

F08FEFP (PDSYTRD)

NAG Parallel Library Routine Document

Note: before using this routine, please read the Users' Note for your implementation to check for implementation-dependent details. You are advised to enclose any calls to NAG Parallel Library routines between calls to Z01AAFP and Z01ABFP.

1 Description

F08FEFP (PDSYTRD) reduces an n by n real symmetric matrix A_s to tridiagonal form T by an orthogonal similarity transformation Q

$$Q^T A_s Q = T,$$

where A_s is a submatrix of an m_A by n_A matrix A , i.e.,

$$A_s \equiv A(i_A : i_A + n - 1, j_A : j_A + n - 1).$$

Note: if $i_A = j_A = 1$ and $n = m_A = n_A$, then $A_s \equiv A$.

Since the matrix A_s is real symmetric, only the the upper triangular part or the lower triangular part is required.

The diagonal elements of the tridiagonal matrix T are represented by a vector d of length n and the off-diagonal elements by a vector e . On exit, the vector d is distributed in the one-dimensional block cyclic form across each logical processor row of the Library Grid. The vector e is similarly distributed.

The orthogonal matrix Q is not formed explicitly but is represented as a product of $n - 1$ elementary reflectors. See the F08 Chapter Introduction for details of the distributions of Q , d and e .

The routine is designed to be used as the first step in computing the eigenvalues of A_s . It should be followed by calls to F01ZFPF to gather the vectors d and e onto each logical processor. Eigenvalues of T , which are the same as the eigenvalues of A_s , can then be computed by calling F08JFFP (PDSTEBZ).

2 Specification

```

SUBROUTINE F08FEFP(UPLO, N, A, IA, JA, IDESCA, D, E, TAU, WORK,
1                LWORK, INFO)
ENTRY          PDSYTRD(UPLO, N, A, IA, JA, IDESCA, D, E, TAU, WORK,
1                LWORK, INFO)
DOUBLE PRECISION A(*), D(*), E(*), TAU(*), WORK(*)
INTEGER         N, IA, JA, IDESCA(*), LWORK, INFO
CHARACTER*1     UPLO

```

The ENTRY statement enables the routine to be called by its ScaLAPACK name.

3 Usage

3.1 Definitions

The following definitions are used in describing the data distribution within this document:

- m_p – the number of rows in the Library Grid.
- n_p – the number of columns in the Library Grid.
- p_r – the row grid coordinate of the calling processor.
- p_c – the column grid coordinate of the calling processor.
- M_b^X – the blocking factor for the distribution of the rows of a matrix X .
- N_b^X – the blocking factor for the distribution of the columns of a matrix X .
- s_r^X – the row coordinate of the processor that possesses the first row of a distributed matrix X .
- s_c^X – the column coordinate of the processor that possesses the first column of a distributed matrix X .

- numroc($\hat{\ell}, L_b^X, p, s^X, \ell_p$) – a function which gives the **number of rows or columns** of a distributed matrix X owned by the processor with the row or column coordinate p (p_r or p_c), where $\hat{\ell}$ is the total number of rows or columns of the matrix, L_b^X is the blocking factor used (M_b^X or N_b^X), s^X is the row or column coordinate (s_r^X or s_c^X) of the processor that possesses the first row or column of the distributed matrix and ℓ_p is either m_p or n_p . The Library provides the function Z01CAFP (NUMROC) for the evaluation of this function.
- indxg2p(k, L_b^X, p, s^X, ℓ_p) – a function which gives the processor row or column coordinate which possess the row or column index k of a distributed matrix. The arguments L_b^X , s^X and ℓ_p have the same meaning as in the function numroc. However, the argument p is a dummy integer. The Library provides the utility function Z01CDFP (INDXG2P) for the evaluation of this function.

3.2 Global and Local Arguments

The following global **input** arguments must have the same value on entry to the routine on each processor and the global **output** arguments will have the same value on exit from the routine on each processor:

Global input arguments: UPLO, N, IA, JA, the array elements IDESCA(1) and IDESCA(3:8)
 Global output arguments: INFO

The remaining arguments are local.

3.3 Distribution Strategy

On entry, the matrix A must be stored in the cyclic two-dimensional block format and its descriptor array IDESCA must contain the details of the distributed matrix. The indices i_A and j_A and the order n identify the submatrix A_s . See the F08 Chapter Introduction for further details. The matrices A and A_s must satisfy the following requirements:

$$\begin{aligned} M_b^A &= N_b^A; \\ \text{mod}(i_A - 1, M_b^A) &= 0; \\ \text{mod}(j_A - 1, N_b^A) &= 0. \end{aligned}$$

Any further constraints, including the above, are stated in Section 4 under each argument.

This routine is the first step in the solution of the eigenvalue problem of a dense matrix A . However, before proceeding to eigenroutines, it is often necessary to gather complete copies of the vectors d and e to every logical processor on the grid. The Library provides the routine F01ZPFP for this particular gathering operation. The following prerequisites are mandatory on F08FEFP for post-processing by F01ZPFP:

$$\begin{aligned} i_A &= j_A = 1; \\ s_r^A &= s_c^A = 0. \end{aligned}$$

It is assumed that the data has already been correctly distributed, and if this is not the case, then this routine will fail to produce correct results.

3.4 Related Routines

The Library provides many support routines for the generation/distribution and input/output of data. The following routines may be used in conjunction with F08FEFP (PDSYTRD):

Real matrix output: X04BFFP
 Real matrix generation: F01ZQFP
 Real matrix gather: Z01ZPFP

4 Arguments

1: UPLO — CHARACTER*1 *Global Input*

On entry: indicates whether the upper or lower triangular part of A_s is stored, as follows:

- if UPLO = 'U', then the upper triangular part of A_s is stored;
- if UPLO = 'L', then the lower triangular part of A_s is stored.

Constraint: UPLO = 'U' or 'L'.

2: N — INTEGER *Global Input*

On entry: n , the order of the matrix A_s .

Constraints: $0 \leq N \leq \min(\text{IDESCA}(3), \text{IDESCA}(4))$.

3: A(*) — DOUBLE PRECISION array *Local Input/Local Output*

Note: array A is formally defined as a vector. However, you may find it more convenient to consider A as a two-dimensional array of dimension $(\text{IDESCA}(9), \gamma)$, where $\gamma \geq \text{numroc}(\text{JA}-1+N, \text{IDESCA}(6), p_c, \text{IDESCA}(8), n_p)$.

On entry: the local part of the n by n matrix A_s to be factorized.

If UPLO = 'U', the upper triangle of A_s must be stored and the elements of the matrix below the diagonal are not referenced;

if UPLO = 'L', the lower triangle of A_s must be stored and the elements of the matrix above the diagonal are not referenced.

On exit: The locally held parts of A corresponding to the matrix A_s are overwritten by the elements of the tridiagonal matrix T and the details of the orthogonal matrix Q .

4: IA — INTEGER *Global Input*

On entry: i_A , the row index of matrix A that identifies the first row of the submatrix A_s to be factorized.

Note: i_A must be equal to 1 if this routine is to be followed by a call to F01ZFPF.

Constraints:

$$\begin{aligned} \text{mod}(\text{IA}-1, \text{IDESCA}(5)) &= 0; \\ 1 \leq \text{IA} &\leq \text{IDESCA}(3) - N + 1. \end{aligned}$$

5: JA — INTEGER *Global Input*

On entry: j_A , the column index of matrix A that identifies the first column of the submatrix A_s to be factorized.

Note: j_A must be equal to 1 if this routine is to be followed by a call to F01ZFPF.

Constraints:

$$\begin{aligned} \text{mod}(\text{JA}-1, \text{IDESCA}(6)) &= 0; \\ 1 \leq \text{JA} &\leq \text{IDESCA}(4) - N + 1. \end{aligned}$$

- 6:** IDESCA(*) — INTEGER array *Local Input*

Note: the dimension of the array IDESCA must be at least 9.

Distribution: the array elements IDESCA(1) and IDESCA(3),...,IDESCA(8) must be global to the processor grid and the elements IDESCA(2) and IDESCA(9) are local to each processor.

On entry: the description array for the matrix A . This array must contain details of the distribution of the matrix A and the logical processor grid.

IDESCA(1), the descriptor type. For this routine, which uses a cyclic two-dimensional block distribution, IDESCA(1) = 1;

IDESCA(2), the Library context, usually returned by a call to the Library Grid initialisation routine Z01AAFP;

IDESCA(3), the number of rows, m_A , of the matrix A ;

IDESCA(4), the number of columns, n_A , of the matrix A ;

IDESCA(5), the blocking factor, M_b^A , used to distribute the rows of the matrix A ;

IDESCA(6), the blocking factor, N_b^A , used to distribute the columns of the matrix A ;

IDESCA(7), the processor row index, s_r^A , over which the first row of the matrix A is distributed;

IDESCA(8), the processor column index, s_c^A , over which the first column of the matrix A is distributed;

IDESCA(9), the leading dimension of the conceptual two-dimensional array A .

Constraints:

IDESCA(1) = 1;

IDESCA(3) \geq 0; IDESCA(4) \geq 0;

IDESCA(5) = IDESCA(6) \geq 1;

0 \leq IDESCA(7) \leq $m_p - 1$; 0 \leq IDESCA(8) \leq $n_p - 1$;

IDESCA(9) \geq max(1,numroc(IDESCA(3),IDESCA(5), p_r ,IDESCA(7), m_p)).

- 7:** D(*) — DOUBLE PRECISION array *Local Output*

Note: the dimension of the array D must be at least numroc(JA-1+N,IDESCA(6), p_c ,IDESCA(8), n_p).

On exit: the local parts of the distributed vector d which contains the diagonal elements of the tridiagonal matrix T . The vector d is distributed in one-dimensional block cyclic form across each logical processor row of the two-dimensional logical processor grid. See the F08 Chapter Introduction for further details.

- 8:** E(*) — DOUBLE PRECISION array *Local Output*

Note: the dimension of the array E must be at least numroc(JA-1+N,IDESCA(6), p_c ,IDESCA(8), n_p).

On exit: the local parts of the distributed vector e , which contains the off-diagonal elements of the tridiagonal matrix T : if UPLO = 'U', $e_1 = 0$ and $e_i = t_{i-1,i}$ for $i = 2, \dots, n$; if UPLO = 'L', $e_i = t_{i+1,i}$ for $i = 1, \dots, n - 1$ and $e_n = 0$. The vector e is distributed in the same manner as d .

- 9:** TAU(*) — DOUBLE PRECISION array *Local Output*

Note: the dimension of the array TAU must be at least numroc(JA-1+N,IDESCA(6), p_c ,IDESCA(8), n_p).

On exit: the scalar factors τ_i of the elementary reflectors which define the orthogonal matrix Q . See the F08 Chapter Introduction for further details.

- 10:** WORK(*) — DOUBLE PRECISION array *Local Workspace/Local Output*

Note: the dimension of WORK must be at least max(1,LWORK). The minimum value of LWORK required to successfully call this routine can be obtained by setting LWORK = -1. The required size is returned in array element WORK(1).

On exit: WORK(1) contains the minimum dimension of the array WORK required to successfully complete the task.

11: LWORK — INTEGER*Local Input*

On entry: either -1 (see WORK) or the dimension of the array WORK required to successfully complete the task. If LWORK is set to -1 on entry this routine simply performs some initial error checking and then, if these checks are successful, calculates the minimum size of LWORK required.

Constraints:

either $LWORK = -1$,
 or $LWORK \geq IDESCA(5) \times \max(\alpha+1, 3)$ or where
 $\alpha = \text{numroc}(N, IDESCA(5), p_r, \beta, m_p)$;
 $\beta = \text{indxg2p}(IA, IDESCA(5), p_r, IDESCA(7), m_p)$.

12: INFO — INTEGER*Global Output*

The NAG Parallel Library provides a mechanism, via the routine Z02EAFP, to reduce the amount of parameter validation performed by this routine. For a full description refer to the Z02 Chapter Introduction.

On exit: INFO = 0 (or -9999 if reduced error checking is enabled) unless the routine detects an error (see Section 5).

5 Errors and Warnings

If INFO $\neq 0$ an explanatory message is output and control returned to the calling program.

INFO < 0

On entry, one of the arguments was invalid:

if the k th argument is a scalar INFO = $-k$;
 if the k th argument is an array and its j th element is invalid, INFO = $-(100 \times k + j)$.

This error occurred either because a global argument did not have the same value on all logical processors, or because its value on one or more processors was incorrect.

6 Further Comments

6.1 Algorithmic Detail

The total number of floating-point operations is approximately $\frac{4}{3}n^3$.

6.2 Parallelism Detail

The BLAS operations used in this routine are carried out in parallel.

6.3 Accuracy

The computed tridiagonal matrix T is exactly similar to a nearby matrix $A + E$, where

$$\|E\|_2 \leq p(n)\epsilon\|A\|_2,$$

$p(n)$ is a modestly increasing function of n , and ϵ is the *machine precision*.

The elements of T themselves may be sensitive to small perturbations in A or to rounding errors in the computation, but this does not affect the stability of the eigenvalues and eigenvectors.

7 References

- [1] 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

- [2] Golub G H and van Loan C F (1996) *Matrix Computations* Johns Hopkins University Press (3rd Edition), Baltimore
- [3] Blackford L S, Choi J, Cleary A, D'Azevedo E, Demmel J, Dhillon I, Dongarra J, Hammarling S, Henry G, Petitet A, Stanley K, Walker D and Whaley R C (1997) *ScaLAPACK Users' Guide* SIAM 3600 University City Science Center, Philadelphia, PA 19104-2688, USA. URL: http://www.netlib.org/scalapack/slug/scalapack_slug.html

8 Example

The example program illustrates the reduction of a 9 by 9 symmetric matrix A_s to tridiagonal form T using F08FEFP (PDSYTRD). The (i, j) th element of this matrix A_s is $\max(i, j)$ and the matrix is generated using the routine F01ZQFP. The argument UPLO is set to 'U' and hence only the upper triangular part of the matrix A_s is used.

The diagonal vector d and the off-diagonal vector e are gathered to every logical processor using the routine F01ZFPF. Finally, the vector d (denoted locally by the array DL) and the the vector e (denoted locally by the array EL) are printed on the root processor.

8.1 Example Text

```
*      F08FEFP Example Program Text
*      NAG Parallel Library Release 3 Revised. NAG Copyright 1999.
*      .. Parameters ..
INTEGER          NOUT
PARAMETER       (NOUT=6)
INTEGER          N
PARAMETER       (N=9)
INTEGER          DT
PARAMETER       (DT=1)
INTEGER          NA
PARAMETER       (NA=20)
INTEGER          NB
PARAMETER       (NB=2)
INTEGER          MP, NP
PARAMETER       (MP=2, NP=2)
INTEGER          LDA, TDA
PARAMETER       (LDA=NA/MP+NB, TDA=NA/NP+NB)
INTEGER          LWORK
PARAMETER       (LWORK=25)
*      .. Local Scalars ..
INTEGER          I, IA, ICNTXT, IFAIL, INFO, JA, MPROC, NPROC
LOGICAL          ROOT
CHARACTER        UPLO
*      .. Local Arrays ..
DOUBLE PRECISION A(LDA,TDA), D(TDA), DL(N), E(TDA), EL(N),
+               TAU(TDA), WORK(LWORK)
INTEGER          IDESCA(9)
*      .. External Functions ..
LOGICAL          Z01ACFP
EXTERNAL         Z01ACFP
*      .. External Subroutines ..
EXTERNAL         F01ZFPF, F01ZQFP, F08FEFP, GMATA, Z01AAFP,
+               Z01ABFP
*      .. Executable Statements ..
ROOT = Z01ACFP()
IF (ROOT) THEN
    WRITE (NOUT,*) 'F08FEFP Example Program Results'
```

```

        WRITE (NOUT,*)
        END IF
*
        MPROC = MP
        NPROC = NP
        IFAIL = 0
        CALL Z01AAFP(ICNTXT,MPROC,NPROC,IFAIL)
*
*      Set the starting address and array descriptor for A
*
        IA = 1
        JA = 1
        IDESCA(1) = DT
        IDESCA(2) = ICNTXT
        IDESCA(3) = NA
        IDESCA(4) = NA
        IDESCA(5) = NB
        IDESCA(6) = NB
        IDESCA(7) = 0
        IDESCA(8) = 0
        IDESCA(9) = LDA
*
*      Generate and distribute the matrix A
*
        IFAIL = 0
        CALL F01ZQFP(GMATA,N,N,A,IA,JA,IDESCA,IFAIL)
*
*      Reduce A to tridiagonal form
*
        UPLO = 'U'
        CALL F08FEFP(UPLO,N,A,IA,JA,IDESCA,D,E,TAU,WORK,LWORK,INFO)
*
*      Gather the diagonal D to each logical processor
*
        IFAIL = 0
        CALL F01ZFPF(N,IA,JA,IDESCA,D,DL,WORK,LWORK,IFAIL)
*
*      Gather the off-diagonal E to each logical processor
*
        IFAIL = 0
        CALL F01ZFPF(N,IA,JA,IDESCA,E,EL,WORK,LWORK,IFAIL)
*
*      Print the diagonal and off-diagonal elements of the tridiagonal
*      matrix on the root processor
*
        IF (ROOT) THEN
            WRITE (NOUT,*) 'Elements of the tridiagonal matrix T'
            WRITE (NOUT,*)
            IF (UPLO.EQ.'U') THEN
                WRITE (NOUT,'(3X,''I'',12X,''DL(I)'',9X,''EL(I)'')/')
                WRITE (NOUT,'(1X,I3,5X,2(F12.4,2X))') 1, DL(1)
                DO 20 I = 2, N
                    WRITE (NOUT,'(1X,I3,5X,2(F12.4,2X))') I, DL(I), EL(I)
20             CONTINUE
            ELSE IF (UPLO.EQ.'L') THEN
                WRITE (NOUT,'(3X,''I'',12X,''DL(I)'',9X,''EL(I)'')/')
                DO 40 I = 1, N - 1

```

```

        WRITE (NOUT,'(1X,I3,5X,2(F12.4,2X))') I, DL(I), EL(I)
40      CONTINUE
        WRITE (NOUT,'(1X,I3,5X,2(F12.4,2X))') N, DL(N)
        END IF
      END IF
*
      IFAIL = 0
      CALL Z01ABFP(ICNTXT,'N',IFAIL)
*
      STOP
      END
*
      SUBROUTINE GMATA(I1,I2,J1,J2,AL,LDAL)
*
*      GMATA generates the block A(I1: I2, J1: J2) of the upper
*      triangular part of the symmetric matrix A such that
*
*       $a(i,j) = \max(i,j)$ 
*
*      in the array AL.
*
*      .. Scalar Arguments ..
      INTEGER          I1, I2, J1, J2, LDAL
      CHARACTER        UPLO
*
*      .. Array Arguments ..
      DOUBLE PRECISION AL(LDAL,*)
*
*      .. Local Scalars ..
      INTEGER          I, J, K, L
*
*      .. Intrinsic Functions ..
      INTRINSIC        MAX
*
*      .. Executable Statements ..
      L = 1
      DO 40 J = J1, J2
        K = 1
        DO 20 I = I1, I2
          IF (I.LE.J) THEN
            AL(K,L) = MAX(I,J)
          END IF
          K = K + 1
        20  CONTINUE
        L = L + 1
      40  CONTINUE
      RETURN
      END

```

8.2 Example Data

None.

8.3 Example Results

F08FEFP Example Program Results

Elements of the tridiagonal matrix T

I	DL(I)	EL(I)
1	-0.2778	
2	-0.3309	-0.0248
3	-0.4176	-0.0583
4	-0.5737	-0.1151
5	-0.9000	-0.2235
6	-1.7857	-0.4737
7	-6.2143	-1.2963
8	46.5000	9.4472
9	9.0000	-25.4558
