

nag_sparse_sym_chol_fac (f11jac)

1. Purpose

nag_sparse_sym_chol_fac (f11jac) 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_sym_chol_sol (f11jcc)**.

2. Specification

```
#include <nag.h>
#include <nagf11.h>

void nag_sparse_sym_chol_fac(Integer n, Integer nnz, double **a, Integer *la,
                             Integer **irow, Integer **icol, Integer lfill, double dtol,
                             Nag_SparseSym_Fact mic, double dscale, Nag_SparseSym_Piv pstrat,
                             Integer ipiv[], Integer istr[], Integer *nnzc, Integer *npivm,
                             Nag_Sparse_Comm *comm, NagError *fail);
```

3. Description

This routine 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 for the symmetric iterative solver **nag_sparse_sym_chol_sol (f11jcc)**.

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 6.

The sparse matrix A is represented in symmetric coordinate storage (SCS) format (see Section 2.1.2 of the Chapter Introduction). The array **a** stores all the non-zero elements of the lower triangular part of A , while arrays **irow** and **icol** store the corresponding row and column indices respectively. Multiple non-zero 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. Parameters

n

Input: the order of the matrix A .

Constraint: $n \geq 1$.

nnz

Input: the number of non-zero elements in the lower triangular part of the matrix A .

Constraint: $1 < \text{nnz} \leq n \times (n+1)/2$.

a[la]

Input: the non-zero 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 routine nag_sparse_sym_sort (f11zbc) may be used to order the elements in this way.

Output: the first **nnz** elements of **a** contain the non-zero 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.

la

Input: the dimension of the arrays **a**, **irow** and **icol** as declared in the calling program. These arrays must be of sufficient size to store both A (**nnz** elements) and C (**nnzc** elements); for this reason the length of the arrays may be changed internally by calls to **realloc**. It is therefore *imperative* that these arrays are *allocated* using *malloc* and **NOT** declared as automatic arrays.

Output: if internal allocation has taken place then **la** is set to **nnz** + **nnzc**, otherwise it remains unchanged.

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

irow[la]**icol[la]**

Input: the row and column indices of the non-zero elements supplied in **a**.

Constraint: **irow** and **icol** must satisfy the following constraints (which may be imposed by a call to nag_sparse_sym_sort (f11zbc)):

$$1 \leq \text{irow}[i] \leq \mathbf{n} \text{ and } 1 \leq \text{icol}[i] \leq \mathbf{n}, \text{ for } i = 0, 1, \dots, \text{nnz}-1.$$

$$\text{irow}[i-1] < \text{irow}[i], \text{ or}$$

$$\text{irow}[i-1] = \text{irow}[i] \text{ and } \text{icol}[i-1] < \text{icol}[i], \text{ for } i = 1, 2, \dots, \text{nnz}-1.$$

Output: the row and column indices of the non-zero elements returned in **a**.

lfill

Input: if **lfill** ≥ 0 its value is the maximum level of fill allowed in the decomposition (see Section 6.2). A negative value of **lfill** indicates that **dtol** will be used to control the fill instead.

dtol

Input: if **lfill** < 0 then **dtol** is used as a drop tolerance to control the fill-in (see Section 6.2). Otherwise **dtol** is not referenced.

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

mic

Input: indicates whether or not the factorization should be modified to preserve row sums (see Section 6.3).

If **mic** = **Nag_SparseSym_ModFact**, the factorization is modified (MIC).

If **mic** = **Nag_SparseSym_UnModFact**, the factorization is not modified.

Constraint: **mic** = **Nag_SparseSym_ModFact** or **Nag_SparseSym_UnModFact**.

dscale

Input: the diagonal scaling parameter. All diagonal elements are multiplied by the factor $(1 + \text{dscale})$ at the start of the factorization. This can be used to ensure that the preconditioner is positive-definite. See Section 6.3.

pstrat

Input: specifies the pivoting strategy to be adopted as follows:

if **pstrat** = **Nag_SparseSym_NoPiv** then no pivoting is carried out;

if **pstrat** = **Nag_SparseSym_MarkPiv** then diagonal pivoting aimed at minimizing fill-in is carried out, using the Markowitz strategy;

if **pstrat** = **Nag_SparseSym_UserPiv** then diagonal pivoting is carried out according to the user-defined input value of **ipiv**.

Suggested value: **pstrat** = **Nag_SparseSym_MarkPiv**.

Constraint: **pstrat** = **Nag_SparseSym_NoPiv**, **Nag_SparseSym_MarkPiv** or **Nag_SparseSym_UserPiv**.

ipiv[n]

Input: if **pstrat** = **Nag_SparseSym_UserPiv**, then **ipiv**[$i-1$] 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** = **Nag_SparseSym_UserPiv**, then **ipiv** must contain a valid permutation of the integers on $[1, n]$.

Output: the pivot indices. If **ipiv**[$i-1$] = j then the diagonal element in row j was used as the pivot at elimination stage i .

istr[n+1]

Output: **istr**[i], for $i = 0, 1, \dots, n-1$ holds the starting address in the arrays **a**, **irow** and **icol** of row i of the matrix C . **istr**[n] holds the address of the last non-zero element in C plus one.

nnzc

Output: the number of non-zero elements in the lower triangular matrix C .

npivm

Output: 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_sym_chol_fac` again with an increased value of either **lfill** or **dscale**.

comm

Input/Output: a pointer to a structure of type **Nag_Sparse_Comm** whose members are used by the iterative solver.

fail

The NAG error parameter, see the Essential Introduction to the NAG C Library.

5. Error Indications and Warnings

NE_BAD_PARAM

On entry, parameter **mic** had an illegal value.

On entry, parameter **pstrat** had an illegal value.

NE_INT_ARG_LT

On entry, **n** must not be less than 1: **n** = $\langle value \rangle$.

NE_INT_2

On entry, **nnz** = $\langle value \rangle$, **n** = $\langle value \rangle$.

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

NE_2_INT_ARG_LT

On entry, **la** = $\langle value \rangle$ while **nnz** = $\langle value \rangle$.

These parameters must satisfy $\mathbf{la} \geq 2 \times \mathbf{nnz}$.

NE_REALINT_ARG_CONS

On entry, **dtol** = $\langle value \rangle$ and **lfill** = $\langle value \rangle$.

These parameters must satisfy $\mathbf{dtol} \geq 0.0$ if **lfill** < 0.,

NE_SYMM_MATRIX_DUP

A non-zero element has been supplied which does not lie in the lower triangular part of the matrix A , is out of order, or has duplicate row and column indices, i.e., one or more of the following constraints has been violated:

$1 \leq \mathbf{irow}[i] \leq \mathbf{n}$ and $1 \leq \mathbf{icol}[i] \leq \mathbf{n}$, for $i = 0, 1, \dots, \mathbf{nnz}-1$.

$\mathbf{irow}[i-1] < \mathbf{irow}[i]$, or

$\mathbf{irow}[i-1] = \mathbf{irow}[i]$ and $\mathbf{icol}[i-1] < \mathbf{icol}[i]$, for $i = 1, 2, \dots, \mathbf{nnz}-1$.

Call `nag_sparse_sym_sort` (f11zbc) to reorder and sum or remove duplicates.

NE_INVALID_ROW_PIVOT

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

NE_INTERNAL_ERROR

An internal error has occurred in this function. Check the function call and any array sizes. If the call is correct then please consult NAG for assistance.

NE_ALLOC_FAIL

Memory allocation failed.

6. Further Comments

The time taken for a call to `nag_sparse_sym_chol_fac` is roughly proportional to $\mathbf{nnzc}^2/\mathbf{n}$.

6.1. 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_sym_chol_fac` is used in combination with `nag_sparse_sym_chol_sol` (`f11jcc`), the more accurate the factorization the fewer iterations will be required. However, the cost of the decomposition will also generally increase.

6.2 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 non-zero elements of A are defined to be of level 0. The fill level of a new non-zero 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 **lfill** < 0 the fill-in is controlled by means of the **drop tolerance** **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 **mic** = **Nag_SparseSym_UnModFact**, or subtracted from the diagonal element in the elimination row if **mic** = **Nag_SparseSym_ModFact**.

6.3 Choice of Parameters

There is unfortunately no choice of the various algorithmic parameters 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 **lfill** = 0, **mic** = **Nag_SparseSym_UnModFact** and **dscale** = 0.0. If the value returned for **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 **lfill** or **dscale** until **npivm** falls to a value close to zero. Once suitable values of **lfill** and **dscale** have been found try setting **mic** = **Nag_SparseSym_ModFact** to see if any improvement can be obtained by using **modified** incomplete Cholesky.

`nag_sparse_sym_chol_fac` is primarily designed for positive-definite matrices, but may work for some mildly indefinite problems. If **npivm** cannot be satisfactorily reduced by increasing **lfill** or **dscale** then A is probably too indefinite for this routine.

If A has non-positive off-diagonal elements, is non-singular, 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 (Meijerink and van der Vorst (1977)). In this case a good preconditioner can generally be expected by setting **lfill** = 0, **mic** = **Nag_SparseSym_ModFact** and **dscale** = 0.0.

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

Some illustrations of the application of `nag_sparse_sym_chol_fac` 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).

6.4 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 TR1/95*, NAG Ltd, Oxford.
- 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* (ed O Axelsson and L Y Kolotilina) **1457** Springer-Verlag.

7. See Also

`nag_sparse_sym_chol_sol` (f11jcc)
`nag_sparse_sym_sort` (f11zbc)

8. Example

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

The call to `nag_sparse_sym_chol_fac` has **lfill** = 0, **mic** = **Nag_SparseSymm_UnmodifiedFactor**, **dscale** = 0.0 and **pstrat** = **Nag_SparseSym_MarkPiv**, giving an unmodified zero-fill factorization of an unperturbed matrix, with Markowitz diagonal pivoting.

8.1. Program Text

```
/* nag_sparse_sym_chol_fac (f11jac) Example Program.
 *
 * Copyright 1998 Numerical Algorithms Group.
 *
 * Mark 5, 1998.
 *
 */

#include <nag.h>
#include <stdio.h>
#include <nag_stdlib.h>
#include <nag_string.h>
#include <nagf11.h>

main()
{
    double dtol;
    double *a;
```

```

double dscale;

Integer *irow, *icol;
Integer *ipiv, nnzc, *istr;
Integer i, n, lfill, npivm;
Integer nnz;
Integer num;

Nag_SparseSym_Piv pstrat;
Nag_SparseSym_Fact mic;
Nag_Sparse_Comm comm;

char char_enum[20];

Vprintf("f11jac Example Program Results\n\n");

/* Skip heading in data file */
Vscanf("%*[^\\n]");

/* Read algorithmic parameters */
Vscanf("%ld",&n);
Vscanf("%*[^\\n]");
Vscanf("%ld%*[^\\n",&nnz);
Vscanf("%ld%lf%*[^\\n", &lfill, &dtol);
Vscanf("%s%lf%*[^\\n",char_enum, &dscale);

if (!strcmp(char_enum, "ModFact"))
    mic = Nag_SparseSym_ModFact;
else if (!strcmp(char_enum, "UnModFact"))
    mic = Nag_SparseSym_UnModFact;
else
{
    Vprintf("Unrecognised string for mic enum representation.\\n");
    exit (EXIT_FAILURE);
}

Vscanf("%s%*[^\\n",char_enum);
if (!strcmp(char_enum, "NoPiv"))
    pstrat = Nag_SparseSym_NoPiv;
else if (!strcmp(char_enum, "MarkPiv"))
    pstrat = Nag_SparseSym_MarkPiv;
else if (!strcmp(char_enum, "UserPiv"))
    pstrat = Nag_SparseSym_UserPiv;
else
{
    Vprintf("Unrecognised string for pstrat enum representation.\\n");
    exit (EXIT_FAILURE);
}

num = 2 * nnz;
ipiv = NAG_ALLOC(n,Integer);
istr = NAG_ALLOC(n+1,Integer);
irow = NAG_ALLOC(num,Integer);
icol = NAG_ALLOC(num,Integer);
a = NAG_ALLOC(num,double);

if (!ipiv || !istr || !irow || !icol || !a)
{
    Vprintf("Allocation failure\\n");
    exit (EXIT_FAILURE);
}

/* Read the matrix a */

for (i = 1; i <= nnz; ++i)
    Vscanf("%lf%ld%ld%*[^\\n",&a[i-1], &irow[i-1], &icol[i-1] );

/* Calculate incomplete Cholesky factorization */

f11jac(n, nnz, &a, &num, &irow, &icol, lfill, dtol, mic,

```

```

        dscale, pstrat, ipiv, istr, &nnzc, &npivm, &comm, NAGERR_DEFAULT);

/* Output original matrix */

Vprintf(" Original Matrix \n");
Vprintf(" n      = %6ld\n",n);
Vprintf(" nnz    = %6ld\n\n",nnz);
for (i = 1; i <= nnz; ++i)
    Vprintf(" %8ld%16.4e%8ld%8ld\n",i,a[i-1],irow[i-1],icol[i-1]);
Vprintf("\n");

/* Output details of the factorization */
Vprintf(" Factorization\n n = %6ld \n nnz = %6ld\n",n,nnzc);
Vprintf(" npivm = %6ld\n\n",npivm);
for (i = nnz + 1; i <= nnz + nnzc; ++i)
    Vprintf(" %8ld%16.4e%8ld%8ld\n",i,a[i-1],irow[i-1],icol[i-1]);

Vprintf("\n      i      ipiv(i) \n");
for (i = 1; i <= n; ++i)
    Vprintf(" %8ld%8ld\n",i,ipiv[i-1]);

NAG_FREE(irow);
NAG_FREE(icol);
NAG_FREE(a);
NAG_FREE(istr);
NAG_FREE(ipiv);
exit (EXIT_SUCCESS);
}

```

8.2. Program Data

f11jac Example Program Data

7	n
16	nnz
0 0.0	lfill, dtol
UnModFact 0.0	mic, dscale
MarkPiv	pstrat
4. 1 1	
1. 2 1	
5. 2 2	
2. 3 3	
2. 4 2	
3. 4 4	
-1. 5 1	
1. 5 4	
4. 5 5	
1. 6 2	
-2. 6 5	
3. 6 6	
2. 7 1	
-1. 7 2	
-2. 7 3	
5. 7 7	a[i-1], irow[i-1], icol[i-1], i=1,...,nnz

8.3. Program Results

f11jac Example Program Results

Original Matrix

n = 7
nnz = 16

1	4.0000e+00	1	1
2	1.0000e+00	2	1
3	5.0000e+00	2	2
4	2.0000e+00	3	3
5	2.0000e+00	4	2
6	3.0000e+00	4	4
7	-1.0000e+00	5	1
8	1.0000e+00	5	4
9	4.0000e+00	5	5

10	1.0000e+00	6	2
11	-2.0000e+00	6	5
12	3.0000e+00	6	6
13	2.0000e+00	7	1
14	-1.0000e+00	7	2
15	-2.0000e+00	7	3
16	5.0000e+00	7	7

Factorization
n = 7
nnz = 16
npivm = 0

17	5.0000e-01	1	1
18	3.3333e-01	2	2
19	3.3333e-01	3	2
20	2.7273e-01	3	3
21	-5.4545e-01	4	3
22	5.2381e-01	4	4
23	-2.7273e-01	5	3
24	2.6829e-01	5	5
25	6.6667e-01	6	2
26	5.2381e-01	6	4
27	2.6829e-01	6	5
28	3.4788e-01	6	6
29	-1.0000e+00	7	1
30	5.3659e-01	7	5
31	-5.3455e-01	7	6
32	9.0461e-01	7	7

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