

Nested Dissection Orderings for LU Factorization with Static Pivoting*

Ali Pinar[†] Manmeet Singh[‡] Esmond Ng[§]

Introduction

We consider the LU factorization of a general sparse matrix A . We assume that the matrix has been preprocessed so that pivoting for numerical stability is not needed during the factorization phase. This is the case, for example, when the matrix has been permuted so that elements with large magnitude appear on the diagonal.

The sparsity of the LU factors depends on how the rows and columns are permuted. A popular approach is to compute a symmetric ordering P , such as nested dissection or minimum degree, using the symmetrized matrix $A + A^T$, and then apply the resulting permutation P to the rows and columns of A symmetrically. However, when A is highly unsymmetric, the quality of P , in terms of the sparsity of the LU factors, may be poor.

When P is computed using nested dissection, it is possible to improve the quality of P by refining the separators in nested dissection, as we will show below.

Structure Prediction for LU Factorization with Static Pivoting

During factorization some zeros of the matrix change to nonzeros, which we call *fill*. Representing the nonzero structure of a sparse matrix as a graph enables us to explain the fill during factorization through paths in graphs. Let $A = (a_{ij})$ be an $N \times N$ matrix with nonzeros on its diagonal. The nonzero structure of A is represented by the directed graph $G(A) = (V, E)$, where $V = \{v_1, v_2, \dots, v_N\}$, and $E = \{(v_i, v_j) : i \neq j \text{ and } a_{ij} \neq 0\}$. The effect of LU factorization on the structure of the matrix A and/or its graph $G(A) = (V, E)$ can be described by the fill graph $G^+(A) = (V, E^+)$, where

$$E^+ = \{(v_i, v_j) : i \neq j \text{ and } v_j \text{ is reachable from } v_i \text{ through vertices numbered lower than } \min(i, j)\}.$$

As can be seen from the definition of $G^+(A)$, the fill is determined by the reachabilities through lowered numbered vertices, which calls for a clever way to number vertices. Minimum degree algorithms greedily order vertices with small degrees first as an attempt to limit reachabilities. Nested dissection algorithms on the other hand find a subset of the vertices, a *separator*, removal of which disconnects the graph into two or more parts. Ordering this separator after all other vertices guarantees that no vertex from one part reaches a vertex from another part, thus there is no fill between these parts. The same procedure is applied recursively on the disconnected parts to limit the fill within the part.

When the matrix is symmetric, the graph is undirected since every edge becomes bidirectional. In this case, a separator S is a subset of the vertices that partitions the remaining vertices of the graph into two parts V_1 and V_2 so that there are no edges between V_1 and V_2 . We refer to such

*Supported by the Director, Office of Science, Division of Mathematical, Information, and Computational Sciences of the U.S. Department of Energy under contract DE-AC03-76SF00098.

[†]Corresponding author. Computational Research Division, Lawrence Berkeley National Lab, 1 Cyclotron Road, Berkeley, CA 94720. Email: apinar@lbl.gov

[‡]Department of Mechanical and Industrial Engineering, University of Illinois, Urbana-Champaign. Email: msingh1@uiuc.edu

[§]Computational Research Division, Lawrence Berkeley National Lab, Email: egng@lbl.gov




Figure 1: (a) Strong separator (b) Directed separator

separators on undirected graphs as *undirected separators*. Keeping the size of the separator small is critical; vertices in the separator are reached by practically all other vertices of the graph. The problem of finding a good separator is well-studied under the are of graph partitioning. In this work, we focus on the definition of a separator for nested dissection orderings of unsymmetric matrices to preserve sparsity during LU factorization with static pivoting.

Separators for Unsymmetric Matrices

A *strong separator* S in a directed graph is a subset of vertices that partitions the remaining vertices into two parts V_1 and V_2 so that there are no edges between V_1 and V_2 , as illustrated in Fig. 1. A strong separator permutes the matrix into the following form

$$\begin{pmatrix} A_1 & & A_{13} \\ & A_2 & A_{23} \\ A_{31} & A_{32} & A_S \end{pmatrix} \quad (1)$$

where vertices in V_1 , V_2 , and S correspond to rows/columns of submatrices A_1 , A_2 , and A_S , respectively. When we apply LU factorization to this matrix, there will be no fill in the (1,2)- and (2,1)-blocks of the matrix, and thus the internal reorderings of rows/columns in A_1 and A_2 will not affect each other. Such a separator can be obtained by finding an undirected separator on $G(A+A^T)$. This has been long used as a heuristic for nested dissection orderings for unsymmetric matrices. We show that the two problems are equivalent.

Theorem 1 *An undirected separator in $G(A+A^T)$ is equivalent to a strong separator in $G(A)$.*

We omit the proof due to space limitations.

Although a strong separator limits the fill to only nonzero blocks in (1), it is possible to reduce the fill further with alternative definitions of a separator. A *directed separator* S_d in a directed graph is a subset of vertices that partitions the remaining vertices into two parts V_1 and V_2 so that there are no edges from V_2 to V_1 . Note that V_1 and V_2 are not decoupled completely as in the case of a strong separator, since there may be edges from V_1 to V_2 .

When we use a directed separator to permute the matrix, we will obtain the following block matrix.

$$\begin{pmatrix} A_1 & A_{12} & A_{13} \\ & A_2 & A_{23} \\ A_{31} & A_{32} & A_S \end{pmatrix}$$

The LU factorization of A_1 does not cause any fill in the (2,1)-block and in A_2 . However, there may be fill in A_{12} . Although this may look like a disadvantage, it is important to note that a directed separator reduces the size of A_S , and the extra fill on A_{12} is no more than the fill that will be incurred if a strong separator is used. The following theorem formalizes how we can reduce the fill by using a directed separator instead of a strong separator.




Figure 2: Translating a strong separator to a directed separator

Theorem 2 *Let S be a strong separator for $G(A)$ and let V_1 and V_2 be the resulting disconnected parts, and let $\alpha(V)$ denote an optimal ordering for vertex set V that minimizes the fill of factorization. Let $S_r \subseteq S$ be the set of vertices that have no edges directed from them to V_1 , as illustrated in Fig. 2 $S_c = S \setminus S_r$, is a directed separator for $G(A)$ with V_1 and $V_2 \cup S_r$ as disconnected parts, and the fill induced by using the ordering $\alpha(V_1), \alpha(V_2 \cup S_r), S_c$ is less than or equal to the fill induced by using the ordering $\alpha(V_1), \alpha(V_2), S_r, S_c$.*

Proof: The fill on V_1 and S_c will not change, and the fill for $\alpha(V_2 \cup S_r)$ is less than or equal to the fill for $\alpha(V_2), \alpha(S_r)$, by definition of α . The only point we have to prove is there is no extra fill from V_1 to V_2 , due to moving S_r vertices up in the ordering. This is only possible if there exists a path between two vertices in $V_2 \cup S_r$ go through vertices in V_1 . This is not possible however because there are no edges from V_2 or S_r to V_1 . ■

It is possible to relax the separator definition even further, but we are not including these discussions due to space restrictions. We are currently implementing the proposed techniques.