F11BRF – NAG Fortran Library Routine Document

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

F11BRF is a set-up routine, the first in a suite of three routines for the iterative solution of a complex general (non-Hermitian) system of simultaneous linear equations. F11BRF must be called before F11BSF, the iterative solver. The third routine in the suite, F11BTF, can be used to return additional information about the computation.

These three routines are suitable for the solution of large sparse general (non-Hermitian) systems of equations.

2 Specification

```
   SUBROUTINE F11BRF(METHOD, PRECON, NORM, WEIGHT, ITERM, N, M, TOL,
  1                   MAXITN, ANORM, SIGMAX, MONIT, LWREQ, WORK,
  2                   LWORK, IFAIL)
   INTEGER           ITERM, N, M, MAXITN, MONIT, LWREQ, LWORK, IFAIL
   real              TOL, ANORM, SIGMAX
   CHARACTER*1       PRECON, NORM, WEIGHT
   CHARACTER*(*)     METHOD
   complex           WORK(LWORK)
```

3 Description

The suite consisting of the routines F11BRF, F11BSF, F11BTF is designed to solve the general (non-Hermitian) system of simultaneous linear equations $Ax = b$ of order n, where n is large and the coefficient matrix A is sparse.

F11BRF is a set-up routine which must be called before F11BSF, the iterative solver. The third routine in the suite, F11BTF, can be used to return additional information about the computation. A choice of methods is available:

restarted generalized minimum residual method (RGMRES);

conjugate gradient squared method (CGS);

bi-conjugate gradient stabilized (ℓ) method (Bi-CGSTAB (ℓ));

transpose-free quasi-minimal residual method (TFQMR).

3.1 Restarted Generalized Minimum Residual Method (RGMRES)

The restarted generalized minimum residual method (RGMRES) (Saad and Schultz [4], Barrett *et al.* [2], Dias da Cunha and Hopkins [3]) starts from the residual $r_0 = b - Ax_0$, where x_0 is an initial estimate for the solution (often $x_0 = 0$). An orthogonal basis for the Krylov subspace span$\{A^k r_0\}$, for $k = 0, 1, 2, \ldots$, is generated explicitly: this is referred to as Arnoldi's method (Arnoldi [8]). The solution is then expanded onto the orthogonal basis so as to minimize the residual norm $\|b - Ax\|_2$. The lack of symmetry of A implies that the orthogonal basis is generated by applying a 'long' recurrence relation, whose length increases linearly with the iteration count. For all but the most trivial problems, computational and storage costs can quickly become prohibitive as the iteration count increases. RGMRES limits these costs by employing a restart strategy: every m iterations at most, the Arnoldi process is restarted from $r_l = b - Ax_l$, where the subscript l denotes the last available iterate. Each group of m iterations is referred to as a 'super-iteration'. The value of m is chosen in advance and is fixed throughout the computation. Unfortunately, an optimum value of m cannot easily be predicted.

3.2   Conjugate Gradient Squared Method (CGS)

The conjugate gradient squared method (CGS) (Sonneveld [5], Barrett *et al.* [2], Dias da Cunha and Hopkins [3]) is a development of the bi-conjugate gradient method where the nonsymmetric Lanczos method is applied to reduce the coefficients matrix to tridiagonal form: two bi-orthogonal sequences of vectors are generated starting from the residual $r_0 = b - Ax_0$, where x_0 is an initial estimate for the solution (often $x_0 = 0$) and from the *shadow residual* $\hat{r}_0$ corresponding to the arbitrary problem $A^H \hat{x} = \hat{b}$, where $\hat{b}$ can be any vector, but in practice is chosen so that $r_0 = \hat{r}_0$. In the course of the iteration, the residual and shadow residual $r_i = P_i(A)r_0$ and $\hat{r}_i = P_i(A^H)\hat{r}_0$ are generated, where P_i is a polynomial of order i, and bi-orthogonality is exploited by computing the vector product $\rho_i = (\hat{r}_i, r_i) = (P_i(A^H)\hat{r}_0, P_i(A)r_0) = (\hat{r}_0, P_i^2(A)r_0)$. Applying the 'contraction' operator $P_i(A)$ twice, the iteration coefficients can still be recovered without advancing the solution of the shadow problem, which is of no interest. The CGS method often provides fast convergence; however, there is no reason why the contraction operator should also reduce the once reduced vector $P_i(A)r_0$: this may well lead to a highly irregular convergence which may result in large cancellation errors.

3.3   Bi-Conjugate Gradient Stabilized (ℓ) Method (Bi-CGSTAB (ℓ))

The bi-conjugate gradient stabilized (ℓ) method (Bi-CGSTAB (ℓ)) (van der Vorst [6], Sleijpen and Fokkema [7], Dias da Cunha and Hopkins [3]) is similar to the CGS method above. However, instead of generating the sequence $\{P_i^2(A)r_0\}$, it generates the sequence $\{Q_i(A)P_i(A)r_0\}$ where the $Q_i(A)$ are polynomials chosen to minimize the residual *after* the application of the contraction operator $P_i(A)$. Two main steps can be identified for each iteration: an OR (Orthogonal Residuals) step where a basis of order ℓ is generated by a Bi-CG iteration and an MR (Minimum Residuals) step where the residual is minimized over the basis generated, by a method akin to GMRES. For $\ell = 1$, the method corresponds to the Bi-CGSTAB method of van der Vorst [6]. For $\ell > 1$, more information about complex eigenvalues of the iteration matrix can be taken into account, and this may lead to improved convergence and robustness. However, as ℓ increases, numerical instabilities may arise. For this reason, a maximum value of $\ell = 10$ is imposed, but probably $\ell = 4$ is sufficient in most cases.

3.4   Transpose-Free Quasi-Minimal Residual Method (TFQMR)

The transpose-free quasi-minimal residual method (TFQMR) (Freund and Nachtigal [10], Freund [11]) is conceptually derived from the CGS method. The residual is minimised over the space of the residual vectors generated by the CGS iterations under the simplifying assumption that residuals are almost orthogonal. In practice, this is not the case but theoretical analysis has proved the validity of the method. This has the effect of remedying the rather irregular convergence behaviour with wild oscillations in the residual norm that can degrade the numerical performance and robustness of the CGS method. In general, the TFQMR method can be expected to converge at least as fast as the CGS method, in terms of number of iterations, although each iteration involves a higher operation count. When the CGS method exhibits irregular convergence, the TFQMR method can produce much smoother, almost monotonic convergence curves. However, the close relationship between the CGS and TFQMR method implies that the *overall* speed of convergence is similar for both methods. In some cases, the TFQMR method may converge faster than the CGS method.

3.5   General Considerations

For each method, a sequence of solution iterates $\{x_i\}$ is generated such that, hopefully, the sequence of the residual norms $\{\|r_i\|\}$ converges to a required tolerance. Note that, in general, convergence, when it occurs, is not monotonic.

In the RGMRES and Bi-CGSTAB (ℓ) methods above, the user's program must provide the **maximum** number of basis vectors used, m or ℓ, respectively; however, a **smaller** number of basis vectors may be generated and used when the stability of the solution process requires this (see Section 8).

Faster convergence can be achieved using a **preconditioner** (Golub and Van Loan [1], Barrett *et al.* [2]). A preconditioner maps the original system of equations onto a different system, say

$$\bar{A}\bar{x} = \bar{b}, \tag{1}$$

with, hopefully, better characteristics with respect to its speed of convergence: for example, the condition number of the coefficients matrix can be improved or eigenvalues in its spectrum can be made to coalesce.

An orthogonal basis for the Krylov subspace span$\{\bar{A}^k \bar{r}_0\}$, for $k = 0, 1, \ldots$, is generated and the solution proceeds as outlined above. The algorithms used are such that the solution and residual iterates of the original system are produced, not their preconditioned counterparts. Note that an unsuitable preconditioner or no preconditioning at all may result in a very slow rate, or lack, of convergence. However, preconditioning involves a trade-off between the reduction in the number of iterations required for convergence and the additional computational costs per iteration. Also, setting up a preconditioner may involve non-negligible overheads.

A *left* preconditioner M^{-1} can be used by the RGMRES, CGS and TFQMR methods, such that $\bar{A} = M^{-1}A \sim I_n$ in (1), where I_n is the identity matrix of order n; a *right* preconditioner M^{-1} can be used by the Bi-CGSTAB (ℓ) method, such that $\bar{A} = AM^{-1} \sim I_n$. These are formal definitions, used only in the design of the algorithms; in practice, only the means to compute the matrix–vector products $v = Au$ and $v = A^H u$ (the latter only being required when an estimate of $\|A\|_1$ or $\|A\|_\infty$ is computed internally), and to solve the preconditioning equations $Mv = u$ are required, i.e., explicit information about M, or its inverse is not required at any stage.

The first termination criterion

$$\|r_k\|_p \leq \tau \left(\|b\|_p + \|A\|_p \|x_k\|_p \right) \tag{2}$$

is also available for all three methods. In (2), $p = 1, \infty$ or 2 and τ denotes a user-specified tolerance subject to $\max(10, \sqrt{n})$, $\epsilon \leq \tau < 1$, where ϵ is the **machine precision**. Facilities are provided for the estimation of the norm of the coefficients matrix $\|A\|_1$ or $\|A\|_\infty$, when this is not known in advance, by applying Higham's method (Higham [9]). Note that $\|A\|_2$ cannot be estimated internally. This criterion uses an error bound derived from **backward** error analysis to ensure that the computed solution is the exact solution of a problem as close to the original as the termination tolerance requires. Termination criteria employing bounds derived from **forward** error analysis are not used because any such criteria would require information about the condition number $\kappa(A)$ which is not easily obtainable.

The second termination criterion

$$\|\bar{r}_k\|_2 \leq \tau \left(\|\bar{r}_0\|_2 + \sigma_1(\bar{A}) \|\Delta \bar{x}_k\|_2 \right) \tag{3}$$

is available for all methods except TFQMR. In (3), $\sigma_1(\bar{A}) = \|\bar{A}\|_2$ is the largest singular value of the (preconditioned) iteration matrix $\bar{A}$. This termination criterion monitors the progress of the solution of the preconditioned system of equations and is less expensive to apply than criterion (2) for the Bi-CGSTAB (ℓ) method with $\ell > 1$. Only the RGMRES method provides facilities to estimate $\sigma_1(\bar{A})$ internally, when this is not supplied (see Section 8).

Termination criterion (2) is the recommended choice, despite its additional costs per iteration when using the Bi-CGSTAB (ℓ) method with $\ell > 1$. Also, if the norm of the initial estimate is much larger than the norm of the solution, that is, if $\|x_0\| \gg \|x\|$, a dramatic loss of significant digits could result in complete lack of convergence. The use of criterion (2) will enable the detection of such a situation, and the iteration will be restarted at a suitable point. No such restart facilities are provided for criterion (3).

Optionally, a vector w of user-specified weights can be used in the computation of the vector norms in termination criterion (2), i.e., $\|v\|_p^{(w)} = \|v^{(w)}\|_p$, where $(v^{(w)})_i = w_i v_i$, for $i = 1, 2, \ldots, n$. Note that the use of weights increases the computational costs.

The sequence of calls to the routines comprising the suite is enforced: first, the set-up routine F11BRF must be called, followed by the solver F11BSF. F11BTF can be called either when F11BSF is carrying out a monitoring step or after F11BSF has completed its tasks. Incorrect sequencing will raise an error condition.

In general, it is not possible to recommend one method in preference to another. RGMRES is often used in the solution of systems arising from PDEs. On the other hand, it can easily stagnate when the size m of the orthogonal basis is too small, or the preconditioner is not good enough. CGS can be the fastest method, but the computed residuals can exhibit instability which may greatly affect the convergence and quality of the solution. Bi-CGSTAB (ℓ) seems robust and reliable, but it can be slower than the other methods: if a preconditioner is used and $\ell > 1$, Bi-CGSTAB (ℓ) computes the solution of the preconditioned system $\bar{x}_k = Mx_k$: the preconditioning equations must be solved to obtain the required solution. The algorithm employed limits to 10% or less, when no intermediate monitoring is requested, the number of times the preconditioner has to be thus applied compared with the total number of applications of the preconditioner. TFQMR can be viewed as a more robust variant of the CGS method: it shares the

CGS method speed but avoids the CGS fluctuations in the residual, which may give rise to instability. Also, when the termination criterion (2) is used, the CGS, Bi-CGSTAB (ℓ) and TFQMR methods will restart the iteration automatically when necessary in order to solve the given problem.

4 References

[1] Golub G H and Van Loan C F (1996) *Matrix Computations* Johns Hopkins University Press (3rd Edition), Baltimore

[2] Barrett R, Berry M, Chan T F, Demmel J, Donato J, Dongarra J, Eijkhout V, Pozo R, Romine C and van der Vorst H (1994) *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods* SIAM, Philadelphia

[3] Dias da Cunha R and Hopkins T (1994) PIM 1.1 — the parallel iterative method package for systems of linear equations user's guide — Fortran 77 version *Technical Report* Computing Laboratory, University of Kent at Canterbury, Kent CT2 7NZ, UK

[4] Saad Y and Schultz M (1986) GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems *SIAM J. Sci. Statist. Comput.* **7** 856–869

[5] Sonneveld P (1989) CGS, a fast Lanczos-type solver for nonsymmetric linear systems *SIAM J. Sci. Statist. Comput.* **10** 36–52

[6] van der Vorst H (1989) Bi–CGSTAB, A fast and smoothly converging variant of Bi–CG for the solution of nonsymmetric linear systems *SIAM J. Sci. Statist. Comput.* **13** 631–644

[7] Sleijpen G L G and Fokkema D R (1993) BiCGSTAB(ℓ) for linear equations involving matrices with complex spectrum *ETNA* **1** 11–32

[8] Arnoldi W (1951) The principle of minimized iterations in the solution of the matrix eigenvalue problem *Quart. Appl. Math.* **9** 17–29

[9] Higham N J (1988) FORTRAN codes for estimating the one-norm of a real or complex matrix, with applications to condition estimation *ACM Trans. Math. Software* **14** 381–396

[10] Freund R W and Nachtigal N (1991) QMR: a Quasi-Minimal Residual Method for Non-Hermitian Linear Systems *Numer. Math.* **60** 315–339

[11] Freund R W (1993) A Transpose-Free Quasi-Mimimal Residual Algorithm for Non-Hermitian Linear Systems *SIAM J. Sci. Comput.* **14** 470–482

5 Parameters

1: METHOD — CHARACTER*(*) *Input*

On entry: the iterative method to be used. The possible choices are:

'RGMRES' restarted generalized minimum residual method;
'CGS' conjugate gradient squared method;
'BICGSTAB' bi-conjugate gradient stabilized (ℓ) method;
'TFQMR' transpose-free quasi-minimal residual method.

Constraint: METHOD = 'RGMRES', 'CGS', 'BICGSTAB' or 'TFQMR'.

2: PRECON — CHARACTER*1 *Input*

On entry: determines whether preconditioning is used. The possible choices are:

'N' no preconditioning;
'P' preconditioning.

Constraint: PRECON = 'N' or 'P'.

3:   NORM — CHARACTER*1   *Input*

On entry: defines the matrix and vector norm to be used inthe termination criteria. The possible choices are:

'1' l_1 norm;
'I' l_∞ norm;
'2' l_2 norm.

Suggested value: NORM = 'I', if ITERM = 1; NORM = '2', if ITERM = 2.

Constraints: if ITERM = 1, then NORM = '1', 'I' or '2'; if ITERM = 2, then NORM = '2'.

4:   WEIGHT — CHARACTER*1   *Input*

On entry: specifies whether a vector w of user-supplied weights is to be used in the computation of the vector norms required in termination criterion (2) (ITERM = 1): $\|v\|_p^{(w)} = \|v^{(w)}\|_p$, where $v_i^{(w)} = w_i v_i$, for $i = 1, 2, \ldots, n$. The suffix $p = 1, 2, \infty$ denotes the vector norm used, as specified by the parameter NORM. Note that weights cannot be used when ITERM = 2, i.e., when criterion (3) is used. The possible choices are:

'W'   user-supplied weights are to be used and must be supplied on initial entry to F11BSF.
'N'   all weights are implicitly set equal to one. Weights do not need to be supplied on initial entry to F11BSF.

Suggested value: WEIGHT = 'N'.

Constraints: if ITERM = 1, then WEIGHT = 'W' or 'N'; if ITERM = 2, then WEIGHT = 'N'.

5:   ITERM — INTEGER   *Input*

On entry: defines the termination criterion to be used. The possible choices are:

1   use the termination criterion defined in (2);
2   use the termination criterion defined in (3).

Suggested value: ITERM = 1.

Constraints: if METHOD = 'RGMRES', 'CGS' or 'BICGSTAB', ITERM = 1 or 2;

if METHOD = 'TFQMR', ITERM = 1.

6:   N — INTEGER   *Input*

On entry: the order n of the matrix A.

Constraint: N > 0.

7:   M — INTEGER   *Input*

On entry: if METHOD = 'RGMRES', M is the dimension m of the restart subspace.

If METHOD = 'BICGSTAB', M is the order ℓ of the polynomial Bi-CGSTAB method;

otherwise, M is not referenced.

Constraints:

    if METHOD = 'RGMRES', $0 < M \le \min(N, 50)$;
    if METHOD = 'BICGSTAB', $0 < M \le \min(N, 10)$.

8:   TOL — ***real***   *Input*

On entry: the tolerance τ for the termination criterion. If TOL ≤ 0.0, $\tau = \max(\sqrt{\epsilon}, \sqrt{n}\,\epsilon)$ is used, where ϵ is the ***machine precision***. Otherwise $\tau = \max(\text{TOL}, 10\epsilon, \sqrt{n}\,\epsilon)$ is used.

Constraint: TOL < 1.0.

9: MAXITN — INTEGER *Input*

On entry: the maximum number of iterations.

Constraint: MAXITN > 0.

10: ANORM — **real** *Input*

On entry: if ANORM > 0.0, the value of $\|A\|_p$ to be used in the termination criterion (2) (ITERM = 1).

If ANORM ≤ 0.0, ITERM = 1 and NORM = '1' or 'I', then $\|A\|_1 = \|A\|_\infty$ is estimated internally by F11BSF.

If ITERM = 2, then ANORM is not referenced.

Constraint: if ITERM = 1 and NORM = '2', then ANORM > 0.0.

11: SIGMAX — **real** *Input*

On entry: if ITERM = 2, the largest singular value σ_1 of the preconditioned iteration matrix; otherwise, SIGMAX is not referenced.

If SIGMAX ≤ 0.0, ITERM = 2 and METHOD = 'RGMRES', then the value of σ_1 will be estimated internally.

Constraint: if METHOD = 'CGS' or 'BICGSTAB' and ITERM = 2, then SIGMAX > 0.0.

12: MONIT — INTEGER *Input*

On entry: if MONIT > 0, the frequency at which a monitoring step is executed by F11BSF: if METHOD = 'CGS' or 'TFQMR', a monitoring step is executed every MONIT iterations; otherwise, a monitoring step is executed every MONIT super-iterations (groups of up to m or ℓ iterations for RGMRES or Bi-CGSTAB (ℓ), respectively).

There are some additional computational costs involved in monitoring the solution and residual vectors when the Bi-CGSTAB (ℓ) method is used with $\ell > 1$.

Constraint: MONIT ≤ MAXITN.

13: LWREQ — INTEGER *Output*

On exit: the minimum amount of workspace required by F11BSF. (See also Section 5 of the document for F11BSF.)

14: WORK (LWORK) — **complex** array *Output*

On exit: the workspace WORK is initialised by F11BRF. It must **not** be modified before calling the next routine in the suite, namely F11BSF.

15: LWORK — INTEGER *Input*

On entry: the dimension of the array WORK as declared in the (sub)program from which F11BRF was called.

Constraint: LWORK ≥ 120.

Note: although the minimum value of LWORK ensures the correct functioning of F11BRF, a larger value is required by the other routines in the suite, namely F11BSF and F11BTF. The required value is as follows:

Method	Requirements
RGMRES	LWORK $= 120 + n(m+3) + m(m+5) + 1$, where m is the dimension of the basis;
CGS	LWORK $= 120 + 7n$;
Bi-CGSTAB (ℓ)	LWORK $= 120 + (2n+\ell)(\ell+2) + p$, where ℓ is the order of the method;
TFQMR	LWORK $= 120 + 10n$;

where

$p = 2n$ if $\ell > 1$ and ITERM = 2 was supplied;

$p = n$ if $\ell > 1$ and a preconditioner is used or ITERM = 2 was supplied;

$p = 0$ otherwise.

16:    IFAIL — INTEGER   *Input/Output*

On entry: IFAIL must be set to 0, −1 or 1. For users not familiar with this parameter (described in Chapter P01) the recommended value is 0.

On exit: IFAIL = 0 unless the routine detects an error (see Section 6).

6   Errors and Warnings

If on entry IFAIL = 0 or −1, explanatory error messages are output on the current error message unit (as defined by X04AAF).

Errors detected by the routine:

IFAIL = −i

On entry, the ith argument had an illegal value.

IFAIL = 1

F11BRF has been called out of sequence.

7   Accuracy

Not applicable.

8   Further Comments

RGMRES can estimate internally the maximum singular value σ_1 of the iteration matrix, using $\sigma_1 \sim \|T\|_1$ where T is the upper triangular matrix obtained by QR factorization of the upper Hessenberg matrix generated by the Arnoldi process. The computational costs of this computation are negligible when compared to the overall costs.

Loss of orthogonality in the RGMRES method, or of bi-orthogonality in the Bi-CGSTAB (ℓ) method may degrade the solution and speed of convergence. For both methods, the algorithms employed include checks on the basis vectors so that the number of basis vectors used for a given super-iteration may be less than the value specified in the input parameter M. Also, the CGS, Bi-CGSTAB (ℓ) and TFQMR methods will restart automatically the computation from the last available iterates, when the stability of the solution process requires it.

Termination criterion (3), when available, involves only the residual (or norm of the residual) produced directly by the iteration process: this may differ from the norm of the true residual $\tilde{r}_k = b - Ax_k$, particularly when the norm of the residual is very small. Also, if the norm of the initial estimate of the solution is much larger than the norm of the exact solution, convergence can be achieved despite very large errors in the solution. On the other hand, termination criterion (3) is cheaper to use and inspects the progress of the actual iteration. Termination criterion (2) should be preferred in most cases, despite its slightly larger costs.

9   Example

The example solves an 8×8 nonsymmetric system of simultaneous linear equations using the bi-conjugate gradient stabilized method of order $\ell = 2$, where the coefficients matrix A has a random sparsity pattern. An incomplete LU preconditioner is used (routines F11DAF and F11DBF).

9.1 Example Text

Note. The listing of the example program presented below uses bold italicised terms to denote precision-dependent details. Please read the Users' Note for your implementation to check the interpretation of these terms. As explained in the Essential Introduction to this manual, the results produced may not be identical for all implementations.

```
*     F11BRF Example Program Text
*     Mark 19 Release. NAG Copyright 1999.
*     .. Parameters ..
      INTEGER          NIN, NOUT
      PARAMETER        (NIN=5,NOUT=6)
      INTEGER          NMAX, LA, LIWORK, LWORK
      PARAMETER        (NMAX=1000,LA=10000,LIWORK=7*NMAX+2,LWORK=6*NMAX)
*     .. Local Scalars ..
      real             ANORM, DTOL, SIGMAX, STPLHS, STPRHS, TOL
      INTEGER          I, IFAIL, IFAIL1, IFAILX, IREVCM, ITERM, ITN,
     +                 LFILL, LWREQ, M, MAXITN, MONIT, N, NNZ, NNZC,
     +                 NPIVM
      LOGICAL          LOOP
      CHARACTER        MILU, NORM, PRECON, PSTRAT, WEIGHT
      CHARACTER*8      METHOD
*     .. Local Arrays ..
      complex          A(LA), B(NMAX), WORK(LWORK), X(NMAX)
      real             WGT(NMAX)
      INTEGER          ICOL(LA), IDIAG(NMAX), IPIVP(NMAX), IPIVQ(NMAX),
     +                 IROW(LA), ISTR(NMAX+1), IWORK(LIWORK)
*     .. External Subroutines ..
      EXTERNAL         F11BRF, F11BSF, F11BTF, F11DNF, F11DPF, F11XNF
*     .. Executable Statements ..
      WRITE (NOUT,*) 'F11BRF Example Program Results'
*
*     Skip heading in data file
*
      READ (NIN,*)
      READ (NIN,*) N
      IF (N.LE.NMAX) THEN
*
*        Read or initialize the parameters for the iterative solver
*
         READ (NIN,*) METHOD
         READ (NIN,*) PRECON, NORM, WEIGHT, ITERM
         READ (NIN,*) M, TOL, MAXITN
         READ (NIN,*) MONIT
         ANORM = 0.0e0
         SIGMAX = 0.0e0
*
*        Read the parameters for the preconditioner
*
         READ (NIN,*) LFILL, DTOL
         READ (NIN,*) MILU, PSTRAT
*
*        Read the number of non-zero elements of the matrix A, then read
*        the non-zero elements
*
         READ (NIN,*) NNZ
         DO 20 I = 1, NNZ
            READ (NIN,*) A(I), IROW(I), ICOL(I)
   20    CONTINUE
*
*        Read right-hand side vector b and initial approximate
```

```
*         solution x
*
          READ (NIN,*) (B(I),I=1,N)
          READ (NIN,*) (X(I),I=1,N)
*
*         Calculate incomplete LU factorization
*
          IFAIL = 0
          CALL F11DNF(N,NNZ,A,LA,IROW,ICOL,LFILL,DTOL,PSTRAT,MILU,IPIVP,
     +               IPIVQ,ISTR,IDIAG,NNZC,NPIVM,IWORK,LIWORK,IFAIL)
*
*         Call F11BRF to initialize the solver
*
          IFAIL = 0
          CALL F11BRF(METHOD,PRECON,NORM,WEIGHT,ITERM,N,M,TOL,MAXITN,
     +               ANORM,SIGMAX,MONIT,LWREQ,WORK,LWORK,IFAIL)
*
*         Call repeatedly F11BSF to solve the equations
*         Note that the arrays B and X are overwritten
*
*         On final exit, X will contain the solution and B the residual
*         vector
*
          IFAIL = 0
          IREVCM = 0
          LOOP = .TRUE.
*
          LWREQ = LWORK
   40     CONTINUE
          CALL F11BSF(IREVCM,X,B,WGT,WORK,LWREQ,IFAIL)
          IF (IREVCM.EQ.-1) THEN
             IFAILX = -1
             CALL F11XNF('Transpose',N,NNZ,A,IROW,ICOL,'No checking',X,B,
     +               IFAILX)
          ELSE IF (IREVCM.EQ.1) THEN
             IFAILX = -1
             CALL F11XNF('No transpose',N,NNZ,A,IROW,ICOL,'No checking',
     +               X,B,IFAILX)
          ELSE IF (IREVCM.EQ.2) THEN
             IFAIL1 = -1
             CALL F11DPF('No transpose',N,A,LA,IROW,ICOL,IPIVP,IPIVQ,
     +               ISTR,IDIAG,'No checking',X,B,IFAILX)
          ELSE IF (IREVCM.EQ.3) THEN
             IFAIL1 = 0
             CALL F11BTF(ITN,STPLHS,STPRHS,ANORM,SIGMAX,WORK,LWREQ,
     +               IFAIL1)
             WRITE (NOUT,99999) ITN, STPLHS
             WRITE (NOUT,99998)
             WRITE (NOUT,99996) (X(I),I=1,N)
             WRITE (NOUT,99997)
             WRITE (NOUT,99996) (B(I),I=1,N)
          ELSE IF (IREVCM.EQ.4) THEN
             LOOP = .FALSE.
          END IF
          IF (LOOP) GO TO 40
*
*         Obtain information about the computation
*
```

```
            IFAIL1 = 0
            CALL F11BTF(ITN,STPLHS,STPRHS,ANORM,SIGMAX,WORK,LWREQ,IFAIL1)
    *
    *          Print the output data
    *
            WRITE (NOUT,99995)
            WRITE (NOUT,99994) 'Number of iterations for convergence:    ',
        +       ITN
            WRITE (NOUT,99993) 'Residual norm:                          ',
        +       STPLHS
            WRITE (NOUT,99993) 'Right-hand side of termination criterion:',
        +       STPRHS
            WRITE (NOUT,99993) '1-norm of matrix A:                     ',
        +       ANORM
    *
    *          Output x
    *
            WRITE (NOUT,99998)
            WRITE (NOUT,99996) (X(I),I=1,N)
            WRITE (NOUT,99997)
            WRITE (NOUT,99996) (B(I),I=1,N)
          END IF
          STOP
    *
    99999 FORMAT (/1X,'Monitoring at iteration no.',I4,/1X,1P,'residual no',
        +         'rm: ',e14.4)
    99998 FORMAT (/2X,'  Solution vector')
    99997 FORMAT (/2X,'  Residual vector')
    99996 FORMAT (1X,'(',1P,e16.4,',',1P,e16.4,')')
    99995 FORMAT (/1X,'Final Results')
    99994 FORMAT (1X,A,I4)
    99993 FORMAT (1X,A,1P,e14.4)
          END
```

9.2   Example Data

```
F11BRF Example Program Data
  8                     N
  'TFQMR'               METHOD
  'P'  '1'  'N'  1      PRECON, NORM, WEIGHT, ITERM
  1  1.0E-8   20        M, TOL, MAXITN
  2                     MONIT
  0  0.0                LFILL, DTOL
  'N'  'C'              MILU, PSTRAT
  24                    NNZ
( 2., 1.)   1    1
(-1., 1.)   1    4
( 1.,-3.)   1    8
( 4., 7.)   2    1
(-3., 0.)   2    2
( 2., 4.)   2    5
(-7.,-5.)   3    3
( 2., 1.)   3    6
( 3., 2.)   4    1
(-4., 2.)   4    3
( 0., 1.)   4    4
( 5.,-3.)   4    7
(-1., 2.)   5    2
```

```
( 8., 6.)   5   5
(-3.,-4.)   5   7
(-6.,-2.)   6   1
( 5.,-2.)   6   3
( 2., 0.)   6   6
( 0.,-5.)   7   3
(-1., 5.)   7   5
( 6., 2.)   7   7
(-1., 4.)   8   2
( 2., 0.)   8   6
( 3., 3.)   8   8      A(I), IROW(I), ICOL(I), I=1,...,NNZ
(  7., 11.)
(  1., 24.)
(-13.,-18.)
(-10.,  3.)
( 23., 14.)
( 17., -7.)
( 15., -3.)
( -3., 20.)            B(I), I=1,...,N
( 0., 0.)
( 0., 0.)
( 0., 0.)
( 0., 0.)
( 0., 0.)
( 0., 0.)
( 0., 0.)
( 0., 0.)              X(I), I=1,...,N
```

9.3   Example Results

```
F11BRF Example Program Results

Monitoring at iteration no.   2
residual norm:     8.2345E+01

    Solution vector
(      6.9055E-01,      1.4236E+00)
(      7.3931E-02,     -1.1880E+00)
(      1.4778E+00,      4.7846E-01)
(      5.6572E+00,     -3.0786E+00)
(      1.4243E+00,     -1.1246E+00)
(      1.0374E-01,      1.9740E+00)
(      4.4985E-01,     -1.2715E+00)
(      2.5704E+00,      1.7578E+00)

    Residual vector
(      1.7772E+00,      4.6797E+00)
(      1.0774E+00,      6.4600E+00)
(     -3.2812E+00,     -1.1314E+01)
(     -3.8698E+00,     -1.6438E+00)
(      8.9912E+00,      1.1100E+01)
(      9.7428E+00,     -4.6218E-01)
(      3.1668E+00,      2.8721E+00)
(     -1.0323E+01,      1.5837E+00)
```

```
    Final Results
    Number of iterations for convergence:        4
    Residual norm:                          1.2260E-11
    Right-hand side of termination criterion:   8.9100E-06
    1-norm of matrix A:                     2.7000E+01

       Solution vector
    (        1.0000E+00,        1.0000E+00)
    (        2.0000E+00,       -1.0000E+00)
    (        3.0000E+00,        1.0000E+00)
    (        4.0000E+00,       -1.0000E+00)
    (        3.0000E+00,       -1.0000E+00)
    (        2.0000E+00,        1.0000E+00)
    (        1.0000E+00,       -1.0000E+00)
    (       -1.0283E-14,        3.0000E+00)

       Residual vector
    (        3.2152E-13,       -7.7094E-13)
    (        2.2187E-12,       -5.4712E-13)
    (        4.9205E-13,        9.8055E-13)
    (        1.0552E-12,        2.2737E-13)
    (       -1.0090E-12,       -9.2015E-13)
    (       -9.0594E-13,        1.7586E-13)
    (       -2.5935E-13,        1.1902E-13)
    (        4.4587E-13,        1.8119E-12)
```