

NAG Fortran Library Routine Document

F11GBF

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

F11GBF is an iterative solver for a symmetric system of simultaneous linear equations; F11GBF is the second in a suite of three routines, where the first routine, F11GAF, must be called prior to F11GBF to setup the suite, and the third routine in the suite, F11GCF, can be used to return additional information about the computation.

These three routines are suitable for the solution of large sparse symmetric systems of equations.

2 Specification

```
SUBROUTINE F11GBF(IREVCM, U, V, WORK, LWORK, IFAIL)
  INTEGER          IREVCM, LWORK, IFAIL
  real            U(*), V(*), WORK(LWORK)
```

3 Description

F11GBF solves the symmetric system of linear simultaneous equations $Ax = b$ using either the preconditioned conjugate gradient method (Hestenes and Stiefel (1952), Golub and van Loan (1996), Barrett *et al.* (1994), Dias da Cunha and Hopkins (1994)) or a preconditioned Lanczos method based upon the algorithm SYMMLQ (Paige and Saunders (1975), Barrett *et al.* (1994)).

For a general description of the methods employed you are referred to Section 3 of the document for F11GAF.

F11GBF can solve the system after the first routine in the suite, F11GAF, has been called to initialise the computation and specify the method of solution. The third routine in the suite, F11GCF, can be used to return additional information generated by the computation during monitoring steps and after F11GBF has completed its tasks.

F11GBF uses **reverse communication**, i.e., F11GBF returns repeatedly to the calling program with the parameter IREVCM (see Section 5) set to specified values which require the calling program to carry out a specific task: either to compute the matrix-vector product $v = Au$; to solve the preconditioning equation $Mv = u$; to notify the completion of the computation; or, to allow the calling program to monitor the solution. Through the parameter IREVCM the calling program can cause immediate or tidy termination of the execution. On final exit, the last iterates of the solution and of the residual vectors of the original system of equations are returned.

Reverse communication has the following advantages.

1. Maximum flexibility in the representation and storage of sparse matrices: all matrix operations are performed outside the solver routine, thereby avoiding the need for a complicated interface with enough flexibility to cope with all types of storage schemes and sparsity patterns. This applies also to preconditioners.
2. Enhanced user interaction: the progress of the solution can be closely monitored by the user and tidy or immediate termination can be requested. This is useful, for example, when alternative termination criteria are to be employed or in case of failure of the external routines used to perform matrix operations.

4 References

- Hestenes M and Stiefel E (1952) Methods of conjugate gradients for solving linear systems *J. Res. Nat. Bur. Stand.* **49** 409–436
- Golub G H and van Loan C F (1996) *Matrix Computations* (3rd Edition) Johns Hopkins University Press, Baltimore
- 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
- 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, UK
- Paige C C and Saunders M A (1975) Solution of sparse indefinite systems of linear equations *SIAM J. Numer. Anal.* **12** 617–629
- 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

5 Parameters

Note: this routine uses **reverse communication**. Its use involves an initial entry, intermediate exits and re-entries, and a final exit, as indicated by the **parameter IREVCM**. Between intermediate exits and re-entries, **all parameters other than IREVCM and V must remain unchanged**.

1: IREVCM – INTEGER *Input/Output*

On initial entry: IREVCM = 0, otherwise an error condition will be raised.

On intermediate re-entry: IREVCM must either be unchanged from its previous exit value, or can have one of the following values.

- 5 Tidy termination: the computation will terminate at the end of the current iteration. Further reverse communication exits may occur depending on when the termination request is issued. F11GBF will then return with the termination code IREVCM = 4. Note that before calling F11GBF with IREVCM = 5 the calling program must have performed the tasks required by the value of IREVCM returned by the previous call to F11GBF, otherwise subsequently returned values may be invalid.
- 6 Immediate termination: F11GBF will return immediately with termination code IREVCM = 4 and with any useful information available. This includes the last iterate of the solution and, for conjugate gradient only, the last iterate of the residual vector. The residual vector is generally not available when the Lanczos method (SYMMLQ) is used. F11GBF will then return with the termination code IREVCM = 4.
Immediate termination may be useful, for example, when errors are detected during matrix-vector multiplication or during the solution of the preconditioning equation.

Changing IREVCM to any other value between calls will result in an error.

On intermediate exit: IREVCM has the following meanings

- 1 The calling program must compute the matrix-vector product $v = Au$, where u and v are stored in U and V, respectively;
- 2 The calling program must solve the preconditioning equation $Mv = u$, where u and v are stored in U and V, respectively;
- 3 Monitoring step: the solution and residual at the current iteration are returned in the arrays U and V, respectively. No action by the calling program is required. F11GCF can be called at this step to return additional information.

On final exit: IREVCM = 4: F11GBF has completed its tasks. The value of IFAIL determines whether the iteration has been successfully completed, errors have been detected or the calling program has requested termination.

Constraints: on initial entry, IREVCM = 0; on re-entry, either IREVCM must remain unchanged or be reset to 5 or 6.

2: U(*) – *real* array *Input/Output*

Note: the dimension of the array U must be at least n .

On initial entry: an initial estimate, x_0 , of the solution of the system of equations $Ax = b$.

On intermediate re-entry: U must remain unchanged.

On intermediate exit: the returned value of IREVCM determines the contents of U in the following way:

IREVCM = 1, 2

U holds the vector u on which the operation specified by IREVCM is to be carried out;

IREVCM = 3

U holds the current iterate of the solution vector.

On final exit: if IFAIL = 3 or < 0 , the array U is unchanged from the initial entry to F11GBF. If IFAIL = 1, the array U is unchanged from the last entry to F11GBF. Otherwise, U holds the last iterate of the solution of the system of equations, for all returned values of IFAIL.

3: V(*) – *real* array *Input/Output*

Note: the dimension of the array V must be at least n .

On initial entry: the right-hand side b of the system of equations $Ax = b$.

On intermediate re-entry: the returned value of IREVCM determines the contents of V in the following way:

IREVCM = 1, 2

V must store the vector v , the result of the operation specified by the value of IREVCM returned by the previous call to F11GBF;

IREVCM = 3

V must remain unchanged.

On intermediate exit: if IREVCM = 3, V holds the current iterate of the residual vector. Note that this is an approximation to the true residual vector. Otherwise, it does not contain any useful information.

On final exit: if IFAIL = 3 or < 0 , the array V is unchanged from the last entry to F11GBF. If IFAIL = 1, the array V is unchanged from the initial entry to F11GBF. If IFAIL = 0 or 2, the array V contains the true residual vector of the system of equations (see also Section 6); otherwise, V stores the last iterate of the residual vector unless the Lanczos method (SYMMLQ) was used and IFAIL ≥ 5 , in which case V is set to 0.0.

4: WORK(LWORK) – *real* array *Input/Output*

On initial entry: if user-supplied weights are used in the computation of the vector norms in the termination criterion (see Sections 3 and 5 of the document for F11GAF), these must be stored in WORK(1 : n).

Otherwise, WORK need not be initialised.

On intermediate re-entry: WORK must remain unchanged.

On final exit: if weights are used, WORK(1 : n) remains unchanged from the values supplied on initial entry.

5: LWORK – INTEGER *Input*

On initial entry: the dimension of the array WORK as declared in the (sub)program from which F11GBF is called. The required amount of workspace is as follows:

Method	Requirements
conjugate gradient	LWORK = $5n$
Lanczos (SYMMLQ)	LWORK = $6n$

to this must be added:

$2 \times (\text{MAXITS} + 1)$	if the largest singular value of the iteration matrix is estimated by F11GBF using bisection (see Sections 3, 5 and 8 of the document for F11GAF);
--------------------------------	--

n	if user-defined weights are used in the computation of vector norms for the termination criterion (see Sections 3 and 5 of the document for F11GAF).
-----	--

Constraint: LWORK \geq LWREQ, where LWREQ is returned by F11GAF.

6: IFAIL – INTEGER *Input/Output*

On entry: IFAIL must be set to 0, -1 or 1. Users who are unfamiliar with this parameter 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, for users 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

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

Errors or warnings detected by the routine:

IFAIL = -*i*

On entry, the *i*th argument had an illegal value.

IFAIL = 1

F11GBF has been called again after returning the termination code IREVCN = 4. No further computation has been carried out and all input data and data stored for access by F11GCF have remained unchanged.

IFAIL = 2

The required accuracy could not be obtained. However, F11GBF has terminated with reasonable accuracy: the last iterate of the residual satisfied the termination criterion but the exact residual $r = b - Ax$, did not. A small number of iterations have been carried out after the iterated residual satisfied the termination criterion, but were unable to improve on the accuracy. This error code usually implies that your problem has been fully and satisfactorily solved to within or close to the accuracy available on your system. Further iterations are unlikely to improve on this situation. You should call F11GCF to check the values of the left- and right-hand sides of the termination condition.

IFAIL = 3

F11GAF was either not called before calling F11GBF or it returned an error. The arguments U and V remain unchanged.

IFAIL = 4

The calling program requested a tidy termination before the solution had converged. The arrays U and V return the last iterates available of the solution and of the residual vector, respectively.

IFAIL = 5

The solution did not converge within the maximum number of iterations allowed. The arrays U and V return the last iterates available of the solution and of the residual vector, respectively.

IFAIL = 6

The preconditioner appears not to be positive-definite. It is likely that your results are meaningless: both methods require a positive-definite preconditioner (see also Section 3). However, the array U returns the last iterate of the solution, the array V returns the last iterate of the residual vector, for the conjugate gradient method only.

IFAIL = 7

The matrix of the coefficients appears not to be positive-definite (conjugate gradient method only). The arrays U and V return the last iterates of the solution and residual vector, respectively. However, you should be warned that the results returned can be in error.

IFAIL = 8

The calling program requested an immediate termination. However, the array U returns the last iterate of the solution, the array V returns the last iterate of the residual vector, for the conjugate gradient method only.

7 Accuracy

On completion, i.e., IREVCM = 4 on exit, the arrays U and V will return the solution and residual vectors, x_k and $r_k = b - Ax_k$, respectively, at the k th iteration, the last iteration performed, unless an immediate termination was requested and the Lanczos method (SYMMLQ) was used.

On successful completion, the termination criterion is satisfied to within the user-specified tolerance, as described in Section 3 of the document for F11GAF. The computed values of the left- and right-hand sides of the termination criterion selected can be obtained by a call to F11GCF.

8 Further Comments

The number of operations carried out by F11GBF for each iteration is likely to be principally determined by the computation of the matrix-vector products $v = Au$ and by the solution of the preconditioning equation $Mv = u$ in the calling program. Each of these operations is carried out once every iteration.

The number of the remaining operations in F11GBF for each iteration is approximately proportional to n . Note that the Lanczos method (SYMMLQ) requires a slightly larger number of operations than the conjugate gradient method.

The number of iterations required to achieve a prescribed accuracy cannot be easily determined at the onset, as it can depend dramatically on the conditioning and spectrum of the preconditioned matrix of the coefficients $\bar{A} = E^{-1}AE^{-T}$.

Additional matrix-vector products are required for the computation of $\|A\|_1$ or $\|A\|_\infty$, when this has not been supplied to F11GAF and is required by the termination criterion employed.

The number of operations required to compute $\sigma_1(\bar{A})$ is negligible for reasonable values of SIGTOL and MAXITS (see Sections 5 and 8 of the document for F11GAF).

If the termination criterion $\|r_k\|_p \leq \tau (\|b\|_p + \|A\|_p \|x_k\|_p)$ is used (see Section 3 of the document for F11GAF) and $\|x_0\| \gg \|x_k\|$, so that because of loss of significant digits the required accuracy could not be obtained, the iteration is restarted automatically at some suitable point: F11GBF sets $x_0 = x_k$ and the computation begins again. For particularly badly scaled problems, more than one restart may be necessary.

Naturally, restarting adds to computational costs: it is recommended that the iteration should start from a value x_0 which is as close to the true solution \tilde{x} as can be estimated. Otherwise, the iteration should start from $x_0 = 0$.

9 Example

See Section 9 of the document for F11GAF.
