



A Combined Unifrontal/Multifrontal Method for Unsymmetric Sparse Matrices

TIMOTHY A. DAVIS

University of Florida

and

IAIN S. DUFF

Rutherford Appleton Laboratory

We discuss the organization of frontal matrices in multifrontal methods for the solution of large sparse sets of unsymmetric linear equations. In the multifrontal method, work on a frontal matrix can be suspended, the frontal matrix can be stored for later reuse, and a new frontal matrix can be generated. There are thus several frontal matrices stored during the factorization, and one or more of these are assembled (summed) when creating a new frontal matrix. Although this means that arbitrary sparsity patterns can be handled efficiently, extra work is required to sum the frontal matrices together and can be costly because indirect addressing is required. The (uni)frontal method avoids this extra work by factorizing the matrix with a single frontal matrix. Rows and columns are added to the frontal matrix, and pivot rows and columns are removed. Data movement is simpler, but higher fill-in can result if the matrix cannot be permuted into a variable-band form with small profile. We consider a combined unifrontal/multifrontal algorithm to enable general fill-in reduction orderings to be applied without the data movement of previous multifrontal approaches. We discuss this technique in the context of a code designed for the solution of sparse systems with unsymmetric pattern.

Categories and Subject Descriptors: G.1.3 [Numerical Analysis]: Numerical Linear Algebra—*linear systems* (direct and iterative methods); *sparse, structured and very large systems* (direct and iterative methods); G.4 [Mathematics of Computing]: Mathematical Software—*algorithm design and analysis*

General Terms: Algorithms, Experimentation, Performance

Additional Key Words and Phrases: Frontal methods, linear equations, multifrontal methods, sparse unsymmetric matrices

This work was supported by National Science Foundation grant DMS-9504974.

Iain S. Duff is also associated with European Center for Research and Advanced Training in Scientific Computation (CERFACS), Toulouse, France.

Authors' addresses: T. A. Davis, Computer and Information Science and Engineering Department, University of Florida, Gainesville, FL; email: davis@cise.ufl.edu.; <http://www.cise.ufl.edu/~davis/>; I. S. Duff, Rutherford Appleton Laboratory, Chilton, Didcot, OXON OX11 0QX, England; email: isd@rl.ac.uk.; <http://www.dci.clrc.ac.uk/Group.asp?DCICSENAG>.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 1999 ACM 0098-3500/99/0300-0001 \$5.00

1. INTRODUCTION

We consider the direct solution of sets of linear equations $\mathbf{Ax} = \mathbf{b}$, where the coefficient matrix \mathbf{A} is sparse, unsymmetric, and has a general nonzero pattern. A permutation of the matrix \mathbf{A} is factorized into its LU factors, $\mathbf{PAQ} = \mathbf{LU}$, where \mathbf{P} and \mathbf{Q} are permutation matrices chosen to preserve sparsity and maintain numerical accuracy. Many recent algorithms and software packages for the direct solution of sparse systems are based on a multifrontal approach [Amestoy and Duff 1989; Davis and Duff 1997; Duff and Reid 1983; Liu 1992]. In this article, we will examine a new frontal matrix strategy to be used within a multifrontal approach. We use the term “unifrontal” for what is usually called the “frontal” method so that the term “frontal” can be used generically for both unifrontal and multifrontal methods.

Unifrontal and multifrontal methods compute the LU factors of \mathbf{A} by using data structures that permit regular access of memory and the use of dense matrix kernels (such as the BLAS) in their innermost loops. On supercomputers and high-performance workstations, this can lead to a significant increase in performance over methods that have irregular memory access and that do not use dense matrix kernels.

We discuss unifrontal methods in Section 2. We summarize the multifrontal method in Section 3, and in particular our earlier work on an unsymmetric-pattern multifrontal method. We refer to this prior method as UMFPACK V1.1 [Davis 1995; Davis and Duff 1991; 1997]. The combination of unifrontal and multifrontal methods is discussed in Section 4. The combined algorithm is based on UMFPACK V1.1 and the new frontal matrix strategy discussed here. This combined algorithm is available in Release 12 of the Harwell Subroutine Library [HSL 1996] as the package MA38. In the remainder of this article, we refer to the combined unifrontal/multifrontal algorithm as MA38. We describe our sparse matrix test set, and how we selected it, in Section 5. In Section 6, we consider the influence of a key parameter that is present in both the UMFPACK V1.1 and the MA38 versions of our unsymmetric-pattern multifrontal method. The performance of MA38 is discussed in Section 7, before a few concluding remarks and information on the availability of our codes are given in Section 8.

2. UNFRONTAL METHODS

In a unifrontal scheme [Duff 1984a; Irons 1970; Zitney and Stadtherr 1993; Zitney et al. 1996] the factorization proceeds as a sequence of partial factorizations and eliminations on dense submatrices, called frontal matrices. Although unifrontal methods were originally designed for the solution of finite-element problems [Irons 1970], they can be used on assembled systems [Duff 1984a], and it is this version that we study in this article. For assembled systems, the frontal matrices can be written as

$$\begin{pmatrix} \mathbf{F}_{11} & \mathbf{F}_{12} \\ \mathbf{F}_{21} & \mathbf{F}_{22} \end{pmatrix} \quad (1)$$

where all rows are fully summed (that is, there are no further contributions to come to the rows in (1)) and where the first block column is fully summed. This means that pivots can be chosen from anywhere in the first block column, and within these columns, numerical pivoting with arbitrary row interchanges can be accommodated, since all rows in the frontal matrix are fully summed. We assume, without loss of generality, that the pivots that have been chosen are in the square matrix \mathbf{F}_{11} of (1). \mathbf{F}_{11} is factorized; the Gaussian elimination multipliers overwrite \mathbf{F}_{21} ; and the Schur complement

$$\mathbf{F}_{22} - \mathbf{F}_{21}\mathbf{F}_{11}^{-1}\mathbf{F}_{12} \quad (2)$$

is formed using dense matrix kernels. The submatrix consisting of the rows and columns of the frontal matrix from which pivots have not yet been selected is called the contribution block. In the case above, this is the same as the Schur complement matrix (2).

At the next stage, further rows from the original matrix are assembled with the Schur complement to form another frontal matrix. The frontal matrix is extended in size, if necessary, to accommodate the incoming rows. The overhead is low (compared to a multifrontal method), since each row is assembled only once and there is never any assembly of two (or more) frontal matrices. The entire sequence of frontal matrices is held in the same working array. Data movement is limited to assembling rows of the original matrix into the frontal matrix, and storing rows and columns as they become pivotal. There is never any need to move or assemble the Schur complement into another working array. One important advantage of the method is that only this single working array need reside in memory. Rows of \mathbf{A} can be read sequentially from disk into the frontal matrix. Entries in \mathbf{L} and \mathbf{U} can be written sequentially to disk in the order they are computed. Zitney [1992] gives a detailed description of frontal methods for assembled problems.

An example is shown in Figure 1, where two pivot steps have already been performed on a 5-by-7 frontal matrix (computing the first two rows of \mathbf{U} and columns of \mathbf{L} , respectively). The columns are in pivotal order. Entries in \mathbf{L} and \mathbf{U} are shown in lower case. Row 6 has just been assembled into the current 4-by-7 frontal matrix (shown as a solid box). Columns 3 and 4 are now fully summed and can be eliminated. After this step, rows 7 and 8 must both be assembled before columns 5 and 6 can be eliminated (the dashed box, a 4-by-6 frontal matrix containing rows 5 through 8, and columns 5, 6, 7, 8, 9, and 12). The frontal matrix is, of course, stored without the zero columns, columns 6 and 7 in the dashed box. The dotted box shows the state of the frontal matrix when the next four pivots can be eliminated. To factorize the 12-by-12 sparse matrix in Figure 1, a (dense) working array of size 5-by-7 is sufficient to hold all frontal matrices.

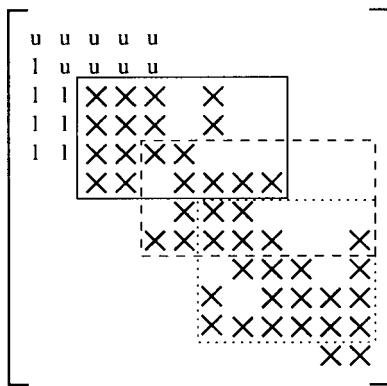


Fig. 1. Frontal method example.

The unifrontal method works well for matrices with small profile, where the profile of a matrix is a measure of how close the nonzero entries are to the diagonal and is given by the expression

$$\sum_{i=1}^n \{ \max_{a_{ij} \neq 0} (i - j) + \max_{a_{ji} \neq 0} (i - j) \},$$

where it is assumed the diagonal is nonzero, so all terms in the summation are nonnegative. For matrices that are symmetric or nearly so, the unifrontal method is typically preceded by an ordering method to reduce the profile such as reverse Cuthill-McKee (RCM) [Chan and George 1980; Cuthill and McKee 1969; Liu and Sherman 1976]. This is typically faster than the sparsity-preserving orderings commonly used by a multifrontal method (such as nested dissection [George and Liu 1981] and minimum degree [Amestoy et al. 1996; George and Liu 1989]). However, for matrices with large profile, the frontal matrix can be large, and an unacceptable amount of fill-in can occur. In particular, we lack effective profile reduction strategies for matrices whose pattern is very unsymmetric.

The unifrontal scheme can easily accommodate numerical pivoting. Because all rows in (1) are fully summed and regular partial pivoting can be performed, it is always possible to choose pivots from all the fully summed columns (unless the matrix is structurally singular).

3. MULTIFRONTAL METHODS

In a multifrontal method [Amestoy and Duff 1989; Duff and Reid 1983; 1984; Liu 1992] for a matrix with a symmetric sparsity pattern, it is common to use an ordering such as minimum degree to reduce the fill-in. An example of a code that is primarily designed for matrices with a symmetric pattern is the Harwell Subroutine Library (HSL) code MA41 [Amestoy and Duff 1989] that will also solve general unsymmetric systems by holding explicit zeros so that the unsymmetric matrix is embedded in

one of symmetric structure. Orderings like minimum degree tend to reduce fill-in much more than profile reduction orderings. The ordering is combined with a symbolic analysis to generate an assembly tree, where each node represents the elimination operations on a frontal matrix and where each edge represents an assembly operation. When using the tree to drive the numerical factorization, the only requirement is that eliminations at any node cannot complete until those at the child nodes have been completed, giving added flexibility for issues such as exploitation of parallelism. As in the unifrontal scheme, the complete frontal matrix (1) cannot normally be factorized, but only a few steps of Gaussian elimination are possible, after which the Schur complement $\mathbf{F}_{22} - \mathbf{F}_{21}\mathbf{F}_{11}^{-1}\mathbf{F}_{12}$ (contribution block) needs to be summed (assembled) with other data at the parent node.

In the unsymmetric-pattern multifrontal method [Davis 1995; Davis and Duff 1991; 1997], the ordering, symbolic analysis, and numerical factorization are performed at the same time. The tree is replaced by a directed acyclic graph (DAG). A contribution block may be assembled into more than one subsequent frontal matrix. For example, consider the LU factors

$$\mathbf{L} + \mathbf{U} - \mathbf{I} = \begin{bmatrix} u_{11} & 0 & 0 & u_{14} & u_{15} & 0 & 0 \\ l_{21} & u_{22} & u_{23} & u_{24} & u_{25} & 0 & u_{27} \\ l_{31} & l_{32} & u_{33} & u_{34} & u_{35} & 0 & u_{37} \\ l_{41} & 0 & 0 & u_{44} & u_{45} & u_{46} & u_{47} \\ 0 & l_{52} & l_{53} & l_{54} & u_{55} & u_{56} & u_{57} \\ 0 & 0 & 0 & 0 & 0 & u_{66} & u_{67} \\ l_{71} & l_{72} & l_{73} & l_{74} & l_{75} & l_{76} & u_{77} \end{bmatrix} \quad (3)$$

of a matrix from Davis and Duff [1997]. The unsymmetric-pattern multifrontal factorization of this matrix is depicted in Figure 2. The heavily shaded regions are the rows and columns of the factors, and the lightly shaded regions are the contribution blocks. The arrows represent the assembly operations from the contribution blocks into solid-lined regions of the same shape in the frontal matrices of the parents. Note that the contribution that the frontal matrix in the lower left of the figure (with the u_{11} pivot) makes to rows 2 and 3 must be assembled into the frontal matrix at the lower right of Figure 3, whereas its contribution to row 4 cannot be assembled into the same frontal matrix.

The first pivot within a frontal matrix (called the “seed” pivot in Davis and Duff [1997]) defines its size. This new frontal matrix is held in a larger working array, to allow room for the assembly of subsequent pivot rows and columns. The next pivots can reside either in the fully summed part or the nonfully summed part. If a potential pivot lies in the nonfully summed part of the frontal matrix, then it is necessary to sum its row and column before it can be used as a pivot. This is possible as long as its fully summed row and column can be accommodated within the larger working array, along with the contribution block and all previous pivot rows and columns of the frontal matrix. This is similar to the node amalgamation [Duff and Reid

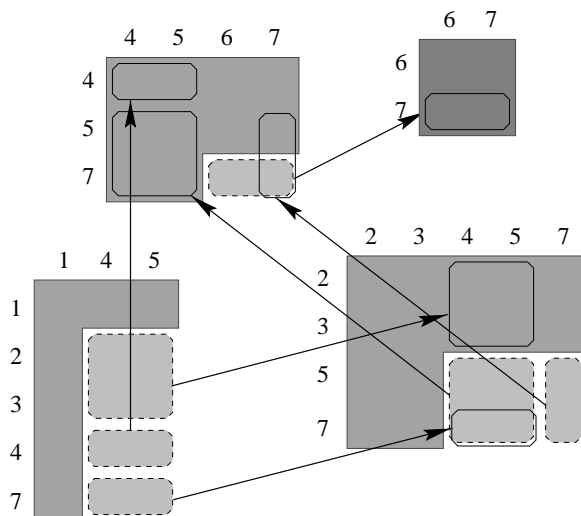


Fig. 2. Assembly DAG for the unsymmetric-pattern multifrontal method.

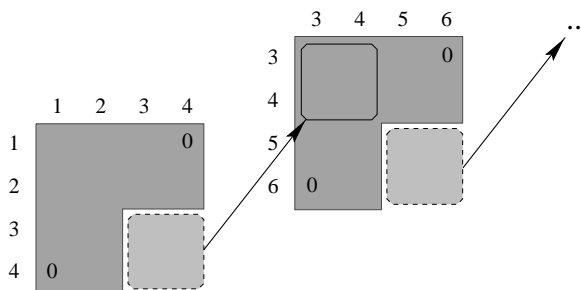


Fig. 3. First two frontal matrices for a pentadiagonal matrix.

1983] used in the symmetric-pattern multifrontal method, except that here we determine the amalgamation during numerical factorization, rather than during the symbolic analysis.

We use the term *normal* multifrontal method to denote a multifrontal method where each frontal matrix is first assembled, then all eliminations are performed, and the contribution block is held (stacked) for assembly at the parent node. Data movement is required both for stacking the contribution block and later assembling it into the parent frontal matrix. Both MA41 and UMFPACK V1.1 are examples of a normal multifrontal method. This terminology is used primarily to distinguish this earlier approach from the new multifrontal approach introduced in this article.

4. COMBINING THE TWO METHODS

Let us now consider an approach that combines some of the best features of the two methods: namely, the lower data movement of unifrontal methods,

and the lower fill-in of normal multifrontal methods. We give an outline of the approach, followed by some implementation details.

4.1 Outline

Assume we have chosen a pivot and determined a frontal matrix as in a normal multifrontal method. At this stage, a normal multifrontal method will select as many pivots as it can from the fully summed part of the frontal matrix, perform the eliminations corresponding to these pivots, store the pivotal rows and columns, and store the contribution block for later assembly at the parent node of the assembly tree.

In UMFPACK V1.1, a potential pivot can be selected from the nonfully summed part. If there is sufficient room in the working array, the row and column of the pivot can be assembled and moved to the fully summed part. Suppose that the rows and columns of such a potential pivot can fit in the working array, but not at the same time as the previous pivotal rows and columns of this frontal matrix. UMFPACK V1.1 would stop the factorization at this point, store the contribution block for later assembly, and continue with a new frontal matrix. In our combined strategy (MA38), we instead perform any arithmetic operations corresponding to the earlier pivots, store their rows and columns (removing them from the frontal matrix), and then assemble the new pivot row and column into the current working array. The contribution block thus need not be stored for later assembly. In this way, we avoid some of the data movement and assemblies of the multifrontal method.

In Figure 3, we show the first two frontal matrices for an n -by- n pentadiagonal matrix. Suppose that UMFPACK V1.1 uses a 4-by-4 working array to hold each frontal matrix and that the pivots are on the diagonal in order. The first two pivot rows and columns fit in the first frontal matrix; this amalgamation causes a_{14} and a_{41} to be treated as “nonzero” entries in the matrix (shown as a zero in the heavily shaded regions). The third pivot row and column do not fit. Once the factorization operations have been completed for the first frontal matrix (using a rank-2 update to the contribution block, a Level 3 BLAS operation) its 2-by-2 contribution block must then be added into the second frontal matrix, in a different working array. A total of $2n$ floating-point values are copied between working arrays, for the $n/2$ contribution blocks.

In MA38, the rank-2 update for the first two pivots is applied, and their pivot rows and columns are stored and are removed from the frontal matrix (just as in the unifrontal method). The 2-by-2 contribution block remains in place, and the third and fourth pivot rows and columns can then be assembled into the working array. The working array now holds what UMFPACK V1.1 would have for its second frontal matrix, but no data has been moved for the 2-by-2 contribution block from the first frontal matrix. As a result, the entire matrix is factorized in a single 4-by-4 working array, and the contribution blocks are held within the working array and are

never stacked, in contrast with the normal multifrontal method discussed in the previous paragraph.

The combined strategy (MA38) allows the use of a general fill-reducing ordering (as does the normal multifrontal method), rather than a profile-reducing ordering. When it encounters submatrices with good profile (as in the pentadiagonal case just considered), it takes advantage of them with a unifrontal strategy and thus has less data movement than the normal multifrontal method.

Although the motivation is different, the idea of continuing with a frontal matrix for some steps before moving to another frontal matrix is similar to recent work in implementing frontal schemes within a domain decomposition environment, for example Duff and Scott [1994], where several fronts are used within a unifrontal context. However, in the case of Duff and Scott [1994], the ordering is done a priori, and no attempt is made to use a minimum-degree ordering.

4.2 Implementation

We now describe how this new frontal matrix strategy is implemented in MA38, which uses a modified version of the minimum-degree algorithm, called the approximate minimum-degree algorithm [Amestoy et al. 1996; Davis and Duff 1997] that uses an upper bound on the degree counts.

MA38 consists of several major steps, each of which comprises several pivot selection and elimination operations. To start a major step, MA38 selects a few (by default 4) columns from those of minimum upper bound degree [Amestoy et al. 1996; Davis and Duff 1997] and computes their patterns, true degrees, and numerical values. A pivot row is selected on the basis of the upper bound on the row degree from those rows with nonzero entries in the selected columns. The pivot must also satisfy a numerical threshold test (it must be at least as large as u times the largest entry in the pivot column, where u is 0.1 by default). Suppose the pivot row and column degrees are r and c , respectively. A k -by- l working array is allocated (k and l are gc and gr , respectively, where by default $g = 2$). The pivot row and column are fully assembled into the working array and define the active frontal matrix. This active frontal matrix is c -by- r but is stored in the k -by- l working array. The approximate-degree update and assembly phase computes the bounds on the degrees of all the rows and columns in this active frontal matrix and assembles previous contribution blocks into the active frontal matrix. A row i in a previous contribution block is assembled into the active frontal matrix if

- (1) the row index i is in the nonzero pattern of the current pivot column, and
- (2) the column indices of the remaining entries in the row are all present in the nonzero pattern of the current pivot row.

Columns of previous contribution blocks are assembled in an analogous manner.

The major step then continues with a sequence of minor steps at each of which another pivot is sought from within the current frontal matrix. These minor steps are repeated until the factorization can no longer continue within the current working array, at which point a new major step is started. When a pivot is chosen in a minor step, its rows and columns are fully assembled into the working array and redefine the active frontal matrix. If there are new rows or columns in the frontal matrix, then the approximate degree update and assembly phase is repeated, as described above. Otherwise, the update and assembly phase can be skipped (analogous to mass elimination in a minimum-degree ordering algorithm [George and Liu 1989]).

After a pivot is selected (at the start of a major step or in a minor step) the corresponding row and column of \mathbf{U} and \mathbf{L} are computed, but the updates from this pivot are not necessarily performed immediately. For efficient use of the Level 3 BLAS, it is better to accumulate a few updates (typically up to 16, if possible) and perform them at the same time.

MA38 and UMFPACK V1.1 differ in how the minor step is performed, although both use the same amalgamation parameter, g . To find a pivot in this minor step, a single candidate column from the current frontal matrix is first selected, choosing one with least value for the upper bound of the column degree, and any pending updates are applied to this column. The column is assembled into a separate work vector, and a pivot row is selected on the basis of the upper bound on the row degrees and a numerical threshold test. Suppose the candidate pivot row and column have degrees r' and c' , respectively. Three conditions apply (where p is the number of pivots currently stored in the active frontal matrix):

- (1) If $r' > l$ or $c' > k$, then factorization can no longer continue within the active frontal matrix. Any pending updates are applied. The LU factors are stored. The active contribution block is saved for later assembly into a subsequent frontal matrix. The major step is now complete.
- (2) If $r' \leq l - p$ and $c' \leq k - p$, then the candidate pivot can fit into the active frontal matrix without removing the p pivots already stored there. Set $p \leftarrow p + 1$. Factorization continues within the active frontal matrix by starting another minor step.
- (3) Otherwise, if $l - p < r' \leq l$ or $k - p < c' \leq k$, then the candidate pivot can fit, but only if some of the previous p pivots are shifted out of the current frontal matrix. Any pending updates are applied. The LU factors corresponding to the pivot rows and columns are removed from the front and stored. The active contribution block is left in place. Set $p \leftarrow 1$. Factorization continues within the active frontal matrix by commencing another minor step.

UMFPACK V1.1 uses the same strategy as MA38, described above, except for Case (3). In Case (3), UMFPACK V1.1 saves the contribution block for later assembly and terminates the major step. We know of no other multifrontal

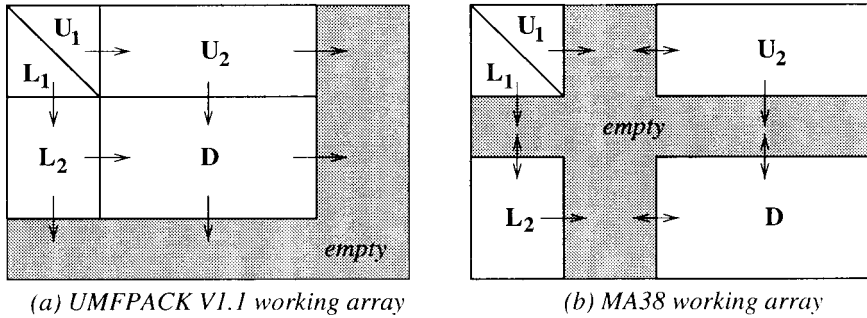


Fig. 4. Data structures for the active frontal matrix.

method that allows the factorization to proceed within the current frontal matrix, as in Case (3). Case (1) does not occur in unifrontal methods, which are given a working array large enough to hold the largest frontal matrix. Cases (2) and (3) do occur in unifrontal methods. Taking simultaneous advantage of all three cases can significantly reduce the memory requirements, data movement, and assembly operations, while still allowing the use of orderings that reduce fill-in.

Figure 4 illustrates how the working array can be organized. The matrices L_1 , L_2 , U_1 , and U_2 in the figure are the columns and rows of the LU factors corresponding to the pivots eliminated within this frontal matrix. The matrix D is the contribution block. The arrows denote how these matrices grow as new pivots are added. When pivots are removed from the working array in Figure 4(b), for Case (3) above, the contribution block does not need to be moved, thus accommodating the more dynamic situation without further data movement. For coding reasons, the working array in MA38 is held in reverse order, but the data movement is similar to that shown in the figure.¹

5. TEST MATRICES AND COMPUTING PLATFORM

In the next two sections, we discuss some experiments on the selection of the amalgamation parameter g for MA38 and UMFPACK V1.1 and compare the performance of MA38 with other sparse matrix codes. In both sections, we use the same set of sparse test matrices.

The sparse matrix collection at the University of Florida [Davis 1997] contains 264 unsymmetric sparse matrices in assembled equation form (as opposed to unassembled finite-element form). This set includes the Harwell-Boeing Sparse Matrix Collection [Duff et al. 1989], Saad's collection [Saad 1994], Bai's collection [Bai et al. 1996], and matrices from other sources [Feldmann et al. 1996; Vavasis 1996; Zitney 1992; Zitney et al. 1996].

¹A similar data organization is employed by the unifrontal HSL code MA42 [Duff and Scott 1993; 1996b].

Table I. Test Matrices

Name	n	Number of Entries	Sym.	Discipline, Comments, and Source
AV41092	41,092	1,683,902	0.001	partial diff. eqn.: 2D, wild coefficients [Vavasis 1996]
TWOTONE	120,750	1,224,224	0.245	circuit simulation: harmonic balance method [Feldman et al. 1996]
PSMIGR 1	3,140	543,162	0.479	demography: US county-to-county migration [Duff et al. 1989]
LHR71C	70,304	1,528,092	0.002	chemical eng.: light hydrocarbon recovery [Zitney et al. 1996]
ONETONE1	36,057	341,088	0.074	circuit simulation: harmonic balance method [Feldman et al. 1996]
ONETONE2	36,057	227,628	0.113	circuit simulation: harmonic balance method [Feldman et al. 1996]
LHR14C	14,270	307,858	0.007	chemical eng.: light hydrocarbon recovery [Zitney et al. 1996]
RW5151	5,151	20,199	0.490	probability: random walk Markov chain [Bai et al. 1996]
ORANI678	2,529	90,158	0.071	economics: Australia [Duff et al. 1989]
RDIST1	4,134	94,408	0.059	chemical eng.: reactive distillation [Zitney 1992]

The 10 matrices that we will use in the next two sections are listed in Table I. The table lists the name, order, number of entries, structural symmetry, the discipline from which the matrix comes, and additional comments. The structural symmetry is the ratio of the number of *matched* off-diagonal entries to the total number of off-diagonal entries. An entry, a_{ij} ($j \neq i$), is matched if a_{ji} is also an entry. This means that a symmetric matrix has a structural symmetry value of 1.0. This measure is one minus the index of asymmetry, originally defined in Duff [1984b].

We chose the test matrices using the following strategy. Since we are considering a method for unsymmetric-patterned matrices, we selected all nonsingular matrices from our collection with structural symmetry less than or equal to 0.5 (89 matrices). We ran all six codes discussed in Section 7 for each of the 89 matrices and selected a matrix for our comparisons in this article if the fastest factorization time of these six codes was greater than 0.5 seconds. This leaves 29 matrices from which we discarded one matrix too large for most methods on the workstation used for these experiments (from the same application as TWOTONE) and 18 matrices from sets already represented (two similar to PSMIGR 1, and all but two of the LHR series²).

²Due to a modeling error, some of the original LHR matrices [Zitney et al. 1996] were extremely ill conditioned. We used the following procedure to correct the matrices. If \mathbf{A} is the original matrix, the corrected matrix is $\mathbf{A} + \mathbf{P}^T \mathbf{F} \mathbf{Q}^T$, where \mathbf{P} and \mathbf{Q} are permutation matrices such that $\mathbf{P} \mathbf{A} \mathbf{Q}$ is in block upper triangular form. The matrix \mathbf{F} is diagonal, with $\mathbf{F}_{ii} = 0.001$ if $[\mathbf{P} \mathbf{A} \mathbf{Q}]_{ii} \geq 0$ and -0.001 otherwise. The corrected matrix has the same nonzero

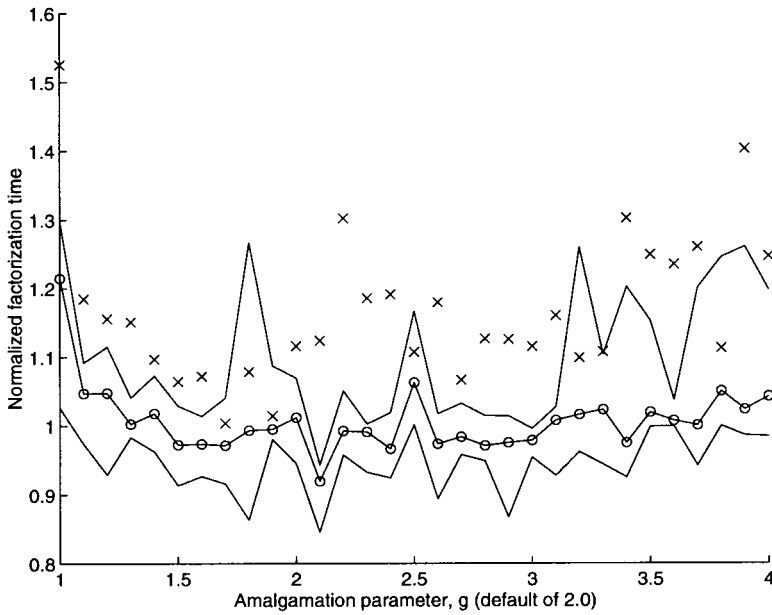


Fig. 5. Normalized run times for 10 matrices, as a function of the amalgamation parameter, g . The median and quartiles for MA38 are given by the solid lines and the median for UMFPACK V1.1 by the crosses.

All experiments reported in this article are on a SUN UltraSparc Model 170, with 256MB of main memory and a single 167MHz processor. Version 4.0 of the SUN Fortran and C compilers were used, with identical optimization parameters (SuperLU is written in C, the other codes in Fortran 77). The BLAS that we use are from Daydé and Duff [1996]. The double-precision matrix-matrix multiply routine, DGEMM, which most of the methods use, runs at about 80Mflops on this workstation. (Note that the single-precision version, SGEMM, runs at about 145Mflops.) The theoretical peak performance is 333Mflops for both double and single precision.

6. AMALGAMATION PARAMETER

We ran MA38 and UMFPACK V1.1 on all matrices in Table I, with g ranging from 1.0 to 4.0 in increments of 0.1. For each matrix, we found the median, with respect to g , of the MA38 run times. The MA38 and UMFPACK V1.1 run times for this matrix were divided by this median to obtain normalized MA38 and UMFPACK V1.1 run times. Then, for each value of g , we found the median (and quartiles) of the normalized MA38 and UMFPACK V1.1 run times for all 10 matrices. Figure 5 shows the median and quartiles for MA38

pattern as the original matrix. We selected the value 0.001 by trial and error. We used the smallest value we found for which Matlab could compute at least two digits of accuracy in the solution.

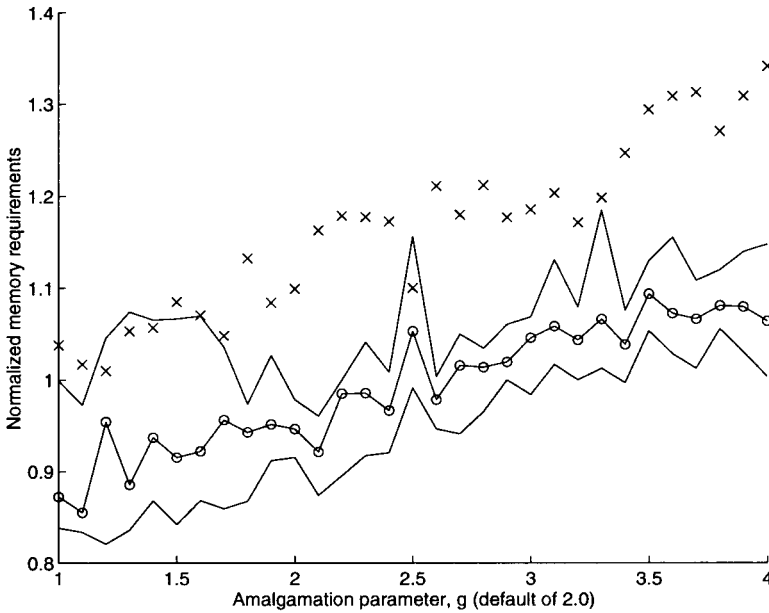


Fig. 6. Normalized memory requirements for 10 matrices, as a function of the amalgamation parameter, g . The median and quartiles for MA38 are given by the solid lines and the median for UMFPACK V1.1 by the crosses.

and the median for UMFPACK V1.1 of the normalized times for the 10 matrices, plotted as a function of g . The MA38 normalized factorization times are plotted as solid lines, with the median time data points circled. The mean normalized factorization time for UMFPACK V1.1 is given by x's. A similar comparison of the normalized memory requirements (excluding fragmentation of the work arrays) of both methods is shown in Figure 6.

Although the fill-in and operation count (not shown) are typically lowest when the minimum amount of memory is allocated for each frontal matrix ($g = 1$), the factorization time is often high because of the additional data movement required to assemble the contribution blocks and the fact that the dense matrix kernels are more efficient for the larger frontal matrices that are produced when $g > 1$.

From the results in Figures 5 and 6, we examine the effect and sensitivity of execution time and memory requirements to changes in the value of g . There are some rapid fluctuations in the graphs, particularly for UMFPACK V1.1. This is caused by changes in pivot order with different values of g . A small change in the working array size can affect whether or not a candidate pivot is selected in a minor step. If it is not selected, we start a new frontal matrix, and a new pivot candidate is found in a major step which considers all columns. This is a very different pivot selection strategy from that used in the minor steps. This high sensitivity is not so noticeable with MA38 because it can continue for longer with the minor step strategy after removing pivot rows and columns from the active frontal matrix.

Although the frontal matrices become larger with larger values of g , there are fewer of them, and thus the median memory requirements increase only slightly as g increases. For these 10 matrices, the median factorization time is lowest when g is 2.1. From experiments with a wider range of matrices, we have found that a default value of $g = 2$ is a reasonable trade-off between memory usage and factorization time.

We also notice, from the results in this section, that exploiting a unifrontal strategy within a multifrontal code (MA38) can reduce both factorization time and memory requirements. Comparing MA38 with its predecessor, UMFPACK V1.1, Figures 5 and 6 show a median improvement of 18% in the run time and 14% in the memory requirements for these 10 matrices. Although not shown in the plots, MA38 is almost four times faster than UMFPACK V1.1 for one matrix (LHR71C).

7. PERFORMANCE

In this section, we compare the performance of the combined unsymmetric-pattern unifrontal/multifrontal code (MA38) with the unsymmetric-pattern multifrontal code UMFPACK V1.1 [Davis 1995; Davis and Duff 1991; 1997], the symmetric-pattern multifrontal code MA41 [Amestoy et al. 1989], the unifrontal code MA42, [Duff 1984a; Duff and Scott 1993; 1996b], and two sparse matrix factorization codes based on partial pivoting: MA48 [Duff and Reid 1993; 1996] and SuperLU Version 1.0 [Demmel et al. 1995; 1997].

Each code can factorize general unsymmetric matrices, and all use dense matrix kernels [Dongarra et al. 1990] to some extent. Each code was given 1.8GB of (virtual) memory to factorize the matrices listed in Table I. Each code has a set of input parameters that control its behavior. We used the recommended defaults for most of these, with a few exceptions that we now indicate.

The default numerical threshold u varies with each code. For MA41, it is 0.01, and the test is by rows, not columns. SuperLU and MA42 select the largest entry in each column (effectively, $u = 1.0$). The rest use $u = 0.1$. We used the default value in each case.

Although MA41 has an option for exploiting shared-memory parallelism, we did not use this option in the runs for this article. Using a nondefault option, MA41 can preorder a matrix, using a maximum transversal algorithm, to ensure the diagonal of the permuted matrix is zero-free. We selected this option because it is recommended for matrices with highly unsymmetric nonzero patterns. This is followed by a default approximate minimum-degree ordering (using the HSL code MC47 [Amestoy et al. 1996]) on the nonzero pattern of $\mathbf{A} + \mathbf{A}^T$. The supernodal partial pivoting code SuperLU preorders the columns via its default method, a multiple minimum-degree ordering (MMD [George and Liu 1989]) on the nonzero pattern of $\mathbf{A}^T\mathbf{A}$.

We used a revised version of the unifrontal code MA42 that takes advantage of explicit zeros in the frontal matrix. MA42 is able to operate

both in-core and out-of-core, using direct-access files. It was primarily designed for a finite-element entry, but we used a simplified driver (MA43), that does not use out-of-core storage and assumes the matrix is held in equation form (although one matrix, AV41092, is obtained from finite-element calculations). In columns “Mem. Used” and “Min. Mem.,” we show the storage required for the factorization if the factors are held in-core and an estimate of the memory required for the out-of-core factorization, respectively. In one case (TWOTONE), there was insufficient storage to hold the factors, but we compute an estimate which we display.

The results, shown in Table II, include the following statistics for each code (all times are CPU times):

- (1) Factorization time, which includes preordering and symbolic factorization, if any.
- (2) Refactorization time, which is the numerical factorization of a matrix whose pivot ordering and symbolic factors are known. It excludes preordering and symbolic factorization.
- (3) Solve time, excluding iterative refinement.
- (4) Solve time, including iterative refinement. MA41, MA38, MA48, and SuperLU all use the same method [Arioli et al. 1989]. MA42 and UMFPACK V1.1 do not provide this option.
- (5) Entries in the LU factors, in millions.
- (6) Memory used to obtain the timings listed in the tables, in millions of eight-byte double-precision words. For MA38, UMFPACK V1.1, and MA48, this includes fragmentation within the work arrays.³ Recall that the workstation used for the experiments has 33.6 million words of main memory (256MB), and about 242 million words of virtual address space were given to each code (1.8GB).
- (7) Minimum memory requirements, in millions of words. For the unifrontal code MA42, this is the in-core memory required for out-of-core factorization (it also includes the memory required to hold the matrix **A**). For MA38, UMFPACK V1.1, and MA48, this is the minimum amount required to guarantee a successful factorization and is the statistic used in Figure 6. This excludes fragmentation in the work arrays, which is removed via garbage collection whenever necessary. The partial pivoting code SuperLU and the sequential option of the symmetric-pattern multifrontal code MA41 do not have fragmentation in their work arrays, and thus do not require garbage collection. For these two methods the minimum required memory and actual memory used are the same, although we needed to add another counter to the HSL version of MA41 to compute this statistic.

³MA48 does not compute this statistic.

Table II. Results

Matrix/ Code	Factor Time (sec.)	Refac. Time (sec.)	Solve Time (sec.)	Solve, w/Iter. (sec.)	Entries in LU (10 ⁶)	Mem. Used (10 ⁶)	Min. Mem. (10 ⁶)	flop Count (10 ⁶)
AV41092								
MA38	1307.9	1618.3	9.88	55.18	39.38	89.58	59.60	57110
UMFPACK	1502.9	1689.0	10.64	—	41.33	102.53	58.49	70290
V1.1								
MA41	<u>296.6</u>	<u>254.0</u>	<u>1.39</u>	<u>7.72</u>	<u>16.53</u>	<u>22.65</u>	22.65	<u>11380</u>
MA42	4236.3	4232.6	32.74	—	135.35	148.66	<u>10.10</u>	235079
MA48	3335.7	1296.1	5.57	50.73	27.31	—	58.46	59040
SuperLU	3799.0	2962.8	14.04	136.43	39.95	52.56	52.56	64970
TWOONE								
MA38	<u>220.7</u>	<u>171.4</u>	1.46	<u>6.24</u>	<u>9.75</u>	<u>26.46</u>	17.94	<u>6988</u>
UMFPACK	<u>228.5</u>	<u>174.9</u>	1.52	—	<u>9.59</u>	35.04	19.31	<u>6493</u>
V1.1								
MA41	817.2	741.1	2.85	13.56	26.44	42.97	42.97	38230
MA42	—	—	—	—	—	333.60	<u>4.64</u>	—
MA48	725.0	306.9	<u>1.29</u>	11.70	10.86	—	24.65	14680
SuperLU	758.0	697.6	8.23	147.22	24.73	37.77	37.77	12420
PSMIGR 1								
MA38	219.0	197.5	<u>0.53</u>	3.01	<u>6.37</u>	46.34	22.86	<u>9412</u>
UMFPACK	<u>207.4</u>	<u>195.6</u>	0.71	—	<u>6.36</u>	26.28	22.02	<u>9428</u>
V1.1								
MA41	<u>191.7</u>	<u>188.3</u>	<u>0.53</u>	<u>2.14</u>	<u>6.28</u>	20.36	20.36	<u>9214</u>
MA42	256.0	255.0	0.66	—	8.27	17.07	<u>8.75</u>	13856
MA48	<u>206.1</u>	<u>178.7</u>	<u>0.51</u>	5.25	<u>6.44</u>	—	14.01	10580
SuperLU	938.7	774.5	1.80	20.21	8.71	<u>11.10</u>	11.10	16630
LHR71C								
MA38	<u>114.5</u>	<u>92.9</u>	<u>1.01</u>	<u>22.25</u>	<u>6.93</u>	<u>11.72</u>	10.87	<u>496</u>
UMFPACK	444.3	422.3	1.96	—	8.35	13.26	12.55	732
V1.1								
MA41	1012.8	995.7	3.63	82.27	18.84	23.75	23.75	4683
MA42	340.6	337.4	1.41	—	12.83	16.15	<u>2.97</u>	1202
MA48	492.3	286.5	<u>1.11</u>	30.51	<u>6.51</u>	—	16.54	695
SuperLU	479.3	451.4	2.26	64.79	<u>7.15</u>	<u>11.57</u>	11.57	<u>487</u>
ONETONE1								
MA38	<u>60.8</u>	<u>57.3</u>	<u>0.54</u>	<u>2.16</u>	<u>4.69</u>	15.92	7.92	<u>2148</u>
UMFPACK	73.7	69.2	<u>0.58</u>	—	<u>5.05</u>	17.34	9.19	2855
V1.1								
MA41	194.7	189.1	0.82	3.88	9.50	12.97	12.97	8169
MA42	164.9	164.2	1.64	—	16.58	18.52	<u>1.49</u>	6012
MA48	323.4	109.7	<u>0.56</u>	4.87	<u>5.11</u>	—	11.29	4564
SuperLU	109.3	97.9	1.16	9.49	<u>4.68</u>	<u>7.10</u>	7.10	2540

(8) Floating-point operation count in the factorization phase, in millions. The values given are obtained from the codes. Operations for the assembly phase are not included.

Results within 10% of the best for each statistic and matrix are underlined. We compared the solutions obtained from each code and found that all codes that provide iterative refinement compute the solutions with comparable accuracy, in terms of the scaled residual and the relative error. When

Table II. *Continued*

Matrix/ Code	Factor Time (sec.)	Refac. Time (sec.)	Solve Time (sec.)	Solve, w/Iter. (sec.)	Entries in LU (10 ⁶)	Mem. Used (10 ⁶)	Min. Mem. (10 ⁶)	flop Count (10 ⁶)
ONETONE2								
MA38	<u>10.7</u>	<u>5.0</u>	0.29	<u>1.25</u>	<u>1.27</u>	4.09	2.54	159
UMFPACK V1.1	12.6	6.8	<u>0.22</u>	—	1.46	4.82	3.33	217
MA41	19.1	14.9	0.29	1.61	2.68	3.90	3.90	601
MA42	35.2	34.8	0.64	—	7.10	8.04	<u>0.73</u>	684
MA48	36.6	8.0	<u>0.21</u>	2.02	<u>1.26</u>	—	3.36	223
SuperLU	<u>9.2</u>	6.6	0.44	3.77	<u>1.31</u>	<u>2.70</u>	2.70	<u>132</u>
LHR14C								
MA38	<u>8.9</u>	<u>4.8</u>	<u>0.17</u>	<u>1.19</u>	<u>1.23</u>	<u>2.21</u>	2.01	<u>64</u>
UMFPACK V1.1	11.4	5.7	0.20	—	1.62	3.40	2.70	122
MA41	28.9	27.1	0.32	1.58	3.14	4.54	4.54	518
MA42	19.2	18.6	0.28	—	2.57	3.43	<u>0.81</u>	235
MA48	21.1	<u>4.7</u>	<u>0.17</u>	1.81	<u>1.22</u>	—	3.15	95
SuperLU	12.0	6.7	0.35	3.81	1.37	<u>2.26</u>	2.26	83
RW5151								
MA38	2.4	<u>1.4</u>	<u>0.06</u>	<u>0.14</u>	<u>0.41</u>	1.24	0.65	47
UMFPACK V1.1	3.1	2.2	<u>0.06</u>	—	0.48	1.39	0.89	69
MA41	2.2	2.1	0.07	<u>0.15</u>	0.67	0.89	0.89	91
MA42	<u>1.6</u>	<u>1.6</u>	0.07	—	0.79	0.89	<u>0.08</u>	63
MA48	6.5	2.5	<u>0.05</u>	0.71	0.47	—	1.04	126
SuperLU	<u>1.8</u>	<u>1.6</u>	0.10	0.73	<u>0.39</u>	<u>0.65</u>	0.65	<u>33</u>
ORANI678								
MA38	<u>1.2</u>	0.4	<u>0.02</u>	<u>0.21</u>	<u>0.11</u>	0.74	0.46	<u>4</u>
UMFPACK V1.1	<u>1.2</u>	0.4	<u>0.02</u>	—	<u>0.11</u>	<u>0.67</u>	0.48	<u>4</u>
MA41	5.3	3.6	0.04	<u>0.21</u>	0.36	2.79	2.79	70
MA42	8.6	8.5	0.08	—	1.00	2.78	1.82	267
MA48	<u>1.1</u>	<u>0.2</u>	<u>0.02</u>	<u>0.20</u>	<u>0.13</u>	—	<u>0.39</u>	10
SuperLU	39.7	1.6	0.16	1.70	0.80	1.00	1.00	31
RDIST1								
MA38	1.6	0.8	<u>0.04</u>	<u>0.24</u>	0.44	0.97	0.74	22
UMFPACK V1.1	2.0	1.1	0.06	—	0.54	1.39	0.97	31
MA41	<u>0.7</u>	<u>0.4</u>	<u>0.03</u>	<u>0.20</u>	<u>0.28</u>	<u>0.49</u>	0.49	<u>10</u>
MA42	2.3	2.2	0.08	—	1.02	1.22	<u>0.20</u>	90
MA48	7.4	1.3	0.06	0.53	0.41	—	1.04	25
SuperLU	1.6	0.8	0.08	0.94	0.39	0.60	0.60	17

iterative refinement is not in use, MA38 and UMFPACK V1.1 produce relative errors with a loss of about two digits of accuracy for six of the matrices when compared with the other methods, and the same accuracy for the other four.

Over all the codes, MA38 has the fastest factorization time for four out of the 10 matrices and is within 10% of the fastest time for two more matrices. Except for one matrix (AV41092) it never takes more than about twice the time of the fastest code. It also has the fastest refactorization time for five

matrices and is within 10% of the fastest time for one more matrix. The solve phase of MA38 is the fastest for four matrices and within 10% of the fastest for two more. Although the time for MA38 solution with iterative refinement seems to rank even better, both the MA48 and SuperLU iterative refinement times include also the calculation of the matrix condition number and an estimate of the forward error. Typically this requires twice as much time as only computing the backward error (as in the iterative refinement runs for MA38 and MA41). The memory requirements of MA38 are rarely the lowest but are usually comparable to the other in-core codes for most matrices. The floating-point operation count for MA38 is often much less than its predecessor, UMFPACK V1.1.

8. SUMMARY

We have demonstrated how the advantages of the unifrontal and multifrontal approaches can be combined. The resulting algorithm (MA38) performs well for unsymmetric matrices from a wide range of disciplines and is an improvement over the previous unsymmetric-pattern multifrontal code (UMFPACK V1.1). Other differences between UMFPACK V1.1 and MA38 include an option of overwriting the matrix \mathbf{A} with its LU factors, printing of input and output parameters, iterative refinement with sparse backward error analysis [Arioli et al. 1989], avoidance of an extra copy of the numerical values of \mathbf{A} when iterative refinement is not in use, more use of Level 3 BLAS within the numerical refactorization routine, and a simpler calling interface. These features improve the robustness of the code and result in a modest decrease in memory use.

Since the codes being compared all offer quite different capabilities and are designed for different environments and different classes of matrices, the results should not be interpreted as a direct comparison between them. For example, MA38 is designed for structurally unsymmetric matrices. A code like MA41 would normally be expected to perform much better on matrices that are symmetrically structured or nearly so. We also note that our current discussion compares performance on only one machine, and as shown in Duff and Scott [1996a], comparative behavior can be strongly influenced by the computing platform being used. However, what we would like to highlight is the improvement that our new technique used in MA38 brings to the unsymmetric-pattern multifrontal method and that MA38 is at least comparable in performance with other sparse matrix codes on our unsymmetric test set.

The combined unifrontal/multifrontal method is available as the Fortran 77 codes, UMFPACK Version 2.2 in Netlib [Dongarra and Grosse 1987],⁴ and MA38 in Release 12 of the Harwell Subroutine Library [HSL 1996].⁵

⁴Versions 1.0 through 2.2 of UMFPACK may only be used for research, education, or benchmarking by academic users and by the U.S. Government. For a copy, send email to netlib@ornl.gov with the message send index from linalg. Other users may use UMFPACK only for benchmarking purposes. Version 2.2 includes complex and complex*16 codes.

ACKNOWLEDGMENTS

We would like to thank Nick Gould, John Reid, and Jennifer Scott from the Rutherford Appleton Laboratory for their helpful comments.

REFERENCES

- AMESTOY, P. R. AND DUFF, I. S. 1989. Vectorization of a multiprocessor multifrontal code. *Int. J. Supercomput. Appl. High Perform. Eng.* 3, 3, 41–59.
- AMESTOY, P., DAVIS, T. A., AND DUFF, I. S. 1996. An approximate minimum degree ordering algorithm. *SIAM J. Matrix Anal. Appl.* 17, 4, 886–905.
- AMESTOY, P. R., DAYDÉ, M., AND DUFF, I. S. 1989. Use of computational kernels in the solution of full and sparse linear equations. In *Proceedings of the International Workshop on Parallel and Distributed Algorithms* (Gers, France, Oct. 3–6), M. Cosnard, Y. Robert, P. Quinton, and M. Raynal, Eds. North-Holland Publishing Co., Amsterdam, The Netherlands, 13–19.
- ARIOLI, M., DEMMEL, J. W., AND DUFF, I. S. 1989. Solving sparse linear systems with sparse backward error. *SIAM J. Matrix Anal. Appl.* 10, 165–190.
- BAI, Z., DAY, D., DEMMEL, J., AND DONGARRA, J. 1996. Test matrix collection (non-Hermitian eigenvalue problems), Release 1. Tech. Rep. (Sept.). University of Kentucky, Lexington, KY. Available at <ftp://ftp.ms.uky.edu/pub/misc/bai/Collection>.
- CHAN, W. M. AND GEORGE, A. 1980. A linear time implementation of the reverse Cuthill-McKee algorithm. *BIT* 20, 8–14.
- CUTHILL, E. AND MCKEE, J. 1969. Reducing the bandwidth of sparse symmetric matrices. In *Proceedings of the 24th National Conference of the Association for Computing Machinery*. Brandon Press, 157–172.
- DAVIS, T. A. 1995. Users' guide to the unsymmetric-pattern multifrontal package (UMFPACK, version 1.1). Tech. Rep. TR-95-004. CISE Department, University of Florida, Gainesville, FL.
- DAVIS, T. A. 1997. University of Florida sparse matrix collection. CISE Department, University of Florida, Gainesville, FL. <http://www.cise.ufl.edu/~davis> and <ftp://ftp.cise.ufl.edu/pub/faculty/davis>.
- DAVIS, T. A. AND DUFF, I. S. 1991. Unsymmetric-pattern multifrontal methods for parallel sparse LU factorization. Tech. Rep. TR-91-023. CISE Department, University of Florida, Gainesville, FL.
- DAVIS, T. A. AND DUFF, I. S. 1997. An unsymmetric-pattern multifrontal method for sparse LU factorization. *SIAM J. Matrix Anal. Appl.* 18, 1, 140–158.
- DAYDÉ, M. J. AND DUFF, I. S. 1996. A blocked implementation of Level 3 BLAS for RISC processors. Tech. Rep. RAL-TR-96-014. Rutherford Appleton Lab., Didcot, Oxon, United Kingdom. Also ENEECIHT-IRIT Tech. Rep. RT/APO/96/1 and CERFACS Rep. TR/PA/96/06.
- DEMMEL, J. W., ESIENSTAT, S. C., GILBERT, J. R., LI, X. S., AND LUI, J. W. H. 1995. A supernodal approach to sparse partial pivoting. Tech. Rep. UCB/CSD-95-883. Computer Science Department, University of California at Berkeley, Berkeley, CA. Also Xerox PARC CSL-95-03 and LAPACK Working Note #103.
- DEMMEL, J. W., GILBERT, J. R., AND LI, X. S. 1997. SuperLU users' guide. Tech. Rep. UC Berkeley, Berkeley, CA.
- DONGARRA, J. J. AND GROSSE, E. 1987. Distribution of mathematical software via electronic mail. *Commun. ACM* 30, 5 (May 1987), 403–407.
- DONGARRA, J. J., DU CROZ, J. J., HAMMARLING, S., AND DUFF, I. S. 1990. A set of level 3 basic linear algebra subprograms. *ACM Trans. Math. Softw.* 16, 1 (Mar. 1990), 1–17.
- DUFF, I. S. 1984a. Design features of a frontal code for solving sparse unsymmetric linear systems out-of-core. *SIAM J. Sci. Comput.* 5, 270–280.

⁵The Harwell Subroutine Library is available from AEA Technology, Harwell; the contact is Dr. Scott Roberts, Harwell Subroutine Library Manager, B 552, AEA Technology Plc, Harwell, Didcot, Oxon OX11 0RA, telephone +44 1235 434988, fax +44 1235 434136, email Scott.Roberts@aeat.co.uk, who will provide details of price and conditions of use.

- DUFF, I. S. 1984b. The solution of nearly symmetric sparse linear systems. In *Computing Methods in Applied Sciences and Engineering*, R. Glowinski and J. Lions, Eds. North-Holland Publishing Co., Amsterdam, The Netherlands, 57–74.
- DUFF, I. S. AND REID, J. K. 1983. The multifrontal solution of indefinite sparse symmetric linear systems. *ACM Trans. Math. Softw.* 9, 302–325.
- DUFF, I. S. AND REID, J. K. 1984. The multifrontal solution of unsymmetric sets of linear equations. *SIAM J. Sci. Comput.* 5, 3, 633–641.
- DUFF, I. S. AND REID, J. K. 1993. MA48, a Fortran code for direct solution of sparse unsymmetric linear systems of equations. Tech. Rep. RAL-93-072. Rutherford Appleton Lab., Didcot, Oxon, United Kingdom.
- DUFF, I. S. AND REID, J. K. 1996. The design of MA48, a code for the direct solution of sparse unsymmetric linear systems of equations. *ACM Trans. Math. Softw.* 22, 2 (June), 187–226.
- DUFF, I. S. AND SCOTT, J. A. 1993. MA42—A new frontal code for solving sparse unsymmetric systems. Tech. Rep. RAL 93-064. Rutherford Appleton Lab., Didcot, Oxon, United Kingdom.
- DUFF, I. S. AND SCOTT, J. A. 1994. The use of multiple fronts in Gaussian elimination. In *Proceedings of the 5th SIAM Conference on Applied Linear Algebra* (Philadelphia, PA), J. Lewis, Ed. SIAM, Philadelphia, PA, 567–571.
- DUFF, I. S. AND SCOTT, J. A. 1996a. A comparison of frontal software with other sparse direct solvers. Tech. Rep. RAL-TR-96-102. Rutherford Appleton Lab., Didcot, Oxon, United Kingdom.
- DUFF, I. S. AND SCOTT, J. A. 1996b. The design of a new frontal code for solving sparse unsymmetric systems. *ACM Trans. Math. Softw.* 22, 1 (Mar.), 30–45.
- DUFF, I. S., GRIMES, R. G., AND LEWIS, J. G. 1989. Sparse matrix test problems. *ACM Trans. Math. Softw.* 15, 1 (Mar.), 1–14.
- FELDMAN, P., MELVILLE, R., AND LONG, D. 1996. Efficient frequency domain analysis of large nonlinear analog circuits. In *Proceedings of the IEEE Conference on Custom Integrated Circuits* (Santa Clara, CA). IEEE Press, Piscataway, NJ.
- GEORGE, A. AND LIU, J. W. H. 1981. *Computer Solution of Large Sparse Positive Definite Systems*. Prentice-Hall, Englewood Cliffs, NJ.
- GEORGE, A. AND LUI, J. W. H. 1989. The evolution of the minimum degree ordering algorithm. *SIAM Rev.* 31, 1, 1–19.
- HSL. 1996. *Harwell Subroutine Library: A Catalogue of Subroutines (Release 12)*. AEA Technology, Didcot, Oxon, United Kingdom.
- IRONS, B. M. 1970. A frontal solution program for finite element analysis. *Int. J. Numer. Method. Eng.* 2, 5–32.
- LIU, J. W. H. 1992. The multifrontal method for sparse matrix solution: Theory and practice. *SIAM Rev.* 34, 1 (Mar. 1992), 82–109.
- LIU, J. W. H. AND SHERMAN, A. H. 1976. Comparative analysis of the Cuthill-McKee and the reverse Cuthill-McKee ordering algorithms for sparse matrices. *SIAM J. Numer. Anal.* 13, 2, 198–213.
- SAAD, Y. 1994. Sparskit: A basic tool kit for sparse matrix computation, version 2. Tech. Rep. Computer Science Department, Univ. of Minnesota, Minneapolis, MN.
- VAVASIS, S. A. 1996. Stable finite elements for problems with wild coefficients. *SIAM J. Numer. Anal.* 33, 3, 890–916.
- ZITNEY, S. E. 1992. Sparse matrix methods for chemical process separation calculations on supercomputers. In *Proceedings of the 1992 Conference on Supercomputing* (Supercomputing '92, Minneapolis, MN, Nov. 16–20), R. Werner, Ed. IEEE Computer Society Press, Los Alamitos, CA, 414–423.
- ZITNEY, S. E. AND STADTHERR, M. A. 1993. Supercomputing strategies for the design and analysis of complex separation systems. *Ind. Eng. Chem. Res.* 32, 604–612.
- ZITNEY, S. E., MALLIA, J., DAVIS, T. A., AND STADTHERR, M. A. 1996. Multifrontal vs. frontal techniques for chemical process simulation on supercomputers. *Comput. Chem. Eng.* 20, 6-7, 641–646.

Received: October 1995; revised: September 1997 and January 1998; accepted: September 1998