

Module 5.1: nag_gen_lin_sys

General Systems of Linear Equations

`nag_gen_lin_sys` provides a procedure for solving general real or complex systems of linear equations with one or many right-hand sides:

$$Ax = b \text{ or } AX = B.$$

It also provides procedures for factorizing A and solving a system of equations when the matrix A has already been factorized.

Contents

Introduction	5.1.3
Procedures	
<code>nag_gen_lin_sol</code>	5.1.5
Solves a general real or complex system of linear equations with one or many right-hand sides	
<code>nag_gen_lin_fac</code>	5.1.9
Performs an LU factorization of a general real or complex matrix	
<code>nag_gen_lin_sol_fac</code>	5.1.13
Solves a general real or complex system of linear equations, with coefficient matrix previously factorized by <code>nag_gen_lin_fac</code>	
Examples	
Example 1: Solution of a General Real System of Linear Equations	5.1.17
Example 2: Factorization of a General Real Matrix and Use of the Factorization to Solve a System of Linear Equations	5.1.19
Additional Examples	5.1.23
References	5.1.24

Introduction

1 Notation and Background

We use the following notation for a system of linear equations:

$Ax = b$, if there is one right-hand side b ;

$AX = B$, if there are many right-hand sides (the columns of the matrix B).

There are options to solve alternative forms of the equations:

$A^T x = b$, $A^T X = B$, $A^H x = b$ or $A^H X = B$.

(If A is real, then $A^H = A^T$.)

In this module, the square matrix A (the *coefficient matrix*) is assumed to be a *general* real or complex matrix, that is, a matrix without any known special properties such as symmetry. If A is symmetric or Hermitian, see the module `nag_sym_lin_sys` (5.2); if A is triangular, see the module `nag_tri_lin_sys` (5.3).

The module provides options to return *forward* or *backward error bounds* on the computed solution. It also provides options to evaluate the *determinant* of A and to estimate the *condition number* of A , which is a measure of the sensitivity of the computed solution to perturbations of the original data or to rounding errors in the computation. For more details on error analysis, see the Chapter Introduction.

To solve the system of equations, the first step is to compute an *LU factorization* of A (with row interchanges to ensure numerical stability). The system of equations can then be solved by forward and backward substitution.

2 Choice of Procedures

The procedure `nag_gen_lin_sol` should be suitable for most purposes; it performs the factorization of A and solves the system of equations in a single call. It also has options to estimate the condition number of A , and to return forward and backward error bounds on the computed solution.

The module also provides lower-level procedures which perform the two computational steps in the solution process:

`nag_gen_lin_fac` computes an *LU* factorization of A , with options to evaluate the determinant and to estimate the condition number;

`nag_gen_lin_sol_fac` solves the system of equations, assuming that A has already been factorized by a call to `nag_gen_lin_fac`. It has options to return forward and backward error bounds on the solution.

These lower-level procedures are intended for more experienced users. For example, they enable a factorization computed by `nag_gen_lin_fac` to be reused several times in repeated calls to `nag_gen_lin_sol_fac`.

Procedure: nag_gen_lin_sol

1 Description

`nag_gen_lin_sol` is a generic procedure which computes the solution of a general real or complex system of linear equations with one or many right-hand sides.

We write:

$Ax = b$, if there is one right-hand side b ;

$AX = B$, if there are many right-hand sides (the columns of the matrix B).

The matrix A (the *coefficient matrix*) is assumed to be a general matrix, without any known special properties such as symmetry.

Optionally, the procedure can solve alternative forms of the system of equations:

$$A^T x = b, \quad A^T X = B, \quad A^H x = b \quad \text{or} \quad A^H X = B.$$

(If A is real, then $A^H = A^T$.)

The procedure also has options to return an estimate of the *condition number* of A , and *forward* and *backward error bounds* for the computed solution or solutions. See the Chapter Introduction for an explanation of these terms. If error bounds are requested, the procedure performs iterative refinement of the computed solution in order to guarantee a small backward error.

2 Usage

USE `nag_gen_lin_sys`

CALL `nag_gen_lin_sol(a, b [, optional arguments])`

2.1 Interfaces

Distinct interfaces are provided for each of the four combinations of the following cases:

Real / complex data

Real data: a and b are of type `real(kind=wp)`.

Complex data: a and b are of type `complex(kind=wp)`.

One / many right-hand sides

One r.h.s.: b is a rank-1 array, and the optional arguments `bwd_err` and `fwd_err` are scalars.

Many r.h.s.: b is a rank-2 array, and the optional arguments `bwd_err` and `fwd_err` are rank-1 arrays.

3 Arguments

Note. All array arguments are assumed-shape arrays. The extent in each dimension must be exactly that required by the problem. Notation such as ' $x(n)$ ' is used in the argument descriptions to specify that the array x must have exactly n elements.

This procedure derives the values of the following problem parameters from the shape of the supplied arrays.

n — the order of the matrix A

r — the number of right-hand sides

3.1 Mandatory Arguments

a(n, n) — real(kind=wp)/complex(kind=wp), intent(inout)

Input: the general matrix A .

Output: A is overwritten by the factors L and U ; the unit diagonal elements of L are not stored.

b(n) / **b**(n, r) — real(kind=wp) / complex(kind=wp), intent(inout)

Input: the right-hand side vector b or matrix B .

Output: overwritten on exit by the solution vector x or matrix X .

Constraints: **b** must be of the same type as **a**.

Note: if optional error bounds are requested then the solution returned is that computed by iterative refinement.

3.2 Optional Arguments

Note. Optional arguments must be supplied by keyword, not by position. The order in which they are described below may differ from the order in which they occur in the argument list.

trans — character(len=1), intent(in), optional

Input: specifies whether the equations involve A or its transpose A^T or its conjugate-transpose A^H ($= A^T$ if A is real).

If **trans** = 'n' or 'N', the equations involve A (i.e., $Ax = b$);

if **trans** = 't' or 'T', the equations involve A^T (i.e., $A^T x = b$);

if **trans** = 'c' or 'C', the equations involve A^H (i.e., $A^H x = b$).

Default: **trans** = 'n'.

Constraints: **trans** = 'n', 'N', 't', 'T', 'c' or 'C'.

bwd_err / **bwd_err**(r) — real(kind=wp), intent(out), optional

Output: if **bwd_err** is a scalar, it returns the component-wise backward error bound for the single solution vector x . Otherwise, **bwd_err**(i) returns the component-wise backward error bound for the i th solution vector, returned in the i th column of **b**, for $i = 1, 2, \dots, r$.

Constraints: if **b** has rank 1, **bwd_err** must be a scalar; if **b** has rank 2, **bwd_err** must be a rank-1 array.

fwd_err / **fwd_err**(r) — real(kind=wp), intent(out), optional

Output: if **fwd_err** is a scalar, it returns an estimated bound for the forward error in the single solution vector x . Otherwise, **fwd_err**(i) returns an estimated bound for the forward error in the i th solution vector, returned in the i th column of **b**, for $i = 1, 2, \dots, r$.

Constraints: if **b** has rank 1, **fwd_err** must be a scalar; if **b** has rank 2, **fwd_err** must be a rank-1 array.

rcond — real(kind=wp), intent(out), optional

Output: an estimate of the reciprocal of the condition number of A ($\kappa_\infty(A)$, if **trans** = 'n' or 'N'; $\kappa_1(A)$ otherwise). **rcond** is set to zero if exact singularity is detected or the estimate underflows. If **rcond** is less than EPSILON(1.0_wp), then A is singular to working precision.

pivot(n) — integer, intent(out), optional

Output: the pivot indices used in the LU factorization; row i was interchanged with row **pivot**(i) for $i = 1, 2, \dots, n$.

error — type(nag_error), intent(inout), optional

The NAG *f190* error-handling argument. See the Essential Introduction, or the module document `nag_error_handling` (1.2). You are recommended to omit this argument if you are unsure how to use it. If this argument is supplied, it *must* be initialized by a call to `nag_set_error` before this procedure is called.

4 Error Codes

Fatal errors (error%level = 3):

error%code	Description
301	An input argument has an invalid value.
302	An array argument has an invalid shape.
303	Array arguments have inconsistent shapes.
320	The procedure was unable to allocate enough memory.

Failures (error%level = 2):

error%code	Description
201	Singular matrix. The matrix A has been factorized, but the factor U has a zero diagonal element, and so is exactly singular. No solutions or error bounds are computed.

Warnings (error%level = 1):

error%code	Description
101	Approximately singular matrix. The estimate of the reciprocal condition number (returned in <code>rcond</code> if present) is less than or equal to <code>EPSILON(1.0_wp)</code> . The matrix is singular to working precision, and it is likely that the computed solution returned in <code>b</code> has no accuracy at all. You should examine the forward error bounds returned in <code>fwd_err</code> , if present.

5 Examples of Usage

A complete example of the use of this procedure appears in Example 1 of this module document.

6 Further Comments

6.1 Algorithmic Detail

The procedure first calls `nag_gen_lin_fac` to factorize A , and to estimate the condition number. It then calls `nag_gen_lin_sol_fac` to compute the solution to the system of equations, and, if required, the error bounds. See the documents for those procedures for more details, and Chapter 3 of Golub and Van Loan [2] for background. The algorithms are derived from LAPACK (see Anderson *et al.* [1]).

6.2 Accuracy

The accuracy of the computed solution is given by the forward and backward error bounds which are returned in the optional arguments `fwd_err` and `bwd_err`.

The backward error bound `bwd_err` is rigorous; the forward error bound `fwd_err` is an estimate, but is almost always satisfied.

The condition number $\kappa_\infty(A)$ gives a general measure of the *sensitivity* of the solution of $Ax = b$, either to uncertainties in the data or to rounding errors in the computation. If the system has one of the alternative forms $A^T x = b$ or $A^H x = b$, the appropriate condition number is $\kappa_1(A)$ ($= \kappa_\infty(A^T) = \kappa_\infty(A^H)$). An estimate of the reciprocal of $\kappa_\infty(A)$ or $\kappa_1(A)$ is returned in the optional argument `rcond`. However, forward error bounds derived using this condition number may be more pessimistic than the bounds returned in `fwd_err`.

6.3 Timing

The time taken is roughly proportional to n^3 . The time taken for complex data is about 4 times as long as that for real data.

Procedure: nag_gen_lin_fac

1 Description

`nag_gen_lin_fac` computes the LU factorization of a general real or complex matrix A of order n , using partial pivoting with row interchanges:

$$A = PLU$$

where

P is a permutation matrix,

L is lower triangular with unit diagonal elements,

U is upper triangular.

This procedure can also return the determinant of A and estimates of the condition numbers $\kappa_1(A)$ and $\kappa_\infty(A)$.

2 Usage

USE `nag_gen_lin_sys`

CALL `nag_gen_lin_fac(a, pivot [, optional arguments])`

3 Arguments

Note. All array arguments are assumed-shape arrays. The extent in each dimension must be exactly that required by the problem. Notation such as ' $\mathbf{x}(n)$ ' is used in the argument descriptions to specify that the array \mathbf{x} must have exactly n elements.

This procedure derives the value of the following problem parameter from the shape of the supplied arrays.

n — the order of the matrix A

3.1 Mandatory Arguments

$\mathbf{a}(n, n)$ — real(kind= wp) / complex(kind= wp), intent(inout)

Input: the general matrix A .

Output: A is overwritten by the factors L and U ; the unit diagonal elements of L are not stored.

$\mathbf{pivot}(n)$ — integer, intent(out)

Output: the pivot indices; row i was interchanged with row $\mathbf{pivot}(i)$ for $i = 1, 2, \dots, n$.

3.2 Optional Arguments

Note. Optional arguments must be supplied by keyword, not by position. The order in which they are described below may differ from the order in which they occur in the argument list.

$\mathbf{rcond_inf}$ — real(kind= wp), intent(out), optional

Output: an estimate of the reciprocal of the condition number of A in the ∞ -norm, $\kappa_\infty(A)$.

$\mathbf{rcond_1}$ — real(kind= wp), intent(out), optional

Output: an estimate of the reciprocal of the condition number of A in the 1-norm, $\kappa_1(A)$.

det_frac — real(kind=wp) / complex(kind=wp), intent(out), optional

det_exp — integer, intent(out), optional

Output: **det_frac** returns the fractional part f , and **det_exp** returns the exponent e , of the determinant of A expressed as $f.b^e$, where b is the base of the representation of the floating point numbers (given by `RADIX(1.0_wp)`), or as `SCALE(det_frac, det_exp)`. The determinant is returned in this form to avoid the risk of overflow or underflow.

Constraints: **det_frac** must be of the same type as **a** and if either **det_frac** or **det_exp** is present the other must also be present.

error — type(nag_error), intent(inout), optional

The NAG *f90* error-handling argument. See the Essential Introduction, or the module document `nag_error_handling` (1.2). You are recommended to omit this argument if you are unsure how to use it. If this argument is supplied, it *must* be initialized by a call to `nag_set_error` before this procedure is called.

4 Error Codes

Fatal errors (error%level = 3):

error%code	Description
302	An array argument has an invalid shape.
303	Array arguments have inconsistent shapes.
305	Invalid absence of an optional argument.
320	The procedure was unable to allocate enough memory.

Failures (error%level = 2):

error%code	Description
201	Singular matrix. The matrix A has been factorized, but the factor U has a zero diagonal element, and so is exactly singular. If the factorization is used to solve a system of linear equations, an error will occur.

5 Examples of Usage

A complete example of the use of this procedure appears in Example 2 of this module document.

6 Further Comments

6.1 Algorithmic Detail

The LU factorization is computed using partial pivoting with row interchanges. See Chapter 3 of Golub and Van Loan [2].

To estimate the condition number $\kappa_1(A)$ ($= \|A\|_1 \|A^{-1}\|_1$), the procedure first computes $\|A\|_1$ directly, and then uses Higham's modification of Hager's method (see Higham [3]) to estimate $\|A^{-1}\|_1$. The procedure returns the reciprocal $\rho = 1/\kappa_1(A)$, rather than $\kappa_1(A)$ itself.

A similar approach is used to estimate $\kappa_\infty(A)$.

The algorithms are derived from LAPACK (see Anderson *et al.* [1]).

6.2 Accuracy

The computed factors L and U are the exact factors of a perturbed matrix $A + E$, such that

$$|E| \leq c(n)\epsilon P|L||U|,$$

where $c(n)$ is a modest linear function of n , and $\epsilon = \text{EPSILON}(1.0_wp)$.

The computed estimate `rcond_inf` or `rcond_1` is never less than the true value ρ , and in practice is nearly always less than 10ρ (although examples can be constructed where the computed estimate is much larger).

Since $\rho = 1/\kappa(A)$, this means that the procedure never overestimates the condition number, and hardly ever underestimates it by more than a factor of 10.

6.3 Timing

The total number of floating-point operations required for the LU factorization is roughly $(2/3)n^3$ for real A , and $(8/3)n^3$ for complex A .

Estimating the condition number involves solving a number of systems of linear equations with A or A^T as the coefficient matrix; the number is usually 4 or 5 and never more than 11. Each solution involves approximately $2n^2$ floating-point operations if A is real, or $8n^2$ if A is complex. Thus, for large n , the cost is much less than that of directly computing A^{-1} and its norm, which would require $O(n^3)$ operations.

Procedure: nag_gen_lin_sol_fac

1 Description

`nag_gen_lin_sol_fac` is a generic procedure which computes the solution of a general real or complex system of linear equations with one or many right-hand sides, assuming that the coefficient matrix has already been factorized by `nag_gen_lin_fac`.

We write:

$Ax = b$, if there is one right-hand side b ;

$AX = B$, if there are many right-hand sides (the columns of the matrix B).

The matrix A (the *coefficient matrix*) is assumed to be a general matrix, which has already been factorized by a call to `nag_gen_lin_fac`.

Optionally, the procedure can solve alternative forms of the system of equations:

$A^T x = b$, $A^T X = B$, $A^H x = b$ or $A^H X = B$.

(If A is real, then $A^H = A^T$.)

The procedure also has options to return *forward* and *backward error bounds* for the computed solution or solutions. See the Chapter Introduction for an explanation of these terms.

If error bounds are requested, the procedure performs iterative refinement of the solution in order to guarantee a small backward error in the computed solution; in this case, a copy of the original matrix A must be supplied in the optional argument `a`, as well as the factorized form in the argument `a_fac`.

2 Usage

USE `nag_gen_lin_sys`

CALL `nag_gen_lin_sol_fac(a_fac, pivot, b [, optional arguments])`

2.1 Interfaces

Distinct interfaces are provided for each of the four combinations of the following cases:

Real / complex data

Real data: `a_fac`, `b` and the optional argument `a` are of type `real(kind=wp)`.

Complex data: `a_fac`, `b` and the optional argument `a` are of type `complex(kind=wp)`.

One / many right-hand sides

One r.h.s.: `b` is a rank-1 array, and the optional arguments `bwd_err` and `fwd_err` are scalars.

Many r.h.s.: `b` is a rank-2 array, and the optional arguments `bwd_err` and `fwd_err` are rank-1 arrays.

3 Arguments

Note. All array arguments are assumed-shape arrays. The extent in each dimension must be exactly that required by the problem. Notation such as '`x(n)`' is used in the argument descriptions to specify that the array `x` must have exactly n elements.

This procedure derives the values of the following problem parameters from the shape of the supplied arrays.

n — the order of the matrix A

r — the number of right-hand sides

3.1 Mandatory Arguments

a_fac(n, n) — real(kind=wp) / complex(kind=wp), intent(in)

Input: the LU factorization of A , as returned by **nag_gen_lin_fac**.

pivot(n) — integer, intent(in)

Input: the pivot indices, as returned by **nag_gen_lin_fac**.

Constraints: $i \leq \text{pivot}(i) \leq n$, for $i = 1, 2, \dots, n$.

b(n) / **b**(n, r) — real(kind=wp) / complex(kind=wp), intent(inout)

Input: the right-hand side vector b or matrix B .

Output: overwritten on exit by the solution vector x or matrix X .

Constraints: **b** must be of the same type as **a_fac**.

Note: if optional error bounds are requested then the solution returned is that computed by iterative refinement.

3.2 Optional Arguments

Note. Optional arguments must be supplied by keyword, not by position. The order in which they are described below may differ from the order in which they occur in the argument list.

trans — character(len=1), intent(in), optional

Input: specifies whether the equations involve A or its transpose A^T or its conjugate-transpose A^H ($= A^T$ if A is real).

If **trans** = 'n' or 'N', the equations involve A (i.e., $Ax = b$);

if **trans** = 't' or 'T', the equations involve A^T (i.e., $A^T x = b$);

if **trans** = 'c' or 'C', the equations involve A^H (i.e., $A^H x = b$).

Default: **trans** = 'n'.

Constraints: **trans** = 'n', 'N', 't', 'T', 'c' or 'C'.

bwd_err / **bwd_err**(r) — real(kind=wp), intent(out), optional

Output: if **bwd_err** is a scalar, it returns the component-wise backward error bound for the single solution vector x . Otherwise, **bwd_err**(i) returns the component-wise backward error bound for the i th solution vector, returned in the i th column of **b**, for $i = 1, 2, \dots, r$.

Constraints: if **bwd_err** is present, the original matrix A must be supplied in **a**; if **b** has rank 1, **bwd_err** must be a scalar; if **b** has rank 2, **bwd_err** must be a rank-1 array.

fwd_err / **fwd_err**(r) — real(kind=wp), intent(out), optional

Output: if **fwd_err** is a scalar, it returns an estimated bound for the forward error in the single solution vector x . Otherwise, **fwd_err**(i) returns an estimated bound for the forward error in the i th solution vector, returned in the i th column of **b**, for $i = 1, 2, \dots, r$.

Constraints: if **fwd_err** is present, the original matrix A must be supplied in **a**; if **b** has rank 1, **fwd_err** must be a scalar; if **b** has rank 2, **fwd_err** must be a rank-1 array.

a(n, n) — real(kind=wp) / complex(kind=wp), intent(in), optional

Input: the original coefficient matrix A .

Constraints: **a** must be present if either **bwd_err** or **fwd_err** is present; **a** must be of the same type as **a_fac**.

error — type(nag_error), intent(inout), optional

The NAG *f90* error-handling argument. See the Essential Introduction, or the module document `nag_error_handling` (1.2). You are recommended to omit this argument if you are unsure how to use it. If this argument is supplied, it *must* be initialized by a call to `nag_set_error` before this procedure is called.

4 Error Codes

Fatal errors (error%level = 3):

error%code	Description
301	An input argument has an invalid value.
302	An array argument has an invalid shape.
303	Array arguments have inconsistent shapes.
305	Invalid absence of an optional argument.
320	The procedure was unable to allocate enough memory.

Failures (error%level = 2):

error%code	Description
201	Singular matrix. In the factorization supplied in <code>a_fac</code> , the factor U has a zero diagonal element, and so is exactly singular. No solutions or error bounds are computed.

5 Examples of Usage

A complete example of the use of this procedure appears in Example 2 of this module document.

6 Further Comments

6.1 Algorithmic Detail

The solution x is computed by forward and backward substitution:

to solve $Ax = b$ (in factorized form, $PLUx = b$), $PLy = b$ is solved for y , and then $Ux = y$ is solved for x ;

to solve $A^T x = b$ (in factorized form, $U^T L^T P^T x = b$), $U^T y = b$ is solved for y , and then $L^T P^T x = y$ is solved for x . A similar approach is used to solve $A^H x = b$.

If error bounds are requested (that is, `fwd_err` or `bwd_err` is present), iterative refinement of the solution is performed (in working precision), to reduce the backward error as far as possible.

The algorithms are derived from LAPACK (see Anderson *et al.* [1]).

6.2 Accuracy

The accuracy of the computed solution is given by the forward and backward error bounds which are returned in the optional arguments `fwd_err` and `bwd_err`.

The backward error bound `bwd_err` is rigorous; the forward error bound `fwd_err` is an estimate, but is almost always satisfied.

For each right-hand side b , the computed solution \hat{x} is the exact solution of a perturbed system of equations $(A + E)\hat{x} = b$, such that

$$|E| \leq c(n)\epsilon P|L||U|,$$

where $c(n)$ is a modest linear function of n , and $\epsilon = \text{EPSILON}(1.0_wp)$.

The condition number $\kappa_\infty(A)$ gives a general measure of the *sensitivity* of the solution of $Ax = b$, either to uncertainties in the data or to rounding errors in the computation. If the system has one of the alternative forms $A^T x = b$ or $A^H x = b$, the appropriate condition number is $\kappa_1(A)$ ($= \kappa_\infty(A^T) = \kappa_\infty(A^H)$). Estimates of the reciprocals of $\kappa_\infty(A)$ and $\kappa_1(A)$ are returned by `nag_gen_lin_fac` in its optional arguments `rcond_inf` and `rcond_1`. However, forward error bounds derived using these condition numbers may be more pessimistic than the bounds returned in `fwd_err`, if present.

If the reciprocal of the condition number $\leq \text{EPSILON}(1.0_wp)$, then A is singular to working precision; if the factorization is used to solve a system of linear equations, the computed solution may have no meaningful accuracy and should be treated with great caution.

6.3 Timing

The number of real floating-point operations required to compute the solutions is roughly $2n^2r$ if A is real, and $8n^2r$ if A is complex.

To compute the error bounds `fwd_err` and `bwd_err` usually requires about 5 times as much work.

Example 1: Solution of a General Real System of Linear Equations

Solve a general real system of linear equations with one right-hand side $Ax = b$, also estimating the condition number of A , and forward and backward error bounds on the computed solution. This example calls the single procedure `nag_gen_lin_sol`.

1 Program Text

Note. The listing of the example program presented below is double precision. Single precision users are referred to Section 5.2 of the Essential Introduction for further information.

```

PROGRAM nag_gen_lin_sys_ex01

! Example Program Text for nag_gen_lin_sys
! NAG f190, Release 4. NAG Copyright 2000.

! .. Use Statements ..
USE nag_examples_io, ONLY : nag_std_in, nag_std_out
USE nag_gen_lin_sys, ONLY : nag_gen_lin_sol
! .. Implicit None Statement ..
IMPLICIT NONE
! .. Intrinsic Functions ..
INTRINSIC KIND
! .. Parameters ..
INTEGER, PARAMETER :: wp = KIND(1.0D0)
! .. Local Scalars ..
INTEGER :: i, n
REAL (wp) :: bwd_err, fwd_err, rcond
CHARACTER (1) :: trans
! .. Local Arrays ..
REAL (wp), ALLOCATABLE :: a(:,,:), b(:)
! .. Executable Statements ..
WRITE (nag_std_out,*) 'Example Program Results for nag_gen_lin_sys_ex01'

READ (nag_std_in,*)          ! Skip heading in data file
READ (nag_std_in,*) n
READ (nag_std_in,*) trans

ALLOCATE (a(n,n),b(n))      ! Allocate storage

READ (nag_std_in,*) (a(i,:),i=1,n)
READ (nag_std_in,*) b

! Solve the system of equations

CALL nag_gen_lin_sol(a,b,trans=trans,bwd_err=bwd_err,fwd_err=fwd_err, &
  rcond=rcond)

WRITE (nag_std_out,*)
WRITE (nag_std_out,'(1X,''kappa(A) (1/rcond)''/2X,ES11.2)') 1/rcond
WRITE (nag_std_out,*)
WRITE (nag_std_out,*) 'Solution'
WRITE (nag_std_out,'(4X,F9.4)') b
WRITE (nag_std_out,*)
WRITE (nag_std_out,*) 'Backward error bound'
WRITE (nag_std_out,'(2X,ES11.2)') bwd_err
WRITE (nag_std_out,*)
WRITE (nag_std_out,*) 'Forward error bound (estimate)'
WRITE (nag_std_out,'(2X,ES11.2)') fwd_err

DEALLOCATE (a,b)           ! Deallocate storage

```

```
END PROGRAM nag_gen_lin_sys_ex01
```

2 Program Data

Example Program Data for nag_gen_lin_sys_ex01

```
4 : Value of n
'N' : Value of trans
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 right-hand side vector b
```

3 Program Results

Example Program Results for nag_gen_lin_sys_ex01

```
kappa(A) (1/rcond)
1.41E+02
```

```
Solution
1.0000
-1.0000
3.0000
-5.0000
```

```
Backward error bound
1.41E-16
```

```
Forward error bound (estimate)
4.68E-14
```

Example 2: Factorization of a General Real Matrix and Use of the Factorization to Solve a System of Linear Equations

Solve a general real system of linear equations with many right-hand sides $AX = B$, with forward and backward error bounds on the computed solutions. This example calls the procedure `nag_gen_lin_fac` to factorize A , and then `nag_gen_lin_sol_fac` to solve the equations using the factorization.

1 Program Text

Note. The listing of the example program presented below is double precision. Single precision users are referred to Section 5.2 of the Essential Introduction for further information.

```

PROGRAM nag_gen_lin_sys_ex02

! Example Program Text for nag_gen_lin_sys
! NAG f190, Release 4. NAG Copyright 2000.

! .. Use Statements ..
USE nag_examples_io, ONLY : nag_std_in, nag_std_out
USE nag_gen_lin_sys, ONLY : nag_gen_lin_fac, nag_gen_lin_sol_fac
USE nag_write_mat, ONLY : nag_write_gen_mat
! .. Implicit None Statement ..
IMPLICIT NONE
! .. Intrinsic Functions ..
INTRINSIC EPSILON, KIND, SCALE
! .. Parameters ..
INTEGER, PARAMETER :: wp = KIND(1.0D0)
! .. Local Scalars ..
INTEGER :: det_exp, i, n, nrhs
REAL (wp) :: det_frac, rcond
CHARACTER (1) :: trans
! .. Local Arrays ..
INTEGER, ALLOCATABLE :: pivot(:)
REAL (wp), ALLOCATABLE :: a(:,,:), a_fac(:,,:), b(:,,:), bwd_err(:), &
  fwd_err(:)
! .. Executable Statements ..
WRITE (nag_std_out,*) 'Example Program Results for nag_gen_lin_sys_ex02'

READ (nag_std_in,*)          ! Skip heading in data file
READ (nag_std_in,*) n, nrhs
READ (nag_std_in,*) trans

ALLOCATE (a(n,n),a_fac(n,n),b(n,nrhs),bwd_err(nrhs),fwd_err(nrhs), &
  pivot(n))                  ! Allocate storage

READ (nag_std_in,*) (a(i,:),i=1,n)
a_fac = a

READ (nag_std_in,*) (b(i,:),i=1,n)

! Carry out the LU factorization

SELECT CASE (trans)
CASE ('C','c','T','t')

  CALL nag_gen_lin_fac(a_fac,pivot,rcond_1=rcond,det_frac=det_frac, &
    det_exp=det_exp)

CASE ('N','n')

  CALL nag_gen_lin_fac(a_fac,pivot,rcond_inf=rcond,det_frac=det_frac, &
    det_exp=det_exp)

```

```

END SELECT

! Print the LU decomposition

WRITE (nag_std_out,*)
WRITE (nag_std_out,*) 'Results of the LU factorization'
WRITE (nag_std_out,*)

CALL nag_write_gen_mat(a_fac,title='Factorized matrix')

WRITE (nag_std_out,*)
WRITE (nag_std_out,*) 'Pivotal sequence (pivot)'
WRITE (nag_std_out,'(2X,10I4:)') pivot
WRITE (nag_std_out,*)
WRITE (nag_std_out, &
' (1X,'determinant = SCALE(det_frac,det_exp) =',2X,ES11.3)') &
SCALE(det_frac,det_exp)
WRITE (nag_std_out,*)
WRITE (nag_std_out,'(1X,'kappa(A) (1/rcond)'/2X,ES11.2)') 1/rcond

IF (rcond<=EPSILON(1.0_wp)) THEN
  WRITE (nag_std_out,*)
  WRITE (nag_std_out,*) ' ** WARNING ** '
  WRITE (nag_std_out,*) 'The matrix is almost singular: the ' // &
'solution may have no accuracy.'
  WRITE (nag_std_out,*) 'Examine the forward error bounds ' // &
'estimates returned in fwd_err.'
END IF

! Solve the system of equations

CALL nag_gen_lin_sol_fac(a_fac,pivot,b,trans=trans,a=a,bwd_err=bwd_err, &
fwd_err=fwd_err)

WRITE (nag_std_out,*)
WRITE (nag_std_out,*) &
'Results of the solution of the simultaneous equations'
WRITE (nag_std_out,*)

CALL nag_write_gen_mat(b,int_col_labels=.TRUE., &
title='Solutions (one solution per column)')

WRITE (nag_std_out,*)
WRITE (nag_std_out,*) 'Backward error bounds'
WRITE (nag_std_out,'(2X,4ES11.2)') bwd_err
WRITE (nag_std_out,*)
WRITE (nag_std_out,*) 'Forward error bounds (estimates)'
WRITE (nag_std_out,'(2X,4ES11.2)') fwd_err

DEALLOCATE (a,a_fac,b,bwd_err,fwd_err,pivot) ! Deallocate storage

END PROGRAM nag_gen_lin_sys_ex02

```

2 Program Data

Example Program Data for nag_gen_lin_sys_ex02

```

4 2 : Values of n, nrhs
'N' : Value of trans
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 18.47
24.35 2.25
0.77 -13.28
-6.22 -6.21 : End of right-hand sides (one rhs per column)

```

3 Program Results

Example Program Results for nag_gen_lin_sys_ex02

Results of the LU factorization

Factorized matrix

```

5.2500 -2.9500 -0.9500 -3.8000
0.3429 3.8914 2.3757 0.4129
0.3010 -0.4631 -1.5139 0.2948
-0.2114 -0.3299 0.0047 0.1314

```

Pivotal sequence (pivot)

```

2 2 3 4

```

determinant = SCALE(det_frac,det_exp) = 4.063E+00

kappa(A) (1/rcond)

```

1.41E+02

```

Results of the solution of the simultaneous equations

Solutions (one solution per column)

```

1 2
1.0000 3.0000
-1.0000 2.0000
3.0000 4.0000
-5.0000 1.0000

```

Backward error bounds

```

1.41E-16 3.73E-17

```

Forward error bounds (estimates)

```

4.68E-14 6.48E-14

```


Additional Examples

Not all example programs supplied with NAG *f*90 appear in full in this module document. The following additional examples, associated with this module, are available.

`nag_gen_lin_sys_ex03`

Solution of a general real system of linear equations with many right-hand sides.

`nag_gen_lin_sys_ex04`

Solution of a general complex system of linear equations with one right-hand side.

`nag_gen_lin_sys_ex05`

Factorization of a general complex matrix and use of the factorization to solve a system of linear equations with many right-hand sides.

`nag_gen_lin_sys_ex06`

Solution of a general complex system of linear equations with many right-hand sides.

References

- [1] Anderson E, Bai Z, Bischof C, Demmel J, Dongarra J J, Du Croz J J, Greenbaum A, Hammarling S, McKenney A, Blackford S and Sorensen D (1999) *LAPACK Users' Guide* (3rd Edition) SIAM, Philadelphia
- [2] Golub G H and Van Loan C F (1989) *Matrix Computations* Johns Hopkins University Press (2nd Edition)
- [3] Higham N J (1988) Algorithm 674: Fortran codes for estimating the one-norm of a real or complex matrix, with applications to condition estimation *ACM Trans. Math. Software* **14** 381–396