

NAG Fortran Library Routine Document

F07GBF (DPPSVX)

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

F07GBF (DPPSVX) uses the Cholesky factorization

$$A = U^T U \quad \text{or} \quad A = LL^T$$

to compute the solution to a real system of linear equations

$$AX = B,$$

where A is an n by n symmetric positive-definite matrix stored in packed format and X and B are n by r matrices. Error bounds on the solution and a condition estimate are also provided.

2 Specification

```

SUBROUTINE F07GBF (FACT, UPLO, N, NRHS, AP, AFP, EQUED, S, B, LDB, X,
1 LDX, RCOND, FERR, BERR, WORK, IWORK, INFO)
    INTEGER          N, NRHS, LDB, LDX, IWORK(*), INFO
    double precision AP(*), AFP(*), S(*), B(LDB,*), X(LDX,*), RCOND,
1 FERR(*), BERR(*), WORK(*)
    CHARACTER*1      FACT, UPLO, EQUED

```

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

3 Description

The following steps are performed:

1. If FACT = 'E', real diagonal scaling factors, D_S , are computed to equilibrate the system:

$$(D_S A D_S)(D_S^{-1} X) = D_S B.$$

Whether or not the system will be equilibrated depends on the scaling of the matrix A , but if equilibration is used, A is overwritten by $D_S A D_S$ and B by $D_S B$.

2. If FACT = 'N' or 'E', the Cholesky decomposition is used to factor the matrix A (after equilibration if FACT = 'E') as $A = U^T U$, if UPLO = 'U', or $A = LL^T$, if UPLO = 'L', where U is an upper triangular matrix and L is a lower triangular matrix.
3. If the leading i by i principal minor is not positive-definite, then the routine returns with INFO = i . Otherwise, the factored 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***, INFO = N + 1 is returned as a warning, but the routine still goes on to solve for X and compute error bounds as described below.
4. The system of equations is solved for X using the factored form of A .
5. Iterative refinement is applied to improve the computed solution matrix and to calculate error bounds and backward error estimates for it.
6. If equilibration was used, the matrix X is premultiplied by D_S 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 URL: <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 factored form of the matrix A is supplied on entry, and if not, whether the matrix A should be equilibrated before it is factored:
 if FACT = 'F' on entry, AFP contains the factored form of A . If EQUED = 'Y', the matrix A has been equilibrated with scaling factors given by S . AP and AFP will not be modified;
 if FACT = 'N', the matrix A will be copied to AFP and factored;
 if FACT = 'E', the matrix A will be equilibrated if necessary, then copied to AFP and factored.
Constraint: FACT = 'F', 'N' or 'E'.
- 2: UPLO – CHARACTER*1 *Input*
On entry: if UPLO = 'U', the upper triangle of A is stored.
 If UPLO = 'L', the lower triangle of A is stored.
Constraint: UPLO = 'U' or 'L'.
- 3: N – INTEGER *Input*
On entry: n , the number of linear equations, i.e., the order of the matrix A .
Constraint: $N \geq 0$.
- 4: 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 .
- 5: AP(*) – **double precision** array *Input/Output*
Note: the dimension of the array AP must be at least $\max(N \times (N + 1)/2)$.
On entry: the upper or lower triangle of the symmetric matrix A , packed columnwise in a linear array, except if FACT = 'F' and EQUED = 'Y', AP must contain the equilibrated matrix $D_S A D_S$. The j th column of A is stored in the array AP as follows:
 if UPLO = 'U', $AP(i + (j - 1) \times j/2) = a_{ij}$ for $1 \leq i \leq j$;
 if UPLO = 'L', $AP(i + (j - 1) \times (2n - j)/2) = a_{ij}$ for $j \leq i \leq n$.
 A is not modified if FACT = 'F' or 'N', or if FACT = 'E' and EQUED = 'N' on exit.
On exit: if FACT = 'E' and EQUED = 'Y', AP is overwritten by $D_S A D_S$.
- 6: AFP(*) – **double precision** array *Input/Output*
Note: the dimension of the array AFP must be at least $\max(1, N)$.
On entry: if FACT = 'F', AFP contains the triangular factor U or L from the Cholesky factorization $A = U^T U$ or $A = L L^T$, in the same storage format as AP. If EQUED \neq 'N', AFP is the factored form of the equilibrated matrix $D_S A D_S$.

On exit: if FACT = 'N', AFP returns the triangular factor U or L from the Cholesky factorization $A = U^T U$ or $A = L L^T$ of the original matrix A .

If FACT = 'E', AFP returns the triangular factor U or L from the Cholesky factorization $A = U^T U$ or $A = L L^T$ of the equilibrated matrix A (see the description of AP for the form of the equilibrated matrix).

7: 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 = 'Y', equilibration was performed, i.e., A has been replaced by $D_S A D_S$.

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

Otherwise, if INFO ≥ 0 , EQUED specifies the form of the equilibration that was performed as specified above.

Constraint: if FACT = 'F', EQUED = 'N' or 'Y'.

8: S(*) – **double precision** array

Input/Output

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

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

If FACT = 'F' and EQUED = 'Y', S must contain the scale factors, D_S , for A ; each element of S must be positive.

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

Otherwise, if INFO ≥ 0 and EQUED = 'Y', S contains the scale factors, D_S , for A ; each element of S is positive.

9: B(LDB,*) – **double precision** 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 EQUED = 'Y', B is overwritten by $D_S B$.

10: LDB – INTEGER

Input

On entry: the first dimension of the array B as declared in the (sub)program from which F07GBF (DPPSVX) is called.

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

11: X(LDX,*) – **double precision** array

Output

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

On exit: if INFO = 0 or INFO = N + 1, the n by r solution matrix X to the original system of equations. Note that if EQUED = 'Y', A and B are modified on exit, and the solution to the equilibrated system is $D_S^{-1} X$.

12: LDX – INTEGER

Input

On entry: the first dimension of the array X as declared in the (sub)program from which F07GBF (DPPSVX) is called.

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

- 13: RCOND – *double precision* *Output*
On exit: if INFO ≥ 0 , an estimate of the reciprocal condition number of the matrix A (after equilibration if that is performed), computed as $\text{RCOND} = 1/(\|A\|_1 \|A^{-1}\|_1)$.
- 14: FERR(*) – *double precision* array *Output*
Note: the dimension of the array FERR must be at least $\max(1, \text{NRHS})$.
On exit: if INFO = 0 or INFO = 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 j th 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.
- 15: BERR(*) – *double precision* array *Output*
Note: the dimension of the array BERR must be at least $\max(1, \text{NRHS})$.
On exit: if INFO = 0 or INFO = N + 1, an estimate of the componentwise 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).
- 16: WORK(*) – *double precision* array *Workspace*
Note: the dimension of the array WORK must be at least $\max(1, 3 \times N)$.
- 17: IWORK(*) – INTEGER array *Workspace*
Note: the dimension of the array IWORK must be at least $\max(1, N)$.
- 18: INFO – INTEGER *Output*
On exit: INFO = 0 unless the routine detects an error (see Section 6).

6 Error Indicators and Warnings

Errors or warnings detected by the routine:

INFO < 0

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

INFO > 0

If INFO = i and $i \leq N$, the leading minor of order i of A is not positive-definite, so the factorization could not be completed, and the solution has not been computed. RCOND = 0 is returned.

If INFO = i and $i = N + 1$, 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.

INFO > 0 and INFO $\leq N$

If INFO = i , the leading minor of order i of A is not positive-definite, so the factorization could not be completed, and the solution has not been computed.

INFO = N + 1

U is nonsingular, but RCOND is less than *machine precision*, so that the matrix A is numerically singular. A solution to the equations $AX = B$, and corresponding error bounds, have nevertheless

been computed because there are some 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 x is the exact solution of a perturbed system of equations $(A + E)x = b$, where

$$|E| \leq c(n)\epsilon|U^T||U|,$$

$c(n)$ is a modest linear function of n , and ϵ is the *machine precision*. See Section 10.1 of Higham (2002) for further details.

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

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

where $\text{cond}(A, \hat{x}, b) = \frac{\|A^{-1}(|A||\hat{x}| + |b|)\|_\infty}{\|\hat{x}\|_\infty} \leq \text{cond}(A) = \|A^{-1}\|_\infty \|A\|_\infty \leq \kappa_\infty(A)$. If \hat{x} is the j th 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 Further Comments

The factorization of A requires approximately $\frac{1}{3}n^3$ floating point operations.

For each right-hand side, computation of the backward error involves a minimum of $4n^2$ floating point operations. Each step of iterative refinement involves an additional $6n^2$ operations. At most 5 steps of iterative refinement are performed, but usually only 1 or 2 steps are required. Estimating the forward error involves solving a number of systems of equations of the form $Ax = b$; the number is usually 4 or 5 and never more than 11. Each solution involves approximately $2n^2$ operations.

The complex analogue of this routine is F07GPF (ZPPSVX).

9 Example

To solve the equations

$$Ax = b,$$

where A is the symmetric positive-definite matrix

$$A = \begin{pmatrix} 4.16 & -3.12 & 0.56 & -0.10 \\ -3.12 & 5.03 & -0.83 & 1.18 \\ 0.56 & -0.83 & 0.76 & 0.34 \\ -0.10 & 1.18 & 0.34 & 1.18 \end{pmatrix}$$

and

$$B = \begin{pmatrix} 8.70 & 8.30 \\ -13.35 & 2.13 \\ 1.89 & 1.61 \\ -4.14 & 5.00 \end{pmatrix}.$$

Error estimates for the solutions, information on equilibration and an estimate of the reciprocal of the condition number of the scaled matrix A are also output.

9.1 Program Text

Note: the listing of the example program presented below uses *bold italicised* terms to denote precision-dependent details. Please read the Users' Note for your implementation to check the interpretation of these terms. As explained in the Essential Introduction to this manual, the results produced may not be identical for all implementations.

```

*      F07GBF Example Program Text
*      Mark 21 Release. NAG Copyright 2004.
*      .. Parameters ..
      INTEGER          NIN, NOUT
      PARAMETER        (NIN=5,NOUT=6)
      INTEGER          NMAX
      PARAMETER        (NMAX=8)
      INTEGER          LDB, LDX, NRHSMX
      PARAMETER        (LDB=NMAX,LDX=NMAX,NRHSMX=NMAX)
      CHARACTER        UPLO
      PARAMETER        (UPLO='U')
*      .. Local Scalars ..
      DOUBLE PRECISION RCOND
      INTEGER          I, IFAIL, INFO, J, N, NRHS
      CHARACTER        EQUED
*      .. Local Arrays ..
      DOUBLE PRECISION AFP((NMAX*(NMAX+1))/2), AP((NMAX*(NMAX+1))/2),
+      B(LDB,NRHSMX), BERR(NRHSMX), FERR(NRHSMX),
+      S(NMAX), WORK(3*NMAX), X(LDX,NRHSMX)
      INTEGER          IWORK(NMAX)
*      .. External Subroutines ..
      EXTERNAL        DPPSVX, X04CAF
*      .. Executable Statements ..
      WRITE (NOUT,*) 'F07GBF Example Program Results'
      WRITE (NOUT,*)
*      Skip heading in data file
      READ (NIN,*)
      READ (NIN,*) N, NRHS
      IF (N.LE.NMAX .AND. NRHS.LE.NRHSMX) THEN
*
*          Read the upper or lower triangular part of the matrix A from
*          data file
*
*          IF (UPLO.EQ.'U') THEN
*              READ (NIN,*) ((AP(I+(J*(J-1))/2),J=I,N),I=1,N)
*          ELSE IF (UPLO.EQ.'L') THEN
*              READ (NIN,*) ((AP(I+((2*N-J)*(J-1))/2),J=1,I),I=1,N)
*          END IF
*
*          Read B from data file
*
*          READ (NIN,*) ((B(I,J),J=1,NRHS),I=1,N)
*
*          Solve the equations AX = B for X
*
*          CALL DPPSVX('Equilibration',UPLO,N,NRHS,AP,AFP,EQUED,S,B,LDB,X,
+          LDX,RCOND,FERR,BERR,WORK,IWORK,INFO)
*
*          IF ((INFO.EQ.0) .OR. (INFO.EQ.N+1)) THEN
*
*              Print solution, error bounds, condition number and the form
*              of equilibration
*
*              IFAIL = 0
*              CALL X04CAF('General',' ',N,NRHS,X,LDX,'Solution(s)',IFAIL)
*
*              WRITE (NOUT,*)
*              WRITE (NOUT,*) 'Backward errors (machine-dependent)'
*              WRITE (NOUT,99999) (BERR(J),J=1,NRHS)
*              WRITE (NOUT,*)
*              WRITE (NOUT,*)
*              WRITE (NOUT,*) 'Estimated forward error bounds (machine-dependent)'
+              WRITE (NOUT,99999) (FERR(J),J=1,NRHS)
*              WRITE (NOUT,*)
*              WRITE (NOUT,*) 'Estimate of reciprocal condition number'

```

```

      WRITE (NOUT,99999) RCOND
      WRITE (NOUT,*)
      IF (EQUED.EQ.'N') THEN
        WRITE (NOUT,*) 'A has not been equilibrated'
      ELSE IF (EQUED.EQ.'S') THEN
        WRITE (NOUT,*)
+       'A has been row and column scaled as diag(S)*A*diag(S)'
      END IF
*
      IF (INFO.EQ.N+1) THEN
        WRITE (NOUT,*)
        WRITE (NOUT,*)
+       'The matrix A is singular to working precision'
      END IF
      ELSE
        WRITE (NOUT,99998) 'The leading minor of order ', INFO,
+       ' is not positive definite'
      END IF
      ELSE
        WRITE (NOUT,*) 'NMAX and/or NRHSMX too small'
      END IF
      STOP
*
99999 FORMAT ((3X,1P,7E11.1))
99998 FORMAT (1X,A,I3,A)
      END

```

9.2 Program Data

F07GBF Example Program Data

4	2			:Values of N and NRHS
4.16	-3.12	0.56	-0.10	
	5.03	-0.83	1.18	
		0.76	0.34	
			1.18	:End of matrix A
8.70	8.30			
-13.35	2.13			
1.89	1.61			
-4.14	5.00			:End of matrix B

9.3 Program Results

F07GBF Example Program Results

Solution(s)

	1	2
1	1.0000	4.0000
2	-1.0000	3.0000
3	2.0000	2.0000
4	-3.0000	1.0000

Backward errors (machine-dependent)

8.3E-17	5.2E-17
---------	---------

Estimated forward error bounds (machine-dependent)

2.4E-14	2.2E-14
---------	---------

Estimate of reciprocal condition number

1.0E-02

A has not been equilibrated
