

# F11BBFP

## NAG Parallel Library Routine Document

**Note:** before using this routine, please read the Users' Note for your implementation to check for implementation-dependent details. You are advised to enclose any calls to NAG Parallel Library routines between calls to Z01AAFP and Z01ABFP.

**Note:** you should read the the F11 Chapter Introduction before using this routine. In particular, some of the notation and terminology used in this document was introduced in Section 2.1 of the F11 Chapter Introduction.

### 1 Description

F11BBFP is an iterative solver for a real general (nonsymmetric) system of simultaneous linear equations; F11BBFP is the second in a suite of three routines, where the first routine, F11BAFP, must be called prior to F11BBFP to set-up the suite, and the third routine in the suite, F11BCFP, can be used to return additional information about the computation either at appropriate monitoring steps or after F11BBFP has completed its tasks.

These three routines are suitable for the solution of large sparse real general (nonsymmetric) systems of equations.

Three iterative methods are available to F11BBFP:

- restarted generalized minimum residual method (RGMRES) (Saad and Schultz [4]),
- conjugate gradient squared method (CGS) (Sonneveld [6]), or
- stabilized bi-conjugate gradient method of order  $\ell$  (Bi-CGSTAB( $\ell$ )) (van der Vorst [7], Sleijpen and Fokkema [5]).

For a general description of the methods employed the user is referred to Section 6.1 of the document for F11BAFP.

F11BBFP uses **reverse communication**, i.e., it returns repeatedly to the calling program with the parameter IREVCM (see Section 4) set to specified values which require the calling program to carry out one of the following tasks: compute the matrix-vector product  $v = Au$  or  $v = A^T u$ ; solve the preconditioning equation  $Mv = u$ ; allow the calling program to monitor the solution; or notify the completion of the computation.

### 2 Specification

```
SUBROUTINE F11BBFP(ICNTXT, IREVCM, U, V, WORK, LWORK, IFAIL)
  DOUBLE PRECISION  U(*), V(*), WORK(LWORK)
  INTEGER           ICNTXT, IREVCM, LWORK, IFAIL
```

### 3 Usage

#### 3.1 Definitions

The following definitions are used in describing the data distribution within this document:

- $m_p$  – the number of rows in the Library Grid.
- $n_p$  – the number of columns in the Library Grid.
- $n_l(i, j)$  – the number of elements of the distributed vectors stored locally on the processor at location  $\{i, j\}$  of the Library Grid.

#### 3.2 Global and Local Arguments

The following global **input** arguments must have the same value on entry to the routine on each processor and the global **output** arguments will have the same value on exit from the routine on each processor:

Global input arguments:        IREVCM, IFAIL

Global output arguments:      IREVCM, IFAIL

The remaining arguments are local.

### 3.3 Distribution Strategy

The vectors in F11BBFP may be distributed in two different ways, as specified by the input parameters of F11BAFP.

First, vectors in F11BBFP may be distributed across all processors in the Library Grid, with different processors holding different parts of the vector. In this case, the output vectors are distributed in the same way as the input vectors.

Second, on initial entry to F11BBFP, input vectors may be distributed along the first column or row of the Library Grid. Each processor in the first column or row then broadcasts the elements stored locally to all the other processors in the same row or column, respectively. Hence, if vectors are distributed by column or row, all processors in the same row or column, respectively, of the grid will hold copies of the same vector elements. The output vectors are distributed in the same way as the input vectors after the initial broadcast.

### 3.4 Related Routines

This is the second in a suite of three routines. The other two routines are:

F11BAFP: to set up the computation;

F11BCFP: to return additional information about the computation.

### 3.5 Requisites

F11BAFP must have been called before F11BBFP.

## 4 Arguments

**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: ICNTXT — INTEGER *Local Input*  
*On initial entry:* the Library context, usually returned by a call to the Library Grid initialization routine Z01AAFP.  
*On intermediate re-entry:* ICNTXT is not referenced. The value supplied on initial entry is used.
- 2: IREVCM — INTEGER *Global Input/Global 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. F11BBFP will then generate the information accessible via F11BCFP and return with the termination code IREVCM = 4 and IFAIL = 4. Note that before calling F11BBFP with IREVCM = 5 the calling program must have performed the tasks required by the value of IREVCM returned by the previous call to F11BBFP, otherwise subsequently returned values may be invalid.
  - 6 Immediate termination: F11BBFP will return immediately with termination code IREVCM = 4 and IFAIL = 8 and with any useful information available. This includes the iterate of the solution vector computed at the end of the last super-iteration performed. Immediate termination may be useful, for example, when errors are detected during a 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:* IREVCVM 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;
- 1 the calling program must compute the matrix–vector product  $v = A^T u$ , where  $u$  and  $v$  are stored in U and V, respectively. This value is returned only during the computation of  $\|A\|_1$  or  $\|A\|_\infty$  by Higham’s method (Higham [3]);
- 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.

*On final exit:* IREVCVM = 4: F11BBFP 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, IREVCVM = 0; on re-entry, IREVCVM must either remain unchanged or be reset to 5 or 6.

**3:** U(\*) — DOUBLE PRECISION array *Local Input/Local Output*

**Note:** the dimension of the array U must be at least  $n_l(i, j)$ , the number of vector elements stored locally, where  $\{i, j\}$  are the coordinates of the calling processor in the Library Grid. The distribution strategies used in F11BBFP are described in Section 3.3.

*On initial entry:*  $x_0$ , the initial estimate.

*On intermediate re-entry:* U must remain unchanged.

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

- |                    |  |
|--------------------|--|
| IREVCVM = –1, 1, 2 | U holds the vector $u$ on which the operations specified by IREVCVM are to be carried out; |
| IREVCVM = 3        | U holds the current iterate of the solution vector.  |

*On final exit:* U holds the last iterate of the solution vector.

**4:** V(\*) — DOUBLE PRECISION array *Local Input/Local Output*

**Note:** the dimension of the array V must be at least  $n_l(i, j)$ , the number of vector elements stored locally, where  $\{i, j\}$  are the coordinates of the calling processor in the Library Grid. The distribution strategies used in F11BBFP are described in Section 3.3.

*On initial entry:*  $b$ , the right-hand side.

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

- |                    |   |
|--------------------|---|
| IREVCVM = –1, 1, 2 | V must store the vector $v$ , the result of the operation specified by the value of IREVCVM returned by the previous call to F11BBFP; |
| IREVCVM = 3        | V must remain unchanged.  |

*On intermediate exit:* if IREVCVM = 3, V holds the current iterate of the residual vector, otherwise it does not contain any useful information.

*On final exit:* V contains the last iterate of the residual vector. The value of IFAIL indicates the success or failure of the solution process. Note that if an immediate termination request was issued, V does not contain any useful information; in this case  $V(1:n_l(i, j)) = 0.0$  is returned.

- 5: WORK(LWORK) — DOUBLE PRECISION array *Local Input/Local Output*

*On initial entry:* if user-supplied weights are used in the computation of the vector norms in the termination criterion (see Section 2 of the F11 Chapter Introduction), these must be stored in WORK(1:n<sub>l</sub>(i, j)).

*On intermediate re-entry:* WORK must remain unchanged.

*On final exit:* if weights are used, WORK(1:n<sub>l</sub>(i, j)) remains unchanged from the values supplied on initial entry.

- 6: LWORK — INTEGER *Local Input*

*On initial entry:* the size of the array WORK as declared in the (sub)program from which F11BBFP was called. The amount of workspace required is as follows:

<b>Method</b>	<b>Requirements</b>
Bi-CGSTAB( $\ell$ )	LWORK = 2 n <sub>l</sub> (i, j) ( $\ell + 2$ ) + $\ell(\ell + 2)$ + p + q, where $\ell$ is the order of the method;
CGS	LWORK = 7n <sub>l</sub> (i, j) + p
RGMRES	LWORK = 3n <sub>l</sub> (i, j) + m (m + n <sub>l</sub> (i, j) + 4) + 1 + p, where m is the dimension of the basis

where

$p = n_l(i, j)$ , if user-defined weights are used in the computation of vector norms for the termination criterion (see Section 4 of the document for F11BAFP);

$p = 0$ , otherwise.

$q = 2n_l(i, j)$ , if  $\ell > 1$  and ITERM = 2 was supplied to F11BAFP (see Section 4 of the document for F11BAFP);

$q = n_l(i, j)$ , if  $\ell > 1$  and a preconditioner is used or ITERM = 2 was supplied to F11BAFP (see Section 4 of the document for F11BAFP).

$q = 0$ , otherwise.

*Constraint:* LWORK  $\geq$  LWREQ, where LWREQ is returned by F11BAFP.

- 7: IFAIL — INTEGER *Global Input/Global Output*

The NAG Parallel Library provides a mechanism, via the routine Z02EAFP, to reduce the amount of parameter validation performed by this routine. For a full description refer to the Z02 Chapter Introduction.

*On initial entry:* IFAIL must be set to 0, 1 or -1. For users not familiar with this parameter, described in the Essential Introduction to the NAG Parallel Library) the recommended values are:

IFAIL = 0, if multigriding is **not** employed;

IFAIL = -1, if multigriding is employed.

IFAIL is stored internally by F11BBFP.

*On intermediate re-entry:* IFAIL is not referenced. The value supplied to F11BBFP on initial entry is used.

*On final exit:* IFAIL = 0, unless the routine detects an error (see Section 5).

## 5 Errors and Warnings

If on entry IFAIL = 0 or -1, explanatory error messages are output from the root processor (or processor {0,0} when the root processor is not available) on the current error message unit (as defined by X04AAF).

## 5.1 Full Error Checking Mode Only

IFAIL = -2000

The routine has been called with an invalid value of ICNTXT on one or more processors.

IFAIL = -1000

The logical processor grid and library mechanism (Library Grid) have not been correctly defined, see Z01AAFP.

IFAIL = - $i$

On entry, the  $i$ th argument had illegal value(s) on one or more processors. For global arguments, this may also be caused by the  $i$ th argument not having the same value on **all** processors (see also Section 3.2).

IFAIL = 1

F11BBFP has been called again after returning the termination code IREVCM = 4. No further computation has been carried out.

IFAIL = 3

F11BAFP was either not called before calling F11BBFP or returned an error.

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 = 8

The calling program requested an immediate termination. The array U returns the last iterate of the solution. Note that if an immediate termination request was issued, V does not contain any useful information; in this case  $V(1:n_l(i, j)) = 0.0$  is returned.

## 5.2 Any Error Checking Mode

IFAIL = 2

The required accuracy could not be obtained. However, F11BBFP has terminated with reasonable accuracy: an internal inexpensive termination criterion, based on quantities readily available during the iteration, was satisfied. However, the termination criterion required, which uses the explicitly computed residual  $r = b - Ax$ , could not be satisfied. A small number of iterations have been carried out after the internal criterion has been satisfied, but have been unable to improve on the accuracy.

IFAIL = 5

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

IFAIL = 6

Algorithmic breakdown. The last available iterates of the solution and residual are returned, although it is possible that they are completely inaccurate. It is an unlikely occurrence, probably due to extreme ill-conditioning. See expert advice.

## 6 Further Comments

### 6.1 Algorithmic Detail

If  $\|A\|_1$  or  $\|A\|_\infty$  are to be computed internally by Higham's method, the calling program must also provide the means to compute the matrix-vector product  $v = A^T u$ . The three iterative methods do not require the computation of  $v = A^T u$  otherwise. This corresponds to a returned value IREVCM = -1.

Further algorithmic details are considered in Section 2 of the F11 Chapter Introduction.

## 6.2 Parallelism Detail

Some parallelism details are considered in Section 2.8 of the F11 Chapter Introduction.

Global operations, excluding the matrix–vector products and preconditioning steps carried out in the calling program, involve **synchronization points**. In the case when the vectors are distributed across all the processor grid, the global operations involve communication between all the participating processors in the grid. However, when the vectors are distributed across the column or row, the global operations involve communications with processors on the same column or row of the grid. In the latter case different columns or rows can then act entirely in parallel, independently of each other.

## 6.3 Accuracy

On completion, the arrays U and V will return the distributed 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, in which case information about the residual vector is generally not available.

On successful completion, the termination criterion is satisfied within the user-specified tolerance, as described in Section 2.1 of the F11 Chapter Introduction and in the documentation of F11BAFP. In any case, the left- and right-hand side of the termination criterion selected can be returned by a call to F11BCFP.

## 6.4 Computational Costs

The number of operations carried out by F11BBFP for each iteration is likely to be determined primarily by the computation of the matrix–vector products  $v = Au$  and by the solution of the preconditioning equation  $Mv = u$  in the calling program.

The number of the remaining operations in F11BBFP for each iteration is approximately proportional to  $\max_{i,j} (n_l(i, j))$ . This is multiplied by a factor which depends linearly on the iteration count within the current super-iteration (see Section 2.1 of the F11 Chapter Introduction).

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  $\bar{A} = M^{-1}A$  (RGMRES and CGS methods) or  $\bar{A} = AM^{-1}$  (Bi-CGSTAB( $\ell$ ) method).

Additional matrix–vector products are required for the estimation of  $\|A\|_1$  or  $\|A\|_\infty$  when this is required by the termination criterion employed, unless the norm has been supplied to F11BAFP.

If the termination criterion  $\|r_k\|_p \leq \tau (\|b\|_p + \|A\|_p \|x_k\|_p)$  is used (see Section 6.1 of the document for F11BAFP) and  $\|x_0\| \gg \|x_k\|$ , then the required accuracy cannot be obtained due to loss of significant digits. The iteration is restarted automatically at some suitable point: F11BBFP sets  $x_0 = x_k$  and the computation begins again. For particularly badly scaled problems, more than one restart may be necessary. This does not apply to the RGMRES method which, by its nature, self-restarts every super-iteration. 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  $x_0 = 0$ .

## 7 References

- [1] 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
- [2] Golub G H and van Loan C F (1996) *Matrix Computations* Johns Hopkins University Press (3rd Edition), Baltimore
- [3] 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
- [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] Sleijpen G L G and Fokkema D R (1993) BiCGSTAB( $\ell$ ) for linear equations involving matrices with complex spectrum *ETNA* **1** 11–32
- [6] Sonneveld P (1989) CGS, a fast Lanczos-type solver for nonsymmetric linear systems *SIAM J. Sci. Statist. Comput.* **10** 36–52
- [7] 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

## 8 Example

See Section 8 of the document for F11BAFP.

---