



# **A note on fast approximate minimum degree orderings for symmetric matrices with some dense rows**

**H. S. Dollar   J. A. Scott**

March 2008 (revised June 2008)

**RAL-TR-2007-022**

© **Science and Technology Facilities Council**

Enquires about copyright, reproduction and requests for additional copies of this report should be addressed to:

Library and Information Services  
SFTC Rutherford Appleton Laboratory  
Harwell Science and Innovation Campus  
Didcot  
OX11 0QX  
UK  
Tel: +44 (0)1235 445384  
Fax: +44(0)1235 446403  
Email: [library@rl.ac.uk](mailto:library@rl.ac.uk)

The STFC ePublication archive (epubs), recording the scientific output of the Chilbolton, Daresbury, and Rutherford Appleton Laboratories is available online at:  
<http://epubs.cclrc.ac.uk/>

**ISSN 1358-6254**

Neither the Council nor the Laboratory accept any responsibility for loss or damage arising from the use of information contained in any of their reports or in any communication about their tests or investigation

## A note on fast approximate minimum degree orderings for symmetric matrices with some dense rows

H. S. Dollar<sup>1</sup>, J. A. Scott<sup>1</sup>

### ABSTRACT

Recently a number of variants of the approximate minimum degree algorithm have been proposed that aim to efficiently order symmetric matrices containing some dense rows. We compare the performance of these variants on a range of problems and highlight their limitations. This leads us to propose a new variant that offers both speed and robustness.

**Keywords:** approximate minimum degree ordering algorithm, sparse symmetric matrices, dense rows, graph algorithms, ordering algorithms.

**AMS(MOS) subject classifications:** 65F50, 65F05.

---

<sup>1</sup> Computational Science and Engineering Department, Rutherford Appleton Laboratory, Chilton, Oxfordshire, OX11 0QX, England, UK.  
Email: s.dollar@rl.ac.uk & j.a.scott@rl.ac.uk  
Current reports available from <http://www.numerical.rl.ac.uk/reports/reports.html>  
This work was supported by EPSRC grant EP/E053351/1.

# 1 Introduction

The efficiency of sparse direct solvers for the solution of symmetric linear systems  $Ax = b$ , in terms of both the storage needed and work performed, is dependent upon the order in which the variables are eliminated, that is, the order in which the pivots are selected. Many solvers include a reordering step that aims to use information on the sparsity pattern of  $A$  to find a permutation so that, if the pivots are chosen in order from the diagonal of the permuted matrix, the computed factors are sparser than if the pivots were chosen in order from the diagonal of the original matrix. If  $A$  is positive definite, the pivot sequence chosen from the sparsity pattern alone can be used by the factorization phase without modification and a Cholesky factorization  $PAP^T = LL^T$ , where  $P$  is a permutation matrix and  $L$  is lower triangular, can be computed. More generally, numerical pivoting must be incorporated during the factorization phase to maintain numerical stability and, in this case, the pivot sequence computed by the symbolic analysis phase may have to be modified.

The problem of finding a permutation  $P$  that results in the smallest amount of fill-in for a Cholesky factorization is NP-complete [1] and so heuristics are used to find a good ordering. Two main classes of methods are widely used: those based on nested dissection [2] and those based on the minimum degree algorithm [3]. In recent years, nested dissection has often been found to be the method of choice for many very large problems (typically those of order greater than 50,000) [4]. However, it can be more expensive than the most efficient implementations of the minimum degree algorithm, which is preferred for more modest size problems and for very sparse problems. Currently, the most successful variant of the minimum degree algorithm is the approximate minimum degree algorithm (AMD) and, in particular, the AMD algorithm introduced by Amestoy, Davis and Duff [5] is widely used. The AMD algorithm is more efficient since it uses computationally cheap bounds on the minimum degree in place of the exact minimum degree and, in practice, it produces orderings that are comparable in quality [5].

Although AMD is generally very successful, an important exception is when the matrix  $A$  has some dense (or almost dense) rows and columns. In this case, the run time for AMD can be high. AMD uses the undirected graph of the matrix and selects each node in turn to have minimum (approximate) degree. Once a node is selected, it is eliminated from the graph and replaced by adding edges between its neighbours so that the neighbours become a clique. If a row is full, the corresponding node will always be adjacent to the eliminated node so its adjacency list has to be scanned and updated, requiring  $\mathcal{O}(n^2)$  operations for a problem with  $n$  variables. This makes the algorithm prohibitively expensive.

We will compare three variants of the AMD algorithm that aim to efficiently form a pivot order when the matrix  $A$  has some dense rows and columns. The first two were proposed by Carmen [6] and Davis [7]: each performs a preprocessing stage during which the rows and columns that it considers to be dense are removed, and then the AMD algorithm is applied to the remaining matrix. The variant that was originally introduced by Amestoy for use within the parallel direct solver MUMPS [8] and was recently discussed by Amestoy *et al.* [9] uses a more sophisticated dynamic partitioning of the rows into dense and sparse rows to efficiently compute the pivot order. Our experiments have shown that all three variants can perform poorly (either in terms of the quality of the pivot order or the time taken to compute it) and this leads us to present a new variant that combines the speed of the methods of Carmen and Davis with the robustness of the Amestoy method.

This paper is organised as follows. We end this section by introducing our test environment then, in Section 2, we review the minimum degree and AMD algorithms. The methods of Carmen, Davis, and Amestoy *et al.* are described and compared in Section 3. The new variant of the AMD algorithm that we propose is introduced in Section 4. In Section 5, we draw our conclusions.

## 1.1 Test environment

In Table 1, we list our test set. Each example is available from the University of Florida Sparse Matrix Collection [10] and, if unsymmetric, has been symmetrized by working with the sum of the given matrix and its transpose. The order  $n$ , the number of off-diagonal nonzero entries  $nz$ , the maximum external

degree  $d_{max}$ , the average external degree  $\mu$ , and the standard deviation  $\sigma$  of the external degrees of each (symmetrized) matrix is given. The listed matrices include those that Amestoy *et al.* found to contain dense (or near dense) rows and have  $\sigma/\mu \geq 1$ . Amestoy *et al.* found that, for matrices with  $\sigma/\mu < 1$ , classical AMD performed well and so we restrict attention to those with  $\sigma/\mu \geq 1$ .

Throughout this paper, we compare the different algorithms by reporting the time to form a pivot order and the forecast number of reals required to hold the matrix factor. The latter piece of information is used as a measure for the quality of the ordering and is generated by passing the pivot order to the analyse phase of MA57. All numerical results were obtained using a 3.6 GHz Intel Xeon dual processor Dell Precision 670 with 4 Gbytes of RAM that was running Red Hat Enterprise Linux Server release 5.1 (kernel 2.6.18-53.1.13.el5). The Nag f95 compiler with the -O4 option was used. All times are CPU times in seconds.

## 2 The minimum degree and AMD algorithms

The minimum degree and AMD algorithms may be presented using the graph model of Rose [11, 12]. The nonzero pattern of a sparse symmetric matrix  $A$  of order  $n$  can be represented by an undirected graph  $G^0 = (V^0, E^0)$  with nodes  $V^0 = \{1, \dots, n\}$  and edges  $E^0$ . An edge  $(i, j)$  is present in  $E^0$  if and only if  $a_{ij} \neq 0$  and  $i \neq j$ . Nodes  $i$  and  $j$  are adjacent to each other (neighbours) in graph  $G^0$  if the edge  $(i, j)$  is present in  $E^0$ .

The elimination graph  $G^k = (V^k, E^k)$  describes the nonzero pattern of the *reduced matrix*  $A^{(k)}$  of order  $n^{(k)}$  still to be factored after  $k$  pivots have been chosen and eliminated. The external degree  $d_i^k$  of node  $i$  is defined to be the number of nodes adjacent to node  $i$  in  $G^k$ . The minimum degree algorithm chooses a node  $p$  that has minimal external degree in the graph  $G^{k-1}$  as the  $k$ th pivot. Algorithm 1 gives an outline of the minimum degree algorithm. The term *variable* is used to indicate a node that has not been removed from the elimination graph.

---



---

### Algorithm 1 The minimum degree algorithm

---



---

```

For each variable  $i \in V^0$ , set  $d_i^0$  to be the external degree of  $i$ 
 $k = 1$ 
while  $k \leq n$  do
  Select variable  $p \in V^{k-1}$  to minimize  $d_p^{k-1}$ 
  Eliminate  $p$ 
  For each variable  $i$  adjacent to  $p$  in  $G^{k-1}$ , update  $d_i^k$ 
   $k = k + 1$ 
end while

```

---



---

The graph  $G^k$  depends on  $G^{k-1}$  and the choice of the  $k$ th pivot.  $G^k$  is found by selecting the  $k$ th pivot from  $V^{k-1}$ , adding edges to  $E^{k-1}$  to make the nodes adjacent to  $p$  in  $G^{k-1}$  a *clique* (a fully connected subgraph) and then removing  $p$  (and its edges) from the graph. The edges added to the graph correspond to fill-in. This addition of edges means that, if  $G^k$  is stored explicitly, we cannot know the storage requirements in advance. To remedy this a quotient graph is used instead of an elimination graph [5]. However, the calculation of the external degree of a variable may then be expensive. Alternatively, the AMD algorithm calculates an upper bound,  $\bar{d}_i$ , on the external degree of a variable and then uses the upper bounds when choosing the next pivot. This generally results in the AMD algorithm being more efficient. Numerical results have shown that the AMD algorithm also produces orderings that are comparable in quality to the best classical minimum degree algorithms [5].

Unfortunately, the AMD algorithm can be very inefficient when the matrix  $A$  has some dense (or almost dense) rows and columns. When a pivot  $p$  is eliminated all of its adjacent neighbours must update

their upper bounds  $\bar{d}_i$ . If  $i$  is (almost) dense, it is (almost) certainly adjacent to  $p$  and updating  $\bar{d}_i$  will involve a large number of comparisons. Thus (almost) every step is expensive. For clarity, throughout the remainder of this paper, we will refer to this variant that takes no account of dense rows and columns as the *classical AMD algorithm*.

We initially compare the performance of the minimum degree algorithm with the classical AMD algorithm. The versions of these algorithms that are in the HSL [13] sparse direct solver MA57 [14] (version 3.2.0) are used. In columns 3 and 4 of Table 2, we give the time to obtain the ordering and the forecast number of reals required to hold the matrix factor. The other variants of AMD given in Table 2 are discussed in Sections 3 and 4. For `rajat29`, the minimum degree algorithm failed to compute a pivot order within 10000 seconds and, hence, no results are given.

As expected, we observe that the classical AMD algorithm is significantly more efficient than the minimum degree algorithm for our test problems while the quality of the orderings is comparable. Observe that the times to obtain the elimination orderings differ by at least an order of magnitude for the `gupta*` problems. However, as Amestoy *et al.* [9] show, for problems in our test set, the times to generate the elimination ordering with the classical AMD algorithm are very high when compared with problems of similar size and number of entries but no dense rows: the algorithms described in Section 3 aim to substantially reduce these times while maintaining the quality of the ordering.

### 3 Variants of AMD for efficiently detecting and treating dense rows

The problem of forming variants of the AMD algorithm that efficiently detect and treat dense rows has recently been considered by Carmen [6], Davis [7] and Amestoy *et al.* [9]. Using the names given by their authors, we refer to these variants as AMDpre, CS\_AMD and QAMD, respectively. The `amd` function of MATLAB<sup>®</sup> (version 7.5) uses the CS\_AMD algorithm.

#### 3.1 The AMDpre and CS\_AMD algorithms

The AMDpre [6] and CS\_AMD [7] algorithms both use a preprocessing step to search the matrix  $A$  for rows that they consider to be dense. The matrix  $A$  is reordered to take the form

$$A = \left[ \begin{array}{c|c} A_1 & A_r^T \\ \hline A_r & A_d \end{array} \right],$$

where  $A_r$  and  $A_d$  are considered dense. The classical AMD algorithm is then applied to  $A_1$  and the dense rows are appended to the end of the resulting pivot order: the authors assume that the dense rows all experience roughly the same amount of fill-in and order them in increasing value of their external degree within  $A$ .

The CS\_AMD algorithm classifies a row in  $A$  as dense if its external degree is greater than  $\alpha\sqrt{n}$ , where  $\alpha > 0$  is a fixed constant (the default is  $\alpha = 10$ ). We remark that Reid incorporated an option that uses a similar idea into the implementation of the minimum degree algorithm within the HSL [13] sparse solver MA27 [15].

The AMDpre algorithm uses a procedure that is equivalent to that given in Algorithm 2. The threshold  $\gamma$  takes the form  $\gamma = \beta\sqrt{n}$ , where  $\beta > 0$  (the default value is  $\beta = 1$ ). The major difference between this and the simpler CS\_AMD variant is that AMDpre updates the external degrees when a row is selected as dense and removed from  $A_1$ . This results in a more complicated implementation but a smaller threshold can be used to remove roughly the same number of dense rows.

---

**Algorithm 2** The AMDpre method for choosing  $A_1$ 

---

Set  $A_1 = A$   
Calculate the external degrees of the rows in  $A_1$   
**while** maximum external degree  $\geq \gamma$  **do**  
    Remove from  $A_1$  the row (and corresponding column) that has largest external degree  
    Update the external degrees of the rows in  $A_1$   
**end while**

---

### 3.2 The QAMD algorithm

The QAMD algorithm was developed independently of the AMDpre and CS\_AMD algorithms. It was originally used by Amestoy in the parallel direct solver MUMPS[8]. It uses a somewhat different approach since the partitioning of the matrix is dynamic, allowing the matrix to be partitioned more than once as the ordering proceeds. In [9], Amestoy *et al.* define a row to be full if all of its entries are (symbolically) nonzero; a row is quasi dense if it has a high proportion of nonzero entries (so that its external degree is large); a row is sparse if it is neither full nor quasi dense. Amestoy *et al.* partition the matrix  $A$  as

$$A = \left[ \begin{array}{c|c|c} A_s & A_{q1}^T & A_{f1}^T \\ \hline A_{q1} & A_q & A_{f2}^T \\ \hline A_{f1} & A_{f2} & A_f \end{array} \right],$$

where the rows in  $A_s$  are sparse, the rows in  $A_q$  are quasi dense, and the rows in  $A_f$  are full. QAMD applies the AMD algorithm to the submatrix  $A_s$  but, if a row becomes quasi dense or full, its variable is removed from  $A_s$  and placed into  $A_q$  or  $A_f$ , respectively. When all of the variables in  $A_s$  have been either eliminated or reclassified, the exact external degrees of the quasi dense rows are calculated and these rows are then reclassified to be either sparse or full and the algorithm restarts. The algorithm can restart a number of times until only full variables remain uneliminated. The corresponding variables are appended to the end of the pivot order (the order in which they are appended is arbitrary).

A threshold  $\tau > 0$  is used to select quasi dense rows. In [9], a variable  $i$  is reclassified as quasi dense if it is not known to be full and  $\bar{d}_i^s + n_q + n_f \geq \tau$ , where  $n_q$  and  $n_f$  are the numbers of quasi dense and dense rows, respectively, and  $\bar{d}_i^s$  is the approximate degree of variable  $i$  with respect to the matrix  $A_s$ . The threshold is updated at the beginning of each restart using a definition of the form  $\tau = \tau(\mu, \sigma)$ . A more detailed description of the QAMD algorithm and the choice of  $\tau$  is given in [9].

We remark that, although the QAMD algorithm can be implemented using the same amount of memory as the classical AMD algorithm, the implementation is complicated because the quotient graph needs to be reformed each time during each restart.

### 3.3 Comparison of AMDpre, CS\_AMD and QAMD

In Table 2, we compare the performance of the AMDpre, CS\_AMD and QAMD methods with that of the classical AMD algorithm for computing a pivot ordering. The default parameters for each method are used. We have highlighted the best times (and those within 10 percent of 0.01 seconds of the best) in bold and the best values (and those within 10 percent of the best) of  $nz(L)$  in italics for each problem. We observe that, as expected, there is generally a significant reduction in the time required to compute the ordering when the problem has some rows that are dense. AMDpre is generally faster than both CS\_AMD and QAMD. For some problems, CS\_AMD fails to detect many of the dense rows and this results in a slow run time. For example, for `gupta2`, CS\_AMD finds 258 dense rows while AMDpre finds 1193, leading to the CS\_AMD time being almost 40 times slower than AMDpre. Since both variants produce orderings with

similar values of  $nz(L)$ , we favour AMDpre and will only consider the AMDpre and QAMD algorithms in the remainder of this section.

For some problems, for example `gupta*`, AMDpre is at least an order of magnitude faster than QAMD. This is because AMDpre only needs to store the structure of the submatrix  $A_1$  but QAMD requires the whole structure to be stored and occasionally searched to allow the method to restart correctly. In terms of the quality of the ordering, for some problems (including `lp11` and `net4-1`), AMDpre outperforms QAMD but for others (including `mip1` and `a0nsdsi1`) the converse is true.

As well as computing a pivot order, QAMD forms the associated assembly tree: this may be of interest for the subsequent factorization. AMDpre and CS\_AMD only compute a pivot order and, hence, if the assembly tree is also required, this will be an additional overhead. Using the analyse phase of MA57 with `gupta3`, we found that it took an additional 0.27 seconds to form the assembly tree after using AMDpre. Therefore, for this problem, AMDpre is still significantly faster than QAMD. Occasionally, the overhead of forming an assembly tree is high relative to the time require to compute the pivot order with AMDpre, for example, the overhead for `mip1` is 0.80 seconds, but, in general, AMDpre remains competitive compared to QAMD. We note that some direct solvers, e.g. the recent out-of-core multifrontal package HSL\_MA77 [16], only require the pivot to be input and they implement their own algorithms for forming the assembly tree.

A potential problem with AMDpre is that it can remove a large number of sparse rows and this will result in the quality of the ordering being lost. To illustrate this, we used the HSL package YM11 to generate a banded random matrix with  $n = 60000$ , bandwidth equal to  $2\sqrt{n}$ ,  $nz = 20436000$ ,  $d_{max} = 405$  and  $\mu = 340$ . The classical AMD and QAMD methods both calculate an ordering in 1.78 seconds with  $nz(L) = 5.36 \times 10^8$ . By comparison, AMDpre removes a total of 18661 rows/columns and takes 1.26 seconds. However, the quality of the pivot order has been lost as  $nz(L) = 4.83 \times 10^9$ .

Conversely, we can produce test problems where the AMDpre method will struggle to detect the dense rows. Suppose  $A$  is of the form

$$A = \begin{bmatrix} S & U^T \\ U & T \end{bmatrix}, \quad (1)$$

where  $S$  is a tridiagonal matrix, only the top right hand entry of  $U$  is nonzero, and  $T$  is a matrix from our test set that contains some dense rows. No matter how large the order of  $S$  becomes, it will be the same rows of  $T$  that cause the slow down in the efficiency of AMD and, hence, roughly the same number of rows should be removed for different dimensions of  $S$ . If we apply AMDpre to  $A$ , then as the order of  $S$  increases, the threshold  $\gamma = \beta\sqrt{n}$  also increases. Therefore, rows that were classified as dense for smaller orders may no longer be classified as dense. Suppose we set  $T$  to be the matrix `gupta1` and let the order of  $A$  be  $mn$  ( $m = 1, 2, \dots$ ). In Figure 1, we compare the time to calculate the elimination ordering of the matrix with different values of  $m$ . As  $m$  increases, the time increases at a faster rate for AMDpre than QAMD. In Figure 2, we plot the number of rows that are detected as dense and, hence, removed during the preprocessing stage of AMDpre. As expected, as  $m$  increases, the number of rows detected as dense drops substantially. Both variants produce orderings of comparable quality. Note that CS\_AMD would suffer in a similar manner as the problem size increases. The behaviour seen in Figure 2 was observed when  $T$  was chosen to be any of the problems in our test set but for some problems, the time difference was less marked than we see in Figure 1.

More generally, Algorithm 2 will fail to correctly detect all of the dense rows in a matrix if the maximum degree is smaller than  $\beta\sqrt{n}$  but still significantly larger than the average external degree of the resulting matrix  $A_1$ . This suggests that we require a method that combines the power of the QAMD method for detecting dense rows with the efficiency of the AMDpre method.

## 4 The AMDD algorithm

The results in the previous section have highlighted the weaknesses of the existing attempts to modify the AMD algorithm to improve performance in the case of some dense rows. In this section, we introduce a






Figure 1: Time to calculate the ordering for the matrix  $A$  defined by (1), where  $T$  is set to the matrix `gupta1` and the dimension of  $A$  is  $nm$  ( $m = 1, 2, \dots$ ) using the AMDpre, QAMD and AMDD methods.




Figure 2: Number of dense rows detected and removed by the AMDpre and AMDD methods for the matrix  $A$  defined by (1), where  $T$  is set to the matrix `gupta1` and the dimension of  $A$  is  $mn$  ( $m = 1, 2, \dots$ ).

new variant that is designed to overcome these weaknesses.

Let us assume that  $k$  rows/columns have been classified as dense and that  $A(k)$  is the matrix that remains after these  $k$  rows and columns have been removed from  $A$ . Let  $\mu(k)$  be the mean of the external degrees of the variables in  $A(k)$  and  $i$  be a dense row in  $A(k)$ . If row and column  $i$  are removed from the matrix  $A(k)$ , we expect the mean of value of the degrees of the variables to be reduced significantly more than if  $i$  were a sparse row. This suggests that, instead of comparing the degree of a variable against the threshold  $\gamma = \beta\sqrt{n}$  as used in the preprocessing stage of the AMDpre method, we should compare  $\mu(k)$  with the value that  $\mu(k+1)$  would take if row/column  $i$  are removed. We declare the row/column as dense if the difference is greater than a threshold  $t(n-k)$ . Consider again the matrix  $A$  given by (1). As the order of  $A$  is increased, the difference in the value of the mean external degree before and after a row (and corresponding column) is removed decreases. Numerical experimentation on the problems in our test set and the additional problems considered in Section 3.3 has shown that defining  $t(n-k)$  as

$$t(n-k) = \frac{\delta \ln(n-k)}{n-k} + \frac{\mu(k)}{n-k-1},$$

with  $\delta = 40$ , produces a variant of AMD that is both efficient and maintains the quality of the ordering. The method was not found to be very sensitive to the choice of  $\delta$ : values between 30 and 50 produced, in general, similar results.

Note that removing row and column  $i$  from  $A(k)$  results in a matrix with mean external degree

$$\mu(k+1) = \frac{\mu(k)(n-k) - 2d_i}{n-k-1}.$$

Therefore, our criteria for declaring a row as dense is

$$\mu(k) - \frac{\mu(k)(n-k) - 2d_i}{n-k-1} \geq \frac{\delta \ln(n-k)}{n-k} + \frac{\mu(k)}{n-k-1}.$$

Rearranging, we obtain

$$d_i - \mu(k) \geq \frac{\delta}{2} \frac{n-k-1}{n-k} \ln(n-k) \approx \frac{\delta}{2} \ln(n-k).$$

Comparing our definition of a dense row with that of AMDpre, we observe that our new definition additionally takes the mean external degree into account, it contains a function that rises less steeply with the order of the matrix, and it also makes allowances for the reduction in the order as rows are removed. We believe that all of these properties are sensible. Our proposed new variant, which we call the AMDD method, is outlined in Algorithm 3. Given an efficient implementation of classical AMD, the AMDD method is straightforward to implement and required no additional memory.

In Figure 1, we observe that the time for AMDD to calculate a pivot order for problem (1) increases with  $m$  at roughly the same rate as for the QAMD method. The key point is that, as  $m$  increases, there is only a small drop in the number of rows detected as dense by AMDD (Figure 2).

In the final column of Table 2, we report the time to compute the elimination ordering and the estimated number of reals in  $L$  generated by orderings from the AMDD algorithm. We observe that, in general, AMDD is significantly faster than QAMD and the quality of the ordering is often better. Note that, if we apply AMDD to the random banded matrix that we reported on in Section 3.3, then no rows/columns are classed as dense and the results are comparable to those of the classical AMD and QAMD methods.

## 5 Conclusions

We have compared a number of variants of AMD that aim to efficiently compute elimination orderings for matrices containing some dense rows and shown that both CS\_AMD and QAMD can be slow compared to AMDpre. Although AMDpre performs well on the majority of our test problems, we have shown that it

---

**Algorithm 3** The AMDD method

---

Set  $A(0) = A$ ,  $k = 0$ , and  $q$  to be an empty list  
For each row  $i \in A(0)$ , set  $d_i$  to be the external degree of  $i$   
Calculate the mean  $\mu(0)$  of the external degrees of the rows in  $A(0)$   
Select row  $i$  from  $A(0)$  that has largest external degree  
**while**  $\mu(k) - \frac{\mu(k)(n-k) - 2d_i}{n-k-1} \geq t(n-k)$  **do**  
    Add  $i$  to the beginning of the list  $q$   
     $A(k+1) = A(k)$  with row and column  $i$  removed  
     $\mu(k+1) = \frac{\mu(k)(n-k) - 2d_i}{n-k-1}$   
     $k = k + 1$   
    For each row  $i \in A(k)$ , update  $d_i$   
    Select row  $i$  from  $A(k)$  that has largest external degree  
**end while**  
Apply classical AMD to  $A(k)$  to form a pivot order  
Append the list  $q$  to the end of this pivot order to form a pivot order for  $A$

---

is possible to construct problems where it fails to correctly detect the dense rows. This led us to propose a new variant, called AMDD, that combines the robustness of QAMD in detecting dense rows with the speed of AMDpre. Our implementation of AMDD requires no extra storage over that required in the implementation of the classical AMD algorithm used within this paper and is much more straightforward to implement than QAMD because it does not need to restart or sub-partition the dense rows. The AMDD variant will be made available within HSL.

## Acknowledgments

We are grateful to Patrick Amestoy for helpful discussions on QAMD and to John Reid for several discussions on implementing minimum degree algorithms and for his careful reading of this paper.

## References

- [1] M. Yannakakis. Computing the minimum fill-in is NP-complete. *SIAM J. Alg. Discrete Meth.*, 2:77–79, 1981.
- [2] A. George. Nested dissection of a regular finite-element mesh. *SIAM J. Num. Anal.*, 10:345–363, 1973.
- [3] W.F. Tinney and J.W. Walker. Direct solutions of sparse network equations by optimally ordered triangular factorization. *Proc. IEEE*, 55:1801–1809, 1967.
- [4] I. S. Duff and J. A. Scott. Towards an automatic ordering for a symmetric sparse direct solver. Technical Report RAL-TR-2006-001, Rutherford Appleton Laboratory, 2006.
- [5] P. R. Amestoy, T. A. Davis, and I. S. Duff. An approximate minimum degree ordering algorithm. *SIAM J. Matrix Anal. Appl.*, 17(4):886–905, 1996.
- [6] J. L. Carmen. An AMD preprocessor for matrices with some dense rows and columns. <http://www.netlib.org/linalg/amd/amdpre.ps>.
- [7] T. A. Davis. *Direct methods for sparse linear systems*, volume 2 of *Fundamentals of Algorithms*. SIAM, Philadelphia, PA, 2006.

- [8] MUMPS. Multifrontal Massively Parallel Solver. See <http://mumps.enseiht.fr> or <http://graal.ens-lyon.fr/MUMPS/>.
- [9] P. R. Amestoy, H. S. Dollar, J. K. Reid, and J. A. Scott. An approximate minimum degree algorithm for matrices with dense rows. Technical Report RAL-TR-2007-020, Rutherford Appleton Laboratory, 2007.
- [10] T. Davis. The University of Florida Sparse Matrix Collection, 2007. <http://www.cise.ufl.edu/davis/techreports/matrices.pdf>.
- [11] D. J. Rose. *Symmetric elimination on sparse positive definite systems and the potential flow network problem*. PhD thesis, Harvard University, 1970.
- [12] D. J. Rose. A graph-theoretic study of the numerical solution of sparse positive definite systems of linear equations. In R. C. Read, editor, *Graph Theory and Computing*, pages 183–217. Academic Press, New York, 1973.
- [13] HSL. A collection of Fortran codes for large-scale scientific computation. See <http://hsl.rl.ac.uk/>, 2007.
- [14] I. S. Duff. MA57– a new code for the solution of sparse symmetric definite and indefinite systems. *ACM Transactions on Mathematical Software*, 30:118–154, 2004.
- [15] I. S. Duff and J. K. Reid. The multifrontal solution of indefinite sparse symmetric linear equations. *ACM Transactions on Mathematical Software*, 9(3):302–325, 1983.
- [16] J. K. Reid and J. A. Scott. An out-of-core sparse cholesky solver. Technical Report RAL-TR-2006-013, Rutherford Appleton Laboratory, 2006.

Problem	$n$	$nz$	$d_{max}$	$\mu$	$\sigma$	$\sigma/\mu$
lp11	32460	328036	253	10.1	15.5	1.53
gupta3	16783	9323427	14672	556	1234	2.22
mip1	66463	10352819	66395	156	351	2.25
net-4	88343	2441727	4791	27.6	85.3	3.09
ckt11752_dc_1	49702	327834	2921	6.60	24.6	3.73
rajat23	110355	559733	3336	5.07	19.3	3.80
rajat22	39899	198081	3336	4.96	24.8	4.99
gupta2	62064	4248286	8413	68.5	356	5.20
gupta1	31802	2164210	8413	68.1	360	5.29
a0nsdsil	80016	355034	5003	4.44	50.0	11.3
blockqp1	60012	640033	40011	10.7	305	28.6
trans5	116835	827879	114191	7.09	371	52.3
ins2	309412	2751484	309412	8.89	590	66.4
rajat29	643994	4026512	454043	6.25	783	125

Table 1: The test set: order  $n$ , number of nonzeros  $nz$ , the maximum external degree  $d_{max}$ , the mean  $\mu$  of the external degrees, and their standard deviation  $\sigma$ .

Problem		MinDeg	Classical	AMDpre	CS_AMD	QAMD	AMDD
lp11	Time	1.22	0.28	0.27	0.28	<b>0.19</b>	0.27
	$nz(L)$	<i>0.97</i>	<i>0.97</i>	<i>0.99</i>	<i>0.97</i>	1.20	<i>0.99</i>
gupta3	Time	115	8.74	<b>0.14</b>	0.55	4.30	0.17
	$nz(L)$	5.72	5.72	5.60	<i>5.35</i>	6.52	<i>5.53</i>
mip1	Time	40.8	10.2	<b>0.40</b>	0.96	0.81	0.46
	$nz(L)$	<i>38.9</i>	<i>39.0</i>	44.8	<i>39.3</i>	<i>39.7</i>	44.1
net4-1	Time	9.28	1.79	<b>0.49</b>	0.58	0.78	<b>0.53</b>
	$nz(L)$	<i>2.46</i>	<i>2.38</i>	<i>2.39</i>	<i>2.37</i>	3.09	<i>2.39</i>
ckt11752_dc_1	Time	0.49	0.19	<b>0.11</b>	0.15	0.15	<b>0.11</b>
	$nz(L)$	<i>0.59</i>	<i>0.56</i>	<i>0.57</i>	<i>0.56</i>	0.63	<i>0.57</i>
rajat23	Time	0.65	0.33	<b>0.19</b>	0.29	0.23	<b>0.19</b>
	$nz(L)$	<i>0.45</i>	<i>0.46</i>	<i>0.47</i>	<i>0.46</i>	<i>0.47</i>	<i>0.47</i>
rajat22	Time	0.29	0.13	<b>0.05</b>	<b>0.06</b>	<b>0.06</b>	<b>0.05</b>
	$nz(L)$	<i>0.15</i>	<i>0.15</i>	<i>0.15</i>	<i>0.15</i>	<i>0.16</i>	<i>0.15</i>
gupta2	Time	2020	88.9	<b>0.37</b>	13.4	5.74	<b>0.34</b>
	$nz(L)$	<i>5.96</i>	<i>5.89</i>	<i>6.30</i>	<i>5.77</i>	<i>5.98</i>	6.41
gupta1	Time	405	28.9	<b>0.09</b>	2.98	2.84	<b>0.09</b>
	$nz(L)$	<i>2.02</i>	<i>2.06</i>	<i>2.08</i>	<i>2.00</i>	<i>2.06</i>	<i>2.07</i>
a0nsdsil	Time	14.6	2.70	<b>0.07</b>	<b>0.07</b>	0.09	<b>0.08</b>
	$nz(L)$	<i>0.34</i>	<i>0.34</i>	0.39	0.39	<i>0.35</i>	0.39
blockqp1	Time	47.4	8.11	<b>0.03</b>	<b>0.03</b>	<b>0.04</b>	<b>0.03</b>
	$nz(L)$	<i>0.38</i>	<i>0.38</i>	<i>0.38</i>	<i>0.38</i>	<i>0.38</i>	<i>0.38</i>
trans5	Time	246	123	<b>0.15</b>	<b>0.16</b>	0.22	<b>0.16</b>
	$nz(L)$	<i>0.68</i>	<i>0.68</i>	<i>0.69</i>	<i>0.69</i>	<i>0.70</i>	<i>0.68</i>
ins2	Time	2080	825	<b>0.31</b>	<b>0.30</b>	0.53	<b>0.31</b>
	$nz(L)$	<i>1.53</i>	<i>1.53</i>	<i>1.59</i>	<i>1.59</i>	<i>1.59</i>	<i>1.59</i>
rajat29	Time	$> 10^4$	4040	<b>2.34</b>	2.62	3.32	<b>2.28</b>
	$nz(L)$	—	<i>9.77</i>	<i>9.97</i>	<i>9.87</i>	11.7	<i>9.93</i>

Table 2: Comparison of the times and the predicted number of reals in  $L$  ( $\times 10^6$ ) generated by orderings from the minimum degree (MinDeg), classical AMD (Classical), AMDpre, CS\_AMD QAMD and AMDD algorithms.