

GENERALIZED NESTED DISSECTION*

RICHARD J. LIPTON,[†] DONALD J. ROSE[‡] AND ROBERT ENDRE TARJANS

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 $O(n \log n)$ and space $O(n^{3/2})$ 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.

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

$$(1) \quad Ax = b$$

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

$$(2) \quad A = LDL^T$$

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 nonzero 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 nonzeros in L (and L^T) in positions where A contains zeros. These new nonzeros 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

$$(3) \quad A' = PAP^T$$

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*¹ $G = (V, E)$. The graph G contains one vertex $i \in V$ for each row (and column) in A , and one edge $\{i, j\} \in E$ for

* Received by the editors December 1, 1977.

[†] Computer Science Department, Yale University, New Haven, Connecticut 06520. The research of this author was partially supported by the U.S. Army Research Office under Grant DAAG 29-76-G-0338.

[‡] Computer Science Department, Vanderbilt University, Nashville, Tennessee 37235. The research of this author was partially supported by the Office of Naval Research under Contract N00014-67-0298-0034. Now at Math Center, Bell Laboratories, Murray Hill, New Jersey 07974.

[§] Computer Science Department, Stanford University, Stanford, California 94305. The research of this author was partially supported by the National Science Foundation under Grant MCS-75-22870 and the Office of Naval Research under Contract N00014-76-C-0388.

¹ The Appendix contains the graph-theoretic definitions used in this paper. It also defines the " O " and " Ω " notations.

each pair of nonzero, off-diagonal P corresponds to a numbering $\pi: V \rightarrow \{1, 2, \dots, n\}$. Corresponding $G_\pi^* = (V, E_\pi^*)$ such that $\{i, j\} \in E_\pi^*$ if $\pi(j)$ is nonzero. See [18], [19] for a graph-theoretic model of sparse matrices. The fill-in E_π^* associated with π is defined as $E_\pi^* = E \cup \{ \{i, j\} \mid \{i, j\} \in E_\pi^* \text{ and } \{i, j\} \notin E \}$.

LEMMA 1 [20]. Assuming $\{v, w\} \in E_\pi^*$ iff $v \neq w$ and there exists a path P from v to w in G_π^* such that $\min \{\pi(v), \pi(w)\} \leq \pi(i) \leq \pi(j)$ for all $\{i, j\} \in P$, then π satisfies this condition.

The running time and storage space are functions of m , the number of nonzeros in A , and n , the number of nodes. Note that $d(i) = \deg(i)$ in G , and that $m = \sum_{i=1}^n d(i)$. For sparse elimination into LDL^T , we divide sparse elimination into the following steps:

- Step 1 Find a good ordering π .
Time: $O(m)$.
Space: $O(m)$.
- Step 2 (Symbolic factorization) Factorize A into LDL^T .
Time: $O(m)$ using no lucky cancellations.
Space: $O(m)$.
- Step 3 (Numeric factorization) Solve $Lz = b$.
Time: $O(\sum_{i=1}^n d(i)^2)$.
Space: $O(m)$.
- Step 4 (Backsolving) Solve $Dy = z$ and $L^T x = y$.
Time: $O(m)$ [19].
Space: $O(m)$.

The reason for separating the steps is that at all known ways of implementing numeric factorization without fill-in require a large overhead which is more than a constant factor. If the system of equations is sparse, it is possible to combine at least parts of the steps.

The efficiency of sparse Gaussian elimination depends on finding an ordering π which minimizes the fill-in. Finding such an ordering is a very hard, perhaps even NP-complete problem. Several ordering schemes are known, including the one by George [11], which allows solving a system on a grid graph in $O(n^{3/2})$ time by the removal of $O(k)$ vertices from the graph at each step, for a total of $k/2 \times k/2$ [21].

In this paper we generalize the graph-theoretic model of sparse matrices to a subgraph relation (i.e., if G_2 is a subgraph of G_1 , then G_2 satisfies an $f(n)$ -separator theorem). We show that a vertex graph G in S has the following property: there are three sets A, B, C such that

DISSECTION*

D. ROBERT ENDRE TARJAN†

Nested dissection, for solving a system of linear equations in n variables and space $O(n^{3/2})$ time. We generalize this to any system of equations defined on a graph. The algorithm applies to any system of equations defined on a graph of two-dimensional finite element problems.

It is efficient for any class of graphs which includes planar graphs and their minors (including "almost all" sparse graphs).

Gaussian elimination the system of equations.

Let x be an $n \times 1$ vector of variables. The process consists of two steps. First, we solve the simplified system of equations.

Next, we solve the simplified system of equations.

Then the time required for solving the simplified systems is $O(n^{3/2})$.

We may be able to save time by using a good ordering of the vertices.

One difficulty with obtaining a good ordering is that it may require creating nonzeros in L (and L^T).

These nonzeros are called *fill-in*.

The rows and columns of A , in the original system, are called *active*.

ordered system. Since A is positive definite with respect to the LDL^T factorization.

a given permutation matrix P and graph $G = (V, E)$. The graph G is defined by the vertices V and edges E in A , and one edge $\{i, j\} \in E$ if and only if $a_{ij} \neq 0$ and $a_{ji} \neq 0$.

Yale University, New Haven, Connecticut 06520. The research of the author was supported in part by a grant under Grant DAAG 29-76-G-0338 from the Army Research Office-Durham, Durham, North Carolina 27706. The research of the author was also supported in part by a grant under Contract N00014-67-0298-0034 from the Naval Research Office, Washington, D.C. 20340.

University of California, San Diego, La Jolla, California 94305. The research of the author was supported in part by a grant under Grant MCS-75-22870 and the Office of Naval Research, Arlington, Virginia 22204.

In this paper, it also defines the "O" and "Θ" notations.

each pair of nonzero, 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 $\pi: V \rightarrow \{1, 2, \dots, n\}$. 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 nonzero. See [18], [19], [20], [25] 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 [20]. *Assuming no cancellation of nonzeros in the factoring of PAP^T , $\{v, w\} \in E_\pi^*$ iff $v \neq w$ and there is a path $v = v_1, v_2, \dots, v_{k+1} = w$ such that $\pi(v_i) < \min\{\pi(v), \pi(w)\}$ for $2 \leq i \leq k$. (Note that a path consisting of a single edge $\{v, w\} \in E$ satisfies this condition.)*

The running time and storage space required by sparse Gaussian elimination are functions of m , the number of nonzeros in L , and of $d(i)$, the number of edges $\{i, j\}$ in G_π^* with $\pi(i) < \pi(j)$. Note that $d(i)$ is the number of nonzeros 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.

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 nonzero positions in L , assuming no lucky cancellation of nonzeros.

Time: $O(m)$ using the algorithm of [20].

Space: $O(m)$.

Step 3 (Numeric factorization.) Compute L .

Time: $O(\sum_{i=1}^{n-1} d(i)(d(i)+3))$ using an algorithm such as described in [6], [12], [22], [25]. The number of multiplications performed during this step is $\frac{1}{2} \sum_{i=1}^{n-1} d(i)(d(i)+3)$ [19].

Space: $O(m)$.

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

Time: $O(m)$ [19].

Space: $O(m)$.

The reason for separating the factorization into two steps (symbolic and numeric) is that all known ways of implementing sparse Gaussian elimination 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 $\frac{1}{2} \sum_{i=1}^{n-1} d(i)(d(i)+3)$. 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 $k/2 \times k/2$ [21].

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 $f(n)$ -separator theorem if there are constants $\frac{1}{2} \leq \alpha < 1$, $\beta > 0$ 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 $|C| \leq \beta f(n)$. Our 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" method [1] to generate the ordering. From separator theorems proved in [16], 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 [26]. 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 § 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 discussed the significance of the results in §§ 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 $l = 0$ 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

$$(4) \quad \begin{aligned} t(m, n) &\leq c_1 n^2 && \text{if } n \leq n_0, \\ t(m, n) &\leq c_2(m+n) + \max \{t(m_1, n_1) + t(m_2, n_2)\} && \text{otherwise.} \end{aligned}$$

² Throughout this paper, $c, c_0, c_1, \dots$, denote suitable positive constants.

where $n_0 = (\beta/(1-\alpha))^2$ and the satisfying

$$(5) \quad \begin{aligned} m_1 + m_2 &\leq m, \\ n &\leq n_1 + n_2 \leq n + \beta \sqrt{n} \end{aligned}$$

A proof by induction similar to (5) is $O((m+n) \log n)$. $\square$

THEOREM 2. Let G be any graph. Let $f(l, n)$ be the fill-in associated with the numbering algorithm.

$$(6) \quad c_3 = \beta^2 \left(\frac{1}{2} \right)$$

Proof. We shall prove that the analysis [15] gives the bound claimed.

Suppose the recursive numbering algorithm has numbered l vertices previously numbered. Assume C is the set of vertices generated by the algorithm. If C contains k vertices, then the number of fill-in edges whose lower endpoint is in C is at most

$$k(k-1)/2.$$

By Lemma 1, two vertices v and w are connected by an edge from v to w through vertices $n_1, \dots, n_k$ if and only if there is a vertex in A with a vertex in B whose lower numbered endpoint is in C (numbered). Then

$$f(l, n) \leq n(n-l)/2$$

$$f(l, n) \leq \beta^2 n^2 / 2$$

otherwise, where the maximum is taken over all possible partitions.

$$l_1 + l_2 \leq l$$

$$n \leq n_1 + n_2$$

$$(1 - \alpha - \beta/\sqrt{n_0})^2$$

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

$$f(l, n) \leq c_4 n + c_5 \log n$$

where c_4 and c_5 are suitably large constants. The bound of $O(n \log n)$ on fill-in size follows.

We prove the claim by induction on n . The claim is true for $n \leq n_0$. Then

$$f(l, n) \leq n(n-1)/2 \leq c_4 n + c_5 \log n$$

$$n \leq n_3/2.$$

Let $n > n_3$ and suppose the claim holds for n_1 and n_2 .

$$f(l, n) \leq \beta l \sqrt{n} + f(l_1, n_1) + f(l_2, n_2) \leq c_4 n + c_5 \log n$$

$$\text{Let } \varepsilon = (1 - \alpha - \beta/\sqrt{n_0} + 1). \text{ Then } \varepsilon_i \leq \varepsilon + 1 < 1, \text{ and } \varepsilon > 0. \text{ Thus } n_i \leq n \text{ and the claim holds for } n_1 \text{ and } n_2 \text{ by induction.}$$

at all systems of equations wh
 ed in $O(n^{3/2})$ time and $O(n \log n)$
 to generate the ordering. F
 method for solving any system
 $O(n^{3/2})$ time and $O(n \log n)$ sp
 onal finite element problems [

and ordering scheme for any cla
 fill-in and multiplication of
 the end of § 2 we list fill-in
 er than $f(n) = \sqrt{n}$.

een Gaussian elimination, g
 lower bond on the cost of Gaus
 ne problem graph. We prove
 now that "almost all" sparse gra
 ction 4 discussed the significant

ss of graphs closed under subgr
 e the constants associated with
 vertex graph in S . The follow
 hat sparse Gaussian eliminatio
 rtices of G are already assign
 e remaining vertices of G are to

ore than $n_0 = (\beta/(1-\alpha))^2$ verti
 to b . Otherwise, find sets A , B
 unnumbered vertices, B cont
 d vertices.

rily from $b-k+1$ to b . Delete
 ithm recursively to the subgr
 es in B from $b-k-j+1$ to b
 nduced by $A \cup C$ to number
 to $a+i-1 = b-k-j$.

pplying this algorithm to G
 from 1 to n . We are intereste
 e size of the fill-in produced by
 of the generated ordering.

fying the $\sqrt{n}$ -separator theorem
 ph. Then the numbering algo

ed by the numbering algorithm

if $n \leq n_0$.

otherwise.

ve constants.

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

$$\begin{aligned}
 (5) \quad & m_1 + m_2 \leq m, \\
 & n \leq n_1 + n_2 \leq n + \beta\sqrt{n}, \quad \text{and} \quad (1-\alpha)n \leq n_i \leq \alpha n + \beta\sqrt{n} \quad \text{for } i = 1, 2.
 \end{aligned}$$

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

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*

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

Proof. We shall prove that the fill-in is $O(n \log n)$. A more careful but lengthier analysis [15] gives the bound claimed in the theorem.

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

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

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 $f(l, n)$ be the maximum number of fill-in edges whose lower numbered endpoint is numbered by the algorithm (and not previously numbered). Then

$$\begin{aligned}
 (8) \quad & f(l, n) \leq n(n-1)/2 \quad \text{if } n \leq n_0, \quad \text{and} \\
 & f(l, n) \leq \beta^2 n/2 + \beta l\sqrt{n} + \max \{f(l_1, n_1) + f(l_2, n_2)\}
 \end{aligned}$$

otherwise, where the maximum is taken over values satisfying

$$\begin{aligned}
 (9) \quad & l_1 + l_2 \leq l + 2\beta\sqrt{n}, \\
 & n \leq n_1 + n_2 \leq n + \beta\sqrt{n}, \\
 & (1-\alpha)n \leq n_i \leq \alpha n + \beta\sqrt{n} \quad \text{for } i = 1, 2.
 \end{aligned}$$

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

$$(10) \quad f(l, n) \leq c_4(l+n) \log_2 n + c_5 l\sqrt{n}$$

where c_4 and c_5 are suitably large positive constants, to be chosen later. The desired bound of $O(n \log n)$ on fill-in size follows from the claim.

We prove the claim by induction on n . Assume $n \leq n_3$, where $n_3 \geq n_0$ is a value to be chosen later. Then

$$(11) \quad f(l, n) \leq n(n-1)/2 \leq (n_3/2)(n-1) \leq c_4(l+n) \log_2 n + c_5 l\sqrt{n},$$

if $c_4 \geq n_3/2$.

Let $n > n_3$ and suppose the claim is true for values smaller than n . Then $f(l, n) \leq \beta^2 n/2 + \beta l\sqrt{n} + f(l_1, n_1) + f(l_2, n_2)$ 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 $n_i \leq \alpha n + \beta\sqrt{n} \leq (\alpha + \beta/\sqrt{n})n \leq (1-\epsilon)n < n$ for $i = 1, 2$, and the claim holds for n_1 and n_2 by the induction hypothesis.

Hence

$$(12) \quad \begin{aligned} f(l, n) &\leq \beta^2 n/2 + \beta l\sqrt{n} + c_4(l_1 + n_1) \log_2 n_1 + (l_2 + n_2) \log_2 n_2 + c_5(l_1\sqrt{n_1} + l_2\sqrt{n_2}) \\ &\leq \beta^2 n/2 + \beta l\sqrt{n} + c_4(n + l + 3\beta\sqrt{n}) \log_2((1 - \varepsilon)n) \\ &\quad + c_5(l + 2\beta\sqrt{n})\sqrt{\alpha n + \beta\sqrt{n}}. \end{aligned}$$

Since $\sqrt{\alpha n + \beta\sqrt{n}} \leq \sqrt{\alpha n} + \beta/(2\sqrt{\alpha}) = \sqrt{\alpha n} + c_6$, we have

$$(13) \quad \begin{aligned} f(l, n) &\leq c_4(n + l) \log_2 n + (c_5\sqrt{\alpha} + \beta)l\sqrt{n} + \beta^2 n/2 + 3c_4\beta\sqrt{n} \log_2 n + 2c_5\beta\sqrt{\alpha n} \\ &\quad + c_5c_6(l + 2\beta\sqrt{n}) + c_4(n + l) \log_2(1 - \varepsilon). \end{aligned}$$

Suppose we choose n_3 large enough so that $n > n_3$ implies $3\beta\sqrt{n} \log_2 n \leq n \log_2(1/(1 - \varepsilon))/2$, choose c_5 large enough so that $c_5\sqrt{\alpha} + \beta \leq c_5$, i.e., $c_5 \geq \beta/(1 - \alpha)$, and choose c_4 large enough so that $\beta^2/2 + 2c_5\beta\sqrt{\alpha} + c_5c_6(1 + 2\beta) \leq c_4 \log_2(1/(1 - \varepsilon))/2$. Then $f(l, n) \leq c_4(n + l) \log_2 n + c_5l\sqrt{n}$ as desired, and the claim follows by induction. $\square$

THEOREM 3. Let G be any n -vertex graph numbered by the algorithm. The total multiplication count associated with the numbering is at most $c_7n^{3/2} + O(n(\log n)^2)$, where

$$(14) \quad c_7 = \beta^2(1/6 + \beta\sqrt{\alpha}(2 + \sqrt{\alpha}/(1 + \sqrt{\alpha}) + 4\alpha/(1 - \alpha))/(1 - \sqrt{\alpha}))/ (1 - \delta)$$

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

Proof. We shall prove that the number of multiplications is $O(n^{3/2})$. A more careful analysis [15] gives the bound claimed in the theorem.

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-numbered 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

$$(15) \quad \begin{aligned} \sum_{i=0}^{\beta\sqrt{n}-1} (i + l)(i + l + 3)/2 &\leq \sum_{i=0}^{\beta\sqrt{n}-1} (i + l)^2/2 + 3\beta^2 n/4 + 3\beta l\sqrt{n}/2 \\ &\leq \beta^3 n^{3/2}/6 + \beta^2 ln/2 + \beta l^2\sqrt{n}/2 + 3\beta^2 n/4 + 3\beta l\sqrt{n}/2. \end{aligned}$$

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

$$(16) \quad \begin{aligned} g(l, n) &\leq n(n - 1)(2n - 1)/12 + 3n(n - 1)/4 = n(n - 1)(n + 4)/6 \quad \text{if } n \leq n_0, \text{ and} \\ g(l, n) &\leq \beta^3 n^{3/2}/6 + \beta^2 ln/2 + \beta l^2\sqrt{n}/2 + 3\beta^2 n/4 + 3\beta l\sqrt{n}/2 \\ &\quad + \max\{g(l_1, n_1) + g(l_2, n_2)\} \end{aligned}$$

otherwise, where the maximum is taken over values satisfying

$$(17) \quad \begin{aligned} l_1 + l_2 &\leq l + 2\beta\sqrt{n}, \\ n &\leq n_1 + n_2 \leq n + \beta\sqrt{n}, \text{ and} \\ (1 - \alpha)n &\leq n_i \leq \alpha n + \beta\sqrt{n} \quad \text{for } i = 1, 2. \end{aligned}$$

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

$$(18) \quad g(l, n) \leq c_8 n^{3/2} + c_9 ln + c_{10} l^2\sqrt{n},$$

where c_8, c_9 , and c_{10} are suitably large constants depending on $O(n^{3/2})$ on multiplications followed by $O(n)$ on additions.

We prove the claim by induction on n . The claim follows for $n \leq n_0$, selected later.

$$(19) \quad g(l, n) \leq n(n - 1)(n + 4)/6$$

if c_8 is sufficiently large.

Let $n > n_4$ and suppose the claim holds for all $n' < n$. In (12), we have

$$(20) \quad \begin{aligned} g(l, n) &\leq c_{11} n^{3/2} \\ &\quad + c_9(l_1 n_1 + l_2 n_2) \end{aligned}$$

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

For fixed $n_1 + n_2$, the function $g(l_1, n_1) + g(l_2, n_2)$ is small as possible and the other is as large as possible.

$$n_1^{3/2} + n_2^{3/2} \leq [(n_1 + n_2)^{3/2} + (n_1 - n_2)^{3/2}]/2$$

$$\leq n^{3/2}$$

$$\leq n^{3/2}$$

$$\leq n^{3/2}$$

$$\leq [\alpha n^{3/2} + (1 - \alpha)n^{3/2}]$$

Let $\alpha \geq 1/2$ implies $\beta(\alpha\sqrt{n}) \leq \beta/(\alpha\sqrt{n})$. Also

$$l_1 n_1 + l_2 n_2 \leq (l + 2\beta\sqrt{n})^2/4$$

$$\leq (l + 2\beta\sqrt{n})^2/4$$

$$l_1^2\sqrt{n_1} + l_2^2\sqrt{n_2} \leq (l + 2\beta\sqrt{n})^2/4$$

$$\leq (l + 2\beta\sqrt{n})^2/4$$

$$\leq \sqrt{\alpha} l^2\sqrt{n} + 4\beta l\sqrt{n} + 4\beta^2 n$$

Letting $\delta = \alpha^{3/2} + (1 - \alpha)^{3/2}$ and $\alpha \geq 1/2$ gives

$$g(l, n) \leq (c_{11} + c_8\delta + 2)$$

$$+ (c_{12} + c_9\alpha + c_{10})$$

$$+ c_{14}(c_8 + c_9 + c_{10})$$

where c_{14} is a suitably large constant.

Suppose we choose n_4 large enough so that

$$(1 - \delta)n^{1/2}/2, (1 - \alpha)n^{1/2}/2, (1 + \alpha)n^{1/2}/2 \leq c_{10},$$

$$c_{10}(1 + \sqrt{\alpha})/2 \leq c_9, \text{ and choose } c_{10} + 4\beta\sqrt{\alpha} \leq c_9, \text{ and choose } c_{10} + 2c_9\alpha\beta + 4c_{10}\beta^2\sqrt{\alpha} \leq c_8.$$

$$c_{10} + 2c_9\alpha\beta + 4c_{10}\beta^2\sqrt{\alpha} \leq c_8$$

$$c_{10} + 2c_9\alpha\beta + 4c_{10}\beta^2\sqrt{\alpha} \leq c_8$$

The claim follows by induction.

where c_8, c_9 , and c_{10} are suitably large constants, to be chosen later. The desired bound of $O(n^{3/2})$ on multiplications follows from the claim.

We prove the claim by induction on n . For $n \leq n_4$, where $n_4 \geq n_0$ is a value to be selected later,

$$(19) \quad g(l, n) \leq n(n-1)(n+4)/6 \leq n_4(n_4-1)(n_4+4)/6 \leq c_8 n^{3/2},$$

if c_8 is sufficiently large.

Let $n > n_4$ and suppose the claim is true for values smaller than n . Then, analogous to (12), we have

$$(20) \quad \begin{aligned} g(l, n) &\leq c_{11} n^{3/2} + c_{12} l n + c_{13} l^2 \sqrt{n} + c_8 (n_1^{3/2} + n_2^{3/2}) \\ &\quad + c_9 (l_1 n_1 + l_2 n_2) + c_{10} (l_1^2 \sqrt{n_1} + l_2^2 \sqrt{n_2}) \end{aligned}$$

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

$$(21) \quad \begin{aligned} n_1^{3/2} + n_2^{3/2} &\leq [(1-\alpha)n]^{3/2} + [\alpha n + \beta \sqrt{n}]^{3/2} \\ &\leq n^{3/2} [(1-\alpha)^{3/2} + \alpha^{3/2} (1 + \beta/(\alpha \sqrt{n}))^{3/2}] \\ &\leq n^{3/2} [(1-\alpha)^{3/2} + \alpha^{3/2} (1 + \beta/\alpha (\sqrt{n}))^2] \\ &\leq n^{3/2} [(1-\alpha)^{3/2} + \alpha^{3/2} (1 + 3\beta/(\alpha \sqrt{n}))] \\ &\leq [\alpha^{3/2} + (1-\alpha)^{3/2}] n^{3/2} + 3\beta \sqrt{\alpha n} \end{aligned}$$

since $\alpha \geq \frac{1}{2}$ implies $\beta(\alpha \sqrt{n}) \leq \beta/(\alpha \sqrt{n_0}) \leq \beta/((1-\alpha)\sqrt{n_0}) = 1$.

Also

$$(22) \quad \begin{aligned} l_1 n_1 + l_2 n_2 &\leq (l + 2\beta \sqrt{n})(\alpha n + \beta \sqrt{n}) \\ &\leq \alpha l n + 2\alpha \beta n^{3/2} + \beta l \sqrt{n} + 2\beta^2 n \end{aligned}$$

and

$$(23) \quad \begin{aligned} l_1^2 \sqrt{n_1} + l_2^2 \sqrt{n_2} &\leq (l + 2\beta \sqrt{n})^2 \sqrt{\alpha n + \beta \sqrt{n}} \\ &\leq (l + 2\beta \sqrt{n})^2 (\sqrt{\alpha n} + \beta/(2\sqrt{\alpha})) \\ &\leq \sqrt{\alpha} l^2 \sqrt{n} + 4\beta \sqrt{\alpha} l n + 4\beta^2 \sqrt{\alpha} n^{3/2} + (l + 2\beta \sqrt{n})^2 \beta/(2\sqrt{\alpha}). \end{aligned}$$

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

$$(24) \quad \begin{aligned} g(l, n) &\leq (c_{11} + c_8 \delta + 2c_9 \alpha \beta + 4c_{10} \beta^2 \sqrt{\alpha}) n^{3/2} \\ &\quad + (c_{12} + c_9 \alpha + 4c_{10} \beta \sqrt{\alpha}) l n + (c_{13} + c_{10} \sqrt{\alpha}) l^2 \sqrt{n} \\ &\quad + c_{14} (c_8 + c_9 + c_{10}) n + c_{14} (c_9 + c_{10}) l \sqrt{n} + c_{14} c_{10} l^2, \end{aligned}$$

where c_{14} is a suitably large constant depending only on α and β .

Suppose we choose n_4 large enough so that $n > n_4$ implies $c_{14} \leq \max \{ (1-\delta)n^{1/2}/2, (1-\alpha)n^{1/2}/2, (1-\sqrt{\alpha})n^{1/2}/2 \}$, choose c_{10} large enough so that $c_{13} + c_{10}(1+\sqrt{\alpha})/2 \leq c_{10}$, choose c_9 large enough so that $c_{12} + c_9(1+\alpha)/2 + c_{10}(c_{14} + 4\beta\sqrt{\alpha}) \leq c_9$, and choose c_8 large enough so that $c_{11} + c_8(1+\delta)/2 + c_{14}(c_9 + c_{10}) + 2c_9 \alpha \beta + 4c_{10} \beta^2 \sqrt{\alpha} \leq c_8$. Then $g(l, n) \leq c_8 n^{3/2} + c_9 l n + c_{10} l^2 \sqrt{n}$ as desired, and the claim follows by induction. $\square$

THEOREM 4. *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 $c_7 n^{3/2} + O(n(\log n)^2)$, where $c_3 \leq 129$ and $c_7 \leq 4002$. Such an ordering can be found in $O(n \log n)$ time.*

Proof. By Corollary 2 of [16], planar graphs satisfy a $\sqrt{n}$ -separator theorem with $\alpha = \frac{2}{3}$ 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. $\square$

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 5. *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 [16], any n -vertex finite element graph with no element having more than k boundary vertices satisfies a $\sqrt{n}$ -separator theorem with $\alpha = \frac{2}{3}$ 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. $\square$

Although planar and almost-planar graphs seem to be the most interesting case, analogues to Theorems 2–5 hold for other classes of graphs. For instance, the following theorems can be proved using the same methods as in the proofs of Theorems 2–5.

THEOREM 6. *Let S be any class of graphs closed under subgraph on which an n^σ separator theorem holds for $\sigma > \frac{1}{2}$. 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 d -dimensional hypercubic grid graphs satisfies Theorem 6 for $\sigma = d - 1/d$.

THEOREM 7. *Let S be any class of graphs closed under subgraph on which an n^σ separator theorem holds for $\frac{1}{3} < \sigma < \frac{1}{2}$. 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 8. *Let S be any class of graphs closed under subgraph on which a $\sqrt{n}$ 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 9. *Let S be any class of graphs closed under subgraph on which an n^σ separator theorem holds for $\sigma < \frac{1}{3}$. 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].

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

Proof. G must have a connected component containing at least $n/3$ vertices. Otherwise there is a set A violating the hypothesis of the lemma, formed as follows. Let

$\emptyset$. Add connected component vertices. Then A contains less than

Let $\pi(v_i) = i$ for $1 \leq i \leq n$ and components of the subgraph of G in smallest integer such that $|C_i^k| \geq$ paragraph. Then $|C_1^{k-1}|, |C_2^{k-1}|,$ in v_k , we can choose q and th

$n/3$

v_k adjacent to some vertex in

Let $A = \bigcup_{i=1}^q C_i^{k-1}$. Let C be in A . By the hypothesis of $(v) \geq k$. Any two vertices v_i are both adjacent to at least $k - \{v_k\}$ are adjacent in G_π^* , since vertex in C_i^{k-1} . Thus C forms a weaker form of Lemma 2, indeed, appears in [5].

THEOREM 10. *Let $G = (V, E)$ any ordering of V produces a of at least $l(l-1)(l+4)/6$.*

Proof. The proof is immediate. Theorem 10 and the results in method of sparse elimination (space) on large classes of ded from [16].



FIG.

THEOREM 11. *For any k , let subset of V such that $n/3 \leq$ adjacent to at least one vertex*

Proof. Without loss of generality vertices in A is at least as in A . Then $n/3 \leq |A| \leq n$ r^* be the number of $|A| \leq 2n/3$, and $r^* \leq 2k/3$. $|C| \geq r - r^* = k - r^* \geq k/3$

G has an elimination ordering with a multiplication count of $O(k^3 n^{3/2})$. Such an ordering can be found in $O(n^2)$ time.

to satisfy a $\sqrt{n}$ -separator theorem. The vertex partition can be found in $O(n^2)$ time. This gives the result. $\square$

from a planar embedding of a graph. (The finite element graph is a planar graph.) The embedding of the graph and each of its faces is a cycle.

element graph with no element graph. Such an ordering can be found in $O(n^2)$ time.

finite element graph with no element graph. Such an ordering can be found in $O(n^2)$ time. This gives the result. $\square$

them to be the most interesting graphs. For instance, the following is in the proofs of Theorems 1 and 2. Let G be a graph with n vertices and m edges. Let S be a subset of V of size $|S| = n/3$. Let C be the set of vertices in $V - S$ adjacent to at least one vertex in S . Then $|C| \geq n/3$. Furthermore, each element $v \in C$ has $\pi(v) \geq k$. Any two vertices $v, w \in C - \{v_k\}$ are adjacent in G_π^* by Lemma 1, since they are both adjacent to at least one vertex in C_j^k . Similarly v_k and any vertex $w \in C - \{v_k\}$ are adjacent in G_π^* , since for some i both v_k and w are adjacent to at least one vertex in C_i^{k-1} . Thus C forms a clique in G_π^* . $\square$

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

THEOREM 10. Let $G = (V, E)$ be a graph satisfying the hypothesis of Lemma 2. Then any ordering of V produces a fill-in of size at least $l(l-1)/2$ and a multiplication count of at least $l(l-1)(l+4)/6$.

Proof. The proof is immediate from Lemma 2. $\square$

Theorem 10 and the results in § 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 we have the following result, adapted from [16].




FIG. 1. A 5×5 square grid graph.

THEOREM 11. For any k , let $G = (V, E)$ be a $k \times k$ square grid graph (Fig. 1). Let A be any subset of V such that $n/3 \leq |A| \leq 2n/3$, where $n = k^2$. Let C be the set of vertices in $V - A$ adjacent to at least one vertex in A . Then $|C| \geq \sqrt{n}/3$.

Proof. Without loss of generality, suppose that the number r of rows of G which contain vertices in A is at least as large as the number c of columns of G which contain vertices in A . Then $n/3 \leq |A| \leq rc \leq r^2$ and $r \geq \sqrt{n}/3$.

Let r^* be the number of rows of G which contain only vertices in A . Then $r^* \leq |A| \leq 2n/3$, and $r^* \leq 2k/3$. If $r^* = 0$, then $|C| \geq r \geq \sqrt{n}/3$. If $r^* \neq 0$, then $k \geq r \geq c = k$ and $|C| \geq r - r^* = k - r^* \geq k/3 = \sqrt{n}/3$. $\square$

By Theorems 10 and 11, $k \times k$ square grid 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, Theorem 10 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 6. 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 12. *Let S be any class of graphs closed under subgraph and satisfying an $n/(\log_2 n)^{1+\epsilon}$ -separator theorem for fixed α , β , and $\epsilon > 0$. If G is a graph in S having n vertices and m edges, then m is $O(n)$.*

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 $n/(\log_2 n)^{1+\epsilon}$ -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 $\beta n/(\log_2 n)^{1+\epsilon}$ 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 $t(n) \leq t(n_1) + t(n_2)$. Hence for any n_3 ,

$$(25) \quad \begin{aligned} t(n) &\leq n(n-1)/2 && \text{if } n \leq n_3, \\ t(n) &\leq \max \{t(n_1) + t(n_2)\} && \text{otherwise,} \end{aligned}$$

where the maximum is taken over values n_1, n_2 satisfying

$$(26) \quad \begin{aligned} n &\leq n_1 + n_2 \leq n + \beta n/(\log_2 n)^{1+\epsilon}, \\ (1-\alpha)n &\leq n_i \leq \alpha n + \beta n/(\log_2 n)^{1+\epsilon} \quad \text{for } i = 1, 2. \end{aligned}$$

We shall show by induction that

$$(27) \quad t(n) \leq c_{15}n - c_{16}n/(\log_2 n)^\epsilon \quad \text{for } n \geq n_0,$$

where c_{15}, c_{16} , and n_0 are suitably large positive constants. The theorem follows.

Let n_0 be large enough so that $(\log_2 n_0)^\epsilon \geq 2\beta/\epsilon$. Choose $n_3 \geq n_0$ large enough so that $(1-\alpha)n_3 \geq n_0$ and $\alpha + \beta/(\log_2 n_3)^{1+\epsilon} < 1$. Choose $c_{15} = n_3 - 1$ and $c_{16} = c_{15}\beta/\epsilon$. Then if $n_0 \leq n \leq n_3$,

$$(28) \quad \begin{aligned} t(n) &\leq n(n-1)/2 \leq c_{15}n/2 \leq c_{15}n - c_{16}n/(\log_2 n_0)^\epsilon \\ &\leq c_{15}n - c_{16}n/(\log_2 n)^\epsilon. \end{aligned}$$

Thus (27) holds for $n_0 \leq n \leq n_3$.

Let $n > n_3$ and suppose (27) holds for values between n_0 and $n-1$. Then $t(n) \leq t(n_1) + t(n_2)$ for some values of n_1, n_2 satisfying (26). We have $n_1, n_2 \geq (1-\alpha)n \geq (1-\alpha)n_3 \geq n_0$. Also, $n_1, n_2 \leq \alpha n + \beta n/(\log_2 n)^{1+\epsilon} \leq n(\alpha + \beta/(\log_2 n)^{1+\epsilon}) \leq n(\alpha + \beta/(\log_2 n_3)^{1+\epsilon}) < n$. Thus, by the induction hypothesis,

$$(29) \quad t(n) \leq c_{15}n_1 - c_{16}n_1/(\log_2 n_1)^\epsilon + c_{15}n_2 - c_{16}n_2/(\log_2 n_2)^\epsilon.$$

By (26), $c_{15}n_1 + c_{15}n_2 \leq c_{15}n + c_{15}\beta n/(\log_2 n)^{1+\epsilon}$. Thus

$$(30) \quad t(n) \leq c_{15}n + c_{15}\beta n/(\log_2 n)^{1+\epsilon} - c_{16}(n_1/(\log_2 n_1)^\epsilon + n_2/(\log_2 n_2)^\epsilon).$$

It remains for us to show that

$$(31) \quad c_{15}\beta n/(\log_2 n)^{1+\epsilon} - c_{16}(n_1/(\log_2 n_1)^\epsilon + n_2/(\log_2 n_2)^\epsilon) \leq -c_{16}n/(\log_2 n)^\epsilon,$$

$$(32) \quad c_{15}\beta n/(\log_2 n)^{1+\epsilon} \leq c_{16}n/(\log_2 n)^\epsilon$$

The right-hand side of (32) is only assumed that $n_1 + n_2 = n$. For fixed n , the right-hand side is minimized when n_1 and n_2 are equal.

$$c_{16}(n_1/(\log_2 n_1)^\epsilon + n_2/(\log_2 n_2)^\epsilon)$$

$$\geq c_{16}n/(\log_2 n)^\epsilon$$

$$= c_{16}n/(\log_2 n)^\epsilon$$

$$\geq c_{16}n/(\log_2 n)^\epsilon$$

$$\geq c_{16}n/(\log_2 n)^\epsilon$$

$$c_{16} = c_{15}\beta/\epsilon. \quad \square$$

Not all sparse graphs have good separators.

Let α be a constant with $0 < \alpha < 1$, there is a constant c such that almost all sparse graphs require a more direct argument, we can prove the following theorem.

THEOREM 13. *For all $\epsilon > 0$ there is a constant $c(\epsilon)$ such that almost all graphs with at least $c(\epsilon)n$ edges have a good separator.*

Proof. We first prove that almost all graphs have the following property:

If A and B are sets of vertices with $|A|, |B| \geq \epsilon n$ and $A \cap B = \emptyset$, then at least $c(\epsilon)n$ edges have one endpoint in A and the other in B .

We prove (P) by an argument like the one used in the proof of Theorem 12. Let G be a graph with n vertices and m edges. Let A and B be two vertex sets of G satisfying (P). Then there are at least $c(\epsilon)n$ edges between A and B . The edges actually occurs in G is at least $c(\epsilon)n$.

Let $n_1 = |A|$ and $n_2 = |B|$. Then $n_1, n_2 \geq \epsilon n$. The edges actually occurs in G is at least $c(\epsilon)n$.

Now we use (P) to prove the theorem. Let A be a set of at least $3\epsilon n/4$ vertices whose induced subgraph is sparse. Then there is a vertex v in A such that $|A \cap N(v)| \geq \epsilon n/4$. Then $\sum_{i=1}^k |A_i| \leq \epsilon n$ for any partition of A into k sets A_i with some vertex v in A_i .

Consider any ordering of the vertices of G . Let B be a subset of A of size $|B| \geq \epsilon n/4$.

By "almost all" we mean that the fraction of graphs for which this is true is at least $1 - o(1)$. We assume that each n -vertex graph has at least $c(\epsilon)n$ edges unless their edge sets are identical. S

hs have an $\Omega(n^{3/2})$ multipli-
ated techniques, one can de-
is [13]. For d -dimensional
 d) lower bound on fill-in
being with the upper bound

separators and sparsity. O-
rators.

d under subgraph and satisfy
> 0. If G is a graph in S ha-

s in any n -vertex graph in S
atisfies an $n/(\log_2 n)^{1+\epsilon}$ -sep-
three sets A, B, C such
an αn vertices, and C conta-
nder subgraph, the subgraph
both in S . If $|A \cup C| = n$
e for any n_3 ,

if $n \leq n_3$,

otherwise,

fying

$)^{1+\epsilon}$,

$2n)^{1+\epsilon}$ for $i = 1, 2$.

for $n \geq n_0$,

stants. The theorem follows

. Choose $n_3 \geq n_0$ large enough

use $c_{15} = n_3 - 1$ and $c_{16} = c_{15}$

$c_{16}n/(\log_2 n_0)^\epsilon$

$c_{16}n/(\log_2 n)^\epsilon$.

s between n_0 and $n-1$.

fying (26). We have n_1

$2n)^{1+\epsilon} \leq n(\alpha + \beta/(\log_2 n)^{1+\epsilon})$

thesis,

$c_{16}n_2/(\log_2 n_2)^\epsilon$.

$n_1)^\epsilon + n_2/(\log_2 n_2)^\epsilon$.

$n_2)^\epsilon) \leq -c_{16}n/(\log_2 n)^\epsilon$,

i.e.

$$(32) \quad c_{15}\beta n/(\log_2 n)^{1+\epsilon} \leq c_{16}(n_1/(\log_2 n_1)^\epsilon + n_2/(\log_2 n_2)^\epsilon - n/(\log_2 n)^\epsilon).$$

The right-hand side of (32) is only reduced by making n_1 and n_2 smaller; thus we can assume that $n_1 + n_2 = n$. For fixed $n_1 + n_2$, the function $n_1/(\log_2 n_1)^\epsilon + n_2/(\log_2 n_2)^\epsilon$ is minimized when n_1 and n_2 are equal. Thus

$$\begin{aligned} & c_{16}(n_1/(\log_2 n_1)^\epsilon + n_2/(\log_2 n_2)^\epsilon - n/(\log_2 n)^\epsilon) \\ & \geq c_{16}(n/(\log_2 n - 1)^\epsilon - n/(\log_2 n)^\epsilon) \\ (33) \quad & = c_{16}n((1/(1 - 1/\log_2 n))^\epsilon - 1)/(\log_2 n)^\epsilon \\ & \geq c_{16}n(1 + 1/\log_2 n)^\epsilon - 1)/(\log_2 n)^\epsilon \\ & \geq c_{16}\epsilon n/(\log_2 n)^{1+\epsilon} = c_{15}\beta n/(\log_2 n)^{1+\epsilon} \end{aligned}$$

since $c_{16} = c_{15}\beta/\epsilon$. $\square$

Not all sparse graphs have good separators. In fact, for fixed α, β such that $\beta < 1 - \alpha \leq \alpha < 1$, there is a constant c such that almost all n -vertex graphs with cn edges have no vertex partition A, B, C satisfying $|A|, |B| \leq \alpha n$, $|C| \leq \beta n$, and C separates A and B . This result is implicit in Theorem 4 of [8]. It follows from Theorem 10 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 13. *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 $m \geq cn$. 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 $(1 - 2c/n)^{\epsilon^2 n^2/16}$. Thus, if c is chosen so that $3^n(1 - 2c/n)^{\epsilon^2 n^2/16} \rightarrow 0$ as $n \rightarrow \infty$, then almost all graphs satisfy (P). Since $(1 - 2c/n)^{\epsilon^2 n^2/16} \rightarrow e^{-c\epsilon^2 n/8}$, choosing $c > (8 \log_e 3)/\epsilon^2$ 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 $A_1, A_2, \dots, A_k$ 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

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 $\{1, 2, \dots, n\}$ and that two graphs are distinct unless their edge sets are identical. See [7] for a thorough discussion of random graphs.

subgraph in G is connected. By property (P) at least $(1 - \varepsilon/2)n$ vertices in $V - B$, and hence at least $(1 - \varepsilon)n$ 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 fill-in clique of at least $(1 - \varepsilon)n$ vertices. $\square$

THEOREM 14. *Almost all n -vertex graphs with $c(\varepsilon)n$ edges have a fill-in of $(1 - \varepsilon)^2 n^2/2 - O(n)$ and a multiplication count of $(1 - \varepsilon)^3 n^3/6 - O(n^2)$, for any ordering.*

Proof. Immediate from Theorem 13. $\square$

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 nonzeros 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], [23]. 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 § 2.

Gaussian elimination can be used to solve systems of linear equations defined over algebras other than the real numbers [2], [4], [24], and the algorithm in § 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 § 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 10 in § 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 "usefully 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 [17].

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 $v = v_0, v_1, v_2, \dots, v_k = w$ such that $\{v_{i-1}, v_i\}$ is an edge for $1 \leq i \leq k$. 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 $G = (V, E)$

is a graph and $V_1 \subseteq V_2$, the graph G_1 is a subgraph of G_2 induced by the vertices V_1 . A graph G is connected if every pair of distinct vertices is joined by a path. The connected components of G are the maximal connected subgraphs of G .

Let A, B, C be a partition of the vertices of G . Let A and B be a partition of the vertices of A and B if no edge joins a vertex in A to a vertex in B .

If f and g are functions of n , $f(n) = O(g(n))$ if there is a constant c , $f(n) \leq cg(n)$ for all but finitely many n . $f(n)$ is $O(f(n))$.

A graph $G = (V, E)$ is planar if it can be drawn in the plane and a map f_2 from E into the plane such that $f_2(\{v, w\})$ has endpoints v and w and $f_2(\{v_1, w_1\})$ and $f_2(\{v_2, w_2\})$ share a point except possibly at v or w . A planar embedding of G is a drawing of G in the plane such that f_2 are deleted from the plane and the boundary of a curve corresponding to f_2 is sometimes not distinguishable from the edge (v, w) such that v and w are adjacent.

A. V. AHO, J. E. HOPCROFT AND J. E. ULLMAN, *Introduction to the Theory of Finite Automata*, Addison-Wesley, Reading, MA, 1975.

R. C. BACKHOUSE AND B. A. CARROLL, *SIAM J. Appl. Math.*, 15 (1975), pp. 161-186.

J. BUNCH AND J. E. HOPCROFT, *SIAM J. Math. Comput.*, 28 (1974), pp. 2-11.

B. A. CARRÉ, *An algebra for network analysis*, Academic Press, New York, 1976, pp. 85-91.

S. C. EISENSTAT, M. H. SCHULTZ AND J. E. HOPCROFT, *Gaussian elimination*, Sparse Matrices, Academic Press, New York, 1976, pp. 85-91.

———, *Considerations in the design of algorithms for sparse matrix computations*, J. R. Bunch and J. E. Hopcroft, Eds., Academic Press, New York, 1976, pp. 1-13.

P. ERDŐS AND J. SPENCER, *Probability in Combinatorics*, Cambridge University Press, Cambridge, 1979.

P. ERDŐS, R. L. GRAHAM AND E. S. SHOSTAK, *Journal of the American Mathematical Society*, 1, 3/4 (1975), pp. 1047-1051.

G. E. FORSYTHE AND C. B. MOLER, *Solution of Sparse Linear Systems*, Englewood Cliffs, NJ, 1967.

A. GEORGE AND J. W. H. LIU, *Algorithms for sparse matrix problems*, this Journal, 15 (1978), pp. 1-13.

J. A. GEORGE, *Nested dissection of a sparse matrix*, this Journal, 15 (1978), pp. 1-13.

F. G. GUSTAVSON, *Some basic techniques for sparse matrix computations and Their Applications*, D. J. Rose, Ed., Academic Press, New York, 1974, pp. 41-52.

A. J. HOFFMAN, M. S. MARTIN AND J. E. HOPCROFT, *Finite element grids*, this Journal, 15 (1978), pp. 1-13.

D. B. JOHNSON, *Efficient algorithms for sparse matrix computations*, (1977), pp. 1-13.

R. J. LIPTON, D. J. ROSE AND R. E. TARJAN, *Lower bounds for the complexity of sparse matrix operations*, STAN-CS-77-645, Computer Science Department, Stanford University, 1977.

R. J. LIPTON AND R. E. TARJAN, *A simple linear time algorithm for the complexity of sparse matrix operations*, Computer Science, University of California, Berkeley, 1979, pp. 177-199.

———, *Applications of a planar separator theorem*, Computer Science, 1977, pp. 16-23.

least $(1 - \epsilon/2)n$ vertices in $V - B$, adjacent to at least one vertex in B as a fill-in edge. Thus the set of vertices is a fill-in clique of at least $(1 - \epsilon)n^2/6 - O(n^2)$, for any order.

the existence of an $O(n^{3/2})$ -time Gaussian elimination on systems of linear equations defined on planar or two-dimensional finite elements. The practicality of the algorithm for large systems, since the worst-case is pessimistic. Experiments by George and his method is practical.

our algorithm to $O(n^{\log_2 7})$ by using matrix factorization [3], [23]. If the system is banded, it is possible to reduce the work of L and recomputing the residuals. These savings in the case of ordering almost-planar graphs is analogous

systems of linear equations defined on planar graphs. The best general

of good separators in a graph and Gaussian elimination is efficient. On which Gaussian elimination is efficient. Good separators in a graph and most all sparse graphs do not have good separators. One should use Gaussian elimination, one should use separators rather than when it has

d. Can generalized nested dissection find good separators? Informally, "goodness" of a separator? Informally, a graph into small pieces. We need to know the trade-offs between complexity and the property of having good separators but in many other problems [1]

consists of a set V of vertices and a set of distinct vertices. If $\{v, w\}$ is an edge, v and w are the endpoints of the edge. v is a sequence of vertices v_i for $1 \leq i \leq k$. If $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$. If $G = (V, E)$

is a graph and $V_1 \subseteq V_2$, the graph $G_1 = (V_1, E_1)$ where $E_1 = E_2 \cap \{\{v, w\} | v, w \in V_1\}$ 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 $f_1(w)$, 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 nonadjacent 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, MA, 1974.
- [2] R. C. BACKHOUSE AND B. A. CARRÉ, *Regular algebra applied to pathfinding problems*, J. Inst. Math. Appl., 15 (1975), pp. 161-186.
- [3] J. BUNCH AND J. E. HOPCROFT, *Triangular factorization and inversion by fast matrix multiplication*, Math. Comput., 28 (1974), pp. 231-236.
- [4] B. A. CARRÉ, *An algebra for network routing problems*, J. Inst. Math. Applic., 7 (1971), pp. 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, pp. 85-96.
- [6] —, *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, pp. 263-273.
- [7] P. ERDŐS AND J. SPENCER, *Probabilistic Methods in Combinatorics*, Academic Press, New York, 1974.
- [8] P. ERDŐS, R. L. GRAHAM AND E. SZEMERÉDI, *On sparse graphs with dense long paths*, Computers and Mathematics, 1, 3/4 (1975), pp. 365-370.
- [9] G. E. FORSYTHE AND C. B. MOLER, *Computer Solution of Linear Algebraic Equations*, Prentice-Hall, Englewood Cliffs, NJ, 1967.
- [10] A. GEORGE AND J. W. H. LIU, *An automatic nested dissection algorithm for irregular finite element problems*, this Journal, 15 (1978), pp. 1053-1069.
- [11] J. A. GEORGE, *Nested dissection of a regular finite element mesh*, this Journal, 10 (1973), pp. 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, pp. 41-52.
- [13] A. J. HOFFMAN, M. S. MARTIN AND D. J. ROSE, *Complexity bounds for regular finite difference and finite element grids*, this Journal, 10 (1973), pp. 364-369.
- [14] D. B. JOHNSON, *Efficient algorithms for shortest paths in sparse networks*, J. Assoc. Comput. Mach., 24 (1977), pp. 1-13.
- [15] R. J. LIPTON, D. J. ROSE AND R. E. TARJAN, *Generalized nested dissection*, Technical Report STAN-CS-77-645, Computer Science Dept., Stanford University, Stanford, CA, 1977.
- [16] R. J. LIPTON AND R. E. TARJAN, *A separator theorem for planar graphs*, Proc. of Conf. on Theoretical Computer Science, University of Waterloo, Waterloo, Ontario, 1977, pp. 1-10; SIAM J. Appl. Math., 36 (1979), pp. 177-199.
- [17] —, *Applications of a planar separator theorem*, Proc. Eighteenth Annual Symp. on Foundations of Computer Science, 1977, pp. 162-170.

- [18] D. J. ROSE, *Triangulated graphs and the elimination process*, J. Math. Anal. Appl., 32 (1970), pp. 597-609.
- [19] ———, *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, pp. 183-217.
- [20] D. J. ROSE, R. E. TARJAN AND G. S. FUEKER, *Algorithmic aspects of vertex elimination on graphs*, SIAM J. Comput., 5 (1976), pp. 266-283.
- [21] D. J. ROSE AND G. F. WHITTEN, *A recursive analysis of dissection strategies*, Sparse Matrix Computations, J. R. Bunch and D. J. Rose, eds., Academic Press, New York, 1976, pp. 59-84.
- [22] A. H. SHERMAN, *On the efficient solution of sparse systems of linear and nonlinear equations*, Ph.D. thesis, Computer Science Dept., Yale University, New Haven, CT, 1975.
- [23] V. STRASSEN, *Gaussian elimination is not optimal*, Numer. Math., 13 (1969), pp. 354-356.
- [24] R. E. TARJAN, *Solving path problems on directed graphs*, Technical Report STAN-CS-75-528, Computer Science Dept., Stanford University, Stanford, CA, 1975.
- [25] ———, *Graph theory and Gaussian elimination*, Sparse Matrix Computations, J. R. Bunch and D. J. Rose, eds., Academic Press, New York, 1976, pp. 3-22.
- [26] O. C. ZIENKIEWICZ, *The Finite Element Method in Engineering Science*, McGraw-Hill, London, 1970.

AN ALGORITHM FOR
BY BLOC

CONNICE A. E

Abstract. This paper describes an algorithm for a similarity transformation. The columns of the transformation are in the subspace of A , and the block is the representation of the transformation. The algorithm attempts to control the condition number of the transformation, so that the well conditioned and the basis vectors are

1. Introduction. The purpose of this paper is to describe an algorithm for reducing a real matrix A of order n to a block triangular form. Specifically, the algorithm finds a matrix X such that $X^{-1}AX$ has the form

$$(1.1) \quad X^{-1}AX = \begin{pmatrix} B_1 & & \\ & \ddots & \\ & & B_k \end{pmatrix}$$

where each matrix B_i is square of order n_i . This is a block triangular decomposition such as (1.1). For small, powers of A can be economical.

$$A^k = X D X^{-1}$$

and this fact can be used to simplify the computation of powers of A (e.g. see [7]). If X is partitioned

$$X = \begin{pmatrix} X_1 \\ X_2 \\ \vdots \\ X_k \end{pmatrix}$$

where each X_i has n_i columns, then the subspace of A , and B_i is the representation of the spectral projector is given by the corresponding rows of X^{-1} (for details see [7]).

There are theoretical and practical considerations. Theoretically, they can be no smaller than the eigenvalues. Practically, they may have to be chosen so that the decompositions such as (1.1) have a good numerical summary.

The principal difficulty is that the eigenvalues are not determined; very small perturbations can cause them to coalesce. Any attempt to separate the eigenvalues by a similarity transformation matrix X whose columns are the eigenvectors of A will be ill-conditioned in the sense of the Frobenius matrix norm). In this case, the eigenvalues of X are not accurate [1]. The eigenvalues; but there are some cases, that can be split completely into real eigenvalues that cannot be split.

Received by the editors October 28, 1978.
Research under Contract N00014-76-C-0001.
Department of Computer Science,