.$$

And if x_i has no sons in the tree:

$$\beta_{N-i+1}(A_r) = N - i + 1 - f_{N-i+1}(A_r) = (\text{the number of nodes strictly between } x_i \text{ and the rightmost node on level } \ell \text{ or } \ell+1 \text{ which is a member of } \text{Adj}(x_i)) + 1.$$

Proof: The proof of this lemma is quite simple. From Figures 5.5 and 5.6, it can easily be seen that for the CM ordering, $x_{f_i(A_c)}$ is the father of x_i in the spanning tree. Since there must be exactly $i - f_i(A_c) - 1$ nodes between $x_{f_i(A_c)}$ and x_i on levels $\ell-1$ and ℓ , the result is immediate for CM. For the RCM ordering, $x_{f_{N-i+1}(A_r)}$ is the rightmost son of x_i , if one exists, so there are $N - i + 1 - f_{N-i+1}(A_r) - 1$ nodes between them on levels ℓ and $\ell+1$. And if x_i has no sons, then $f_{N-i+1}(A_r)$ is the index in the RCM ordering of that node in $\text{Adj}(x_i)$ which is labelled last by CM. Since the CM spanning tree is labelled top-to-bottom and left-to-right, this node must be the rightmost member of $\text{Adj}(x_i)$ on either level ℓ or level $\ell+1$.

Using the simple counting argument of Lemma 5.4, we can obtain the following results, which are stated without proof.

Theorem 5.5: For a regular right triangular mesh with $(n+1)^2$ vertex nodes and p_\emptyset interior nodes in each triangular element:

$$(i) \quad |\text{Env}(A_c)| = \frac{2}{3} (2p_\emptyset + 1)^2 n^3 + O(n^2);$$

$$(ii) \quad |\text{Env}(A_r)| = \frac{2}{3} (2p_\emptyset + 1) n^3 + O(n^2);$$

$$(iii) \quad \frac{|\text{Env}(A_r)|}{|\text{Env}(A_c)|} \approx \frac{1}{2p_\emptyset + 1}.$$

From these results we see that the use of RCM leads to very impressive storage reductions for even small values of p_\emptyset . For $p_\emptyset = 1$, as in Example 3, the reduction would be nearly 70% for large values of n . Even larger savings may be expected in the operation counts for factorization. Analysis suggests that the additional savings due to the placement of edge nodes in a triangular finite element mesh are not so dramatic as those for interior nodes, and the results given in Tables 5.3 and 5.5 bear this out.

6. Conclusion

In [3] George advocates the use of envelope methods rather than band methods, and this may lead to substantial savings in many cases of practical interest. In the implementation of envelope methods, the use of Jennings' envelope storage scheme is quite attractive. The scheme stores the rows of the envelope of the lower triangle of the N by N coefficient matrix A in a linear array. An additional N address pointers are used to locate the positions of the diagonal elements in the main storage array. From Lemma 5.4, we see that these pointers may be obtained as a direct by-product of the CM ordering process, for to compute them we need only the value of f_i for each i . In order to get the pointers for the RCM ordering, we merely note that for each i , $f_{N-i+1}(A_r) = \min(N-i+1, N-k+1)$, where k is the largest label assigned by CM to a member of $\text{Adj}(x_i)$. For non-interior nodes, x_k will usually be the rightmost son of x_i in the CM spanning tree.

In terms of both storage and computation, we have seen that the RCM ordering algorithm is superior to the original CM scheme for envelope methods. Intuitively, we can attribute this to the fact that CM attempts to minimize the ordering distance between a graph node and its unordered neighbors, while RCM attempts to minimize the distance between a vertex and its ordered neighbors. Hence it appears that CM tends to minimize $|\text{Tenv}(A)|$, while RCM tends to minimize $|\text{Env}(A)|$. For systems to be solved with envelope methods, then, RCM is a better ordering to use than CM. However, we should note that for some finite element structures (e.g. the nine-point finite element grid) neither RCM nor CM is the best ordering to use.

The savings found in Example 3 and the following analysis are dramatic. This demonstrates that in practical applications where the CM ordering scheme is popular (such as structural analysis), it is possible to substantially reduce the amount of storage and the number of multiplicative operations required for the direct solution of systems (1.1). And although all of our examples were based on symmetric systems, it is clear that envelope methods with the RCM ordering may be applied to systems (1.1) in which only the zero structure of A is symmetric. Thus our results may have an important effect on the solution of linear systems of equations which arise in many physical problems to which envelope methods are not currently applied.

Acknowledgments

The authors would like to thank Professor J.A. George of the University of Waterloo and Professors M.H. Schultz and S.C. Eisenstat of Yale University for their helpful discussions and advice.

References


- [1] E. Cuthill, Several Strategies for Reducing the Bandwidth of Matrices, Sparse Matrices and their Applications, D.J. Rose and R.A. Willoughby, eds., Plenum Press, New York, 1972, pp. 157-166.
- [2] E. Cuthill and J. McKee, Reducing the Bandwidth of Sparse Symmetric Matrices, Proc. 24-th Nat. Conf. of the ACM, ACM Publication P-69, 1122 Ave. of the Americas, New York, 1969, pp. 157-172.
- [3] J.A. George, Computer Implementation of the Finite Element Method, Stanford Computer Science Dept., Technical Report STAN-CS-71-208, Stanford, California, 1971.
- [4] J.A. George, A Survey of Sparse Matrix Methods in the Direct Solution of Finite Element Equations, Proc. Summer Computer Simulation Conf., Montreal, Canada, July 17-19, 1973, pp. 15-20.
- [5] J.A. George and W.H. Liu, Some Results on Fill for Sparse Matrices, submitted to SIAM J. Numer. Anal.
- [6] A. Jennings, A Compact Storage Scheme for the Solution of Symmetric Simultaneous Equations, Comput. J. 9(1966), pp. 281-285.
- [7] R.S. Martin, G. Peters, and J.H. Wilkinson, Symmetric Decomposition of a Positive Definite Matrix, Numer. Math. 7(1965), pp. 362-383.
- [8] J.H. Wilkinson, The Algebraic Eigenvalue Problem, Clarendon Press, London, 1965.

$$A = \begin{bmatrix} X & 0 & X & 0 & X & 0 & 0 \\ 0 & X & 0 & X & 0 & X & 0 \\ X & 0 & X & 0 & X & 0 & 0 \\ 0 & X & 0 & X & X & 0 & X \\ X & 0 & X & X & X & 0 & 0 \\ 0 & X & 0 & 0 & 0 & X & 0 \\ 0 & 0 & 0 & X & 0 & 0 & X \end{bmatrix}$$

\underline{i}	$\underline{f}_i(A)$	$\underline{\beta}_i(A)$	$\underline{\omega}_i(A)$
1	1	0	2
2	2	0	4
3	1	2	3
4	2	2	3
5	1	4	2
6	2	4	1
7	4	3	0


$$b(A) = 4 \quad |B(A)| = 25 \quad |Env(A)| = 22 \quad |Tenv(A)| = 20$$

Figure 2.1



CM Ordering


$$A_C = \begin{bmatrix} X & X & & & & & \\ X & X & X & X & X & X & X \\ & X & X & & & & \\ X & & X & & & & \\ X & & & X & & & \\ X & & & & X & & \\ X & & & & & & X \end{bmatrix}$$




RCM Ordering

$$A_R = \begin{bmatrix} X & & & & & & X \\ & X & & & & & X \\ & & X & & & & X \\ & & & X & & & X \\ & & & & X & X & \\ X & X & X & X & X & X & X \\ & & & & & & X & X \end{bmatrix}$$

Figure 4.1: Orderings for Star Graph with N = 7 Nodes




CM




RCM

Figure 4.2: Orderings for a Five-point Finite Difference Mesh



CM



RCM

Figure 4.3: Orderings for a Nine-point Finite Element Mesh




Figure 5.1: 3 by 3 Regular Right Triangular Mesh
Ordered by the CM Algorithm




Figure 5.2: 2 by 2 Regular Right Triangular Mesh with
Element 2-6¹ Ordered by the CM Algorithm

¹We adopt the notation in [3], where the two parts of the hyphenated name refer respectively to the degree of the interpolating polynomial and to the number of nodes associated with the element.




Figure 5.3: 2 by 2 Regular Right Triangular Mesh with Element
3-10 Ordered by CM Algorithm




Figure 5.4: 3 by 3 Regular Right Triangular Mesh with
 $p_0 = 1$ Interior Nodes, Ordered by CM




Figure 5.5: Spanning Tree Generated by the CM Algorithm for a 3 by 3 Right Triangular Mesh with $p_0 = 1$.




Figure 5.6: Spanning Tree Associated with the CM Ordering for the 2 by 2 Regular Square Mesh

Table 5.1

Theoretical Amount Saved for the Regular Square Mesh

Problem			Storage			Operation Count		
n	$N = (n+1)^2$	Nonzeros	CM	RCM	RCM/CM	CM	RCM	RCM/CM
4	25	97	171	147	.860	726	530	.730
8	81	353	997	885	.888	7324	5812	.793
16	269	1345	6665	6185	.928	89336	77736	.870
32	1089	4749	48401	46417	.959	1231088	1140816	.927

Table 5.2

Experimental Savings for the Regular Right

Triangular Mesh with Element 2-6

Problem			Storage			Operation Count		
n	$N = (2n+1)^2$	Nonzeros	CM	RCM	RCM/CM	CM	RCM	RCM/CM
4	81	360	1078	755	.700	8758	4183	.478
5	121	555	1971	1310	.665	19214	8324	.433
6	169	792	3244	2077	.640	36808	14857	.404
7	225	1071	4961	3088	.622	64040	24506	.383
8	289	1392	7186	4375	.609	103914	38115	.367
9	361	1755	9983	5970	.598	159698	56600	.354

Table 5.3

Experimental Savings for the Regular Right
 Triangular Mesh with Element 3-10

Problem			Storage			Operation Count		
n	$N = (3n+1)^2$	Nonzeros	CM	RCM	RCM/CM	CM	RCM	RCM/CM
3	100	684	2002	1252	.625	25002	9429	.377
4	168	1200	4558	2518	.552	74528	22046	.296
5	256	1860	8626	4396	.510	172453	43624	.253
6	361	2664	14530	6994	.481	342564	77574	.226

Table 5.4

Experimental Results for the Regular Right
 Triangular Mesh with Element 1-3

Problem			Storage			Operation Count		
n	$N = (n+1)^2$	Nonzeros	CM	RCM	RCM/CM	CM	RCM	RCM/CM
4	25	81	115	115	1.00	320	320	1.00
8	81	289	597	597	1.00	2616	2616	1.00
16	289	1089	3689	3689	1.00	27472	27472	1.00
32	1089	4225	25553	25553	1.00	344608	344608	1.00

Table 5.5

Experimental Results for the Regular Right

Triangular Mesh with $p_0 = 1$

Problem			Storage			Operation Count		
n	$N = (n+1)^2 + 2n^2$	Nonzeros	CM	RCM	RCM/CM	CM	RCM	RCM/CM
4	57	177	529	323	.611	2975	1088	.366
8	209	673	3687	1781	.483	38037	8808	.232
16	801	2625	27139	11177	.412	527081	89200	.169
32	3137	10369	207099	77393	.374	7761201	1083232	.140