NAG Toolbox for MATLAB

f11gb

1 Purpose

f11gb is an iterative solver for a symmetric system of simultaneous linear equations; f11gb is the second in a suite of three functions, where the first function, f11ga, must be called prior to f11gb to setup the suite, and the third function in the suite, f11gc, can be used to return additional information about the computation.

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

2 Syntax

```
[irevcm, u, v, work, ifail] = f11gb(irevcm, u, v, work, 'n', n, 'lwork',
lwork)
```

3 Description

f11gb solves the symmetric system of linear simultaneous equations $Ax = b$ using either the preconditioned conjugate gradient method (see Hestenes and Stiefel 1952, Golub and Van Loan 1996, Barrett *et al.* 1994 and Dias da Cunha and Hopkins 1994) or a preconditioned Lanczos method based upon the algorithm SYMMLQ (see Paige and Saunders 1975 and Barrett *et al.* 1994).

For a general description of the methods employed you are referred to Section 3 of the document for f11ga.

f11gb can solve the system after the first function in the suite, f11ga, has been called to initialize the computation and specify the method of solution. The third function in the suite, f11gc, can be used to return additional information generated by the computation during monitoring steps and after f11gb has completed its tasks.

f11gb uses **reverse communication**, i.e., f11gb 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 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

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

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

Hestenes M and Stiefel E 1952 Methods of conjugate gradients for solving linear systems *J. Res. Nat. Bur. Stand.* **49** 409–436

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

Paige C C and Saunders M A 1975 Solution of sparse indefinite systems of linear equations *SIAM J. Numer. Anal.* **12** 617–629

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 IREVCM**. Between intermediate exits and re-entries, **all parameters other than irevcm and v must remain unchanged**.

5.1 Compulsory Input Parameters

1: **irevcm – int32 scalar**

On initial entry: **irevcm** $= 0$, 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. f11gb will then return with the termination code **irevcm** $= 4$. Note that before calling f11gb with **irevcm** $= 5$ the calling program must have performed the tasks required by the value of **irevcm** returned by the previous call to f11gb, otherwise subsequently returned values may be invalid.

6 Immediate termination: f11gb 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. f11gb 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.

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

2: **u(n) – 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(n) – 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:

irevcm $= 1$ 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 f11gb.

irevcm $= 3$

v must remain unchanged.

4: **work**(**lwork**) – **double array**

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 f11ga), these must be stored in **work**$(1:,n)$.

Otherwise, **work** need not be initialized.

On intermediate re-entry: must remain unchanged.

5.2 Optional Input Parameters

1: **n – int32 scalar**

Default: The dimension of the arrays **u**, **v**. (An error is raised if these dimensions are not equal.)

n, the order of the matrix A.

Constraint: **n** > 0.

2: **lwork – int32 scalar**

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

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

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

to this must be added:

$2 \times ($**maxits** $+ 1)$	if the largest singular value of the iteration matrix is estimated by f11gb using bisection (see Sections 3, 5 and 8 of the document for f11ga).
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 f11ga).

Constraint: **lwork** $\geq$ **lwreq**, where **lwreq** is returned by f11ga.

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 $v = Au$, 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. f11gc can be called at this step to return additional information.

On final exit: **irevcm** $= 4$: f11gb 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(n) – 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:

irevcm $= 1$ or 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 $-i$, the array **u** is unchanged from the initial entry to f11gb. If **ifail** $= 1$, the array **u** is unchanged from the last entry to f11gb. Otherwise, **u** holds the last iterate of the solution of the system of equations, for all returned values of **ifail**.

3: **v(n) – 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 last entry to f11gb. If **ifail** $= 1$, the array **v** is unchanged from the initial entry to f11gb. 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) – double array**

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

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: **n**, 3: **u**, 4: **v**, 5: **work**, 6: **lwork**, 7: **ifail**.

ifail $= 1$

f11gb has been called again after returning the termination code **irevcm** $= 4$. No further computation has been carried out and all input data and data stored for access by f11gc have remained unchanged.

ifail $= 2$

The required accuracy could not be obtained. However, f11gb 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 satisfactory solved to within or close to the accuracy available on your system. Further iterations are unlikely to improve on this situation. You should call f11gc to check the values of the left- and right-hand sides of the termination condition.

ifail $= 3$

f11ga was either not called before calling f11gb 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 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 kth 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 f11ga. The computed values of the left- and right-hand sides of the termination criterion selected can be obtained by a call to f11gc.

8 Further Comments

The number of operations carried out by f11gb 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 f11gb 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 $\bar{A} = E^{-1}AE^{-\mathrm{T}}$.

Additional matrix-vector products are required for the computation of $\|A\|_1$ or $\|A\|_\infty$, when this has not been supplied to f11ga 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 f11ga).

If the termination criterion $\|r_k\|_p \leq \tau\left(\|b\|_p + \|A\|_p, \|x_k\|_p\right)$ is used (see Section 3 of the document for f11ga) 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: f11gb 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

```
method = 'CG';
precon = 'P';
n = int32(7);
tol = 1e-06;
maxitn = int32(20);
anorm = 0;
sigmax = 0;
maxits = int32(7);
monit = int32(2);
lfill = int32(0);
dtol = 0;
mic = 'N';
dscale = 0;
pstrat = 'M';
ipiv = zeros(n, 1, 'int32');
nnz = int32(16);
a = zeros(100, 1);
a(1:nnz) = [4, 1, 5, 2, 2, 3, -1, 1, 4, 1, -2, 3, 2, -1, -2, 5];
irow = zeros(100, 1, 'int32');
irow(1:nnz) = [int32(1), int32(2), int32(2), int32(3), int32(4), ...
               int32(4), int32(5), int32(5), int32(5), int32(6), ...
                      int32(6),  int32(6),  int32(7),  int32(7),  int32(7),
int32(7)];
icol = zeros(100, 1, 'int32');
icol(1:nnz) = [int32(1), int32(1), int32(2), int32(3), int32(2), ...
               int32(4), int32(1), int32(4), int32(5), int32(2), ...
                      int32(5),  int32(6),  int32(1),  int32(2),  int32(3),
int32(7)];
v = [15, 18, -8, 21, 11, 10, 29];
u = zeros(7, 1);
irevcm = int32(0);

% Calculate incomplete Cholesky factorisation
[a, irow, icol, ipiv, istr, nnzc, npivm, ifail] = ...
```

```
    f11ja(nnz, a, irow, icol, lfill, dtol, mic, dscale, ipiv);

% Initialise the solver
[lwreq, ifail] = ...
     f11ga(method, precon, n, tol, maxitn, anorm, sigmax, maxits, monit,
'norm_p', '1', 'sigcmp', 's');

work = zeros(lwreq, 1);

while (irevcm ~= 4)
  [irevcm, u, v, work, ifail] = f11gb(irevcm, u, v, work);

  if (irevcm == 1)
    % Calculate v = Au
    [v, ifail] = f11xe(a, irow, icol, 'No Checking', u, 'nnz', nnz);
    if (ifail ~= 0)
      irevcm = int32(6);
    end
  elseif (irevcm == 2)
    % Solve Mv = u
    [v, ifail] = f11jb(a, irow, icol, ipiv, istr, 'No Checking', u);
    if (ifail ~= 0)
      irevcm = int32(6);
    end
  elseif (irevcm == 3)
    % Display monitoring information
    [itn, stplhs, stprhs, anorm, sigmax, its, sigerr, ifail] = f11gc();
    if (ifail == 0)
      fprintf('\nMonitoring at iteration no. %d\n', itn);
      fprintf('residual norm: %14.4f\n', stplhs);
      fprintf('  Solution Vector  Residual Vector\n');
      for i=1:n
        fprintf('%16.4f %16.4f\n', u(i), v(i));
      end
    end
  end
end

% Obtain information about the computation
[itn, stplhs, stprhs, anorm, sigmax, its, sigerr, ifail] = f11gc();

% Print the output data
fprintf('\nFinal results\n');
fprintf('Number 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('1-norm of matrix A:                        %14.4e\n', anorm);
fprintf('Largest singular value of A_bar:          %14.4e\n', sigmax);
fprintf('  Solution Vector  Residual Vector\n');
for i=1:n
  fprintf('%16.4f %16.4f\n', u(i), v(i));
end
```

```
Monitoring at iteration no. 2
residual norm:         1.9938
  Solution Vector  Residual Vector
          0.9632          -0.2296
          1.9934           0.2225
          3.0583           0.0958
          4.1453          -0.2515
          4.8289          -0.1716
          5.6630           0.6753
          7.1062          -0.3474

Monitoring at iteration no. 4
residual norm:         0.0067
  Solution Vector  Residual Vector
          0.9994          -0.0011
          2.0011          -0.0025
          3.0008          -0.0000
```

```
          3.9996          0.0000
          4.9991          0.0021
          5.9993         -0.0009
          7.0007          0.0001
Final results
Number of iterations for convergence:        5
Residual norm:                               2.0428e-14
Right-hand side of termination criteria:     3.9200e-04
1-norm of matrix A:                          1.0000e+01
Largest singular value of A_bar:             1.3596e+00
Solution Vector  Residual Vector
1.0000           0.0000
          2.0000          0.0000
          3.0000         -0.0000
          4.0000         -0.0000
          5.0000         -0.0000
          6.0000          0.0000
          7.0000          0.0000
```