

## NAG Library Function Document

### nag\_sparse\_nsym\_basic\_setup (f11bdc)

#### 1 Purpose

nag\_sparse\_nsym\_basic\_setup (f11bdc) is a setup function, the first in a suite of three functions for the iterative solution of a real general (nonsymmetric) system of simultaneous linear equations. nag\_sparse\_nsym\_basic\_setup (f11bdc) must be called before nag\_sparse\_nsym\_basic\_solver (f11bec), the iterative solver. The third function in the suite, nag\_sparse\_nsym\_basic\_diagnostic (f11bfc), 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 Specification

```
#include <nag.h>
#include <nagf11.h>

void nag_sparse_nsym_basic_setup (Nag_SparseNsym_Method method,
    Nag_SparseNsym_PrecType precon, Nag_NormType norm,
    Nag_SparseNsym_Weight weight, Integer iterm, Integer n, Integer m,
    double tol, Integer maxitn, double anorm, double sigmax, Integer monit,
    Integer *lwreq, double work[], Integer lwork, NagError *fail)
```

#### 3 Description

The suite consisting of the functions nag\_sparse\_nsym\_basic\_setup (f11bdc), nag\_sparse\_nsym\_basic\_solver (f11bec) and nag\_sparse\_nsym\_basic\_diagnostic (f11bfc) is designed to solve the general (nonsymmetric) system of simultaneous linear equations  $Ax = b$  of order  $n$ , where  $n$  is large and the coefficient matrix  $A$  is sparse.

nag\_sparse\_nsym\_basic\_setup (f11bdc) is a setup function which must be called before nag\_sparse\_nsym\_basic\_solver (f11bec), the iterative solver. The third function in the suite, nag\_sparse\_nsym\_basic\_diagnostic (f11bfc), can be used to return additional information about the computation. A choice of methods is available:

- restarted generalized minimum residual method (RGMRES);
- conjugate gradient squared method (CGS);
- bi-conjugate gradient stabilized ( $\ell$ ) method (Bi-CGSTAB( $\ell$ ));
- transpose-free quasi-minimal residual method (TFQMR).

##### 3.1 Restarted Generalized Minimum Residual Method (RGMRES)

The restarted generalized minimum residual method (RGMRES) (see Saad and Schultz (1986), Barrett *et al.* (1994) and Dias da Cunha and Hopkins (1994)) starts from the residual  $r_0 = b - Ax_0$ , where  $x_0$  is an initial estimate for the solution (often  $x_0 = 0$ ). An orthogonal basis for the Krylov subspace  $\text{span}\{A^k r_0\}$ , for  $k = 0, 1, \dots$ , is generated explicitly: this is referred to as Arnoldi's method (see Arnoldi (1951)). The solution is then expanded onto the orthogonal basis so as to minimize the residual norm  $\|b - Ax\|_2$ . The lack of symmetry of  $A$  implies that the orthogonal basis is generated by applying a 'long' recurrence relation, whose length increases linearly with the iteration count. For all but the most trivial problems, computational and storage costs can quickly become prohibitive as the iteration count increases. RGMRES limits these costs by employing a restart strategy: every  $m$  iterations at most, the Arnoldi process is restarted from  $r_l = b - Ax_l$ , where the subscript  $l$  denotes the last available iterate. Each group of  $m$  iterations is referred to as a 'super-iteration'. The value of  $m$  is chosen in advance and is fixed throughout the computation. Unfortunately, an optimum value of  $m$  cannot easily be predicted.

### 3.2 Conjugate Gradient Squared Method (CGS)

The conjugate gradient squared method (CGS) (see Sonneveld (1989), Barrett *et al.* (1994) and Dias da Cunha and Hopkins (1994)) is a development of the bi-conjugate gradient method where the nonsymmetric Lanczos method is applied to reduce the coefficients matrix to real tridiagonal form: two bi-orthogonal sequences of vectors are generated starting from the residual  $r_0 = b - Ax_0$ , where  $x_0$  is an initial estimate for the solution (often  $x_0 = 0$ ) and from the *shadow residual*  $\hat{r}_0$  corresponding to the arbitrary problem  $A^T \hat{x} = \hat{b}$ , where  $\hat{b}$  can be any vector, but in practice is chosen so that  $r_0 = \hat{r}_0$ . In the course of the iteration, the residual and shadow residual  $r_i = P_i(A)r_0$  and  $\hat{r}_i = P_i(A^T)\hat{r}_0$  are generated, where  $P_i$  is a polynomial of order  $i$ , and bi-orthogonality is exploited by computing the vector product  $\rho_i = (\hat{r}_i, r_i) = (P_i(A^T)\hat{r}_0, P_i(A)r_0) = (\hat{r}_0, P_i^2(A)r_0)$ . Applying the ‘contraction’ operator  $P_i(A)$  twice, the iteration coefficients can still be recovered without advancing the solution of the shadow problem, which is of no interest. The CGS method often provides fast convergence; however, there is no reason why the contraction operator should also reduce the once reduced vector  $P_i(A)r_0$ : this may well lead to a highly irregular convergence which may result in large cancellation errors.

### 3.3 Bi-Conjugate Gradient Stabilized ( $\ell$ ) Method (Bi-CGSTAB( $\ell$ ))

The bi-conjugate gradient stabilized ( $\ell$ ) method (Bi-CGSTAB( $\ell$ )) (see Van der Vorst (1989), Sleijpen and Fokkema (1993) and Dias da Cunha and Hopkins (1994)) is similar to the CGS method above. However, instead of generating the sequence  $\{P_i^2(A)r_0\}$ , it generates the sequence  $\{Q_i(A)P_i(A)r_0\}$ , where the  $Q_i(A)$  are polynomials chosen to minimize the residual *after* the application of the contraction operator  $P_i(A)$ . Two main steps can be identified for each iteration: an OR (Orthogonal Residuals) step where a basis of order  $\ell$  is generated by a Bi-CG iteration and an MR (Minimum Residuals) step where the residual is minimized over the basis generated, by a method akin to GMRES. For  $\ell = 1$ , the method corresponds to the Bi-CGSTAB method of Van der Vorst (1989). For  $\ell > 1$ , more information about complex eigenvalues of the iteration matrix can be taken into account, and this may lead to improved convergence and robustness. However, as  $\ell$  increases, numerical instabilities may arise. For this reason, a maximum value of  $\ell = 10$  is imposed, but probably  $\ell = 4$  is sufficient in most cases.

### 3.4 Transpose-free Quasi-minimal Residual Method (TFQMR)

The transpose-free quasi-minimal residual method (TFQMR) (see Freund and Nachtigal (1991) and Freund (1993)) is conceptually derived from the CGS method. The residual is minimized over the space of the residual vectors generated by the CGS iterations under the simplifying assumption that residuals are almost orthogonal. In practice, this is not the case but theoretical analysis has proved the validity of the method. This has the effect of remedying the rather irregular convergence behaviour with wild oscillations in the residual norm that can degrade the numerical performance and robustness of the CGS method. In general, the TFQMR method can be expected to converge at least as fast as the CGS method, in terms of number of iterations, although each iteration involves a higher operation count. When the CGS method exhibits irregular convergence, the TFQMR method can produce much smoother, almost monotonic convergence curves. However, the close relationship between the CGS and TFQMR method implies that the *overall* speed of convergence is similar for both methods. In some cases, the TFQMR method may converge faster than the CGS method.

### 3.5 General Considerations

For each method, a sequence of solution iterates  $\{x_i\}$  is generated such that, hopefully, the sequence of the residual norms  $\{\|r_i\|\}$  converges to a required tolerance. Note that, in general, convergence, when it occurs, is not monotonic.

In the RGMRES and Bi-CGSTAB( $\ell$ ) methods above, your program must provide the **maximum** number of basis vectors used,  $m$  or  $\ell$ , respectively; however, a **smaller** number of basis vectors may be generated and used when the stability of the solution process requires this (see Section 9).

Faster convergence 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

$$\bar{A}\bar{x} = \bar{b}, \quad (1)$$

with, hopefully, better characteristics with respect to its speed of convergence: for example, the condition number of the coefficients matrix can be improved or eigenvalues in its spectrum can be made to coalesce. An orthogonal basis for the Krylov subspace  $\text{span}\{\bar{A}^k \bar{r}_0\}$ , for  $k = 0, 1, \dots$ , 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 *left* preconditioner  $M^{-1}$  can be used by the RGMRES, CGS and TFQMR methods, such that  $\bar{A} = M^{-1}A \sim I_n$  in (1), where  $I_n$  is the identity matrix of order  $n$ ; a *right* preconditioner  $M^{-1}$  can be used by the Bi-CGSTAB( $\ell$ ) method, such that  $\bar{A} = AM^{-1} \sim I_n$ . These are formal definitions, used only in the design of the algorithms; in practice, only the means to compute the matrix–vector products  $v = Au$  and  $v = A^T u$  (the latter only being required when an estimate of  $\|A\|_1$  or  $\|A\|_\infty$  is computed internally), and to solve the preconditioning equations  $Mv = u$  are required, i.e., explicit information about  $M$ , or its inverse is not required at any stage.

The first termination criterion

$$\|r_k\|_p \leq \tau \left( \|b\|_p + \|A\|_p \times \|x_k\|_p \right) \quad (2)$$

is available for all four 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 coefficients matrix  $\|A\|_1$  or  $\|A\|_\infty$ , when this is not known in advance, 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 are not used because any such criteria would require information about the condition number  $\kappa(A)$  which is not easily obtainable.

The second termination criterion

$$\|\bar{r}_k\|_2 \leq \tau \left( \|\bar{r}_0\|_2 + \sigma_1(\bar{A}) \times \|\Delta \bar{x}_k\|_2 \right) \quad (3)$$

is available for all methods except TFQMR. 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) for the Bi-CGSTAB( $\ell$ ) method with  $\ell > 1$ . Only the RGMRES method provides facilities to estimate  $\sigma_1(\bar{A})$  internally, when this is not supplied (see Section 9).

Termination criterion (2) is the recommended choice, despite its additional costs per iteration when using the Bi-CGSTAB( $\ell$ ) method with  $\ell > 1$ . Also, if the norm of the initial estimate is much larger than the norm of the solution, that is, if  $\|x_0\| \gg \|x\|$ , 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.,  $\|v\|_p^{(w)} = \|v^{(w)}\|_p$ , where  $(v^{(w)})_i = w_i v_i$ , for  $i = 1, 2, \dots, n$ . Note that the use of weights increases the computational costs.

The sequence of calls to the functions comprising the suite is enforced: first, the setup function `nag_sparse_nsym_basic_setup` (f11bdc) must be called, followed by the solver `nag_sparse_nsym_basic_solver` (f11bec). `nag_sparse_nsym_basic_diagnostic` (f11bfc) can be called either when `nag_sparse_nsym_basic_solver` (f11bec) is carrying out a monitoring step or after `nag_sparse_nsym_basic_solver` (f11bec) has completed its tasks. Incorrect sequencing will raise an error condition.

In general, it is not possible to recommend one method in preference to another. RGMRES is often used in the solution of systems arising from PDEs. On the other hand, it can easily stagnate when the size  $m$  of the orthogonal basis is too small, or the preconditioner is not good enough. CGS can be the fastest

method, but the computed residuals can exhibit instability which may greatly affect the convergence and quality of the solution. Bi-CGSTAB( $\ell$ ) seems robust and reliable, but it can be slower than the other methods: if a preconditioner is used and  $\ell > 1$ , Bi-CGSTAB( $\ell$ ) computes the solution of the preconditioned system  $\bar{x}_k = Mx_k$ : the preconditioning equations must be solved to obtain the required solution. The algorithm employed limits to 10% or less, when no intermediate monitoring is requested, the number of times the preconditioner has to be thus applied compared with the total number of applications of the preconditioner. TFQMR can be viewed as a more robust variant of the CGS method: it shares the CGS method speed but avoids the CGS fluctuations in the residual, which may give rise to instability. Also, when the termination criterion (2) is used, the CGS, Bi-CGSTAB( $\ell$ ) and TFQMR methods will restart the iteration automatically when necessary in order to solve the given problem.

## 4 References

- Arnoldi W (1951) The principle of minimized iterations in the solution of the matrix eigenvalue problem *Quart. Appl. Math.* **9** 17–29
- 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
- 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
- Golub G H and Van Loan C F (1996) *Matrix Computations* (3rd Edition) Johns Hopkins University Press, Baltimore
- 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( $\ell$ ) 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 Arguments

- 1: **method** – Nag\_SparseNsym\_Method *Input*  
*On entry:* the iterative method to be used.
- method** = Nag\_SparseNsym\_RGMRES  
 Restarted generalized minimum residual method.
- method** = Nag\_SparseNsym\_CGS  
 Conjugate gradient squared method.
- method** = Nag\_SparseNsym\_BiCGSTAB  
 Bi-conjugate gradient stabilized ( $\ell$ ) method.

**method** = Nag\_SparseNsym\_TFQMR  
 Transpose-free quasi-minimal residual method.

*Constraint:* **method** = Nag\_SparseNsym\_RGMRES, Nag\_SparseNsym\_CGS, Nag\_SparseNsym\_BiCGSTAB or Nag\_SparseNsym\_TFQMR.

2: **precon** – Nag\_SparseNsym\_PrecType *Input*

*On entry:* determines whether preconditioning is used.

**precon** = Nag\_SparseNsym\_NoPrec  
 No preconditioning.

**precon** = Nag\_SparseNsym\_Prec  
 Preconditioning.

*Constraint:* **precon** = Nag\_SparseNsym\_NoPrec or Nag\_SparseNsym\_Prec.

3: **norm** – Nag\_NormType *Input*

*On entry:* defines the matrix and vector norm to be used in the termination criteria.

**norm** = Nag\_OneNorm  
 $l_1$  norm.

**norm** = Nag\_InfNorm  
 $l_\infty$  norm.

**norm** = Nag\_TwoNorm  
 $l_2$  norm.

*Suggested value:*

if **iterm** = 1, **norm** = Nag\_InfNorm;  
 if **iterm** = 2, **norm** = Nag\_TwoNorm.

*Constraints:*

if **iterm** = 1, **norm** = Nag\_OneNorm, Nag\_InfNorm or Nag\_TwoNorm;  
 if **iterm** = 2, **norm** = Nag\_TwoNorm.

4: **weight** – Nag\_SparseNsym\_Weight *Input*

*On entry:* specifies whether a vector  $w$  of user-supplied weights is to be used in the computation of the vector norms required in termination criterion (2) (**iterm** = 1):  $\|v\|_p^{(w)} = \|v^{(w)}\|_p$ , where  $v_i^{(w)} = w_i v_i$ , for  $i = 1, 2, \dots, n$ . 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** = Nag\_SparseNsym\_Weighted  
 User-supplied weights are to be used and must be supplied on initial entry to nag\_sparse\_nsym\_basic\_solver (f11bec).

**weight** = Nag\_SparseNsym\_UnWeighted  
 All weights are implicitly set equal to one. Weights do not need to be supplied on initial entry to nag\_sparse\_nsym\_basic\_solver (f11bec).

*Suggested value:* **weight** = Nag\_SparseNsym\_UnWeighted.

*Constraints:*

if **iterm** = 1, **weight** = Nag\_SparseNsym\_Weighted or Nag\_SparseNsym\_UnWeighted;  
 if **iterm** = 2, **weight** = Nag\_SparseNsym\_UnWeighted.

- 5: **iterm** – Integer *Input*  
*On entry:* defines the termination criterion to be used.  
**iterm** = 1  
 Use the termination criterion defined in (2).  
**iterm** = 2  
 Use the termination criterion defined in (3).  
*Suggested value:* **iterm** = 1.  
*Constraints:*  
 if **method** = Nag\_SparseNsym\_TFQMR or **weight** = Nag\_SparseNsym\_Weighted or  
**norm**  $\neq$  Nag\_TwoNorm, **iterm** = 1;  
 otherwise **iterm** = 1 or 2.  
**Note:** **iterm** = 2 is only appropriate for a restricted set of choices for **method**, **norm** and **weight**;  
 t h a t i s **norm** = Nag\_TwoNorm, **weight** = Nag\_SparseNsym\_UnWeighted a n d  
**method**  $\neq$  Nag\_SparseNsym\_TFQMR.
- 6: **n** – Integer *Input*  
*On entry:*  $n$ , the order of the matrix  $A$ .  
*Constraint:* **n** > 0.
- 7: **m** – Integer *Input*  
*On entry:* if **method** = Nag\_SparseNsym\_RGMRES, **m** is the dimension  $m$  of the restart subspace.  
 If **method** = Nag\_SparseNsym\_BiCGSTAB, **m** is the order  $\ell$  of the polynomial Bi-CGSTAB  
 method.  
 Otherwise, **m** is not referenced.  
*Constraints:*  
 if **method** = Nag\_SparseNsym\_RGMRES,  $0 < \mathbf{m} \leq \min(\mathbf{n}, 50)$ ;  
 if **method** = Nag\_SparseNsym\_BiCGSTAB,  $0 < \mathbf{m} \leq \min(\mathbf{n}, 10)$ .
- 8: **tol** – double *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  $\tau = \max(\mathbf{tol}, 10\epsilon, \sqrt{n}\epsilon)$  is used.  
*Constraint:* **tol** < 1.0.
- 9: **maxitn** – Integer *Input*  
*On entry:* the maximum number of iterations.  
*Constraint:* **maxitn** > 0.
- 10: **anorm** – double *Input*  
*On entry:* if **anorm** > 0.0, the value of  $\|A\|_p$  to be used in the termination criterion (2)  
 (**iterm** = 1).  
 If **anorm**  $\leq$  0.0, **iterm** = 1 and **norm** = Nag\_OneNorm or Nag\_InfNorm, then  $\|A\|_1 = \|A\|_\infty$  is  
 estimated internally by nag\_sparse\_nsym\_basic\_solver (f11bec).  
 If **iterm** = 2, **anorm** is not referenced.  
*Constraint:* if **iterm** = 1 and **norm** = Nag\_TwoNorm, **anorm** > 0.0.

- 11: **sigmax** – double *Input*  
*On entry:* if **iterm** = 2, the largest singular value  $\sigma_1$  of the preconditioned iteration matrix; otherwise, **sigmax** is not referenced.  
 If **sigmax**  $\leq$  0.0, **iterm** = 2 and **method** = Nag\_SparseNsym\_RGMRES, then the value of  $\sigma_1$  will be estimated internally.  
*Constraint:* if **method** = Nag\_SparseNsym\_CGS or Nag\_SparseNsym\_BiCGSTAB and **iterm** = 2, **sigmax** > 0.0.
- 12: **monit** – Integer *Input*  
*On entry:* if **monit** > 0, the frequency at which a monitoring step is executed by nag\_sparse\_nsym\_basic\_solver (f11bec): if **method** = Nag\_SparseNsym\_CGS or Nag\_SparseNsym\_TFQMR, a monitoring step is executed every **monit** iterations; otherwise, a monitoring step is executed every **monit** super-iterations (groups of up to  $m$  or  $\ell$  iterations for RGMRES or Bi-CGSTAB( $\ell$ ), respectively).  
 There are some additional computational costs involved in monitoring the solution and residual vectors when the Bi-CGSTAB( $\ell$ ) method is used with  $\ell > 1$ .  
*Constraint:* **monit**  $\leq$  **maxitn**.
- 13: **lwreq** – Integer \* *Output*  
*On exit:* the minimum amount of workspace required by nag\_sparse\_nsym\_basic\_solver (f11bec). (See also Section 5 in nag\_sparse\_nsym\_basic\_solver (f11bec).)
- 14: **work[lwork]** – double *Communication Array*  
*On exit:* the array **work** is initialized by nag\_sparse\_nsym\_basic\_setup (f11bdc). It must **not** be modified before calling the next function in the suite, namely nag\_sparse\_nsym\_basic\_solver (f11bec).
- 15: **lwork** – Integer *Input*  
*On entry:* the dimension of the array **work**.  
*Constraint:* **lwork**  $\geq$  100.  
**Note:** although the minimum value of **lwork** ensures the correct functioning of nag\_sparse\_nsym\_basic\_setup (f11bdc), a larger value is required by the other functions in the suite, namely nag\_sparse\_nsym\_basic\_solver (f11bec) and nag\_sparse\_nsym\_basic\_diagnostic (f11bfc). The required value is as follows:
- | <b>Method</b>       | <b>Requirements</b>   |
|---------------------|---|
| RGMRES              | <b>lwork</b> = $100 + n(m + 3) + m(m + 5) + 1$ , where $m$ is the dimension of the basis.   |
| CGS                 | <b>lwork</b> = $100 + 7n$ .   |
| Bi-CGSTAB( $\ell$ ) | <b>lwork</b> = $100 + (2n + \ell)(\ell + 2) + p$ , where $\ell$ is the order of the method. |
| TFQMR               | <b>lwork</b> = $100 + 10n$ ,  |
- 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.
- 16: **fail** – NagError \* *Input/Output*  
 The NAG error argument (see Section 3.6 in the Essential Introduction).

## 6 Error Indicators and Warnings

### NE\_BAD\_PARAM

On entry, argument  $\langle value \rangle$  had an illegal value.

### NE\_CONSTRAINT

On entry, **item** = 1, **norm** = Nag\_TwoNorm and **anorm** =  $\langle value \rangle$ .  
Constraint: if **item** = 1 and **norm** = Nag\_TwoNorm, **anorm** > 0.0.

On entry, **item** = 2, **method** =  $\langle value \rangle$  and **sigmax** =  $\langle value \rangle$ .  
Constraint: if **item** = 2 and **method** = Nag\_SparseNsym\_CGS or Nag\_SparseNsym\_BiCGSTAB, **sigmax** > 0.0.

### NE\_ENUM\_3\_INT

On entry, **method** =  $\langle value \rangle$ , **weight** =  $\langle value \rangle$ , **norm** =  $\langle value \rangle$  and **item** =  $\langle value \rangle$ .  
Constraint: if **method** = Nag\_SparseNsym\_TFQMR or **weight** = Nag\_SparseNsym\_Weighted or **norm**  $\neq$  Nag\_TwoNorm, **item** = 1. Otherwise, **item** = 1 or 2.

### NE\_ENUM\_INT\_2

On entry, **m** =  $\langle value \rangle$  and **method** =  $\langle value \rangle$ .  
Constraint: if **method** = Nag\_SparseNsym\_RGMRES or Nag\_SparseNsym\_BiCGSTAB, **m** > 0.

On entry, **m** =  $\langle value \rangle$ , **n** =  $\langle value \rangle$  and **method** =  $\langle value \rangle$ .  
Constraint: if **method** = Nag\_SparseNsym\_RGMRES, **m**  $\leq$  min(**n**, 50). If **method** = Nag\_SparseNsym\_BiCGSTAB, **m**  $\leq$  min(**n**, 10).

### NE\_INT

On entry, **lwork** =  $\langle value \rangle$ .  
Constraint: **lwork**  $\geq$  100.

On entry, **maxitn** =  $\langle value \rangle$ .  
Constraint: **maxitn** > 0.

On entry, **n** =  $\langle value \rangle$ .  
Constraint: **n** > 0.

### NE\_INT\_2

On entry, **monit** =  $\langle value \rangle$  and **maxitn** =  $\langle value \rangle$ .  
Constraint: **monit**  $\leq$  **maxitn**.

### NE\_INTERNAL\_ERROR

An internal error has occurred in this function. Check the function call and any array sizes. If the call is correct then please contact NAG for assistance.

### NE\_OUT\_OF\_SEQUENCE

nag\_sparse\_nsym\_basic\_setup (f11bdc) has been called out of sequence: either nag\_sparse\_nsym\_basic\_setup (f11bdc) has been called twice or nag\_sparse\_nsym\_basic\_solver (f11bec) has not terminated its current task.

### NE\_REAL

On entry, **tol** =  $\langle value \rangle$ .  
Constraint: **tol** < 1.0.

## 7 Accuracy

Not applicable.



## 8 Parallelism and Performance

Not applicable.

## 9 Further Comments

RGMRES can estimate internally the maximum singular value  $\sigma_1$  of the iteration matrix, using  $\sigma_1 \sim \|T\|_1$ , where  $T$  is the upper triangular matrix obtained by  $QR$  factorization of the upper Hessenberg matrix generated by the Arnoldi process. The computational costs of this computation are negligible when compared to the overall costs.

Loss of orthogonality in the RGMRES method, or of bi-orthogonality in the Bi-CGSTAB( $\ell$ ) method may degrade the solution and speed of convergence. For both methods, the algorithms employed include checks on the basis vectors so that the number of basis vectors used for a given super-iteration may be less than the value specified in the input argument **m**. Also, if termination criterion (2) is used the CGS, Bi-CGSTAB( $\ell$ ) and TFQMR methods will restart automatically the computation from the last available iterates, when the stability of the solution process requires it.

Termination criterion (3), when available, involves only the residual (or norm of the residual) produced directly by the iteration process: this may differ from the norm of the true residual  $\tilde{r}_k = b - Ax_k$ , particularly when the norm of the residual is very small. Also, if the norm of the initial estimate of the solution is much larger than the norm of the exact solution, convergence can be achieved despite very large errors in the solution. On the other hand, termination criterion (3) is cheaper to use and inspects the progress of the actual iteration. Termination criterion (2) should be preferred in most cases, despite its slightly larger costs.

## 10 Example

This example solves an  $8 \times 8$  nonsymmetric system of simultaneous linear equations using the bi-conjugate gradient stabilized method of order  $\ell = 1$ , where the coefficients matrix  $A$  has a random sparsity pattern. An incomplete  $LU$  preconditioner is used (routines `nag_sparse_nsym_fac` (f11dac) and `nag_sparse_nsym_precon_ilu_solve` (f11dbc)).

### 10.1 Program Text

```

/* nag_sparse_nsym_basic_setup (f11bdc) Example Program.
 *
 * Copyright 2011, Numerical Algorithms Group.
 *
 * Mark 23, 2011.
 */

#include <nag.h>
#include <nag_stdlib.h>
#include <nagf11.h>
int main(void)
{
    /* Scalars */
    Integer          exit_status = 0;
    double           anorm, dtol, sigmax, stplhs, stprhs, tol;
    Integer          i, irevcn, iterm, itn, j, la, lfill, lwork,
                   lwreq, m, maxitn, monit, n, nnz, nnzc, npivm;

    /* Arrays */
    char             nag_enum_arg[100];
    double           *a = 0, *b = 0, *wgt = 0, *work = 0, *x = 0;
    Integer          *icol = 0, *idiag = 0, *ipivp = 0, *ipivq = 0,
                   *irow = 0, *istr = 0;

    /* NAG types */
    Nag_SparseNsym_Piv      pstrat;
    Nag_SparseNsym_Fact     milu;
    Nag_SparseNsym_Method   method;
    Nag_SparseNsym_PrecType precon;
    Nag_NormType            norm;
    Nag_SparseNsym_Weight   weight;

```

```

Nag_TransType          trans;
Nag_SparseNsym_CheckData check = Nag_SparseNsym_NoCheck;
Nag_Error              fail, fail1;

INIT_FAIL(fail);
INIT_FAIL(fail1);

printf("nag_sparse_nsym_basic_setup (f11bdc) Example Program Results\n");
/* Skip heading in data file */
scanf("%*[\n]");
scanf("%ld%*[\n]", &n);
scanf("%ld%*[\n]", &nnz);
la = 3 * nnz;
lwork = 200;
if (
    ! (a = NAG_ALLOC((la), double)) ||
    ! (b = NAG_ALLOC((n), double)) ||
    ! (work = NAG_ALLOC((lwork), double)) ||
    ! (x = NAG_ALLOC((n), double)) ||
    ! (wgt = NAG_ALLOC((n), double)) ||
    ! (icol = NAG_ALLOC((la), Integer)) ||
    ! (idiag = NAG_ALLOC((n), Integer)) ||
    ! (ipivp = NAG_ALLOC((n), Integer)) ||
    ! (ipivq = NAG_ALLOC((n), Integer)) ||
    ! (irow = NAG_ALLOC((la), Integer)) ||
    ! (istr = NAG_ALLOC((n + 1), Integer))
) {
    printf("Allocation failure\n");
    exit_status = -1;
    goto END;
}

/* Read or initialize the parameters for the iterative solver */
scanf("%99s%*[\n]", nag_enum_arg);
/* nag_enum_name_to_value (x04nac).
 * Converts NAG enum member name to value
 */
method = (Nag_SparseNsym_Method) nag_enum_name_to_value(nag_enum_arg);
scanf("%99s%*[\n]", nag_enum_arg);
precon = (Nag_SparseNsym_PrecType) nag_enum_name_to_value(nag_enum_arg);
scanf("%99s%*[\n]", nag_enum_arg);
norm = (Nag_NormType) nag_enum_name_to_value(nag_enum_arg);
scanf("%99s%*[\n]", nag_enum_arg);
weight = (Nag_SparseNsym_Weight) nag_enum_name_to_value(nag_enum_arg);
scanf("%ld%*[\n]", &iterm);
scanf("%ld%lf%ld%*[\n]", &m, &tol, &maxitn);
scanf("%ld%*[\n]", &monit);

/* Read the parameters for the preconditioner */
scanf("%ld%lf%*[\n]", &lfill, &dtol);
scanf("%99s%*[\n]", nag_enum_arg);
milu = (Nag_SparseNsym_Fact) nag_enum_name_to_value(nag_enum_arg);
scanf("%99s%*[\n]", nag_enum_arg);
pstrat = (Nag_SparseNsym_Piv) nag_enum_name_to_value(nag_enum_arg);

/* Read the non-zero elements of the matrix A */
for (i = 0; i <= nnz-1; i++)
    scanf("%lf%NAG_IFMT %ld%*[\n]", &a[i], &irow[i], &icol[i]);

/* Read right-hand side vector B and initial approximate solution */
for (j = 0; j < n; j++)
    scanf("%lf", &b[j]);
scanf("%*[\n] ");
for (j = 0; j < n; j++)
    scanf("%lf", &x[j]);
scanf("%*[\n] ");

/* nag_sparse_nsym_fac (f11dac).
 * Incomplete LU factorization (nonsymmetric)
 */
nag_sparse_nsym_fac(n, nnz, &a, &la, &irow, &icol, lfill, dtol, pstrat,

```

```

        milu, ipivp, ipivq, istr, idiag, &nnzc, &npivm, &fail);
if (fail.code != NE_NOERROR)
{
    printf("Error from nag_sparse_nsym_fac (f11dac)\n%s\n", fail.message);
    exit_status = 1;
    goto END;
}

/* Initialize the solver using nag_sparse_nsym_basic_setup (f11bdc)
 * Real sparse nonsymmetric linear systems, setup routine
 */
anorm = 0.0;
sigmax = 0.0;
nag_sparse_nsym_basic_setup(method, precon, norm, weight, iterm, n, m, tol,
                           maxitn, anorm, sigmax, monit, &lwreq, work,
                           lwork, &fail);
if (fail.code != NE_NOERROR)
{
    printf("Error from nag_sparse_nsym_basic_setup (f11bdc)\n%s\n",
          fail.message);
    exit_status = 2;
    goto END;
}

/* Call solver repeatedly 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;
lwreq = lwork;
/* First call to nag_sparse_nsym_basic_solver (f11bec)
 * Real sparse nonsymmetric linear systems, solver routine
 * preconditioned RGMRES, CGS, Bi-CGSTAB or TFQMR method
 */
nag_sparse_nsym_basic_solver(&irevcm, x, b, wgt, work, lwreq, &fail);
while (irevcm != 4) {
    switch (irevcm) {
        case -1:
            /* nag_sparse_nsym_matvec (f11xac)
             * Real sparse nonsymmetric matrix vector multiply
             */
            trans = Nag_Trans;
            nag_sparse_nsym_matvec(trans, n, nnz, a, irow, icol, check, x, b,
                                   &fail1);
            break;
        case 1:
            trans = Nag_NoTrans;
            nag_sparse_nsym_matvec(trans, n, nnz, a, irow, icol, check, x, b,
                                   &fail1);
            break;
        case 2:
            /* nag_sparse_nsym_precon_ilu_solve (f11dbc)
             * Solution of real linear system involving incomplete LU
             * preconditioning matrix
             */
            trans = Nag_NoTrans;
            nag_sparse_nsym_precon_ilu_solve(trans, n, a, la, irow, icol, ipivp,
                                             ipivq, istr, idiag, check, x, b,
                                             &fail1);
            break;
        case 3:
            /* nag_sparse_nsym_basic_diagnostic (f11bfc)
             * Real sparse nonsymmetric linear systems, diagnostic routine
             */
            nag_sparse_nsym_basic_diagnostic(&itn, &stplhs, &stprhs, &anorm,
                                             &sigmax, work, lwreq, &fail1);
            printf("\nMonitoring at iteration no.%4ld\n", itn);
            printf("residual norm:%14.4e\n\n", stplhs);
            printf("%16s%16s\n", "Current Sol", "Current Res");
            for (i = 0; i < n; i++)
                printf("%16.4e%16.4e\n", x[i], b[i]);
    }
}

```

```

        printf("\n");
    }
    if (fail1.code != NE_NOERROR) irevcm = 6;
    /* Next call to nag_sparse_nsym_basic_solver (f11bec) */
    nag_sparse_nsym_basic_solver(&irevcm, x, b, wgt, work, lwreq, &fail);
}
if (fail.code != NE_NOERROR)
{
    printf("Error from nag_sparse_nsym_basic_solver (f11bec)\n%s\n",
        fail.message);
    exit_status = 3;
    goto END;
}

/* Obtain information about the computation using
 * nag_sparse_nsym_basic_diagnostic (f11bfc).
 */
nag_sparse_nsym_basic_diagnostic(&itn, &stplhs, &stprhs, &anorm, &sigmax,
    work, lwreq, &fail);
if (fail.code != NE_NOERROR)
{
    printf("Error from nag_sparse_nsym_basic_diagnostic (f11bfc) \n%s\n",
        fail.message);
    exit_status = 4;
    goto END;
}
/* Print the output data */
printf("Final Results\n");
printf("Number of iterations for convergence:      %5ld\n", itn);
printf("Residual norm:                             %14.4e\n", stplhs);
printf("Right-hand side of termination criterion: %14.4e\n", stprhs);
printf("1-norm of matrix A:                         %14.4e\n", anorm);
/* Output x */
printf("\n%16s%16s\n", "Solution", "Residual");
for (i = 0; i < n; i++)
    printf("%16.4e%16.4e\n", x[i], b[i]);
printf("\n");
END:

    NAG_FREE(a);
    NAG_FREE(b);
    NAG_FREE(work);
    NAG_FREE(x);
    NAG_FREE(wgt);
    NAG_FREE(icol);
    NAG_FREE(idiag);
    NAG_FREE(ipivp);
    NAG_FREE(ipivq);
    NAG_FREE(irow);
    NAG_FREE(istr);

    return exit_status;
}

```

## 10.2 Program Data

```

nag_sparse_nsym_basic_setup (f11bdc) Example Program Data
  8                               : n
 24                               : nnz
Nag_SparseNsym_BiCGSTAB          : method
Nag_SparseNsym_Prec              : precon
Nag_OneNorm                      : norm
Nag_SparseNsym_UnWeighted        : weight
  1                               : iterm
  1  1.0e-8 20                   : m, tol, maxitn
  1                               : monit
  0  0.0                          : lfill, dtol
Nag_SparseNsym_UnModFact         : milu
Nag_SparseNsym_CompletePiv       : pstrat
 2.0  1  1

```

```

-1.0  1  4
 1.0  1  8
 4.0  2  1
-3.0  2  2
 2.0  2  5
-7.0  3  3
 2.0  3  6
 3.0  4  1
-4.0  4  3
 5.0  4  4
 5.0  4  7
-1.0  5  2
 8.0  5  5
-3.0  5  7
-6.0  6  1
 5.0  6  3
 2.0  6  6
-5.0  7  3
-1.0  7  5
 6.0  7  7
-1.0  8  2
 2.0  8  6
 3.0  8  8
 6.0  8.0 -9.0 46.0 : a[i], irow[i], icol[i], i=0,...,nnz-1
17.0 21.0 22.0 34.0 : b[i], i=0,...,n-1
 0.0  0.0  0.0  0.0
 0.0  0.0  0.0  0.0 : x[i], i=0,...,n-1

```

### 10.3 Program Results

nag\_sparse\_nsym\_basic\_setup (f11bdc) Example Program Results

Monitoring at iteration no. 1  
residual norm: 1.4059e+02

Current Sol	Current Res
-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 no. 2  
residual norm: 3.2742e+01

Current Sol	Current Res
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

Final Results  
Number of iterations for convergence: 3  
Residual norm: 1.0373e-08  
Right-hand side of termination criterion: 5.5900e-06  
1-norm of matrix A: 1.1000e+01

Solution	Residual
1.0000e+00	-1.3554e-09
2.0000e+00	-2.6109e-09
3.0000e+00	2.2471e-10
4.0000e+00	-3.2203e-09

5.0000e+00	6.3045e-10
6.0000e+00	-5.2431e-10
7.0000e+00	9.5771e-10
8.0000e+00	-8.4890e-10

---