



1 SUMMARY

To **solve a sparse symmetric system of linear equations**. Given a sparse symmetric matrix $\mathbf{A} = \{a_{ij}\}_{n \times n}$ and an n -vector \mathbf{b} (or an $n \times s$ matrix \mathbf{B}), this subroutine solves the system $\mathbf{Ax}=\mathbf{b}$ ($\mathbf{AX}=\mathbf{B}$). The matrix \mathbf{A} need not be definite.

The multifrontal method is used. It is a direct method based on a sparse variant of Gaussian elimination and is discussed further by Duff and Reid, ACM Trans. Math. Software **9** (1983), 302-325. A detailed discussion on the MA57 strategy and performance is given by Duff, ACM Trans. Math. Software **30** (2004), 118-144. Relevant work on pivoting and scaling strategies is given by Duff and Pralet, SIAM Journal Matrix Analysis and Applications **27** (2005), 313-340. More recent work on static pivoting is given in the Technical Report RAL-TR-2005-007 by Duff and Pralet. This will be published in the SIAM Journal on Matrix Analysis and Applications and can be obtained from the web site

<http://www.numerical.rl.ac.uk/reports/reports.html>

The MA57 package has a range of options including several sparsity orderings, multiple right-hand sides, partial solutions, error analysis, scaling, a matrix modification facility, and a stop and restart facility. Although the default settings should work well in general, there are several parameters available to enable the user to tune the code for his or her problem class or computer architecture. A recent modification (in Version 3.2.0) includes an option to increase the efficiency for matrices that are severely rank deficient.

ATTRIBUTES — **Version:** 3.7.0. **Types:** Real (single, double) **Calls:** FD15, MC64, MC71, MC47, BLAS routines _GEMM, _TPSV, and _GEMV, and (optionally) METIS_NODEND from the MeTiS package. **Remark:** Supersedes MA27. **Original date:** September 2000, Version 3.7.0: May 2011. **Origin:** I. S. Duff, Rutherford Appleton Laboratory.

2 HOW TO USE THE PACKAGE

2.1 Argument lists and calling sequences

There are six entries:

- MA57I/ID sets default values for the components of the arrays that hold control parameters. Normally the user will call MA57I/ID prior to any call to MA57A/AD. If non-default values for any of the control parameters are required, they should be set immediately after the call to MA57I/ID.
- MA57A/AD accepts the pattern of \mathbf{A} and chooses pivots for Gaussian elimination using a selection criterion to preserve sparsity. It subsequently constructs subsidiary information for the actual factorization by MA57B/BD. The user may provide the pivot sequence; in which case only the necessary information for MA57B/BD will be generated.
- MA57B/BD factorizes a matrix \mathbf{A} using the information from a previous call to MA57A/AD. The actual pivot sequence used may differ from that of MA57A/AD if \mathbf{A} is not definite.
- MA57C/CD uses the factors generated by MA57B/BD to solve a system of equations $\mathbf{Ax}=\mathbf{b}$ ($\mathbf{AX}=\mathbf{B}$).
- MA57D/DD uses the factors generated by MA57B/BD to solve a system of equations $\mathbf{Ax}=\mathbf{b}$ using iterative refinement and (optionally) returning estimates of the error.
- MA57E/ED maps the data into new larger size arrays following a failure of MA57B/BD through insufficient storage so that the numerical factorization can be continued using these larger arrays.

A call to MA57C/CD or MA57D/DD must be preceded by a call to MA57B/BD which in turn must be preceded by a call to MA57A/AD. Since the information passed from one subroutine to the next is not corrupted by the second, several calls to MA57B/BD for matrices with the same sparsity pattern but different values may follow a single call to

MA57A/AD, and similarly MA57C/CD or MA57D/DD can be used repeatedly to solve for different right-hand sides \mathbf{b} (\mathbf{B}). Note that it would be possible to use MA57A/AD on several matrices before calling MA57B/BD. When we state that parameters “must be unchanged since the call to *subroutine*”, we mean a successful call of the routine on the same matrix.

In MA57I, MA57A, MA57B, MA57C, MA57D, and MA57E, all REALs are default reals; in MA57ID, MA57AD, MA57BD, MA57CD, MA57DD, and MA57ED, all REALs are default double precision.

2.1.1 To set default values of controlling parameters

The single precision version

```
CALL MA57I(CNTL, ICNTL)
```

The double precision version

```
CALL MA57ID(CNTL, ICNTL)
```

CNTL is a REAL array of length 5 that need not be set by the user. On return it contains default values. For further information see Section 2.2.

ICNTL is an INTEGER array of length 20 that need not be set by the user. On return it contains default values. For further information see Section 2.2.

2.1.2 To perform symbolic manipulations

The single precision version

```
CALL MA57A(N, NE, IRN, JCN, LKEEP, KEEP, IWORK, ICNTL, INFO, RINFO)
```

The double precision version

```
CALL MA57AD(N, NE, IRN, JCN, LKEEP, KEEP, IWORK, ICNTL, INFO, RINFO)
```

N is an INTEGER variable that must be set by the user to the order n of the matrix \mathbf{A} . It is not altered by the subroutine. **Restriction:** $N \geq 1$.

NE is an INTEGER variable that must be set by the user to the number of entries input in IRN and JCN. It is not altered by the subroutine. **Restriction:** $NE \geq 0$.

IRN and JCN are INTEGER arrays of length NE. The user must set them so that each diagonal entry a_{ii} is represented by $IRN(k)=i$ and $JCN(k)=i$ and each pair of off-diagonal entries a_{ij} and a_{ji} is represented by $IRN(k)=i$ and $JCN(k)=j$ or by $IRN(k)=j$ and $JCN(k)=i$. If a diagonal entry is zero, it can be omitted. Multiple entries are permitted. If $IRN(k)$ or $JCN(k)$ are less than 1 or greater than N, the entry is ignored. These arrays are not altered by any of the calls to the MA57 package. They must be preserved by the user between this call and a call to MA57D/DD for the same matrix.

LKEEP is an INTEGER variable that must be set by the user to the length of array KEEP. It might be more efficient to allocate more than the minimum required, say about N to $2*N$ more space. **Restriction:** $LKEEP \geq 5*N+NE+MAX(N, NE)+42$.

KEEP is an INTEGER array of length LKEEP. It need not be set by the user and must be preserved between a call to MA57A/AD and subsequent calls to MA57B/BD. If the user wishes to input the pivot sequence, the position of variable i in the pivot order should be placed in $KEEP(I)$, $I = 1, 2, \dots, N$ and $ICNTL(6)$ should be set to 1. The subroutine may replace the given order by another that gives the same fill-in pattern and virtually identical numerical results.

IWORK is an INTEGER array of length $5*N$ that is used as workspace.

ICNTL is an INTEGER array of length 20 that contains control parameters and must be set by the user. Default values for the components may be set by a call to MA57I/ID. Details of the control parameters are given in Section 2.2.

INFO is an INTEGER array of length 40 that need not be set by the user. On return from MA57A/AD, a non-negative value for INFO(1) indicates that the subroutine has performed successfully. For nonzero values, see Section 2.3. For the meanings of the values of other components of INFO, see Section 2.2.

RINFO is a REAL array of length 20 that need not be set by the user. This array supplies information on the execution of the subroutine. For the meanings of the values of the components of RINFO, see Section 2.2.

2.1.3 To factorize a matrix (or continue factorization after restart)

The single precision version

```
CALL MA57B(N,NE,A,FACT,LFACT,IFACT,LIFACT,LKEEP,KEEP,IWORK,ICNTL,CNTL,
*          INFO,RINFO)
```

The double precision version

```
CALL MA57BD(N,NE,A,FACT,LFACT,IFACT,LIFACT,LKEEP,KEEP,IWORK,ICNTL,CNTL,
*          INFO,RINFO)
```

N is an INTEGER variable that must be set by the user to the order n of the matrix **A**. It must be unchanged since the call to MA57A/AD and is not altered by the subroutine. **Restriction:** $N \geq 1$.

NE is an INTEGER variable that must be set by the user to the number of entries in the matrix **A**. It must be unchanged since the call to MA57A/AD and is not altered by the subroutine. **Restriction:** $NE \geq 0$.

A is a REAL array of length NE that must be set by the user so that $A(k)$ holds the value of the diagonal entry or pair of off-diagonal entries whose indices were held in IRN(k) and JCN(k) on entry to MA57A/AD, for $k = 1, 2, \dots, NE$. Multiple entries are summed and any that correspond to an IRN(k) or JCN(k) value that was out of range are ignored. The array A is not altered by the subroutine and must be preserved by the user between this call and a call to MA57D/DD for the same matrix.

FACT is a REAL array of length LFACT. It need not be set by the user and, on exit, will hold the entries of the factors of the matrix **A**. These entries in FACT must be preserved by the user between calls to this subroutine and subsequent calls to MA57C/CD or MA57D/DD. If MA57B/BD is called with ICNTL(7) set to 4, the entries LFACT-N+1 to LFACT of FACT will hold the changes made to the pivot entries by the matrix modification. If MA57B/BD is being called after a call to MA57E/ED, then FACT must be the same array as the array NEWFAC as output from MA57E/ED.

LFACT is an INTEGER variable that must be set by the user to the length of array FACT. It must be at least as great as INFO(9) as output from MA57A/AD (see Section 2.2). It is advisable to allow a slightly greater value because numerical pivoting might increase storage requirements. In extreme cases, for example when there are many zeros on the diagonal, a value of more than twice INFO(9) might be needed. It is not altered by the subroutine.

IFACT is an INTEGER array of length LIFACT. It need not be set by the user and, on exit, holds integer indexing information on the matrix factors. It must be preserved by the user between calls to this subroutine and subsequent calls to MA57C/CD or MA57D/DD. If MA57B/BD is being called after a call to MA57E/ED, then IFACT must be the same array as the array NEWIFC as output from MA57E/ED.

LIFACT is an INTEGER variable that must be set by the user to the length of array IFACT. It must be at least as great as INFO(10) as output from MA57A/AD (see Section 2.2). A slightly greater value is recommended because numerical pivoting may increase storage requirements. It is not altered by the subroutine.

LKEEP is an INTEGER variable that must be set by the user to the length of array KEEP. **Restriction:** $LKEEP \geq 5*N+NE+MAX(N,NE)+42$.

KEEP is an INTEGER array of length LKEEP that must be unchanged since the call to MA57A/AD. It is not altered by MA57B/BD.

IWORK is an INTEGER array of length N. It is used as workspace by the subroutine.

ICNTL is an INTEGER array of length 20 that contains control parameters and must be set by the user. Default values for the components may be set by a call to MA57I/ID. Details of the control parameters are given in Section 2.2.

CNTL is a REAL array of length 5 that contains control parameters and must be set by the user. Default values for the components may be set by a call to MA57I/ID. Details of the control parameters are given in Section 2.2.

INFO is an INTEGER array of length 40 that need not be set by the user. On return from MA57B/BD, a non-negative value for INFO(1) indicates that the subroutine has performed successfully. For nonzero values, see Section 2.3. For the meanings of the values of other components of INFO, see Section 2.2.

RINFO is a REAL array of length 20 that need not be set by the user. This array supplies information on the execution of the subroutine. For the meanings of the values of the components of RINFO, see Section 2.2.

2.1.4 To solve equations using the factors from MA57B/BD without iterative refinement

The single precision version

```
CALL MA57C(JOB,N,FACT,LFACT,IFACT,LIFACT,NRHS,RHS,LRHS,WORK,LWORK,IWORK,
*          ICNTL,INFO)
```

The double precision version

```
CALL MA57CD(JOB,N,FACT,LFACT,IFACT,LIFACT,NRHS,RHS,LRHS,WORK,LWORK,IWORK,
*           ICNTL,INFO)
```

JOB is an INTEGER variable that must be set by the user to determine the system of equations being solved. If JOB is less than or equal to 1, then the system(s)

$$\mathbf{Ax}=\mathbf{b} \text{ or } \mathbf{AX}=\mathbf{B}$$

is solved.

For JOB equal to 2, 3, or 4, a partial solution is obtained. Specifically, if the factorization is expressed as

$$\mathbf{PSASP}^T = \mathbf{LDL}^T,$$

where \mathbf{S} is the diagonal scaling matrix from the scaling when ICNTL(15) equals 1 and \mathbf{P} the permutation matrix, then the systems solved are:

$$\text{JOB} = 2 \quad \mathbf{LSx} = \mathbf{PSb} \quad \text{or} \quad \mathbf{LSX} = \mathbf{PSB}$$

$$\text{JOB} = 3 \quad \mathbf{DS}^{-1}\mathbf{x} = \mathbf{Sb} \quad \text{or} \quad \mathbf{DS}^{-1}\mathbf{X} = \mathbf{SB}$$

$$\text{JOB} \geq 4 \quad \mathbf{L}^T\mathbf{P}^{-1}\mathbf{S}^{-1}\mathbf{x} = \mathbf{S}^{-1}\mathbf{b} \quad \text{or} \quad \mathbf{L}^T\mathbf{P}^{-1}\mathbf{S}^{-1}\mathbf{X} = \mathbf{S}^{-1}\mathbf{B}$$

where the right-hand side(s) \mathbf{b} or \mathbf{B} is input in array RHS.

N is an INTEGER variable that must be set by the user to the order n of the matrix \mathbf{A} . It must be unchanged since the call to MA57B/BD and is not altered by the subroutine.

FACT is a REAL array of length LFACT that must be unchanged since the call to MA57B/BD. It is not altered by the subroutine.

LFACT is an INTEGER variable that must be set by the user to the length of array FACT. It is not altered by the subroutine.

IFACT is an INTEGER array of length LIFACT that must be unchanged since the call to MA57B/BD. It is not altered by the subroutine.

LIFACT is an INTEGER variable that must be set by the user to the length of array IFACT. It is not altered by the subroutine.

NRHS is an INTEGER variable that must be set by the user to the number of right-hand sides, s , being solved. It is not

altered by the subroutine. **Restriction:** $\text{NRHS} > 0$.

RHS is a REAL array of dimensions LRHS by NRHS . On entry, $\text{RHS}(I, J)$ must hold the I th component of the J th right-hand side of the equations being solved. On exit it will hold the corresponding entry of the J th solution vector.

LRHS is an INTEGER variable that must be set by the user to the leading dimension of array **RHS**. It is not altered by the subroutine. **Restriction:** $\text{LRHS} \geq \text{N}$.

WORK is a REAL array of length **LWORK** that is used as workspace by MA57C/CD.

LWORK is an INTEGER variable that must be set by the user to the length of array **WORK**. It must be at least as large as $\text{N} * \text{NRHS}$. It is not altered by MA57C/CD. **Restriction:** $\text{LWORK} \geq \text{N} * \text{NRHS}$.

IWORK is an INTEGER array of length **N** that is used as workspace by MA57C/CD.

ICNTL is an INTEGER array of length 20 that contains control parameters and must be set by the user. Default values for the components may be set by a call to MA57I/ID. Details of the control parameters are given in Section 2.2.

INFO is an INTEGER array of length 40 that need not be set by the user. On return from MA57E/ED, a non-negative value for **INFO**(1) indicates that the subroutine has performed successfully. For nonzero values, see Section 2.3.

2.1.5 To solve equations using iterative refinement

The single precision version

```
CALL MA57D(JOB,N,NE,A,IRN,JCN,FACT,LFACT,IFACT,LIFACT,RHS,X,RESID,WORK,IWORK,
           ICNTL,CNTL,INFO,RINFO)
```

The double precision version

```
CALL MA57DD(JOB,N,NE,A,IRN,JCN,FACT,LFACT,IFACT,LIFACT,RHS,X,RESID,WORK,IWORK,
            ICNTL,CNTL,INFO,RINFO)
```

JOB must be set by the user to determine what action is desired by the user. Values of **JOB** and their effect are:
If **ICNTL**(9) > 1, then **JOB**=0 if no estimate of solution in **X**, and **JOB**=2 if estimate of solution in **X**.

If **ICNTL**(9)=1, then:

0: Solve $\mathbf{Ax}=\mathbf{b}$, calculate residual $\mathbf{r}=\mathbf{b}-\mathbf{Ax}$ and exit.

(Note that MA57C/CD should be used if solution without residual is required)

1: Solve $\mathbf{Ax}=\mathbf{b}$, calculate residual $\mathbf{r}=\mathbf{b}-\mathbf{Ax}$, solve $\mathbf{A}(\mathbf{dx})=\mathbf{r}$, update solution and exit (that is, we perform one step of iterative refinement).

If **JOB** > 1, an estimate of the solution must be input in **X**.

2: Calculate residual $\mathbf{r}=\mathbf{b}-\mathbf{Ax}$, solve $\mathbf{A}(\mathbf{dx})=\mathbf{r}$, update solution and exit.

If **JOB** > 2, the residual for this estimate must also be input.

3: Solve $\mathbf{A}(\mathbf{dx})=\mathbf{r}$, update solution and exit.

4: Solve $\mathbf{A}(\mathbf{dx})=\mathbf{r}$, update solution, calculate residual for new solution and exit.

Restriction: $0 \leq \text{JOB} \leq 4$ and $\text{JOB} = 0$ or 2 if **ICNTL**(9) > 1.

N is an INTEGER variable that must be set by the user to the order n of the matrix **A**. It must be unchanged since the call to MA57B/BD and is not altered by the subroutine. **Restriction:** $\text{N} \geq 1$.

NE is an INTEGER variable that must be set by the user to the number of entries in the matrix **A**. It must be unchanged since the call to MA57A/AD and is not altered by the subroutine. **Restriction:** $\text{NE} \geq 0$.

A is a REAL array of length **NE** that must be unchanged since the call to MA57B/BD. It is not altered by the

subroutine.

IRN and JCN are INTEGER arrays of length NE that must be unchanged since the call to MA57A/AD. They are not altered by the subroutine.

FACT is a REAL array of length LFACT that must be unchanged since the call to MA57B/BD. It is not altered by the subroutine.

LFACT is an INTEGER variable that must be set by the user to the length of array FACT. It is not altered by the subroutine.

IFACT is an INTEGER array of length LIFACT that must be unchanged since the call to MA57B/BD. It is not altered by the subroutine.

LIFACT is an INTEGER variable that must be set by the user to the length of array IFACT. It is not altered by the subroutine.

RHS is a REAL array of length N that must be set by the user to the right-hand side for the equation being solved. It is not altered by the subroutine.

X is a REAL array of length N. If $JOB \geq 2$, it must be set by the user to an estimated solution. Otherwise, it need not be set by the user. On exit, it holds the improved solution vector.

RESID is a REAL array of length N. If $JOB > 2$, it must be set on entry to the value of the residual for the current solution estimate held in X. Otherwise, it need not be set by the user on entry. If $JOB = 0$ or 4 or if $ICNTL(9) > 1$, on exit it will hold the residual vector for the equations being solved. If $1 \leq JOB \leq 3$, then RESID will hold on exit the last correction vector added to the solution X.

WORK is a REAL work array. If $ICNTL(9) = 1$, it must be of length at least N. If $ICNTL(9) > 1$, it must be of length at least $3 \times N$. If $ICNTL(9) > 1$ and $ICNTL(10) > 0$, then WORK must be of length at least $4 \times N$.

IWORK is an INTEGER of length N that is used as a work array if $ICNTL(1) > 9$. It is not accessed if $ICNTL(9) = 1$.

ICNTL is an INTEGER array of length 20 that contains control parameters and must be set by the user. Default values for the components may be set by a call to MA57I/ID. Details of the control parameters are given in Section 2.2.

CNTL is a REAL array of length 5 that contains control parameters and must be set by the user. Default values for the components may be set by a call to MA57I/ID. Details of the control parameters are given in Section 2.2.

INFO is an INTEGER array of length 40 that need not be set by the user. On return from MA57D/DD, a non-negative value for $INFO(1)$ indicates that the subroutine has performed successfully. For nonzero values, see Section 2.3.

RINFO is a REAL array of length 20 that need not be set by the user. This array supplies information on the execution of MA57D/DD. For the meanings of the values of the components of RINFO, see Section 2.2.

2.1.6 The use of the ordering from MeTiS

The MA57 package uses the MeTiS graph partitioning library available from the University of Minnesota website (see parameter $ICNTL(6)$ in Section 2.2). If MeTiS is not available then the user must compile with the supplied replacement routine METIS_NODEND.

2.1.7 To copy data into a larger array.

If the arrays FACT or IFACT are not of sufficient length for MA57B/BD to complete the factorization, the contents of these arrays can be moved to larger arrays (NEWFAC and NEWIFC, respectively) by calling MA57E/ED. MA57B/BD can then be recalled using these arrays in place of FACT and IFACT, and the factorization will start from where it finished during the previous factorization attempt without the need to re-perform the computations up to that point.

The single precision version

```
CALL MA57E(N,IC,KEEP,FACT,LFACT,NEWFAC,LNEW,IFACT,LIFACT,NEWIFC,LNEW,INFO)
```

The double precision version

```
CALL MA57ED(N,IC,KEEP,FACT,LFACT,NEWFAC,LNEW,IFACT,LIFACT,NEWIFC,LNEW,INFO)
```

N is an INTEGER variable that must be set by the user to the order n of the matrix **A**. It must be unchanged since the failed call to MA57B/BD and is not altered by the subroutine.

IC is an INTEGER variable that must be set by the user to 0 if the real array is being copied, and to ≥ 1 if the integer array is being copied. Values less than 0 are treated as 0.

KEEP is an INTEGER array of length $5*N+NE+MAX(N,NE)+42$. This array should be unchanged since the failed call to MA57B/BD and should be preserved by the user between this call and a subsequent call to MA57B/BD.

FACT is a REAL array of length **LFACT** that must be unchanged since the failed call to MA57B/BD. It is unchanged by MA57E/ED.

LFACT is an INTEGER variable that must be set to the length of array **FACT**. It is unchanged by MA57E/ED.

NEWFAC is a REAL array of length **LNEW** that need not be set by the user and is only accessed if $IC < 1$. On a successful exit from MA57E/ED, **NEWFAC** will hold the factors and other data generated by MA57B/BD up to the point at which it failed. MA57B/BD can now be called using this array in place of **FACT**.

LNEW is an INTEGER variable that must be set to the length of array **NEWFAC**. It is unchanged by MA57E/ED.

Restriction: $LNEW > LFACT$.

IFACT is an INTEGER array of length **LIFACT** that must be unchanged since the failed call to MA57B/BD. It is unchanged by MA57E/ED.

LIFACT is an INTEGER variable that must be set to the length of array **IFACT**. It is unchanged by MA57E/ED.

NEWIFC is an INTEGER array of length **LINEW** that need not be set by the user and is only accessed if $IC \geq 1$. On a successful exit from MA57E/ED, **NEWIFC** will hold the integer information generated by MA57B/BD up to the point at which it failed. MA57B/BD can now be called using this array in place of **IFACT**.

LINEW is an INTEGER variable that must be set to the length of array **NEWIFC**. It is unchanged by MA57E/ED.

Restriction: $LINEW > LIFACT$.

INFO is an INTEGER array of length 40 that need not be set by the user. On return from MA57E/ED, a non-negative value for **INFO**(1) indicates that the subroutine has performed successfully. For nonzero values, see Section 2.3.

2.2 Arrays for control and information

The components of the arrays **CNTL** and **ICNTL** control the action of MA57. Default values are set by MA57I/ID. Only the first five components of **CNTL** and the first 16 components of **ICNTL** are presently used by MA57. The remaining components are set to zero by MA57I/ID and are not accessed in MA57. The components of the arrays **RINFO** and **INFO** provide information on the action of MA57.

CNTL(1) has default value 0.01 and is used for threshold pivoting. Values less than 0.0 will be treated as 0.0 and values greater than 0.5 as 0.5. Values near 0.0 may perhaps give faster factorization times and less entries in the factors but may result in a less stable factorization. This parameter is only accessed if **ICNTL**(7) is equal to 1.

CNTL(2) has default value 10^{-20} . MA57B/BD will treat any pivot whose modulus is less than **CNTL**(2) as zero. If **ICNTL**(16) = 1, then blocks of entries less than **CNTL**(2) can be discarded during the factorization and the corresponding pivots are placed at the end of the ordering. In this case, a normal value for **CNTL**(2) could be 10^{-12} .

CNTL(3) has default value 0.5. It is used by MA57D/DD to monitor the convergence of the iterative refinement. If the

norm of the scaled residuals does not decrease by a factor of at least $CNTL(3)$, convergence is deemed to be too slow and MA57/DD terminates with $INFO(1)$ set to -8.

$CNTL(4)$ has default value 0.0. It is used by MA57B/BD to control the static pivoting (see Section 4). If $CNTL(4) > 0.0$, then small pivots may be replaced by entries of value $CNTL(4)$ so that the factorization could be inaccurate. If $CNTL(5)$ is also greater than zero, then $CNTL(4)$ is treated as zero (that is uneliminated variables are delayed) until $CNTL(5) * N$ fully summed variables have been delayed. If static pivots are used, it is recommended that iterative refinement is used when computing the solution (see Section 2.1.5).

$CNTL(5)$ has default value 0.0. Static pivoting is invoked if $CNTL(4)$ is greater than zero and the accumulated number of delayed pivots exceeds $CNTL(5) * N$.

$ICNTL(1)$ has default value 6. It is the output stream for error messages. If it is negative, these messages will be suppressed.

$ICNTL(2)$ has default value 6. It is the output stream for warning messages. If it is negative, these messages are suppressed.

$ICNTL(3)$ has default value 6. It is the output stream for monitoring printing. If it is negative, this printing is suppressed.

$ICNTL(4)$ has default value -1. It is the output stream for the printing of statistics. If it is negative, the statistics are not printed. This control parameter is not operative in the present version.

$ICNTL(5)$ is used by MA57 to control printing of error, warning, and monitoring messages. It has default value 2. Possible values are:

- < 1 No messages output.
- 1 Only error messages printed.
- 2 Errors and warnings printed.
- 3 Errors and warnings and terse monitoring (only first ten entries of arrays printed).
- ≥ 4 Errors and warnings and all information on input and output parameters printed.

$ICNTL(6)$ has default value 5 and must be set by the user to 1 if the pivot order in KEEP is to be used by MA57A/AD. For any other value of $ICNTL(6)$, a suitable pivot order will be chosen automatically. Our experience is that the choice of ordering can be quite crucial to performance so our default invokes logic that chooses automatically between the MeTiS ordering and a version of approximate minimum degree. In general, the MeTiS ordering gives the best results for large matrices and so we recommend that the user should ensure that this library is available. It can be obtained from:

<http://www-users.cs.umn.edu/~karypis/metis/metis/download.html>

Currently the following choices are available:

- 0 AMD ordering using MC47 (without the detection of dense rows).
- 2 AMD ordering using MC47.
- 3 Minimum degree ordering as generated by the MA27 code.
- 4 MeTiS ordering is used. Note that the user needs to supply the MeTiS library. If it is not supplied and this option is requested, the routine will return immediately with $INFO(1)$ set to -18.
- 5 An automatic choice is made between the MeTiS ordering and MC47. If the MeTiS library is not supplied and this option is requested, the routine uses MC47.
- ≥ 6 Currently is equivalent to setting $ICNTL(6)=5$, but may be used for alternative orderings in later releases of MA57.

ICNTL(7) has default value 1 and is used to control numerical pivoting. It is used in MA57A/AD in the automatic selection of an ordering (if ICNTL(6) = 5) and by MA57B/BD to determine the pivoting strategy. Values out of range cause an error return with INFO(1) equal to -10. Possible values are:

- 1 Numerical pivoting is performed using the threshold value in CNTL(1).
- 2 No pivoting will be performed and an error exit will occur immediately a sign change or a zero is detected among the pivots. This is suitable for cases when **A** is thought to be definite and is likely to decrease the factorization time while still providing a stable decomposition.
- 3 No pivoting will be performed and an error exit will occur if a pivot of modulus less than or equal to CNTL(2) is detected. This is likely to decrease the factorization time, but may be unstable if there is a sign change among the pivots.
- 4 No pivoting will be performed but the matrix will be altered so that all pivots are of the same sign. The alterations to the pivots will be returned by MA57B/BD in entries LFACT-N to LFACT-1 of array FACT, if scaling is switched off (ICNTL(15) ≠ 1), and from positions LFACT-2*N to LFACT-N-1, otherwise. ICNTL(7) must also be set to 4 on entry to MA57A/AD so that the analysis can predict the space required when using this option.

ICNTL(8) has default value 0. If MA57B/BD is called with ICNTL(8) = 0 and runs out of space, then the factorization will discard factors and try to continue the factorization to determine the amount of space needed for a successful factorization. In this case, a factorization will not have been produced. If ICNTL(8) ≠ 0, and the factorization stops because of lack of space, the user must allocate a larger array to replace FACT or IFACT and transfer data from FACT or IFACT by using MA57E/ED. MA57B/BD can then be called again using these larger arrays, and the factorization will continue from the point at which it failed.

ICNTL(9) has default value 10. It corresponds to the maximum permitted number of steps of iterative refinement and is only operational on a call to MA57D/DD. If ICNTL(9) = 1, one step of iterative refinement is performed but no error analysis is done. If ICNTL(9) > 1, iterative refinement is performed and an estimate of the backward error is returned in RINFO(6) and RINFO(7), and the infinity norm of the input matrix, the infinity norm of the computed solution, and the norm of the scaled residual in RINFO(8) to RINFO(10), respectively. If ICNTL(10) is also positive, an estimate of the condition number and error are also returned. If ICNTL(9) < 0, then MA57D/DD will return with INFO(1) = -13.

ICNTL(10) has default value 0. On a call to MA57D/DD, if ICNTL(9) > 1, a positive value for ICNTL(10) will return estimates of the condition number of the system in RINFO(11) and RINFO(12) and an estimate of the error in RINFO(13).

ICNTL(11) has default value 16 and is the block size used by the Level 3 BLAS in MA57B/BD.

ICNTL(12) has default value 16. Two nodes of the assembly tree are merged only if both involve less than ICNTL(12) eliminations.

ICNTL(13) has default value 10 and is a control on the number of columns in a block pivot above which the Level 2 and Level 3 BLAS are used in MA57C/CD and MA57D/DD.

ICNTL(14) has default value 100 and is the percentage density of a row that is regarded as dense during the pivot selection strategy in MA57A/AD that uses the MA27 minimum degree ordering (ICNTL(6) equal to 3). This is useful since the ordering strategy can be slow on matrices with dense rows if this option is not used.

ICNTL(15) has default value 1 and indicates that the matrix will be scaled using a symmetrized version of the HSL code MC64. Setting ICNTL(15) to any other value will suppress this option. Scaling factors will be returned by MA57B/BD in entries LFACT-N to LFACT-1 of array FACT. Although the scaling factors are computed in the call to MA57B/BD, this parameter should also be set for the call to MA57A/AD so that the correct amount of storage can be predicted.

ICNTL(16) has default value 0. If ICNTL(16) is set to 1, then when small entries (defined by CNTL(2)) are

detected, they are removed and the corresponding pivots placed at the end of the factorization. This can be particularly efficient if the matrix is highly rank deficient.

ICNTL(17) to ICNTL(20) are set to zero by MA57I/ID but are not currently used by MA57.

INFO(1) is zero if the routine is successful, is negative if an error occurred, and is positive for a warning (see Section 2.3).

INFO(2) holds additional information concerning the error or warning (see Section 2.3).

Statistics produced after analysis phase (MA57A/AD). The forecasts will be correct if the matrix is definite.

INFO(3) Number of indices out of range.

INFO(4) Number of off-diagonal duplicates.

INFO(5) Forecast number of reals in the factors.

INFO(6) Forecast number of integers in the factors.

INFO(7) Forecast maximum frontal matrix size.

INFO(8) Number of nodes in the assembly tree.

INFO(9) and INFO(10) are lengths of arrays FACT and IFACT required for a successful completion of MA57B/BD without the need for data compression provided no numerical pivoting is performed. The actual amount required may be higher because of numerical pivoting.

INFO(11) and INFO(12) are lengths of arrays FACT and IFACT required for a successful completion of MA57B/BD allowing data compression, again provided no numerical pivoting is performed. Numerical pivoting may cause a higher value to be required.

INFO(13) Number of data compresses. If this is high (say > 10), the performance of MA57A/AD may be improved by increasing the length of array KEEP.

INFO(36) Ordering actually used (refer to numbering scheme for ICNTL(6)).

Statistics produced after factorization (MA57B/BD)

INFO(14) Number of entries in factors.

INFO(15) Storage for real data for factors.

INFO(16) Storage for integer data for factors.

INFO(17) and INFO(18). The minimum length of arrays FACT and IFACT, respectively, that would be sufficient for a successful factorization of this matrix. In the event of a failure with INFO(1) equal to -3 (-4), INFO(17)(INFO(18)) is set to the estimated real (integer) space required.

INFO(19) and INFO(20). The minimum length of arrays FACT and IFACT, respectively, that would be sufficient for a successful factorization of this matrix without any compresses on the data structures.

INFO(21) Order of largest frontal matrix.

INFO(22) Number of 2x2 pivots used in the factorization.

INFO(23) Total number of fully-summed variables that were passed to the father node because of pivoting considerations (called delayed pivots. If a variable is passed further up the tree, it will be counted again. This will be zero for a definite system, if ICNTL(7) > 1 or if static pivoting is used (CNTL(4) > 0.0 and CNTL(5) equal to zero).

INFO(24) Number of negative eigenvalues in the factorization of **A**.

INFO(25) Rank of the factorization of **A**. Note that, if static pivoting is invoked (see CNTL(4) in Section 2.2), then

the rank of the factorization will always be n .

INFO(26) Number of sign changes of pivot when ICNTL(7) = 3.

INFO(27) Pivot step at which matrix modification commences when ICNTL(7) = 4. INFO(27) is set to 0 if no modification performed.

INFO(28) and INFO(29) Number of data compresses on real and integer data structures respectively. If either of these is high (say > 10), then the speed of the factorization may be increased by allocating more space to the arrays IFACT or FACT as appropriate.

INFO(31) Number of block pivots in factors.

INFO(32) Number of zeros in the triangle of the factors.

INFO(33) Number of zeros in the rectangle of the factors.

INFO(34) Number of zero columns in rectangle of the factors.

INFO(35) Number of pivots chosen in static pivoting phase.

Statistics produced after solution using MA57D/DD with ICNTL(9) > 1

INFO(30) Number of solution steps in iterative refinement.

Not currently used by MA57.

INFO(37) to INFO(40).

Statistics produced after analysis phase

RINFO(1) The forecast number of floating-point additions for the assembly processes.

RINFO(2) The forecast number of floating-point operations to perform the elimination operations (counting a multiply-add pair as two operations).

Statistics produced after factorization

RINFO(3) Number of floating-point additions for the assembly process.

RINFO(4) Number of floating-point operations for the elimination.

RINFO(5) Additional number of floating-point operations when using _GEMM for the elimination.

RINFO(14) Maximum change to pivot when ICNTL(7) = 4.

RINFO(15) Smallest pivot in modified matrix when ICNTL(7) = 4.

RINFO(16) Minimum value of scaling factor when ICNTL(15) = 1.

RINFO(17) Maximum value of scaling factor when ICNTL(15) = 1.

RINFO(18) Maximum modulus of matrix entry (held in FACT(LFACT)).

Statistics produced after solve using MA57D/DD with ICNTL(9) > 1.

RINFO(6) and RINFO(7) are used to hold information on the backward error. See Section 2.4.

RINFO(8) Infinity norm of the input matrix, where

$$\|\mathbf{A}\|_{\infty} = \max_{1 \leq i \leq n} \sum_{j=1}^n |a_{ij}|.$$

RINFO(9) Infinity norm of the computed solution, where

$$\|\mathbf{x}\|_{\infty} = \max_{1 \leq i \leq n} |x_i|.$$

RINFO(10) Norm of scaled residuals

$$\frac{\max_{1 \leq i \leq n} \left| \sum_{j=1}^n a_{ij} x_j - b_i \right|}{\|A\|_{\infty} \|x\|_{\infty}}.$$

The following are output after a call to MA57D/DD with ICNTL(9) > 1 and ICNTL(10) > 0.

RINFO(11) and RINFO(12) are the condition number for the system as used by Arioli, Demmel, and Duff.

RINFO(13) is an upper bound for the infinity norm of the error in the solution.

Not currently used by MA57.

RINFO(19) to RINFO(20).

2.3 Error diagnostics

A successful return is indicated by a value of INFO(1) equal to zero. Possible nonzero values for INFO(1) are given below. In each case an identifying message is output on unit ICNTL(1) (errors) or ICNTL(2) (warnings).

A negative value of INFO(1) is associated with an error message that will be output on unit ICNTL(2).

- 1 Value of N is out of range on a call to MA57A/AD, MA57B/BD, MA57C/CD, or MA57D/DD. Value given is held in INFO(2).
- 2 Value of NE is out of range on a call to MA57A/AD, MA57B/BD, or MA57D/DD. Value given is held in INFO(2).
- 3 Failure due to insufficient REAL space on a call to MA57B/BD. INFO(17) is set to a value that may suffice for LFACT. INFO(2) is set to value of LFACT. The user can allocate a larger array and copy the contents of FACT into it using MA57E/ED, and recall MA57B/BD.
- 4 Failure due to insufficient INTEGER space on a call to MA57B/BD. INFO(18) is set to a value that may suffice for LIFACT. INFO(2) is set to value of LIFACT. The user can allocate a larger array and copy the contents of IFACT into it using MA57E/ED, and recall MA57B/BD.
- 5 A pivot with magnitude less than or equal to CNTL(2) was found at pivot step INFO(2) when calling MA57B/BD with ICNTL(7) = 2 or 3, or the correction obtained when using matrix modification does not give a pivot greater than CNTL(2) when ICNTL(7) = 4.
- 6 A change of sign of pivots has been detected when ICNTL(7) = 2. INFO(2) is set to the pivot step at which the change was detected on a call to MA57B/BD.
- 7 Either LNEW < LFACT or LINEW < LIFACT on a call to MA57E/ED. INFO(2) is set to LNEW or LINEW as appropriate.
- 8 Iterative refinement fails to converge in specified number of iterations on a call to MA57D/DD.
- 9 Error in permutation array when ICNTL(6)=1 on a call to MA57A/AD. INFO(2) holds first component at which error was detected.
- 10 Value of ICNTL(7) out of range on a call to MA57B/BD. Value given held in INFO(2).
- 11 LRHS < N on a call to MA57C/CD. INFO(2) holds value of LRHS.
- 12 Invalid value for JOB on a call to MA57D/DD. Value given held in INFO(2).
- 13 Invalid value of ICNTL(9) on a call to MA57D/DD. Value given held in INFO(2).
- 14 Failure of MC71A/AD on a call to MA57D/DD with ICNTL(10)> 0.
- 15 LKEEP less than 5*N+NE+MAX(N,NE)+42 on a call to MA57A/AD or MA57B/BD. INFO(2) holds value of LKEEP.
- 16 NRHS less than 1 on call to MA57C/CD. INFO(2) holds value of NRHS.

- 17 LWORK too small on entry to MA57C/CD. INFO(2) holds minimum length required.
- 18 An ordering using MeTiS was requested in the call to MA57A/AD but the MeTiS package was not linked in.
A positive value of INFO(1) is associated with a warning message that will be output on unit ICNTL(2).
- +1 Index (in IRN or JCN) out of range on call to MA57A/AD or MA57D/DD. Action taken by subroutine is to ignore any such entries and continue. INFO(3) is set to the number of out-of-range entries. Details of the first ten are printed on unit ICNTL(2).
- +2 Duplicate indices on call to MA57A/AD or MA57D/DD. Action taken by subroutine is to keep the duplicates and then to sum corresponding reals when MA57B/BD is called. INFO(4) is set to the number of duplicate entries. Details of the first ten are printed on unit ICNTL(2).
- +3 Both out-of-range indices and duplicates exist.
- +4 Matrix is rank deficient on exit from MA57B/BD. In this case, a decomposition will still have been produced that will enable the subsequent solution of consistent equations. INFO(25) will be set to the rank of the factorized matrix.
- +5 Pivots have different signs when factorizing a supposedly definite matrix (ICNTL(7) = 3) on call to MA57B/BD. INFO(26) is set to the number of sign changes.
- +8 During error analysis the infinity norm of the computed solution was found to be zero.
- +10 Insufficient real space to complete factorization when MA57B/BD called with ICNTL(8) $\neq 0$. User can copy real values to a longer array using MA57E/ED and recall MA57B/BD using this longer array to continue the factorization.
- +11 Insufficient integer space to complete factorization when MA57B/BD called with ICNTL(8) $\neq 0$. User can copy integer values to a longer array using MA57E/ED and recall MA57B/BD using this longer array to continue the factorization.

2.4 Error estimates

We calculate an estimate of the sparse backward error using the theory and measure developed by Arioli, Demmel, and Duff (1989). We use the notation $\bar{\mathbf{x}}$ for the computed solution and a modulus sign on a vector or matrix to indicate the vector or matrix obtained by replacing all entries by their moduli. The scaled residual

$$\frac{|\mathbf{b} - \mathbf{A}\bar{\mathbf{x}}|_i}{(|\mathbf{b}| + |\mathbf{A}||\bar{\mathbf{x}}|)_i} \quad (1)$$

is calculated for all equations except those for which the numerator is nonzero and the denominator is small. For the exceptional equations,

$$\frac{|\mathbf{b} - \mathbf{A}\bar{\mathbf{x}}|_i}{(|\mathbf{A}||\bar{\mathbf{x}}|)_i + \|\mathbf{A}_i\|_\infty \|\bar{\mathbf{x}}\|_\infty} \quad (2)$$

is used instead, where \mathbf{A}_i is row i of \mathbf{A} . The largest scaled residual (1) is returned in RINFO(6) and the largest scaled residual (2) is returned in RINFO(7). If all equations are in category (1), zero is returned in RINFO(7). The computed solution $\bar{\mathbf{x}}$ is the exact solution of the equation

$$(\mathbf{A} + \delta\mathbf{A})\mathbf{x} = (\mathbf{b} + \delta\mathbf{b})$$

where $\delta\mathbf{A}_{ij} \leq \max(\text{RINFO}(6), \text{RINFO}(7))|\mathbf{A}_{ij}|$ and $\delta\mathbf{b}_i \leq \max(\text{RINFO}(6)|\mathbf{b}_i|, \text{RINFO}(7)\|\mathbf{A}_i\|_\infty \|\bar{\mathbf{x}}\|_\infty)$. Note that $\delta\mathbf{A}$ respects the sparsity in \mathbf{A} . An upper bound for the forward error is returned in RINFO(13) that is computed as $\text{RINFO}(6) * \text{RINFO}(11) + \text{RINFO}(7) * \text{RINFO}(12)$ where RINFO(11) and RINFO(12) are condition numbers corresponding to κ_{ω_1} and κ_{ω_2} , respectively as defined in equation (15) of Arioli, Demmel, and Duff (1989).

Reference

Arioli, M. Demmel, J. W., and Duff, I. S. (1989). Solving sparse linear systems with sparse backward error. *SIAM J. Matrix Anal. and Applies.* **10**, 165-190.

3 GENERAL INFORMATION

Use of common: None.

Other routines called directly:

BLAS **routines:** STPSV/DTPSV, SGEMM/DGEMM, and SGEMV/DGEMV.

Routine from MeTiS package (optional): METIS_NODEND.

HSL routines: FD15A/AD, MC64W/WD, MC71A/AD, and MC47B/BD.

All the other subroutines called by the principal subroutines are in the MA57 package. They are called MA57F/FD, MA57G/GD, MA57H/HD, MA57J/JD, MA57K/KD, MA57L/LD, MA57M/MD, MA57N/ND, MA57O/OD, MA57P/PD, MA57Q/QD, MA57R/RD, MA57S/SD, MA57T/TD, MA57U/UD, MA57V/VD, MA57W/WD, MA57X/XD, MA57Y/YD, and MA57Z/ZD.

Input/output: Error, warning, monitoring, and statistical messages only. Error messages on unit ICNTL(1), which has default value 6. Warning and monitoring messages on unit ICNTL(2) and ICNTL(3), respectively, which have default value 6. Printing of these messages is controlled by ICNTL(5) and all messages are suppressed if ICNTL(5) is less than one or if the appropriate unit number is set to a negative value.

Restrictions: $N \geq 1$, $NE \geq 0$ on entry to MA57A/AD, MA57B/BD, MA57D/DD.

$0 \leq \text{JOB} \leq 4$ and $\text{JOB} = 0$ or 2 if ICNTL(9) > 1 on entry to MA57D/DD.

LNEW > LFACT and LINEW > LFACT on entry to MA57E/ED.

NRHS > 0 and LRHS $\geq N$ and LWORK $\geq \text{NRHS} * N$ on entry to MA57C/CD.

Changes between Version 1.0.0 and Version 2.0.0:

Version 2.0.0 incorporates several additional features to those of Version 1.0.0. We give information on these in this section and indicate where default values of control parameters have changed from the earlier version.

Built-in scaling is now available through a symmetrized version of MC64. This is controlled by the new control parameter ICNTL(15) whose default is set so that scaling is performed. The matrix is explicitly scaled internally to the package as are the right-hand side and the solution so that the user need not be concerned with this. Iterative refinement, if requested, is based on the original unscaled matrix. The minimum and maximum scale factors are returned in RINFO(16) and RINFO(17), respectively.

The user can now use an ordering from the MeTiS package. This option is invoked by setting ICNTL(6) equal to 4 or 5. Since the user must load the MeTiS library separately from HSL, a dummy routine has been included that will return the error INFO(1) = -18 if the MeTiS library is not available and ICNTL(6) is equal to 4. If the MeTiS library is not available and ICNTL(6) is equal to 5, then action is taken as if ICNTL(6) had been set to 2.

Static pivoting is now used so that the factorization can be performed using the same storage as predicted by the analysis even if the matrix is not positive definite. This is invoked by setting CNTL(4) to a nonzero value. A certain amount of additional storage can be allowed by setting CNTL(5) to a nonzero value.

Further testing on large problems has determined that a better value for ICNTL(12) is 16 so that this value has now been set as the default for this parameter.

Changes between Version 2.0.0 and Version 3.0.0:

An additional option has been added to the choice of ordering strategies. If ICNTL(6) is set to 5 (now the default), then MA57A/AD will choose automatically between using MeTiS or AMD (avoiding problems with dense rows); first based on matrix order and density and then, if necessary, by running both ordering strategies.

The actual ordering used is returned in `INFO(36)`.

The call to the AMD ordering was switched to use the new MC47 code in Version 3.1.0 and, in Version 3.2.0, an additional option has been added (`ICNTL(16)`) to remove small entries from the factorization and place corresponding pivots at the end of the factorization. This can result in far more efficient code (time and memory) if the matrix is highly rank deficient.

4 METHOD

Technical Report RAL-TR-2002-024 “MA57 – a new code for the solution of sparse symmetric indefinite systems” by Duff (published as ACM Trans. Math. Software **30** (2004), 118-144) describes the code and illustrates its performance. Some of the features that are new to Version 2.0.0 are described in the report RAL-TR-2004-020 “Strategies for scaling and pivoting for sparse symmetric indefinite problems” by Duff and Pralet (published as SIAM Journal Matrix Analysis and Applications **27** (2005), 313-340). RAL reports can be obtained from the web site

<http://www.numerical.rl.ac.uk/reports/reports.html>

A version of sparse Gaussian elimination is used.

The MA57 package supersedes MA27 for solving symmetric indefinite systems. In addition to being more efficient largely through its use of the Level 3 BLAS, it has many added features. Among these are: a fast mapping of data prior to a numerical factorization, the ability to modify pivots if the matrix is not definite, the efficient solution of several right-hand sides, a routine implementing iterative refinement, and the possibility of restarting the factorization should it run out of space.

The MA57A/AD entry (with `ICNTL(6)≠1`) chooses a pivot ordering based on either nested dissection or minimum degree using a generalized element model of the elimination to avoid storing the filled-in pattern explicitly. The elimination is represented as an assembly tree with the order of elimination determined by a depth-first search of the tree.

The MA57B/BD entry factorizes the matrix by using the assembly and elimination ordering generated by MA57A/AD. By default, the input matrix is first scaled using a symmetrized version of the HSL code MC64. At each stage in the multifrontal approach pivoting and elimination are performed on full submatrices and, when diagonal 1×1 pivots would be numerically unstable, 2×2 diagonal blocks are used. Thus MA57B/BD can be used to factor indefinite systems and will perform well on machines with caches. If during the factorization the reduced matrix has entries all less than a tolerance (default zero) the matrix is declared singular with deficiency equal to the size of the reduced matrix. Pivots corresponding to remaining rows are set to 1.

If `CNTL(4)` is set to a value greater than zero, static pivoting is invoked (RAL-TR-2005-007 “Towards a stable static pivoting strategy for the sequential and parallel solution of sparse symmetric indefinite systems” by Duff and Pralet, to appear in SIAM J Matrix Analysis and Applications). This means that if, at any stage of the multifrontal elimination, there are fully summed rows and columns from which it is not possible to choose pivots because of the threshold criterion defined by `CNTL(1)`, then we first try to get pivots using a weaker threshold by successively trying values one tenth of the previous until a threshold of $\sqrt{(\text{CNTL}(1) * \text{CNTL}(4))}$ is reached. If there are still fully summed rows and columns left, then the diagonal entries are replaced by `CNTL(4)` times the largest modulus of an entry in the scaled matrix and are used as 1×1 pivots in the factorization. If, furthermore, `CNTL(5)` is greater than zero, then `CNTL(4)` is treated as 0 (uneliminated variables are delayed) until `CNTL(5)*N` fully-summed variables have been delayed.

The MA57C/CD entry uses the factors from MA57B/BD to solve systems of equations either by loading the appropriate parts of the vectors into an array of the current front size and using full matrix code or by indirect addressing at each stage, depending on the value of `ICNTL(13)` and the size of the frontal matrix.

The MA57D/DD entry uses MA57C/CD to solve the system using either one step of iterative refinement or iterating until a measure of the backward error is reduced sufficiently according to the work of Arioli, Demmel, and Duff.

The MA57E/ED entry just copies the contents of arrays and to larger arrays and resets some pointers so that the factorization can continue using these larger arrays.

This multifrontal technique is described in detail by Duff and Reid (AERE-R.10533, 1982) and Duff and Reid, ACM Trans. Math. Software **9** (1983), 302-325.

5 EXAMPLE OF USE

We illustrate the use of the package on the solution of the single set of equations

$$\begin{pmatrix} 2 & 3 & & & \\ 3 & 0 & 4 & & 6 \\ & 4 & 1 & 5 & \\ & & 5 & 0 & \\ 6 & & & & 1 \end{pmatrix} \mathbf{x} = \begin{pmatrix} 8 \\ 45 \\ 31 \\ 15 \\ 17 \end{pmatrix}$$

We run two programs on this same matrix. The first uses the default ICNTL(6)=5, and the second requests a MeTiS ordering (ICNTL(6)=4). On this small example, the first program will choose an AMD ordering whether or not the MeTiS package is available, and we have set ICNTL(5) to 4 so that all the information passed to and from the package is displayed. The second will return with an error flag (INFO(1)=-18) if MeTiS is not available. We show the output from this program both when MeTiS is available and when it is not. Note that these examples do not illustrate all the facilities.

Program 1

```
C Simple example of use of MA57 package
  INTEGER LRHS,LFACT,LKEEP,LIFACT,LWORK
  PARAMETER (LKEEP=100,LRHS=10,LWORK=12,LFACT=1000,LIFACT=500)
  INTEGER IRN(10),JCN(10),IWORK(40),KEEP(LKEEP),IFACT(LIFACT)
  INTEGER ICNTL(40),INFO(40),N,NE,I,JOB
  DOUBLE PRECISION A(30),WORK(LWORK),FACT(LFACT),RHS(LRHS,1),
+                  CNTL(5),RINFO(20)
  EXTERNAL MA57ID,MA57AD,MA57BD,MA57CD

C Set default values for control parameters.
  CALL MA57ID(CNTL,ICNTL)

C Ask for full printing from MA57 package
  ICNTL(5) = 4

C Read matrix and right-hand side
  READ (5,*) N,NE
  READ (5,*) (IRN(I),JCN(I),A(I),I=1,NE)
  READ (5,*) (RHS(I,1),I=1,N)

C Analyse sparsity pattern
  CALL MA57AD(N,NE,IRN,JCN,LKEEP,KEEP,IWORK,ICNTL,INFO,RINFO)

C Factorize matrix
  CALL MA57BD(N,NE,A,FACT,LFACT,IFACT,LIFACT,LKEEP,KEEP,
+          IWORK,ICNTL,CNTL,INFO,RINFO)

C Solve the equations
  JOB = 1
  CALL MA57CD(JOB,N,FACT,LFACT,IFACT,LIFACT,1,RHS,LRHS,
+          WORK,LWORK,IWORK,ICNTL,INFO)
  END
```

Data


```

5 7
1 1 2.0
1 2 3.0
2 3 4.0
2 5 6.0
3 3 1.0
3 4 5.0
5 5 1.0
8. 45. 31. 15. 17.

```

Output

Entering analysis phase (MA57AD) with ...

| | | | |
|----------|---|---|----|
| N | Order of matrix | = | 5 |
| NE | Number of entries | = | 7 |
| ICNTL(1) | Stream for errors | = | 6 |
| --- | (2) Stream for warnings | = | 6 |
| --- | (3) Stream for monitoring | = | 6 |
| --- | (4) Stream for statistics | = | -1 |
| --- | (5) Level of diagnostic printing | = | 4 |
| --- | (6) Flag for input pivot order | = | 5 |
| --- | (7) Numerical pivoting control (st est) | = | 1 |
| --- | (12) Node amalgamation parameter | = | 16 |
| --- | (15) Scaling control (storage estimate) | = | 1 |

Matrix entries:

| | | | | | | | | |
|------|---|----|------|---|----|------|---|----|
| 1: (| 1 | 1) | 2: (| 1 | 2) | 3: (| 2 | 3) |
| 4: (| 2 | 5) | 5: (| 3 | 3) | 6: (| 3 | 4) |
| 7: (| 5 | 5) | | | | | | |

Leaving analysis phase (MA57AD) with ...

| | | | |
|----------|---|---|-------------|
| INFO(1) | Error indicator | = | 0 |
| | Number of entries in matrix with diagonal | = | 9 |
| INFO(3) | Number of out-of-range indices | = | 0 |
| INFO(4) | Number of off-diagonal duplicates | = | 0 |
| INFO(5) | Forecast real storage for factors | = | 12 |
| ---- | (6) Forecast integer storage for factors | = | 13 |
| ---- | (7) Forecast maximum front size | = | 4 |
| ---- | (8) Number of nodes in assembly tree | = | 2 |
| ---- | (9) Size of FACT without compress | = | 48 |
| ---- | (10) Size of IFACT without compress | = | 53 |
| ---- | (11) Size of FACT with compress | = | 48 |
| ---- | (12) Size of IFACT with compress | = | 53 |
| ---- | (13) Number of compresses | = | 0 |
| ---- | (36) Ordering strategy used by code | = | 2 |
| RINFO(1) | Forecast additions for assembly | = | 1.00000D+01 |
| RINFO(2) | Forecast ops for elimination | = | 3.40000D+01 |

Permutation array:

| | | | | |
|---|---|---|---|---|
| 5 | 4 | 3 | 2 | 1 |
|---|---|---|---|---|

Number of entries in rows of permuted matrix:

| | | | | |
|---|---|---|---|---|
| 2 | 2 | 2 | 2 | 1 |
|---|---|---|---|---|

Column indices of permuted matrix:

| | | | | |
|---|---|---|---|---|
| 5 | 2 | 4 | 3 | 3 |
| 2 | 2 | 1 | 1 | |

Tree nodes at which variables eliminated:

| | | | | |
|---|---|---|---|---|
| 2 | 2 | 2 | 2 | 1 |
|---|---|---|---|---|

Map array:

| | | | | | |
|--|---|---|---|---|---|
| | 9 | 8 | 6 | 2 | 5 |
| | 4 | 1 | | | |

Entering factorization phase (MA57BD) with ...

| | | | |
|----------|-------------------------------------|---|-------------|
| N | Order of input matrix | = | 5 |
| NE | Entries in input matrix | = | 7 |
| | Entries in input matrix (inc diags) | = | 9 |
| ICNTL(1) | Stream for errors | = | 6 |
| --- | (2) Stream for warnings | = | 6 |
| --- | (3) Stream for monitoring | = | 6 |
| --- | (4) Stream for statistics | = | -1 |
| --- | (5) Level of diagnostic printing | = | 4 |
| --- | (7) Numerical pivoting control | = | 1 |
| --- | (8) Restart or discard factors | = | 0 |
| --- | (11) Block size for Level 3 BLAS | = | 16 |
| --- | (15) Scaling control (1 on) | = | 1 |
| --- | (16) Dropping control (1 on) | = | 0 |
| LFACT | Size of real working space | = | 1000 |
| LIFACT | Size of integer working space | = | 500 |
| | Number nodes in assembly tree | = | 2 |
| CNTL(1) | Value of threshold parameter | = | 0.10000D-01 |
| CNTL(2) | Threshold for zero pivot | = | 0.10000D-19 |
| CNTL(4) | Control for value of static pivots | = | 0.00000D+00 |
| CNTL(5) | Control for number delayed pivots | = | 0.00000D+00 |

Matrix entries:

| | | | | | |
|----|-------------------|----|-------------------|----|-------------------|
| 1: | (2.00000000D+00) | 2: | (3.00000000D+00) | 3: | (4.00000000D+00) |
| 4: | (6.00000000D+00) | 5: | (1.00000000D+00) | 6: | (5.00000000D+00) |
| 7: | (1.00000000D+00) | | | | |

Permutation array:

| | | | | |
|---|---|---|---|---|
| 5 | 4 | 3 | 2 | 1 |
|---|---|---|---|---|

Number of entries in rows of permuted matrix:

| | | | | |
|---|---|---|---|---|
| 2 | 2 | 2 | 2 | 1 |
|---|---|---|---|---|

Tree nodes at which variables eliminated:

| | | | | |
|---|---|---|---|---|
| 2 | 2 | 2 | 2 | 1 |
|---|---|---|---|---|

Number of assemblies at each tree node:

| | | | | |
|---|---|--|--|--|
| 0 | 1 | | | |
|---|---|--|--|--|

Map array:

| | | | | |
|---|---|---|---|---|
| 9 | 8 | 6 | 2 | 5 |
| 4 | 1 | | | |

Column indices of permuted matrix:

| | | | | |
|---|---|---|---|---|
| 5 | 2 | 4 | 3 | 3 |
| 2 | 2 | 1 | 1 | |

Leaving factorization phase (MA57BD) with ...

| | | | |
|----------|------------------------------------|---|----|
| INFO (1) | = | 0 | |
| --- | (2) | = | 0 |
| --- | (14) Number of entries in factors | = | 12 |
| --- | (15) Real storage for factors | = | 13 |
| --- | (16) Integer storage for factors | = | 13 |
| --- | (17) Min LFACT with compresses | = | 48 |
| --- | (18) Min LIFACT with compresses | = | 53 |
| --- | (19) Min LFACT without compresses | = | 48 |
| --- | (20) Min LIFACT without compresses | = | 53 |

```

--- (21) Order of largest frontal matrix      =      4
--- (22) Number of 2x2 pivots                 =      1
--- (23) Number of delayed pivots             =      0
--- (24) Number of negative eigenvalues       =      2
--- (25) Rank of factorization                 =      5
--- (28) Number compresses on real data       =      0
--- (29) Number compresses on integer data    =      0
--- (31) Number of block pivots in factors     =      2
--- (32) Number of zeros factors triangle     =      4
--- (33) Number of zeros factors rectangle    =      0
--- (34) Number of zero cols factors rect     =      0
--- (35) Number of static pivots               =      0
RINFO(3) Operations during node assembly      = 1.000D+01
----- (4) Operations during node elimination = 3.900D+01
----- (5) Extra operations because of BLAS   = 0.000D+00
----- (18) Largest modulus of entry in matrix = 6.000D+00
RINFO(16) Minimum value of scaling factor     = 4.082D-01
----- (17) Maximum value of scaling factor   = 7.071D-01

```

For each block, the following information is provided:

1. Block number, number of rows, number of columns
2. List of indices for the pivot, each negated if part of a 2x2 pivot
3. The factorized block pivot
It has the form

$$\begin{array}{ccccc} & & & & -1 & & T \\ & & & & L & D & L \end{array}$$
 and is printed as D and L packed together.
4. List of indices for the non-pivot columns
5. The non-pivot part as rectangular block by rows

```

Block pivot      1 with      1 rows and      2 columns
      5
      6.0000D+00
      2
      -6.0000D+00

```

```

Block pivot      2 with      4 rows and      4 columns
      -4      -3      2      1
      -2.0000D-01  1.0000D+00  7.3030D-01  0.0000D+00
      -0.0000D+00  0.0000D+00  0.0000D+00  0.0000D+00
      -1.6667D-01 -1.4434D-01
      8.8889D-01

```

Scaling factors:

```
0.70711D+00 0.40825D+00 0.44721D+00 0.44721D+00 0.40825D+00
```

Entering solution phase (MA57CD) with ...

```

JOB      Control on coefficient matrix      =      1
N        Order of matrix                    =      5
ICNTL(1) Stream for errors                   =      6
--- (2)  Stream for warnings                  =      6
--- (3)  Stream for monitoring                 =      6
--- (4)  Stream for statistics                 =     -1
--- (5)  Level of diagnostic printing          =      4
LFACT    Length of array FACT                =     1000
LIFACT   Length of array IFACT               =      500
NRHS     Number of right-hand sides          =      1
LRHS     Leading dimension of RHS array      =     10

```

```
LW          Leading dimension of work array      =          12
ICNTL(13) Threshold for Level 2 and 3 BLAS      =          10
```

For each block, the following information is provided:

1. Block number, number of rows, number of columns
2. List of indices for the pivot, each negated if part of a 2x2 pivot
3. The factorized block pivot
It has the form

$$\begin{array}{ccccc} & & & & -1 & T \\ & & & & L & D & L \end{array}$$
 and is printed as D and L packed together.
4. List of indices for the non-pivot columns
5. The non-pivot part as rectangular block by rows

```
Block pivot      1 with      1 rows and      2 columns
      5
      6.0000D+00
      2
      -6.0000D+00
```

```
Block pivot      2 with      4 rows and      4 columns
      -4      -3      2      1
      -2.0000D-01  1.0000D+00  7.3030D-01  0.0000D+00
      -0.0000D+00  0.0000D+00  0.0000D+00  0.0000D+00
      -1.6667D-01 -1.4434D-01
      8.8889D-01
```

Scaling factors:

```
0.70711D+00 0.40825D+00 0.44721D+00 0.44721D+00 0.40825D+00
```

```
Right-hand side      1
      8.000D+00      4.500D+01      3.100D+01      1.500D+01      1.700D+01
```

Leaving solution phase (MA57CD) with ...

```
Solution      1
      1.000D+00      2.000D+00      3.000D+00      4.000D+00      5.000D+00
```

Program 2

```
C Simple example of use of MA57 package
      INTEGER LRHS,LFACT,LKEEP,LIFACT,LWORK
      PARAMETER (LKEEP=100,LRHS=10,LWORK=12,LFACT=1000,LIFACT=500)
      INTEGER IRN(10),JCN(10),IWORK(40),KEEP(LKEEP),IFACT(LIFACT)
      INTEGER ICNTL(40),INFO(40),N,NE,I,JOB
      DOUBLE PRECISION A(30),WORK(LWORK),FACT(LFACT),RHS(LRHS,1),
      +      CNTL(5),RINFO(20)
      EXTERNAL MA57ID,MA57AD,MA57BD,MA57CD

C Set default values for control parameters.
      CALL MA57ID(CNTL,ICNTL)

C Request MeTiS ordering
      ICNTL(6) = 4
```

```
C Read matrix and right-hand side
  READ (5,*) N,NE
  READ (5,*) (IRN(I),JCN(I),A(I),I=1,NE)
  READ (5,*) (RHS(I,1),I=1,N)

C Analyse sparsity pattern
  CALL MA57AD(N,NE,IRN,JCN,LKEEP,KEEP,IWORK,ICNTL,INFO,RINFO)

  IF (INFO(1).LT.0) GO TO 500

C Factorize matrix
  CALL MA57BD(N,NE,A,FACT,LFACT,IFACT,LIFACT,LKEEP,KEEP,
+           IWORK,ICNTL,CNTL,INFO,RINFO)

C Solve the equations
  JOB = 1
  CALL MA57CD(JOB,N,FACT,LFACT,IFACT,LIFACT,1,RHS,LRHS,
+           WORK,LWORK,IWORK,ICNTL,INFO)
  WRITE(6,100) 'Solution:',(RHS(I,1),I=1,N)
100 FORMAT(A/(1P,5D12.3))
500 STOP
END
```

Data

```
5 7
1 1 2.0
1 2 3.0
2 3 4.0
2 5 6.0
3 3 1.0
3 4 5.0
5 5 1.0
8. 45. 31. 15. 17.
```

Output (MeTiS not available)

```
**** Error return from MA57AD ****  INFO(1) =-18
MeTiS ordering requested but MeTiS not linked
```

Output (MeTiS available)

```
Solution:
  1.000D+00   2.000D+00   3.000D+00   4.000D+00   5.000D+00
```