

NAG Toolbox

nag_sparse_real_symm_precon_ichol (f11ja)

1 Purpose

nag_sparse_real_symm_precon_ichol (f11ja) computes an incomplete Cholesky factorization of a real sparse symmetric matrix, represented in symmetric coordinate storage format. This factorization may be used as a preconditioner in combination with nag_sparse_real_symm_basic_solver (f11ge) or nag_sparse_real_symm_solve_ichol (f11jc).

2 Syntax

```
[a, irow, icol, ipiv, istr, nnzc, npivm, ifail] =
nag_sparse_real_symm_precon_ichol(nnz, a, irow, icol, lfill, dtol, mic, dscale,
ipiv, 'n', n, 'la', la, 'pstrat', pstrat)

[a, irow, icol, ipiv, istr, nnzc, npivm, ifail] = f11ja(nnz, a, irow, icol,
lfill, dtol, mic, dscale, ipiv, 'n', n, 'la', la, 'pstrat', pstrat)
```

3 Description

nag_sparse_real_symm_precon_ichol (f11ja) computes an incomplete Cholesky factorization (see Meijerink and Van der Vorst (1977)) of a real sparse symmetric n by n matrix A . It is designed specifically for positive definite matrices, but may also work for some mildly indefinite cases. The factorization is intended primarily for use as a preconditioner with one of the symmetric iterative solvers nag_sparse_real_symm_basic_solver (f11ge) or nag_sparse_real_symm_solve_ichol (f11jc).

The decomposition is written in the form

$$A = M + R$$

where

$$M = PLDL^T P^T$$

and P is a permutation matrix, L is lower triangular with unit diagonal elements, D is diagonal and R is a remainder matrix.

The amount of fill-in occurring in the factorization can vary from zero to complete fill, and can be controlled by specifying either the maximum level of fill **lfill**, or the drop tolerance **dtol**. The factorization may be modified in order to preserve row sums, and the diagonal elements may be perturbed to ensure that the preconditioner is positive definite. Diagonal pivoting may optionally be employed, either with a user-defined ordering, or using the Markowitz strategy (see Markowitz (1957)), which aims to minimize fill-in. For further details see Section 9.

The sparse matrix A is represented in symmetric coordinate storage (SCS) format (see Section 2.1.2 in the F11 Chapter Introduction). The array **a** stores all the nonzero elements of the lower triangular part of A , while arrays **irow** and **icol** store the corresponding row and column indices respectively. Multiple nonzero elements may not be specified for the same row and column index.

The preconditioning matrix M is returned in terms of the SCS representation of the lower triangular matrix

$$C = L + D^{-1} - I.$$

4 References

Chan T F (1991) Fourier analysis of relaxed incomplete factorization preconditioners *SIAM J. Sci. Statist. Comput.* **12**(2) 668–680

Markowitz H M (1957) The elimination form of the inverse and its application to linear programming *Management Sci.* **3** 255–269

Meijerink J and Van der Vorst H (1977) An iterative solution method for linear systems of which the coefficient matrix is a symmetric M-matrix *Math. Comput.* **31** 148–162

Salvini S A and Shaw G J (1995) An evaluation of new NAG Library solvers for large sparse symmetric linear systems *NAG Technical Report TRI/95*

Van der Vorst H A (1990) The convergence behaviour of preconditioned CG and CG-S in the presence of rounding errors *Lecture Notes in Mathematics* (eds O Axelsson and L Y Kolotilina) **1457** Springer–Verlag

5 Parameters

5.1 Compulsory Input Parameters

1: **nz** – INTEGER

The number of nonzero elements in the lower triangular part of the matrix A .

Constraint: $1 \leq \mathbf{nz} \leq \mathbf{n} \times (\mathbf{n} + 1)/2$.

2: **a(la)** – REAL (KIND=nag_wp) array

The nonzero elements in the lower triangular part of the matrix A , ordered by increasing row index, and by increasing column index within each row. Multiple entries for the same row and column indices are not permitted. The function `nag_sparse_real_symm_sort (f11zb)` may be used to order the elements in this way.

3: **irow(la)** – INTEGER array

4: **icol(la)** – INTEGER array

The row and column indices of the nonzero elements supplied in **a**.

Constraints:

irow and **icol** must satisfy these constraints (which may be imposed by a call to `nag_sparse_real_symm_sort (f11zb)`):

$$1 \leq \mathbf{irow}(i) \leq \mathbf{n} \text{ and } 1 \leq \mathbf{icol}(i) \leq \mathbf{irow}(i), \text{ for } i = 1, 2, \dots, \mathbf{nz};$$

$$\mathbf{irow}(i-1) < \mathbf{irow}(i) \text{ or } \mathbf{irow}(i-1) = \mathbf{irow}(i) \text{ and } \mathbf{icol}(i-1) < \mathbf{icol}(i), \text{ for } i = 2, 3, \dots, \mathbf{nz}.$$

5: **lfill** – INTEGER

If **lfill** ≥ 0 its value is the maximum level of fill allowed in the decomposition (see Section 9.2). A negative value of **lfill** indicates that **dtol** will be used to control the fill instead.

6: **dtol** – REAL (KIND=nag_wp)

If **lfill** < 0 , **dtol** is used as a drop tolerance to control the fill-in (see Section 9.2); otherwise **dtol** is not referenced.

Constraint: if **lfill** < 0 , **dtol** ≥ 0.0 .

7: **mic** – CHARACTER(1)

Indicates whether or not the factorization should be modified to preserve row sums (see Section 9.3).

mic = 'M'

The factorization is modified.

mic = 'N'

The factorization is not modified.

Constraint: **mic** = 'M' or 'N'.

8: **dscale** – REAL (KIND=nag_wp)

The diagonal scaling parameter. All diagonal elements are multiplied by the factor $(1 + \mathbf{dscale})$ at the start of the factorization. This can be used to ensure that the preconditioner is positive definite. See Section 9.3.

9: **ipiv**(**n**) – INTEGER array

If **pstrat** = 'U', then **ipiv**(*i*) must specify the row index of the diagonal element used as a pivot at elimination stage *i*. Otherwise **ipiv** need not be initialized.

Constraint: if **pstrat** = 'U', **ipiv** must contain a valid permutation of the integers on [1,**n**].

5.2 Optional Input Parameters1: **n** – INTEGER

Default: the dimension of the array **ipiv**.

n, the order of the matrix *A*.

Constraint: $\mathbf{n} \geq 1$.

2: **la** – INTEGER

Default: the dimension of the arrays **a**, **irow**, **icol**. (An error is raised if these dimensions are not equal.)

The dimension of the arrays **a**, **irow** and **icol**. these arrays must be of sufficient size to store both *A* (**nnz** elements) and *C* (**nnzc** elements).

Constraint: $\mathbf{la} \geq 2 \times \mathbf{nnz}$.

3: **pstrat** – CHARACTER(1)

Suggested value: **pstrat** = 'M'.

Default: 'M'

Specifies the pivoting strategy to be adopted.

pstrat = 'N'

No pivoting is carried out.

pstrat = 'M'

Diagonal pivoting aimed at minimizing fill-in is carried out, using the Markowitz strategy.

pstrat = 'U'

Diagonal pivoting is carried out according to the user-defined input value of **ipiv**.

Constraint: **pstrat** = 'N', 'M' or 'U'.

5.3 Output Parameters

1: **a(la)** – REAL (KIND=nag_wp) array

The first **nnz** elements of **a** contain the nonzero elements of A and the next **nnzc** elements contain the elements of the lower triangular matrix C . Matrix elements are ordered by increasing row index, and by increasing column index within each row.

2: **irow(la)** – INTEGER array

3: **icol(la)** – INTEGER array

The row and column indices of the nonzero elements returned in **a**.

4: **ipiv(n)** – INTEGER array

The pivot indices. If **ipiv**(i) = j then the diagonal element in row j was used as the pivot at elimination stage i .

5: **istr(n + 1)** – INTEGER array

istr(i), for $i = 1, 2, \dots, n$, is the starting address in the arrays **a**, **irow** and **icol** of row i of the matrix C . **istr**($n + 1$) is the address of the last nonzero element in C plus one.

6: **nnzc** – INTEGER

The number of nonzero elements in the lower triangular matrix C .

7: **npivm** – INTEGER

The number of pivots which were modified during the factorization to ensure that M was positive definite. The quality of the preconditioner will generally depend on the returned value of **npivm**. If **npivm** is large the preconditioner may not be satisfactory. In this case it may be advantageous to call `nag_sparse_real_symm_precon_ichol` (f11ja) again with an increased value of either **lfill** or **dscale**. See also Section 9.4.

8: **ifail** – INTEGER

ifail = 0 unless the function detects an error (see Section 5).

6 Error Indicators and Warnings

Errors or warnings detected by the function:

ifail = 1

On entry, **n** < 1,
 or **nnz** < 1,
 or **nnz** > $n \times (n + 1)/2$,
 or **la** < $2 \times \text{nnz}$,
 or **dtol** < 0.0,
 or **mic** ≠ 'M' or 'N',
 or **pstrat** ≠ 'N', 'M' or 'U',
 or $liwork < 2 \times la - 3 \times nnz + 7 \times n + 1$, and **lfill** ≥ 0,
 or $liwork < la - nnz + 7 \times n + 1$, and **lfill** < 0.

ifail = 2

On entry, the arrays **irow** and **icol** fail to satisfy the following constraints:

$1 \leq \text{irow}(i) \leq n$ and $1 \leq \text{icol}(i) \leq \text{irow}(i)$, for $i = 1, 2, \dots, \text{nnz}$;

$\text{irow}(i - 1) < \text{irow}(i)$, or $\text{irow}(i - 1) = \text{irow}(i)$ and $\text{icol}(i - 1) < \text{icol}(i)$, for $i = 2, 3, \dots, \text{nnz}$.

Therefore a nonzero element has been supplied which does not lie in the lower triangular part of A , is out of order, or has duplicate row and column indices. Call `nag_sparse_real_symm_sort` (f11zb) to reorder and sum or remove duplicates.

ifail = 3

On entry, **pstrat** = 'U', but **ipiv** does not represent a valid permutation of the integers in $[1, n]$. An input value of **ipiv** is either out of range or repeated.

ifail = 4

la is too small, resulting in insufficient storage space for fill-in elements. The decomposition has been terminated before completion. Either increase **la** or reduce the amount of fill by setting **pstrat** = 'M', reducing **lfill**, or increasing **dtol**.

ifail = 5 (`nag_sparse_real_symm_sort` (f11zb))

A serious error has occurred in an internal call to the specified function. Check all function calls and array sizes. Seek expert help.

ifail = -99

An unexpected error has been triggered by this routine. Please contact NAG.

ifail = -399

Your licence key may have expired or may not have been installed correctly.

ifail = -999

Dynamic memory allocation failed.

7 Accuracy

The accuracy of the factorization will be determined by the size of the elements that are dropped and the size of any modifications made to the diagonal elements. If these sizes are small then the computed factors will correspond to a matrix close to A . The factorization can generally be made more accurate by increasing **lfill**, or by reducing **dtol** with **lfill** < 0.

If `nag_sparse_real_symm_precon_ichol` (f11ja) is used in combination with `nag_sparse_real_symm_basic_solver` (f11ge) or `nag_sparse_real_symm_solve_ichol` (f11jc), the more accurate the factorization the fewer iterations will be required. However, the cost of the decomposition will also generally increase.

8 Further Comments

8.1 Timing

The time taken for a call to `nag_sparse_real_symm_precon_ichol` (f11ja) is roughly proportional to $(nnzc)^2/n$.

8.2 Control of Fill-in

If **lfill** ≥ 0 the amount of fill-in occurring in the incomplete factorization is controlled by limiting the maximum **level** of fill-in to **lfill**. The original nonzero elements of A are defined to be of level 0. The fill level of a new nonzero location occurring during the factorization is defined as

$$k = \max(k_e, k_c) + 1,$$

where k_e is the level of fill of the element being eliminated, and k_c is the level of fill of the element causing the fill-in.

If $\mathbf{lfill} < 0$ the fill-in is controlled by means of the **drop tolerance** \mathbf{dtol} . A potential fill-in element a_{ij} occurring in row i and column j will not be included if

$$|a_{ij}| < \mathbf{dtol} \times \sqrt{|a_{ii}a_{jj}|}.$$

For either method of control, any elements which are not included are discarded if $\mathbf{mic} = 'N'$, or subtracted from the diagonal element in the elimination row if $\mathbf{mic} = 'M'$.

8.3 Choice of Arguments

There is unfortunately no choice of the various algorithmic arguments which is optimal for all types of symmetric matrix, and some experimentation will generally be required for each new type of matrix encountered.

If the matrix A is not known to have any particular special properties the following strategy is recommended. Start with $\mathbf{lfill} = 0$, $\mathbf{mic} = 'N'$ and $\mathbf{dscale} = 0.0$. If the value returned for \mathbf{npivm} is significantly larger than zero, i.e., a large number of pivot modifications were required to ensure that M was positive definite, the preconditioner is not likely to be satisfactory. In this case increase either \mathbf{lfill} or \mathbf{dscale} until \mathbf{npivm} falls to a value close to zero. Once suitable values of \mathbf{lfill} and \mathbf{dscale} have been found try setting $\mathbf{mic} = 'M'$ to see if any improvement can be obtained by using **modified** incomplete Cholesky.

`nag_sparse_real_symm_precon_ichol` (f11ja) is primarily designed for positive definite matrices, but may work for some mildly indefinite problems. If \mathbf{npivm} cannot be satisfactorily reduced by increasing \mathbf{lfill} or \mathbf{dscale} then A is probably too indefinite for this function.

If A has non-positive off-diagonal elements, is nonsingular, and has only non-negative elements in its inverse, it is called an 'M-matrix'. It can be shown that no pivot modifications are required in the incomplete Cholesky factorization of an M-matrix (see Meijerink and Van der Vorst (1977)). In this case a good preconditioner can generally be expected by setting $\mathbf{lfill} = 0$, $\mathbf{mic} = 'M'$ and $\mathbf{dscale} = 0.0$.

For certain mesh-based problems involving M-matrices it can be shown in theory that setting $\mathbf{mic} = 'M'$, and choosing \mathbf{dscale} appropriately can reduce the order of magnitude of the condition number of the preconditioned matrix as a function of the mesh steplength (see Chan (1991)). In practise this property often holds even with $\mathbf{dscale} = 0.0$, although an improvement in condition can result from increasing \mathbf{dscale} slightly (see Van der Vorst (1990)).

Some illustrations of the application of `nag_sparse_real_symm_precon_ichol` (f11ja) to linear systems arising from the discretization of two-dimensional elliptic partial differential equations, and to random-valued randomly structured symmetric positive definite linear systems, can be found in Salvini and Shaw (1995).

8.4 Direct Solution of positive definite Systems

Although it is not their primary purpose, `nag_sparse_real_symm_precon_ichol` (f11ja) and `nag_sparse_real_symm_precon_ichol_solve` (f11jb) may be used together to obtain a **direct** solution to a symmetric positive definite linear system. To achieve this the call to `nag_sparse_real_symm_precon_ichol_solve` (f11jb) should be preceded by a **complete** Cholesky factorization

$$A = PLDL^T P^T = M.$$

A complete factorization is obtained from a call to `nag_sparse_real_symm_precon_ichol` (f11ja) with $\mathbf{lfill} < 0$ and $\mathbf{dtol} = 0.0$, provided $\mathbf{npivm} = 0$ on exit. A nonzero value of \mathbf{npivm} indicates that A is not positive definite, or is ill-conditioned. A factorization with nonzero \mathbf{npivm} may serve as a preconditioner, but will not result in a direct solution. It is therefore **essential** to check the output value of \mathbf{npivm} if a direct solution is required.

The use of `nag_sparse_real_symm_precon_ichol` (f11ja) and `nag_sparse_real_symm_precon_ichol_solve` (f11jb) as a direct method is illustrated in Section 10 in `nag_sparse_real_symm_precon_ichol_solve` (f11jb).

9 Example

This example reads in a symmetric sparse matrix A and calls `nag_sparse_real_symm_precon_ichol` (f11ja) to compute an incomplete Cholesky factorization. It then outputs the nonzero elements of both A and $C = L + D^{-1} - I$.

The call to `nag_sparse_real_symm_precon_ichol` (f11ja) has `lfill = 0`, `mic = 'N'`, `dscale = 0.0` and `pstrat = 'M'`, giving an unmodified zero-fill factorization of an unperturbed matrix, with Markowitz diagonal pivoting.

9.1 Program Text

```
function f11ja_example

fprintf('f11ja example results\n\n');

% Sparse matrix A
n    = nag_int(7);
nz   = nag_int(16);
a    = zeros(1000,1);
irow = zeros(1000,1,nag_int_name);
icol = irow;

a(1:16) = [4  1  5  2  2  3 -1  1  4  1 -2  3  2 -1 -2  5];
irow(1:16) = [1  2  2  3  4  4  5  5  5  6  6  6  7  7  7  7];
icol(1:16) = [1  1  2  3  2  4  1  4  5  2  5  6  1  2  3  7];

% Incomplete Cholesky factorization, zero fill
lfill = nag_int(0);
dtol  = 0;
mic   = 'N';
dscale = 0;
ipiv  = zeros(n, 1, nag_int_name);

[a, irow, icol, ipiv, istr, nnzc, npivm, ifail] = ...
f11ja( ...
    nz, a, irow, icol, lfill, dtol, mic, dscale, ipiv);

% Display details
fprintf(' Original Matrix\n');
inda = 1:nz;
amat = [inda' a(inda) irow(inda) icol(inda)];
fprintf('n      = %4d\n', n);
fprintf('nz     = %4d\n', nz);
fprintf('\n      a      irow      icol\n');
fprintf('%4d %11.4f%8d%8d\n',amat');

fprintf('\n Factorization\n');
inda = nz+1:nz+nnzc;
amat = [inda' a(inda) irow(inda) icol(inda)];
fprintf('n      = %4d\n', n);
fprintf('nz     = %4d\n', nnzc);
fprintf('npivm = %4d\n', npivm);
fprintf('\n      a      irow      icol\n');
fprintf('%4d %11.4f%8d%8d\n',amat');
fprintf('\n      i      ipiv(i)\n');
fprintf('%4d %8d\n',[[1:n]'; ipiv]);
```

9.2 Program Results

```
f11ja example results

Original Matrix
n      =      7
nz     =     16

      a      irow      icol
1      4.0000      1      1
2      1.0000      2      1
```

3	5.0000	2	2
4	2.0000	3	3
5	2.0000	4	2
6	3.0000	4	4
7	-1.0000	5	1
8	1.0000	5	4
9	4.0000	5	5
10	1.0000	6	2
11	-2.0000	6	5
12	3.0000	6	6
13	2.0000	7	1
14	-1.0000	7	2
15	-2.0000	7	3
16	5.0000	7	7

Factorization

n = 7
 nz = 16
 npivm = 0

	a	irow	icol
17	1.0000	1	1
18	0.0000	2	2
19	0.0000	3	2
20	0.0000	3	3
21	-1.0000	4	3
22	1.0000	4	4
23	0.0000	5	3
24	0.0000	5	5
25	1.0000	6	2
26	1.0000	6	4
27	0.0000	6	5
28	0.0000	6	6
29	-1.0000	7	1
30	1.0000	7	5
31	-1.0000	7	6
32	1.0000	7	7

i	ipiv(i)
1	2
3	4
5	6
7	3
4	5
6	1
2	7
