GENERALIZED NESTED DISSECTION

by

R. J. Lipton, D. J. Rose and R. E. Tarjan

STAN-CS-77-645
DECEMBER 1977

COMPUTER SCIENCE DEPARTMENT
School of Humanities and Sciences
STANFORD UNIVERSITY

Generalized Nested Dissection

Richard J. Lipton */
Donald J. Rose **/
Robert Endre Tarjan ***/

August, 1977

Abstract.

J. A. George has discovered a method, called nested dissection, for solving a system of linear equations defined on an $n = k \times k$ square grid in $O(n \log n)$ space and $O(n^{3/2})$ time. We generalize this method without degrading the time and space bounds so that it applies to any system-of equations defined on a planar or almost-planar graph. Such systems arise in the solution of two-dimensional finite element problems. Our method uses the fact that planar graphs have good separators.

More generally, we show that sparse Gaussian elimination is efficient for any class of graphs which have good separators, and conversely that graphs without good separators (including almost all sparse graphs) are not amenable to sparse Gaussian elimination.

Keywords: finite element method, Gaussian elimination, nested dissection, planar graph, sparse matrix.

*/ Computer Science Dept., Yale University, New Haven, Connecticut 06520.
Research partially supported by the U.S. Army Research Office, grant no. DAAG 29-76-G-0338.

**/ Computer Science Dept., Vanderbilt University, Nashville, Tennessee 37235.
Research partially supported by Office of Naval Research contract N00014-67-0298-0034.

***/ Computer Science Dept., Stanford University, Stanford, California 94305.
Research partially supported by National Science Foundation grant MCS-75-22870 and Office of Naval Research contract N00014-76-C-0688.

1. Introduction.

Suppose we wish to solve by Gaussian elimination the system of linear equations

(1) $Ax = b$

where A is an $n \times n$ symmetric positive definite matrix, x is an $n \times 1$ vector of variables, and b is an $n \times 1$ vector of constants. The solution process consists of two steps. First, we factor A by means of row operations into

(2) $A = LDL^T$

where L is lower triangular and D is diagonal. Next, we solve the simplified systems $Lz = b$, $Dy = z$, and $L^T x = y$.

If A is dense (i.e., A contains mostly non-zero elements) then the time required for factoring A is $O(n^3)$ and the time required for solving the simplified systems is $O(n^2)$. If A is sparse (i.e., A contains mostly zero elements), we may be able to save time and storage space by avoiding explicit manipulation of zeros. One difficulty with obtaining such a savings is that the factoring process may create non-zeros in L (and L^T) in positions where A contains zeros. These new non-zeros are called fill-in.

One way to reduce the fill-in is to permute the rows and columns of A , i.e., to transform A into

(3) $A' = PAP^T$

where P is a permutation matrix, and to solve the reordered system. Since A is positive definite, the reordered system is numerically stable with respect to the LDL^T factorization [9].

In order to characterize the fill-in associated with a given permutation matrix P , we represent the class of matrices PAP^T by an undirected graph*/ $G = (V,E)$. The graph G contains one vertex $i \in V$ for each row (and column) in A , and one- edge $\{i,j\} \in E$ for each pair of non-zero, off-diagonal elements $a_{ij} = a_{ji} \neq 0$ in A . Each permutation matrix P corresponds to a numbering of the vertices of G ; i.e., to a one-to-one mapping $\pi: V \rightarrow \{1,2,\ldots,n\}$. Corresponding to the factorization $PAP^T = LDL^T$ is a graph $G_\pi^* = (V,E_\pi^*)$ such that $\{i,j\} \in E_\pi^*$ iff $i > j$ and the element of L in row $\pi(i)$ and column $\pi(j)$ is non-zero. See [17, 18,19, 23] for a discussion of the properties of this graph-theoretic model of sparse Gaussian elimination. The following lemma characterizes the fill-in E_π^* associated with an ordering π .

Lemma 1 [19]. Assuming no cancellation of non-zeros in the factoring of PAP^T, $\{v, w\} \in E_\pi^*$ iff $v \neq w$ and there is a path $v = v_1, v_2, \ldots, v_{k+1} = w$ such that $\pi(v_i) < \min\{\pi(v), \pi(w)\}$ for $2 \leq i \leq k$.

The running time and storage space required by sparse Gaussian elimination are functions of m , the number of non-zeros in L , and of d(i) , the number of edges $\{i,j\}$ in G_π^* with $n(i) < n(j)$. Note that d(i) is the number of non-zeros in column i of L (and row i of L^T), and that $m = \sum_{i=1}^{n-1} d(i)$. For purposes of analysis and implementation, we can divide sparse elimination into the following four steps.

*/ The appendix contains the graph-theoretic definitions used in this paper. It also defines the "O" and "Ω" notations.

Step 1. Find a good ordering π .

The time and space required by this step depend upon the method used.

Step 2. (Symbolic factorization.) Compute the non-zero positions in L, assuming no lucky cancellation of non-zeros.

Time: O(m) using the algorithm of [19].

Space : O(m) .

Step 3. (Numeric factorization.) Compute L .

Time: $O\left(\sum_{i=1}^{n-1} d(i)(d(i)+3)\right)$ using an algorithm such as described in [6,12,20,23]. The number of multiplications performed during this step is $\frac{1}{2}\sum_{i=1}^{n-1} d(i)(d(i)+3)$ [18].

Space: O(m) .

Step 4. (Backsolving.) Solve $Lz = b$, $Dy = z$ and $L^T x = y$.

Time: O(m) [18] .

Space : O(m) .

The reason for separating the factorization into two steps (symbolic and numeric) is that all known methods which compute the numeric factorization without first finding the fill-in positions have a time bound for overhead which is more than a constant factor greater than the number of multiplications. If the system of equations is to be solved for only one right-hand side b , it is possible to combine at least part of Step 4 (solving $Lz = b$ and $Dy = z$) with Step 3.

The efficiency of sparse Gaussian elimination depends upon Step 1; that is, upon finding an ordering π which reduces the size of the fill-in m and the multiplication count $\frac{1}{2}\sum_{i=1}^{n} d(i)(d(i)+3)$. Finding such a good ordering for an arbitrary graph seems to be a very hard, perhaps even NP-complete problem. However, for some special cases good ordering schemes are known. One such scheme is the nested dissection method of J. A. George [11], which allows the solution of systems whose graph is an $n = k \times k$ square grid graph in $O(n^{3/2})$ time and $O(n \log n)$ space. George's scheme uses the fact that removal of $O(k)$ vertices from a $k \times k$ square grid leaves four square grids, each roughly $k/2 \times k/2$.

In this paper we generalize George's idea. Let S be a class of graphs closed under the subgraph relation (i.e., if $G_2 \in S$ and G_1 is a subgraph of G_2 then $G_1 \in S$). The class S satisfies an <u>f(n) -separator theorem</u> if there are constants $\alpha < 1$, $\beta > 0$ for which any n-vertex graph G in S has the following property: the vertices of G can be partitioned into three sets A, B, C such that no vertex in A is adjacent to any vertex in B , $|A|, |B| \leq \alpha n$, and $|C| \leq \beta f(n)$. Cur main result is that all systems of equations whose graphs satisfy a $\sqrt{n}$ -separator theorem can be solved in $O(n^{3/2})$ time and $O(n \log n)$ space using a "divide and conquer" [1] method to generate the ordering. From separator theorems proved in [15], we obtain a method for solving any system of equations whose graph is planar or almost-planar in $O(n^{3/2})$ time and $O(n \log n)$ space. Such systems arise in the solution of two-dimensional finite element problems [24]. Section 2 presents these results.

More generally, divide and conquer gives a good ordering scheme for any class of graphs satisfying an f(n) -separator theorem; the fill-in and multiplication count produced by the ordering depend upon $f(n)$.

At the end of Section 2 we list fill-in and multiplication bounds for various values of f(n) other than $f(n) = \sqrt{n}$.

Section 3 presents some relationships between Gaussian elimination, good separators, sparsity, and random graphs. We give a lower bound on the cost of Gaussian elimination in terms of the size of separators in the problem graph. We prove that graphs with good separators are sparse. Finally, we show that almost all sparse graphs have no good ordering for Gaussian elimination. Section 4 discusses the significance of the results in Sections 2 and 3.

2. Generalized Nested Dissection.

Let S be a class of graphs closed under subgraph on which a $\sqrt{n}$-separator theorem holds, let α, β be the constants associated with the separator theorem, and let $G = (V,E)$ be an n-vertex graph in S. The following recursive algorithm numbers the vertices of G so that sparse Gaussian elimination is efficient. The algorithm assumes that ℓ of the vertices of G are already assigned numbers, each of which is greater than b, and that the remaining vertices of G are to be numbered consecutively from a to b.

Numbering Algorithm.

If G contains no more than $n_0 = (\beta/(1-\alpha)^2)$ vertices, number the unnumbered vertices arbitrarily from a to b. Otherwise, find sets A, B, C satisfying the $\sqrt{n}$-separator theorem. Let A contain i unnumbered vertices, B contain j unnumbered vertices, and C contain k unnumbered vertices.

Number the unnumbered vertices in C arbitrarily from $b-k+1$ to b. Delete all edges with both endpoints in C. Apply the algorithm recursively to the subgraph induced by $B \cup C$ to number the unnumbered vertices in B from $b-k-j+1$ to $b-k$. Apply the algorithm recursively to the subgraph induced by $A \cup C$ to number the unnumbered vertices in A from $a = b-k-j-i+1$ to $a+i-1 = b-k-j$.

If G initially has no numbered vertices, then applying this algorithm to G with $a = 1$, $b = n$, and $\ell = 0$ will number the vertices of G from 1 to n. We are interested in three properties of this algorithm: its running time, the size of the fill-in produced by the ordering it generates, and the multiplication count of the generated ordering.

Theorem 1. Suppose that a vertex partition satisfying the $\sqrt{n}$ -separator theorem can be found in $O(m+n)$ time on an n-vertex, m-edge graph. Then the numbering algorithm requires $O((m+n)\log n)$ time.

Proof. Let $t(m,n)$ be the maximum time required by the numbering algorithm on any graph in S with n vertices and m edges. Then

(4) $$t(m,n) \le c_1 \text{ */} \quad \text{if } n \le n_0 ,$$

$$t(m,n) \le c_2(m+n) + \max\{t(m_1,n_1) + t(m_2,n_2)\} \quad \text{otherwise,}$$

where $n_0 = (\beta/(1-\alpha))^2$ and the maximum is taken over values of m_1 , n_1 , m_2 , n_2 satisfying

(5) $$m_1 + m_2 \le m ,$$

$$n \le n_1 + n_2 < n + \beta\sqrt{n} , \quad \text{and}$$

$$(1-\beta)n \le n_1, n_2 \le \alpha n + \beta\sqrt{n} .$$

A proof by induction similar to the one below for the fill-in bound shows that $t(m,n)$ is $O((m+n)\log n)$. □

Theorem 2. Let G be any n-vertex graph numbered by the algorithm. The total size of the fill-in associated with the numbering is at most $c_3 n \log_2 n + O(n)$, where

(6) $$c_3 = \beta^2(1/2 + 2\sqrt{\alpha}/(1-\sqrt{\alpha}))/\log_2(1/\alpha) .$$

*/ Throughout this paper, $c, c_0, c_1, \ldots$ denote suitable positive constants.

Proof. Suppose the recursive numbering algorithm is applied to an n-vertex graph G with ℓ vertices previously numbered. Assume $n > n_0$ and let A, B, C be the vertex partition generated by the algorithm. If C contains k unnumbered vertices, then the maximum number of fill-in edges whose lower numbered endpoint is in C is

(7) $$k(k-1)/2 + k\ell < \beta^2 n/2 + \beta \ell \sqrt{n} .$$

By Lemma 1, two vertices v and w are joined by a fill-in edge if and only if there is a path from v to w through vertices numbered less than both v and w . Thus no fill-in edge joins a vertex in A with a vertex in B . Let $f(\ell,n)$ be the maximum number of fill-in edges whose lower numbered endpoint is numbered by the algorithm (and not previously numbered). Then

(8) $$f(\ell,n) \leq n(n-1)/2 \qquad \text{if } n \leq n_0 \text{, and}$$

$$f(\ell,n) \leq \beta^2 n/2 + \beta \ell\sqrt{n} + \max\{f(\ell_1,n_1) + f(\ell_2,n_2)\}$$

otherwise, where the maximum is taken over values satisfying

(9) $$\ell_1 + \ell_2 \leq \ell + 2\beta\sqrt{n} ,$$

$$n \leq n_1 + n_2 \leq n + \beta\sqrt{n} , \text{ and}$$

$$(1-\alpha)n \leq n_1, n_2 \leq \alpha n + \beta\sqrt{n} .$$

We claim that for all $n > 1$,

(10) $$f(\ell,n) \leq c_3 n \log_2 n + c_4 \ell\sqrt{n} + c_5 \ell \log_2 n + c_6 n - c_7\sqrt{n} \log_2 n ,$$

where

(11) $c_3 = \beta^2[1/2 + 2\sqrt{\alpha}/(1-\sqrt{\alpha})] / \log_2(1/\alpha)$,

$c_4 = \beta/(1-\sqrt{\alpha})$,

$c_5 = c_2\beta / [2\sqrt{\alpha} \log_2(1/\alpha)]$, and

c_6 and c_7 are suitably large positive constants, to be chosen later.

We have attempted to minimize c_3 , c_4 , and c_5 in this bound, but have chosen c_6 and c_7 to make the proof easy, The theorem's bound on fill-in size follows from the claim.

Proof of claim. Let

(12) $g(\ell,n) = c_3 n \log_2 n + c_4 \ell\sqrt{n} + c_5 \ell \log_2 n$ and

$h(n) = c_6 n - c_7\sqrt{n} \log_2 n$.

We prove the claim by induction on n . Assume $h(n) \geq n_3^2/2$, where $n_3 \geq n_0$ is a value to be chosen later. Then $n \leq 5$ implies $f(\ell,n) < n(n-1)/2 \leq n_3^2/2 \leq h(n) \leq g(\ell,n)+h(n)$.

Let $n > n_3$ and suppose the claim is true for values smaller than n . Then $f(\ell,n) \leq \beta^2 n/2 + \beta \ell\sqrt{n} + f(\ell_1,n_1) + f(\ell_2,n_2)$ for suitable values of ℓ_1 , n_1 , ℓ_2 , n_2 .

Let $\epsilon = (1-\alpha-\beta/\sqrt{n_0+1})$. Since $\sqrt{n_0+1} > \sqrt{n_0} \geq \beta/(1-\alpha)$, we have $\alpha+\beta/\sqrt{n_0+1} < 1$, and $\epsilon > 0$. Thus $n_1,n_2 < \alpha n + \beta\sqrt{n} \leq (\alpha+\beta/\sqrt{n})n < (1-\epsilon)n < n$, and the claim holds for n_1 and n_2 by the induction hypothesis.

Hence

(13) $$f(\ell,n) \le \beta^2 n/2 + \beta \ell\sqrt{n} + g(\ell_1,n_1) + g(\ell_2,n_2) + h(n_1) + h(n_2)$$

$$< \beta^2 n/2 + \beta \ell\sqrt{n}$$

$$+ [c_3(n+\beta\sqrt{n}) + c_5(\ell + 2\beta\sqrt{n})] \log_2(\alpha n + \beta\sqrt{n})$$

$$+ c_4(\ell + 2\beta\sqrt{n}) \sqrt{\alpha n + \beta\sqrt{n}}$$

$$+ h(n_1) + h(n_2) \; .$$

Now

(14) $$\log_2(\alpha n + \beta\sqrt{n}) = \log 2\, n + \log_2 \alpha + \log_2(1 + \beta/(\alpha\sqrt{n}))$$

$$< \log 2\, n + \log 2\, \alpha + (\beta \log_2 e)/(\alpha\sqrt{n})$$

since $\log_e(1+x) < x$ for $x < 0$.

Also

(15) $$\sqrt{\alpha n + \beta\sqrt{n}} \le \sqrt{\alpha n} + \beta/(2\sqrt{\alpha}) \; .$$

Substituting into the bound on $f(\ell,n)$, we find that

(16) $$f(\ell,n) \le \beta^2 n/2 + \beta \ell\sqrt{n}$$

$$+ (c_3 n + c_5 \ell)[\log_2 n + \log_2 \alpha + (\beta \log_2 e)/(\alpha\sqrt{n})] + (c_3 + 2c_5)\beta\sqrt{n} \log 2\, n$$

$$+ c_4(\ell + 2\beta\sqrt{n})[\sqrt{\alpha n} + \beta/(2\sqrt{\alpha})]$$

$$+ h(n_1) + h(n_2)$$

$$\le c_3 n \log_2 n$$

$$+ \beta^2 n/2 + c_3 n \log 2\, \alpha + 2c_4 \beta\sqrt{\alpha}\, n$$

$$+ \beta \ell\sqrt{n} + c_4\sqrt{\alpha}\, \ell\sqrt{n}$$

$$+ c_5 \ell \log_2 n$$

$$+ c_5 \ell \log_2 \alpha + c_4 \beta \, \ell/(2\sqrt{\alpha})$$

$$+ (c_3 \beta \log_2 e)\sqrt{n}/\alpha + (c_5 \beta \log_2 e)\, \ell/(\alpha\sqrt{n})$$

$$+ (c_3 + 2c_5)\beta \sqrt{n} \log_2 n + 2c_4 \beta^2 \sqrt{n}/(2\sqrt{\alpha})$$

$$+ h(n_1) + h(n_2)$$

$$\leq c_3 n \log_2 n + c_4 \ell \sqrt{n} + c_5 \ell \log_2 n + c_8 \sqrt{n} \log_2 n + h(n_1) + h(n_2)$$

by the choice of c_3, c_4, and c_5, where

(17) $$c_8 = (c_3 + c_5)\beta(\log_2 e)/\alpha + (c_3 + 2c_5)\beta + 2c_4 \beta^2/(2\sqrt{\alpha}) .$$

All that remains is to show

(18) $$h(n) \geq n_3^2/2 \quad \text{for } 1 < n, \text{ and}$$

$$h(n) \geq c_8 \sqrt{n} \log_2 n + h(n_1) + h(n_2) \quad \text{for } n > n_3$$

and. any n_1, n_2 satisfying

(19) $$n \leq n_1 + n_2 \leq n + \beta\sqrt{n} \quad \text{and}$$

$$(1-\alpha)n \leq n_1, n_2 \leq \alpha n + \beta\sqrt{n} .$$

Choose n_3 such that

(20) $$n_3 \geq n_0 \quad \text{and}$$

$$28 \log_2 e < (y-1) \log_2 n_3 - \gamma \log_2(1/(1-\alpha)) ,$$

where $y = \sqrt{\alpha} + \sqrt{1-\alpha}$. Let $c_9 = y - 1 - (\gamma \log_2(1/1-\alpha))/\log_2 n_3$.

Choose c_6 such that

(21) $$c_6 \geq \max\{n_3^2/2 , c_8/(c_9/(2 \log_2 e) - \beta/\log_2 n_3)\} .$$

Finally, choose

$$(22) \qquad c_7 = (c_8 \;\square\; c_6\beta/\log_2 n_3)/c_9 \;.$$

Then $h(1) = c_6 \geq n_3^2/2$. Furthermore

$$(23) \qquad \frac{d}{dn}\,\frac{h(n)}{\sqrt{n}} = \frac{d}{dn}(c_6\sqrt{n} - c_7 \log_2 n)$$

$$= c_6/(2\sqrt{n}) - (c_7 \log_2 e)/\log_2 n \;.$$

By the choice of c_6 and c_7 , $c_6/2 \geq c_7 \log_2 e$, which implies $\frac{d}{dn}\,\frac{h(n)}{\sqrt{n}} \geq 0$ for $n > 1$. Hence $h(n) \geq n_3^2/2$ for $n \geq 1$.

We also have

$$(24) \qquad c_8\sqrt{n}\,\log_2 n + h(n_1) + h(n_2)$$

$$\leq c_8\sqrt{n}\,\log_2 n + c_6(n + \beta\sqrt{n}) - c_7(\sqrt{n_1} + \sqrt{n_2})\log_2((1-\alpha)n) \;.$$

For fixed $n_1 + n_2$, the function $\sqrt{n_1} + \sqrt{n_2}$ is minimized when one of n_1 , n_2 is as large as possible and the other is as small as possible. Thus

$$(25) \qquad \sqrt{n_1} + \sqrt{n_2} \geq \sqrt{\alpha n} + \sqrt{(1-\alpha)n} \geq \gamma\sqrt{n} \;.$$

Hence

$$(26) \qquad c_8\sqrt{n}\,\log_2 n + h(n_1) + h(n_2)$$

$$\leq c_8\sqrt{n}\,\log_2 n + c_6 n + c_6\beta\sqrt{n} - c_7\beta\sqrt{n}(\log_2 n - \log_2(1/(1a))$$

$$\geq c_6 n - c_7\gamma\sqrt{n}\,\log_2 n + c_8\sqrt{n}\,\log_2 n + c_6\beta\sqrt{n} + c_7\gamma\sqrt{n}\,\log_2(1/(1-\alpha)) .$$

By selection of c_7 ,

(27) $$c_7 = (c_8 + c_6\beta/\log_2 n_3) / (\gamma - 1 - \gamma \log_2(1/(1-\alpha))/\log_2 n_3)$$

$$\geq (c_8 + c_6\beta/\log_2 n) / (\gamma - 1 - \gamma \log_2(1/(1-\alpha))/\log_2 n)$$

if $n > n_3$. Thus

(28) $$c_7 (\gamma - 1 - \gamma \log_2(1/(1-\alpha))/\log_2 n) \geq c_8 + c_6\beta/\log_2 n \quad \text{and}$$

$$-c_7\sqrt{n} \log_2 n \geq c_8\sqrt{n} \log_2 n + c_6\beta\sqrt{n} - c_7\gamma\sqrt{n} \log_2 n + c_7\gamma \sqrt{n} \log_2 (1/(1-\alpha)) .$$

This means

(29) $$c_8\sqrt{n} \log_2 n + h(n_1) + h(n_2) \leq c_6 n - c_7\sqrt{n} \log_2 n .$$

This completes the proof of the claim. □

Theorem 2. Let G be any n-vertex graph numbered by the algorithm. The total multiplication count associated with the numbering is at most $c_{11}n^{3/2} + O(n(\log n)^2)$, where

(30) $$c_{11} = \beta^2(1/6 + \beta\sqrt{\alpha}(2+\sqrt{\alpha}/(1+\sqrt{\alpha}) + 4\alpha/(1-\alpha)) / (1-\sqrt{\alpha})) / (1-\delta)$$

with $\delta = \alpha^{3/2} + (1-\alpha)^{3/2}$.

Proof, Consider the number of multiplications associated with the ordering. The number of multiplications associated with a given vertex V is $d(v)(d(v)+3)/2$, where $d(v)$ is the number of fill-in edges whose lower-nurribered vertex is v . Thus a bound on the number of multiplications associated with a separator C generated by one call of the recursive numbering algorithm is

$$(31)\qquad \sum_{i=0}^{\beta\sqrt{n}-1} (i+\ell)(i+\ell+3)/2$$

$$< \sum_{i=0}^{\beta\sqrt{n}-1} (i+\ell)^2/2 + 3\beta^2 n/4 + 3\beta\ell\sqrt{n}/2$$

$$\leq \beta^3 n^{3/2}/6 + \beta^2\ell n/2 + \beta\ell^2\sqrt{n}/2 + 3\beta^2 n/4 + 3\beta\ell\sqrt{n}/2 \quad .$$

Let $q(\ell,n)$ be the maximum number of **multiplications** associated with vertices not previously numbered when the recursive numbering algorithm is applied to a graph in C having n vertices, of which ℓ are previously numbered. Then

$$(32)\qquad q(\ell,n)$$

$$< n(n-1)(2n-1)/12 + 3n(n-1)/4 = n(n-1)(n+4)/6 \quad \text{if} \quad n \leq n_0 \text{, and}$$

$$q(\ell,n) \leq \beta^3 n^{3/2}/6 + \beta^2\ell n/2 + \beta\ell^2\sqrt{n}/2 + 3\beta^2 n/4 + 3\beta\ell\sqrt{n}/2 + \max\{q(\ell_1,n_1) + q(\ell_2,n_2)\}$$

otherwise, where the maximum is taken over values satisfying

$$(33)\qquad \ell_1 + \ell_2 \leq \ell + 2\beta\sqrt{n} \quad ,$$

$$n \leq n_1+n_2 \leq n + \beta\sqrt{n} \quad , \quad \text{and}$$

$$(1-\alpha)n \leq n_1, n_2 \leq \alpha n + \beta\sqrt{n} \quad .$$

We claim that for all $n \geq 1$,

$$(34)\qquad q(\ell,n)$$

$$\leq c_{11}n^{3/2} + c_{12}\ell n + c_{13}\ell^2\sqrt{n} + c_{14}n(\log_2 n)^2 + c_{15}\ell^2 \log_2 n + c_{16}\ell\sqrt{n} \quad ,$$

where

$$(35)\quad c_{11} = \beta^2\{1/6 + \beta\sqrt{\alpha}\,[2+\sqrt{\alpha}/(1+\sqrt{\alpha})+4\alpha/(1-\alpha)]/(1-\sqrt{\alpha})\}/(1-\delta)\ ,$$

$$c_{12} = \beta^2[1/2 + 2\sqrt{\alpha}/(1-\sqrt{\alpha})]/(1-\alpha)\ ,$$

$$c_{13} = \beta/[2(1-\sqrt{\alpha})]\ , \text{ and}$$

c_{14} , c_{15} , c_{16} are suitably large positive constants.

The theorem's bound on multiplications follows from the claim.

Proof of claim. Let

$$(36)\quad r(\ell,n) = c_{11}n^{3/2} + c_{12}\ell n + c_{13}\ell^2\sqrt{n} \quad \text{and}$$

$$s(\ell,n) = c_{14}n(\log_2 n)^2 + c_{15}\ell^2 \log_2 n + c_{16}\ell\sqrt{n}\ .$$

We prove the claim by induction on n . For $n \le n_4$, where $n_4 \ge n_0$ is a value to be selected later, $q(\ell,n) \le n(n-1)(n+4)/6 \le n_4(n_4-1)(n_4+4)/6 < s(\ell,n)$ if c_{14} is sufficiently large.

Let $n > n_4$ and suppose the claim is true for values smaller than n . Then

$$(37)\quad q(\ell,n) \le \beta^3n^{3/2}/6 + \beta^2\ell n/2 + \beta\ell^2\sqrt{n}/2 + 3\beta^2n/4 + 3\beta\ell\sqrt{n}/2$$
$$+ q(\ell_1,n_1) + q(\ell_2,n_2)$$
$$< \beta^3n^{3/2}/6 + \beta^2\ell n/2 + \beta\ell^2\sqrt{n}/2 + 3\beta^2n/4 + 3\beta\ell\sqrt{n}/2$$
$$+ c_{11}(n_1^{3/2}+n_2^{3/2}) + c_{12}(\ell_1n_1+\ell_2n_2) + c_{13}(\ell_1^2\sqrt{n_1} + \ell_2^2\sqrt{n_2})$$
$$+ s(\ell_1,n_1) + s(\ell_2,n_2)$$

for suitable values of ℓ_1 , n_1 , ℓ_2 , n_2 .

For fixed n_1+n_2 , the function $n_1^{3/2} + n_2^{3/2}$ is maximized when one of n_1, n_2 is as small as possible and the other is as large as possible. Thus

$$(38)\quad n_1^{3/2} + n_2^{3/2} \le [(1-\alpha)n]^{3/2} + [an + \beta\sqrt{n}\,]^{3/2}$$

$$< n^{3/2}[(1-\alpha)^{3/2} + \alpha^{3/2}(1+\beta/(\alpha\sqrt{n}\,))^{3/2}]$$

$$< n^{3/2}[(1-\alpha)^{3/2} + \alpha^{3/2}(1+\beta/(\alpha\sqrt{n}\,))^{2}]$$

$$\le n^{3/2}[(1-\alpha)^{3/2} + \alpha^{3/2}(1+3\beta/(\alpha\sqrt{n}\,))]$$

$$< [\alpha^{3/2} + (1-\alpha)^{3/2}]n^{3/2} + 3\beta\sqrt{\alpha}\; n$$

since $\beta/(\alpha\sqrt{n}\,) \le \beta/(\alpha\sqrt{n_0}\,) \le \beta/((1-\alpha)\sqrt{n_0}\,) = 1$.

Also

$$(39)\quad \ell_1 n_1 + \ell_2 n_2 \le (\ell + 2\beta\sqrt{n}\,)(\alpha n + \beta\sqrt{n}\,)$$

$$\le \alpha\ell n + 2\alpha\beta n^{3/2} + \beta\ell\sqrt{n} + 2\beta^2 n$$

and

$$(40)\quad \ell_1^2\sqrt{n_1} + \ell_2^2\sqrt{n_2} \le (\ell + 2\beta\sqrt{n}\,)^2\sqrt{\alpha n + \beta\sqrt{n}}$$

$$\le (\ell + 2\beta\sqrt{n}\,)^2(\sqrt{\alpha n} + \beta/(2\sqrt{\alpha}\,))$$

$$\le \sqrt{\alpha}\;\ell^2\sqrt{n} + 4\beta\sqrt{\alpha}\,\ell n + 4\beta^2\sqrt{\alpha}\; n^{3/2} + (\ell + 2\beta\sqrt{n}\,)^2\;\beta/(2\sqrt{\alpha}\,)\ .$$

Letting $\delta = \alpha^{3/2} + (1-\alpha)^{3/2}$ and combining the above inequalities with the bound on $q(\ell,n)$ gives

$$(41)\quad q(\ell,n) \le \beta^2 n^{3/2}/6 + c_{11}\delta n^{3/2} + 2c_{12}\alpha\;\beta n^{3/2} + 4c_{13}\beta^2\sqrt{\alpha}\; n^{3/2}$$

$$+ \beta^2\ell n/2 + c_{12}\alpha\ell n + 4c_{13}\beta\sqrt{\alpha}\;\ell n$$

$$+ \beta\ell^2\sqrt{n}/2 + c_{13}\sqrt{\alpha}\;\ell^2\sqrt{n}$$

$$+ 3\beta^2 n/4 + 3c_{11}\beta\sqrt{\alpha}\,n + 2c_{12}\beta^2 n + 2c_{13}\beta^3 n/\sqrt{\alpha}$$

$$+ 3\beta\ell\sqrt{n}/2 + c_{12}\beta\ell\sqrt{n} + 2c_{13}\beta^2\ell\sqrt{n}/\sqrt{\alpha}$$

$$+ c_{13}\beta\ell^2/(2\sqrt{\alpha}\,)$$

$$+ s(\ell_1,n_1) + s(\ell_2,n_2)$$

$$\leq c_{11}n^{3/2} + c_{12}\ell n + c_{13}\ell^2\sqrt{n}$$

$$+ c_{17}n + c_{18}\ell\sqrt{n} + c_{19}\ell^2 + s(\ell_1,n_1) + s(\ell_2,n_2)$$

where

(42)
$$c_{17} = 3\beta^2/4 + 3c_{11}\beta\sqrt{\alpha} + 2c_{12}\beta^2 + 2c_{13}\beta^3/\sqrt{\alpha} ,$$
$$c_{18} = 3\beta/2 + c_{12}\beta + 2c_{13}\beta^2/\sqrt{\alpha} , \text{ and}$$
$$c_{19} = c_{13}\beta/(2\sqrt{\alpha}) .$$

All that remains is to show that

(43)
$$s(\ell,n) \leq c_{17}n + c_{18}\ell\sqrt{n} + c_{19}\ell^2 + s(\ell_1,n_1) + s(\ell_2,n_2)$$

if c_{14}, c_{15}, c_{16}, and n_4 are chosen sufficiently large. This derivation is similar to that for the fill-in bound and we shall not go through it here. The claim follows by induction on n . □

Theorem 3. Let G be any planar graph. Then G has an elimination ordering which produces a fill-in of size $c_3 n\log n + O(n)$ and a multiplication count of $c_{11}n^{3/2} + O(n(\log n)^2)$, where $c_3 \leq 128.5$ and $c_{11} < 4002$. Such an ordering can be found in $O(n \log n)$ time.

Proof. By Corollary 2 of [15], planar graphs satisfy a $\sqrt{n}$ -separator theorem with $\alpha = 2/3$ and $\beta = 2\sqrt{2}$. Furthermore the appropriate vertex partition can be found in $O(n)$ time. Plugging into the bounds of Theorems 1-3 gives the result. □

A finite element graph is any graph formed from a planar embedding of a planar graph by adding all possible diagonals to each face. (The finite element graph has a clique corresponding to each face of the embedded planar graph.) The embedded planar graph is called the skeleton of the finite element graph and each of its faces is an element of the finite element graph.

Theorem 4. Let G be any n-vertex finite element graph with no element having more than k boundary vertices. Then G has an elimination ordering which produces a fill-in of size $O(k^2 n \log n)$ and multiplication count $O(k^3 n^{3/2})$. Such an ordering can be found in $O(n \log n)$ time.

Proof. By Corollary 4 of [15], any n-vertex finite element graph with no element having more than k boundary vertices satisfies a $\sqrt{n}$-separator theorem with $\alpha = 2/3$ and $\beta = 4\lfloor k/2 \rfloor$. Furthermore the appropriate vertex partition can be found in $O(n)$ time. Plugging into the bounds of Theorems 1-3 gives the result. □

Although planar and almost-planar graphs seem to be the most interesting case, analogues to Theorems 2- 4 hold for other classes of graphs. For instance, the following theorems can be proved using the same methods as in the proofs of Theorems 2 -4.

Theorem 5. Let S be any class of graphs closed under subgraph on which an n^σ separator theorem holds for $\sigma > 1/2$. Then for any n-vertex graph G in S, there is an elimination ordering with $O(n^{2\sigma})$ fill-in size and $O(n^{3\sigma})$ multiplication count.

The class of k-dimensional hypercubic grid graphs satisfies Theorem 6 for $\sigma = k-1/k$.

Theorem 6. Let S be any class of graphs closed under subgraph on which an n^{σ} separator theorem holds for $\sigma < 1/3 < 1/2$. Then for any n-vertex graph G in S there is an elimination ordering with O(n) fill-in size and $O(n^{3\sigma})$ multiplication count.

Theorem 7. Let S be any class of graphs closed under subgraph on which a $\sqrt[3]{n}$ separator theorem holds. Then for any n-vertex graph G in S, there is an elimination ordering with O(n) fill-in size and $O(n \log_2 n)$ multiplication count.

Theorem 8. Let S be any class of graphs closed under subgraph on which a n n^{σ} separator theorem holds for $\sigma < 1/3$. Then for any n-vertex graph G in S, there is an elimination ordering with O(n) fill-in . size and multiplication count.

3. Gaussian Elimination, Separators, and Sparsity,

In this section we explore additional relationships between sparse Gaussian elimination, good separators, and sparse graphs. We have shown that the existence of good separators in a graph and its subgraphs allows us to carry out sparse Gaussian elimination efficiently. It is natural to ask whether the converse is true; that is, whether the existence of good separators is necessary for efficient sparse elimination. To prove a result of this kind, we need a strengthened version of a lemma in [5]

Let $G = (V,E)$ be an undirected graph with an ordering π . Our proof technique makes use of the following algorithm, which adds edges to G and eventually produces a graph which contains the fill-in graph G^* . Associated with the graph during execution of the algorithm is a subset of its cliques, called elements. Initially the set of elements consists of the edges of the graph.

Element Merging Algorithm.

Repeat the following step for each vertex v from $\pi^{-1}(1)$ to $\pi^{-1}(n)$.

General step. Choose two elements e_1 and e_2 containing v . Add to the graph all edges not already present which join a vertex in e_1 and a vertex in e_2 ; simultaneously delete elements e_1 and e_2 and add a new element consisting of their union. Repeat until v is contained in only one element. Mark v eliminated.

Let G_k be the graph existing after k executions of the general step. We note the following properties of the algorithm, which are easy to verify.

(i) At all times during execution of the algorithm, every edge is contained in at least one element.

(ii) The number of elements containing a vertex never increases.

(iii) At the end of the algorithm, each connected component of the original graph comprises a single element.

(iv) After a vertex v is eliminated, v is contained in only one element.

(V) An edge {v,w} is a fill-in edge if and only if {v,w} is added to the graph before either v or w is eliminated. (In general G_n properly contains G_π^*.)

Property (v) follows from the definition of vertex elimination on a graph, which models Gaussian elimination on the corresponding matrix. 'See [5,17,18, 19,23].

Lemma 2. Let G = (V,E) be an n-vertex graph satisfying the following property for some $\ell < n/3$ and g : every set of vertices A such that $I \leq |A| \leq n-\ell$ is adjacent to at least g vertices in V-A . Then if π is any ordering of V , G_π^* contains a clique of at least g vertices.

Proof. G must have a connected component containing at least ℓ vertices. Otherwise there is a set A violating the hypothesis of the lemma, formed as follows. Let A = $\emptyset$. Add connected components to A one at a time until A contains at least ℓ vertices. Then A contains less than $21 < n-\ell$ vertices.

Apply the element merging algorithm to G with ordering π . Let e be the first element formed which contains at least ℓ vertices. Then e contains no more than 2ℓ vertices, since it is composed of two previously formed elements. Let A be the set of vertices not in e. A contains at least n-21 $\geq \ell$ and at most $n-\ell$ vertices. Let C be the set of vertices in e adjacent to at least one vertex in A. By the hypothesis of the lemma C contains at least g vertices. When e is formed, each vertex in C is in some element other than e by (i). Thus by (iv) each vertex in C is uneliminated when e is formed. By (v) the clique formed by C is contained in G_π^*. □

A weaker form of Lemma 2 and its proof, in Which the degrees of all vertices are assumed to be bounded, appears in [5].

Theorem 9. Let $G = (V,E)$ be a graph satisfying the hypothesis of Lemma2. Then any ordering of V produces a fill-in of size at least $g(g-1)/2$ and a multiplication count of at least $g(g-1)(g+4)/6$.

Proof. Immediate from Lemma 2. □

Theorem 9 and the results in Section 2 imply that generalized nested dissection is the best method of sparse elimination (to within a constant factor in running time and storage space) on large classes of graphs. For instance $n = k \times k$ square grid graphs satisfy the hypothesis of Lemma 2 for $\ell = n/3$ and $g = \sqrt{n}/3$ [15]. Thus such graphs have an $\Omega(n^{3/2})$ multiplication count for any ordering [13]. By using more sophisticated techniques, one can derive an $\Omega(n \log n)$ lower bound on the fill-in for such graphs [13]. For d-dimensional hypercubic grid graphs, Lemma 2 gives an $\Omega(n^{2(d-1)/d})$ lower bound on fill-in and an $\Omega(n^{3(d-1)/d})$ lower bound on multiplications, agreeing with the upper bounds in Theorem 5. See [5].

We turn now to the relationship between good separators and sparsity. Our first result shows that only sparse graphs have good separators.

Theorem 10. Let S be any class of graphs closed under subgraph and satisfying an $n/(\log_2 n)^2$ -separator theorem for fixed α, β . If G is a graph in S having n vertices and m edges, then $m \leq c_{21} n$.

Proof. Let t(n) be the maximum number of edges in any n-vertex graph in S . Let G be an n-vertex graph in S with t(n) edges. Since S satisfies an $n/(\log_2 n)^2$ -separator theorem, the vertices of G can be partitioned into three sets A, B, C such that C separates A and B , A and B each contain no more than αn vertices, and C contains no more than $\beta n/(\log_2 n)^2$ vertices. Since S is closed under subgraph, the subgraphs of G induced by the vertex sets $A \cup C$ and $B \cup C$ are both in S . If $|A \cup C| = n_1$ and $|B \cup C| = n_2$, it follows that $t(n) < t(n_1) + t(n_2)$. Hence

(44) $$t(n) \leq n(n-1)/2 \quad \text{if} \quad n \leq n_0, \text{ and}$$

$$t(n) \leq \max\{t(n_1) + t(n_2)\} \quad \text{otherwise,}$$

where the maximum is taken over values n_1, n_2 satisfying

(45) $$n < n_1 + n_2 \leq n + \beta n/(\log_2 n)^2 \text{ , and}$$

$$(1-\alpha)n \leq n_1 , n_2 \leq \alpha n + \beta n/(\log_2 n)^2 .$$

An inductive proof like those in Section 2 shows that

(46) $$t(n) \leq c_{21} n - c_{22} n/\log_2 n ,$$

where c_{21} and c_{22} are suitably large positive constants. □

Not all sparsegraphshave good separators. In fact, for fixed a, β such that $\beta <$ la! $\leq \alpha < 1$, there is a constant c such that almost all*/ n-vertex graphs with cn edges have no vertex partition A, B, C satisfying $|A|, |B| \leq \alpha n$, $|C| \leq$ Bn.., and C separates A and B . This result is implicit in Theorem 4 of [8]. It follows from Theorem 9 that almost all sparse graphs require $\Omega(n^2)$ fill-in and $\Omega(n^3)$ multiplication count. By using a more direct argument, we can obtain a stronger result.

Theorem 11. For all $\epsilon > 0$ there is a constant $c(\epsilon)$ such that almost all n-vertex graphs with at least $c(\epsilon)n$ edges have a fill-in clique of at least $(1-\epsilon)n$ vertices for any ordering.

Proof. We first prove that almost all n-vertex graphs with at least cn edges have the following property:

(P) If A and B are sets of vertices such that $|A|, |B| \geq \epsilon n/4$ and $A \cap B = \emptyset$, then at least one edge joins A and B .

We prove (P) by an argument like that used to prove Theorem 4 of [8]. Consider a random graph G with n vertices and m edges, where $m > cn$. The number of ways to choose two vertex sets A, B satisfying $|A|, |B| \geq \epsilon n/4$, $A \cap B = \emptyset$ is less than 3^n . Between A and B there are at least $\epsilon^2 n^2/16$ potential edges. The probability that none of these edges actually occurs in G is less than $(1 - 2c/n)^{\epsilon^2 n^2/16}$. This, if c is

*/ By "almost all" we mean that the fraction of n-vertex graphs satisfying the property tends with increasing n to one. We assume that each n-vertex graph has vertex set $\{1,2,\ldots,n\}$ and that two graphs are distinct unless their edge sets are identical. See [7] for a thorough discussion of random graphs.

chosen so that $3^n(1-2c/n)^{\epsilon^2 n^2/16} \to 0$ as $n \to \infty$, then almost all graphs satisfy (P). Since $(1-2c/n)^{\epsilon^2 n^2/16} \to e^{-c\epsilon^2 n/8}$, choosing $c > (8 \log, 3)/\epsilon^2$ gives the result.

Now we use (P) to prove the theorem." Let $G = (V,E)$ be any graph satisfying (P). Consider any set A of at least $3\epsilon n/4$ vertices in G , A contains a subset B of at least $\epsilon n/4$ vertices whose induced subgraph in G is connected. Otherwise, we can derive a contradiction as follows, Let $A_1, A_2,$ ● ✡🔔 be the vertex sets of the connected components of the subgraph of G induced by A . Let j be the minimum index such that $\sum_{i=1}^{j} |A_i| \geq \epsilon n/4$. Then $\sum_{i=1}^{j} |A_i| \leq \epsilon n/2$. By (P) there must be an edge joining some vertex in $\bigcup_{i=1}^{j} A_i$ with some vertex in $\bigcup_{i=j+1}^{k} A_i$. This is impossible by the definition of the A_i's.

Consider any ordering of the vertices of G . Let A be the first $3\epsilon n/4$ vertices in the ordering. Let B be a subset of A containing at least $\epsilon n/4$ vertices whose induced subgraph in G is connected. By property (P) at least $(1-\epsilon/2)n$ vertices in V-B , and hence at least $(1-\epsilon)n$ vertices in V-A , must be adjacent to at least one vertex in B . By Lemma 1, any pair of such vertices are joined by a fill-in edge. Thus the set of vertices in V-B adjacent to at least one vertex in B is a fili-in clique of at least $(1-\epsilon)n$ vertices. ☐

Theorem 12. Almost all n-vertex graphs with $c(\epsilon)n$ edges have a fill-in of $(1-\epsilon)^2 n^2/2 - O(n)$ and a multiplication count of $(1-\epsilon)^3 n^3/6 - O(n^2)$, for any ordering.

Proof. Immediate from Theorem 11. ☐

4. Remarks.

We have demonstrated the existence of an $O(n^{3/2})$ -time, $O(n \log n)$ -space method for carrying out sparse Gaussian elimination on systems whose pattern of non-zeros corresponds to a planar or two-dimensional finite element graph. Such systems arise often in real problems. The practicality of the algorithm remains to be tested, and the constants in Theorem 3 are large. However, we believe that the algorithm is potentially useful for solving large systems, since the worst-case bounds derived here are probably much too pessimistic. Experiments by George and Liu [10] with a similar algorithm suggest that our method is practical.

It is possible to reduce the running time of our algorithm to $O(n^{\log_2 7})$ by using Strassen's algorithm for matrix multiplication and factorization [3,21]. If the system of equations is to be solved for just one right-hand side b, it is possible to reduce the storage required to $O(n)$ by storing only part of L and recomputing the rest as necessary. Reference [5] describes how to achieve these savings in the case of ordinary nested dissection; the generalization to planar and almost-planar graphs is analogous to the results in Section 2.

Gaussian elimination can be used to solve systems of linear equations defined over algebras other than the real numbers [2,4,22], and the algorithm in Section 2 applies to these other situations. For instance, the single-source shortest paths problem with negative-weight edges can be solved in $O(n^{3/2})$ time on planar graphs. The best general sparse algorithm [14] requires $O(n^2 \log n)$ time.

The results in Section 2 show that the existence of good separators in a graph and its subgraphs is enough to guarantee that sparse Gaussian elimination is efficient. Conversely, Theorem 9 in Section 3 shows that a graph for which Gaussian elimination is efficient must have a good separator. The existence of good separators in a graph and its subgraphs implies that the graph is sparse, but almost all sparse graphs do not have good separators. These results suggest that when studying Gaussian elimination, one should regard a graph as "sparse" when it has good separators rather than when it has a small edge/vertex ratio.

A number of questions remain to be explored, Can generalized nested dissection be implemented efficiently? Is it practical? How does one find good separators in a graph? What is a useful definition of the "goodness" of a separator? Informally, a separator is good if it is small and divides the graph into small pieces. We need a quantitative definition which embodies this idea. What are the trade-offs between the size of the separator and the size of the pieces it produces? The property of having good separators is crucial not only in Gaussian elimination but in many other problems [16].

Appendix: Definitions

A graph $G = (V,E)$ consists of a set V of vertices and a set E of edges. Each edge is an unordered pair $\{v,w\}$ of distinct vertices. If $\{v,w\}$ is an edge, v and w are adjacent, v and w are incident to $\{v,w\}$, and v and w are the endpoints of $\{v,w\}$. A path of length k with endpoints v, w is a sequence of vertices $v = v_0, v_1, v_2, \ldots, v_k = w$ such that $\{v_{i-1}, v_i\}$ is an edge for $1 \leq i < k$. If $G_1 = (V_1,E_1)$ and $G_2 = (V_2,E_2)$ are graphs, G_1 is a subgraph of G_2 if $V_1 \subseteq V_2$ and $E_1 \subseteq E_2$. If $G = (V_2,E_2)$ is a graph and $V_1 \subseteq V_2$, the graph $G_1 = (V_1,E_1)$ where $E_1 = E_2 \cap \{\{v,w\} \mid v,w \in v\}$ is the subgraph of G_2 induced by the vertex set V_1. A clique is a graph in which an edge joins every pair of distinct vertices. A graph is connected if every pair of its vertices are joined by a path, The connected components of a graph are its maximal connected subgraphs, Let A, B, C be a partition of the vertices of a graph $G = (V,E)$. We say C separates A and B if no edge joins a vertex in A with a vertex in B.

If f and g are functions of n, " $f(n)$ is $O(g(n))$ " means that for some positive constant c, $f(n) \leq cg(n)$ for all but finitely many values of n; " $f(n)$ is $\Omega(g(n))$ " means $g(n)$ is $O(f(n))$.

A graph $G = (V,E)$ is planar if there is a one-to-one map f_1 from v into points in the plane and a map f_2 from E into simple curves in the plane such that, for each edge $\{v,w\} \in E$, $f_2(\{v,w\})$ has endpoints $f_1(v)$ and $f_2(w)$, and no two curves $f_2(\{v_1,w_1\})$, $f_2(\{v_2,w_2\})$ share a point except possibly a common endpoint. Such a pair of maps f_1, f_2 is a planar embedding of G. The connected planar regions formed when the ranges of f_1 and f_2 are deleted from the plane are called the faces of the embedding. Each face is bounded by a curve corresponding to a cycle of G, called the boundary of the face. We shall sometimes not distinguish between a face and its boundary. A diagonal of a face is an edge (v,w) such that v and w are non-adjacent vertices on the boundary of the face.

References

[1] A. V. Aho, J. E. Hopcroft, and J. D. Ullman, The Design and Analysis of Computer Algorithms, Addison-Wesley, Reading, Mass., 1974.

[2] R. C. Backhouse and B. A, Carré, "Regular algebra applied to path-finding problems," J. Inst. Maths. Applics. 15 (1975), 161-186.

[3] J. Bunch and J. E. Hopcroft, "Triangular factorization and inversion by fast matrix multiplication," Math. Comp. 28 (1974), 231-236.

[4] B. A. Carré, "An algebra for network routing problems," J. Inst. Maths. Applics. 7 (1971), 273-294.

[5] S. C. Eisenstat, M. H. Schultz, and A. H. Sherman, "Applications of an element model for Gaussian elimination,' Sparse Matrix Computations, J. R. Bunch and D. J. Rose, eds., Academic Press, New York, 1976, 85-96.

[6] S. C. Eisenstat, M. H. Schultz, and A. H. Sherman, "Considerations in the design of software for sparse Gaussian elimination," Sparse Matrix Computations, J. R. Bunch and D. J. Rose, eds., Academic Press, New York, 1976, 263-273.

[7] P. Erdös and J. Spencer, Probabilistic Methods in Combinatorics, Academic Press, New York, 1974.

[8] P. Erdös, R. L. Graham, and E. Szemerédi, "On sparse graphs with dense long paths," Computers and Mathematics, 1, 3/4 (1975), 365-370.

[9] G. E. Forsythe and C. B. Moler, Computer Solution of Linear Algebraic Equations, Prentice-Hall, Englewood Cliffs, N. J., 1967.

[10] A. George and J. W. H. Liu, "An automatic nested dissection algorithm for irregular finite element problems," Research report CS-76-38, Department of Computer Science, University of Waterloo (1976).

[11] J. A. George, "Nested dissection of a regular finite element mesh," SIAM J. Num. Anal. 10 (1973), 345-367.

[12] F. G. Gustavson, "Some basic techniques for solving sparse systems of linear equations," Sparse Matrices and Their Applications, D. J. Rose and R. A. Willoughby, eds., Plenum Press, New York, 1972, 41-52.

[13] A. J. Hoffman, M. S. Martin, and D. J. Rose, "Complexity bounds for regular finite difference and finite element grids," SIAM J. Numer. Anal. 10 (1973), 364-369.

[14] D. B. Johnson, "Efficient algorithms for shortest paths in sparse networks," Journal ACM 24 (1977), 1-13.

[15] R. J. Lipton and R. E. Tarjan, "A separator theorem for planar graphs," to appear.

[16] R. J. Lipton and R. E. Tarjan, "Applications of a planar separator theorem," to appear.

[17] D. J. Rose, "Triangulated graphs and the elimination process," J. Math. Anal. Appl. 32 (1970), 597-609.

[18] D. J. Rose, "A graph-theoretic study of the numerical solution of sparse positive definite systems of linear equations," Graph Theory and Computing, R. Read, ed., Academic Press, New York, 1973, 183-217.

[19] D. J, Rose, R. E. Tarjan, and G. S. Lueker, "Algorithmic aspects of vertex elimination on graphs," SIAM J. Comput. 5 (1976), 266-283.

[20] A. H. Sherman, "On the efficient solution of sparse systems of linear and nonlinear equations," Ph.D. thesis, Computer Science Dept., Yale University (1975).

[21] V, Strassen, "Gaussian elimination is not optimal," Numerische Mathematik 13 (1969), 354-356.

[22] R. E. Tarjan, "Solving path problems on directed graphs," Technical Report STAN-CS-75 -528, Computer Science Dept., Stanford University (1975).

[23] R. E. Tarjan, "Graph theory and Gaussian elimination," Sparse Matrix Computations, J. R. Bunch and D. J. Rose, eds., Academic Press, New York, 1976, 3-22.

[24] O. C. Zienkiewicz, The Finite Element Method in Engineering Science, McGraw-Hill, London, 1970.