NAG Library Routine Document

F07BBF (DGBSVX)

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

F07BBF (DGBSVX) uses the LU factorization to compute the solution to a real system of linear equations

$$AX = B \quad \text{or} \quad A^{\mathrm{T}}X = B,$$

where A is an n by n band matrix with k_l subdiagonals and k_u superdiagonals, and X and B are n by r matrices. Error bounds on the solution and a condition estimate are also provided.

2 Specification

```
SUBROUTINE F07BBF (FACT, TRANS, N, KL, KU, NRHS, AB, LDAB, AFB, LDAFB,    &
                   IPIV, EQUED, R, C, B, LDB, X, LDX, RCOND, FERR, BERR,  &
                   WORK, IWORK, INFO)
INTEGER            N, KL, KU, NRHS, LDAB, LDAFB, IPIV(*), LDB, LDX,       &
                   IWORK(N), INFO
REAL (KIND=nag_wp) AB(LDAB,*), AFB(LDAFB,*), R(*), C(*), B(LDB,*),        &
                   X(LDX,*), RCOND, FERR(NRHS), BERR(NRHS),               &
                   WORK(max(1,3*N))
CHARACTER(1)       FACT, TRANS, EQUED
```

The routine may be called by its LAPACK name **dgbsvx**.

3 Description

F07BBF (DGBSVX) performs the following steps:

1. **Equilibration**

   The linear system to be solved may be badly scaled. However, the system can be equilibrated as a first stage by setting FACT = 'E'. In this case, real scaling factors are computed and these factors then determine whether the system is to be equilibrated. Equilibrated forms of the systems $AX = B$ and $A^{\mathrm{T}}X = B$ are

   $$(D_R A D_C)(D_C^{-1}X) = D_R B$$

   and

   $$(D_R A D_C)^{\mathrm{T}}(D_R^{-1}X) = D_C B,$$

   respectively, where D_R and D_C are diagonal matrices, with positive diagonal elements, formed from the computed scaling factors.

   When equilibration is used, A will be overwritten by $D_R A D_C$ and B will be overwritten by $D_R B$ (or $D_C B$ when the solution of $A^{\mathrm{T}}X = B$ is sought).

2. **Factorization**

   The matrix A, or its scaled form, is copied and factored using the LU decomposition

   $$A = PLU,$$

   where P is a permutation matrix, L is a unit lower triangular matrix, and U is upper triangular.

This stage can be by-passed when a factored matrix (with scaled matrices and scaling factors) are supplied; for example, as provided by a previous call to F07BBF (DGBSVX) with the same matrix A.

3. **Condition Number Estimation**

The LU factorization of A determines whether a solution to the linear system exists. If some diagonal element of U is zero, then U is exactly singular, no solution exists and the routine returns with a failure. Otherwise the factorized form of A is used to estimate the condition number of the matrix A. If the reciprocal of the condition number is less than **machine precision** then a warning code is returned on final exit.

4. **Solution**

The (equilibrated) system is solved for X ($D_C^{-1}X$ or $D_R^{-1}X$) using the factored form of A ($D_R A D_C$).

5. **Iterative Refinement**

Iterative refinement is applied to improve the computed solution matrix and to calculate error bounds and backward error estimates for the computed solution.

6. **Construct Solution Matrix X**

If equilibration was used, the matrix X is premultiplied by D_C (if TRANS = 'N') or D_R (if TRANS = 'T' or 'C') so that it solves the original system before equilibration.

4   References

Anderson E, Bai Z, Bischof C, Blackford S, Demmel J, Dongarra J J, Du Croz J J, Greenbaum A, Hammarling S, McKenney A and Sorensen D (1999) *LAPACK Users' Guide* (3rd Edition) SIAM, Philadelphia http://www.netlib.org/lapack/lug

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

Higham N J (2002) *Accuracy and Stability of Numerical Algorithms* (2nd Edition) SIAM, Philadelphia

5   Parameters

1:     FACT – CHARACTER(1)   *Input*

On entry: specifies whether or not the factorized form of the matrix A is supplied on entry, and if not, whether the matrix A should be equilibrated before it is factorized.

FACT = 'F'
      AFB and IPIV contain the factorized form of A. If EQUED $\neq$ 'N', the matrix A has been equilibrated with scaling factors given by R and C. AB, AFB and IPIV are not modified.

FACT = 'N'
      The matrix A will be copied to AFB and factorized.

FACT = 'E'
      The matrix A will be equilibrated if necessary, then copied to AFB and factorized.

Constraint: FACT = 'F', 'N' or 'E'.

2:     TRANS – CHARACTER(1)   *Input*

On entry: specifies the form of the system of equations.

TRANS = 'N'
      $AX = B$ (No transpose).

TRANS = 'T' or 'C'
$$A^T X = B \text{ (Transpose)}.$$

Constraint: TRANS = 'N', 'T' or 'C'.

3:    N – INTEGER   *Input*

On entry: n, the number of linear equations, i.e., the order of the matrix A.

Constraint: N ≥ 0.

4:    KL – INTEGER   *Input*

On entry: k_l, the number of subdiagonals within the band of the matrix A.

Constraint: KL ≥ 0.

5:    KU – INTEGER   *Input*

On entry: k_u, the number of superdiagonals within the band of the matrix A.

Constraint: KU ≥ 0.

6:    NRHS – INTEGER   *Input*

On entry: r, the number of right-hand sides, i.e., the number of columns of the matrix B.

Constraint: NRHS ≥ 0.

7:    AB(LDAB, ∗) – REAL (KIND=nag_wp) array   *Input/Output*

Note: the second dimension of the array AB must be at least $\max(1, N)$.

On entry: the n by n coefficient matrix A.

The matrix is stored in rows 1 to $k_l + k_u + 1$, more precisely, the element A_{ij} must be stored in

$$AB(k_u + 1 + i - j, j) \quad \text{for } \max(1, j - k_u) \le i \le \min(n, j + k_l).$$

See Section 9 for further details.

If FACT = 'F' and EQUED ≠ 'N', A must have been equilibrated by the scaling factors in R and/or C.

On exit: if FACT = 'F' or 'N', or if FACT = 'E' and EQUED = 'N', AB is not modified.

If EQUED ≠ 'N' then, if no constraints are violated, A is scaled as follows:

>   if EQUED = 'R', $A = D_r A$;

>   if EQUED = 'C', $A = A D_c$;

>   if EQUED = 'B', $A = D_r A D_c$.

8:    LDAB – INTEGER   *Input*

On entry: the first dimension of the array AB as declared in the (sub)program from which F07BBF (DGBSVX) is called.

Constraint: LDAB ≥ KL + KU + 1.

9:    AFB(LDAFB, ∗) – REAL (KIND=nag_wp) array   *Input/Output*

Note: the second dimension of the array AFB must be at least $\max(1, N)$.

On entry: if FACT = 'N' or 'E', AFB need not be set.

If FACT = 'F', details of the LU factorization of the n by n band matrix A, as computed by F07BDF (DGBTRF).

The upper triangular band matrix U, with $k_l + k_u$ superdiagonals, is stored in rows 1 to $k_l + k_u + 1$ of the array, and the multipliers used to form the matrix L are stored in rows $k_l + k_u + 2$ to $2k_l + k_u + 1$.

If EQUED $\neq$ 'N', AFB is the factorized form of the equilibrated matrix A.

On exit: if FACT = 'F', AFB is unchanged from entry.

Otherwise, if no constraints are violated, then if FACT = 'N', AFB returns details of the LU factorization of the band matrix A, and if FACT = 'E', AFB returns details of the LU factorization of the equilibrated band matrix A (see the description of AB for the form of the equilibrated matrix).

10:   LDAFB – INTEGER   *Input*

On entry: the first dimension of the array AFB as declared in the (sub)program from which F07BBF (DGBSVX) is called.

Constraint: LDAFB $\geq 2 \times$ KL + KU + 1.

11:   IPIV(∗) – INTEGER array   *Input/Output*

Note: the dimension of the array IPIV must be at least $\max(1, \text{N})$.

On entry: if FACT = 'N' or 'E', IPIV need not be set.

If FACT = 'F', IPIV contains the pivot indices from the factorization $A = LU$, as computed by F07BDF (DGBTRF); row i of the matrix was interchanged with row IPIV(i).

On exit: if FACT = 'F', IPIV is unchanged from entry.

Otherwise, if no constraints are violated, IPIV contains the pivot indices that define the permutation matrix P; at the ith step row i of the matrix was interchanged with row IPIV(i). IPIV(i) = i indicates a row interchange was not required.

If FACT = 'N', the pivot indices are those corresponding to the factorization $A = LU$ of the original matrix A.

If FACT = 'E', the pivot indices are those corresponding to the factorization of $A = LU$ of the equilibrated matrix A.

12:   EQUED – CHARACTER(1)   *Input/Output*

On entry: if FACT = 'N' or 'E', EQUED need not be set.

If FACT = 'F', EQUED must specify the form of the equilibration that was performed as follows:

      if EQUED = 'N', no equilibration;

      if EQUED = 'R', row equilibration, i.e., A has been premultiplied by D_R;

      if EQUED = 'C', column equilibration, i.e., A has been postmultiplied by D_C;

      if EQUED = 'B', both row and column equilibration, i.e., A has been replaced by $D_R A D_C$.

On exit: if FACT = 'F', EQUED is unchanged from entry.

Otherwise, if no constraints are violated, EQUED specifies the form of equilibration that was performed as specified above.

Constraint: if FACT = 'F', EQUED = 'N', 'R', 'C' or 'B'.

13:   R(∗) – REAL (KIND=nag_wp) array   *Input/Output*

Note: the dimension of the array R must be at least $\max(1, \text{N})$.

On entry: if FACT = 'N' or 'E', R need not be set.

If FACT = 'F' and EQUED = 'R' or 'B', R must contain the row scale factors for A, D_R; each element of R must be positive.

On exit: if FACT = 'F', R is unchanged from entry.

Otherwise, if no constraints are violated and EQUED = 'R' or 'B', R contains the row scale factors for A, D_R, such that A is multiplied on the left by D_R; each element of R is positive.

14:    C(∗) – REAL (KIND=nag_wp) array   *Input/Output*

   Note: the dimension of the array C must be at least $\max(1, N)$.

   On entry: if FACT = 'N' or 'E', C need not be set.

   If FACT = 'F' or EQUED = 'C' or 'B', C must contain the column scale factors for A, D_C; each element of C must be positive.

   On exit: if FACT = 'F', C is unchanged from entry.

   Otherwise, if no constraints are violated and EQUED = 'C' or 'B', C contains the row scale factors for A, D_C; each element of C is positive.

15:    B(LDB, ∗) – REAL (KIND=nag_wp) array   *Input/Output*

   Note: the second dimension of the array B must be at least $\max(1, \text{NRHS})$.

   On entry: the n by r right-hand side matrix B.

   On exit: if EQUED = 'N', B is not modified.

   If TRANS = 'N' and EQUED = 'R' or 'B', B is overwritten by $D_R B$.

   If TRANS = 'T' or 'C' and EQUED = 'C' or 'B', B is overwritten by $D_C B$.

16:    LDB – INTEGER   *Input*

   On entry: the first dimension of the array B as declared in the (sub)program from which F07BBF (DGBSVX) is called.

   Constraint: $\text{LDB} \geq \max(1, N)$.

17:    X(LDX, ∗) – REAL (KIND=nag_wp) array   *Output*

   Note: the second dimension of the array X must be at least $\max(1, \text{NRHS})$.

   On exit: if INFO = 0 or $N + 1$, the n by r solution matrix X to the original system of equations. Note that the arrays A and B are modified on exit if EQUED $\neq$ 'N', and the solution to the equilibrated system is $D_C^{-1} X$ if TRANS = 'N' and EQUED = 'C' or 'B', or $D_R^{-1} X$ if TRANS = 'T' or 'C' and EQUED = 'R' or 'B'.

18:    LDX – INTEGER   *Input*

   On entry: the first dimension of the array X as declared in the (sub)program from which F07BBF (DGBSVX) is called.

   Constraint: $\text{LDX} \geq \max(1, N)$.

19:    RCOND – REAL (KIND=nag_wp)   *Output*

   On exit: if no constraints are violated, an estimate of the reciprocal condition number of the matrix A (after equilibration if that is performed), computed as $\text{RCOND} = 1.0 / \left(\|A\|_1 \|A^{-1}\|_1 \right)$.

20:    FERR(NRHS) – REAL (KIND=nag_wp) array   *Output*

   On exit: if INFO = 0 or $N + 1$, an estimate of the forward error bound for each computed solution vector, such that $\|\hat{x}_j - x_j\|_\infty / \|x_j\|_\infty \leq \text{FERR}(j)$ where $\hat{x}_j$ is the jth column of the computed solution returned in the array X and x_j is the corresponding column of the exact solution X. The estimate is as reliable as the estimate for RCOND, and is almost always a slight overestimate of the true error.

21:   BERR(NRHS) – REAL (KIND=nag_wp) array   *Output*

   On exit: if INFO = 0 or N + 1, an estimate of the component-wise relative backward error of each computed solution vector $\hat{x}_j$ (i.e., the smallest relative change in any element of A or B that makes $\hat{x}_j$ an exact solution).

22:   WORK(max$(1, 3 \times$ N$))$ – REAL (KIND=nag_wp) array   *Output*

   On exit: if INFO = 0, WORK(1) contains the reciprocal pivot growth factor $\max\left|a_{ij}\right|/\max\left|u_{ij}\right|$. If WORK(1) is much less than 1, then the stability of the LU factorization of the (equilibrated) matrix A could be poor. This also means that the solution X, condition estimator RCOND, and forward error bound FERR could be unreliable. If the factorization fails with INFO > 0 and INFO ≤ N, WORK(1) contains the reciprocal pivot growth factor for the leading INFO columns of A.

23:   IWORK(N) – INTEGER array   *Workspace*

24:   INFO – INTEGER   *Output*

   On exit: INFO = 0 unless the routine detects an error (see Section 6).

6   Error Indicators and Warnings

INFO < 0

   If INFO = $-i$, argument i had an illegal value. An explanatory message is output, and execution of the program is terminated.

INFO > 0 and INFO ≤ N

   Element $\langle value \rangle$ of the diagonal is exactly zero. The factorization has been completed, but the factor U is exactly singular, so the solution and error bounds could not be computed. RCOND = 0.0 is returned.

INFO = N + 1

   U is nonsingular, but RCOND is less than ***machine precision***, meaning that the matrix is singular to working precision. Nevertheless, the solution and error bounds are computed because there are a number of situations where the computed solution can be more accurate than the value of RCOND would suggest.

7   Accuracy

For each right-hand side vector b, the computed solution $\hat{x}$ is the exact solution of a perturbed system of equations $(A + E)\hat{x} = b$, where

$$|E| \leq c(n)\epsilon P|L||U|,$$

$c(n)$ is a modest linear function of n, and ϵ is the ***machine precision***. See Section 9.3 of Higham (2002) for further details.

If x is the true solution, then the computed solution $\hat{x}$ satisfies a forward error bound of the form

$$\frac{\|x - \hat{x}\|_\infty}{\|\hat{x}\|_\infty} \leq w_c \operatorname{cond}(A, \hat{x}, b)$$

where $\operatorname{cond}(A, \hat{x}, b) = \left\||A^{-1}|(|A||\hat{x}| + |b|)\right\|_\infty / \|\hat{x}\|_\infty \leq \operatorname{cond}(A) = \left\||A^{-1}||A|\right\|_\infty \leq \kappa_\infty(A)$. If $\hat{x}$ is the jth column of X, then w_c is returned in BERR(j) and a bound on $\|x - \hat{x}\|_\infty / \|\hat{x}\|_\infty$ is returned in FERR(j). See Section 4.4 of Anderson *et al.* (1999) for further details.

8   Parallelism and Performance

F07BBF (DGBSVX) is threaded by NAG for parallel execution in multithreaded implementations of the NAG Library.

F07BBF (DGBSVX) makes calls to BLAS and/or LAPACK routines, which may be threaded within the vendor library used by this implementation. Consult the documentation for the vendor library for further information.

Please consult the X06 Chapter Introduction for information on how to control and interrogate the OpenMP environment used within this routine. Please also consult the Users' Note for your implementation for any additional implementation-specific information.

9   Further Comments

The band storage scheme for the array AB is illustrated by the following example, when $n = 6$, $k_l = 1$, and $k_u = 2$. Storage of the band matrix A in the array AB:

$$
\begin{array}{cccccc}
* & * & a_{13} & a_{24} & a_{35} & a_{46} \\
* & a_{12} & a_{23} & a_{34} & a_{45} & a_{56} \\
a_{11} & a_{22} & a_{33} & a_{44} & a_{55} & a_{66} \\
a_{21} & a_{32} & a_{43} & a_{54} & a_{65} & *
\end{array}
$$

The total number of floating-point operations required to solve the equations $AX = B$ depends upon the pivoting required, but if $n \gg k_l + k_u$ then it is approximately bounded by $O(nk_l(k_l + k_u))$ for the factorization and $O(n(2k_l + k_u)r)$ for the solution following the factorization. The condition number estimation typically requires between four and five solves and never more than eleven solves, following the factorization. The solution is then refined, and the errors estimated, using iterative refinement; see F07BHF (DGBRFS) for information on the floating-point operations required.

In practice the condition number estimator is very reliable, but it can underestimate the true condition number; see Section 15.3 of Higham (2002) for further details.

The complex analogue of this routine is F07BPF (ZGBSVX).

10   Example

This example solves the equations

$$AX = B,$$

where A is the band matrix

$$
A = \begin{pmatrix}
-0.23 & 2.54 & -3.66 & 0 \\
-6.98 & 2.46 & -2.73 & -2.13 \\
0 & 2.56 & 2.46 & 4.07 \\
0 & 0 & -4.78 & -3.82
\end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix}
4.42 & -36.01 \\
27.13 & -31.67 \\
-6.14 & -1.16 \\
10.50 & -25.82
\end{pmatrix}.
$$

Estimates for the backward errors, forward errors, condition number and pivot growth are also output, together with information on the equilibration of A.

10.1   Program Text

```
Program f07bbfe

!     F07BBF Example Program Text

!     Mark 25 Release. NAG Copyright 2014.

!     .. Use Statements ..
      Use nag_library, Only: dgbsvx, nag_wp, x04caf
!     .. Implicit None Statement ..
      Implicit None
!     .. Parameters ..
      Integer, Parameter              :: nin = 5, nout = 6
```

```
!       .. Local Scalars ..
        Real (Kind=nag_wp)                     :: rcond
        Integer                                :: i, ifail, info, j, k, kl, ku, ldab,  &
                                                  ldafb, ldb, ldx, n, nrhs
        Character (1)                          :: equed
!       .. Local Arrays ..
        Real (Kind=nag_wp), Allocatable :: ab(:,:), afb(:,:), b(:,:), berr(:),  &
                                           c(:), ferr(:), r(:), work(:), x(:,:)
        Integer, Allocatable            :: ipiv(:), iwork(:)
!       .. Intrinsic Procedures ..
        Intrinsic                       :: max, min
!       .. Executable Statements ..
        Write (nout,*) 'F07BBF Example Program Results'
        Write (nout,*)
        Flush (nout)
!       Skip heading in data file
        Read (nin,*)
        Read (nin,*) n, nrhs, kl, ku
        ldb = n
        ldx = n
        ldab = kl + ku + 1
        ldafb = ldab + kl
        Allocate (ab(ldab,n),afb(ldafb,n),b(ldb,nrhs),berr(nrhs),c(n), &
          ferr(nrhs),r(n),work(3*n),x(ldx,nrhs),ipiv(n),iwork(n))

!       Read the band matrix A and B from data file

        k = ku + 1
        Read (nin,*)((ab(k+i-j,j),j=max(i-kl,1),min(i+ku,n)),i=1,n)
        Read (nin,*)(b(i,1:nrhs),i=1,n)

!       Solve the equations AX = B for X
!       The NAG name equivalent of dgbsvx is f07bbf
        Call dgbsvx('Equilibration','No transpose',n,kl,ku,nrhs,ab,ldab,afb, &
          ldafb,ipiv,equed,r,c,b,ldb,x,ldx,rcond,ferr,berr,work,iwork,info)

        If ((info==0) .Or. (info==n+1)) Then

!         Print solution, error bounds, condition number, the form
!         of equilibration and the pivot growth factor

!         ifail: behaviour on error exit
!                =0 for hard exit, =1 for quiet-soft, =-1 for noisy-soft
          ifail = 0
          Call x04caf('General',' ',n,nrhs,x,ldx,'Solution(s)',ifail)

          Write (nout,*)
          Write (nout,*) 'Backward errors (machine-dependent)'
          Write (nout,99999) berr(1:nrhs)
          Write (nout,*)
          Write (nout,*) 'Estimated forward error bounds (machine-dependent)'
          Write (nout,99999) ferr(1:nrhs)
          Write (nout,*)
          Write (nout,*) 'Estimate of reciprocal condition number'
          Write (nout,99999) rcond
          Write (nout,*)
          If (equed=='N') Then
            Write (nout,*) 'A has not been equilibrated'
          Else If (equed=='R') Then
            Write (nout,*) 'A has been row scaled as diag(R)*A'
          Else If (equed=='C') Then
            Write (nout,*) 'A has been column scaled as A*diag(C)'
          Else If (equed=='B') Then
            Write (nout,*) &
              'A has been row and column scaled as diag(R)*A*diag(C)'
          End If
          Write (nout,*)
          Write (nout,*) 'Estimate of reciprocal pivot growth factor'
          Write (nout,99999) work(1)

          If (info==n+1) Then
```

```
        Write (nout,*)
        Write (nout,*) 'The matrix A is singular to working precision'
      End If
    Else
      Write (nout,99998) 'The (', info, ',', info, ')', &
        ' element of the factor U is zero'
    End If

99999 Format ((3X,1P,7E11.1))
99998 Format (1X,A,I3,A,I3,A,A)
    End Program f07bbfe
```

10.2 Program Data

```
F07BBF Example Program Data
  4  2  1  2                    :Values of N, NRHS, KL and KU
 -0.23   2.54  -3.66
 -6.98   2.46  -2.73  -2.13
         2.56   2.46   4.07
               -4.78  -3.82    :End of matrix A
  4.42 -36.01
 27.13 -31.67
 -6.14  -1.16
 10.50 -25.82                   :End of matrix B
```

10.3 Program Results

```
 F07BBF Example Program Results

 Solution(s)
             1          2
 1    -2.0000     1.0000
 2     3.0000    -4.0000
 3     1.0000     7.0000
 4    -4.0000    -2.0000

 Backward errors (machine-dependent)
      1.1E-16    9.9E-17

 Estimated forward error bounds (machine-dependent)
      1.6E-14    1.9E-14

 Estimate of reciprocal condition number
      1.8E-02

 A has not been equilibrated

 Estimate of reciprocal pivot growth factor
      1.0E+00
```