

Reversing the row order for the row-by-row frontal method

by

J. K. Reid and J. A. Scott

Abstract

The efficiency of the row-by-row frontal method is dependent on the row ordering used. Numerical experience has shown us that it can be advantageous to reverse a given row ordering. We present two results on invariances under the reversal of the ordering and use real applications to illustrate the variations that can take place.

Keywords: ordering rows, frontal method.

Computational Science and Engineering Department,
Atlas Centre, Rutherford Appleton Laboratory,
Oxon OX11 0QX, England.

May 28, 1999.

1 Introduction

The row-by-row frontal method (see, for example, Duff, Erisman and Reid, 1986, Section 10.6) solves a large sparse unsymmetric set of n linear equations by Gaussian elimination with the help of a full rectangular matrix held in memory and known as the *frontal matrix*. The size of the frontal matrix varies during the elimination. Rows are *assembled* (added) into the frontal matrix one by one. Whenever a column becomes *fully summed*, that is, the last row in which it has an entry is assembled, a pivot is chosen in the column and is used to eliminate the column and the row containing the pivot. The eliminated row and column are stored for use in the back-substitution or in the solution of further systems of equations.

Since an elimination can only take place after a column becomes fully summed, the order in which the rows are assembled will determine both how long a variable remains in the front and the order in which the variables are eliminated (apart from the order among columns that become fully summed at the same assembly step, which has no effect on the computational requirements). For efficiency, in terms of both storage and arithmetic operations, the rows need to be assembled in an order that keeps the size of the frontal matrix as small as possible. Scott (1998) has considered a number of strategies for determining such an ordering and has found that the results of using the different orderings on practical problems can vary enormously. While developing row ordering software for the Harwell Subroutine Library (1995), Scott tried experimenting with reversing any given row order. She found that this makes no difference to some of the usual measures of the quality of the ordering while it can make enormous changes to others.

These numerical results motivated us to prove the results on invariances under the reversal of the ordering that we present in Section 2. In Section 3, we illustrate the variations that can take place in real applications.

2 Two theorems on invariance

We follow Scott (1998) and use the terms *row frontsize* and *column frontsize* for the number of rows and columns in the front. Our invariance results pertain to the column frontsizes.

The row and column frontsizes *after assembly* refer to the frontsizes after a row has been added to the front. When a row is added, more than one column may become fully summed; for each such column, a pivot is chosen and an elimination is performed. The row and column frontsizes *before elimination* refer to the frontsizes immediately prior to an elimination. The *frontal matrix size* refers to the product of the row and column frontsizes.

We use the term *forward* row order to refer to the row ordering $1, 2, \dots, n$ (that is, to the given ordering), while the *reverse* ordering is $n, n - 1, \dots, 1$. We assume that the matrix (in the forward order) is $A = \{a_{ij}\}$ and that a reference to ‘row i ’ refers to the row $a_{i\bullet}$.

Before presenting our results, we introduce the small illustrative example

$$\begin{pmatrix} x & & & & & \\ x & x & & x & & \\ x & x & x & x & & \\ & x & & x & & \\ x & & & & & x \end{pmatrix}.$$

With the forward order we have:

Assembly 1: Frontal matrix increases to 1×1 . No eliminations.

Assembly 2: Frontal matrix increases to 2×3 . No eliminations.

Assembly 3: Frontal matrix increases to 3×4 . Columns 1 and 3 eliminated.
Frontal matrix reduces to 1×2 .

Assembly 4: Frontal matrix increases to 2×2 . Column 4 eliminated.
Frontal matrix reduces to 1×1 .

Assembly 5: Frontal matrix increases to 2×2 . Columns 2 and 5 eliminated.

Thus the row frontsize, column frontsize, and the frontal matrix size before each of the five eliminations and their mean values are

row frontsize	3	2	2	2	1	2.0
column frontsize	4	3	2	2	1	2.4
frontal matrix size	12	6	4	4	1	5.4

When the order is reversed, the corresponding statistics are

row frontsize	1	2	2	1	1	1.4
column frontsize	2	4	3	2	1	2.4
frontal matrix size	2	8	6	2	1	3.8

The first result is about the column frontsize after each assembly. This is important since the size of the frontal matrix after assembly determines how much memory is needed by the frontal solver.

Theorem 1 *If a given row order is reversed, the sequence of n column frontsizes after the assemblies is reversed.*

Proof After assembly of row i , in the forward or reverse row order, the column frontsize is the number of columns with first entry at row i or earlier and last entry at row i or later, that is, the number of columns j for which $\min\{k : a_{kj} \neq 0\} \leq i$ and $\max\{k : a_{kj} \neq 0\} \geq i$.

Corollary 1 *The maximum column frontsize is invariant.*

Corollary 2 *The mean column frontsize after assembly is invariant.*

Corollary 3 *The sum of the column frontsizes after assembly is invariant.*

Note that the sum of the column frontsizes after assembly is equal to the sum of the lifetimes, where the lifetime of a variable is defined to be the number of assembly steps for which the variable is in the front. The lifetime of variable j is also the length of column j , that is, $\max\{k : a_{kj} \neq 0\} - \min\{k : a_{kj} \neq 0\}$. The sum of the lifetimes is used by Camarda (1997) to compare the quality of different row orderings.

Our second result concerns the column frontsize before each elimination. This is important since the row and column frontsizes before an elimination determine how much work is associated with the elimination and how much storage is required for the pivotal column and row.

Theorem 2 *If a given row order is reversed, the sequence of n column frontsizes before the eliminations is permuted.*

Proof At assembly step i in the forward order, the column frontsize increases by one for each column with its first entry in row i , that is, each j such that $\min\{k : a_{kj} \neq 0\} = i$. Following this, it decreases by one for each column with its last entry in row i , that is, each j such that $\max\{k : a_{kj} \neq 0\} = i$, as that column is eliminated. There are n occasions when it increases by one and n occasions when it decreases by one. These $2n$ events begin and end with the column frontsize of zero. For each increase from $f - 1$ to f there is a corresponding later decrease from f to $f - 1$, and to establish a one-one correspondence, we take the first such decrease if there is more than one. Each decrease corresponds to an elimination with column frontsize f using the forward row order. Each increase corresponds to an elimination with column frontsize f using the reverse row order. The result follows.

For our small example, the sequence of column frontsizes for the forward order is 0, 1, 2, 3, 4, 3, 2, 1, 2, 1, 0. Note that there are two increases from 1 to 2 and two decreases from 2 to 1; the first increase is taken to correspond with the first decrease and the second increase to correspond with the second decrease.

Corollary 4 *The mean column frontsize before elimination is invariant.*

Note that the mean values in Corollaries 2 and 4 usually differ, while there is only one maximum value.

3 Numerical experience

We now present some numerical results that illustrate our theoretical results. All the frontsizes quoted in this section are frontsizes *before elimination*. The test problems are those used by Scott (1998). Each problem arises from a real engineering or industrial application. Our results are presented in Tables 3.1

Table 3.1: The maximum (max.) and mean row and column frontsizes before elimination for the forward and reverse row ordering.

Identifier	Forward				Reverse			
	max. row	max. col.	mean row	mean col.	max. row	max. col.	mean row	mean col.
bayer04	349	501	174	275	263	501	105	275
bayer09	83	153	38	62	78	153	26	62
bp1600	242	324	91	142	143	324	32	142
extr1	32	46	15	27	21	46	13	27
gre1107	95	208	55	125	123	208	72	125
hydr1	45	86	22	43	42	86	22	43
lhr07c	169	225	51	111	127	225	66	111
lhr14c	164	366	69	188	255	366	125	188
meg1	700	1150	368	639	517	1150	235	639
onetone2	297	658	198	405	365	658	209	405
orani678	1333	1576	604	771	320	1576	70	771
rdist1	40	81	28	60	49	81	32	60
west2021	38	52	18	29	22	52	12	29

Table 3.2: The maximum and mean frontal matrix size and the sum of lifetimes ($\cdot 10^3$) for the forward and reverse row ordering.

Identifier	Forward			Reverse		
	Max. frontal matrix size	Mean frontal matrix size	Sum of lifetimes	Max. frontal matrix size	Mean frontal matrix size	Sum of lifetimes
bayer04	17365	5341	5708	12387	3340	5708
bayer09	1237	283	194	1193	197	194
bp1600	6607	1833	100	4633	659	100
extr1	147	43	76	84	37	76
gre1107	1925	814	140	2489	1055	140
hydr1	387	104	226	361	104	226
lhr07c	3489	622	858	2331	791	858
lhr14c	5953	1530	2755	8992	2849	2755
meg1	77970	28129	1750	54395	18375	1750
onetone2	19483	8646	14637	23980	9264	14637
orani678	210080	64517	1704	48247	8679	1704
rdist1	275	169	242	397	202	242
west2021	198	57	58	86	36	58

and 3.2. In each case, the problem is reordered using our new code MC62 (Scott, 1999), with default values for all control variables.

From the tables, we see that for some problems, including `bp1600` and `orani678`, it can be extremely advantageous to reverse the row order. For other problems, such as `hydr1`, reversal has little effect. Note also that reversing the order can reduce the maximum row frontsize while increasing the mean row frontsize or the mean frontal matrix size. This is illustrated by problem `lhr07c`. In this case, which is the better ordering depends upon whether the prime concern is to reduce main memory requirements (choose the smaller maximum frontal matrix size), to minimise factor storage (choose the smaller mean row frontsize), or to minimise the operation count (choose the smaller mean frontal matrix size). Since the cost of computing the frontsizes for the forward and reverse orders is negligible compared with the cost of using a frontal solver, MC62 computes a new row ordering and then automatically reverses it and selects the better of the two orderings. By default, MC62 chooses the ordering for which the mean frontal matrix size is smaller.

References

- K.V. Camarda. *Ordering strategies for sparse matrices in chemical process simulation*. PhD thesis, University of Illinois at Urbana-Champaign, 1997.
- I.S. Duff, A.M. Erisman, and J.K. Reid. *Direct Methods for Sparse Matrices*. Oxford University Press, 1986.
- Harwell Subroutine Library. *A Catalogue of Subroutines (Release 12)*. Advanced Computing Department, AEA Technology, Harwell Laboratory, Oxfordshire, England, 1995.
- J.A. Scott. A new row ordering strategy for frontal solvers. Technical Report RAL-TR-1998-056, Rutherford Appleton Laboratory, 1998. To appear in *Numerical Linear Algebra and Applications*.
- J.A. Scott. Row ordering for frontal solvers in chemical process engineering. Technical Report RAL-TR-1999-035, Rutherford Appleton Laboratory, 1999.