

NAG Fortran Library Routine Document

F04QAF

Note: before using this routine, please read the Users' Note for your implementation to check the interpretation of *bold italicised* terms and other implementation-dependent details.

1 Purpose

F04QAF solves sparse non-symmetric equations, sparse linear least-squares problems and sparse damped linear least-squares problems, using a Lanczos algorithm.

2 Specification

```

SUBROUTINE F04QAF (M, N, B, X, SE, APROD, DAMP, ATOL, BTOL, CONLIM,
1 ITNLIM, MSGLVL, ITN, ANORM, ACOND, RNORM, ARNORM,
2 XNORM, WORK, RUSER, LRUSER, IUSER, LIUSER, INFORM,
3 IFAIL)

INTEGER M, N, ITNLIM, MSGLVL, ITN, LRUSER, IUSER(LIUSER),
1 LIUSER, INFORM, IFAIL
double precision B(M), X(N), SE(N), DAMP, ATOL, BTOL, CONLIM, ANORM,
1 ACOND, RNORM, ARNORM, XNORM, WORK(N,2), RUSER(LRUSER)
EXTERNAL APROD

```

3 Description

F04QAF can be used to solve a system of linear equations

$$Ax = b \quad (1)$$

where A is an n by n sparse non-symmetric matrix, or can be used to solve linear least-squares problems, so that F04QAF minimizes the value ρ given by

$$\rho = \|r\|, \quad r = b - Ax \quad (2)$$

where A is an m by n sparse matrix and $\|r\|$ denotes the Euclidean length of r so that $\|r\|^2 = r^T r$. A damping parameter, λ , may be included in the least-squares problem in which case F04QAF minimizes the value ρ given by

$$\rho^2 = \|r\|^2 + \lambda^2 \|x\|^2. \quad (3)$$

λ is supplied as the parameter DAMP and should of course be zero if the solution to problems (1) or (2) is required. Minimizing ρ in (3) is often called ridge regression.

F04QAF is based upon algorithm LSQR (see Paige and Saunders (1982a) and Paige and Saunders (1982b)) and solves the problems by an algorithm based upon the Lanczos process. The routine does not require A explicitly, but A is specified via a user-supplied (sub)program APROD which must perform the operations $(y + Ax)$ and $(x + A^T y)$ for a given n element vector x and m element vector y . A parameter to user-supplied (sub)program APROD specifies which of the two operations is required on a given entry.

The routine also returns estimates of the standard errors of the sample regression coefficients (x_i , for $i = 1, 2, \dots, n$) given by the diagonal elements of the estimated variance-covariance matrix V . When problem (2) is being solved and A is of full rank, then V is given by

$$V = s^2 (A^T A)^{-1}, \quad s^2 = \rho^2 / (m - n), \quad m > n$$

and when problem (3) is being solved then V is given by

$$V = s^2 (A^T A + \lambda^2 I)^{-1}, \quad s^2 = \rho^2 / m, \quad \lambda \neq 0.$$

Let \bar{A} denote the matrix

$$\bar{A} = A, \quad \lambda = 0; \quad \bar{A} = \begin{pmatrix} A \\ \lambda I \end{pmatrix}, \quad \lambda \neq 0, \quad (4)$$

let \bar{r} denote the residual vector

$$\bar{r} = r, \quad \lambda = 0; \quad \bar{r} = \begin{pmatrix} b \\ 0 \end{pmatrix} - \bar{A}x, \quad \lambda \neq 0 \quad (5)$$

corresponding to an iterate x , so that $\rho = \|\bar{r}\|$ is the function being minimized, and let $\|A\|$ denote the Frobenius (Euclidean) norm of A . Then the routine accepts x as a solution if it is estimated that one of the following two conditions is satisfied:

$$\rho \leq tol_1 \|\bar{A}\| \|x\| + tol_2 \|b\| \quad (6)$$

$$\|\bar{A}^T \bar{r}\| \leq tol_1 \|\bar{A}\| \rho \quad (7)$$

where tol_1 and tol_2 are user-supplied tolerances which estimate the relative errors in A and b respectively. Condition (6) is appropriate for compatible problems where, in theory, we expect the residual to be zero and will be satisfied by an acceptable solution x to a compatible problem. Condition (7) is appropriate for incompatible systems where we do not expect the residual to be zero and is based upon the observation that, in theory,

$$\bar{A}^T \bar{r} = 0$$

when x is a solution to the least-squares problem, and so (7) will be satisfied by an acceptable solution x to a linear least-squares problem.

The routine also includes a test to prevent convergence to solutions, x , with unacceptably large elements. This can happen if A is nearly singular or is nearly rank deficient. If we let the singular values of \bar{A} be

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0$$

then the condition number of \bar{A} is defined as

$$\text{cond}(\bar{A}) = \sigma_1 / \sigma_k$$

where σ_k is the smallest non-zero singular value of \bar{A} and hence k is the rank of \bar{A} . When $k < n$, then \bar{A} is rank deficient, the least-squares solution is not unique and F04QAF will normally converge to the minimal length solution. In practice \bar{A} will not have exactly zero singular values, but may instead have small singular values that we wish to regard as zero.

The routine provides for this possibility by terminating if

$$\text{cond}(\bar{A}) \geq c_{\text{lim}} \quad (8)$$

where c_{lim} is a user-supplied limit on the condition number of \bar{A} . For problem (1) termination with this condition indicates that A is nearly singular and for problem (2) indicates that A is nearly rank deficient and so has near linear dependencies in its columns. In this case inspection of $\|r\|$, $\|A^T r\|$ and $\|x\|$, which are all returned by the routine, will indicate whether or not an acceptable solution has been found. Condition (8), perhaps in conjunction with $\lambda \neq 0$, can be used to try and 'regularise' least-squares solutions. A full discussion of the stopping criteria is given in Section 6 of Paige and Saunders (1982a).

Introduction of a non-zero damping parameter λ tends to reduce the size of the computed solution and to make its components less sensitive to changes in the data, and F04QAF is applicable when a value of λ is known *a priori*. To have an effect, λ should normally be at least $\sqrt{\epsilon} \|A\|$ where ϵ is the *machine precision*. For further discussion see Paige and Saunders (1982b) and the references given there.

Whenever possible the matrix A should be scaled so that the relative errors in the elements of A are all of comparable size. Such a scaling helps to prevent the least-squares problem from being unnecessarily sensitive to data errors and will normally reduce the number of iterations required. At the very least, in the absence of better information, the columns of A should be scaled to have roughly equal column length.

4 References

Paige C C and Saunders M A (1982a) LSQR: An algorithm for sparse linear equations and sparse least-squares *ACM Trans. Math. Software* **8** 43–71

Paige C C and Saunders M A (1982b) Algorithm 583 LSQR: Sparse linear equations and least-squares problems *ACM Trans. Math. Software* **8** 195–209

5 Parameters

- 1: M – INTEGER *Input*
On entry: m , the number of rows of the matrix A .
Constraint: $M \geq 1$.
- 2: N – INTEGER *Input*
On entry: n , the number of columns of the matrix A .
Constraint: $N \geq 1$.
- 3: B(M) – **double precision** array *Input/Output*
On entry: the right-hand side vector b .
On exit: is overwritten.
- 4: X(N) – **double precision** array *Output*
On exit: the solution vector x .
- 5: SE(N) – **double precision** array *Output*
On exit: the estimates of the standard errors of the components of x . Thus $SE(i)$ contains an estimate of $\sqrt{v_{ii}}$, where v_{ii} is the i th diagonal element of the estimated variance-covariance matrix V . The estimates returned in SE will be the lower bounds on the actual estimated standard errors, but will usually be correct to at least one significant figure.
- 6: APROD – SUBROUTINE, supplied by the user. *External Procedure*
 APROD must perform the operations $y := y + Ax$ and $x := x + A^T y$ for given vectors x and y .
 Its specification is:

<pre> SUBROUTINE APROD (MODE, M, N, X, Y, RUSER, LRUSER, IUSER, LIUSER) INTEGER MODE, M, N, LRUSER, IUSER(LIUSER), LIUSER double precision X(N), Y(M), RUSER(LRUSER) </pre>
<p>1: MODE – INTEGER <i>Input/Output</i></p> <p><i>On entry:</i> specifies which operation is to be performed.</p> <p>MODE = 1 APROD must compute $y + Ax$.</p> <p>MODE = 2 APROD must compute $x + A^T y$.</p> <p><i>On exit:</i> may be used as a flag to indicate a failure in the computation of $y + Ax$ or $x + A^T y$. If MODE is negative on exit from APROD, F04QAF will exit immediately with IFAIL set to MODE.</p>

2:	M – INTEGER	<i>Input</i>
	<i>On entry:</i> m , the number of rows of A .	
3:	N – INTEGER	<i>Input</i>
	<i>On entry:</i> n , the number of columns of A .	
4:	X(N) – double precision array	<i>Input/Output</i>
	<i>On entry:</i> the vector x .	
	<i>On exit:</i> if MODE = 1, X must be unchanged.	
	If MODE = 2, X must contain $x + A^T y$.	
5:	Y(M) – double precision array	<i>Input/Output</i>
	<i>On entry:</i> the vector y .	
	<i>On exit:</i> if MODE = 1, Y must contain $y + Ax$.	
	If MODE = 2, Y must be unchanged.	
6:	RUSER(LRUSER) – double precision array	<i>User Workspace</i>
7:	LRUSER – INTEGER	<i>Input</i>
8:	IUSER(LIUSER) – INTEGER array	<i>User Workspace</i>
9:	LIUSER – INTEGER	<i>Input</i>
	APROD is called from F04QAF with the parameters RUSER, LRUSER, IUSER and LIUSER as supplied to F04QAF. You are free to use the arrays RUSER and IUSER to supply information to APROD as an alternative to using COMMON.	

APROD must be declared as EXTERNAL in the (sub)program from which F04QAF is called. Parameters denoted as *Input* must **not** be changed by this procedure.

- 7: DAMP – **double precision** *Input*
On entry: the value λ . If either problem (1) or problem (2) is to be solved, then DAMP must be supplied as zero.
- 8: ATOL – **double precision** *Input*
On entry: the tolerance, tol_1 , of the convergence criteria (6) and (7); it should be an estimate of the largest relative error in the elements of A . For example, if the elements of A are correct to about 4 significant figures, then ATOL should be set to about 5×10^{-4} . If ATOL is supplied as less than ϵ , where ϵ is the **machine precision**, then the value ϵ is used instead of ATOL.
- 9: BTOL – **double precision** *Input*
On entry: the tolerance, tol_2 , of the convergence criterion (6); it should be an estimate of the largest relative error in the elements of B . For example, if the elements of B are correct to about 4 significant figures, then BTOL should be set to about 5×10^{-4} . If BTOL is supplied as less than ϵ , then the value ϵ is used instead of BTOL.
- 10: CONLIM – **double precision** *Input*
On entry: the value c_{lim} of equation (8); it should be an upper limit on the condition number of \bar{A} . CONLIM should not normally be chosen much larger than $1.0/ATOL$. If CONLIM is supplied as zero, then the value $1.0/\epsilon$ is used instead of CONLIM.

- 11: ITNLIM – INTEGER *Input/Output*
On entry: an upper limit on the number of iterations. If $ITNLIM \leq 0$, then the value N is used in place of ITNLIM, but for ill-conditioned problems a higher value of ITNLIM is likely to be necessary.
On exit: unchanged unless $ITNLIM \leq 0$ on entry, in which case it is set to N.
- 12: MSGLVL – INTEGER *Input*
On entry: the level of printing from F04QAF. If $MSGLVL \leq 0$, then no printing occurs, but otherwise messages will be output on the advisory message channel (see X04ABF). A description of the printed output is given in Section 8.1. The level of printing is determined as follows:
 $MSGLVL \leq 0$
 No printing.
 $MSGLVL = 1$
 A brief summary is printed just prior to return from F04QAF.
 $MSGLVL \geq 2$
 A summary line is printed periodically to monitor the progress of F04QAF, together with a brief summary just prior to return from F04QAF.
- 13: ITN – INTEGER *Output*
On exit: the number of iterations performed.
- 14: ANORM – *double precision* *Output*
On exit: an estimate of $\|\bar{A}\|$ for the matrix \bar{A} of (4).
- 15: ACOND – *double precision* *Output*
On exit: an estimate of $\text{cond}(\bar{A})$ which is a lower bound.
- 16: RNORM – *double precision* *Output*
On exit: an estimate of $\|\bar{r}\|$ for the residual, \bar{r} , of (5) corresponding to the solution x returned in X. Note that $\|\bar{r}\|$ is the function being minimized.
- 17: ARNORM – *double precision* *Output*
On exit: an estimate of the $\|\bar{A}^T \bar{r}\|$ corresponding to the solution x returned in X.
- 18: XNORM – *double precision* *Output*
On exit: an estimate of $\|x\|$ for the solution x returned in X.
- 19: WORK(N,2) – *double precision* array *Workspace*
- 20: RUSER(LRUSER) – *double precision* array *User Workspace*
 RUSER is not used by F04QAF, but is passed directly to user-supplied (sub)program APROD and may be used to pass information to that routine.
- 21: LRUSER – INTEGER *Input*
On entry: the dimension of the array RUSER as declared in the (sub)program from which F04QAF is called.
Constraint: $LRUSER \geq 1$.

- 22: IUSER(LIUSER) – INTEGER array *User Workspace*
 IUSER is not used by F04QAF, but is passed directly to user-supplied (sub)program APROD and may be used to pass information to that routine.
- 23: LIUSER – INTEGER *Input*
On entry: the dimension of the array IUSER as declared in the (sub)program from which F04QAF is called.
Constraint: LIUSER \geq 1.
- 24: INFORM – INTEGER *Output*
On exit: the reason for termination of F04QAF.
 INFORM = 0
 The exact solution is $x = 0$. No iterations are performed in this case.
 INFORM = 1
 The termination criterion of (6) has been satisfied with tol_1 and tol_2 as the values supplied in ATOL and BTOL respectively.
 INFORM = 2
 The termination criterion of (7) has been satisfied with tol_1 as the value supplied in ATOL.
 INFORM = 3
 The termination criterion of (6) has been satisfied with tol_1 and/or tol_2 as the value ϵ , where ϵ is the *machine precision*. One or both of the values supplied in ATOL and BTOL must have been less than ϵ and was too small for this machine.
 INFORM = 4
 The termination criterion of (7) has been satisfied with tol_1 as the value ϵ . The value supplied in ATOL must have been less than ϵ and was too small for this machine.
 The values INFORM = 5, 6 and 7 correspond to failure with IFAIL = 2, 3 and 4 respectively (see Section 6) and when IFAIL is negative INFORM will be set to the same negative value.
- 25: IFAIL – INTEGER *Input/Output*
On entry: IFAIL must be set to 0, -1 or 1. If you are unfamiliar with this parameter you should refer to Chapter P01 for details.
On exit: IFAIL = 0 unless the routine detects an error (see Section 6).
 For environments where it might be inappropriate to halt program execution when an error is detected, the value -1 or 1 is recommended. If the output of error messages is undesirable, then the value 1 is recommended. Otherwise, if you are not familiar with this parameter the recommended value is 0. **When the value -1 or 1 is used it is essential to test the value of IFAIL on exit.**

6 Error Indicators and Warnings

Errors detected by the routine:

IFAIL < 0

A negative value of IFAIL indicates an exit from F04QAF because you have set MODE negative in the user-supplied (sub)program APROD. The value of IFAIL will be the same as your setting of MODE.

IFAIL = 1

On entry, $M < 1$,
 or $N < 1$,
 or $LRUSER < 1$,
 or $LIUSER < 1$.

IFAIL = 2

The condition of (8) has been satisfied for the value of c_{lim} supplied in CONLIM. If this failure is unexpected you should check that (sub)program APROD is working correctly. Although conditions (6) or (7) have not been satisfied, the values returned in RNORM, ARNORM and XNORM may nevertheless indicate that an acceptable solution has been reached.

IFAIL = 3

The condition of (8) has been satisfied for the value $c_{lim} = 1.0/\epsilon$, where ϵ is the *machine precision*. The matrix \bar{A} is nearly singular or rank deficient and the problem is too ill-conditioned for this machine. If this failure is unexpected, you should check that (sub)program APROD is working correctly.

IFAIL = 4

The limit on the number of iterations has been reached. The number of iterations required by F04QAF and the condition of the matrix \bar{A} can depend strongly on the scaling of the problem. Poor scaling of the rows and columns of A should be avoided whenever possible.

7 Accuracy

When the problem is compatible, the computed solution x will satisfy the equation

$$r = b - Ax,$$

where an estimate of $\|r\|$ is returned in the parameter RNORM. When the problem is incompatible, the computed solution x will satisfy the equation

$$\bar{A}^T \bar{r} = e,$$

where an estimate of $\|e\|$ is returned in the parameter ARNORM. See also Section 6.2 of Paige and Saunders (1982b).

8 Further Comments

The time taken by F04QAF is likely to be principally determined by the time taken in the user-supplied (sub)program APROD, which is called twice on each iteration, once with MODE = 1 and once with MODE = 2. The time taken per iteration by the remaining operations in F04QAF is approximately proportional to $\max(m, n)$.

The Lanczos process will usually converge more quickly if A is pre-conditioned by a non-singular matrix M that approximates A in some sense and is also chosen so that equations of the form $My = c$ can efficiently be solved for y . For a discussion of pre-conditioning, see the F11 Chapter Introduction. In the context of F04QAF, problem (1) is equivalent to

$$(AM^{-1})y = b, \quad Mx = y$$

and problem (2) is equivalent to minimizing

$$\rho = \|r\|, \quad r = b - (AM^{-1})y, \quad Mx = y.$$

Note that the normal matrix $(AM^{-1})^T(AM^{-1}) = M^{-T}(A^T A)M^{-1}$ so that the pre-conditioning AM^{-1} is equivalent to the pre-conditioning $M^{-T}(A^T A)M^{-1}$ of the normal matrix $A^T A$.

Pre-conditioning can be incorporated into F04QAF simply by coding the user-supplied (sub)program APROD to compute $y + AM^{-1}x$ and $x + M^{-T}A^T y$ in place of $y + Ax$ and $x + A^T y$ respectively, and then solving the equations $Mx = y$ for x on return from F04QAF. The quantity $y + AM^{-1}x$ should be computed by solving $Mz = x$ for z and then computing $y + Az$, and $x + M^{-T}A^T y$ should be computed by solving $M^T z = A^T y$ for z and then forming $x + z$.

8.1 Description of the Printed Output

When MSGLVL > 0, then F04QAF will produce output (except in the case where the routine fails with IFAIL = 1) on the advisory message channel (see X04ABF).

When MSGLVL ≥ 2 then a summary line is printed periodically giving the following information:

Output	Meaning
ITN	Iteration number, k .
X(1)	The first element of the current iterate x_k .
FUNCTION	The current value of the function, ρ , being minimized.
COMPAT	An estimate of $\ \bar{r}_k\ /\ b\ $, where \bar{r}_k is the residual corresponding to x_k . This value should converge to zero (in theory) if and only if the problem is compatible. COMPAT decreases monotonically.
INCOMPAT	An estimate of $\ \bar{A}^T \bar{r}_k\ /(\ \bar{A}\ \ \bar{r}_k\)$ which should converge to zero if and only if at the solution ρ is non-zero. INCOMPAT is not usually monotonic.
NRM(ABAR)	A monotonically increasing estimate of $\ \bar{A}\ $.
COND(ABAR)	A monotonically increasing estimate of the condition number $\text{cond}(\bar{A})$.

9 Example

This example solves the linear least-squares problem

$$\min \rho = \|r\|, \quad r = b - Ax$$

where A is the 13 by 12 matrix and b is the 13 element vector given by

$$A = \begin{pmatrix} 1 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & 0 & -1 & 4 & -1 & 0 & 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & -1 & 4 & -1 & 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & 0 & -1 & 4 & -1 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 & 0 & -1 & 4 & -1 & 0 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 1 \end{pmatrix}, \quad b = -h^2 \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \\ 1 \\ 1 \\ 0 \\ 1 \\ 0 \\ 0 \\ -h^{-3} \end{pmatrix}$$

with $h = 0.1$.

Such a problem can arise by considering the Neumann problem on a rectangle

$$\begin{aligned} \frac{\delta u}{\delta n} &= 0 \\ \frac{\delta u}{\delta n} &= 0 \quad \boxed{\nabla^2 u = g(x,y)} \quad \frac{\delta u}{\delta n} = 0 \quad \int_c u = 1 \\ \frac{\delta u}{\delta n} &= 0 \end{aligned}$$

where C is the boundary of the rectangle, and discretizing as illustrated below with the square mesh




Figure 1

The 12 by 12 symmetric part of A represents the difference equations and the final row comes from the normalizing condition. The example program has $g(x,y) = 1$ at all the internal mesh points, but apart from this is written in a general manner so that the number of rows (NROWS) and columns (NCOLS) in the grid can readily be altered.

9.1 Program Text

```
*      F04QAF Example Program Text
*      Mark 14 Revised. NAG Copyright 1989.
*      .. Parameters ..
INTEGER          MMAX, NMAX, LIWORK, LRWORK
PARAMETER       (MMAX=100,NMAX=99,LIWORK=1,LRWORK=1)
INTEGER          NOUT
PARAMETER       (NOUT=6)
*      .. Scalars in Common ..
INTEGER          NCOLS, NROWS
*      .. Local Scalars ..
DOUBLE PRECISION ACOND, ANORM, ARNORM, ATOL, BTOL, C, CONLIM,
+              DAMP, H, RNORM, XNORM
INTEGER          I, I1, IFAIL, INFORM, ITN, ITNLIM, K, M, MSGLVL,
+              N
*      .. Local Arrays ..
DOUBLE PRECISION B(MMAX), RWORK(LRWORK), SE(NMAX), WORK(NMAX,2),
+              X(NMAX)
INTEGER          IWORK(LIWORK)
*      .. External Subroutines ..
EXTERNAL         APROD, F04QAF, X04ABF
*      .. Common blocks ..
COMMON          /USER/NROWS, NCOLS
*      .. Executable Statements ..
WRITE (NOUT,*) 'F04QAF Example Program Results'
WRITE (NOUT,*)
CALL X04ABF(1,NOUT)
NROWS = 4
NCOLS = 4
H = 0.1D0
N = NCOLS*NROWS - 4
M = N + 1
IF (NROWS.LT.3 .OR. NCOLS.LT.3 .OR. M.GT.MMAX) THEN
  WRITE (NOUT,99998) 'NROWS or NCOLS is out of range: NROWS = ',
+  NROWS, ' NCOLS = ', NCOLS
ELSE
*
*      Initialize RHS and other quantities required by F04QAF.
*      Convergence will be sooner if we do not regard A as exact,
*      so ATOL is not set to zero.
*
DO 20 I = 1, N
  B(I) = 0.0D0
20 CONTINUE
C = -H**2
I1 = NROWS
DO 60 K = 3, NCOLS
```

```

      DO 40 I = I1, I1 + NROWS - 3
        B(I) = C
40     CONTINUE
      I1 = I1 + NROWS
60     CONTINUE
      B(M) = 1.0D0/H
      DAMP = 0.0D0
      ATOL = 1.0D-5
      BTOL = 1.0D-4
      CONLIM = 1.0D0/ATOL
      ITNLIM = 100
*     * Set MSGVLV to 2 to get output at each iteration *
      MSGVLV = 1
      IFAIL = 1
*
      CALL F04QAF(M,N,B,X,SE,APROD,DAMP,ATOL,BTOL,CONLIM,ITNLIM,
+           MSGVLV,ITN,ANORM,ACOND,RNORM,ARNORM,XNORM,WORK,
+           RWORK,LRWORK,IWORK,LIWORK,INFORM,IFAIL)
*
      WRITE (NOUT,*)
      IF (IFAIL.NE.0) THEN
        WRITE (NOUT,99999) 'F04QAF fails. IFAIL =', IFAIL
      ELSE
        WRITE (NOUT,*) 'Solution returned by F04QAF'
        WRITE (NOUT,99997) (X(I),I=1,N)
        WRITE (NOUT,*)
        WRITE (NOUT,99996) 'Norm of the residual = ', RNORM
      END IF
    END IF
  END STOP
*
99999 FORMAT (1X,A,I3)
99998 FORMAT (1X,A,I6,A,I6)
99997 FORMAT (1X,5F9.3)
99996 FORMAT (1X,A,1P,E12.2)
  END
*
  SUBROUTINE APROD(MODE,M,N,X,Y,RWORK,LRWORK,IWORK,LIWORK)
*   APROD returns
*     Y = Y + A*X           when MODE = 1
*     X = X + (A**T)*Y     when MODE = 2
*   for a given X and Y.
*   .. Scalar Arguments ..
  INTEGER          LIWORK, LRWORK, M, MODE, N
*   .. Array Arguments ..
  DOUBLE PRECISION RWORK(LRWORK), X(N), Y(M)
  INTEGER          IWORK(LIWORK)
*   .. Scalars in Common ..
  INTEGER          NCOLS, NROWS
*   .. Local Scalars ..
  INTEGER          J, J1, J2
*   .. External Subroutines ..
  EXTERNAL        ATIMES
*   .. Common blocks ..
  COMMON          /USER/NROWS, NCOLS
*   .. Executable Statements ..
  IF (MODE.NE.2) THEN
    CALL ATIMES(NROWS,NCOLS,N,X,Y)
    DO 20 J = 1, NROWS - 2
      Y(M) = Y(M) + X(J)
20    CONTINUE
    DO 40 J = 1, NCOLS - 2
      Y(M) = Y(M) + X(J*NROWS-1) + X(J*NROWS+NROWS-2)
40    CONTINUE
    DO 60 J = M - NROWS + 2, N
      Y(M) = Y(M) + X(J)
60    CONTINUE
  ELSE
    CALL ATIMES(NROWS,NCOLS,N,Y,X)
    DO 80 J = 1, NROWS - 2
      X(J) = X(J) + Y(M)
80    CONTINUE
  END IF

```

```

80     CONTINUE
      DO 100 J = 1, NCOLS - 2
        J1 = J*NROWS - 1
        J2 = J1 + NROWS - 1
        X(J1) = X(J1) + Y(M)
        X(J2) = X(J2) + Y(M)
100    CONTINUE
      DO 120 J = M - NROWS + 2, N
        X(J) = X(J) + Y(M)
120    CONTINUE
      END IF
      RETURN
      END
*
      SUBROUTINE ATIMES(NROWS,NCOLS,N,X,Y)
*     ATIMES is called by routine APROD and returns
*     Y = Y + ANN*X,
*     where ANN is the N by N symmetric part of the matrix A.
*     .. Scalar Arguments ..
      INTEGER          N, NCOLS, NROWS
*     .. Array Arguments ..
      DOUBLE PRECISION X(N), Y(N)
*     .. Local Scalars ..
      INTEGER          I, I1, I2, I3, IL, J
*     .. Executable Statements ..
      DO 20 J = 1, NROWS - 2
        Y(J) = Y(J) + X(J) - X(J+NROWS-1)
20     CONTINUE
      DO 60 J = 1, NCOLS - 2
        I = J*NROWS - 1
        Y(I) = Y(I) + X(I) - X(I+1)
        I1 = I + 1
        IL = I1 + NROWS - 3
        DO 40 I = I1, IL
          I2 = I - NROWS
          IF (J.EQ.1) I2 = I2 + 1
          I3 = I + NROWS
          IF (J.EQ.NCOLS-2) I3 = I3 - 1
          Y(I) = Y(I) - X(I2) - X(I-1) + 4.0D0*X(I) - X(I+1) - X(I3)
40     CONTINUE
        I = IL + 1
        Y(I) = Y(I) - X(I-1) + X(I)
60     CONTINUE
      DO 80 J = N - NROWS + 3, N
        Y(J) = Y(J) - X(J-NROWS+1) + X(J)
80     CONTINUE
      RETURN
      END

```

9.2 Program Data

None.

9.3 Program Results

F04QAF Example Program Results

OUTPUT FROM F04QAF.

LEAST-SQUARES SOLUTION OF $A \cdot X = B$

THE MATRIX A HAS 13 ROWS AND 12 COLS
 THE DAMPING PARAMETER IS DAMP = 0.00E+00

ATOL = 1.00E-05 CONLIM = 1.00E+05
 BTOL = 1.00E-04 ITNLIM = 100

NO. OF ITERATIONS = 2
 STOPPING CONDITION = 2
 (THE LEAST-SQRS SOLN IS GOOD ENOUGH, GIVEN ATOL)

ACTUAL	NORM(RBAR),	NORM(X)	1.15E-02	4.33E+00
	NORM(TRANSPOSE(ABAR)*RBAR)		9.16E-15	
ESTIMATES OF	NORM(ABAR),	COND(ABAR)	4.12E+00	2.45E+00

Solution returned by F04QAF

1.250	1.250	1.250	1.247	1.247
1.250	1.250	1.247	1.247	1.250
1.250	1.250			

Norm of the residual = 1.15E-02
