

F02WQFP

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

F02WQFP computes the Singular Value Decomposition (SVD) of a real matrix whose columns are distributed on a two-dimensional Library Grid.

The SVD of an m by n real rectangular matrix A (where $m \geq n$) is defined as

$$A = \begin{bmatrix} U | \tilde{U} \end{bmatrix} \begin{bmatrix} \Sigma \\ 0 \end{bmatrix} V^T = U \Sigma V^T$$

where U is an $m \times n$ matrix of left singular vectors, \tilde{U} is an $m \times (m - n)$ matrix, V is an $n \times n$ orthogonal matrix of right singular vectors, and Σ is an $n \times n$ diagonal matrix of singular values. The singular values $\sigma_1, \sigma_2, \dots, \sigma_n$ are non-negative and in non-increasing order of magnitude. The matrix $\begin{bmatrix} U | \tilde{U} \end{bmatrix}$ is orthogonal.

For the case, $m < n$ the SVD may be defined in the form

$$A = U \begin{bmatrix} \Sigma | 0 \end{bmatrix} \begin{bmatrix} V | \tilde{V} \end{bmatrix}^T = U \Sigma V^T$$

where U is an $m \times m$ orthogonal matrix, V is an $n \times m$ matrix of right singular vectors, \tilde{V} is an $n \times (n - m)$ matrix, and Σ is an $m \times m$ diagonal matrix of singular values. The matrix $\begin{bmatrix} V | \tilde{V} \end{bmatrix}$ is orthogonal.

In general, matrices \tilde{V} and \tilde{U} are not unique. For convenience in the description of the routine, the definition of singular values is extended when $m < n$ by defining $\sigma_{m+1} = \dots = \sigma_n = 0$. In that case, the columns of \tilde{V} may be considered as right singular vectors corresponding to zero singular values.

F02WQFP computes only the left singular vectors U which correspond to non-zero singular values. Optionally all right singular vectors V (including \tilde{V} if $m < n$) are computed.

2 Specification

```
SUBROUTINE F02WQFP(ICNTXT, M, N, A, LDA, VWANT, NX, NR, IFAIL)
DOUBLE PRECISION  A(0:LDA-1,0:*)
INTEGER            ICNTXT, M, N, LDA, NX, NR, IFAIL
LOGICAL           VWANT
```

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 – $m_p \times n_p$, the total number of processors in the Library Grid.
- p_d – the number of logical processors which hold columns of the matrix A .
- N_b – the maximum number of columns of the matrix A held locally on a logical processor.
- N_x – the actual number of columns of the matrix A held locally on a logical processor where $0 \leq N_x \leq N_b$.
- n_r – the number of non-zero singular values computed by a logical processor, where $0 \leq n_r \leq N_x$.
- $[x]$ – the ceiling function of x , which gives the smallest integer which is not less than x .

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: M, N, VWANT, IFAIL

Global output arguments: IFAIL

The remaining arguments are local.

3.3 Distribution Strategy

Columns of the matrix A are allocated to logical processors on the Library Grid row by row (i.e., in row major ordering of the grid) starting from the $\{0,0\}$ logical processor. Each logical processor that contains columns of the matrix contains $N_b = \lceil n/p \rceil$ columns, except the last processor that actually contains data, for which the number of columns held may be less than N_b . This processor will contain $\text{mod}(n, N_b)$ columns if $\text{mod}(n, N_b) \neq 0$, and will contain N_b columns otherwise. Some logical processors may not contain any columns of the matrix if n is not large relative to p , but if $n > (p-1)^2$ then all processors will certainly contain columns of the matrix.

The number of logical processors that contain columns of the matrix is given by $p_d = \lceil n/N_b \rceil$.

The following example illustrates a case where the last processor with data is not the last processor of the grid. Furthermore the number of columns on the last processor with data is not equal to the number of columns on other processors.

If $m_p = 2$, $n_p = 4$ then $p = m_p \times n_p = 8$. If $n = 41$ then $N_b = \lceil n/p \rceil = \lceil 5.125 \rceil = 6$, $\text{mod}(n, N_b) = 5 \neq 0$ and $p_d = \lceil n/N_b \rceil = \lceil 6.833 \rceil = 7$.

processor $\{0,0\}$ $N_x = 6$ columns (1:6)	processor $\{0,1\}$ $N_x = 6$ columns (7:12)	processor $\{0,2\}$ $N_x = 6$ columns (13:18)	processor $\{0,3\}$ $N_x = 6$ columns (19:24)
processor $\{1,0\}$ $N_x = 6$ columns (25:30)	processor $\{1,1\}$ $N_x = 6$ columns (31:36)	processor $\{1,2\}$ $N_x = 5$ columns (37:41)	processor $\{1,3\}$ $N_x = 0$

If the data is distributed incorrectly, the routine may fail to produce correct results or will exit with an error flag.

3.4 Related Routines

The Library provides support routines for the generation, scattering/gathering and input/output of matrices whose columns are distributed on the Library Grid. The following routines may be used in conjunction with F02WQFP:

Real matrix generation: F01ZRFP

Real matrix output: X04BFFP

4 Arguments

1: ICNTXT — INTEGER *Local Input*

On entry: the Library context, usually returned by a call to the Library Grid initialisation routine Z01AAFP.

Note: the value of ICNTXT **must not** be changed.

2: M — INTEGER *Global Input*

On entry: m , the number of rows of A .

Constraint: $M \geq 0$.

- 3:** N — INTEGER *Global Input*
On entry: n , the number of columns of A .
Constraint: $N \geq 0$.
- 4:** A(0:LDA-1,0:*) — DOUBLE PRECISION array *Local Input/Local Output*
Note: the size of the second dimension of the array A must be at least $N_x + 1$ where N_x is the number of columns of A held locally by the logical processor. The array A is not referenced if $N_x = 0$.
On entry: A(1:m,1: N_x) must contain columns of the matrix A , as defined by the distribution strategy (see Section 3.3).
On exit: A(0,1: N_x) contains N_x singular values of the matrix A stored on this logical processor. They are ordered locally and globally (in the row major ordering of the processors) in non-increasing order of magnitude.
A(1:m,1: n_r) contains the left singular vectors corresponding to non-zero singular values. A left singular vector is not computed if the corresponding singular value and the corresponding elements of A are set to zero.
If VWANT = .TRUE., then A($m+1 : m+n$, 1 : N_x) contains the right singular vectors corresponding to the singular values held on this logical processor.
The remainder of the array is used as workspace and contains no useful information.
- 5:** LDA — INTEGER *Local Input*
On entry: the size of the first dimension of the array A as declared in the (sub)program from which F02WQFP is called.
Constraint: $LDA \geq M + N + 2$ if VWANT = .TRUE.; otherwise $LDA \geq M + 2$.
- 6:** VWANT — LOGICAL *Global Input*
On entry: VWANT must be set to .TRUE. if right singular vectors are required.
- 7:** NX — INTEGER *Local Output*
On exit: N_x , the actual number of columns of the matrix A held on the logical processor.
- 8:** NR — INTEGER *Local Output*
On exit: n_r , the number of non-zero singular values held on the logical processor.
- 9:** IFAIL — INTEGER *Global Input/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 entry: IFAIL must be set to 0, -1 or 1. For users not familiar with this argument (described in the Essential Introduction) the recommended values are:
 IFAIL = 0, if multigridding is **not** employed;
 IFAIL = -1, if multigridding is employed.
On exit: IFAIL = 0 (or -9999 if reduced error checking is enabled) unless the routine detects an error (see Section 5).

5 Errors and Warnings

If on entry $IFAIL = 0$ or -1 , explanatory error messages are output from the root processor (or processor $\{0,0\}$ when the root processor is not available) on the current error message unit (as defined by X04AAF).

5.1 Full Error Checking Mode Only

$IFAIL = -2000$

The routine has been called with an invalid value of ICNTXT on one or more processors.

$IFAIL = -1000$

The logical processor grid and library mechanism (Library Grid) have not been correctly defined, see Z01AAFP.

$IFAIL = -i$

On entry, the i th argument was invalid. 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. An explanatory message distinguishes between these two cases.

5.2 Any Error Checking Mode

$IFAIL = 1$

The Jacobi algorithm has not converged.

6 Further Comments

6.1 Algorithmic Detail

The algorithm is based on a one-sided Jacobi method, see Hestenes [1].

6.2 Parallelism Detail

The algorithm uses a linear array of logical processors. This linear array is mapped to the Library Grid based on the row major ordering beginning from the $\{0,0\}$ logical processor on the two-dimensional array. Most of the communication is between neighbours on the linear array of processors.

6.3 Accuracy

The computed factors U , Σ and V satisfy the relation

$$U\Sigma V^T = A + E,$$

where

$$\|E\| \leq c\epsilon\|A\|,$$

ϵ being the *machine precision*, c is a modest function of m and n and $\|\cdot\|$ denotes the 2-norm.

7 References

- [1] Hestenes M R (1958) Inversion of matrices by biorthogonalization and related results *J. SIAM* **6** 51–90

8 Example

To find the singular value decomposition of the 4 by 7 matrix A given by

$$A = \left(\begin{array}{cc|cc|cc|c} 1.0 & 1.0 & 1.0 & 1.0 & 1.0 & 1.0 & 1.0 \\ 0.0 & 1.0 & 1.0 & 1.0 & 1.0 & 1.0 & 1.0 \\ 0.0 & 0.0 & 1.0 & 1.0 & 1.0 & 1.0 & 1.0 \\ 0.0 & 0.0 & 0.0 & 1.0 & 1.0 & 1.0 & 1.0 \end{array} \right)$$

and to print results on the root processor. The routine F01ZRFP is used to generate the matrix A on a 2 by 2 logical processor grid. The number of columns of the matrix A on each logical processor, N_x , is equal to 2 on logical processors {0,0}, {0,1}, and {1,0}. On the final logical processor {1,1}, $N_x = 1$. This blocking is indicated by the vertical lines in the matrix A above.

The routine X04BFFP is used to bring results to the root processor and print the results. The results are the non-zero singular values, and the corresponding left and right singular vectors.

8.1 Example Text

```
*      F02WQFP Example Program Text
*      NAG Parallel Library Release 2. NAG Copyright 1996.
*      .. Parameters ..
      INTEGER          NOUT
      PARAMETER       (NOUT=6)
      INTEGER          M, N, MM, NM
      PARAMETER       (M=4,N=7,MM=20,NM=20)
      INTEGER          MG, NG
      PARAMETER       (MG=2,NG=2)
      INTEGER          LDA, TDA
      PARAMETER       (LDA=MM+NM+2,TDA=(NM/(MG*NG)+2))
      CHARACTER*25     FORMT
      PARAMETER       (FORMT='F12.4')
      LOGICAL          VWANT
      PARAMETER       (VWANT=.TRUE.)
*      .. Local Scalars ..
      INTEGER          ICNTXT, ICOFF, IFAIL, MP, NP, NR, NX
      LOGICAL          ROOT
      CHARACTER        CNUMOP, TITOP
*      .. Local Arrays ..
      DOUBLE PRECISION A(0:LDA-1,0:TDA-1), W(LDA,TDA)
*      .. External Functions ..
      INTEGER          OFFSET
      LOGICAL          Z01ACFP
      EXTERNAL         OFFSET, Z01ACFP
*      .. External Subroutines ..
      EXTERNAL         F01ZRFP, F02WQFP, GMATA, X04BFFP, Z01AAFP,
+                     Z01ABFP
*      .. Executable Statements ..
      ROOT = Z01ACFP()
*
      IF (ROOT) THEN
          WRITE (NOUT,*) 'F02WQFP Example Program Results'
          WRITE (NOUT,*)
      END IF
*
*      Define the 2D processor grid
*
      MP = MG
      NP = NG
      IFAIL = 0
```

```

*
  CALL Z01AAFP(ICNTXT,MP,NP,IFAIL)
*
*   Generate the matrix A
*
  IFAIL = 0
*
  CALL F01ZRFP(ICNTXT,GMATA,M,N,A(1,1),LDA,NX,IFAIL)
*
*   Compute the SVD
*
  IFAIL = 0
*
  CALL F02WQFP(ICNTXT,M,N,A,LDA,VWANT,NX,NR,IFAIL)
*
*   Print singular values
*
  IF (ROOT) THEN
    WRITE (NOUT,*) 'Singular values'
    WRITE (NOUT,*)
    TITOP = 'N'
    CNUMOP = 'G'
  END IF
  ICOFF = 0
  IFAIL = 0
*
  CALL X04BFFP(ICNTXT,NOUT,1,NR,A(0,1),LDA,FORMAT,TITOP,CNUMOP,ICOFF,
+             W,LDA,IFAIL)
*
*   Print left singular vectors
*
  IF (ROOT) THEN
    WRITE (NOUT,*) 'Left singular vectors'
    WRITE (NOUT,*)
  END IF
  IFAIL = 0
*
  CALL X04BFFP(ICNTXT,NOUT,M,NR,A(1,1),LDA,FORMAT,TITOP,CNUMOP,ICOFF,
+             W,LDA,IFAIL)
*
*   Print right singular vectors (full set)
*
  IF (ROOT) THEN
    WRITE (NOUT,*) 'Right singular vectors (full set)'
    WRITE (NOUT,*)
  END IF
  IFAIL = 0
*
  CALL X04BFFP(ICNTXT,NOUT,N,NX,A(M+1,1),LDA,FORMAT,TITOP,CNUMOP,
+             ICOFF,W,LDA,IFAIL)
*
*   Print right singular vectors (which define the null space)
*
  IF (ROOT) THEN
    WRITE (NOUT,*) 'Right singular vectors (null space only)'
    WRITE (NOUT,*)
    CNUMOP = 'L'
  END IF

```

```

*
*   (Integer function OFFSET gives the additive offset required to
*   transform the local column number to the global column number.)
*
      ICOFF = OFFSET(ICNTXT,N)
      IFAIL = 0
*
      CALL X04BFFP(ICNTXT,NOUT,N,NX-NR,A(M+1,NR+1),LDA,FORMAT,TITOP,
+             CNUMOP,ICOFF,W,LDA,IFAIL)
*
      IFAIL = 0
*
      Undefine the grid
*
      CALL Z01ABFP(ICNTXT,'N',IFAIL)
*
      STOP
      END
*
      SUBROUTINE GMATA(M,J1,J2,AL,LDAL)
*
*   GMATA generates the block A( 1: M, J1: J2 ) of the matrix A such
*   that
*
*       a(i,j) = 0.0 if i > j
*       a(i,j) = 1.0 else
*
*   in the array AL.
*
*   .. Scalar Arguments ..
      INTEGER          J1, J2, LDAL, M
*   .. Array Arguments ..
      DOUBLE PRECISION AL(LDAL,*)
*   .. Local Scalars ..
      INTEGER          I, J, L
*   .. Executable Statements ..
      L = 1
      DO 40 J = J1, J2
          DO 20 I = 1, M
              IF (J.GE.I) THEN
                  AL(I,L) = 1.0
              ELSE
                  AL(I,L) = 0.0
              END IF
          20 CONTINUE
          L = L + 1
      40 CONTINUE
*
*   End of GMATA.
*
      RETURN
      END
*
      INTEGER FUNCTION OFFSET(ICNTXT,N)
*
*   Integer function OFFSET gives the (additive) offset required to
*   transform the local column number to the global column number if
*   the distribution corresponds to the data structure

```

```

*   in F01ZRFP, F02FQFP or F02WQFP.
*
*   .. Scalar Arguments ..
INTEGER          ICNTXT, N
*   .. Local Scalars ..
INTEGER          ME, MP, MYCOL, MYROW, NB, NP, P
*   .. External Subroutines ..
EXTERNAL         Z01ZAFP
*   .. Intrinsic Functions ..
INTRINSIC        MOD
*   .. Executable Statements ..
CALL Z01ZAFP(ICNTXT,MP,NP,MYROW,MYCOL)
P = MP*NP
IF (P.GT.N) THEN
    NB = 1
ELSE
    NB = N/P
    IF (MOD(N,P).NE.0) THEN
        NB = NB + 1
    END IF
END IF
ME = MP*MYROW + MYCOL
OFFSET = ME*NB
*
*   End of OFFSET.
*
RETURN
END

```

8.2 Example Data

None.

8.3 Example Results

F02WQFP Example Program Results

Singular values

1	2
4.4485	1.1983
3	4
0.6953	0.5396

Left singular vectors

1	2
0.5653	-0.6145
0.5368	-0.1866
0.4811	0.3713
0.4011	0.6706
3	4
0.4832	0.2631
-0.5164	-0.6406
-0.4478	0.6559
0.5471	-0.3002

Right singular vectors (full set)

1	2
0.1271	-0.5128
0.2477	-0.6685
0.3559	-0.3587
0.4460	0.2009
0.4460	0.2009
0.4460	0.2009
0.4460	0.2009

3	4
0.6950	0.4877
-0.0477	-0.6996
-0.6917	0.5159
0.0951	-0.0405
0.0951	-0.0405
0.0951	-0.0405
0.0951	-0.0405

5	6
0.0000	0.0000
0.0000	0.0000
0.0000	0.0000
0.8573	-0.1227
-0.3436	-0.3632
-0.3436	-0.3632
-0.1700	0.8492

7
0.0000
0.0000
0.0000
0.0000
0.0000
-0.7071
0.7071
0.0000

Right singular vectors (null space only)

5	6
0.0000	0.0000
0.0000	0.0000
0.0000	0.0000
0.8573	-0.1227
-0.3436	-0.3632
-0.3436	-0.3632
-0.1700	0.8492

7
0.0000
0.0000
0.0000
0.0000
0.0000
-0.7071
0.7071
0.0000