Numer. Math. 50, 377–404 (1987)

Numerische
Mathematik

The Analysis of a Nested Dissection Algorithm

John R. Gilbert[1,*] and Robert Endre Tarjan[2,**]

[1] Computer Science Department, Cornell University, Ithaca, NY 14853, USA
[2] Computer Science Department, Princeton University, Princeton, NJ 08544, USA

Summary. Nested dissection is an algorithm invented by Alan George for preserving sparsity in Gaussian elimination on symmetric positive definite matrices. Nested dissection can be viewed as a recursive divide-and-conquer algorithm on an undirected graph; it uses *separators* in the graph, which are small sets of vertices whose removal divides the graph approximately in half. George and Liu gave an implementation of nested dissection that used a heuristic to find separators. Lipton and Tarjan gave an algorithm to find $n^{1/2}$-separators in planar graphs and two-dimensional finite element graphs, and Lipton, Rose, and Tarjan used these separators in a modified version of nested dissection, guaranteeing bounds of $O(n \log n)$ on fill and $O(n^{3/2})$ on operation count. We analyze the combination of the original George-Liu nested dissection algorithm and the Lipton-Tarjan planar separator algorithm. This combination is interesting because it is easier to implement than the Lipton-Rose-Tarjan version, especially in the framework of existing sparse matrix software. Using some topological graph theory, we prove $O(n \log n)$ fill and $O(n^{3/2})$ operation count bounds for planar graphs, two-dimensional finite element graphs, graphs of bounded genus, and graphs of bounded degree with $n^{1/2}$-separators. For planar and finite element graphs, the leading constant factor is smaller than that in the Lipton-Rose-Tarjan analysis. We also construct a class of graphs with $n^{1/2}$-separators for which our algorithm does not achieve an $O(n \log n)$ bound on fill.
Subject Classifications: AMS(MOS): 05C10, 65F05, 65F50, CR: G.1.3, G.2.2

1. Introduction

Suppose that we want to solve a sparse system of n linear equations in n unknowns,

* The work of this author was supported in part by the Hertz Foundation under a graduate fellowship and by the National Science Foundation under Grant MCS 82-02948
** The work of this author was supported in part by the National Science Foundation under Grant MCS 78-26858 and by the Office of Naval Research under Contract N00014-76-C-0688

$$M\mathbf{x}=\mathbf{b},$$

where M is an n by n symmetric, positive definite matrix. We can use a version of Gaussian elimination to find the Cholesky factorization $M=LL^T$, where L is a lower triangular matrix with positive diagonal. We then solve for $\mathbf{x}$ by solving the two triangular systems $L\mathbf{y}=\mathbf{b}$ and $L^T\mathbf{x}=\mathbf{y}$.

The complexity of this procedure depends on the sparsity of the matrices M and L. Suppose column j of L contains d_j nonzeros. Using algorithms and data structures described in George and Liu [4], we can factor and solve the system in space proportional to $\sum_j d_j$ (which is the number of nonzeros in L) and time proportional to $\sum_j d_j^2$. Ignoring cancellation due to numerical coincidence, L will have nonzeros below the diagonal everywhere that M does, and also some other places. We define the *fill* to be the set of below-diagonal positions in which L is nonzero and M is zero.

If P is a permutation matrix, PMP^T is a symmetric, positive definite matrix obtained by permuting the rows and columns of M. The fill in the triangular factor of PMP^T may be drastically different for different choices of P. We can think of P as a choice of an order in which to eliminate the variables of the system.

Finding the order that gives the smallest possible fill is an NP-complete problem [22]. Most sparse matrices do not have elimination orders with small fill: For any positive ε there is a constant $c(\varepsilon)$ such that almost all n by n symmetric matrices with $c(\varepsilon)n$ nonzeros have at least $(1-\varepsilon)^2n^2/2-O(n)$ fill for every order [14]. Gilbert [5] presents a class of symmetric matrices with four nonzeros per row that have $\Theta(n^2)$ fill for every elimination order.

Although the outlook is gloomy for general elimination algorithms with low fill, good elimination orders can be found for some classes of problems. George [2] invented an algorithm called *nested dissection* for ordering the variables in a system that comes from finite differences on a regular square grid. George and Liu [3] gave a heuristic nested dissection algorithm for general matrices. Lipton, Rose, and Tarjan used the planar separator theorem [14, 15] in a modified version of the George-Liu algorithm, which they called *generalized nested dissection*. Their order gives $O(n\log n)$ fill and $O(n^{3/2})$ operation count on any system whose subgraphs have $n^{1/2}$-separators, which includes planar graphs and two-dimensional finite element meshes. (See Sect. 2 for definitions.) These bounds are within a constant factor of the best possible.

In this paper, we analyze what happens when the planar separator theorem is applied to the original George-Liu algorithm. This combination is easier to implement than the LRT algorithm, since it fits nicely into George and Liu's Sparspak package [4]. We prove $O(n\log n)$ fill and $O(n^{3/2})$ operation count bounds on a large class of graphs, including planar graphs and two-dimensional finite element meshes but not including all graphs whose subgraphs have $n^{1/2}$-separators. For planar graphs, the leading constant in the analysis is smaller than that in Lipton, Rose, and Tarjan's analysis of the LRT algorithm. The analysis itself is interesting because it uses more topological information than does the LRT analysis.

The remainder of the paper is organized as follows. Section 2 gives graph-theoretic definitions and lemmas. Section 3 presents the nested dissection algorithm we analyze in the rest of the paper. Sections 4 through 7 give upper bounds on the performance of this algorithm. Section 8 discusses how tight this analysis is. Section 9 presents two examples to show that the hypotheses of the fill bounds are necessary. Section 10 consists of remarks and conclusions. The appendix solves a recurrence relation that is used in the analysis.

We use the following notation to describe the asymptotic behavior of nonnegative functions of nonnegative integers. We say $f(n)=O(g(n))$ if there is a constant c such that $f(n)\leq cg(n)$ for all but finitely many n. We say $f(n)=\Omega(g(n))$ if $g(n)=O(f(n))$. We say $f(n)=\Theta(g(n))$ if $f(n)=O(g(n))$ and $g(n)=O(f(n))$. We use $\lg x$ to denote the logarithm of x to the base 2.

A preliminary version of this paper has appeared as a technical report [5, Sects. 2.1–2.11].

2. Graph Theory and Gaussian Elimination

We shall state and analyze our algorithm by using a graph theoretic model that was proposed by Parter [18] and studied in detail by Rose [20]. Recently this model has seen wide use [4, 21].

Let $M=(m_{ij})$ be an n by n symmetric, positive definite matrix. The graph $G=G(M)$ associated with M is an undirected graph with vertex set $V=\{v_1, \dots, v_n\}$ and edge set

$$E=\{\{v_i, v_j\}: i\neq j \text{ and } m_{ij}\neq 0\}.$$

Thus G has a vertex for each variable in the system $M\mathbf{x}=\mathbf{b}$ and an edge for each symmetric pair of nonzero coefficients. Figure 1 shows a matrix and its associated graph.

If we use the i-th equation to eliminate the i-th variable from the system $M\mathbf{x}=\mathbf{b}$ - that is, we pivot on m_{ii} - then the $n-1$ by $n-1$ matrix of the coefficients of the remaining variables in the remaining equations is still symmetric and positive definite. Its graph, which has $n-1$ vertices, is obtained from G by first adding edges to make all of v_i's neighbors mutually adjacent, and then deleting v_i and all edges incident on v_i. (Here and henceforth we assume that no zeros are created by numerical cancellation.)

An *elimination order* on G is a permutation of the vertices, which is a bijection π: $\{1, \dots, n\}\rightarrow V$. Reducing G to the null graph by sucessively eliminating vertices $\pi(1), \dots, \pi(n)$ is precisely analogous to performing Gaussian elimination on M, choosing as pivots the diagonal elements that correspond to $\pi(1), \dots, \pi(n)$. Figure 2 is an example. The zeros of M that become nonzero during this elimination correspond to the edges that are added to the graph at each step. These are the *fill* edges, and the number of such edges is the *size* of the fill, or simply the *fill*. The *filled graph* $G^*_\pi(M)$, or just G^*, is the graph obtained from G by adding the fill due to π, as in Fig. 3. The filled graph is the graph of the matrix $L+L^T$.

$$\begin{pmatrix} 9 & 3 & 2 & 1 & \\ 3 & 9 & & 4 & \\ 2 & & 9 & & 2 \\ 1 & 4 & & 9 & \\ & & 2 & & 9 \end{pmatrix}$$

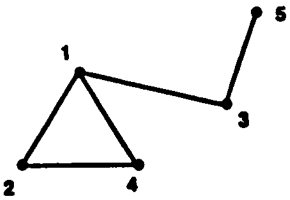

Fig. 1. A matrix M and its graph $G(M)$

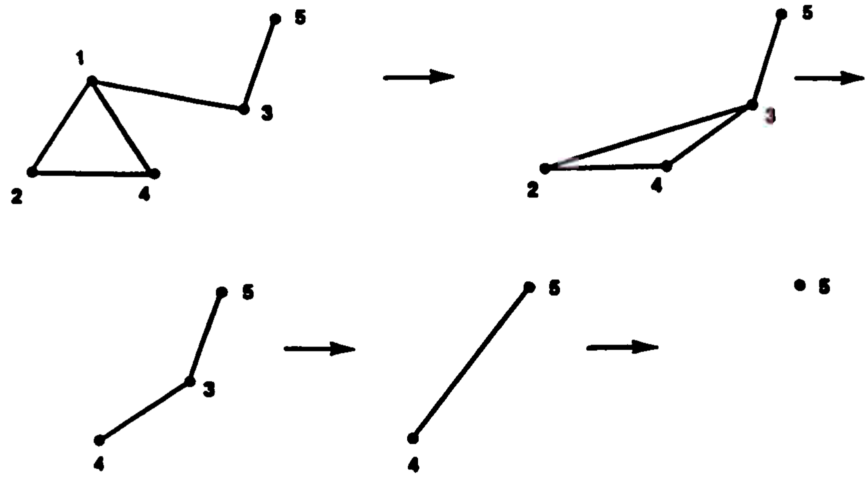

Fig. 2. Eliminating the graph

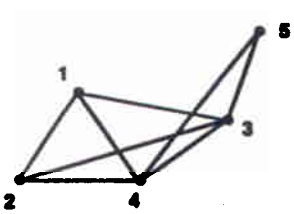

Fig. 3. The filled graph for the matrix in Fig. 1

The problem of finding a permutation of M that gives a sparse factor L is therefore the same as the problem of finding an elimination order for the vertices of G that gives small fill. One way to decide which edges will fill in without actually performing the elimination is given by a lemma of Rose, Tarjan, and Lueker, which says that edge $\{v, w\}$ fills in if and only if there is a path from v to w in G that contains only vertices eliminated earlier than both v and w.

Lemma 1 [21]. *If G is a graph with elimination order π, then $\{v, w\}$ is an edge of G_π^* if and only if there is a path $v = v_1, v_2, \ldots, v_k = w$ in G such that*

$$\pi^{-1}(v_i) < \min\{\pi^{-1}(v), \pi^{-1}(w)\} \quad \textit{for } 1 < i < k. \quad \square$$

Dissection algorithms are based on *separators* in graphs: The idea is to find a set of vertices that separates the graph and eliminate them last. Following Lipton and Tarjan [15], we say that a class S of graphs satisfies an $f(n)$-*separator theorem* for constants $\alpha<1$ and $\beta>0$ if every n-vertex graph in S has a vertex partition $A\cup B\cup C$ such that

$$|A|,\ |B|\leq \alpha n,$$
$$|C|\leq \beta f(n),$$

and no edge has one endpoint in A and the other in B.

Most sparse graphs do not have nontrivial separator theorems (in a sense made precise in Lipton et al. [14]), but some useful classes of graphs do. Separator theorems are known for trees [11], outerplanar graphs [13], graphs of bounded genus [6], hypercubes [5], chordal graphs [7], and several graphs that are useful in parallel computation [10, 13]. Lipton and Tarjan proved a $n^{1/2}$-separator theorem for planar graphs.

Theorem 1 [15]. *Planar graphs satisfy a $n^{1/2}$-separator theorem with constants $\alpha=\frac{2}{3}$ and $\beta=8^{1/2}$.* □

Lipton and Tarjan also gave a linear-time algorithm to find such a separator in a planar graph. Djidjev [1] improved the constant $\beta=8^{1/2}$ to $6^{1/2}$ and gave a lower bound of $(4\pi 3^{1/2}/9)^{1/2}$; the best possible value is not known. Miller [16] proved a $n^{1/2}$-separator theorem for maximal planar graphs in which the separator is a single simple cycle in the graph.

We are most interested in the graphs of matrices that arise when using finite element methods on two-dimensional surfaces. These are planar graphs and *finite element graphs*. A finite element graph is obtained from a planar graph as follows: Embed the graph in the plane. Identify certain points (vertices, points on edges, points in faces) as "nodes". Add edges between all nodes that share a face. If the number of nodes per face is bounded by k, finite element graphs satisfy a $kn^{1/2}$-separator theorem [15].

For a recursive divide-and-conquer algorithm based on separators to work, the subgraphs into which the original graph is separated must themselves have separators, and so on. We say that a class of graphs is *closed under subgraph* if it contains all subgraphs of all its members. The class of planar graphs is closed under subgraph.

We say that a graph G has a $n^{1/2}$-*separator decomposition* (with constants α and β) if G has a $n^{1/2}$-separator C with those constants and every connected component of $G-C$ has a $n^{1/2}$-separator decomposition. Having a $n^{1/2}$-separator decomposition is a weaker condition than having $n^{1/2}$-separators for all subgraphs. (An example is the graph with a $n^{1/2}$-vertex clique and $n-n^{1/2}$ isolated vertices.) Leighton [12] discusses this difference for separators that consist of edges rather than vertices. A $n^{1/2}$-separator decomposition (for edges) is a $(\beta n^{1/2}, 1/\alpha^{1/2})$-*bifurcator* in his terminology.

The fill bounds for planar graphs (and, more generally, graphs of bounded genus) follow from two facts about such graphs. First, they are sparse.

Lemma 2 [9]. *If G is a graph of genus g with $n>2$ vertices and m edges, then $m \leqq 3n-6+6g$. If in addition G is bipartite, then $m \leqq 2n-4+4g$.* □

Second, a planar graph remains planar when two adjacent vertices and the edge between them are contracted into a single vertex. Let G be a graph, and let $\{v, w\}$ be an edge of G. Let G' be the graph that is obtained from G by replacing v and w with a single vertex adjacent to every vertex that is adjacent to either v or w in G. We say that G was transformed into G' by *contracting* the edge $\{v, w\}$. A *contraction* of G is any graph obtained from G by contracting a set of edges. Equivalently, a contraction can be got by selecting a set of connected subgraphs of G and shrinking each to a single vertex. A contraction of a planar graph is still planar; in fact, any contraction of a graph embedded in a surface can be embedded in the same surface, since an edge can be shrunk continuously without disturbing the embedding.

The essential property for the fill analysis is that contractions must be sparse. A class S of graphs is said to be *sparse-contractible* (with density $\delta>0$) if every n-vertex contraction of a graph G in S has at most $\delta n+O(1)$ edges. Thus planar graphs are sparse-contractible with density 3, and trees are sparse-contractible with density 1.

3. The Nested Dissection Algorithm

Nested dissection operates by finding a separator in the graph, ordering its vertices last, and then recursively ordering the vertices in the subgraphs left by removal of the separator. In this section we present two nested dissection algorithms. The first is due to Lipton, Rose, and Tarjan; the second is the one we analyze in the remainder of this paper.

Let $\alpha<1$ and $\beta>0$ be the constants in a $n^{1/2}$-separator theorem, and let n_0 be a positive constant. The first algorithm assumes that all of G's subgraphs satisfy a $n^{1/2}$-separator theorem, while the second just assumes that G has a $n^{1/2}$-separator decomposition.

Algorithm 1. Given an integer a and an n-vertex graph G whose subgraphs all satisfy a $n^{1/2}$-separator theorem, number the vertices of G up to a. In general this algorithm assumes that some of the vertices of G (say l of them) are already numbered, and numbers the other $n-l$ from $a-n+l+1$ up to a.

If n is not more than n_0, number the unnumbered vertices arbitrarily. Otherwise, proceed as follows.

1. [separate] Choose a set C of at most $\beta n^{1/2}$ vertices whose removal divides the rest of G into two (not necessarily connected) components A and B with at most αn vertices each. Suppose that C has s unnumbered vertices. Number them arbitrarily from $a-s+1$ up to a.

2. [form components] Let G_1 be $(A \cup C, E(A \cup C)-E(C))$, that is, the subgraph of G induced by the vertices in A and C less any edges with both endpoints in C. Similarly, let G_2 be $(B \cup C, E(B \cup C)-E(C))$. Suppose that G_1 and G_2 have s_1 and s_2 unnumbered vertices respectively.

3. [number components recursively] Call the algorithm recursively twice to

number G_1 from $a-s-s_1+1$ up to $a-s$, and to number G_2 from $a-s-s_1-s_2+1$ up to $a-s-s_1$.

To begin, call algorithm with all vertices unnumbered and $a=n$. □

This is Lipton, Rose, and Tarjan's original version of the generalized nested dissection algorithm [14]. We shall call this the LRT algorithm and refer to the elimination order it produces as the LRT order. It guarantees $O(n \lg n)$ fill and $O(n^{3/2})$ operation count for a graph all of whose subgraphs satisfy a $n^{1/2}$-separator theorem. Notice that the algorithm is not called recursively on every connected component of $G-C$, but that the components are divided into two groups and exactly two recursive calls are made at each step. Also notice that the vertices of the separator are included in the recursive calls, but are not renumbered.

Algorithm 2. Given an integer a and an n-vertex graph G with a $n^{1/2}$-separator decomposition, number the vertices of G up to a.

If n is not more than n_0, number the unnumbered vertices arbitrarily. Otherwise, proceed as follows.

1. [separate] Find a separator C with $s \leqq \beta n^{1/2}$ vertices that divides the rest of G into connected components $A_1, \ldots, A_k$, where A_i has $s_i \leqq \alpha n$ vertices. Number the vertices of C arbitrarily from $a-s+1$ up to a.

2. [number components recursively] Call the algorithm recursively k times for $i=1, 2, \ldots, k$ to number the vertices of A_i up to $a-s-\sum_{j<i} s_j$.

To begin, call the algorithm with $a=n$. □

This algorithm, which we call the ND algorithm, leaves the vertices of C out of the recursive call and does one recursive call per component. The ND algorithm does not give the same fill bounds as the LRT algorithm for all classes of graphs with $n^{1/2}$-separators; in Sect. 9 we shall present a class for which LRT gives $O(n \lg n)$ fill but ND can give $\Theta(n^{5/4})$ fill. However, ND does give $O(n \lg n)$ fill and $O(n^{3/2})$ operation count for planar graphs, finite element graphs, graphs of bounded genus, and graphs of bounded degree with $n^{1/2}$-separators.

We feel it is interesting to study the ND algorithm for a number of reasons. First is the theoretical question whether including the separator in the recursive calls of LRT is really necessary. The analysis showing that the answer is "sometimes, but not on the graphs we are most interested in" is rather different in flavor from the analysis of the LRT algorithm. The ND version of the algorithm should be a little easier to implement than the LRT version; in particular, it fits nicely into the nested dissection routines in the Waterloo Sparspak sparse matrix package [4]. Finally, the constants in the fill bounds for ND are somewhat smaller than those in LRT.

4. Separator Trees

Suppose the graph G has a $n^{1/2}$-separator decomposition with constants $\alpha<1$ and $\beta>0$. The recursion in the ND algorithm decomposes G into a tree of separators. Figure 4 shows a graph and the separators used by the algorithm,

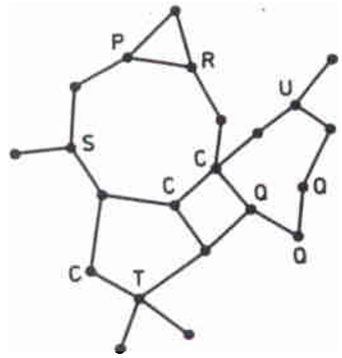

Fig. 4. Separators used by the ND algorithm

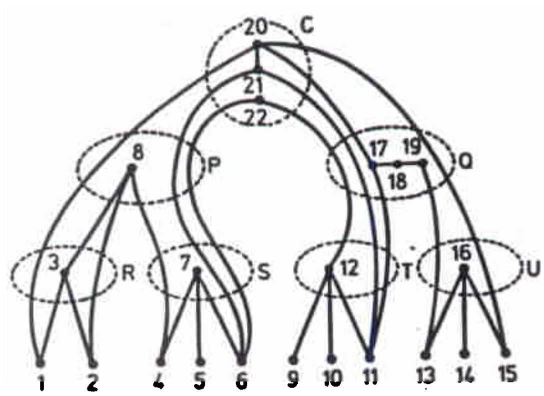

Fig. 5. The separator tree

and Fig. 5 shows the same graph drawn to exhibit the tree structure. The *separator tree* produced by nested dissection for a graph G is the tree whose internal nodes are the separators and whose external nodes are the bottom-level divisions of G (with at most n_0 vertices each); each separator has as children the separators of the parts into which it separates its subgraph. To keep things straight, the vertices of G are called *vertices*, and the vertices of the separator tree are called *nodes*. Thus a node is a subgraph of G, and may contain many vertices.

Much of the structure of G is reflected in the separator tree. An edge $\{v, w\}$ can be in G only if the node containing v is an ancestor or descendant of the node containing w. (A node is its own ancestor.) The ND elimination order is a postorder on the tree, so all the vertices in one node are eliminated before any vertex in that node's parent. The vertices of each node are ordered consecutively.

Lemma 1 says that if $\{v, w\}$ is a fill edge, there is a path from v to w through vertices with lower numbers than either v or w. Since the vertices in the separator of a subgraph have higher numbers than the other vertices in the subgraph, this means that no fill edge can cross a separator. This in turn implies that fill edges, like edges of G, must follow tree paths.

Lemma 3. *In the ND order on G, if $\{v, w\}$ is a fill edge and v has a higher number than w, then the node of the separator tree containing v is an ancestor of the node containing w.* □

This lemma limits the number of possible fill edges to $O(n^{3/2})$ (since there are $O(n^{3/2})$ possible edges to the top-level separator, and the sum of this over the whole tree is $O(n^{3/2})$ by Lemma 12), but we can do a lot better. The key observation is the following lemma.

Lemma 4. *In the ND order on G, if $\{v, w\}$ is a fill edge and v's node is an ancestor of w's node, then there is an edge of G from v to a vertex in some node that is a descendant of w's node.*

Proof. By Lemma 1, there is a path from v to w through vertices with lower numbers than w. Let $\{v, x\}$ be the first edge on that path. Vertex x cannot be in a node that is a proper ancestor of w's node, since x has lower number than w. If x's node were neither an ancestor nor a descendant of w's node, the path from x to w would have to include a vertex in a node that is an ancestor of both x's node and w's node; but such a vertex would have a higher number than w. Therefore x's node must be a descendant of w's node, and $\{v, x\}$ satisfies the statement of the lemma. □

The levels of the separator tree are numbered from zero, which is the level of the root. A basic property of the separator tree is that a subtree rooted on level k has at most $\max\{n_0, \alpha^k n\}$ vertices of G in it, and hence the root of such a subtree has at most $\alpha^{k/2}\beta n^{1/2}$ vertices if it is an internal node.

Lemma 5. *Let G be as above, and let T be the separator tree produced for G by the ND algorithm. Let N_i be a level-i node of T for $0 \leq i \leq m$. The number of vertices of G in $N_0 \cup N_1 \cup \ldots \cup N_m$ is less than*

$$\frac{\beta\alpha^{1/2}}{1-\alpha^{1/2}} n^{1/2} + n_0.$$

Proof. The size of a separator at the k-th level of the tree is at most $\alpha^{k/2}\beta n^{1/2}$, so the total size of nodes $N_0, \ldots, N_m$ is less than

$$n_0 + \alpha^{1/2}\beta n^{1/2} + \alpha\beta n^{1/2} + \ldots = n_0 + \frac{\beta\alpha^{1/2}}{1-\alpha^{1/2}} n^{1/2},$$

where the first term is for an external node. □

5. Planar and Sparse-Contractible Graphs

The main result of this section is that if G is a planar graph, then fill for the ND elimination order is $O(n \lg n)$. Planar graphs do not necessarily have bounded degree, but they are sparse; a planar graph with n vertices has at most $3n-6$ edges. Thus the average degree of the vertices is bounded. This alone is not enough to prove an $O(n \lg n)$ fill bound, as the example in Sect. 9 shows. However, planar graphs are also sparse-contractible: They remain

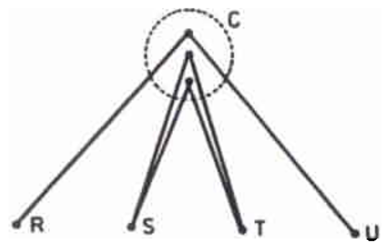

Fig. 6. Contracted graph for Lemma 6

sparse when two adjacent vertices and the edge between them are contracted into a single vertex. This and a $n^{1/2}$-separator theorem are enough to imply the fill bound.

Sparsity, contractibility, and separability are related in devious ways. If a class S of graphs is closed under subgraph and is sparse-contractible with density δ, then any subgraph of a contraction of a graph in S also has at most $\delta n + O(1)$ edges. This follows because any subgraph of a contraction of G can be obtained by deleting edges from a contraction of a subgraph of G. If S is closed under subgraph and satisfies a $n^{1/2}$-separator theorem, then graphs in S are sparse [15]. Therefore if S satisfies a $n^{1/2}$-separator theorem and is closed under subgraph and contraction, it is also sparse-contractible. It is not known whether sparse contractibility implies a nontrivial separator theorem.

The first bound on fill in this section will be stated for a class of graphs that satisfies a $n^{1/2}$-separator theorem and is closed under subgraph and contraction, but the proof will use only the facts that it has a $n^{1/2}$-separator decomposition and is sparse-contractible.

Lemma 6. *Let S be a class of graphs that satisfies a $n^{1/2}$-separator theorem with constants $\alpha < 1$ and $\beta > 0$ and is closed under subgraph and contraction. As remarked above, this implies that graphs in S are sparse. Suppose no n-vertex graph in S has more than $\delta n + c$ edges. When the ND algorithm is applied to a graph G in S with $n > n_0$ vertices, the number of fill edges with at least one endpoint in C, the top-level separator, is $O(\delta n)$.*

Proof. Let $\mathcal{N}$ be the set of nodes of the separator tree for G, and let $\mathcal{N}_k$ be the set of nodes on level k of the tree. Thus $\mathcal{N}_0 = \{C\}$, and $\mathcal{N} = \mathcal{N}_0 \cup \mathcal{N}_1 \cup \ldots$. For a given node N, let s_N be the number of vertices in N.

We begin by counting fill to the root C of the separator tree from the nodes on level k of the tree. Each subtree whose root is on level k is connected. Consider the graph that is obtained from G by contracting each such subtree into a single vertex. Throw out all the vertices of this graph except contracted vertices and vertices in C. Also throw out edges between vertices in C. Call the resulting graph G_k. Figure 6 shows G_2 from the graph in Figs. 4 and 5. Graph G_k is in S, so it has at most $\delta|G_k| + c$ edges.

By Lemma 4 there can be fill to a vertex $v \in C$ from a level-k node N only if there is an edge in G_k from v to the contracted vertex corresponding to N.

Each such edge accounts for at most one fill edge from each vertex of G in N, or s_N fill edges in all. If f_k is the size of the fill to C from level-k nodes, and e_N is the degree in G_k of the contracted vertex corresponding to node N, this means that

$$f_k \leq \sum_{N \in \mathcal{N}_k} e_N s_N.$$

e_N

Let $\mathcal{M}_k$ be $\{N \in \mathcal{N}_k : s_N > \delta\}$, the set of level-$k$ nodes with degree greater than δ in the contracted graph. Then

$$\begin{aligned} f_k &\leq \sum_{\mathcal{N}_k} \delta s_N + \sum_{\mathcal{M}_k} (e_N - \delta) s_N \\ &\leq \delta \sum_{\mathcal{N}_k} s_N + \hat{s}_k \sum_{\mathcal{M}_k} (e_N - \delta), \end{aligned} \tag{1}$$

where $\hat{s}_k = \max_{N \in \mathcal{N}_k} s_N$.

Consider the subgraph of G_k that is induced by the vertices of C and the contracted vertices in $\mathcal{M}_k$. This graph is in S, and it has at most $\beta n^{1/2} + |\mathcal{M}_k|$ vertices, so

$$\sum_{\mathcal{M}_k} e_N < \delta(\beta n^{1/2} + |\mathcal{M}_k|) + c,$$

and

$$\sum_{\mathcal{M}_k} (e_N - \delta) < \delta \beta n^{1/2} + c. \tag{2}$$

Equations (1) and (2) imply

$$f_k < \delta \sum_{\mathcal{N}_k} s_N + \delta \beta n^{1/2} \hat{s}_k + c \hat{s}_k. \tag{3}$$

Therefore the total fill to C is

$$\sum_{k \geq 0} f_k < \binom{|C|}{2} + \delta \sum_{\mathcal{N}} s_N + \delta \beta n^{1/2} \sum_{k > 0} \hat{s}_k + c \sum_{k > 0} \hat{s}_k. \tag{4}$$

Certainly $\sum_{\mathcal{N}} s_N = n$, and $\binom{|C|}{2} < \beta^2 n/2$. By Lemma 5,

$$\sum_{k > 0} \hat{s}_k < \frac{\beta \alpha^{1/2}}{1 - \alpha^{1/2}} n^{1/2} + n_0.$$

Substituting these estimates into equation (4) gives a bound on fill to C of

$$\frac{\beta^2 n}{2} + \delta n + \frac{\delta \beta^2 \alpha^{1/2}}{1 - \alpha^{1/2}} n + O(n^{1/2}). \tag{5}$$

This is the bound claimed in the statement of the lemma. □

Theorem 2. *Let S be a class of graphs that satisfies a $n^{1/2}$-separator theorem and is closed under contraction and subgraph. Suppose that no n-vertex graph in S has more than $\delta n + c$ edges. If G in S has $n > n_0$ vertices, the ND order causes $O(\delta n \lg n)$ fill.*

Proof. The fill when eliminating G is the union over every internal separator tree node of the fill edges whose higher-numbered vertex is in that node, plus the fill edges within the external nodes of the tree. A fill edge whose higher-numbered endpoint is in a given internal node has its other endpoint in a descendant of that node. Thus if a given internal node is the root of a subtree containing m vertices, by Lemma 6 the number of fill edges with higher-numbered endpoints in that node is at most

$$\left(\frac{\beta^2}{2}+\delta+\frac{\delta\beta^2\alpha^{1/2}}{1-\alpha^{1/2}}\right)m+O(m^{1/2}).$$

By Lemma 12 (in the Appendix), the sum of this function over the internal nodes of the separator tree is at most

$$\frac{\beta^2/2+\delta+\delta\beta^2\alpha^{1/2}/(1-\alpha^{1/2})}{-\alpha\lg\alpha-(1-\alpha)\lg(1-\alpha)}\,n\lg n+O(n).$$

Fill within an external node is at most $\binom{n_0}{2}$ edges, for a total over the whole graph of $O(n)$ edges. Thus the expression above bounds the fill for the entire graph. □

Planar graphs are the most interesting graphs to which Theorem 3 applies. To finish this section we shall compute the leading coefficient of the bound for planar graphs, after first tightening the analysis slightly.

Theorem 3. *If G is a planar graph with n vertices then the ND order causes fill at most*

$$\frac{\beta^2/2+2+\beta^2 2^{1/2}/(1-\alpha^{1/2})}{-\alpha\lg\alpha-(1-\alpha)\lg(1-\alpha)}\,n\lg n+O(n).$$

Proof. This is a slight improvement of the bound in Theorem 3, and we get it by tightening the proof of Lemma 6. First, the graph whose edges we counted in the proof of the lemma (the subgraph of G_k induced by the vertices of C and the contracted vertices) is planar and bipartite, so by Lemma 2 we can take $\delta=2$ and $c=0$.

Now we can be more careful in the analysis of $\sum_{\mathcal{M}_k}(e_N-2)s_N$ in Eq. (1) of Lemma 6. Since there are at most $\beta n^{1/2}$ vertices in C, no single e_N can be greater than $\beta n^{1/2}$. Thus if $\hat{s}_k$ is the size of the largest node on level k, and $\hat{\hat{s}}_k$ is the size of the second-largest node on level k, then

$$\begin{aligned}\sum_{\mathcal{M}_k}(e_N-2)s_N &\le \beta n^{1/2}\hat{s}_k+\Big(\sum_{\mathcal{M}_k}(e_N-2)-\beta n^{1/2}\Big)\hat{\hat{s}}_k\\ &<\beta n^{1/2}(\hat{s}_k+\hat{\hat{s}}_k).\end{aligned}\tag{6}$$

Lemma 6 used the bound $\hat{s}_k\le\alpha^{k/2}\beta n^{1/2}$. Now, however, we can reason as follows. Suppose that N is the largest node on level k and M is the second-largest. Nodes N and M have a lowest common ancestor P in the separator tree. Suppose that P is on level l. The number of vertices in P's subtree is at most $\alpha^l n$. If P has children P_1 and P_2, the numbers of vertices in their subtrees are at most $\alpha^l \gamma n$ and $\alpha^l(1-\gamma)n$ respectively, for some γ between 0 and 1. Following on down to level k, this means that the numbers of vertices in N's and M's subtrees are at most $\alpha^{k-1}\gamma n$ and $\alpha^{k-1}(1-\gamma)n$ respectively. Thus $\hat{s}_k+\hat{\hat{s}}_k$, which is s_N+s_M, is at most $\beta\alpha^{(k-1)/2}(\gamma^{1/2}+(1-\gamma)^{1/2})n^{1/2}$. This is largest when $\gamma=\frac{1}{2}$, so

$$\hat{s}_k+\hat{\hat{s}}_k\leqq 2^{1/2}\beta\alpha^{(k-1)/2}n^{1/2}.$$

With Eq. (6) this gives

$$\sum_{\mathcal{M}_k}(e_N-2)s_N\leqq\beta^2 2^{1/2}\alpha^{(k-1)/2}n. \tag{7}$$

Summing this for $k>0$ as in Lemma 5 yields $\beta^2 2^{1/2}n/(1-\alpha^{1/2})$, and putting it all together gives a bound on fill to C of

$$\left(\frac{\beta^2}{2}+2+\frac{\beta^2 2^{1/2}}{1-\alpha^{1/2}}\right)n+O(n).$$

This is summed by Lemma 12 to give the bound in the statement of the theorem. □

The constants in Djidjev's version of the planar separator theorem [1] are $\alpha=2/3$ and $\beta=6^{1/2}$. Plugging these into the leading coefficient of the bound in Theorem 4 yields about 55.8. The coefficient derived by Lipton et al. [14] for their algorithm, which includes the separator in both recursive calls, is about 96.4 (using Djidjev's separators). It seems likely that both numbers are somewhat larger than the best possible bounds. One reason is that the coefficients are proportional to the square of β, and $\beta=6^{1/2}$ is probably an overestimate of the best possible value. Section 8 contains further discussion of the leading coefficient in the fill bound.

6. Graphs of Bounded Degree

Let G be a graph with a $n^{1/2}$-separator decomposition and maximum vertex degree d. Techniques similar to those in Sect. 5 can be used to prove the following.

Theorem 4. *If G is as above, then the ND order causes fill at most*

$$\frac{\beta^2/2+d\beta^2\alpha^{1/2}/(1-\alpha^{1/2})}{-\alpha\lg\alpha-(1-\alpha)\lg(1-\alpha)}n\lg n+O(dn).\quad\square$$

Details are given in Gilbert [5]. For $\alpha=2/3$ and $\beta=6^{1/2}$ the leading constant is approximately $29d+3$. Roman [19] independently obtained this result with the leading constant $(\beta^2/2+d\beta^2\alpha^{1/2}/(1-\alpha^{1/2}))/-\lg\alpha$, which is approximately $46d+5$.

7. Bounds on Operation Count

Until now we have been concerned only with the fill incurred by an elimination order, which is a measure of the space required to factor the matrix. Here we analyze the time taken by the factorization in ND order. The entire factorization can be done in time proportional to the number of arithmetic operations performed on matrix elements [4], so we will just analyze the operation count.

In practice storage is more expensive than time. We might therefore concentrate on fill bounds and have faith that, since operation count is at least loosely related to fill, a low fill algorithm will have a small operation count. However, it is comforting to know at least the order of growth of the operation count. Dense Gaussian elimination requires $\Theta(n^3)$ operations on an n by n matrix. George's nested dissection on a square grid requires $O(n^{3/2})$ operations, as does the Lipton-Rose-Tarjan algorithm on a graph whose subgraphs satisfy a $n^{1/2}$-separator theorem. We will show that the ND algorithm also requires $O(n^{3/2})$ operations on all the classes of graphs discussed in Sects. 5 and 6.

Pivoting on the diagonal element corresponding to vertex v in a graph G (that is, eliminating vertex v) requires arithmetic operations proportional to the square of the degree of v [4]. One way to count operations for an elimination order is based on the filled graph G^*. The filled graph has G's edges plus the fill edges. Let us make G^* a directed graph by orienting each edge from the lower-numbered to the higher-numbered endpoint. Then the cost of eliminating v is $O(\mathrm{od}(v)^2)$, where $\mathrm{od}(v)$ is the out-degree of v in this directed version of G^* (and also one less than the number of nonzeros in the column of L corresponding to v). The operation count for the whole elimination is $O(\sum_v \mathrm{od}(v)^2)$.

Notice that the space required for elimination is on the order of the number of edges in G^*, which is $\sum_v \mathrm{od}(v)$. Thus space is a first moment of the out-degree, and time is a second moment. This gives an easy way to get a rough bound on operation count. Suppose G is a graph with a $n^{1/2}$-separator decomposition and suppose the ND algorithm gives $O(n \lg n)$ fill on G. If (v, w) is a directed edge in G^*, the node of the separator tree containing w is an ancestor of the node containing v. Therefore $\mathrm{od}(v)$ is at most the number of vertices on a path in the tree from the root to a leaf, which is $O(n^{1/2})$ by Lemma 5. Then the operation count is at most

$$\sum_v \mathrm{od}(v)^2 \le \max_v \mathrm{od}(v) \sum_v \mathrm{od}(v) = O(n^{3/2} \lg n).$$

Getting rid of the extra $\lg n$ in this loose bound requires some careful argument along lines similar to those of the $O(n \lg n)$ fill bound. Again we will contract subtrees of the separator tree and use the sparsity of the resulting graph. In the space bound we counted fill edges by their higher-numbered vertices, which amounts to computing fill as $\sum_v \mathrm{id}(v)$ in G^*. Using the lower-numbered vertices makes things look a little bit different. To keep the proof from getting too unwieldy, we shall begin by stating it for planar graphs; we shall then indicate how to modify it for degree-bounded and sparse-contractible graphs.

A planar bipartite graph has fewer than twice as many edges as vertices. The first lemma says that we can actually associate each edge with one of its endpoints so that no vertex is associated with more than two edges.

Lemma 7. *Let G be a planar bipartite graph with n vertices (and hence at most $2n-4$ edges). There is a function ϕ from the edges of G to the vertices of G such that for all edges e, $\phi(e)$ is an endpoint of e; and for all vertices v, $\phi(e)=v$ for at most two different edges e.*

Proof. Define the *arboricity* of a graph to be the minimum number of edge-disjoint spanning forests into which the graph can be decomposed. Thus if the arboricity of $G=(V,E)$ is k, it is possible to write $E=E_1\cup\ldots\cup E_k$ in such a way that (V,E_i) is acyclic for $1\leqq i\leqq k$. A theorem of Nash-Williams [17] is that if q_r is the maximum number of edges in any r-vertex subgraph of H, then the arboricity of H is $\max_r\,[q_r/(r-2)]$. Any subgraph of a planar bipartite graph is planar and bipartite, so the arboricity of G is at most 2. Therefore G can be written as $(V,E_1\cup E_2)$ where $G_1=(V,E_1)$ and $G_2=(V,E_2)$ are forests. To find the required function ϕ, proceed as follows. Any forest with at least one edge has at least one vertex of degree one. Choose an edge $\{v,w\}$ in E_1 such that v has degree one in G_1. Set $\phi(\{v,w\})=v$, and delete $\{v,w\}$ from E_1. Repeat until E_1 is empty, and then carry out the same process with G_2. This assigns a value to $\phi(e)$ for every edge e of G, and uses each vertex at most twice. □

Incidentally, if we only require that $\phi(e)=v$ for at most three different edges e, the result follows immediately from the fact that the average vertex degree of every subgraph of G is less than four. Simply find a vertex of degree three or less, associate with it its incident edges, delete that vertex and those edges from the graph, and repeat until the graph is empty.

Lemma 8. *Let G be a planar graph, and let G^* be the filled graph consisting of G plus the fill from the ND order. Direct the edges of G^* from lower to higher numbered vertices. Then* $\sum_v \mathrm{od}(v)^2=O(n^{3/2})$.

Proof. This proof will use much of the same notation as in the proof of the fill bound in Lemma 6.

Let $\mathcal{N}$ be the set of nodes of the separator tree for G, and let $\mathcal{N}_k$ be the set of nodes on level k of the tree. We will bound the operation count separately for each level of the tree. Let $p_k=\sum_{v\in N\in\mathcal{N}_k}\mathrm{od}(v)^2$ be the sum over all vertices on level k of the square of the out-degree.

Each subtree whose root is on level k is a connected subgraph of G. Let G_k be the graph obtained from G by contracting each such subgraph into a single vertex, and deleting any edges of G that are not incident on contracted vertices. Figure 7 shows G_2 from the graph in Figs. 4 and 5. (This is not quite the same G_k as we defined in the proof of Lemma 6; there we also deleted all vertices on

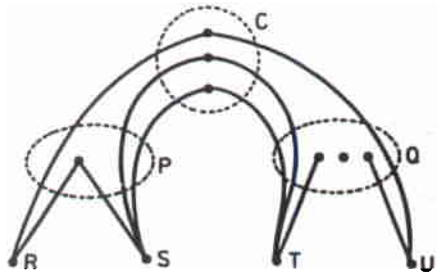

Fig. 7. Contracted graph for Lemma 8

levels 1 through $k-1$.) This graph is planar and bipartite. Now let v be a vertex of G in node N on level k of the separator tree, and let (v, w) be an edge of G^* (thus $\{v, w\}$ is an edge of G or a fill edge, and v is eliminated before w). The node containing w is an ancestor of N, and by Lemma 4 there is an edge in G joining w and some vertex in the subtree rooted at N. Therefore either v and w are both in node N, or in G_k there is an edge joining w and the contracted vertex corresponding to N. If s_N is the number of vertices in node N and e_N is the number of edges incident on contracted vertex N in G_k, then $\mathrm{od}(v)$ is at most $s_N + e_N$, so

$$p_k \leq \sum_{N \in \mathcal{N}_k} s_N(s_N + e_N)^2. \tag{1}$$

Lemma 7 says we can associate each edge of G_k with one of its endpoints in such a way that at most two edges are associated with each vertex. Let us call those edges associated with contracted vertices "red" edges and those edges associated with vertices of G on levels 0 through $k-1$ of the separator tree "blue" edges. Then the e_N edges incident on contracted vertex N consist of at most two red edges and some blue edges. Suppose that there are r_N red edges and b_N blue edges, so $r_N \leq 2$ and $e_N = r_N + b_N$. Equation (1) can be written as

$$p_k \leq \sum_{N \in \mathcal{N}_k} s_N(s_N + r_N + b_N)^2.$$

Now if a, b, and c are real numbers, then

$$(a+b+c)^2 \leq (a+b+c)^2 + (a-b)^2 + (b-c)^2 + (c-a)^2 = 3(a^2+b^2+c^2),$$

so

$$\begin{aligned} p_k &\leq 3 \sum_{N \in \mathcal{N}_k} s_N^3 + 3 \sum_{N \in \mathcal{N}_k} s_N r_N^2 + 3 \sum_{N \in \mathcal{N}_k} s_N b_N^2 \\ &\leq 3 \sum_{N \in \mathcal{N}_k} s_N^3 + 12 \sum_{N \in \mathcal{N}_k} s_N + 3\hat{s}_k \sum_{N \in \mathcal{N}_k} b_N^2, \end{aligned} \tag{2}$$

where $\hat{s}_k = \max_{N \in \mathcal{N}_k} s_N$.

The first two terms of expression (2) are easy to handle. The main job is to bound

$$\sum_{N \in \mathcal{N}_k} b_N^2. \tag{3}$$

To do this we will pair each contracted vertex N with a node on some level from 0 to $k-1$, and then dominate sum (3) by a sum over the nodes on levels 0 through $k-1$. Consider some node M on level $r<k$ of the separator tree. The vertices in M are vertices of G_k, each with at most two incident blue edges. The other endpoints of these blue edges are contracted vertices. We shall argue that it suffices to bound sum (3) under the assumption that all the blue edges with one endpoint in M have the same contracted vertex as their other endpoint, possibly allowing multiple edges. Suppose that the blue edges from vertices in M are incident on at least two different contracted vertices N and N', with $b_N \geq b_{N'}$. Then sum (3) can only be made larger if each blue edge $\{v, N'\}$ with $v \in M$ is replaced by a blue edge $\{v, N\}$. Thus we shall assume henceforth that all the blue edges coming from the vertices in the same node go to the same contracted vertex.

Now let N be a contracted vertex. Blue edges incident on N may come from nodes on several levels. Let M be the node closest to the root such that a blue edge $\{v, N\}$ exists for $v \in M$. Then all the blue edges incident on N come from nodes on the tree path from M to N. If the number of vertices of G in the subtree rooted at M is n_M, the number of vertices of G on this tree path is at most

$$\beta n_M^{1/2} + \beta(\alpha n_M)^{1/2} + \beta(\alpha^2 n_M)^{1/2} + \ldots = O(n_M^{1/2}),$$

so b_N, the number of blue edges incident on N, is also $O(n_M^{1/2})$.

The mapping that takes a contracted vertex N to the node M described above is one-to-one since each node M on levels 0 through $k-1$ has blue edges to only one contracted vertex. Therefore, since $b_N^2 = O(n_M)$,

$$\sum_{N \in \mathcal{N}_k} b_N^2 \leq \sum_{\substack{M \in \mathcal{N}_r \\ 0 \leq r < k}} c n_M$$

for some $c>0$. Lemma 13 (in the Appendix) says that the second sum is at most ckn, so

$$\sum_{N \in \mathcal{N}_k} b_N^2 \leq ckn. \tag{4}$$

Now we can put it all together. Substituting Eq. (4) into Eq. (2) and summing over all levels k yields

$$\sum_{v \in G^*} \mathrm{od}(v)^2 \leq 3 \sum_{N \in \mathcal{N}} s_N^3 + 12 \sum_{N \in \mathcal{N}} s_N + 3 \sum_{k \geq 0} \hat{s}_k ckn. \tag{5}$$

Since $s_N \leq \beta n_N^{1/2}$ and $\hat{s}_k \leq \alpha^{k/2} \beta n^{1/2}$, this is at most

$$3\beta^3 \sum_{N \in \mathcal{N}} n_N^{3/2} + 12\beta \sum_{N \in \mathcal{N}} n_N^{1/2} + 3\beta c n^{3/2} \sum_{k \geq 0} k\alpha^{k/2}.$$

By Lemma 12, the first sum is $O(n^{3/2})$ and the second is $O(n)$. The third sum converges to a constant, so the entire expression is $O(n^{3/2})$, as we set out to prove. □

Theorem 5. *If G is a planar graph with n vertices then the ND order gives an $O(n^{3/2})$ operation count.* □

This bound can be extended to the other classes of graphs for which we have proved fill bounds. We sketch the proofs below.

Corollary 1. *Let S be a class of graphs that is closed under contraction and subgraph, and satisfies a $n^{1/2}$-separator theorem. If G is an n-vertex graph in S then the ND order gives an $O(n^{3/2})$ operation count.*

Proof. Recall from Sect. 6 that graphs in S must be sparse; suppose that no n-vertex graph in S has more than $\delta(n-1)$ edges. Then the arboricity of any member of S is at most δ, so we can associate each edge of a member of S with one of its endpoints so that no vertex is associated with more than δ edges. The proof of Lemma 8 now applies, with at most δ red edges incident on each contracted vertex of G_k and at most δ blue edges incident on each noncontracted vertex. □

Corollary 2. *Let G be an n-vertex graph with a $n^{1/2}$-separator decomposition and with vertex degree bounded by d. The ND order gives an $O(n^{3/2})$ operation count.*

Proof. The proof of Lemma 8 applies again with minor changes. The contracted graphs G_k do not necessarily have bounded degee. However, the degrees of the noncontracted vertices are still at most d. We can apply the analysis in the lemma by coloring all the edges blue. Each noncontracted vertex has at most d incident blue edges, and each contracted vertex has no incident red edges. □

Jean Roman [19] has independently proved an $O(n^{3/2})$ bound on operation count for the ND order on graphs of bounded degree.

8. Remarks on the Constants

In Sect. 6 we saw that the ND order on a planar graph gives less than $56n \lg n + O(n)$ fill. Here we exhibit a class of planar graphs for which fill is about $32n \lg n$ in the worst case.

First let us make clear what we mean by "the worst case". The ND algorithm treats the planar separator algorithm as a black box that returns a set of at most $\beta n^{1/2}$ vertices. The graph we construct below has large fill for one specific elimination order that comes from one specific $n^{1/2}$-separator decomposition of the graph. In this sense the separator algorithm is considered an adversary.

Let n_0 and β be fixed. We shall construct a family of graphs (with separator trees) with parameter k, the number of levels in the separator tree. Each graph has one or two distinguished vertices called *terminal vertices*. For $k=0$ the graph is a star with n_0 vertices. The vertex of degree n_0-1 is the terminal vertex, and it is eliminated first. Now suppose we have constructed the graph whose tree has k levels, and that it has n_k vertices. To construct the tree with $k+1$ levels, take two copies of the k-level tree and add s_{k+1} new vertices to

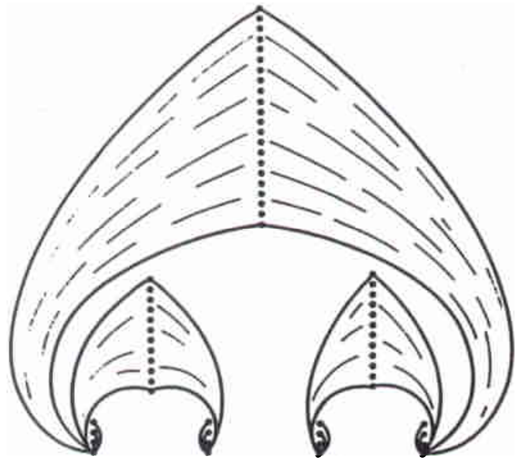

Fig. 8. A parachute graph

form a top-level separator, where s_{k+1} is maximum subject to $s_{k+1} \leqq \beta(2n_k + s_{k+1})^{1/2}$. Add two edges incident on each top-level vertex, one to a terminal vertex of each subgraph. These two terminal vertices of the subgraphs are the terminal vertices of the $k+1$-level graph. Figure 8 shows the graph for $n_0 = 4$, $\beta = 6^{1/2}$, and $k=2$. We call this the *parachute graph*.

It is not immediately clear how to count the vertices in the parachute graph. The leading coefficient of the fill can be computed without knowing exactly how big the graph is, however. Let n_k be the number of vertices in the k-level parachute, and let s_k be the number of vertices in its top-level separator. (Notice that we are now numbering levels from the bottom of the separator tree rather than, as usual, from the top. This is useful because now, in any parachute graph, s_l is the size of a separator at level l and n_l is the size of the subtree rooted there.) Solving $s_{k+1} \leqq \beta(2n_k + s_{k+1})^{1/2}$ for s_{k+1} leads to the following recurrence, which holds for $k \geqq 0$.

$$\begin{aligned} s_{k+1} &= [(2\beta^2 n_k + \beta^4/4)^{1/2} + \beta^2/2] \\ n_{k+1} &= 2n_k + s_{k+1}. \end{aligned} \tag{1}$$

Lemma 9. *Let s_k be given by recurrence* (1) *above. Then there is a positive constant γ such that $s_k = \gamma 2^{k/2} + O(1)$.*

Proof. First we eliminate n_k from the recurrence. Since $s_k = [\beta n_k^{1/2}]$, we can substitute $(s_k + O(1))^2$ for $\beta^2 n_k$ in the recurrence. Then some simplification yields

$$s_{k+1} = 2^{1/2} s_k + O(1).$$

Defining $t_k = 2^{-k/2} s_k$, we can multiply this equation by $2^{-(k+1)/2}$ to get

$$t_{k+1} = t_k + O(2^{-k/2}).$$

From this it follows that $\lim_{k\to\infty} t_k = \gamma$ exists, and that $|\gamma - t_k|$ is $O(\sum_{i\geq k} 2^{-i/2})$, which is $O(2^{-k/2})$. Thus $t_k = \gamma + O(2^{-k/2})$ and $s_k = \gamma 2^{k/2} + O(1)$. □

Now we use this estimate of s_k to get an estimate of the fill for a k-level graph.

Lemma 10. *If a k-level parachute graph is eliminated in ND order, the size of the fill is* $\gamma^2(\frac{5}{2} + 8^{1/2})k2^k + O(2^k)$, *where* γ *is the constant from Lemma* 9.

Proof. First consider fill edges that are incident on vertices in the top-level separator. A vertex in the top-level separator fills in to every other vertex in that separator, and also to every vertex in every separator on the two tree paths to the terminal vertices. Thus fill to the top-level separator is $\binom{s_k}{2}$ within that separator, plus $2s_k \sum_{0<i<k} s_i + n_0 - 1$ to vertices in other separators. This all sums to $\gamma^2(\frac{5}{2} + 8^{1/2})2^k + O(2^{k/2})$, using the estimate of s_k from Lemma 9.

It remains only to sum this over the whole separator tree. For $1 \leq i \leq k$ there are 2^{k-i} separators on level i. Fill within level 0 is less than $2^k n_0^2$, which is $O(2^k)$. Therefore the total fill is

$$\sum_{1\leq i\leq k} 2^{k-i}(\gamma^2(\tfrac{5}{2} + 8^{1/2})2^i + O(2^{i/2})) + O(2^k) = \gamma^2(\tfrac{5}{2} + 8^{1/2})k2^k + O(2^k). \quad \square$$

The estimates from the last two lemmas contain a constant γ whose value we do not know how to compute. Fortunately, this does not matter. We can compute the leading coefficient of the fill without knowing the number of vertices in the graph.

Theorem 6. *Let G be a k-level parachute graph. If G has* n_k *vertices, the fill caused by the ND order is*

$$(\tfrac{5}{2} + 8^{1/2})\beta^2 n_k \lg n_k + O(n_k).$$

Proof. First, n_k is $(s_k + O(1))^2/\beta^2$. The estimate of s_k in Lemma 9 then implies that $n_k = \gamma^2 2^k/\beta^2 + O(2^{k/2})$, and therefore $\lg n_k = k + O(1)$ and $n_k \lg n_k = (\gamma^2 k 2^k/\beta^2)(1 + O(1/k))$. Dividing this into the fill estimate from Lemma 10 yields $(\frac{5}{2} + 8^{1/2})\beta^2 + O(1/k)$. Since $1/k = O(1/\lg n_k)$, this proves the theorem. □

The leading coefficient of the parachute graph's fill is independent of n_0, the size below which fragments are not separated further. The constant α that bounds the unevenness of the split does not appear because all the splits are exactly in half; it might be possible to get slightly larger fill by splitting the graph unevenly, but the calculations become messy.

For $\beta = 6^{1/2}$ the parachute fill comes to about $32.0 n \lg n$, as compared to our upper bound of about $55.8 n \lg n$.

Of course, this is "worst-case" behavior in its most negative sense. We are assuming not only a particularly bad graph as input, but also the worst possible behavior of the separator algorithm. (We note, though, that the George-Liu automatic nested dissection algorithm would produce this worst-case order for the parachute graph.) The top-level separator in the parachute is a minimal separator but it is far from a minimum separator: The terminal vertices are a two-vertex separator. A nested dissection order using these separators would give only $O(n)$ fill. This suggests that adding some heuristics to the separator algorithm might improve the constant in the worst case, or at least in many cases.

9. Graphs with Larger Fill

The hypotheses of the $O(n \lg n)$ fill bound for the ND algorithm are a $n^{1/2}$-separator decomposition and either bounded degree or sparse contractibility. In this section we show that a $n^{1/2}$-separator decomposition alone is not enough to imply the fill bound, and that in fact even $n^{1/2}$-separators for all subgraphs are not enough.

Any graph with a $n^{1/2}$-separator decomposition has no more than $O(n^{3/2})$ fill under the ND order, because fill edges must follow tree paths. However, having a $n^{1/2}$-separator decomposition does not imply that a graph is sparse, much less that its filled graph has $O(n \lg n)$ edges. Consider the complete bipartite graph with $n^{1/2}$ vertices in one part and $n - n^{1/2}$ vertices in the other part. This graph has a $n^{1/2}$-separator decomposition, but it has $\Theta(n^{3/2})$ edges. We conclude that an $O(n \lg n)$ fill bound does not follow from a $n^{1/2}$-separator decomposition.

We suspect that there are sparse graphs with $n^{1/2}$-separator decompositions for which every elimination order gives $\Omega(n^{3/2})$ fill. Indeed, we conjecture that for any $\alpha < 1$, $\beta > 0$, and $\gamma > 0$ there exists $c > 0$ such that almost all cn-edge graphs that have a $n^{1/2}$-separator decomposition for constants α and β will have at least $\gamma n^{3/2}$ fill for every elimination order.

The Lipton-Rose-Tarjan generalized nested dissection algorithm, which includes the separator vertices in each recursive call, gives $O(n \lg n)$ fill on every graph whose subgraphs all satisfy a $n^{1/2}$-separator theorem. This is a stronger hypothesis than the existence of a $n^{1/2}$-separator decomposition. Does the ND algorithm give $O(n \lg n)$ fill on every graph whose subgraphs all satisfy a $n^{1/2}$-separator theorem? The answer is no; below we present a class of such graphs for which, in one particular ND order, fill is $\Omega(n^{5/4})$.

Let k be a positive integer. We define graph V_k to have k^2 vertices $v_{11}, v_{12}, \ldots, v_{kk}$. The edges are $\{v_{ij}, v_{ij+1}\}$ for $1 \leq i \leq k$ and $1 \leq j \leq k-1$, and $\{v_{1j}, v_{ij}\}$ for $2 \leq i \leq k$ and $1 \leq j \leq k$. This graph is a k by k grid graph in which the top end of each vertical edge has been detached from its vertex and reattached to the vertex at the head of its column. Figure 9 shows V_5. "Column c" of V_k means the vertices v_{ic} for $1 \leq i \leq k$, and "row r" means the vertices v_{rj} for $1 \leq j \leq k$. This graph is sparse, but not sparse-contractible: If each row except the first is contracted into a single vertex, the result contains as a subgraph the complete bipartite with k vertices in one part and $k-1$ in the other.

Theorem 7. *All subgraphs of the graph V_k defined above satisfy a $n^{1/2}$-separator theorem for constants $\alpha = 2/3$ and $\beta = 3$.*

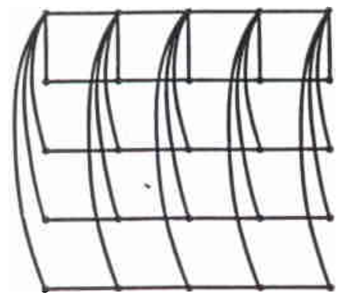

Fig. 9. Graph V_5 for Theorem 7

Proof. Let G be a subgraph of V_k with n vertices. There is at least one column c such that removal of all the vertices of G in that column separates G into two parts with at most $n/2$ vertices in each. If column c has fewer than $n^{1/2}$ vertices of G, it is the required separator. Otherwise let $c_1 \leq c$ be the minimum and $c_2 \geq c$ be the maximum such that every column from c_1 through c_2 contains at least $n^{1/2}$ vertices of G. If columns $c_1 - 1$ and $c_2 + 1$ are deleted (which deletes at most $2n^{1/2} - 2$ vertices of G), then G falls into parts A_1, B, and A_2, with at most $n/2$ vertices in each of A_1 and A_2. If A_1 has at least $n/3$ vertices then column $c_1 - 1$ is the required separator; similarly for A_2 and column $c_2 + 1$. Suppose that A_1 and A_2 each have less than $n/3$ vertices. Removal of vertices v_{1c_1} through v_{1c_2} causes B to fall into separate rows, and removal of at most one more vertex from one of those rows ensures that no remaining component of B has more than $n/2$ vertices. Since G has only n vertices and each of columns c_1 through c_2 has at least $n^{1/2}$ of them, $c_2 - c_1$ is less than $n^{1/2}$ and this last step deletes at most $n^{1/2} + 1$ vertices. Thus in all at most $3n^{1/2} - 1$ vertices of G have been deleted. These vertices are the separator. □

Theorem 8. *Let k be a positive integer, let $n = k^2$, and let V_k be the n-vertex graph defined above. If the ND algorithm is run on V_k and an appropriate adversary is allowed to choose the separators and the order of elimination within each separator, then $\Omega(n^{5/4})$ fill will occur.*

Proof. The adversary's choice of a top-level separator is the first row of V_k. Removal of these $n^{1/2}$ vertices leaves $k-1$ components, each of which is a row with k vertices. The adversary chooses to separate each row by removing its middle $k^{1/2}$ vertices; more precisely, row i is separated by removing $v_{is}, \ldots, v_{it}$ where $s = \lceil (k - k^{1/2})/2 \rceil$ and $t = \lceil (k + k^{1/2})/2 \rceil$. (This is definitely not a good separator.) These vertices are eliminated in order from v_{is} to v_{it}. The rest of the adversary's choices do not matter.

Consider the fill between vertices in the first row and vertices in the separator of another row, say row 2. Let v_{1i} be a vertex in the first row with $i < s$, and let v_{2j} be a vertex in the separator of the second row, so $s \leq j \leq t$. All of $v_{2i}, \ldots, v_{2j-1}$ are eliminated before v_{1i} or v_{2j}, so $\{v_{1i}, v_{2j}\}$ is a fill edge. The number of such edges is $(s-1)(t-s+1)$, which is $\Theta(k^{3/2})$. Similar fill edges appear from each of rows 3 through k, for a total of $\Theta(k^{5/2})$ edges. Since $n = k^2$, this is $\Theta(n^{5/4})$ fill just between separators on levels 0 and 1. □

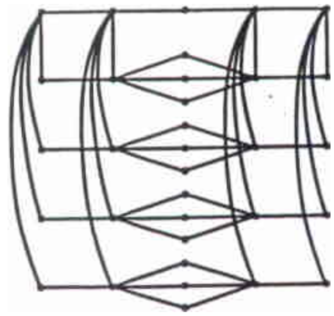

Fig. 10. Graph V_5 modified to use minimal separators

This example is a little bit fishy. A minimal separator for an individual row is a single vertex, and with this choice the fill is only $O(n \lg n)$. We can modify the example by duplicating the middle vertex of each row $k^{1/2}$ times as shown in Fig. 10. This graph and its subgraphs still satisfy a $n^{1/2}$-separator theorem. Now each row has a minimal separator with $k^{1/2}$ vertices, and fill is still $\Omega(n^{5/4})$ using these separators. However, it is still clear that the best row separator to choose has only one or two vertices.

The Lipton-Rose-Tarjan algorithm will give $O(n \lg n)$ fill for this graph, since its subgraphs all satisfy a $n^{1/2}$-separator theorem. Recall that that algorithm does not do a recursive call for every component into which the separator divides the graph, but collects the components into two pieces, adds the separator vertices to each piece, and does exactly two recursive calls. If the LRT algorithm called itself for each component, the adversary could make the same choice of separators for this example and fill would still be $\Omega(n^{5/4})$. Therefore the fact that there are exactly two recursive calls is essential to the fill bound for the LRT algorithm. On the other hand, making the ND algorithm do exactly two recursive calls would not limit fill to $O(n \lg n)$: The adversary could choose row 1 as the top-level separator, and then choose about $\lg k$ levels of empty separators until each subgraph was a single row of V_k.

As we observed above, the ND algorithm gives $O(n^{3/2})$ fill on any graph with a $n^{1/2}$-separator decomposition. We do not know whether there are such graphs for which fill is asymptotically more than $n^{5/4}$.

10. Conclusion

We have presented and analyzed a nested dissection algorithm that can be considered either as an extension of the George-Liu heuristic nested dissection, or as a variation of the Lipton-Rose-Tarjan generalized nested dissection. The LRT algorithm achieves $O(n \lg n)$ fill on graphs whose subgraphs all have $n^{1/2}$-separators; the ND algorithm presented here achieves $O(n \lg n)$ fill bounds for graphs with $n^{1/2}$-separator decompositions and either bounded degree or sparse contractions. Both classes include planar and finite-element graphs, which are probably the most useful applications of nested dissection with $n^{1/2}$-separators.

The constant factor in the ND fill bound is a bit less than 2/3 of that in Lipton, Rose, and Tarjan's analysis of the LRT algorithm [14]. Both constants are probably overestimates, and which (if either) version will be more practical can probably be decided only by experiment. Experiments on the ND algorithm are underway, using the Waterloo Sparspak sparse matrix package developed by Alan George and his colleagues [4]. Jean Roman [19] independently proved an $O(n \lg n)$ fill bound for the ND algorithm on planar graphs of bounded degree; his experiments suggest that the algorithm may be practical.

Another class of graphs with sparse contractions and $n^{1/2}$-separators is graphs of genus bounded by a constant g. Such graphs may be useful in using finite element methods to solve problems on the surface of a three-dimensional object with holes in it. Graphs of genus at most g satisfy a $(gn)^{1/2}$-separator theorem [6], and the separator can be found in $O(n+g)$ time from an embedding of the graph. This leads immediately to an $O(gn \lg n)$ fill bound for the ND algorithm. A graph of genus g has a *plane reducer* (which is a set of vertices whose removal leaves a planar graph) of size $O((gn)^{1/2} \lg g)$ [6]. This gives an $O(n \lg n + gn \lg^2 g)$ fill bound by a proof like that in Gilbert [5, Sect. 2.9]. If, as seems likely, graphs of genus g actually have plane reducers of $O((gn)^{1/2})$ vertices, this fill bound could be improved to $O(n \lg n + gn)$.

We do not know of a class of graphs that is sparse-contractible and does not satisfy a $n^{1/2}$-separator theorem. It is interesting to ask whether sparse contractibility implies a nontrivial separator theorem.

Miller [16] has proved a new version of the planar separator theorem: A maximal planar graph (one in which every face has three sides) has a $n^{1/2}$-separator that is a simple cycle. This may have theoretical or practical implications for nested dissection algorithms.

Throughout this paper we have avoided discussing numerical problems by assuming that the coefficient matrix is symmetric and positive definite. Gilbert and Schreiber [5, 8] have investigated an algorithm that combines a version of nested dissection with partial pivoting for stability. They show that this dissection pivoting algorithm limits fill to $O(n \lg n)$ on some classes of indefinite matrices with symmetric structure, including planar and finite-element graphs of bounded degree.

Appendix. Recurrences on the Separator Tree

The following lemmas solve a recurrence relation that will let us bound sums over the nodes of the separator tree.

Lemma 11. *Let $f(x)$ be a real-valued function that is continuous on the closed interval $[0, 1]$ and twice differentiable on the open interval $(0, 1)$, and suppose that $f(0)=0$ and $f''(x)>0$ for $0<x<1$. Let α be a real number with $\frac{1}{2} \leqq \alpha < 1$.*

Consider sequences $\{x_1, \dots, x_k\}$ of real numbers that satisfy $0 \leqq x_i \leqq \alpha$ for $1 \leqq i \leqq k$ and $\sum x_i = 1$. The maximum over all k and all such sequences $\{x_i\}$ of

$$\sum_{1 \leq i \leq k} f(x_i)$$

is attained when $k=2$, $x_1=1-\alpha$, *and* $x_2=\alpha$.

Remark. If $f''(x)<0$, the minimum is attained when $k=2$, $x_1=1-\alpha$, and $x_2=\alpha$.

Proof. The proof is based on the following Fact. If $0 \leqq a \leqq b \leqq c \leqq d \leqq 1$ and $a+d=b+c \leqq 1$, then $f(a)+f(d) \geqq f(b)+f(c)$. To prove the Fact, let $g(x)=f(x)+f(a+d-x)$. Then $g''(x)=f''(x)+f''(a+d-x)$, which is positive for $a \leqq x \leqq d$. Therefore the only maxima of $g(x)$ in this interval are at its endpoints, $x=a$ and $x=d$, and by definition $g(a)=g(d)$. Therefore $g(a) \geqq g(b)$. Substituting into the definition of $g(x)$ gives $f(a)+f(d) \geqq f(b)+f(c)$.

The proof of the lemma consists of three applications of the Fact. Let $\{x_1, \ldots, x_k\}$ be a sequence of real numbers with $0 \leqq x_1 \leqq \ldots \leqq x_k \leqq \alpha$ and $\sum_i x_i = 1$. We shall shorten the sequence without increasing $\sum_i f(x_i)$ until $k=2$, and then we shall see that $x_1=1-\alpha$ and $x_2=\alpha$ minimize the sum in that case.

First, suppose that $k>2$ and $x_1+x_2 \leqq \frac{1}{2}$. The Fact says that $f(x_1)+f(x_2)$ is at most $f(0)+f(x_1+x_2)$, which is equal to $f(x_1+x_2)$. Therefore we can get a shorter sequence by replacing x_1 and x_2 with x_1+x_2.

Second, suppose that $k>2$ and $x_1+x_2>\frac{1}{2}$. If k were at least 4, this would imply that $x_1+x_2+x_3+x_4>1$, contrary to the hypothesis of the lemma. Thus $k=3$, $x_1 \leqq x_2 \leqq x_3 \leqq \frac{1}{2}$, and $x_1+x_2+x_3=1$. Now $x_1+x_3-\frac{1}{2} \geqq 0$, so the Fact says that $f(x_1)+f(x_3) \leqq f(x_1+x_3-\frac{1}{2})+f(\frac{1}{2})$. The Fact also says that $f(x_2)+f(x_1+x_3-\frac{1}{2}) \leqq f(0)+f(\frac{1}{2})$. Adding the last two inequalities and cancelling terms gives $f(x_1)+f(x_2)+f(x_3) \leqq 2f(\frac{1}{2})$. Therefore we can get a shorter sequence by replacing $\{x_1, x_2, x_3\}$ with $\{\frac{1}{2}, \frac{1}{2}\}$.

Finally, suppose that $k=2$. Then $1-\alpha \leqq x_1 \leqq x_2 \leqq \alpha$, so the Fact says that $f(x_1)+f(x_2) \leqq f(1-\alpha)+f(\alpha)$. □

Lemma 12. *Let G be an n-vertex graph with a $n^{1/2}$-separator decomposition (for constants $\alpha<1$, $\beta>0$, and $n_0>0$), and let f be a nonnegative real-valued function on the nodes of the separator tree for G. Suppose that f is zero on external nodes, and $f(N) \leqq cm^\lambda$ whenever N is an internal node whose subtree contains m vertices (not nodes), for some $\lambda>0$. Then the sum of $f(N)$ over all nodes N of the separator tree is at most*

$$\begin{cases} \dfrac{c}{1-\alpha^{\lambda-1}} n^\lambda & \text{if } \lambda>1, \\[2ex] \dfrac{c}{-\alpha \lg \alpha - (1-\alpha)\lg(1-\alpha)} n \lg n & \text{if } \lambda=1, \\[2ex] O(n) & \text{if } \lambda<1. \end{cases}$$

Proof. Let $p(n)$ be the maximum of $\sum f(N)$ over all separator trees for all such graphs with n vertices. Then

$$p(n)=0 \qquad \text{for } 0\leq n\leq n_0,$$
$$p(n)\leq \max\{p(n_1)+\ldots+p(n_k)\}+cn^\lambda \qquad \text{for } n>n_0, \tag{1}$$

where the maximum is taken over all k and $n_1, \ldots, n_k$ satisfying

$$n_1+\ldots+n_k=n, \quad 0\leq n_i\leq \alpha n \quad \text{for } 1\leq i\leq k. \tag{2}$$

We may assume that the sum of the n_i is exactly n because $p(n)$ is nonnegative, so the sum in Eq. (1) is not made smaller by adding more terms. Now we consider the recurrence case by case.

Case 1. $\lambda>1$. Proof is by induction on n. The lemma holds for $0\leq n\leq n_0$ since then $p(n)=0$. If $n>n_0$, there is a set $\{n_i\}$ satisfying Eq. (2) such that

$$p(n)\leq \frac{c}{1-\alpha^{\lambda-1}} \sum_{1\leq i\leq k} n_i^\lambda+cn^\lambda. \tag{3}$$

The sum is at most $(\max_i n_i)^{\lambda-1}\sum_i n_i$, which is at most $(\alpha n)^{\lambda-1}n$ or $\alpha^{\lambda-1}n^\lambda$. Substituting this into Eq. (3) gives $p(n)\leq cn^\lambda/(1-\alpha^{\lambda-1})$.

Case 2. $\lambda=1$. Proof is again by induction, and the case $n\leq n_0$ is again trivial. If $n>n_0$, there is a set $\{n_i\}$ satisfying Eq. (2) such that

$$p(n)\leq \frac{c}{-\alpha\lg\alpha-(1-\alpha)\lg(1-\alpha)} \sum_{1\leq i\leq k} n_i\lg n_i+cn. \tag{4}$$

Define the function $f(x)=xn\lg xn$ to be 0 at $x=0$, which makes it continuous there. Then Lemma 11 applies, and says that

$$\sum_{1\leq i\leq k} n_i\lg n_i\leq \alpha n\lg\alpha n+(1-\alpha)n\lg(1-\alpha)n.$$

Substituting this into Eq. (4) gives the desired result.

Case 3. $\lambda<1$. If $n>n_0$, there is a set $\{n_i\}$ satisfying Eq. (2) such that

$$p(n)\leq \sum_{1\leq i\leq k} p(n_i)+cn^\lambda. \tag{5}$$

In this case we transform Eq. (5) to get rid of the cn^λ term. The second derivative of $f(x)=x^\lambda$ is negative. Thus by the remark in Lemma 11, $n_1^\lambda+\ldots+n_k^\lambda$ is at least $(\alpha n)^\lambda+((1-\alpha)n)^\lambda$. Let $\delta=\alpha^\lambda+(1-\alpha)^\lambda-1$. Then $\delta>0$, and $n_1^\lambda+\ldots+n_k^\lambda\geq(1+\delta)n^\lambda$. This inequality can be written as

$$\frac{c}{\delta}n^\lambda\leq \sum_{1\leq i\leq k}\frac{c}{\delta}n_i^\lambda-cn^\lambda. \tag{6}$$

Now we add Eqs. (5) and (6), and define $q(n)=p(n)+cn^\lambda/\delta$ in the sum. The result is

$$q(n) \leq \sum_{1 \leq i \leq k} q(n_i) \quad \text{for } n > n_0. \tag{7}$$

The solution to recurrence (7) is linear, so $q(n) = O(n)$ and $p(n) = O(n)$. □

Lemma 13. *Let G and f be as in Lemma 12, and suppose $f(N) \leq cm$ (that is, $\lambda = 1$). Then the sum of $f(N)$ over all nodes N on levels 0 through $k-1$ of the separator tree is at most ckn.*

Proof. Let $\mathcal{N}_i$ be the set of nodes on level i of the separator tree, and let n_N be the number of vertices of G in the subtree rooted at N. Then the sum in the statement of the lemma is equal to.

$$\sum_{0 \leq i < k} \sum_{N \in \mathcal{N}_i} c n_N.$$

The subtrees rooted on level i are disjoint, so the inner sum above is at most cn. Therefore the whole sum is at most ckn. □

References

1. Djidjev, H.N.: On the problem of partitioning planar graphs. SIAM J. Algebraic Discrete Methods 3, 229–240 (1982)
2. George, A.: Nested dissection of a regular finite element mesh. SIAM J. Numer. Anal. 10, 345–363 (1973)
3. George, A., Liu, J.W.-H.: An automatic nested dissection algorithm for irregular finite element problems. SIAM J. Numer. Anal. 15, 1053–1069 (1978)
4. George, A., Liu, J.W.-H.: Computer Solution of Large Sparse Positive Definite Systems. Englewood Cliffs, NJ: Prentice-Hall 1981
5. Gilbert, J.R.: Graph Separator Theorems and Sparse Gaussian Elimination. Ph.D. thesis, Stanford University 1980
6. Gilbert, J.R., Hutchinson, J.P., Tarjan, R.E.: A separator theorem for graphs of bounded genus. J. Algorithms 5, 391–407 (1984)
7. Gilbert, J.R., Rose, D.J., Edenbrandt, A.: A separator theorem for chordal graphs. SIAM J. Algebraic Discrete Methods 5, 306–313 (1984)
8. Gilbert, J.R., Schreiber, R.: Nested dissection with partial pivoting. Sparse Matrix Symposium, Fairfield Glade, Tennessee 1982
9. Harary, F.: Graph Theory. Reading, MA: Addison-Wesley 1969
10. Hoey, D., Leiserson, C.E.: A layout for the shuffle-exchange network. Carnegie-Mellon University Computer Science Department technical report CMU-CS-80-139 (1980)
11. Jordan, C.: Sur les assemblages de lignes. J. Reine Angew. Math. 70, 185–190 (1869)
12. Leighton, F.T.: A layout strategy for VLSI which is provably good. Proc. 14th Ann. ACM Symp. Theory Comput. pp. 85–98 (1982)
13. Leiserson, C.E.: Area-efficient graph layouts (for VLSI). Proc. 21st Ann. Symp. Found. Comput. Sci. pp. 270–281 (1980)
14. Lipton, R.J., Rose, D.J., Tarjan, R.E.: Generalized nested dissection. SIAM J. Numer. Anal. 16, 346–358 (1979)
15. Lipton, R.J., Tarjan, R.E.: A separator theorem for planar graphs. SIAM J. Appl. Math. 36, 177–189 (1979)
16. Miller, G.L.: Finding small simple cycle separators for 2-connected planar graphs. Proc. 16th Ann. ACM Symp. Theory Comput. pp. 376–382 (1984)

17. Nash-Williams, C.St.J.A.: Decomposition of finite graphs into forests. J. Lond. Math. Soc. 39, 12 (1964)
18. Parter, S.: The use of linear graphs in Gauss elimination. SIAM Rev. 3, 119–130 (1961)
19. Roman, J.: Calculs de complexité relatifs à une méthode de dissection emboîtée. Numer. Math. 47, 175–190 (1985)
20. Rose, D.J.: A graph-theoretic study of the numerical solution of sparse positive definite systems of linear equations. In: Graph Theory and Computing (R.C. Read, ed.), pp. 183–217. New York: Academic Press 1972
21. Rose, D.J., Tarjan, R.E., Lueker, G.S.: Algorithmic aspects of vertex elimination on graphs. SIAM J. Comput. 5, 266–283 (1976)
22. Yannakakis, M.: Computing the minimum fill-in is NP-complete. SIAM J. Algebraic Discrete Methods 2, 77–79 (1981)

Received August 15, 1984/August 12, 1986