

NAG Toolbox for MATLAB

f11be

1 Purpose

f11be is an iterative solver for a real general (nonsymmetric) system of simultaneous linear equations; f11be is the second in a suite of three functions, where the first function, f11bd, must be called prior to f11be to set up the suite, and the third function in the suite, f11bf, can be used to return additional information about the computation.

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

2 Syntax

```
[irevcn, u, v, work, ifail] = f11be(irevcn, u, v, wgt, work, 'lwork',  
lwork)
```

3 Description

f11be solves the general (nonsymmetric) system of linear simultaneous equations $Ax = b$ using one of four available methods: RGMRES, the preconditioned restarted generalized minimum residual method (see Saad and Schultz 1986); CGS, the preconditioned conjugate gradient squared method (see Sonneveld 1989); Bi-CGSTAB(ℓ), the bi-conjugate gradient stabilized method of order ℓ (see Van der Vorst 1989 and Sleijpen and Fokkema 1993); or TFQMR, the quasi-minimal transpose-free quasi-minimum residual method (see Freund and Nachtigal 1991 and Freund 1993).

For a general description of the methods employed you are referred to Section 3 of the document for f11bd.

f11be can solve the system after the first function in the suite, f11bd, has been called to initialize the computation and specify the method of solution. The third function in the suite, f11bf, can be used to return additional information generated by the computation, during monitoring steps and after f11be has completed its tasks.

f11be uses **reverse communication**, i.e., it returns repeatedly to the calling program with the parameter **irevcn** (see Section 5) 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 $v = A^T u$ (the four methods require the matrix transpose-vector product only if $\|A\|_1$ or $\|A\|_\infty$ is estimated internally by Higham's method (see Higham 1988));

solve the preconditioning equation $Mv = u$;

notify the completion of the computation;

allow the calling program to monitor the solution.

Through the parameter **irevcn** 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 function, 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 you 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 functions used to perform matrix operations.

4 References

- Freund R W 1993 A transpose-free quasi-minimal residual algorithm for non-Hermitian linear systems *SIAM J. Sci. Comput.* **14** 470–482
- Freund R W and Nachtigal N 1991 QMR: a Quasi-Minimal Residual Method for Non-Hermitian Linear Systems *Numer. Math.* **60** 315–339
- 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
- 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
- Sleijpen G L G and Fokkema D R 1993 BiCGSTAB(ℓ) for linear equations involving matrices with complex spectrum *ETNA* **1** 11–32
- Sonneveld P 1989 CGS, a fast Lanczos-type solver for nonsymmetric linear systems *SIAM J. Sci. Statist. Comput.* **10** 36–52
- 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

5 Parameters

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

5.1 Compulsory Input Parameters

1: **irevcn** – int32 scalar

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

On intermediate re-entry: 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. f11be will then return with the termination code **irevcn** = 4. Note that before calling f11be with **irevcn** = 5 the calling program must have performed the tasks required by the value of **irevcn** returned by the previous call to f11be, otherwise subsequently returned values may be invalid.
- 6 Immediate termination: f11be will return immediately with termination code **irevcn** = 4 and with any useful information available. This includes the last iterate of the solution. The residual vector is generally not available. f11be will then return with the termination code **irevcn** = 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 **irevcn** to any other value between calls will result in an error.

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

2: **u(*)** – double array

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: must remain unchanged.

3: **v(*) – double array**

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:

if **irevcm** = -1, 1 or 2, **v** must store the vector v , the result of the operation specified by the value of **irevcm** returned by the previous call to f11be;
if **irevcm** = 3, **v** must remain unchanged.

4: **wgt(*) – double array**

Note: the dimension of the array **wgt** must be at least $\max(1, n)$.

The user-supplied weights, if these are to be used in the computation of the vector norms in the termination criterion (see Sections 3 and 5 of the document for f11bd).

5: **work(lwork) – double array**

On initial entry: the workspace **work** as returned by f11bd (see also Section 5 of the document for f11bd).

On intermediate re-entry: must remain unchanged.

5.2 Optional Input Parameters

1: **lwork – int32 scalar**

Default: The dimension of the array **work**.

On initial entry: The required amount of workspace is as follows:

Method	Requirements
RGMRES	lwork = $100 + n(m + 3) + m(m + 5) + 1$, where m is the dimension of the basis.
CGS	lwork = $100 + 7n$.
Bi-CGSTAB(ℓ)	lwork = $100 + (2n + \ell)(\ell + 2) + p$, where ℓ is the order of the method.
TFQMR	lwork = $100 + 10n$,

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.

Constraint: **lwork** \geq **lwreq**, where **lwreq** is returned by f11bd.

5.3 Input Parameters Omitted from the MATLAB Interface

None.

5.4 Output Parameters

1: **irevcm – int32 scalar**

On intermediate exit: has the following meanings.

-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; RGMRES, CGS and Bi-CGSTAB(ℓ) methods return

irevcm = -1 only if the matrix norm $\|A\|_1$ or $\|A\|_\infty$ is estimated internally using Higham's method. This can only happen if **iterm** = 1 in f11bd.

- 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. f11bf can be called at this step to return additional information.

On final exit: **irevcm** = 4: f11be 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.

2: **u(*)** – double array

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

On intermediate exit: the returned value of **irevcm** determines the contents of **u** in the following way:

if **irevcm** = -1, 1 or 2, **u** holds the vector u on which the operation specified by **irevcm** is to be carried out;
if **irevcm** = 3, **u** holds the current iterate of the solution vector.

On final exit: if **ifail** = 3 or $-i$, the array **u** is unchanged from the initial entry to f11be. If **ifail** = 1, the array **u** is unchanged from the last entry to f11be. Otherwise, **u** holds the last available iterate of the solution of the system of equations, for all returned values of **ifail**.

3: **v(*)** – double array

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

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 $-i$, the array **v** is unchanged from the initial entry to f11be. If **ifail** = 1, the array **v** is unchanged from the last entry to f11be. 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 available iterate of the residual vector unless **ifail** = 8 is returned on last entry, in which case **v** is set to 0.0.

4: **work(lwork)** – double array

5: **ifail** – int32 scalar

0 unless the function detects an error (see Section 6).

6 Error Indicators and Warnings

Errors or warnings detected by the function:

ifail = $-i$

If **ifail** = $-i$, parameter i had an illegal value on entry. The parameters are numbered as follows:

1: **irevcm**, 2: **u**, 3: **v**, 4: **wgt**, 5: **work**, 6: **lwork**, 7: **ifail**.

ifail = 1

f11be 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 f11bf have remained unchanged.

ifail = 2

The required accuracy could not be obtained. However, f11be has terminated with reasonable accuracy: the last iterate of the residual satisfied the termination criterion but the exact residual $r = b - Ax$, did not. After the first occurrence of this situation, the iteration was restarted once, but f11be could not 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 f11bf to check the values of the left- and right-hand sides of the termination condition.

ifail = 3

f11bd was either not called before calling f11be 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

Algorithmic breakdown. The last available iterates of the solution and residuals are returned, although it is possible that they are completely inaccurate.

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 CGS and TFQMR methods only.

7 Accuracy

On completion, i.e., **irevcn** = 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.

On successful completion, the termination criterion is satisfied to within the user-specified tolerance, as described in Section 3 of the document for f11bd. The computed values of the left- and right-hand sides of the termination criterion selected can be obtained by a call to f11bf.

8 Further Comments

The number of operations carried out by f11be 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 f11be for each iteration is approximately proportional to n .

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

Additional matrix-vector products are required for the computation of $\|A\|_1$ or $\|A\|_\infty$, when this has not been supplied to f11bd and is required by the termination criterion employed.

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 f11bd) 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: f11be 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 own 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 $x_0 = 0$.

9 Example

```

method = 'BICGSTAB';
precon = 'P';
n = int32(8);
m = int32(1);
nnz = int32(24);
a = zeros(10000, 1);
a(1:24) = [2; -1; 1; 4; -3; 2; -7; 2; 3; -4; 5; 5; -1; 8; -3; -6; 5; 2; -
5; -1; 6; -1; 2; 3];
irow = zeros(10000, 1, 'int32');
irow(1:24) = [int32(1);
              int32(1);
              int32(2);
              int32(2);
              int32(2);
              int32(3);
              int32(3);
              int32(4);
              int32(4);
              int32(4);
              int32(4);
              int32(5);
              int32(5);
              int32(5);
              int32(6);
              int32(6);
              int32(6);
              int32(7);
              int32(7);
              int32(7);
              int32(8);
              int32(8);
              int32(8)];
icol = zeros(10000, 1, 'int32');
icol(1:24) = [int32(1);
              int32(4);
              int32(8);
              int32(1);
              int32(2);
              int32(5);
              int32(3);
              int32(6);
              int32(1);
              int32(3);
              int32(4);
              int32(7);
              int32(2);
              int32(5);
              int32(7);
              int32(1);
              int32(3);
              int32(3);
              int32(3);
              int32(3);
              int32(3);
              int32(3);
              int32(3)];

```

```

        int32(6);
        int32(3);
        int32(5);
        int32(7);
        int32(2);
        int32(6);
        int32(8)];
lfill = int32(0);
ipivp = zeros(8, 1, 'int32');
ipivq = zeros(8, 1, 'int32');
dtol = 0;
milu = 'N';
tol = 1e-08;
maxitn = int32(20);
anorm = 0;
sigmax = 0;
monit = int32(1);
lwork = int32(6000);
irevcm = int32(0);
u = zeros(8, 1);
v = [6;
      8;
     -9;
     46;
     17;
     21;
     22;
     34];
wgt = zeros(8, 1);
[a, irow, icol, ipivp, ipivq, istr, idiag, nnzc, npivm, ifail] = ...
    fllda(nnz, a, irow, icol, lfill, dtol, milu, ipivp, ipivq);
[lwreq, work, ifail] = ...
    flldb(method, precon, n, m, tol, maxitn, anorm, sigmax, monit,
lwork, 'norm_p', '1');
while (irevcm ~= 4)
    [irevcm, u, v, work, ifail] = f11be(irevcm, u, v, wgt, work);

    if (irevcm == -1)
        [v, ifail] = fl1xa('T', a(1:nnz), irow(1:nnz), icol(1:nnz), 'N', u);
    elseif (irevcm == 1)
        [v, ifail] = fl1xa('N', a(1:nnz), irow(1:nnz), icol(1:nnz), 'N', u);
    elseif (irevcm == 2)
        [v, ifail] = fl1db('N', a, irow, icol, ipivp, ipivq, istr, idiag,
'N', u);
        if (ifail ~= 0)
            irevcm = 6;
        end
    elseif (irevcm == 3)
        [itn, stplhs, stprhs, anorm, sigmax, ifail] = f11bf(work);
        fprintf('\nMonitoring at iteration number %d\nresidual norm:
%14.4e\n', itn, stplhs);
        fprintf('\n    Solution Vector    Residual Vector\n');
        for i = 1:n
            fprintf('%+16.4e %+16.4e\n', u(i), v(i));
        end
    end
end
% Get information about the computation
[itn, stplhs, stprhs, anorm, sigmax, ifail] = f11bf(work);
fprintf('\nNumber of iterations for convergence: %d\n', itn);
fprintf('Residual norm: %14.4e\n', stplhs);
fprintf('Right-hand side of termination criteria: %14.4e\n', stprhs);
fprintf('i-norm of matrix a: %14.4e\n', anorm);
fprintf('\n    Solution Vector    Residual Vector\n');
for i = 1:n
    fprintf('%+16.4e %+16.4e\n', u(i), v(i));
end
end

```

```

Monitoring at iteration number 1
residual norm:      1.4059e+02

```

Solution Vector	Residual Vector
-4.5858e+00	+1.5256e+01
+1.0154e+00	+2.6624e+01
-2.2234e+00	-8.7498e+00
+6.0097e+00	+1.8602e+01
+1.3827e+00	+8.2821e+00
-7.9070e+00	+2.0416e+01
+4.4270e-01	+9.6094e+00
+5.9248e+00	+3.3055e+01

Monitoring at iteration number 2
residual norm: 3.2742e+01

Solution Vector	Residual Vector
+4.1642e+00	-2.9585e+00
+4.9370e+00	-5.5523e+00
+4.8101e+00	+8.2070e-01
+5.4324e+00	-1.6828e+01
+5.8531e+00	+5.5975e-01
+1.1925e+01	-1.9150e+00
+8.4826e+00	+1.0081e+00
+6.0625e+00	-3.1004e+00

Number of iterations for convergence: 3
Residual norm: 1.0474e-08
Right-hand side of termination criteria: 5.5900e-06
i-norm of matrix a: 1.1000e+01

Solution Vector	Residual Vector
+1.0000e+00	-1.3686e-09
+2.0000e+00	-2.6364e-09
+3.0000e+00	+2.2691e-10
+4.0000e+00	-3.2517e-09
+5.0000e+00	+6.3660e-10
+6.0000e+00	-5.2943e-10
+7.0000e+00	+9.6704e-10
+8.0000e+00	-8.5717e-10