X04YAFP

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

X04YAFP outputs to an external file a real dense vector x of length n, distributed conformally to a sparse matrix A (see Section 2.5 of the F11 Chapter Introduction).

One of the routines F11ZBFP or F11ZPFP must be called prior to X04YAFP to set up auxiliary information about the sparse matrix A in the array IAINFO.

2 Specification

```
SUBROUTINE X04YAFP(ICNTXT, NOUT, N, X, FORMAT, IAINFO, WORK, IFAIL)
DOUBLE PRECISION   X(*), WORK(*)
INTEGER            ICNTXT, NOUT, N, IAINFO(*), IFAIL
CHARACTER*(*)      FORMAT
```

3 Usage

3.1 Definitions

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

m_l –   the number of vector elements stored on the calling processor (m_l = IAINFO(3), see IAINFO).

$m_l^{\max}$ –   the maximum number of vector elements stored on any processor of the Library Grid ($m_l^{\max}$ = IAINFO(5), see IAINFO).

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: N, IFAIL

Global output arguments: IFAIL

Note: NOUT and FORMAT are only referenced on the root (or {0,0}) processor since it is only the root (or {0,0}) processor which performs output.

3.3 Distribution Strategy

The vector x must be distributed conformally to the sparse matrix A, i.e., x must be distributed across the Library Grid in the same way as each of the columns of the matrix A. This data distribution is described in more detail in Section 2.5 of the F11 Chapter Introduction.

3.4 Related Routines

Some Library routines can be used to generate or distribute real dense vectors conformally to a given real or complex sparse matrix.

Real vector generation: F01YEFP

Real vector scatter: F01XEFP

3.5 Requisites

The sparse matrix A must have been preprocessed to set up the auxiliary information vector IAINFO by F11ZBFP or F11ZPFP.

Cyclic row block distribution: F11ZBFP or F11ZPFP for real or complex A, respectively.

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:   NOUT — INTEGER   *Local Input*

On entry: the unit number to which the output will be directed.

Constraint: $1 \leq \text{NOUT} \leq 99$.

3:   N — INTEGER   *Global Input*

On entry: n, the order of the vector x. It must contain the same value as the parameter N used in a prior call to F11ZBFP or F11ZPFP in which the array IAINFO was initialised.

Constraint: $\text{N} \geq 1$.

4:   X($*$) — DOUBLE PRECISION array   *Local Input*

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

On entry: the local part of the vector x.

5:   FORMAT — CHARACTER*($*$)   *Local Input*

On entry: the format which will be used for output of the elements of x.

Constraint: any legal Fortran format for the output of floating-point numbers.

6:   IAINFO($*$) — INTEGER array   *Local Input*

Note: the dimension of the array IAINFO must be at least $\max(200, \text{IAINFO}(2))$.

On entry: the first IAINFO(2) elements of IAINFO contain information about the matrix A. The array IAINFO must have been initialised by a prior call to an appropriate Chapter F11 routine. The first IAINFO(2) elements of IAINFO **must not** be changed between successive calls to library routines involving the sparse matrix A.

Note: on exit from the Chapter F11 routine the element IAINFO(3) contains m_l, the number of rows of the matrix stored locally, and IAINFO(5) contains $m_l^{\max}$, the maximum number of vector elements stored on any processor of the Library Grid.

7:   WORK($*$) — DOUBLE PRECISION array   *Local Workspace*

Note: the dimension of the array WORK must be at least $\max(1, \text{IAINFO}(3) + \text{IAINFO}(5))$ on the root (or $\{0,0\}$) processor, and $\max(1, 2 \times \text{IAINFO}(3))$ on all other processors.

8:   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 ith argument was invalid. This error occured 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

>   IAINFO was not initialised by a prior call to F11ZBFP or F11ZPFP.

IFAIL = 2

>   An error has occurred in writing to unit NOUT. The file may already have been opened for reading only or the user may have run out of disk space.

5.3   Errors Detected in Reduced Error Checking Mode

The following errors conditions are guaranteed to be detected properly when the error checking mode was set to reduced error checking by a prior call to Z02EAFP with LEVEL = ± 1: IFAIL = 2.

6   Further Comments

The vector is first redistributed, using the workspace WORK, such that each logical processor holds a contiguous part of the vector. Output is then performed by the root (or {0,0}) processor if the root processor is not available. All other processors communicate their local portion of the redistributed vector to the root (or {0,0}) processor.

7   References

None.

8   Example

See Section 8 of the document for F11BAFP.