

NAG Library Function Document

nag_dsgesv (f07acc)

1 Purpose

nag_dsgesv (f07acc) computes the solution to a real system of linear equations

$$AX = B,$$

where A is an n by n matrix and X and B are n by r matrices.

2 Specification

```
#include <nag.h>
#include <nagf07.h>

void nag_dsgesv (Nag_OrderType order, Integer n, Integer nrhs, double a[],
                Integer pda, Integer ipiv[], const double b[], Integer pdb, double x[],
                Integer pdx, Integer *iter, NagError *fail)
```

3 Description

nag_dsgesv (f07acc) first attempts to factorize the matrix in single precision and use this factorization within an iterative refinement procedure to produce a solution with full double precision accuracy. If the approach fails the method switches to a double precision factorization and solve.

The iterative refinement process is stopped if

$$\mathbf{iter} > \mathit{itermax},$$

where \mathbf{iter} is the number of iterations carried out thus far and $\mathit{itermax}$ is the maximum number of iterations allowed, which is fixed at 30 iterations. The process is also stopped if for all right-hand sides we have

$$\|\mathit{resid}\| < \sqrt{\mathbf{n}}\|x\|\|A\|\epsilon,$$

where $\|\mathit{resid}\|$ is the ∞ -norm of the residual, $\|x\|$ is the ∞ -norm of the solution, $\|A\|$ is the ∞ -operator-norm of the matrix A and ϵ is the *machine precision* returned by nag_machine_precision (X02AJC).

The iterative refinement strategy used by nag_dsgesv (f07acc) can be more efficient than the corresponding direct full precision algorithm. Since this strategy must perform iterative refinement on each right-hand side, any efficiency gains will reduce as the number of right-hand sides increases. Conversely, as the matrix size increases the cost of these iterative refinements become less significant relative to the cost of factorization. Thus, any efficiency gains will be greatest for a very small number of right-hand sides and for large matrix sizes. The cut-off values for the number of right-hand sides and matrix size, for which the iterative refinement strategy performs better, depends on the relative performance of the reduced and full precision factorization and back-substitution. For now, nag_dsgesv (f07acc) always attempts the iterative refinement strategy first; you are advised to compare the performance of nag_dsgesv (f07acc) with that of its full precision counterpart nag_dgesv (f07aac) to determine whether this strategy is worthwhile for your particular problem dimensions.

4 References

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 <http://www.netlib.org/lapack/lug>

Buttari A, Dongarra J, Langou J, Langou J, Luszczek P and Kurzak J (2007) Mixed precision iterative refinement techniques for the solution of dense linear systems *International Journal of High Performance Computing Applications*

Golub G H and Van Loan C F (1996) *Matrix Computations* (3rd Edition) Johns Hopkins University Press, Baltimore

5 Arguments

1: **order** – Nag_OrderType *Input*

On entry: the **order** argument specifies the two-dimensional storage scheme being used, i.e., row-major ordering or column-major ordering. C language defined storage is specified by **order** = Nag_RowMajor. See Section 3.2.1.3 in the Essential Introduction for a more detailed explanation of the use of this argument.

Constraint: **order** = Nag_RowMajor or Nag_ColMajor.

2: **n** – Integer *Input*

On entry: *n*, the number of linear equations, i.e., the order of the matrix *A*.

Constraint: **n** ≥ 0.

3: **nrhs** – Integer *Input*

On entry: *r*, the number of right-hand sides, i.e., the number of columns of the matrix *B*.

Constraint: **nrhs** ≥ 0.

4: **a**[*dim*] – double *Input/Output*

Note: the dimension, *dim*, of the array **a** must be at least max(1, **pda** × **n**).

The (*i*, *j*)th element of the matrix *A* is stored in

$$\begin{aligned} &\mathbf{a}[(j-1) \times \mathbf{pda} + i - 1] \text{ when } \mathbf{order} = \text{Nag_ColMajor}; \\ &\mathbf{a}[(i-1) \times \mathbf{pda} + j - 1] \text{ when } \mathbf{order} = \text{Nag_RowMajor}. \end{aligned}$$

On entry: the *n* by *n* coefficient matrix *A*.

On exit: if iterative refinement has been successfully used (i.e., if **fail.code** = NE_NOERROR and **iter** ≥ 0), then *A* is unchanged. If double precision factorization has been used (when **fail.code** = NE_NOERROR and **iter** < 0), *A* contains the factors *L* and *U* from the factorization $A = PLU$; the unit diagonal elements of *L* are not stored.

5: **pda** – Integer *Input*

On entry: the stride separating row or column elements (depending on the value of **order**) in the array **a**.

Constraint: **pda** ≥ max(1, **n**).

6: **ipiv**[**n**] – Integer *Output*

On exit: if no constraints are violated, the pivot indices that define the permutation matrix *P*; at the *i*th step row *i* of the matrix was interchanged with row **ipiv**[*i* - 1]. **ipiv**[*i* - 1] = *i* indicates a row interchange was not required. **ipiv** corresponds either to the single precision factorization (if **fail.code** = NE_NOERROR and **iter** ≥ 0) or to the double precision factorization (if **fail.code** = NE_NOERROR and **iter** < 0).

- 7: **b**[*dim*] – const double *Input*
- Note:** the dimension, *dim*, of the array **b** must be at least
 $\max(1, \mathbf{pdb} \times \mathbf{nrhs})$ when **order** = Nag_ColMajor;
 $\max(1, \mathbf{n} \times \mathbf{pdb})$ when **order** = Nag_RowMajor.
- The (*i*, *j*)th element of the matrix *B* is stored in
 $\mathbf{b}[(j - 1) \times \mathbf{pdb} + i - 1]$ when **order** = Nag_ColMajor;
 $\mathbf{b}[(i - 1) \times \mathbf{pdb} + j - 1]$ when **order** = Nag_RowMajor.
- On entry:* the *n* by *r* right-hand side matrix *B*.
- 8: **pdb** – Integer *Input*
- On entry:* the stride separating row or column elements (depending on the value of **order**) in the array **b**.
- Constraints:*
 if **order** = Nag_ColMajor, **pdb** $\geq \max(1, \mathbf{n})$;
 if **order** = Nag_RowMajor, **pdb** $\geq \max(1, \mathbf{nrhs})$.
- 9: **x**[*dim*] – double *Output*
- Note:** the dimension, *dim*, of the array **x** must be at least
 $\max(1, \mathbf{pdx} \times \mathbf{nrhs})$ when **order** = Nag_ColMajor;
 $\max(1, \mathbf{n} \times \mathbf{pdx})$ when **order** = Nag_RowMajor.
- The (*i*, *j*)th element of the matrix *X* is stored in
 $\mathbf{x}[(j - 1) \times \mathbf{pdx} + i - 1]$ when **order** = Nag_ColMajor;
 $\mathbf{x}[(i - 1) \times \mathbf{pdx} + j - 1]$ when **order** = Nag_RowMajor.
- On exit:* if **fail.code** = NE_NOERROR, the *n* by *r* solution matrix *X*.
- 10: **pdx** – Integer *Input*
- On entry:* the stride separating row or column elements (depending on the value of **order**) in the array **x**.
- Constraints:*
 if **order** = Nag_ColMajor, **pdx** $\geq \max(1, \mathbf{n})$;
 if **order** = Nag_RowMajor, **pdx** $\geq \max(1, \mathbf{nrhs})$.
- 11: **iter** – Integer * *Output*
- On exit:* if **iter** > 0, iterative refinement has been successfully used and **iter** is the number of iterations carried out.
- If **iter** < 0, iterative refinement has failed for one of the reasons given below and double precision factorization has been carried out instead.
- iter** = -1
 Taking into account machine parameters, and the values of **n** and **nrhs**, it is not worth working in single precision.
- iter** = -2
 Overflow of an entry occurred when moving from double to single precision.
- iter** = -3
 An intermediate single precision factorization failed.
- iter** = -31
 The maximum permitted number of iterations was exceeded.

12: **fail** – NagError *

Input/Output

The NAG error argument (see Section 3.6 in the Essential Introduction).

6 Error Indicators and Warnings

NE_ALLOC_FAIL

Dynamic memory allocation failed.

NE_BAD_PARAM

On entry, argument $\langle value \rangle$ had an illegal value.

NE_INT

On entry, $\mathbf{n} = \langle value \rangle$.

Constraint: $\mathbf{n} \geq 0$.

On entry, $\mathbf{nrhs} = \langle value \rangle$.

Constraint: $\mathbf{nrhs} \geq 0$.

On entry, $\mathbf{pda} = \langle value \rangle$.

Constraint: $\mathbf{pda} > 0$.

On entry, $\mathbf{pdb} = \langle value \rangle$.

Constraint: $\mathbf{pdb} > 0$.

On entry, $\mathbf{pdx} = \langle value \rangle$.

Constraint: $\mathbf{pdx} > 0$.

NE_INT_2

On entry, $\mathbf{pda} = \langle value \rangle$ and $\mathbf{n} = \langle value \rangle$.

Constraint: $\mathbf{pda} \geq \max(1, \mathbf{n})$.

On entry, $\mathbf{pdb} = \langle value \rangle$ and $\mathbf{n} = \langle value \rangle$.

Constraint: $\mathbf{pdb} \geq \max(1, \mathbf{n})$.

On entry, $\mathbf{pdb} = \langle value \rangle$ and $\mathbf{nrhs} = \langle value \rangle$.

Constraint: $\mathbf{pdb} \geq \max(1, \mathbf{nrhs})$.

On entry, $\mathbf{pdx} = \langle value \rangle$ and $\mathbf{n} = \langle value \rangle$.

Constraint: $\mathbf{pdx} \geq \max(1, \mathbf{n})$.

On entry, $\mathbf{pdx} = \langle value \rangle$ and $\mathbf{nrhs} = \langle value \rangle$.

Constraint: $\mathbf{pdx} \geq \max(1, \mathbf{nrhs})$.

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 contact NAG for assistance.

NE_SINGULAR

$U(\langle value \rangle, \langle value \rangle)$ is exactly zero. The factorization has been completed, but the factor U is exactly singular, so the solution could not be computed.

7 Accuracy

The computed solution for a single right-hand side, \hat{x} , satisfies the equation of the form

$$(A + E)\hat{x} = b,$$

where

$$\|E\|_1 = O(\epsilon)\|A\|_1$$

and ϵ is the *machine precision*. An approximate error bound for the computed solution is given by

$$\frac{\|\hat{x} - x\|_1}{\|x\|_1} \leq \kappa(A) \frac{\|E\|_1}{\|A\|_1}$$

where $\kappa(A) = \|A^{-1}\|_1 \|A\|_1$, the condition number of A with respect to the solution of the linear equations. See Section 4.4 of Anderson *et al.* (1999) for further details.

8 Parallelism and Performance

nag_dsgesv (f07acc) is threaded by NAG for parallel execution in multithreaded implementations of the NAG Library.

nag_dsgesv (f07acc) makes calls to BLAS and/or LAPACK routines, which may be threaded within the vendor library used by this implementation. Consult the documentation for the vendor library for further information.

Please consult the Users' Note for your implementation for any additional implementation-specific information.

9 Further Comments

The complex analogue of this function is nag_zcgesv (f07aqc).

10 Example

This example solves the equations

$$Ax = b,$$

where A is the general matrix

$$A = \begin{pmatrix} 1.80 & 2.88 & 2.05 & -0.89 \\ 5.25 & -2.95 & -0.95 & -3.80 \\ 1.58 & -2.69 & -2.90 & -1.04 \\ -1.11 & -0.66 & -0.59 & 0.80 \end{pmatrix} \quad \text{and} \quad b = \begin{pmatrix} 9.52 \\ 24.35 \\ 0.77 \\ -6.22 \end{pmatrix}.$$

10.1 Program Text

```

/* nag_dsgesv (f07acc) Example Program.
 *
 * Copyright 2009, Numerical Algorithms Group.
 *
 * Mark 23, 2011.
 */

#include <stdio.h>
#include <math.h>
#include <nag.h>
#include <nag_stdlib.h>
#include <nagf07.h>

int main(void)
{
    /* Scalars */
    Integer    exit_status = 0;
    Integer    i, iter, j, n, nrhs, pda, pdb, pdx;

    /* Arrays */
    Integer    *ipiv = 0;
    double     *a = 0, *b = 0, *x = 0;

```

```

/* Nag Types */
NagError      fail;
Nag_OrderType order;

INIT_FAIL(fail);

printf("nag_dsgesv (f07acc) Example Program Results\n\n");
/* Skip heading in data file*/
scanf("%*[\n]");
scanf("%ld%ld%*[\n]", &n, &nrhs);
if (n < 0 || nrhs < 0)
{
    printf("Invalid n or nrhs\n");
    exit_status = 1;
    return exit_status;
}

pda = n;
#ifdef NAG_COLUMN_MAJOR
    pdb = n;
    pdx = n;
#define A(I, J) a[(J-1)*pda + I-1]
#define B(I, J) b[(J-1)*pdb + I-1]
    order = Nag_ColMajor;
#else
    pdb = nrhs;
    pdx = nrhs;
#define A(I, J) a[(I-1)*pda + J-1]
#define B(I, J) b[(I-1)*pdb + J-1]
    order = Nag_RowMajor;
#endif

/* Allocate memory */
if (!(a = NAG_ALLOC(n*n, double)) ||
    !(b = NAG_ALLOC(n*nrhs, double)) ||
    !(x = NAG_ALLOC(n*nrhs, double)) ||
    !(ipiv = NAG_ALLOC(n, Integer)))
{
    printf("Allocation failure\n");
    exit_status = -1;
    goto END;
}

/* Read A and B from data file*/
for (i = 1; i <= n; i++)
    for (j = 1; j <= n; j++) scanf("%lf", &A(i, j));
scanf("%*[\n] ");
for (i = 1; i <= n; i++)
    for (j = 1; j <= nrhs; j++) scanf("%lf", &B(i, j));
scanf("%*[\n] ");

/* Solve the equations Ax = b for x using nag_dsgesv (f07acc)
 * Mixed precision real system solver.
 */
nag_dsgesv(order, n, nrhs, a, pda, ipiv, b, pdb, x, pdx, &iter, &fail);
if (fail.code != NE_NOERROR)
{
    printf("Error from nag_dsgesv (f07acc).\n%s\n", fail.message);
    exit_status = 1;
    goto END;
}

/* Print solution using
 * nag_gen_real_mat_print (x04cac)
 */
fflush(stdout);
nag_gen_real_mat_print(order, Nag_GeneralMatrix, Nag_NonUnitDiag, n, nrhs, x,
                      pdx, "Solution(s)", 0, &fail);
if (fail.code != NE_NOERROR)
{
    printf("Error from nag_gen_real_mat_print (x04cac).\n%s\n", fail.message);
}

```

```

        exit_status = 1;
        goto END;
    }

    /* Print pivot indices*/
    printf("\nPivot indices\n");
    for (i = 0; i < n; i++)
        printf("%11ld%s", ipiv[i], (i+1)%7?" ":"\n");
    printf("\n");

END:
    NAG_FREE(a);
    NAG_FREE(b);
    NAG_FREE(x);
    NAG_FREE(ipiv);

    return exit_status;
}
#undef A
#undef B

```

10.2 Program Data

nag_dsgesv (f07acc) Example Program Data

```

4 1                                     :Value of n and nrhs

1.80  2.88  2.05 -0.89
5.25 -2.95 -0.95 -3.80
1.58 -2.69 -2.90 -1.04
-1.11 -0.66 -0.59  0.80 :End of matrix A

9.52  24.35  0.77 -6.22 :End of vector b

```

10.3 Program Results

nag_dsgesv (f07acc) Example Program Results

```

Solution(s)
1
1      1.0000
2     -1.0000
3      3.0000
4     -5.0000

Pivot indices
      2          3          4

```
