

**New Pivot Selection
for Sparse Symmetric Indefinite Factorization**

by

DUANGPEN JETPIPATTANAPONG AND GUN SRIJUNTONGSIRI

Abstract

We propose a new pivot selection technique for symmetric indefinite factorization of sparse matrices. Such factorization should maintain both sparsity and numerical stability of the factors, both of which depend solely on the choices of the pivots. Our method is based on the minimum degree algorithm and also considers the stability of the factors at the same time. Our experiments show that our method produces factors that are sparser than the factors computed by MA57 [1] and are stable.

Key Words: symmetric indefinite factorization, pivot selection, sparse matrices

2010 Mathematics Subject Classification: Primary 15A23,
Secondary 65F50

1 Introduction

There are two general approaches to solving linear systems: direct and iterative methods. Iterative methods start with an initial guess and successively generate better approximate solutions at each iteration. The running time of an iterative method depends directly on the required accuracy of the solution. Direct methods, on the other hand, attempt to find the solution by a finite number of operations and usually involve factoring the matrix.

Our work here is on direct methods for solving linear systems where the matrix is symmetric indefinite and sparse, which has many applications in linear and nonlinear optimization and finite element computation, for example. Symmetric indefinite factorization (SIF) is not unique as the resulting factors depend on the choices of the pivots during the factorization. Pivots should be chosen such that the resulting factors are stable and do not have many fill-ins—the entries that are zeros in the original matrix but are nonzeros in the factors.

There are many heuristic techniques for selecting pivots to minimize the number of fill-ins for the related problem of Cholesky factorization, which is the most suitable factorization for symmetric positive definite matrices, in literature. We briefly discuss a few such well-known techniques here since some of their ideas are also applicable to SIF. These ordering algorithms can be classified into three classes: local, global, and hybrid approaches. Local approach such as the minimum degree and the minimum fill algorithms [2, 3, 4, 5, 6] selects the pivot that is expected to minimize the number of fill-ins at each factorization step in a greedy fashion. Global approach such as Cuthill-McKee and nested dissection methods [7, 8, 9] selects pivots by considering the overall structure of the matrix. Hybrid approach, on the other hand, combines the ideas from both local and global approaches.

The well-known minimum degree algorithm [2] chooses the column that has the minimum off-diagonal nonzero elements in the remaining matrix as the pivot for the current step. Different improvements of the minimum degree algorithm have been proposed [5] such as multiple minimum degree [10] and approximate minimum degree algorithms [11] and become the practical standard in the implementations.

Another famous pivot selection algorithm is nested dissection [8]. By defining a graph whose vertices represent each column of the matrix and whose edges represent nonzero entries in the matrix, nested dissection recursively find a separator—a set of vertices that partitions the graph into two disconnected subgraphs—and ordering the pivots recursively with the two subgraphs first followed by the separator vertices. Cuthill-McKee [7] propose another pivot selection algorithm that aims to reduce the bandwidth of the matrix based on breadth first search of the structure graph.

The main difference between Cholesky factorization and SIF is in the size of pivots. For SIF, each pivot can be either a scalar or a 2-by-2 matrix while pivots in Cholesky factorization are all scalar. Moreover, unlike Cholesky factorization, the choice of pivots in SIF also affects the stability of the resulting factors [12].

There are many pivot selection algorithms proposed specifically for SIF such as Bunch-Parlett [13], Bunch-Kaufman [14], and bounded Bunch Kaufman (BBK) [15] algorithms. Bunch-Parlett method searches the whole remaining submatrix at each stage for the largest-magnitude diagonal and the largest-magnitude off-diagonal. It chooses the largest-magnitude diagonal as the 1-by-1 pivot if the resulting growth rate is acceptable. Otherwise, it selects the largest-magnitude off-diagonal and its relative diagonal elements as the 2-by-2 pivot block. This method requires $O(n^3)$ comparisons and yields a matrix L whose maximum element is bounded by 2.781. Bunch-Kaufman pivoting strategy searches for the largest-magnitude off-diagonal elements of at most two columns for each iteration. It requires $O(n^2)$ comparisons, but the elements in L are unbounded. BBK combines the two above strategies. By monitoring the size of the elements in L , BBK uses the Bunch-Kaufman strategy when it yields modest element growth. Otherwise, it repeatedly searches for an acceptable pivot. In average cases, the total cost of BBK is the same as Bunch-Kaufman, but in the worst cases its cost can be the same as that of the Bunch-Parlett strategy.

Additionally, there are other types of techniques for solving symmetric indefinite linear systems. Paige and Saunders [16] propose two algorithms, SYMMLQ and MINRES, for solving such systems. The algorithms apply orthogonal factorization together with the conjugate gradient method to solve the system. Duff et al. [17] propose a pivotal strategy for decomposing sparse symmetric indefinite matrices that limits the magnitude of the element in the factors for stability. Duff and Reid [18] propose a multifrontal method to solve indefinite sparse symmetric linear systems based on minimum-degree ordering. The multifrontal approach is widely used in many sparse direct solvers such as MA57 and MUMPS [1, 19].

We propose a new pivot selection algorithm for sparse SIF. Our algorithm applies the idea of minimum-degree ordering to consider both 1-by-1 and 2-by-2 pivots while also considers the stability of the resulting factors. Our experiments show that our algorithm produces stable factors that are sparser than the factors produced by [1].

For the rest of the article, we describe symmetric indefinite factorization in Section 2. Section 3 explains the minimum-degree ordering algorithm. Section 4 describes our algorithm. Section 5 shows our experiment and the results. Finally, we conclude the article

in Section 6.

2 Symmetric Indefinite Factorization

Solving a linear system

$$Ax = b, \quad (2.1)$$

where $A \in \mathbb{R}^{n \times n}$ is symmetric indefinite, is generally done by first obtaining the symmetric indefinite factorization

$$P^T AP = LBL^T, \quad (2.2)$$

where P is a permutation matrix, L is a unit lower triangular matrix, and B is a block diagonal matrix

$$B = \begin{bmatrix} B^{(1)} & 0 & \cdots & 0 \\ 0 & B^{(2)} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & B^{(K)} \end{bmatrix},$$

where each block $B^{(k)}$ is either a 1-by-1 or 2-by-2 matrix and is nonsingular. Matrix P represents the pivoting—the exchanging of rows and columns of A —during the factorization in order to maintain both sparsity and stability of the factor. After obtaining the factorization, back and forward substitutions are used to compute the solution of (2.1) by the following steps:

- (i) Solve $z : Lz = P^T b$.
- (ii) Solve $\hat{z} : B\hat{z} = z$.
- (iii) Solve $\bar{z} : L^T \bar{z} = \hat{z}$.
- (iv) Set $x : x = P\bar{z}$.

Recall that Steps (ii) and (iv) are trivial (due to the structure of P and B) and therefore the computational time for solving the linear system depends solely on the factorization and back and forward substitutions in Steps (i) and (iii), which in turn depend on the sparsity of L .

To perform symmetric indefinite factorization, let $A^{(k)}$ be the (smaller) matrix that remains to be factorized in the k th iteration. The algorithm starts with $A^{(1)} = A$. For each iteration, we first identify a submatrix $B^{(k)}$ from elements of $A^{(k)}$ that is suitable to be used as the pivot block. The submatrix $B^{(k)}$ is either a single diagonal element of $A^{(k)}$ ($[a_{ll}^{(k)}]$) or a 2-by-2 block with two diagonal elements of $A^{(k)}$ ($\begin{bmatrix} a_{ll}^{(k)} & a_{lr}^{(k)} \\ a_{rl}^{(k)} & a_{rr}^{(k)} \end{bmatrix}$). Note that there are many methods to select the pivot (We explain pivot selection algorithms and our proposed pivot selection algorithm in the next section). Next, we find the permutation matrix $P^{(k)}$ satisfying

$$(P^{(k)})^T A^{(k)} P^{(k)} = \begin{bmatrix} B^{(k)} & (C^{(k)})^T \\ C^{(k)} & Z^{(k)} \end{bmatrix}. \quad (2.3)$$

That is, $P^{(k)}$ is the permutation matrix corresponding to the exchanges of rows and columns that move the chosen pivot to the top-left corner. The right-hand side of (2.3) can be factorized as

$$(P^{(k)})^T A^{(k)} P^{(k)} = \begin{bmatrix} I & 0 \\ C^{(k)}(B^{(k)})^{-1} & I \end{bmatrix} \cdot \begin{bmatrix} B^{(k)} & 0 \\ 0 & Z^{(k)} - C^{(k)}(B^{(k)})^{-1}(C^{(k)})^T \end{bmatrix} \cdot \begin{bmatrix} I & (B^{(k)})^{-1}(C^{(k)})^T \\ 0 & I \end{bmatrix}. \quad (2.4)$$

Let $L^{(k)} = C^{(k)}(B^{(k)})^{-1}$ and $A^{(k+1)} = Z^{(k)} - C^{(k)}(B^{(k)})^{-1}(C^{(k)})^T$. Equation (2.4) can be rewritten as

$$(P^{(k)})^T A^{(k)} P^{(k)} = \begin{bmatrix} I & 0 \\ L^{(k)} & I \end{bmatrix} \cdot \begin{bmatrix} B^{(k)} & 0 \\ 0 & A^{(k+1)} \end{bmatrix} \cdot \begin{bmatrix} I & (L^{(k)})^T \\ 0 & I \end{bmatrix}. \quad (2.5)$$

The same process can be repeated recursively on the matrix $A^{(k+1)}$. Note that the dimension of $A^{(k+1)}$ is less than the dimension of $A^{(k)}$ by either one or two depending on the dimension of $B^{(k)}$.

3 Pivot selection with Minimum Degree

Finding the optimal ordering that minimizes fill-in is NP-hard [6] therefore a heuristic is often used for pivot selection. Choosing pivot at each step should be inexpensive, lead to at most modest growth in the elements of the remaining matrix, and not cause L to be too much denser than the original matrix. One of the well-known and efficient pivot selection techniques is the minimum degree algorithm [2, 3, 4]. The algorithm considers the pivot based on the following graph model. Define an undirected graph $G = (V, E)$, where $V = \{1, \dots, n\}$ and $E = \{\{i, j\} : i \neq j \text{ and } a_{ij} \neq 0\}$. Observe that the degree of v ($\deg(v)$), where $v \in V$, is the number of nonzero off-diagonal elements on the v th row. The vertex v with minimum $\deg(v)$ is chosen as the pivot.

Define the elimination graph $G_v = (V \setminus \{v\}, E')$, where $E' = E \cup \{\{i, j\} : \{i, v\} \in E \text{ and } \{v, j\} \in E\} \setminus \{\{v, i\} : i = 1, 2, \dots, n\}$. Graph G_v is used to choose the next pivot, and so on. That is, the minimum degree algorithm is as follows.

Algorithm 1. *Minimum Degree Algorithm*

Define G *as described above.*
while $G \neq \emptyset$ **do**
 $v =$ *the vertex with minimum* $\deg(v)$
 $G = G_v$
end while

Note that the minimum degree algorithm identifies the pivot at each step without any numerical calculation. For this reason, it can be used as the ordering step before factorizing the matrix.

Many improvements of the minimum degree algorithm and its implementation have been proposed [5] such as decreasing the computation time for the degree update by considering

the indistinguishable nodes [20] or minimum degree independent nodes [10], reducing the computation cost by using an approximate minimum degree [11], and saving space by using the quotient graph model [21]. Heggernes et al. [22] show that the time complexity of the minimum degree algorithm is $O(mn^2)$, the time complexity of the multiple minimum degree is also $O(mn^2)$, and the time complexity of the approximate minimum degree is $O(mn)$, where m is the number of edges in the elimination graph and n is the number of vertices in the graph.

4 Our pivot selection algorithm

Unlike in Cholesky factorization, pivots in symmetric indefinite factorization can be either a scalar or a 2-by-2 matrix therefore the minimum degree algorithm cannot be used as is in this case.

The stability condition that our algorithm uses is proposed by Duff et al. [23] and also used as a thresholding test for 1-by-1 and 2-by-2 pivots in MA57 [1]. We consider a 1-by-1 pivot a_{ii} to be *acceptably stable* if

$$|a_{ii}| \geq \alpha \max_{r \neq i} |a_{ri}|. \quad (4.1)$$

Similarly, a 2-by-2 pivot $\begin{bmatrix} a_{ii} & a_{ij} \\ a_{ji} & a_{jj} \end{bmatrix}$ is considered to be *acceptably stable* if

$$\left| \begin{bmatrix} a_{ii} & a_{ij} \\ a_{ji} & a_{jj} \end{bmatrix}^{-1} \right| \cdot \begin{bmatrix} \max_{r \neq i, r \neq j} |a_{ri}| \\ \max_{r \neq i, r \neq j} |a_{rj}| \end{bmatrix} \leq \begin{bmatrix} \alpha^{-1} \\ \alpha^{-1} \end{bmatrix}. \quad (4.2)$$

Conditions (4.1) and (4.2) limit the magnitudes of the entries of L to $1/\alpha$ at most. The appropriate value of α is $0 < \alpha \leq 0.5$. The default value of α in MA57 is 0.01 [1].

Let us call the column with the fewest number of off-diagonal nonzeros the *minimum degree column*. Let i be the minimum degree column of the matrix A . We accept a_{ii} as the 1-by-1 pivot ($B^{(k)}$) if a_{ii} satisfies (4.1). Otherwise, we proceed to search for a suitable 2-by-2 pivot $\begin{bmatrix} a_{ii} & a_{ij} \\ a_{ji} & a_{jj} \end{bmatrix}$ that satisfies (4.2) as follows. Let

$$Z_i = \{z | a_{iz} \neq 0 \text{ and } z \neq i\}. \quad (4.3)$$

Consider all submatrices $\begin{bmatrix} a_{ii} & a_{iz} \\ a_{zi} & a_{zz} \end{bmatrix}$, where $z \in Z_i$, as the candidates for a 2-by-2 pivot. The degree of each candidate $\deg(i, z)$ is the number of rows l where $l \neq i, z$ and at least one of a_{li} and a_{lz} is nonzero. To compute $\deg(i, z)$, define

$$d(i, z, l) = \begin{cases} 0, & \text{if } a_{li} = 0 \text{ and } a_{lz} = 0, \\ 1, & \text{otherwise.} \end{cases} \quad (4.4)$$

Hence,

$$\deg(i, z) = \sum_{l \neq i, z} d(i, z, l). \quad (4.5)$$

Our algorithm then considers all of the candidates with the minimum out-degree. Specifically, $\begin{bmatrix} a_{ii} & a_{ij} \\ a_{ji} & a_{jj} \end{bmatrix}$ is qualified if

$$\deg(i, j) = \min_{z \in Z_i} \deg(i, z). \quad (4.6)$$

If a qualified candidate also satisfies (4.2), it is chosen as a pivot. Otherwise, we remove j from the Z_i and repeat the process of selecting a 2-by-2 pivot until we either find a qualified candidate that also satisfies (4.2) or Z_i becomes empty. In the latter case, we set i to be the next next minimum degree column and repeat the process from the beginning (from testing whether a_{ii} is a suitable 1-by-1 pivot). The algorithm is as shown in Algorithm 2 below.

Algorithm 2. *Our Pivot Selection Algorithm*

```

// A is a n-by-n symmetric indefinite matrix
Let M = {1, 2, ..., n}
while a suitable pivot is not yet found and M is not empty do
  Let i be the minimum degree column among all column indices in M
  if  $a_{ii}$  is accepted then
    Use  $a_{ii}$  as the 1-by-1 pivot
  else
    Let  $Z_i = \{z | a_{iz} \neq 0 \text{ and } z \neq i\}$ 
    while a suitable pivot is not yet found and  $Z_i$  is not empty do
      Let j be such that  $\begin{bmatrix} a_{ii} & a_{ij} \\ a_{ji} & a_{jj} \end{bmatrix}$  has the minimum out-degree and  $j \in Z_i$ 
      if  $\begin{bmatrix} a_{ii} & a_{ij} \\ a_{ji} & a_{jj} \end{bmatrix}$  satisfies (4.2) then
        Use  $\begin{bmatrix} a_{ii} & a_{ij} \\ a_{ji} & a_{jj} \end{bmatrix}$  as the 2-by-2 pivot
      else
        Remove j from  $Z_i$ 
      end if
    end while
    Remove i from M
  end if
end while

```

Lastly, when the remaining matrix is fully dense, we continue with a conventional pivot selection algorithm such as BBK instead.

5 Experiments and results

This section compares the efficiency of our algorithm with MA57, which is based on the multifrontal method. The experiments are performed in Matlab 2011a on matrices of varying dimensions from 100 to 5000. For each dimension, we vary the percentage of nonzeros in the matrices from 5 to 30 percent. We test with 20 different instances for problems with

Table 1: Average percentage of nonzeros in the factor L produced by MA57 and our algorithm for problems with 100, 300, 500, 1000, 3000, and 5000 dimensions and 30, 20, 10, and 5 percent of nonzeros in the matrix. The percentage of nonzeros in L is computed by dividing the number of nonzeros in L by n^2 and then multiplying the result by 100.

n	Percentage of nonzeros in L							
	30		20		10		5	
	MA57	Our method	MA57	Our method	MA57	Our method	MA57	Our method
100	46.20	45.54	40.90	39.24	22.68	18.73	11.02	6.60
300	46.07	45.39	43.03	41.89	35.76	33.15	25.17	21.23
500	47.37	46.98	45.26	44.52	39.98	38.17	17.42	12.04
1000	48.53	48.35	47.39	47.00	44.01	43.02	38.46	36.36
3000	49.46	49.37	49.00	48.84	47.56	47.21	45.11	44.19
5000	49.64	49.61	49.36	49.26	48.47	48.22	46.86	46.23

Table 2: Average residuals of the factorization produced by MA57 and our algorithm for problems with 300, 500, 1000, and 2000 dimensions and 30, 20, 10 and 5 percent of nonzeros in the matrix.

n	Residual ($\times 10^{-10}$)							
	30		20		10		5	
	MA57	Our method	MA57	Our method	MA57	Our method	MA57	Our method
100	0.00339	0.00018	0.00410	0.00022	0.00190	0.00016	0.00045	0.00006
300	0.03072	0.00077	0.02634	0.00083	0.02378	0.00083	0.01039	0.00059
500	0.08489	0.00128	0.06665	0.00161	0.04665	0.00169	0.02199	0.00076
1000	0.20679	0.00342	0.21691	0.00374	0.17399	0.00333	0.10509	0.00355
3000	1.63656	0.01312	1.80491	0.01281	1.32961	0.02150	1.13003	0.02160
5000	4.45974	0.02488	3.49949	0.02361	2.51524	0.03264	2.20916	0.03152

100, 300, and 500 dimensions and 10 different instances for problems with 1000, 3000, and 5000 dimensions. We show the percentage of nonzeros in the factor L of the two methods in Table 1, which shows that our method produces sparser factors than MA57 in all cases. Note that the small percentage improvement for large matrices are not insignificant as small decrease in nonzeros does lead to significantly faster factorization time. Finally, Table 2 shows the residuals $\|P^T AP - LBL^T\|$ of the results of both methods. The result shows that our method produces more accurate factors than MA57.

6 Conclusion

In this article, we propose a new pivot selection algorithm for symmetric indefinite factorization. Our method is based on the minimum degree algorithm but is able to select both 1-by-1 and 2-by-2 pivots that are stable. Our experimental results show that our algorithm produces factors that are stable and also sparser than MA57.

Acknowledgement. This work is partially supported by Center of Excellence in Intelligent Informatics, Speech and Language Technology, and Service Innovation (CILS), Thammasat University and National Research Universities, Thailand.

References

- [1] I. S. DUFF, MA57 – A Code for the Solution of Sparse Symmetric Definite and Indefinite Systems, *ACM Transactions on Mathematical Software* **30(2)** (2004), 118–144
- [2] H. M. MARKOWITZ, The elimination form of the inverse and its application to linear programming, *Management Science* **3(3)** (1957), 255–269
- [3] W. F. TINNEY AND J. W. WALKER, Direct Solution of Sparse Network Equations by Optimally Ordered Triangular Factorization, *in: Proceedings of the IEEE* **55(11)** (1967), 1801–1809
- [4] D. J. ROSE, *A graph-theoretic study of the numerical solution of sparse positive definite systems of linear equations*, R.C. Read (Ed.), Graph Theory and Computing, Academic Press, New York, 1972.
- [5] A. GEORGE AND J. W. H. LIU, The Evolution of the Minimum Degree Ordering Algorithm, *SIAM Review* **31(1)** (1989), 1–19
- [6] M. YANNAKAKIS, Computing the Minimum Fill-In is NP-Complete, *SIAM Journal on Algebraic Discrete Methods* **2(1)** (1981), 77–79
- [7] E. CUTHILL AND J. MCKEE, Reducing the Bandwidth of Sparse Symmetric Matrices, *in: ACM '69 Proceedings of the 1969 24th national conference* (1969), 157–172
- [8] A. GEORGE, Nested Dissection of a Regular Finite Element Mesh, *SIAM Journal on Numerical Analysis* **10(2)** (1973), 345–363
- [9] R. J. LIPTON, D. J. ROSE, AND R. ENDRE TARJAN, Generalized Nested Dissection, *SIAM Journal on Numerical Analysis* **16(2)** (1979), 346–358
- [10] J. W. H. LIU, Modification of The Minimum-Degree Algorithm by Multiple Elimination, *ACM Transactions on Mathematical Software* **11(2)** (1985), 141–153
- [11] P. R. AMESTOY, T. A. DAVIS, AND I. S. DUFF, An Approximate Minimum Degree Ordering Algorithm, *SIAM Journal Matrix Analysis and Applications* **17(4)** (1996), 886–905
- [12] P. E. GILL, W. MURRAY, AND M.H. WRIGHT, *Numerical Linear Algebra and Optimization*, Vol.1, Addison-Wesley, 1991.
- [13] J. R. BUNCH AND B. N. PARLETT, Direct Methods for Solving Symmetric Indefinite Systems of Linear Equations, *SIAM Journal on Numerical Analysis* **8(4)** (1971), 639–655
- [14] J. R. BUNCH AND L. KAUFMAN, Some Stable Methods for Calculating Inertia and Solving Symmetric Linear Systems, *Mathematics of Computation* **31(137)** (1977), 163–179

- [15] C. ASHCRAFT, R. G. GRIMES, AND J. G. LEWIS, Accurate Symmetric Indefinite Linear Equation Solvers, *SIAM Journal on Matrix Analysis and Applications* **20(2)** (1995), 513–561
- [16] C. C. PAIGE AND M. A. SAUNDERS, Solution of Sparse Indefinite Systems of Linear Equations, *SIAM Journal on Numerical Analysis* **12(4)** (1974), 617–629
- [17] I. S. DUFF, J. K. REID, N. MUNKSGAARD, AND H. B. NIELSEN, Direct Solution of Sets of Linear Equations whose Matrix is Sparse, Symmetric and Indefinite, *IMA Journal of Applied Mathematics* **23(2)** (1979), 235–250
- [18] I. S. DUFF AND J. K. REID, The Multifrontal Solution of Indefinite Sparse Symmetric Linear Equations, *ACM Transactions on Mathematical Software* **9(3)** (1983), 302–325
- [19] N. I. M. GOULD, J. A. SCOTT, AND Y. HU, A Numerical Evaluation of Sparse Direct Solvers for the Solution of Large Sparse Symmetric Linear Systems of Equations, *ACM Transactions on Mathematical Software* **33(2)** (2007), 118–144
- [20] A. GEORGE AND D. R. MCINTYRE, On the Application of the Minimum Degree Algorithm to Finite Element Systems, *SIAM Journal on Numerical Analysis* **15(1)** (1978), 90–112
- [21] A. GEORGE AND J. W. H. LIU, A Fast Implementation of the Minimum Degree Algorithm Using Quotient Graphs, *ACM Transactions on Mathematical Software* **6(3)** (1980), 337–358
- [22] P. HEGGERNES, S. C. EISENSTAT, G. KUMFERT, AND A. POTHEN, The Computational Complexity of the Minimum Degree Algorithm, *in: Proceedings of NIK 2001 - 14th Norwegian Computer Science Conference* (2001), 98–109
- [23] I. S. DUFF, N. I. M. GOULD, J. K. REID, AND J. A. SCOTT, The Factorization of Sparse Symmetric Indefinite Matrices, *IMA Journal of Numerical Analysis* **11(2)** (1991), 181–204

Received: 29.01.2016

Revised: 12.04.2016

Accepted: 13.04.2016

School of Information, Computer
and Communication Technology,
Sirindhorn International
Institute of Technology,
Thammasat University, Thailand
E-mail: duangpenj@gmail.com,
gun@siit.tu.ac.th