Some Nested Dissection Order is Nearly Optimal

John R. Gilbert*

TR 86-767
August 1986

Department of Computer Science
Cornell University
Ithaca, NY 14853

* Computer Science Department, Cornell University, Ithaca, NY 14853. Research partially supported by the National Science Foundation under grant DCR-8451385, and by IBM Corporation under research contract 573616.

Some Nested Dissection Order is Nearly Optimal

John R. Gilbert*
August 1986

Abstract. The *minimum fill problem* is to reorder the rows and columns of a given sparse symmetric matrix so that its triangular factor is as sparse as possible. Equivalently, it is to find the smallest set of edges whose addition makes a given undirected graph chordal. The problem is known to be NP-complete, and no polynomial-time approximation algorithms are known that provide any nontrivial guarantee for arbitrary graphs (matrices), although some heuristics perform well in practice.

Nested dissection is one such heuristic. In this note we prove that every graph with a fixed bound on vertex degree has a nested dissection order that achieves fill within a factor of $O(\log n)$ of minimum. This does not lead to a polynomial-time approximation algorithm, however, because the proof does not give an efficient method for finding the separators required by nested dissection.

---

* Computer Science Department, Cornell University, Ithaca, NY 14853. Research partially supported by the National Science Foundation under grant DCR-8451385, and by IBM Corporation under research contract 573616.

1. Introduction. Let G be a connected, undirected graph with n vertices. An *elimination order* on G is an ordering of the vertices which we will write as a one-to-one function $\alpha : V(G) \rightarrow 1, \ldots, n$. The *filled graph* G^*_α of G with respect to α is the graph whose vertices are the vertices of G, and whose edges are all those edges $\{v, w\}$ such that there is a path $v = v_1, v_2, \ldots, v_k = w$ in G with $\alpha(v_i) < \min(\alpha(v), \alpha(w))$ for $1 < i < k$. Thus G^*_α is G plus some extra edges. The *fill* due to α is the number of edges in G^*_α.

The filled graph is useful in algorithms for sparse Cholesky factorization, which is a version of Gaussian elimination used to solve systems of linear equations whose coefficient matrices are symmetric and positive definite [15]. Suppose that G is the graph whose adjacency matrix is the coefficient matrix A, with vertex v corresponding to row and column $\alpha(v)$, and that L is the Cholesky factor of A (that is, L is the lower triangular matrix with $LL^T = A$). Then $L + L^T$ is the adjacency matrix of the filled graph G^*_α. To perform Cholesky factorization efficiently in time and space, it is desirable to order the rows and columns of A so that L will have as few nonzeros as possible; this corresponds to choosing an order α for which G^*_α has as few edges as possible.

The *minimum fill problem* is, given G, to find α such that the number of edges in G^*_α is minimum. Yannakakis [17] proved this problem NP-complete. Little is known about approximation algorithms for the minimum fill problem. Like graph coloring and the maximum clique problem, a trivial algorithm gets within a factor of n of minimum; unlike those problems, no polynomial-time algorithm is known that guarantees a solution within even a factor of $O(n/\log n)$ of minimum. A greedy heuristic called *minimum degree* [3] performs well in practice, but is known not to guarantee fill within a constant factor of minimum on arbitrary graphs [8]. Another heuristic called *nested dissection* [3] performs well in practice on some kinds of matrices, and guarantees fill at most $O(n \log n)$ on planar graphs, two-dimensional finite element graphs, or graphs with a fixed bound on genus [7, 16]. (Incidentally, it remains an open problem to determine whether minimum degree guarantees $O(n \log n)$ fill on all planar graphs, or even on a regular square grid graph.)

Nested dissection is a divide-and-conquer ordering scheme that uses cutsets called *separators*, which are defined below. The theorem in this paper is that in any graph of bounded vertex degree, the correct choice of separators leads to a nested dissection ordering that guarantees fill within a factor of $O(d \log n)$ of minimum, where d is the degree bound. Unfortunately, this does not lead to a polynomial-time approximation algorithm, because the proof does not give an efficient algorithm to find the separators in question. Various heuristics are known to find separators, however [1, 2, 10], and this result at least suggests that good separator heuristics are the right direction to look for good fill heuristics. Section 3 discusses directions in which it would be useful to strengthen this result.

If G is a graph with n vertices, a *separator* for G is a set of vertices whose removal leaves no connected component with more than $n/2$ vertices. Every graph has trivial separators with $O(n)$ vertices; classes of graphs with nontrivial separators include trees [9], planar graphs [13], graphs of bounded genus [5], chordal graphs [6], and hypercubes [4]. A *balanced separator decomposition* is a partition of the vertices of G into a tree of separators, as follows. The root of the tree is a separator C for G. The subtrees are balanced separator decompositions of the connected components of $G - C$. (The notation

$G - C$ means the subgraph of G induced by the vertices that are not members of C.)

Separators can be used to construct elimination orders in a divide-and-conquer algorithm called *nested dissection*, proposed by Alan George and extended and analyzed by George, Liu, Lipton, Rose, Tarjan, and Gilbert [3, 7, 12]. A *nested dissection order* based on a balanced separator decomposition with root C is an order that assigns the highest numbers to the vertices of C (in any order), and numbers each subtree in nested dissection order (recursively) with consecutive numbers. Intuitively, nested dissection orders limit fill because no fill edges can join vertices in different subtrees of the tree of separators. For example, planar graphs have nested dissection orders that limit fill to $O(n \log n)$ [7, 12]. Nested dissection orders are also useful for parallel sparse Cholesky factorization, since they divide the problem into independent subproblems coupled only through the separators [14].

A graph is *chordal* if every cycle of length at least 4 has a *chord*, which is an edge joining two vertices that are not consecutive on the cycle. Rose [15] showed that a graph G is chordal if and only if it has an order α for which $G_\alpha^* = G$, that is, an order for which no fill occurs. Finding a minimum fill order for an arbitrary graph is the same as finding the smallest set of edges whose addition makes the graph chordal.

Gilbert, Rose and Edenbrandt [6] proved several separator theorems for chordal graphs and weighted chordal graphs. The version we need is a special case of their Theorem 3.

Theorem 1 [6]. *Let G be a chordal graph with n vertices and m edges. Then G has a separator C of $O(\sqrt{m})$ vertices such that every connected component of $G - C$ has at most $n/2$ vertices and at most $2m/3$ edges.* ∎

(The theorem holds with ϵn in place of $n/2$ and ϵm in place of $2m/3$, for any fixed positive ϵ.)

2. Main result.

Theorem 2. *Let G be a connected, undirected graph with n vertices and maximum degree d. There exists a balanced separator decomposition for G such that any nested dissection order based on that decomposition gives a fill within a factor of $O(d \log n)$ of the minimum fill for G.*

Proof. Let α be a minimum fill order for G. Let $H = G_\alpha^*$ be the filled graph for α, and suppose H has h edges.

Form a separator decomposition for H as follows. Use Theorem 1 to find a separating set C of size $O(\sqrt{h})$ for H, such that every component of $H - C$ has at most $n/2$ vertices and at most $2h/3$ edges. Let C be the root of the decomposition, and form similar decompositions of the components of $H - C$ recursively. This decomposition is a balanced separator decomposition for G. Let β be a nested dissection order based on it. Let $K = G_\beta^*$ be the filled graph for β, and suppose K has k edges.

We can bound k by an argument similar to that in Gilbert [4, Section 2.6]. It is easy to show [7] that if $\{v, w\}$ is an edge of K, with v in node A_v of the separator decomposition and w in node A_w and $\beta(v) < \beta(w)$, then A_v is a descendant of A_w in the tree of separators, and there is an edge $\{x, w\}$ of G with x in a node A_x that is a descendant of A_v. (We include a node itself among its descendants.)

First we count edges whose higher-numbered endpoints are in C. Each edge $\{v,w\}$ of K with $w \in C$ corresponds to some edge $\{x,w\}$ of G with A_x a descendant of A_v. Each such $\{x,w\}$ of G corresponds to at most one $\{v,w\}$ for each vertex v of G in a node on the tree path from A_x to $C = A_w$. This tree path contains at most one node on each level of the tree, so the total number of vertices of G on the path is

$$O(\sqrt{h} + \sqrt{2h/3} + \sqrt{4h/9} + \cdots) = O(\sqrt{h}).$$

Thus one edge $\{x,w\}$ of G corresponds to $O(\sqrt{h})$ edges $\{v,w\}$ of K. There are $O(\sqrt{h})$ vertices in C, each with at most d incident edges, so there are at most $O(dh)$ edges of K with higher-numbered endpoints in C.

Each edge of K either has an endpoint in C or is an edge of a component of $K - C$, so k is bounded by the solution to the recurrence

$$k(h,n) = O(dh) + \sum_i k(h_i, n_i),$$

where the sum is over the components of $K - C$, which satisfy $h_i \leq 2h/3$, $\sum_i h_i \leq h$, $n_i \leq n/2$, and $\sum_i n_i \leq n$. This is actually a one-variable recurrence—it does not depend on n—and its solution is $k(h,n) = O(dh \log h)$. Since h is less than n^2, this implies $k = O(dh \log n)$, so k is within $O(d \log n)$ of minimum. ∎

3. Remarks. Many open questions remain, both about this result and about approximate minimum fill algorithms in general. Can this result be tightened by a factor of $\log n$? Not every graph admits a balanced decomposition for which nested dissection gives exactly the minimum possible fill; an example is an n-vertex path, for which nested dissection gives a filled graph with nearly twice as many edges as minimum. However, I do not know of a class of graphs for which every balanced decomposition gives a filled graph that is too large by more than a constant factor. We may also ask whether the degree bound is necessary in the statement of the theorem.

The main disadvantage of this result is that it gives no guidance in finding the separators that give a good nested dissection order. A slightly stronger result than this one would be to prove that an approximation algorithm for the minimum-size separator problem would necessarily give an approximation algorithm for the minimum fill problem. This would be a step closer to finding a polynomial-time algorithm that guaranteed fill within a nontrivial factor (i.e., $o(n)$) of minimum.

Fill is not the only criterion for a good ordering for sparse Cholesky factorization. Another is arithmetic operation count, which is proportional to

$$\sum_{\text{vertices } v} (|\{w : \alpha(w) > \alpha(v) \text{ and } \{v,w\} \text{ is an edge of } G^*_\alpha\}|)^2.$$

It may be possible to prove a result similar to this one for the minimization problem for this measure; in practice, low fill and low operation count usually go together.

A third criterion is *parallel depth*, which is a measure of the degree of serial dependency in a parallel elimination ordering. Parallel depth has been studied by Liu [14]; it can be defined as the length of the longest path in G^*_α that is monotone increasing in α. In contrast with operation count, orders that give small fill do not necessarily give small depth. For example, the minimum-fill orders for a path of n vertices give depth at least $n/2$, while a nested dissection order gives depth $O(\log n)$. The parallel depth of an order is at least as large as the largest clique in G^*_α; this can be used to show that every graph admits a balanced decomposition whose nested dissection orderings have parallel depth within a factor of $O(\log n)$ of minimum. A more detailed investigation of minimizing parallel depth would be interesting.

References.

[1] Thang Bui, Soma Chaudhuri, Tom Leighton, and Mike Sipser. Graph bisection algorithms with good average case behavior. *Proceedings of the 25th Annual Symposium on Foundations of Computer Science*, pages 181–192, 1984.

[2] Alan George and Joseph W. H. Liu. An automatic nested dissection algorithm for irregular finite element problems. *SIAM Journal on Numerical Analysis* **15**: 1053–1069, 1978.

[3] Alan George and Joseph W-H Liu. *Computer Solution of Large Sparse Positive Definite Systems*. Prentice-Hall, 1981.

[4] John Russell Gilbert. *Graph Separator Theorems and Sparse Gaussian Elimination*. Ph. D. thesis, Stanford University, 1980.

[5] John R. Gilbert, Joan P. Hutchinson, and Robert Endre Tarjan. A separator theorem for graphs of bounded genus. *Journal of Algorithms* **5**: 391–407, 1984.

[6] John R. Gilbert, Donald J. Rose, and Anders Edenbrandt. A separator theorem for chordal graphs. *SIAM Journal on Algebraic and Discrete Methods* **5**: 306–313, 1984.

[7] John R. Gilbert and Robert Endre Tarjan. The analysis of a nested dissection algorithm. To appear in *Numerische Mathematik*.

[8] Christopher Hempel. Minimum degree can miss by more than a constant factor. Unpublished manuscript, Cornell University, 1985.

[9] Camille Jordan. Sur les assemblages de lignes. *Journal für die Reine und Angewandte Mathematik* **70**: 185–190, 1869.

[10] B. W. Kernighan and S. Lin. An efficient heuristic procedure for partitioning graphs. *Bell System Technical Journal* **49**: 291–307, 1970.

[11] M. Leuze and L. Saxton. On minimum parallel computing times for Gaussian elimination. *Congressus Numerantium* **40**: 169–179, 1983.

[12] Richard J. Lipton, Donald J. Rose, and Robert Endre Tarjan. Generalized nested dissection. *SIAM Journal on Numerical Analysis* **16**: 346–358, 1979.

[13] Richard J. Lipton and Robert Endre Tarjan. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics* **36**: 177–189, 1979.

[14] Joseph W. H. Liu. Computational models and task scheduling for parallel sparse Cholesky factorization. Computer Science Department report CS-85-01, York University, Downsview, Ontario, 1985.

[15] Donald J. Rose. A graph-theoretic study of the numerical solution of sparse positive definite systems of linear equations. In Ronald C. Read (editor), *Graph Theory and Computing*, pages 183–217, Academic Press, 1972.

[16] Donald J. Rose, R. Endre Tarjan, and George S. Lueker. Algorithmic aspects of vertex elimination on graphs. *SIAM Journal on Computing* **5**: 266–283, 1976.

[17] Mihalis Yannakakis. Computing the minimum fill-in is NP-complete. *SIAM Journal on Algebraic and Discrete Methods* **2**: 77–79, 1981.