# NAG Library Routine Document <br> F11GDF 

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

F11GDF is a setup routine, the first in a suite of three routines for the iterative solution of a symmetric system of simultaneous linear equations. F11GDF must be called before the iterative solver, F11GEF. The third routine in the suite, F11GFF, 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 F11GDF (METHOD, PRECON, SIGCMP, NORM, WEIGHT, ITERM, N, TOL, &
    MAXITN, ANORM, SIGMAX, SIGTOL, MAXITS, MONIT, LWREQ, &
    WORK, LWORK, IFAIL)
INTEGER ITERM, N, MAXITN, MAXITS, MONIT, LWREQ, LWORK, IFAIL
REAL (KIND=nag_wp) TOL, ANORM, SIGMAX, SIGTOL, WORK(LWORK)
CHARACTER(*) METHOD
CHARACTER(1) PRECON, SIGCMP, NORM, WEIGHT
```


## 3 Description

The suite consisting of the routines F11GDF, F11GEF and F11GFF is designed to solve the symmetric system of simultaneous linear equations $A x=b$ of order $n$, where $n$ is large and the matrix of the coefficients $A$ is sparse.

F11GDF is a setup routine which must be called before F11GEF, the iterative solver. The third routine in the suite, F11GFF can be used to return additional information about the computation. One of the following methods can be used:

1. Conjugate Gradient Method (CG)

For this method (see Hestenes and Stiefel (1952), Golub and Van Loan (1996), Barrett et al. (1994) and Dias da Cunha and Hopkins (1994)), the matrix $A$ should ideally be positive definite. The application of the Conjugate Gradient method to indefinite matrices may lead to failure or to lack of convergence.

## 2. Lanczos Method (SYMMLQ)

This method, based upon the algorithm SYMMLQ (see Paige and Saunders (1975) and Barrett et al. (1994)), is suitable for both positive definite and indefinite matrices. It is more robust than the Conjugate Gradient method but less efficient when $A$ is positive definite.

## 3. Minimum Residual Method (MINRES)

This method may be used when the matrix is indefinite. It seeks to reduce the norm of the residual at each iteration and often takes fewer iterations than the other methods. It does however require slightly more memory.
The CG and SYMMLQ methods start from the residual $r_{0}=b-A x_{0}$, where $x_{0}$ is an initial estimate for the solution (often $x_{0}=0$ ), and generate an orthogonal basis for the Krylov subspace span $\left\{A^{k} r_{0}\right\}$, for $k=0,1, \ldots$, by means of three-term recurrence relations (see Golub and Van Loan (1996)). A sequence of symmetric tridiagonal matrices $\left\{T_{k}\right\}$ is also generated. Here and in the following, the index $k$ denotes the iteration count. The resulting symmetric tridiagonal systems of equations are usually more easily solved than the original problem. A sequence of solution iterates $\left\{x_{k}\right\}$ is thus generated such that the
sequence of the norms of the residuals $\left\{\left\|r_{k}\right\|\right\}$ converges to a required tolerance. Note that, in general, the convergence is not monotonic.

In exact arithmetic, after $n$ iterations, this process is equivalent to an orthogonal reduction of $A$ to symmetric tridiagonal form, $T_{n}=Q^{\mathrm{T}} A Q$; the solution $x_{n}$ would thus achieve exact convergence. In finite-precision arithmetic, cancellation and round-off errors accumulate causing loss of orthogonality. These methods must therefore be viewed as genuinely iterative methods, able to converge to a solution within a prescribed tolerance.
The orthogonal basis is not formed explicitly in either method. The basic difference between the Conjugate Gradient and Lanczos methods lies in the method of solution of the resulting symmetric tridiagonal systems of equations: the conjugate gradient method is equivalent to carrying out an $L D L^{\mathrm{T}}$ (Cholesky) factorization whereas the Lanczos method (SYMMLQ) uses an $L Q$ factorization.
Faster convergence for all the methods can be achieved using a preconditioner (see Golub and Van Loan (1996) and Barrett et al. (1994)). A preconditioner maps the original system of equations onto a different system, say

$$
\begin{equation*}
\bar{A} \bar{x}=\bar{b}, \tag{1}
\end{equation*}
$$

with, hopefully, better characteristics with respect to its speed of convergence: for example, the condition number of the matrix of the coefficients can be improved or eigenvalues in its spectrum can be made to coalesce. An orthogonal basis for the Krylov subspace $\operatorname{span}\left\{\bar{A}^{k} \bar{r}_{0}\right\}$, 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 preconditioner must be symmetric and positive definite, i.e., representable by $M=E E^{\mathrm{T}}$, where $M$ is nonsingular, and such that $\bar{A}=E^{-1} A E^{-\mathrm{T}} \sim I_{n}$ in (1), where $I_{n}$ is the identity matrix of order $n$. Also, we can define $\bar{r}=E^{-1} r$ and $\bar{x}=E^{\mathrm{T}} x$. These are formal definitions, used only in the design of the algorithms; in practice, only the means to compute the matrix-vector products $v=A u$ and to solve the preconditioning equations $M v=u$ are required, that is, explicit information about $M, E$ or their inverses is not required at any stage.
The first termination criterion

$$
\begin{equation*}
\left\|r_{k}\right\|_{p} \leq \tau\left(\|b\|_{p}+\|A\|_{p} \times\left\|x_{k}\right\|_{p}\right) \tag{2}
\end{equation*}
$$

is available for both conjugate gradient and Lanczos (SYMMLQ) methods. In (2), $p=1, \infty$ or 2 and $\tau$ denotes a user-specified tolerance subject to $\max (10, \sqrt{n}) \epsilon \leq \tau<1$, where $\epsilon$ is the machine precision. Facilities are provided for the estimation of the norm of the matrix of the coefficients $\|A\|_{1}=\|A\|_{\infty}$, when this is not known in advance, used in (2), by applying Higham's method (see Higham (1988)). 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 could be used, but any such criteria would require information about the condition number $\kappa(A)$ which is not easily obtainable.
The second termination criterion

$$
\begin{equation*}
\left\|\bar{r}_{k}\right\|_{2} \leq \tau \max \left(1.0,\|b\|_{2} /\left\|r_{0}\right\|_{2}\right)\left(\left\|\bar{r}_{0}\right\|_{2}+\sigma_{1}(\bar{A}) \times\left\|\Delta \bar{x}_{k}\right\|_{2}\right) \tag{3}
\end{equation*}
$$

is available only for the Lanczos method (SYMMLQ). 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). When $\sigma_{1}(\bar{A})$ is not supplied, facilities are provided for its estimation by $\sigma_{1}(\bar{A}) \sim \max _{k} \sigma_{1}\left(T_{k}\right)$. The interlacing property $\sigma_{1}\left(T_{k-1}\right) \leq \sigma_{1}\left(T_{k}\right)$ and Gerschgorin's theorem provide lower and upper bounds from which $\sigma_{1}\left(T_{k}\right)$ can be easily computed by bisection. Alternatively, the less expensive estimate
$\sigma_{1}(\bar{A}) \sim \max _{k}\left\|T_{k}\right\|_{1}$ can be used, where $\sigma_{1}(\bar{A}) \leq\left\|T_{k}\right\|_{1}$ by Gerschgorin's theorem. Note that only order of magnitude estimates are required by the termination criterion.
Termination criterion (2) is the recommended choice, despite its (small) additional costs per iteration when using the Lanczos method (SYMMLQ). Also, if the norm of the initial estimate is much larger than the norm of the solution, that is, if $\left\|x_{0}\right\| \gg\|x\|$, 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., $\|v\|_{p}^{(w)}=\left\|v^{(w)}\right\|_{p}$, where $\left(v^{(w)}\right)_{i}=w_{i} v_{i}$, for $i=1,2, \ldots, n$. Note that the use of weights increases the computational costs.
The MINRES algorithm terminates when the norm of the residual of the preconditioned system $F$, $\|F\|_{2} \leq \tau \times\|\bar{A}\|_{2} \times\left\|x_{k}\right\|_{2}$, where $\bar{A}$ is the preconditioned matrix.
The termination criteria discussed are not robust in the presence of a non-trivial nullspace of $A$, i.e., when $A$ is singular. It is then possible for $\left\|x_{k}\right\|_{p}$ to grow without limit, spuriously satisfying the termination criterion. If singularity is suspected, more robust routines can be found in Chapter E04.
The sequence of calls to the routines comprising the suite is enforced: first, the setup routine F11GDF must be called, followed by the solver F11GEF. The diagnostic routine F11GFF can be called either when F11GEF is carrying out a monitoring step or after F11GEF has completed its tasks. Incorrect sequencing will raise an error condition.

## 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 Arguments

1: METHOD - CHARACTER(*)
On entry: the iterative method to be used.

## METHOD = 'CG'

Conjugate gradient method (CG).
METHOD = 'SYMMLQ'
Lanczos method (SYMMLQ).

## METHOD = 'MINRES'

Minimum residual method (MINRES).
Constraint: METHOD = 'CG', 'SYMMLQ' or 'MINRES'.

PRECON - CHARACTER(1)
On entry: determines whether preconditioning is used.
PRECON $=$ ' N '
No preconditioning.
PRECON $=$ ' $\mathrm{P}^{\prime}$
Preconditioning.
Constraint: PRECON $=$ ' N ' or ' P '.
3: SIGCMP - CHARACTER(1)
Input
On entry: determines whether an estimate of $\sigma_{1}(\bar{A})=\left\|E^{-1} A E^{-T}\right\|_{2}$, the largest singular value of the preconditioned matrix of the coefficients, is to be computed using the bisection method on the sequence of tridiagonal matrices $\left\{T_{k}\right\}$ generated during the iteration. Note that $\bar{A}=A$ when a preconditioner is not used.

If SIGMAX $>0.0$ (see below), i.e., when $\sigma_{1}(\bar{A})$ is supplied, the value of SIGCMP is ignored.
SIGCMP = 'S'
$\sigma_{1}(\bar{A})$ is to be computed using the bisection method.
SIGCMP $={ }^{\prime} \mathrm{N}^{\prime}$
The bisection method is not used.
If the termination criterion (3) is used, requiring $\sigma_{1}(\bar{A})$, an inexpensive estimate is computed and used (see Section 3).
It is not used if METHOD $=$ 'MINRES'.
Suggested value: SIGCMP $=$ ' N '.
Constraint: SIGCMP $=$ 'S' or 'N'.
NORM - CHARACTER(1)
Input
On entry: if METHOD = 'CG' or 'SYMMLQ', NORM defines the matrix and vector norm to be used in the termination criteria.

NORM $=1$ '
Use the $l_{1}$ norm.
NORM = 'I'
Use the $l_{\infty}$ norm.
NORM $=$ ' 2 '
Use the $l_{2}$ norm.
It has no effect if METHOD = 'MINRES'.
Suggested value:

> if ITERM $=1$, NORM $=$ 'I';
> if ITERM $=2$, NORM $=$ ' 2 '.

Constraints:

> if $\operatorname{ITERM}=1$, NORM $=$ ' 1 ', 'I' or ' 2 ';
> if $\operatorname{ITERM}=2$, NORM $=$ ' 2 '.

5: WEIGHT - CHARACTER(1)
Input
On entry: specifies whether a vector $w$ of user-supplied weights is to be used in the vector norms used in the computation of termination criterion (2) (ITERM $=1$ ): $\|v\|_{p}^{(w)}=\left\|v^{(w)}\right\|_{p}$, 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 argument NORM. Note that weights cannot be used when ITERM $=2$, i.e., when criterion (3) is used.

WEIGHT $=$ ' W '
User-supplied weights are to be used and must be supplied on initial entry to F11GEF.
WEIGHT = ' $\mathrm{N}^{\prime}$
All weights are implicitly set equal to one. Weights do not need to be supplied on initial entry to F11GEF.
It has no effect if METHOD = 'MINRES'.
Suggested value: WEIGHT $=$ ' N '.
Constraints:
if $\operatorname{ITERM}=1$, WEIGHT $=$ ' W ' or ' N ';
if $\operatorname{ITERM}=2$, WEIGHT $=$ ' N '.
6: ITERM - INTEGER
Input
On entry: defines the termination criterion to be used.
ITERM $=1$
Use the termination criterion defined in (2) (both conjugate gradient and Lanczos (SYMMLQ) methods).

## ITERM $=2$

Use the termination criterion defined in (3) (Lanczos method (SYMMLQ) only).
It has no effect if METHOD $=$ 'MINRES'.
Suggested value: ITERM $=1$.

## Constraints:

```
    if METHOD = 'CG', ITERM = 1;
    if METHOD = 'SYMMLQ', ITERM = 1 or 2.
```

7: N - INTEGER Input
On entry: $n$, the order of the matrix $A$.
Constraint: $\mathrm{N}>0$.

8: $\quad$ TOL - REAL (KIND=nag_wp)
Input
On entry: the tolerance $\tau$ for the termination criterion.
If TOL $\leq 0.0, \tau=\max (\sqrt{\epsilon}, \sqrt{n} \epsilon)$ is used, where $\epsilon$ is the machine precision.
Otherwise $\tau=\max (\mathrm{TOL}, 10 \epsilon, \sqrt{n} \epsilon)$ is used.
Constraint: TOL $<1.0$.
9: MAXITN - INTEGER
On entry: the maximum number of iterations.
Constraint: MAXITN $>0$.
10: $\quad$ ANORM - REAL (KIND=nag_wp)
Input
On entry: if ANORM $>0.0$, the value of $\|A\|_{p}$ to be used in the termination criterion (2) (ITERM = 1).
If $\operatorname{ANORM} \leq 0.0$, ITERM $=1$ and $\operatorname{NORM}=$ ' 1 ' or 'I', then $\|A\|_{1}=\|A\|_{\infty}$ is estimated internally by F11GEF.
If $\operatorname{ITERM}=2$, then $A N O R M$ is not referenced.

It has no effect if METHOD = 'MINRES'.
Constraint: if ITERM $=1$ and NORM $=2$, ANORM $>0.0$.

11: SIGMAX - REAL (KIND=nag_wp)
Input
On entry: if SIGMAX $>0.0$, the value of $\sigma_{1}(\bar{A})=\left\|E^{-1} A E^{-\mathrm{T}}\right\|_{2}$.
If SIGMAX $\leq 0.0, \sigma_{1}(\bar{A})$ is estimated by F11GEF when either SIGCMP $=$ ' S ' or termination criterion (3) (ITERM $=2$ ) is employed, though it will be used only in the latter case.
Otherwise, or if METHOD = 'MINRES', SIGMAX is not referenced.
12: $\operatorname{SIGTOL}-\operatorname{REAL}(\mathrm{KIND}=$ nag_wp $)$
Input
On entry: the tolerance used in assessing the convergence of the estimate of $\sigma_{1}(\bar{A})=\|\bar{A}\|_{2}$ when the bisection method is used.

If SIGTOL $\leq 0.0$, the default value $\operatorname{SIGTOL}=0.01$ is used. The actual value used is $\max ($ SIGTOL,$\epsilon)$.
If SIGCMP $=$ ' N ' or SIGMAX $>0.0$, then SIGTOL is not referenced.
It has no effect if METHOD $=$ 'MINRES'.
Suggested value: SIGTOL $=0.01$ should be sufficient in most cases.
Constraint: if SIGCMP $=$ 'S' and SIGMAX $\leq 0.0$, SIGTOL $<1.0$.
13: MAXITS - INTEGER
Input
On entry: the maximum iteration number $k=$ MAXITS for which $\sigma_{1}\left(T_{k}\right)$ is computed by bisection (see also Section 3). If SIGCMP $=$ 'N' or SIGMAX $>0.0$, or if METHOD $=$ 'MINRES', then MAXITS is not referenced.
Suggested value: MAXITS $=\min (10, n)$ when SIGTOL is of the order of its default value (0.01).
Constraint: if SIGCMP $=$ 'S' and SIGMAX $\leq 0.0,1 \leq$ MAXITS $\leq$ MAXITN.
14: MONIT - INTEGER
Input
On entry: if MONIT $>0$, the frequency at which a monitoring step is executed by F11GEF: the current solution and residual iterates will be returned by F11GEF and a call to F11GFF made possible every MONIT iterations, starting from the (MONIT)th. Otherwise, no monitoring takes place.

There are some additional computational costs involved in monitoring the solution and residual vectors when the Lanczos method (SYMMLQ) is used.

Constraint: MONIT $\leq$ MAXITN.
15: LWREQ - INTEGER
Output
On exit: the minimum amount of workspace required by F11GEF. (See also Section 5 in F11GEF.)

16: $\quad$ WORK(LWORK) - REAL (KIND=nag_wp) array
On exit: the array WORK is initialized by F11GDF. It must not be modified before calling the next routine in the suite, namely F11GEF.

17: LWORK - INTEGER
Input
On entry: the dimension of the array WORK as declared in the (sub)program from which F11GDF is called.

Constraint: LWORK $\geq 120$.

Note: although the minimum value of LWORK ensures the correct functioning of F11GDF, a larger value is required by the other routines in the suite, namely F11GEF and F11GFF. The required value is as follows:

## Method Requirements

CG $\quad$ LWORK $=120+5 n+p$.
SYMMLQ $\quad$ LWORK $=120+6 n+p$,
MINRES $\quad$ LWORK $=120+9 n$,
where
$p=2 *($ MAXITS +1$)$, when an estimate of $\sigma_{1}(A)$ (SIGMAX) is computed;
$p=0$, otherwise.
18: IFAIL - INTEGER

## Input/Output

On entry: IFAIL must be set to $0,-1$ or 1 . If you are unfamiliar with this argument you should refer to Section 3.4 in How to Use the NAG Library and its Documentation for details.

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, if you are not familiar with this argument, the recommended value is 0 . When the value -1 or 1 is used it is essential to test the value of IFAIL on exit.

On exit: IFAIL $=0$ unless the routine detects an error or a warning has been flagged (see Section 6).

## 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 $=1$
F11GDF has been called out of sequence.
IFAIL $=-99$
An unexpected error has been triggered by this routine. Please contact NAG.
See Section 3.9 in How to Use the NAG Library and its Documentation for further information.
IFAIL $=-399$
Your licence key may have expired or may not have been installed correctly.
See Section 3.8 in How to Use the NAG Library and its Documentation for further information.
IFAIL $=-999$
Dynamic memory allocation failed.
See Section 3.7 in How to Use the NAG Library and its Documentation for further information.

## 7 Accuracy

Not applicable.

## 8 Parallelism and Performance

F11GDF is not threaded in any implementation.

## 9 Further Comments

When $\sigma_{1}(\bar{A})$ is not supplied (SIGMAX $\leq 0.0$ ) but it is required, it is estimated by F11GEF using either of the two methods described in Section 3, as specified by the argument SIGCMP. In particular, if SIGCMP $=$ 'S', then the computation of $\sigma_{1}(\bar{A})$ is deemed to have converged when the differences between three successive values of $\sigma_{1}\left(T_{k}\right)$ differ, in a relative sense, by less than the tolerance SIGTOL, i.e., when

$$
\max \left(\frac{\left|\sigma_{1}^{(k)}-\sigma_{1}^{(k-1)}\right|}{\sigma_{1}^{(k)}}, \frac{\left|\sigma_{1}^{(k)}-\sigma_{1}^{(k-2)}\right|}{\sigma_{1}^{(k)}}\right) \leq \text { SIGTOL }
$$

The computation of $\sigma_{1}(\bar{A})$ is also terminated when the iteration count exceeds the maximum value allowed, i.e., $k \geq$ MAXITS.

Bisection is increasingly expensive with increasing iteration count. A reasonably large value of SIGTOL, of the order of the suggested value, is recommended and an excessive value of MAXITS should be avoided. Under these conditions, $\sigma_{1}(\bar{A})$ usually converges within very few iterations.

## 10 Example

This example solves a symmetric system of simultaneous linear equations using the conjugate gradient method, where the matrix of the coefficients $A$, has a random sparsity pattern. An incomplete Cholesky preconditioner is used (F11JAF and F11JBF).

### 10.1 Program Text

```
    Program fllgdfe
    F11GDF Example Program Text
    Mark 26 Release. NAG Copyright 2016.
    .. Use Statements ..
    Use nag_library, Only: fllgdf, fllgef, fllgff, flljaf, flljbf, fllxef, &
                    nag_wp
    .. Implicit None Statement ..
    Implicit None
! .. Parameters ..
    Integer, Parameter :: nin = 5, nout = 6
! .. Local Scalars ..
    Real (Kind=nag_wp) :: anorm, dscale, dtol, sigerr, sigmax, &
    sigtol, stplhs, stprhs, tol
    Integer :: i, ifail, ifaill, irevcm, iterm, &
                                    itn, its, la, lfill, liwork, lwork, &
                                    lwreq, maxitn, maxits, monit, n, &
                                    nnz, nnzc, npivm
    Character (6) :: method
    Character (1) :: mic, norm, precon, pstrat, sigcmp, &
                                    weight
! .. Local Arrays ..
    Real (Kind=nag_wp), Allocatable :: a(:), b(:), wgt(:), work(:), x(:)
    Integer, Allocatable :: icol(:), ipiv(:), irow(:), istr(:), &
                                    iwork(:)
! .. Executable Statements ..
    Write (nout,*) 'F11GDF Example Program Results'
    Skip heading in data file
    Read (nin,*)
```

```
Read (nin,*) n
Read (nin,*) nnz
la = 2*nnz
liwork = 2*la + 7*n + 1
lwork = 120
Allocate (a(la),b(n),wgt(n),work(lwork),x(n),icol(la),ipiv(n),irow(la), &
    istr(n+1),iwork(liwork))
Read or initialize the parameters for the iterative solver
Read (nin,*) method
Read (nin,*) precon, sigcmp, norm, weight, iterm
Read (nin,*) tol, maxitn
Read (nin,*) monit
anorm = O.OEO_nag_wp
sigmax = 0.0EO_nag_wp
sigtol = 1.OE-2_nag_wp
maxits = n
! Read the parameters for the preconditioner
Read (nin,*) lfill, dtol
Read (nin,*) mic, dscale
Read (nin,*) pstrat
Read the nonzero elements of the matrix A
Do i = 1, nnz
    Read (nin,*) a(i), irow(i), icol(i)
End Do
Read right-hand side vector b and initial approximate solution x
Read (nin,*) b(1:n)
Read (nin,*) x(1:n)
If (method=='CG') Then
    Write (nout,99999)
Else If (method=='SYMMLQ') Then
    Write (nout,99998)
Else If (method=='MINRES') Then
    Write (nout,99997)
End If
Calculate incomplete Cholesky factorization
ifail: behaviour on error exit
                =O for hard exit, =1 for quiet-soft, =-1 for noisy-soft
ifail = 0
Call flljaf(n,nnz,a,la,irow,icol,lfill,dtol,mic,dscale,pstrat,ipiv,istr, &
    nnzc,npivm,iwork,liwork,ifail)
Call F11GDF to initialize the solver
Do
    ifail = 0
    Call fllgdf(method,precon,sigcmp,norm,weight,iterm,n,tol,maxitn,anorm, &
            sigmax,sigtol,maxits,monit,lwreq,work,lwork,ifail)
        If (lwork>=lwreq) Then
            Exit
        Else
            Deallocate (work)
            lwork = lwreq
            Allocate (work(lwork))
        End If
End Do
Call repeatedly F11GEF 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
    irevcm = 0
    ifail = 1
loop: Do
        Call fllgef(irevcm,x,b,wgt,work,lwork,ifail)
        If (irevcm/=4) Then
        ifail1 = -1
        Select Case (irevcm)
        Case (1)
                    Call fllxef(n,nnz,a,irow,icol,'No checking',x,b,ifail1)
            Case (2)
                Call flljbf(n,a,la,irow,icol,ipiv,istr,'No checking',x,b,ifail1)
            Case (3)
                ifail1 = 0
                Call fllgff(itn,stplhs,stprhs,anorm,sigmax,its,sigerr,work,lwork, &
                        ifail1)
                Write (nout,99996) itn, stplhs
                Write (nout,99995)
                Write (nout,99994)(x(i),b(i),i=1,n)
            End Select
        If (ifaill/=0) Then
            irevcm = 6
        End If
        Else If (ifail/=0) Then
            Write (nout,99990) ifail
            Go To 100
        Else
            Exit loop
        End If
    End Do loop
! Obtain information about the computation
    ifail1 = 0
    Call fllgff(itn,stplhs,stprhs,anorm,sigmax,its,sigerr,work,lwork,ifaill)
! Print the output data
    Write (nout,99993)
    Write (nout,99992) 'Number of iterations for convergence: ', itn
    Write (nout,99991) 'Residual norm: ', stplhs
    Write (nout,99991) 'Right-hand side of termination criterion:', stprhs
    Write (nout,99991) '1-norm of matrix A: , anorm
    Write (nout,99991) 'Largest singular value of A_bar: ', sigmax
    Output x
    Write (nout,99995)
    Write (nout,99994)(x(i),b(i),i=1,n)
    Continue
99999 Format (/,1X,'Solve a system of linear equations using the conjug',
    'ate gradient method')
99998 Format (/,1X,'Solve a system of linear equations using the Lanczo', &
    's method (SYMMLQ)')
99997 Format (/,1X,'Solve a system of linear equations using the minimu', &
    'm residual method (MINRES)')
99996 Format (/,1X,'Monitoring at iteration no.',I4,/,1X,1P,'residual no', &
    'rm: ',E14.4)
99995 Format (2X,'Solution vector',2X,'Residual vector')
99994 Format (1X,1P,E16.4,1X,E16.4)
```

```
99993 Format (/,1X,'Final Results')
99992 Format (1X,A,I4)
99991 Format (1X,A,1P,E14.4)
99990 Format (IX,/,1X,' ** F11GEF returned with IFAIL = ',I5)
    End Program fllgdfe
```


### 10.2 Program Data



### 10.3 Program Results

```
F11GDF Example Program Results
Solve a system of linear equations using the conjugate gradient method
Monitoring at iteration no. 2
residual norm: 1.9938E+00
    Solution vector Residual vector
        9.6320E-01 -2.2960E-01
        1.9934E+00 2.2254E-01
        3.0583E+00 9.5827E-02
        4.1453E+00 -2.5155E-01
        4.8289E+00 -1.7160E-01
        5.6630E+00 6.7533E-01
        7.1062E+00 -3.4737E-01
Monitoring at iteration no. 4
residual norm: 6.6574E-03
    Solution vector Residual vector
            9.9940E-01 -1.0551E-03
            2.0011E+00 -2.4675E-03
            3.0008E+00 -1.7116E-05
            3.9996E+00 4.4929E-05
            4.9991E+00 2.1359E-03
            5.9993E+00 -8.7482E-04
            7.0007E+00 6.2045E-05
Final Results
Number of iterations for convergence: 5
Residual norm:
2.0428E-14
```

| Right-hand side of termination criterion: | $3.9200 \mathrm{E}-04$ |  |
| ---: | ---: | ---: |
| l-norm of matrix A: | $1.0000 \mathrm{E}+01$ |  |
| Largest singular value of A_bar: | $1.3596 \mathrm{E}+00$ |  |
| Solution vector | Residual vector |  |
| $1.0000 \mathrm{E}+00$ | $0.0000 \mathrm{E}+00$ |  |
| $2.0000 \mathrm{E}+00$ | $0.0000 \mathrm{E}+00$ |  |
| $3.0000 \mathrm{E}+00$ | $-2.6645 \mathrm{E}-15$ |  |
| $4.0000 \mathrm{E}+00$ | $-3.5527 \mathrm{E}-15$ |  |
| $5.0000 \mathrm{E}+00$ | $-5.3291 \mathrm{E}-15$ |  |
| $6.0000 \mathrm{E}+00$ | $1.7764 \mathrm{E}-15$ |  |
| $7.0000 \mathrm{E}+00$ | $7.1054 \mathrm{E}-15$ |  |

