

Fill Bounds for Nested Dissection

CS 219: Sparse Matrix Algorithms

April 15, 2013

These notes give proofs, or references to proofs, of the theorems from the April 8 lecture.

1 Preliminaries

We consider an n -by- n symmetric, positive definite matrix A and its undirected graph $G = G(A)$, which has vertices 1 through n and an edge (i, j) for each symmetric pair of off-diagonal nonzero elements $A(i, j) = A(j, i)$. The *filled graph* $G^+ = G^+(A)$ is formed by adding edges to G as follows:

```
begin with all vertices "unmarked"
for j = 1 to n
  mark vertex j
  add edges between j's unmarked neighbors
```

The edges of G^+ are called *fill edges*, whether or not they were originally edges of G . The filled graph describes the nonzero structure of the Cholesky factor of A . If the Cholesky factorization is $A = LL^T$, then $G^+(A) = G(L + L^T)$. Equality holds provided there is no coincidental cancellation during the factorization, that is, adding two nonzeros always yields a nonzero. This is usually true in floating-point arithmetic. If there is coincidental cancellation, then $G(L + L^T)$ is a subgraph of $G^+(A)$.

Path Lemma. *Let G be a graph with vertices 1 through n , and let G^+ be the filled graph. Then (v, w) is an edge of G^+ if and only if G contains a path from v to w of the form $(v, u_1, u_2, \dots, u_k, w)$ with $u_i < \min(v, w)$ for each i . This includes the possibility $k = 0$, in which case (v, w) is an edge of G and therefore of G^+ .*

Davis [1] Theorem 4.1 is the Path Lemma in different notation.

The *elimination tree* (or *etree*) $T = T(G) = T(A)$ is a rooted spanning tree of the filled graph G^+ ; it has the same vertices 1 through n , and each vertex v except the root has as its parent $p(v)$ the smallest higher-numbered neighbor of v . That is,

$$p(v) = \min\{w : w > v \text{ and } (v, w) \text{ is an edge of } G^+\}.$$

Equivalently, the parent $p(j)$ is the row number i of the first nonzero $L(i, j)$ below the diagonal in column j of the Cholesky factor L .

Liu [7] is a good survey of elimination tree theory. The following facts about etrees follow directly from the analysis in Section 4.1 of Davis.

- If G is connected, then T is a connected tree; otherwise, T is a forest with one tree for each connected component of G .
- The vertices of T are numbered in *topological order*, that is, $p(v) > v$ for all v except the root.
- T is a *depth-first spanning tree* of G^+ , that is, every edge of G^+ joins an ancestor and a descendant in T —there are no *cross edges* joining different subtrees.
- A *dag* is a directed acyclic graph: edge is directed from one endpoint to the other, and there is no directed cycle of edges. The *transitive reduction* of a dag is the subgraph created by deleting every directed edge (v, w) for which there is a directed path of two or more edges from v to w ; that is, we delete any edge that shortcuts a path.

Consider the directed graph $G(L^T)$, which is G^+ with each edge directed from its smaller to its larger endpoint. Note that $G(L^T)$ is a dag—it can have no directed cycle because L is triangular. It is a fact that the transitive reduction of $G(L^T)$ is the elimination tree T (turned into a dag by directing each edge of T from child to parent).

Etree Fill Lemma. *Let G be a graph with vertices 1 through n , let G^+ be the filled graph, and let $T = T(G)$ be the elimination tree. Then G^+ contains edge (v, w) with $v < w$ if and only if G contains some edge (u, w) with $u \leq w$ and u a descendant of v .*

Davis Theorem 4.5 is the Etree Fill Lemma in different notation. Informally, the Etree Fill Lemma says that all the fill edges are “caused by” edges of G , via the following process: Start with an edge (u, w) of G with $u < w$. Put your right finger on u , and your left finger on w . Move your right finger up the tree from u to w one vertex at a time, creating a fill edge to w from every vertex v you encounter.

A *chordal graph* is an undirected graph in which for every cycle of length more than three there is an edge joining two nonconsecutive vertices. Rose [9] showed that chordal graphs are the same as *perfect elimination graphs*, those that admit a vertex permutation for which the graph game adds no new fill edges. Every filled graph G^+ is chordal, and conversely every chordal graph arises as the filled graph of a symmetric positive definite matrix. Chordal graphs have a rich theory, some of which is useful in designing efficient sparse matrix factorization algorithms.

We want to study the asymptotic scaling of fill as n increases, simplifying the analysis by suppressing constant factors. Our notation for comparing the growth rate of functions of n is $O()$ “order at most,” $\Omega()$ “order at least,” and $\Theta()$ “order exactly.” Formally, for nonnegative functions $p(n)$ and $q(n)$, we say that $p(n) = O(q(n))$ if there is some constant $a > 0$ such that $p(n) \leq aq(n)$ for all sufficiently large n ; $p(n) = \Omega(q(n))$ if there is some constant $a > 0$ such that $aq(n) \leq p(n)$ for all sufficiently large n ; and $p(n) = \Theta(q(n))$ if there are some constants $a > 0$ and $b > 0$ such that $aq(n) \leq p(n) \leq bq(n)$ for all sufficiently large n . Notice that $p(n) = O(q(n))$ means that $q(n) = \Omega(p(n))$, and $p(n) = \Theta(q(n))$ means that both $p(n) = O(q(n))$ and $p(n) = \Omega(q(n))$.

2 Theorems

Theorem 1. *With the natural permutation, the n -vertex model problem has $\Theta(n^{3/2})$ fill.*

Proof. Let G be the graph of the model problem, a k -by- k square grid with $n = k^2$ vertices numbered from 1 to n by rows.

We first show that the elimination tree $T(G)$ is a simple path from vertex 1 (the only leaf) to vertex n (the root). Vertex v is always adjacent to vertex $v + 1$ in G , unless $v = rk$ is the last vertex in some row r of the grid. In the latter case, $v + 1 = rk + 1$ is the first vertex in row $r + 1$. Then vertices v and $v + 1$ are joined in G by the path

$$(v = rk, rk - 1, rk - 2, \dots, r(k - 1) + 1, rk + 1 = v + 1),$$

whose intermediate vertices are all lower numbered than either endpoint. The Path Lemma then implies that $(rk, rk + 1)$ is an edge of the filled graph G^+ . We conclude that $(v, v + 1)$ is an edge of G^+ for every vertex $v < n$. The etree parent of v is its smallest higher-numbered neighbor in G^+ ; thus the parent of v is $v + 1$ except for the root $v = n$, and $T(G)$ is just a path.

Now we use the Etree Fill Lemma to count the fill caused by edges of G . There are two kinds of edges in G : “horizontal” edges $(v, v + 1)$ and “vertical” edges $(v, v + k)$. Horizontal edge $(v, v + 1)$ causes no fill, because $v + 1$ is the parent of v . Vertical edge $(v, v + k)$ causes $k - 1$ fill edges $(v + 1, v + k), (v + 2, v + k), \dots, (v + k - 1, v + k)$. There are $k^2 - k$ vertical edges $(v, v + k)$ (one for each vertex v not in the last row of the grid), so in all they cause $(k^2 - k)(k - 1) = \Theta(k^3) = \Theta(n^{3/2})$ fill edges. Thus the total number of edges in G^+ is $\Theta(n^{3/2})$. \square

Theorem 2 (George [2]). *With a nested dissection permutation, the n -vertex model problem has $O(n \log n)$ fill.*

Proof. The elimination tree for a nested dissection ordering is as shown in Figure 1. The k vertices of the top-level separator are a path hanging from the root of the whole tree; then the approximately $k/2$ vertices of each second-level separator hang from the roots of the two subtrees; then beginning with the third level the entire structure repeats itself recursively, with four subtrees each corresponding to a nearly square grid half as wide as the original, each with about $n/4$ vertices in all, including $k/2$ vertices in its first separator. Since we’re ignoring the constant factor in $O()$ it suffices the case where $k = 2^d - 1$ one less than a power of 2, so that the the subgrids are exactly square with dimension $2^{d-1} - 1$.

The height of the elimination tree is about

$$k + k/2 + k/2 + k/4 + k/4 + k/8 + k/8 + \dots = 3k.$$

We begin by counting fill edges incident on the $2k - 1$ vertices in the first two levels of separators. Let v be one such vertex. The Etree Fill Lemma says that any such fill edge is caused by some edge (u, v) of G with $u < v$, and that u is a descendant of v in T . Edge (u, v) causes one fill edge for each vertex on the path up the tree from u to v ; the path is not longer than the tree height, so this is not more than $3k$ fill edges. Vertex v has at most 4 neighbors in G , so v is an endpoint of




Figure 1: Elimination tree for a nested dissection ordering on a k -by- k mesh, with $k = 10$ and $n = k^2 = 100$.

not more than $12k$ fill edges. In all, then, the $2k - 1$ vertices in the top two levels of separators are incident on at most $12k \cdot 2k = 24n$ fill edges.

Now we can write a divide-and-conquer recurrence for an upper bound $f(n)$ on the total number of fill edges. Removing the top two levels of separators leaves four subgrids with at about $n/4$ vertices each. Therefore

$$f(n) = 4f(n/4) + 24n.$$

We can use the “master method” (www.cs.ucsb.edu/~gilbert/cs140/slides/cs140-cilklecture2.pdf) to solve this recurrence, or just iterate it. Either way, the result is $f(n) = O(n \log n)$. \square

George [2] also computed the constant factors in the $O()$ notation.

As usual, it’s trickier to prove lower bounds than upper bounds. The lower bound in the next theorem says that nested dissection is asymptotically the best possible elimination order for the model problem. It’s a nice counting argument that uses the Fill Path Lemma.

Theorem 3 (Hoffman, Martin, and Rose [4]). *With any permutation, the n -vertex model problem has $\Omega(n \log n)$ fill.*

Proof. Consider an arbitrary j -by- j subgrid of the k -by- k grid (with $n = k^2$), as shown in Figure 2. Define a “special edge” as an edge of the filled graph G^+ that joins two parallel boundaries of the j -subgrid, either top and bottom or left and right (or both).




Figure 2: An example of a special edge. Left, a j -subgrid within a k -grid, with $j = 5$, $k = 20$, and $n = k^2 = 400$. Right, detail of the 5-subgrid. The dotted line is a special edge that appears at the time the striped vertex is marked in the Cholesky graph game, assuming the black vertices were marked earlier and the white vertices are still unmarked.

We claim that regardless of the elimination order, every j -subgrid for $j \geq 2$ has a special edge in G^+ . Proof: consider the first step in the Cholesky game at which some boundary of the j -subgrid has all its internal vertices (not counting corners of the grid) marked. If a vertical (left or right) boundary just had all internal vertices marked, then a special edge is created between the top and bottom boundaries; if a horizontal (top or bottom) boundary just had all internal vertices marked, then a special edge is created between the left and right boundaries.

Now we show that there are enough different special edges to give us our $\Omega(n \log n)$ lower bound. Any particular edge of G^+ is special for only one value of j . And any particular j -special edge can be special in at most j different j -subgrids. The number of j -subgrids is $(k + 1 - j)^2$; therefore, the number of j -special edges is at least $(k + 1 - j)^2/j$, and the number of special edges overall is at least

$$\sum_{j=2}^k (k + 1 - j)^2/j.$$

This is

$$(k + 1)^2 \sum_{j=2}^k 1/j - 2(k + 1) \sum_{j=2}^k 1 + \sum_{j=2}^k j,$$

which is $\Omega(k^2 \log k - 2k^2 + k^2/2) = \Omega(k^2 \log k) = \Omega(n \log n)$. □

Theorem 4. *With nested dissection, any planar graph with n vertices has $O(n \log n)$ fill.*

The bound in Theorem 4 [3, 5] follows from the Lipton/Tarjan planar separator theorem [6].

The fill bound for the 3D model problem, a k -by- k -by- k grid of $n = k^3$ vertices, comes from solving a recurrence equation similar to that in the proof of Theorem 2. That bound also generalizes to 3D finite element meshes:

Theorem 5. *With a nested dissection permutation, any “well-behaved” 3D finite element mesh with n vertices has $O(n^{4/3})$ fill.*

The tricky part of generalizing the 3D model problem to 3D finite element meshes is to define the right class of graphs, that is, to define a broad class of graphs that represent reasonable discretizations of 3D physical simulations and that also share with the 3D model problem the existence of $O(n^{2/3})$ separators. Every graph (including the complete graph K_n) can be embedded in \mathbb{R}^3 without intersecting edges—there is plenty of room in three dimensions to route edges around each other—so the obvious generalization of planar graphs to “3D embeddable graphs” doesn’t work. Miller et al. [8] finally came up with the right notion of “well-behaved” meshes in any fixed number of geometric dimensions, along with the appropriate separator theorem and an efficient algorithm to find the separators.

References

- [1] T.A. Davis. *Direct Methods for Sparse Linear Systems*. SIAM, Philadelphia, PA, 2006.
- [2] J. A. George. Nested dissection of a regular finite element mesh. *SIAM J. Numerical Analysis*, 10:345–363, 1973.
- [3] J. R. Gilbert and R. E. Tarjan. The analysis of a nested dissection algorithm. *Numer. Math.*, 50:377–404, 1987.
- [4] A. J. Hoffman, M. S. Martin, and D. J. Rose. Complexity bounds for regular finite difference and finite element grids. *SIAM J. Numer. Anal.*, 10:364–369, 1973.
- [5] R. J. Lipton, D. J. Rose, and R. E. Tarjan. Generalized nested dissection. *SIAM J. Numer. Anal.*, 16:346–358, 1979.
- [6] R. J. Lipton and R. E. Tarjan. A separator theorem for planar graphs. *SIAM J. Appl. Math.*, 36:177–189, 1979.
- [7] J. W. H. Liu. The role of elimination trees in sparse factorization. *SIAM J. Matrix Analysis and Applications*, 11:134–172, 1990.
- [8] Gary L. Miller, Shang-Hua Teng, William Thurston, and Stephen A. Vavasis. Geometric separators for finite-element meshes. *SIAM J. on Scientific Computing*, 19(2):364–386, 1998.
- [9] Donald J. Rose. Triangulated graphs and the elimination process. *Journal of Mathematical Analysis and Applications*, 32(3):597–609, 1970.