

nag_opt_lin_lsq (e04ncc)

1. Purpose

nag_opt_lin_lsq solves linearly constrained linear least-squares problems and convex quadratic programming problems. It is not intended for large sparse problems.

2. Specification

```
#include <nag.h>
#include <nage04.h>

void nag_opt_lin_lsq(Integer m, Integer n, Integer nclin, double a[],
                    Integer tda, double bl[], double bu[], double cvec[],
                    double b[], double h[], Integer tdh, Integer kx[],
                    double x[], double *objf, Nag_E04_Opt *options,
                    Nag_Comm *comm, NagError *fail)
```

3. Description

nag_opt_lin_lsq is designed to solve a class of quadratic programming problems stated in the following general form:

$$\underset{x \in R^n}{\text{minimize}} \quad F(x) \quad \text{subject to} \quad l \leq \begin{Bmatrix} x \\ Ax \end{Bmatrix} \leq u, \quad (1)$$

where A is an n_L by n matrix and the objective function $F(x)$ may be specified in a variety of ways depending upon the particular problem to be solved. The available forms for $F(x)$ are listed in Table 1 below, in which the prefixes FP, LP, QP and LS stand for ‘feasible point’, ‘linear programming’, ‘quadratic programming’ and ‘least-squares’ respectively, c is an n element vector, b is an m element vector, and $\|x\|$ denotes the Euclidean length of x .

Problem Type	$F(x)$	Matrix H
FP	Not applicable	Not applicable
LP	$c^T x$	Not applicable
QP1	$\frac{1}{2}x^T H x$	n by n symmetric positive semi-definite
QP2	$c^T x + \frac{1}{2}x^T H x$	n by n symmetric positive semi-definite
QP3	$\frac{1}{2}x^T H^T H x$	m by n upper trapezoidal
QP4	$c^T x + \frac{1}{2}x^T H^T H x$	m by n upper trapezoidal
LS1	$\frac{1}{2}\ b - Hx\ ^2$	m by n
LS2	$c^T x + \frac{1}{2}\ b - Hx\ ^2$	m by n
LS3	$\frac{1}{2}\ b - Hx\ ^2$	m by n upper trapezoidal
LS4	$c^T x + \frac{1}{2}\ b - Hx\ ^2$	m by n upper trapezoidal

Table 1

For problems of type LS, H is referred to as the *least-squares* matrix, or the *matrix of observations*, and b as the *vector of observations*. The default problem type is LS1, and other objective functions are selected by using the optional parameter **prob** (see Section 8.2).

When H is upper trapezoidal it will usually be the case that $m = n$, so that H is upper triangular, but full generality has been allowed for in the specification of the problem. The upper trapezoidal form is intended for cases where a previous factorization, such as a QR factorization, has been performed.

The constraints involving A are called the *general* constraints. Note that upper and lower bounds are specified for all the variables and for all the general constraints. An equality constraint can be specified by setting $l_i = u_i$. If certain bounds are not present, the associated elements of l or u can be set to special values that will be treated as $-\infty$ or $+\infty$. (See the description of the optional parameter **inf_bound** in Section 8.2.)

The function $F(x)$ is a quadratic function, whose defining feature is that its second-derivative matrix $\nabla^2 F(x)$ (the *Hessian matrix*) is constant. For the LP case, $\nabla^2 F(x) = 0$; for QP1 and QP2, $\nabla^2 F(x) = H$; and for QP3, QP4 and LS problems, $\nabla^2 F(x) = H^T H$ and the Hessian matrix is positive semi-definite (positive definite if H is full rank), so that $F(x)$ is convex. If H is defined as the zero matrix, nag_opt_lin_lsq will solve the resulting linear programming problem; however, this can be accomplished more efficiently by using nag_opt_lp (e04mfc).

Problems of type QP3 and QP4 for which H is not in upper trapezoidal form should be solved as problems of type LS1 and LS2 respectively, with $b = 0$.

The user must supply an initial estimate of the solution.

If H is of full rank then nag_opt_lin_lsq will obtain the unique (global) minimum. If H is not of full rank then the solution may still be a global minimum if all active constraints have non-zero Lagrange multipliers. Otherwise the solution obtained will be either a weak minimum (i.e., with a unique optimal objective value, but an infinite set of optimal x), or else the objective function is unbounded below in the feasible region. The last case can only occur when $F(x)$ contains an explicit linear term (as in problems LP, QP2, QP4, LS2 and LS4).

The method used by nag_opt_lin_lsq is described in detail in Section 7.

4. Parameters

m

Input: m , the number of rows in the matrix H . If the problem is of type FP or LP, **m** is not referenced and is assumed to be zero. The default type is LS1; other problem types can be specified using the optional parameter **prob**, see Section 8.2.

If the problem is of type QP, **m** will usually be n , the number of variables. However, a value of **m** less than n is appropriate for problem type QP3 or QP4 if H is an upper trapezoidal matrix with m rows. Similarly, **m** may be used to define the dimension of a leading block of non-zeros in the Hessian matrices of QP1 or QP2. In QP cases, m should not be greater than n ; if it is, the last $(m - n)$ rows of H are ignored.

If the problem is a least-squares problem (in particular, the default type LS1), **m** is also the dimension of the array **b**. Note that all possibilities ($m < n$, $m = n$ and $m > n$) are allowed in this case.

Constraint: **m** > 0 if problem is not FP or LP.

n

Input: n , the number of variables.

Constraint: **n** > 0.

nclin

Input: n_L , the number of general linear constraints.

Constraint: **nclin** \geq 0.

a[nclin][tda]

Input: the i th row of **a** must contain the coefficients of the i th general linear constraint (the i th row of A), for $i = 1, 2, \dots, n_L$.

If **nclin** = 0 then the array **a** is not referenced.

tda

Input: the second dimension of the array **a** as declared in the function from which nag_opt_lin_lsq is called.

Constraint: **tda** \geq **n** if **nclin** > 0.

bl[n+nclin]

bu[n+nclin]

Input: **bl** must contain the lower bounds and **bu** the upper bounds, for all the constraints in the following order. The first n elements of each array must contain the bounds on the variables, and the next n_L elements the bounds for the general linear constraints (if any). To specify a non-existent lower bound (i.e., $l_j = -\infty$), set **bl**[$j - 1$] \leq **inf_bound**, and to specify a non-existent upper bound (i.e., $u_j = +\infty$), set **bu**[$j - 1$] \geq **inf_bound**, where **inf_bound** is one

of the optional parameters (default value 10^{20} (see Section 8.2)). To specify the j th constraint as an equality, set $\mathbf{bl}[j - 1] = \mathbf{bu}[j - 1] = \beta$, say, where $|\beta| < \mathbf{inf_bound}$.

Constraints:

$$\mathbf{bl}[j] \leq \mathbf{bu}[j], \text{ for } j = 0, 1, \dots, \mathbf{n+nc_lin}-1,$$

$$|\beta| < \mathbf{inf_bound} \text{ when } \mathbf{bl}[j] = \mathbf{bu}[j] = \beta.$$

cvec[n]

Input: the coefficients of the explicit linear term of the objective function when the problem is of type LP, QP2, QP4, LS2 or LS4.

If the problem is of type FP, QP1, QP3, LS1 (the default) or LS3, **cvec** is not referenced and may be set to the null pointer.

b[m]

Input: the m elements of the vector of observations.

Output: the transformed residual vector of equation (10) (see Section 7.3).

b is referenced only in the case of least-squares problem types (in particular, the default type LS1. For other problem types, **b** may be set to the null pointer.

h[m][tdh]

Input: the array **h** must contain the matrix H as specified in Table 1 (see Section 3).

For problems QP1 and QP2, the first m rows and columns of **h** must contain the leading m by m rows and columns of the symmetric Hessian matrix. Only the diagonal and upper triangular elements of the leading m rows and columns of **h** are referenced. The remaining elements are assumed to be zero and need not be assigned.

For problems QP3, QP4, LS3 and LS4, the first m rows of **h** must contain an m by n upper trapezoidal factor of either the Hessian or the least-squares matrix, ordered according to the array **kx** (see below). The factor need not be of full rank, i.e., some of the diagonals may be zero. However, as a general rule, the larger the dimension of the leading non-singular sub-matrix of H , the fewer iterations will be required. Elements outside the upper trapezoidal part of the first m rows of H are assumed to be zero and need not be assigned.

If a constrained least-squares problem contains a very large number of observations, storage limitations may prevent storage of the entire least-squares matrix. In such cases, the user should transform the original H into a triangular matrix before the call to `nag_opt_lin_lsq` and solve as type LS3 or LS4.

Output: by default, **h** contains the upper triangular Cholesky factor R of equation (8) (see Section 7.3), with columns ordered as indicated by **kx** (see below). If the optional parameter **hessian** = **TRUE** (see Section 8.2), and the problem is one of the LS or QP types, **h** contains the upper triangular Cholesky factor of the Hessian matrix $\nabla^2 F$, with columns ordered as indicated by **kx** (see below). In either case, this matrix may be used to obtain the variance-covariance matrix or to recover the upper triangular factor of the original least-squares matrix.

If the problem is of type FP or LP, **h** is not referenced and may be set to the null pointer.

tdh

Input: the second dimension of the array **h** as declared in the function from which `nag_opt_lin_lsq` is called.

Constraint: **tdh** \geq **n**.

kx[n]

Input: for problems of type QP3, QP4, LS3 or LS4 the array **kx** must specify the order of the columns of the matrix H with respect to the ordering of **x**. Thus if column j of H is the column associated with the variable x_i then $\mathbf{kx}[j - 1] = i$.

If the problem is of any other type then the array **kx** need not be initialized.

Constraints:

$$1 \leq \mathbf{kx}[i] \leq \mathbf{n}, \text{ for } i = 0, 1, \dots, \mathbf{n}-1,$$

$$\mathbf{kx}[i] \neq \mathbf{kx}[j] \text{ whenever } i \neq j.$$

Output: **kx** defines the order of the columns of H with respect to the ordering of **x**, as described above.

x[n]

Input: an initial estimate of the solution.

Output: the point at which nag_opt_lin_lsq terminated. If **fail.code** = **NE_NOERROR**, **NW_SOLN_NOT_UNIQUE** or **NW_NOT_FEASIBLE**, **x** contains an estimate of the solution.**objf**Output: the value of the objective function at x if x is feasible, or the sum of infeasibilities at x otherwise. If the problem is of type FP and x is feasible, **objf** is set to zero.**options**Input/Output: a pointer to a structure of type Nag_E04_Opt whose members are optional parameters for nag_opt_lin_lsq. These structure members offer the means of adjusting some of the parameter values of the algorithm and on output will supply further details of the results. A description of the members of **options** is given below in Section 8. Some of the results returned in **options** can be used by nag_opt_lin_lsq to perform a ‘warm start’ (see the member **start** in Section 8.2).If any of these optional parameters are required then the structure **options** should be declared and initialized by a call to nag_opt_init (e04xxc) and supplied as an argument to nag_opt_lin_lsq. However, if the optional parameters are not required the NAG defined null pointer, **E04_DEFAULT**, can be used in the function call.**comm**Input/Output: structure containing pointers for communication with an optional user-defined printing function; see Section 8.3.1 for details. If the user does not need to make use of this communication feature the null pointer **NAGCOMM_NULL** may be used in the call to nag_opt_lin_lsq; **comm** will then be declared internally for use in calls to user-supplied functions.**fail**

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

Users are recommended to declare and initialize **fail** and set **fail.print** = **TRUE** for this function.

4.1. Description of Printed Output

Intermediate and final results are printed out by default. The level of printed output can be controlled by the user with the structure member **options.print_level** (see Section 8.2). The default print level of **Nag_Soln_Iter** provides a single line of output at each iteration and the final result. This section describes the default printout produced by nag_opt_lin_lsq.The convention for numbering the constraints in the iteration results is that indices 1 to n refer to the bounds on the variables, and indices $n + 1$ to $n + n_L$ refer to the general constraints.The following line of output is produced at every iteration. In all case, the values of the quantities printed are those in effect *on completion* of the the given iteration.

Itn	is the iteration count.
Step	is the step taken along the computed search direction. If a constraint is added during the current iteration, Step will be the step to the nearest constraint. During the optimality phase, the step can be greater than 1.0 only if the factor R_z is singular (see Section 7.3).
Ninf	is the number of violated constraints (infeasibilities). This will be zero during the optimality phase.
Sinf/Objective	is the value of the current objective function. If x is not feasible, Sinf gives a weighted sum of the magnitudes of constraint violations. If x is feasible, Objective is the value of the objective function. The output line for the final iteration of the feasibility phase (i.e., the first iteration for which Ninf is zero) will give the value of the true objective at the first feasible point.

During the optimality phase, the value of the objective function will be non-increasing. During the feasibility phase, the number of constraint infeasibilities

will not increase until either a feasible point is found, or the optimality of the multipliers implies that no feasible point exists. Once optimal multipliers are obtained, the number of infeasibilities can increase, but the sum of infeasibilities will either remain constant or be reduced until the minimum sum of infeasibilities is found.

Norm Gz $\|Z_1^T g_{FR}\|$, the Euclidean norm of the reduced gradient with respect to Z_1 (see Section 7.3). During the optimality phase, this norm will be approximately zero after a unit step.

The printout of the final result consists of:

Varbl gives the name (V) and index j , for $j = 1, 2, \dots, n$ of the variable.

State gives the state of the variable (FR if neither bound is in the working set, EQ if a fixed variable, LL if on its lower bound, UL if on its upper bound, TF if temporarily fixed at its current value). If **Value** lies outside the upper or lower bounds by more than the optional parameter **ftol** (default value $\sqrt{\epsilon}$, where ϵ is the **machine precision**; see Section 8.2), **State** will be ++ or -- respectively.

A key is sometimes printed before **State** to give some additional information about the state of a variable.

A *Alternative optimum possible.* The variable is active at one of its bounds, but its Lagrange Multiplier is essentially zero. This means that if the variable were allowed to start moving away from its bound, there would be no change to the objective function. The values of the other free variables *might* change, giving a genuine alternative solution. However, if there are any degenerate variables (labelled D), the actual change might prove to be zero, since one of them could encounter a bound immediately. In either case, the values of the Lagrange multipliers might also change.

D *Degenerate.* The variable is free, but it is equal to (or very close to) one of its bounds.

I *Infeasible.* The variable is currently violating one of its bounds by more than **ftol**.

Value is the value of the variable at the final iteration.

Lower bound is the lower bound specified for variable j . (**None** indicates that $\mathbf{bl}[j-1] \leq -\mathbf{inf_bound}$, where **inf_bound** is the optional parameter.)

Upper bound is the upper bound specified for variable j . (**None** indicates that $\mathbf{bu}[j-1] \geq \mathbf{inf_bound}$, where **inf_bound** is the optional parameter.)

Lagr mult is the value of the Lagrange multiplier for the associated bound. This will be zero if **State** is FR unless $\mathbf{bl}[j-1] \leq -\mathbf{inf_bound}$ and $\mathbf{bu}[j-1] \geq \mathbf{inf_bound}$, in which case the entry will be blank. If x is optimal, the multiplier should be non-negative if **State** is LL, and non-positive if **State** is UL.

Residual is the difference between the variable **Value** and the nearer of its (finite) bounds $\mathbf{bl}[j-1]$ and $\mathbf{bu}[j-1]$. A blank entry indicates that the associated variable is not bounded (i.e., $\mathbf{bl}[j-1] \leq -\mathbf{inf_bound}$ and $\mathbf{bu}[j-1] \geq \mathbf{inf_bound}$).

The meaning of the printout for general constraints is the same as that given above for variables, with ‘variable’ replaced by ‘constraint’, $\mathbf{bl}[j-1]$ and $\mathbf{bu}[j-1]$ replaced by $\mathbf{bl}[n+j-1]$ and $\mathbf{bu}[n+j-1]$ respectively, and with the following change in the heading:

L Con the name (L) and index j , for $j = 1, 2, \dots, n_L$ of the linear constraint.

Note that movement off a constraint (as opposed to a variable moving away from its bound) can be interpreted as allowing the entry in the **Residual** column to become positive.

Numerical values are output with a fixed number of digits; they are not guaranteed to be accurate to this precision.

5. Comments

A list of possible error exits and warnings from nag_opt_lin_lsq is given in Section 9. Scaling and accuracy are considered in Section 10.

6. Example 1

To minimize the function $\frac{1}{2}\|b - Hx\|^2$, where

$$H = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 2 & 1 & 1 & 1 & 1 & 2 & 0 & 0 \\ 1 & 1 & 3 & 1 & 1 & 1 & -1 & -1 & -3 \\ 1 & 1 & 1 & 4 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 3 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 2 & 1 & 1 & 0 & 0 & 0 & -1 \\ 1 & 1 & 1 & 1 & 0 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 1 & 1 & 1 & 2 & 2 & 3 \\ 1 & 0 & 1 & 1 & 1 & 1 & 0 & 2 & 2 \end{pmatrix} \quad \text{and} \quad b = \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}$$

subject to the bounds

$$\begin{aligned} 0 &\leq x_1 \leq 2 \\ 0 &\leq x_2 \leq 2 \\ -\infty &\leq x_3 \leq 2 \\ 0 &\leq x_4 \leq 2 \\ 0 &\leq x_5 \leq 2 \\ 0 &\leq x_6 \leq 2 \\ 0 &\leq x_7 \leq 2 \\ 0 &\leq x_8 \leq 2 \\ 0 &\leq x_9 \leq 2 \end{aligned}$$

and to the general constraints

$$\begin{aligned} 2.0 &\leq x_1 + x_2 + x_3 + x_4 + x_5 + x_6 + x_7 + x_8 + 4x_9 \leq \infty \\ -\infty &\leq x_1 + 2x_2 + 3x_3 + 4x_4 - 2x_5 + x_6 + x_7 + x_8 + x_9 \leq 2.0 \\ 1.0 &\leq x_1 - x_2 + x_3 - x_4 + x_5 + x_6 + x_7 + x_8 + x_9 \leq 4.0 \end{aligned}$$

The initial point, which is infeasible, is

$$x_0 = (1.0, 0.5, 0.3333, 0.25, 0.2, 0.1667, 0.1428, 0.125, 0.1111)^T,$$

and $F(x_0) = 9.4746$ (to five figures).

The optimal solution (to five figures) is

$$x^* = (0.0, 0.041526, 0.58718, 0.0, 0.099643, 0.0, 0.04906, 0.0, 0.30565)^T,$$

and $F(x^*) = 0.081341$. Four bound constraints and all three general constraints are active at the solution.

This example shows the simple use of nag_opt_lin_lsq where default values are used for all optional parameters. An example showing the use of optional parameters is given in Section 13. There is one example program file, the main program of which calls both examples. The main program and Example 1 are given below.

6.1. Program Text

```

/* nag_opt_lin_lsq(e04ncc) Example Program.
 *
 * Copyright 1998 Numerical Algorithms Group.
 *
 * Mark 5, 1998.
 */
#include <nag.h>
#include <stdio.h>
#include <nag_stdlib.h>
#include <nage04.h>

#ifdef NAG_PROTO
static void ex1(void);
static void ex2(void);
#else
static void ex1();
static void ex2();
#endif

main()
{
    Vprintf("e04ncc Example Program Results.\n");
    ex1();
    ex2();
    exit(EXIT_SUCCESS);
}

#ifdef NAG_PROTO
static void ex1(void)
#else
static void ex1()
#endif
{
#define MMAX 10
#define NMAX 10
#define NCLIN 10
#define MAXBND NMAX+NCLIN

    /* Local variables */
    double a[NCLIN][NMAX], b[MMAX], bl[MAXBND], bu[MAXBND];
    double h[MMAX][NMAX], x[NMAX];
    double objf;
    Integer kx[NMAX];
    Integer i, j;
    Integer m, n, nclin;
    Integer tda, tdh;
    static NagError fail;

    fail.print = TRUE;

    Vprintf("\nExample 1: default options\n");
    Vscanf(" %*[^\\n]"); /* Skip heading in data file */
    Vscanf(" %*[^\\n]");

    /* Read problem dimensions */
    Vscanf(" %*[^\\n]");
    Vscanf("%ld%ld%ld%*[^\\n]", &m, &n, &nclin);

    if (m <= MMAX && n <= NMAX && nclin <= NCLIN)
    {
        tda = NMAX;
        tdh = NMAX;

        /* Read h, b, a, bl, bu and x from data file */

        Vscanf(" %*[^\\n]");

```

```

    for (i = 0; i < m; ++i)
        for (j = 0; j < n; ++j)
            Vscanf("%lf",&h[i][j]);

    Vscanf(" %*[\n]");
    for (i = 0; i < m; ++i)
        Vscanf("%lf",&b[i]);

    if (nclin > 0)
    {
        Vscanf(" %*[\n]");
        for (i = 0; i < nclin; ++i)
            for (j = 0; j < n; ++j)
                Vscanf("%lf",&a[i][j]);
    }

    /* Read lower bounds */
    Vscanf(" %*[\n]");
    for (i = 0; i < n + nclin; ++i)
        Vscanf("%lf",&bl[i]);

    /* Read upper bounds */
    Vscanf(" %*[\n]");
    for (i = 0; i < n + nclin; ++i)
        Vscanf("%lf",&bu[i]);

    /* Read the initial point x */
    Vscanf(" %*[\n]");
    for (i = 0; i < n; ++i)
        Vscanf("%lf",&x[i]);

    e04ncc(m, n, nclin, (double*)a, tda, bl, bu, (double*)0, b,
          (double*)h, tdh, kx, x, &objf,
          E04_DEFAULT, NAGCOMM_NULL, &fail);
}
} /* ex1 */

```

6.2. Program Data

e04ncc Example Program Data

Data for example 1

Values of m, n, nclin

10 9 3

Objective function matrix H

1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
1.0	2.0	1.0	1.0	1.0	1.0	2.0	0.0	0.0
1.0	1.0	3.0	1.0	1.0	1.0	-1.0	-1.0	-3.0
1.0	1.0	1.0	4.0	1.0	1.0	1.0	1.0	1.0
1.0	1.0	1.0	3.0	1.0	1.0	1.0	1.0	1.0
1.0	1.0	2.0	1.0	1.0	0.0	0.0	0.0	-1.0
1.0	1.0	1.0	1.0	0.0	1.0	1.0	1.0	1.0
1.0	1.0	1.0	0.0	1.0	1.0	1.0	1.0	1.0
1.0	1.0	0.0	1.0	1.0	1.0	2.0	2.0	3.0
1.0	0.0	1.0	1.0	1.0	1.0	0.0	2.0	2.0

Vector of observations - array b

1.0 1.0 1.0 1.0 1.0 1.0 1.0 1.0 1.0 1.0

Linear constraint matrix A

1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	4.0
1.0	2.0	3.0	4.0	-2.0	1.0	1.0	1.0	1.0
1.0	-1.0	1.0	-1.0	1.0	1.0	1.0	1.0	1.0

Lower bounds

0.0	0.0	-1.0e+25	0.0	0.0	0.0	0.0	0.0	0.0	0.0
2.0	-1.0e+25	1.0							


```

Upper bounds
  2.0      2.0      2.0      2.0      2.0      2.0      2.0      2.0
  1.0e+25  2.0      4.0

Initial estimate of x
  1.0      0.5      0.3333  0.25      0.2      0.1667      0.1428      0.125      0.1111

```

6.3. Program Results

e04ncc Example Program Results.

Example 1: default options

Parameters to e04ncc

```

Linear constraints..... 3      Number of variables..... 9
Objective matrix rows..... 10

```

```

prob..... Nag_LS1      start..... Nag_Cold
ftol..... 1.05e-08      rank_tol..... 1.11e-14
crash_tol..... 1.00e-02      hessian..... FALSE
inf_bound..... 1.00e+20      inf_step..... 1.00e+20
fmax_iter..... 60      max_iter..... 60
machine precision..... 1.11e-16
print_level..... Nag_Soln_Iter
outfile..... stdout

```

Memory allocation:

```

state..... Nag
ax..... Nag      lambda..... Nag

```

Rank of the objective function data matrix = 6

It	n	Step	Ninf	Sinf/Objective	Norm Gz
0		0.0e+00	1	2.145500e+00	0.0e+00
1		2.5e-01	1	1.145500e+00	0.0e+00
2		3.8e-01	0	6.595685e+00	2.2e+01
3		1.8e-02	0	6.365562e+00	1.6e+01
4		3.3e-01	0	2.869136e+00	8.5e+00
5		1.4e-01	0	2.256269e+00	7.1e+00
6		1.2e-01	0	1.899743e+00	5.4e+00
7		1.1e-01	0	1.770819e+00	0.0e+00
8		1.0e+00	0	1.390865e+00	2.7e-15
9		1.7e-01	0	9.782704e-01	3.3e-01
10		1.0e+00	0	9.766042e-01	2.5e-15
11		6.3e-01	0	2.190278e-01	5.9e-01
12		1.0e+00	0	1.652065e-01	3.4e-15
13		1.0e+00	0	9.605160e-02	6.5e-15
14		3.0e-02	0	9.236999e-02	4.5e-01
15		1.0e+00	0	8.134082e-02	6.0e-15

Exit from LS problem after 15 iterations.

Varbl	State	Value	Lower Bound	Upper Bound	Lagr Mult	Residual
V 1	LL	0.00000e+00	0.00000e+00	2.00000e+00	1.5715e-01	0.0000e+00
V 2	FR	4.15261e-02	0.00000e+00	2.00000e+00	0.0000e+00	4.1526e-02
V 3	FR	5.87176e-01	None	2.00000e+00	0.0000e+00	1.4128e+00
V 4	LL	0.00000e+00	0.00000e+00	2.00000e+00	8.7817e-01	0.0000e+00
V 5	FR	9.96432e-02	0.00000e+00	2.00000e+00	0.0000e+00	9.9643e-02
V 6	LL	0.00000e+00	0.00000e+00	2.00000e+00	1.4728e-01	0.0000e+00
V 7	FR	4.90578e-02	0.00000e+00	2.00000e+00	0.0000e+00	4.9058e-02
V 8	LL	0.00000e+00	0.00000e+00	2.00000e+00	8.6026e-01	0.0000e+00
V 9	FR	3.05649e-01	0.00000e+00	2.00000e+00	0.0000e+00	3.0565e-01

L Con	State	Value	Lower Bound	Upper Bound	Lagr Mult	Residual
L 1	LL	2.00000e+00	2.00000e+00	None	3.7775e-01	0.0000e+00
L 2	UL	2.00000e+00	None	2.00000e+00	-5.7914e-02	4.4409e-16
L 3	LL	1.00000e+00	1.00000e+00	4.00000e+00	1.0753e-01	4.4409e-16

Exit after 15 iterations.

Optimal LS solution found.

Final LS objective value = 8.1340823e-02

7. Further Description

This section gives a detailed description of the algorithm used in nag_opt_lin_lsq. This, and possibly the next section, Section 8, may be omitted if the more sophisticated features of the algorithm and software are not currently of interest.

7.1. Overview

nag_opt_lin_lsq is based on an inertia-controlling method that maintains a Cholesky factorization of the reduced Hessian (see below). The method is based on that of Gill and Murray (1978) and is described in detail by Gill *et al* (1981). Here we briefly summarize the main features of the method.

nag_opt_lin_lsq uses essentially the same algorithm as the subroutine LSSOL described in Gill *et al* (1986). It is based on a two-phase (primal) quadratic programming method with features to exploit the convexity of the objective function due to Gill *et al* (1984). (In the full-rank case, the method is related to that of Stoer, see Stoer (1971).) nag_opt_lin_lsq has two phases: finding an initial feasible point by minimizing the sum of infeasibilities (the *feasibility phase*), and minimizing the quadratic objective function within the feasible region (the *optimality phase*). The two-phase nature of the algorithm is reflected by changing the function being minimized from the sum of infeasibilities to the quadratic objective function. The feasibility phase does *not* perform the standard simplex method (i.e., it does not necessarily find a vertex), except in the LP case when $n_L \leq n$. Once any iterate is feasible, all subsequent iterates remain feasible.

nag_opt_lin_lsq has been designed to be efficient when used to solve a *sequence* of related problems — for example, within a sequential quadratic programming method for nonlinearly constrained optimization (e.g., nag_opt_nlp (e04ucc)). In particular, the user may specify an initial working set (the indices of the constraints believed to be satisfied exactly at the solution); see the discussion of the optional parameter **start** in Section 8.2.

In general, an iterative process is required to solve a quadratic program. (For simplicity, we shall always consider a typical iteration and avoid reference to the index of the iteration.) Each new iterate \bar{x} is defined by

$$\bar{x} = x + \alpha p, \quad (2)$$

where the *step length* α is a non-negative scalar, and p is called the *search direction*.

At each point x , a *working set* of constraints is defined to be a linearly independent subset of the constraints that are satisfied ‘exactly’ (to within the tolerance defined by the optional parameter **ftol**; see Section 8.2). The working set is the current prediction of the constraints that hold with equality at a solution of (1). The search direction is constructed so that the constraints in the working set remain *unaltered* for any value of the step length. For a bound constraint in the working set, this property is achieved by setting the corresponding element of the search direction to zero. Thus, the associated variable is *fixed*, and specification of the working set induces a partition of x into *fixed* and *free* variables. During a given iteration, the fixed variables are effectively removed from the problem; since the relevant elements of the search direction are zero, the columns of A corresponding to fixed variables may be ignored.

Let n_W denote the number of general constraints in the working set and let n_{FX} denote the number of variables fixed at one of their bounds (n_W and n_{FX} are the quantities **Lin** and **Bnd** in the extended iteration printout from nag_opt_lin_lsq; see Section 8.3). Similarly, let n_{FR} ($n_{FR} = n - n_{FX}$) denote the number of free variables. At every iteration, *the variables are re-ordered so that the last n_{FX} variables are fixed*, with all other relevant vectors and matrices ordered accordingly. The order of the variables is indicated by the contents of the array **kx** on exit (see Section 4).

7.2 Definition of the Search Direction

Let A_{FR} denote the n_W by n_{FR} sub-matrix of general constraints in the working set corresponding to the free variables, and let p_{FR} denote the search direction with respect to the free variables only.

The general constraints in the working set will be unaltered by any move along p if

$$A_{\text{FR}} p_{\text{FR}} = 0. \quad (3)$$

In order to compute p_{FR} , the TQ factorization of A_{FR} is used:

$$A_{\text{FR}} Q_{\text{FR}} = (0 \ T) \quad (4)$$

where T is a non-singular n_{W} by n_{W} reverse-triangular matrix (i.e., $t_{ij} = 0$ if $i + j < n_{\text{W}}$), and the non-singular n_{FR} by n_{FR} matrix Q_{FR} is the product of orthogonal transformations (see Gill *et al* (1984)). If the columns of Q_{FR} are partitioned so that

$$Q_{\text{FR}} = (Z \ Y), \quad (5)$$

where Y is n_{FR} by n_{W} , then the n_Z ($n_Z = n_{\text{FR}} - n_{\text{W}}$) columns of Z form a basis for the null space of A_{FR} . Let n_R be an integer such that $0 \leq n_R \leq n_Z$, and let Z_1 denote a matrix whose n_R columns are a subset of the columns of Z . (The integer n_R is the quantity **Zr** in the extended iteration printout from `nag_opt_lin_lsq`; see Section 8.3. In many cases, Z_1 will include *all* the columns of Z .) The direction p_{FR} will satisfy (3) if

$$p_{\text{FR}} = Z_1 p_Z \quad (6)$$

where p_Z is any n_R -vector.

7.3 The Main Iteration

Let Q denote the n by n matrix

$$Q = \begin{pmatrix} Q_{\text{FR}} & \\ & I_{\text{FX}} \end{pmatrix} \quad (7)$$

where I_{FX} is the identity matrix of order n_{FX} . Let R denote an n by n upper triangular matrix (the *Cholesky factor*) such that

$$Q^T \widetilde{\nabla^2 F} Q \equiv H_Q = R^T R, \quad (8)$$

and let the matrix of the first n_Z rows and columns of R be denoted by R_Z . (The matrix $\widetilde{\nabla^2 F}$ in (8) is the Hessian with its rows and columns permuted so that the free variables come first.)

The definition of p_Z in (6) depends on whether or not the matrix R_Z is singular at x . In the non-singular case, p_Z satisfies the equations

$$R_Z^T R_Z p_Z = -g_Z \quad (9)$$

where g_Z denotes the vector $Z^T g_{\text{FR}}$ and g denotes the objective gradient. (The norm of g_{FR} is the printed quantity **Norm Gf**; see Section 8.3.) When p_Z is defined by (9), $x + p$ is the minimizer of the objective function subject to the constraints (bounds and general) in the working set treated as equalities. In general, a vector f_Z is available such that $R_Z^T f_Z = -g_Z$, which allows p_Z to be computed from a single back-substitution $R_Z p_Z = f_Z$. For example, when solving problem LS1, f_Z comprises the first n_Z elements of the *transformed residual vector*

$$f = P(b - Hx) \quad (10)$$

which is recurred from one iteration to the next, where P is an orthogonal matrix.

In the singular case, p_Z is defined such that

$$R_Z p_Z = 0 \quad \text{and} \quad g_Z^T p_Z < 0. \quad (11)$$

This vector has the property that the objective function is linear along p and may be reduced by any step of the form $x + \alpha p$, where $\alpha > 0$.

The vector $Z^T g_{\text{FR}}$ is known as the *projected gradient* at x . If the projected gradient is zero, x is a constrained stationary point in the subspace defined by Z . During the feasibility phase, the projected gradient will usually be zero only at a vertex (although it may be zero at non-vertices in the presence of constraint dependencies). During the optimality phase, a zero projected gradient implies that x minimizes the quadratic objective when the constraints in the working set are treated as equalities. At a constrained stationary point, Lagrange multipliers λ_A and λ_B for the general and bound constraints are defined from the equations

$$A_{\text{FR}}^T \lambda_A = g_{\text{FR}} \quad \text{and} \quad \lambda_B = g_{\text{FX}} - A_{\text{FX}}^T \lambda_A. \quad (12)$$

Given a positive constant δ of the order of the **machine precision**, the Lagrange multiplier λ_j corresponding to an inequality constraint in the working set is said to be *optimal* if $\lambda_j \leq \delta$ when the associated constraint is at its *upper bound*, or if $\lambda_j \geq -\delta$ when the associated constraint is at its *lower bound*. If a multiplier is non-optimal, the objective function (either the true objective or the sum of infeasibilities) can be reduced by deleting the corresponding constraint (with index **Jdel**; see Section 8.3) from the working set.

If optimal multipliers occur during the feasibility phase and the sum of infeasibilities is non-zero, there is no feasible point, and nag_opt_lin_lsq will continue until the minimum value of the sum of infeasibilities has been found. At this point, the Lagrange multiplier λ_j corresponding to an inequality constraint in the working set will be such that $-(1 + \delta) \leq \lambda_j \leq \delta$ when the associated constraint is at its *upper bound*, and $-\delta \leq \lambda_j \leq (1 + \delta)$ when the associated constraint is at its *lower bound*. Lagrange multipliers for equality constraints will satisfy $|\lambda_j| \leq 1 + \delta$.

The choice of step length is based on remaining feasible with respect to the satisfied constraints. If R_Z is non-singular and $x + p$ is feasible, α will be taken as unity. In this case, the projected gradient at \bar{x} will be zero, and Lagrange multipliers are computed. Otherwise, α is set to α_M , the step to the ‘nearest’ constraint (with index **Jadd**; see Section 8.3), which is added to the working set at the next iteration.

If H is not input as a triangular matrix, it is overwritten by a triangular matrix R satisfying (8) obtained using the Cholesky factorization in the QP case, or the QR factorization in the LS case. Column interchanges are used in both cases, and an estimate is made of the rank of the triangular factor. Thereafter, the dependent rows of R are eliminated from the problem.

Each change in the working set leads to a simple change to A_{FR} : if the status of a general constraint changes, a *row* of A_{FR} is altered; if a bound constraint enters or leaves the working set, a *column* of A_{FR} changes. Explicit representations are recurred of the matrices T , Q_{FR} and R ; and of vectors $Q^T g$, $Q^T c$ and f , which are related by the formulae

$$f = Pb - \begin{pmatrix} R \\ 0 \end{pmatrix} Q^T x, \quad (b \equiv 0 \text{ for the QP case}),$$

and

$$Q^T g = Q^T c - R^T f.$$

Note that the triangular factor R associated with the Hessian of the original problem is updated during both the optimality *and* the feasibility phases.

The treatment of the singular case depends critically on the following feature of the matrix updating schemes used in nag_opt_lin_lsq: if a given factor R_Z is non-singular, it can become singular during subsequent iterations only when a constraint leaves the working set, in which case only its last diagonal element can become zero. This property implies that a vector satisfying (11) may be found using the single back-substitution $\bar{R}_Z p_Z = e_Z$, where \bar{R}_Z is the matrix R_Z with a unit last diagonal, and e_Z is a vector of all zeros except in the last position. If the Hessian matrix $\nabla^2 F$ is singular, the matrix R (and hence R_Z) may be singular at the start of the optimality phase. However, R_Z will be non-singular if enough constraints are included in the initial working set. (The matrix with no rows and columns is positive-definite by definition, corresponding to the case when A_{FR} contains n_{FR} constraints.) The idea is to include as many general constraints as necessary to ensure a non-singular R_Z .

At the beginning of each phase, an upper triangular matrix R_1 is determined that is the largest non-singular leading sub-matrix of R_Z . The use of interchanges during the factorization of H tends to maximize the dimension of R_1 . (The rank of R_1 is estimated using the optional parameter **rank_tol**; see Section 8.2.) Let Z_1 denote the columns of Z corresponding to R_1 , and let Z be partitioned as $Z = (Z_1 \ Z_2)$. A working set for which Z_1 defines the null space can be obtained by including the rows of Z_2^T as ‘artificial constraints’. Minimization of the objective function then proceeds within the subspace defined by Z_1 .

The artificially augmented working set is given by

$$\bar{A}_{\text{FR}} = \begin{pmatrix} A_{\text{FR}} \\ Z_2^T \end{pmatrix}, \tag{13}$$

so that p_{FR} will satisfy $A_{\text{FR}}p_{\text{FR}} = 0$ and $Z_2^T p_{\text{FR}} = 0$. By definition of the TQ factorization, \bar{A}_{FR} automatically satisfies the following:

$$\bar{A}_{\text{FR}} Q_{\text{FR}} = \begin{pmatrix} A_{\text{FR}} \\ Z_2^T \end{pmatrix} Q_{\text{FR}} = \begin{pmatrix} A_{\text{FR}} \\ Z_2^T \end{pmatrix} (Z_1 \ Z_2 \ Y) = (0 \ \bar{T}),$$

where

$$\bar{T} = \begin{pmatrix} 0 & T \\ I & 0 \end{pmatrix},$$

and hence the TQ factorization of (13) requires no additional work.

The matrix Z_2 need not be kept fixed, since its role is purely to define an appropriate null space; the TQ factorization can therefore be updated in the normal fashion as the iterations proceed. No work is required to ‘delete’ the artificial constraints associated with Z_2 when $Z_1^T g_{\text{FR}} = 0$, since this simply involves repartitioning Q_{FR} . When deciding which constraint to delete, the ‘artificial’ multiplier vector associated with the rows of Z_2^T is equal to $Z_2^T g_{\text{FR}}$, and the multipliers corresponding to the rows of the ‘true’ working set are the multipliers that would be obtained if the temporary constraints were not present.

The number of columns in Z_2 and Z_1 , the Euclidean norm of $Z_1^T g_{\text{FR}}$, and the condition estimator of R_1 appear in the extended iteration printout as **Art**, **Zr**, **Norm Gz** and **Cond Rz** respectively (see Section 8.3).

Although the algorithm of `nag_opt_lin_lsq` does not perform simplex steps in general, there is one exception: a linear program with fewer general constraints than variables (i.e., $n_L \leq n$). (Use of the simplex method in this situation leads to savings in storage.) At the starting point, the ‘natural’ working set (the set of constraints exactly or nearly satisfied at the starting point) is augmented with a suitable number of ‘temporary’ bounds, each of which has the effect of temporarily fixing a variable at its current value. In subsequent iterations, a temporary bound is treated as a standard constraint until it is deleted from the working set, in which case it is never added again.

One of the most important features of `nag_opt_lin_lsq` is its control of the conditioning of the working set, whose nearness to linear dependence is estimated by the ratio of the largest to smallest diagonals of the TQ factor T (the printed value **Cond T**; see Section 8.3). In constructing the initial working set, constraints are excluded that would result in a large value of **Cond T**. Thereafter, `nag_opt_lin_lsq` allows constraints to be violated by as much as a user-specified feasibility tolerance (see **ftol**, Section 8.2) in order to provide, whenever possible, a *choice* of constraints to be added to the working set at a given iteration. Let α_M denote the maximum step at which $x + \alpha_M p$ does not violate any constraint by more than its feasibility tolerance. All constraints at distance α ($\alpha \leq \alpha_M$) along p from the current point are then viewed as acceptable candidates for inclusion in the working set. The constraint whose normal makes the largest angle with the search direction is added to the working set. In order to ensure that the new iterate satisfies the constraints in the working set as accurately as possible, the step taken is the exact distance to the newly added constraint. As a consequence, negative steps are occasionally permitted, since the current iterate may violate the constraint to be added by as much as the feasibility tolerance.

8. Optional Parameters

A number of optional input and output parameters to `nag_opt_lin_lsq` are available through the structure argument **options**, type `Nag_E04_Opt`. A parameter may be selected by assigning an appropriate value to the relevant structure member; those parameters not selected will be assigned default values. If no use is to be made of any of the optional parameters the user should use the NAG defined null pointer, `E04_DEFAULT`, in place of **options** when calling `nag_opt_lin_lsq`; the default settings will then be used for all parameters.

Before assigning values to **options** directly the structure **must** be initialized by a call to the function `nag_opt_init` (`e04xxc`). Values may then be assigned to the structure members in the normal C manner.

Option settings may also be read from a text file using the function `nag_opt_read` (`e04xyc`) in which case initialization of the **options** structure will be performed automatically if not already done. Any subsequent direct assignment to the **options** structure **must not** be preceded by initialization.

If assignment of functions and memory to pointers in the **options** structure is required, then this must be done directly in the calling program; they cannot be assigned using `nag_opt_read` (e04xyc).

8.1. Optional Parameter Checklist and Default Values

For easy reference, the following list shows the members of **options** which are valid for `nag_opt_lin_lsq` together with their default values where relevant. The number ϵ is a generic notation for *machine precision* (see `nag_machine_precision` (X02AJC)).

<code>Nag_ProblemType prob</code>	Nag_LS1
<code>Nag_Start start</code>	Nag_Cold
<code>Boolean list</code>	TRUE
<code>Nag_PrintType print_level</code>	Nag_Soln_Iter
<code>char outfile[80]</code>	<code>stdout</code>
<code>void (*print_fun)()</code>	<code>NULL</code>
<code>Integer fmax_iter</code>	<code>max(50, 5(n+nclin))</code>
<code>Integer max_iter</code>	<code>max(50, 5(n+nclin))</code>
<code>double crash_tol</code>	<code>0.01</code>
<code>double ftol</code>	$\sqrt{\epsilon}$
<code>double inf_bound</code>	10^{20}
<code>double inf_step</code>	<code>max(inf_bound, 10²⁰)</code>
<code>double rank_tol</code>	<code>100ϵ or 10$\sqrt{\epsilon}$</code>
<code>Integer *state</code>	<code>size n+nclin</code>
<code>double *ax</code>	<code>size nclin</code>
<code>double *lambda</code>	<code>size n+nclin</code>
<code>Boolean hessian</code>	FALSE
<code>Integer iter</code>	

8.2. Description of Optional Parameters

prob – `Nag_ProblemType` Default = **Nag_LS1**

Input: specifies the type of objective function to be minimized during the optimality phase. The following are the ten possible values of **prob** and the size of the arrays **h**, **kx**, **b** and **cvec** that are required to define the objective function:

Nag_FP	h , b and cvec not referenced;
Nag_LP	h and b not referenced, cvec[n] ;
Nag_QP1	h[m][tdh] symmetric, b and cvec not referenced;
Nag_QP2	h[m][tdh] symmetric, b not referenced, cvec[n] ;
Nag_QP3	h[m][tdh] upper trapezoidal, b and cvec not referenced;
Nag_QP4	h[m][tdh] upper trapezoidal, b not referenced, cvec[n] .
Nag_LS1	h[m][tdh] , b[m] , cvec not referenced;
Nag_LS2	h[m][tdh] , b[m] , cvec[n] ;
Nag_LS3	h[m][tdh] upper trapezoidal, b[m] , cvec not referenced;
Nag_LS4	h[m][tdh] upper trapezoidal, b[m] , cvec[n] .

The array **kx[n]** must be supplied for all problem types but need only be initialized for types **Nag_QP3**, **Nag_QP4**, **Nag_LS3** and **Nag_LS4**. If $H = 0$, i.e., the objective function is purely linear, the efficiency of `nag_opt_lin_lsq` may be increased by specifying **options.prob** = **Nag_LP**. Constraint: **options.prob** = **Nag_FP**, **Nag_LP**, **Nag_QP1**, **Nag_QP2**, **Nag_QP3**, **Nag_QP4**, **Nag_LS1**, **Nag_LS2**, **Nag_LS3** or **Nag_LS4**.

start – Nag_Start Default = **Nag_Cold**

Input: specifies how the initial working set is chosen. With **start** = **Nag_Cold**, nag_opt_lin_lsq chooses the initial working set based on the values of the variables and constraints at the initial point. Broadly speaking, the initial working set will include equality constraints and bounds or inequality constraints that violate or ‘nearly’ satisfy their bounds (to within the value of the optional parameter **crash_tol**; see below).

With **start** = **Nag_Warm**, the user must provide a valid definition of every array element of the optional parameter **state** (see below). nag_opt_lin_lsq will override the user’s specification of **state** if necessary, so that a poor choice of the working set will not cause a fatal error. For instance, any elements of **state** which are set to -2 , -1 or 4 will be reset to zero, as will any elements which are set to 3 when the corresponding elements of **bl** and **bu** are not equal. A warm start will be advantageous if a good estimate of the initial working set is available – for example, when nag_opt_lin_lsq is called repeatedly to solve related problems.

Constraint: **options.start** = **Nag_Cold** or **Nag_Warm**.

list – Boolean Default = **TRUE**

Input: if **options.list** = **TRUE** the parameter settings in the call to nag_opt_lin_lsq will be printed.

print_level – Nag_PrintType Default = **Nag_Soln_Iter**

Input: the level of results printout produced by nag_opt_lin_lsq. The following values are available.

Nag_NoPrint	No output.
Nag_Soln	The final solution.
Nag_Iter	One line of output for each iteration.
Nag_Iter_Long	A longer line of output for each iteration with more information (line exceeds 80 characters).
Nag_Soln_Iter	The final solution and one line of output for each iteration.
Nag_Soln_Iter_Long	The final solution and one long line of output for each iteration (line exceeds 80 characters).
Nag_Soln_Iter_Const	As Nag_Soln_Iter_Long with the Lagrange multipliers, the variables x , the constraint values Ax and the constraint status also printed at each iteration.
Nag_Soln_Iter_Full	As Nag_Soln_Iter_Const with the diagonal elements of the matrix T associated with the TQ factorization (see (4) in Section 7.2) of the working set, and the diagonal elements of the upper triangular matrix R printed at each iteration.

Details of each level of results printout are described in Section 8.3.

Constraint: **options.print_level** = **Nag_NoPrint**, **Nag_Soln**, **Nag_Iter**, **Nag_Soln_Iter**, **Nag_Iter_Long**, **Nag_Soln_Iter_Long**, **Nag_Soln_Iter_Const** or **Nag_Soln_Iter_Full**.

outfile – char[80] Default = **stdout**

Input: the name of the file to which results should be printed. If **options.outfile**[0] = ‘\0’ then the **stdout** stream is used.

print_fun – pointer to function Default = **NULL**

Input: printing function defined by the user; the prototype of **print_fun** is

```
void (*print_fun)(const Nag_Search_State *st, Nag_Comm *comm);
```

See Section 8.3.1 below for further details.

fmax_iter – Integer Default = $\max(50, 5(\mathbf{n} + \mathbf{nclin}))$
max_iter – Integer Default = $\max(50, 5(\mathbf{n} + \mathbf{nclin}))$

Input: **fmax_iter** and **max_iter** specify the maximum number of iterations allowed in the feasibility and optimality phase, respectively.

If the user wishes to check that a call to nag_opt_lin_lsq is correct before attempting to solve the problem in full then **fmax_iter** may be set to 0. No iterations will then be performed but all initialization prior to the first iteration will be done and a listing of parameter settings will be output, if optional parameter **list** = **TRUE** (the default setting).

Constraints:

$$\mathbf{options.fmax_iter} \geq 0,$$

$$\mathbf{options.max_iter} \geq 0.$$

crash_tol – double Default = 0.01

Input: **crash_tol** is used when optional parameter **start** = **Nag_Cold** (the default) and nag_opt_lin_lsq selects an initial working set. The initial working set will include (if possible) bounds or general inequality constraints that lie within **crash_tol** of their bounds. In particular, a constraint of the form $a_j^T x \geq l$ will be included in the initial working set if $|a_j^T x - l| \leq \mathbf{crash_tol} \times (1 + |l|)$.

Constraint: $0.0 \leq \mathbf{options.crash_tol} \leq 1.0$.

ftol – double Default = $\sqrt{\epsilon}$

Input: defines the maximum acceptable *absolute* violation in each constraint at a ‘feasible’ point. For example, if the variables and the coefficients in the general constraints are of order unity, and the latter are correct to about 6 decimal digits, it would be appropriate to specify **ftol** as 10^{-6} .

nag_opt_lin_lsq attempts to find a feasible solution before optimizing the objective function. If the sum of infeasibilities cannot be reduced to zero, nag_opt_lin_lsq finds the minimum value of the sum. Let **Sinf** be the corresponding sum of infeasibilities. If **Sinf** is quite small, it may be appropriate to raise **ftol** by a factor of 10 or 100. Otherwise, some error in the data should be suspected.

Note that a ‘feasible solution’ is a solution that satisfies the current constraints to within the feasibility tolerance **ftol**.

Constraint: **options.ftol** > 0.0.

inf_bound – double Default = 10^{20}

Input: **inf_bound** defines the ‘infinite’ bound in the definition of the problem constraints. Any upper bound greater than or equal to **inf_bound** will be regarded as plus infinity (and similarly any lower bound less than or equal to $-\mathbf{inf_bound}$ will be regarded as minus infinity).

Constraint: **options.inf_bound** > 0.0.

inf_step – double Default = $\max(\mathbf{inf_bound}, 10^{20})$

Input: specifies the magnitude of the change in variables that will be considered a step to an unbounded solution. (Note that an unbounded solution can occur only when the Hessian is singular and the objective contains an explicit linear term.) If the change in x during an iteration would exceed the value of **inf_step**, the objective function is considered to be unbounded below in the feasible region.

Constraint: **options.inf_step** > 0.0.

rank_tol – double Default = 100ϵ or $10\sqrt{\epsilon}$

The default value is 100ϵ for problem types QP1, LS1 and LS3 but is $10\sqrt{\epsilon}$ for other QP and LS problem types. This option does not apply to FP or LP problem types.

Input: **rank_tol** enables the user to control the estimate of the triangular factor R_1 (see Section 7.3). If ρ_i denotes the function $\rho_i = \max\{|R_{11}|, |R_{22}|, \dots, |R_{ii}|\}$, the rank of R is defined to be smallest index i such that $|R_{i+1, i+1}| \leq \mathbf{rank_tol} \times |\rho_{i+1}|$.

Constraint: $0.0 < \mathbf{options.rank_tol} < 1.0$.

state – Integer * Default memory = **n+nclin**

Input: **state** need not be set if the default option of **options.start** = **Nag_Cold** is used as **n+nclin** values of memory will be automatically allocated by nag_opt_lin_lsqr.

If the option **start** = **Nag_Warm** has been chosen, **state** must point to a minimum of **n+nclin** elements of memory. This memory will already be available if the **options** structure has been used in a previous call to nag_opt_lin_lsqr from the calling program, with **start** = **Nag_Cold** and the same values of **n** and **nclin**. If a previous call has not been made sufficient memory must be allocated to **state** by the user.

When a warm start is chosen **state** should specify the status of the constraints at the start of the feasibility phase. More precisely, the first n elements of **state** refer to the upper and lower bounds on the variables, and the next n_L elements refer to the general linear constraints (if any). Possible values for **state**[j] are as follows:

state [j]	Meaning
0	The constraint should <i>not</i> be in the initial working set.
1	The constraint should be in the initial working set at its lower bound.
2	The constraint should be in the initial working set at its upper bound.
3	The constraint should be in the initial working set as an equality. This value should only be specified if bl [j] = bu [j].

The values -2 , -1 and 4 are also acceptable but will be reset to zero by the function, as will any elements which are set to 3 when the corresponding elements of **bu** and **bl** are not equal. If nag_opt_lin_lsqr has been called previously with the same values of **n** and **nclin**, **state** already contains satisfactory information. (See also the description of the optional parameter **start**). The function also adjusts (if necessary) the values supplied in **x** to be consistent with the values supplied in **state**.

Constraint: $-2 \leq \mathbf{state}[j-1] \leq 4$, for $j = 1, 2, \dots, \mathbf{n+nclin}-1$.

Output: the status of the constraints in the working set at the point returned in **x**. The significance of each possible value of **state**[j] is as follows:

state [j]	Meaning
-2	The constraint violates its lower bound by more than the feasibility tolerance.
-1	The constraint violates its upper bound by more than the feasibility tolerance.
0	The constraint is satisfied to within the feasibility tolerance, but is not in the working set.
1	This inequality constraint is included in the working set at its lower bound.
2	This inequality constraint is included in the working set at its upper bound.
3	This constraint is included in the working set as an equality. This value of state can occur only when bl [j] = bu [j].
4	This corresponds to optimality being declared with x [j] being temporarily fixed at its current value. This value of state can only occur when fail.code = NW_SOLN_NOT_UNIQUE .

ax – double * Default memory = **nclin**

Input: **nclin** values of memory will be automatically allocated by nag_opt_lin_lsqr and this is the recommended method of use of **options.ax**. However a user may supply memory from the calling program.

Output: If **nclin** > 0, **ax** points to the final values of the linear constraints Ax .

lambda – double *Default memory = **n+nclin**

Input: **n+nclin** values of memory will be automatically allocated by nag_opt_lin_lsq and this is the recommended method of use of **options.lambda**. However a user may supply memory from the calling program.

Output: the values of the Lagrange multipliers for each constraint with respect to the current working set. The first n elements contain the multipliers for the bound constraints on the variables, and the next n_L elements contain the multipliers for the general linear constraints (if any). If **state**[$j - 1$] = 0 (i.e., constraint j is not in the working set), **lambda**[$j - 1$] is zero. If x is optimal, **lambda**[$j - 1$] should be non-negative if **state**[$j - 1$] = 1, non-positive if **state**[$j - 1$] = 2 and zero if **state**[$j - 1$] = 4.

hessian – BooleanDefault = **FALSE**

Input: controls the contents of the parameter **h** on return from nag_opt_lin_lsq. nag_opt_lin_lsq works exclusively with the transformed and reordered matrix H_Q (8), and hence extra computation is required to form the Hessian itself. If the optional parameter **hessian** = **FALSE**, **h** contains the Cholesky factor of the matrix H_Q with columns ordered as indicated by **kx** (see Section 4). If **hessian** = **TRUE**, **h** contains the Cholesky factor of the Hessian matrix $\nabla^2 F$, with columns ordered as indicated by **kx**.

iter – Integer

Output: the total number of iterations performed in the feasibility phase and (if appropriate) the optimality phase.

8.3. Description of Printed Output

The level of printed output can be controlled by the user with the structure members **options.list** and **options.print_level** (see Section 8.2). If **list** = **TRUE** then the parameter values to nag_opt_lin_lsq are listed, whereas the printout of results is governed by the value of **print_level**. The default of **print_level** = **Nag_Soln_Iter** provides a single line of output at each iteration and the final result. This section describes all of the possible levels of results printout available from nag_opt_lin_lsq.

To aid interpretation of the printed results, the following convention is used for numbering the constraints: indices 1 to n refer to the bounds on the variables, and indices $n + 1$ to $n + n_L$ refer to the general constraints.

When **print_level** = **Nag_Iter** or **Nag_Soln_Iter** the following line of output is produced at every iteration. In all cases, the values of the quantities printed are those in effect *on completion* of the given iteration.

Itn	is the iteration count.
Step	is the step taken along the computed search direction. If a constraint is added during the current iteration, Step will be the step to the nearest constraint. During the optimality phase, the step can be greater than 1.0 only if the factor R_Z is singular (see Section 7.3).
Ninf	is the number of violated constraints (infeasibilities). This will be zero during the optimality phase.
Sinf/Objective	is the value of the current objective function. If x is not feasible, Sinf gives a weighted sum of the magnitudes of constraint violations. If x is feasible, Objective is the value of the objective function. The output line for the final iteration of the feasibility phase (i.e., the first iteration for which Ninf is zero) will give the value of the true objective at the first feasible point.

During the optimality phase, the value of the objective function will be non-increasing. During the feasibility phase, the number of constraint infeasibilities will not increase until either a feasible point is found, or the optimality of the multipliers implies that no feasible point exists. Once optimal multipliers are obtained, the number of infeasibilities can increase, but the sum of infeasibilities will either remain constant or be reduced until the minimum sum of infeasibilities is found.

Norm Gz $\|Z_1^T g_{FR}\|$, the Euclidean norm of the reduced gradient with respect to Z_1 (see Section 7.3). During the optimality phase, this norm will be approximately zero after a unit step.

If **print_level** = **Nag_Iter_Long**, **Nag_Soln_Iter_Long**, **Nag_Soln_Iter_Const** or **Nag_Soln_Iter_Full** the line of printout is extended to give the following additional information. (Note that this longer line extends over more than 80 characters.)

Jdel is the index of the constraint deleted from the working set, along with the designation **L** (lower bound), **U** (upper bound), **E** (equality), **F** (temporarily fixed variable) or **A** (artificial constraint). If **Jdel** is zero, no constraint was deleted.

Jadd is the index of the constraint added to the working set, along with a designation as for **Jdel**. If **Jadd** is zero, no constraint was added.

Bnd is the number of simple bound constraints in the current working set.

Lin is the number of general linear constraints in the current working set.

Art is the number of artificial constraints in the working set, i.e., the number of columns of Z_2 (see Section 7.3).

Zr is the number of columns of Z_1 (see Section 7.2). **Zr** is the dimension of the subspace in which the objective function is currently being minimized. The value of **Zr** is the number of variables minus the number of constraints in the working set; i.e., $\mathbf{Zr} = n - (\mathbf{Bnd} + \mathbf{Lin} + \mathbf{Art})$.

The value of n_Z , the number of columns of Z (see Section 7) can be calculated as $n_Z = n - (\mathbf{Bnd} + \mathbf{Lin})$. A zero value of n_Z implies that x lies at a vertex of the feasible region.

Norm Gf is the Euclidean norm of the gradient function with respect to the free variables, i.e., variables not currently held at a bound.

Cond T is a lower bound on the condition number of the working set.

Cond Rz is a lower bound on the condition number of the triangular factor R_1 (the first **Zr** rows and columns of the factor R_Z).

When **print_level** = **Nag_Soln_Iter_Const** or **Nag_Soln_Iter_Full** more detailed results are given at each iteration. For the setting **Nag_Soln_Iter_Const** additional values output are:

Value of x is the value of x currently held in **x**.

State is the current value of **options.state** associated with x .

Value of Ax is the value of Ax currently held in **options.ax**.

State is the current value of **options.state** associated with Ax .

Also printed are the Lagrange Multipliers for the bound constraints, linear constraints and artificial constraints.

If **print_level** = **Nag_Soln_Iter_Full** then the diagonals of T and R are also output at each iteration.

When **print_level** = **Nag_Soln**, **Nag_Soln_Iter**, **Nag_Soln_Iter_Long**, **Nag_Soln_Iter_Const** or **Nag_Soln_Iter_Full** the final printout from **nag_opt_lin_lsq** includes a listing of the status of every variable and constraint. The following describes the printout for each variable.

Varbl gives the name (**V**) and index j , for $j = 1, 2, \dots, n$ of the variable.

State gives the state of the variable (**FR** if neither bound is in the working set, **EQ** if a fixed variable, **LL** if on its lower bound, **UL** if on its upper bound, **TF** if temporarily fixed at its current value). If **Value** lies outside the upper or lower bounds by more than the optional parameter **ftol** (default value $\sqrt{\epsilon}$, where ϵ is the **machine precision**; see Section 8.2), **State** will be **++** or **--** respectively.

A key is sometimes printed before **State** to give some additional information about the state of a variable.

- A *Alternative optimum possible.* The variable is active at one of its bounds, but its Lagrange Multiplier is essentially zero. This means that if the variable were allowed to start moving away from its bound, there would be no change to the objective function. The values of the other free variables *might* change, giving a genuine alternative solution. However, if there are any degenerate variables (labelled D), the actual change might prove to be zero, since one of them could encounter a bound immediately. In either case, the values of the Lagrange multipliers might also change.
- D *Degenerate.* The variable is free, but it is equal to (or very close to) one of its bounds.
- I *Infeasible.* The variable is currently violating one of its bounds by more than **ftol**.

Value	is the value of the variable at the final iteration.
Lower bound	is the lower bound specified for variable j . (None indicates that $\mathbf{bl}[j-1] \leq -\mathbf{inf_bound}$, where inf_bound is the optional parameter.)
Upper bound	is the upper bound specified for variable j . (None indicates that $\mathbf{bu}[j-1] \geq \mathbf{inf_bound}$, where inf_bound is the optional parameter.)
Lagr mult	is the value of the Lagrange multiplier for the associated bound. This will be zero if State is FR unless $\mathbf{bl}[j-1] \leq -\mathbf{inf_bound}$ and $\mathbf{bu}[j-1] \geq \mathbf{inf_bound}$, in which case the entry will be blank. If x is optimal, the multiplier should be non-negative if State is LL, and non-positive if State is UL.
Residual	is the difference between the variable Value and the nearer of its (finite) bounds $\mathbf{bl}[j-1]$ and $\mathbf{bu}[j-1]$. A blank entry indicates that the associated variable is not bounded (i.e., $\mathbf{bl}[j-1] \leq -\mathbf{inf_bound}$ and $\mathbf{bu}[j-1] \geq \mathbf{inf_bound}$).

The meaning of the printout for general constraints is the same as that given above for variables, with 'variable' replaced by 'constraint', $\mathbf{bl}[j-1]$ and $\mathbf{bu}[j-1]$ replaced by $\mathbf{bl}[n+j-1]$ and $\mathbf{bu}[n+j-1]$ respectively, and with the following change in the heading:

L Con the name (L) and index j , for $j = 1, 2, \dots, n_L$ of the linear constraint.

Note that movement off a constraint (as opposed to a variable moving away from its bound) can be interpreted as allowing the entry in the **Residual** column to become positive.

Numerical values are output with a fixed number of digits; they are not guaranteed to be accurate to this precision.

If **options.print_level = Nag_NoPrint** then printout will be suppressed; the user can print the final solution when nag_opt_lin_lsq returns to the calling program.

8.3.1. Output of Results Via a User-defined Printing Function

The user may also specify their own print function for output of iteration results and the final solution by use of the **options.print_fun** function pointer, which has prototype

```
void (*print_fun)(const Nag_Search_State *st, Nag_Comm *comm);
```

The rest of this section can be skipped by a user who only wishes to use the default printing facilities.

When a user-defined function is assigned to **options.print_fun** this will be called in preference to the internal print function of nag_opt_lin_lsq. Calls to the user-defined function are again controlled by means of the **options.print_level** member. Information is provided through **st** and **comm**, the two structure arguments to **print_fun**.

If **comm->it_prt = TRUE** then the results from the last iteration of nag_opt_lin_lsq are provided through **st**. Note that **print_fun** will be called with **comm->it_prt = TRUE** only if **print_level = Nag_Iter, Nag_Iter_Long, Nag_Soln_Iter, Nag_Soln_Iter_Long, Nag_Soln_Iter_Const** or **Nag_Soln_Iter_Full**. The following members of **st** are set:

n – Integer
the number of variables.

nclin – Integer
the number of linear constraints.

iter – Integer
the iteration count.

jdel – Integer
index of constraint deleted from the working set.

jadd – Integer
index of constraint added to the working set.

step – double
the step taken along the computed search direction.

ninf – Integer
the number of violated constraints (infeasibilities).

f – double
the current value of the objective function if **st->ninf** = 0; otherwise, **f** is a weighted sum of the magnitudes of constraint violations.

bnd – Integer
number of bound constraints in the working set.

lin – Integer
number of general linear constraints in the working set.

nart – Integer
number of artificial constraints in the working set (see Section 7.3).

nrnk – Integer
the rank of the upper triangular matrix R (see Section 7.3).

nrz – Integer
number of columns of Z_1 (see Section 7.2).

norm_gz – double
Euclidean norm of the reduced gradient, $\|Z_1^T g_{FR}\|$ (see Section 7.3).

norm_gf – double
Euclidean norm of the gradient function with respect to the free variables.

cond_t – double
a lower bound on the condition number of the working set.

cond_r – double
a lower bound on the condition number of the triangular factor R_1 (see Section 7.3).

x – double *
the components $\mathbf{x}[j-1]$ of the current point x , $j = 1, 2, \dots, \mathbf{st}->\mathbf{n}$.

ax – double *
if **st->nclin** > 0, the **st->nclin** components of the linear constraints Ax .

state – Integer *
state contains the status of the **st->n** variables and **st->nclin** general linear constraints. See Section 8.2 for a description of the possible status values.

diagt – double *
if **st->lin** > 0, the **st->lin** elements in the diagonal of the matrix T .

diagr – double *
if **st->nrnk** > 0, the first **st->nrnk** elements of the diagonal of the upper triangular matrix R .

If **comm->new_lm** = **TRUE** then the Lagrange multipliers have been updated and the following members of **st** are set:

bnd – Integer
the number of bound constraints in the working set.

kx – Integer *

bclambda – double *
indices of the bound constraints in the working set, with associated multipliers.
st->kx[*i*] is the index of the constraint with multiplier **st->bclambda**[*i*], for $i = 0, 1, \dots, \text{st->bnd}-1$.

lin – Integer
the number of linear constraints in the working set.

kactive – Integer *

lambda – double *
indices of the linear constraints in the working set, with associated multipliers.
st->kactive[*i*] is the index of the constraint with multiplier **st->lambda**[**st->bnd** + *i*] for $i = 0, 1, \dots, \text{st->lin}-1$.

nart – Integer
the number of artificial constraints in the working set (see Section 7.3).

gq – double *
st->gq[*i*], for $i = 0, 1, \dots, \text{st->nart}-1$, hold the multipliers for the artificial constraints.

If **comm->sol_prt** = **TRUE** then the final result from nag_opt_lin_lsq is available and the following members of **st** are set:

n – Integer
the number of variables.

nclin – Integer
the number of linear constraints.

iter – Integer
the iteration count.

x – double *
the components **x**[*j* - 1] of the final point *x*, for $j = 1, 2, \dots, \text{st->n}$.

feasible – Boolean
will be **TRUE** if the final point is feasible.

f – double
the final value of the objective function if **st->feasible** is **TRUE**; otherwise, the sum of infeasibilities. If the problem is of type FP and *x* is feasible then **f** is set to zero.

ax – double *
if **st->nclin** > 0, the **st->nclin** components of the final linear constraint activities, *Ax*.

state – Integer *
contains the final status of the **st->n** variables and **st->nclin** general linear constraints. See Section 8.2 for a description of the possible status values.

lambda – double *
contains the **st->n**+**st->nclin** final values of the Lagrange multipliers.

bl – double *
contains the **st->n**+**st->nclin** lower bounds.

bu – double *
contains the **st->n**+**st->nclin** upper bounds.

endstate – Nag_EndState
the state of termination of nag_opt_lin_lsq. Possible values of **endstate** and their correspondence to the exit value of **fail.code** are:

Value of endstate	Value of fail.code
Nag_Feasible or Nag_Optimal	NE_NOERROR
Nag_Weakmin	NW_SOLN_NOT_UNIQUE
Nag_Unbounded	NE_UNBOUNDED
Nag_Infeasible	NW_NOT_FEASIBLE
Nag_Too_Many_Iter	NW_TOO_MANY_ITER
Nag_Cycling	NE_CYCLING

The relevant members of the structure **comm** are:

it_prt – Boolean

will be **TRUE** when the print function is called with the result of the current iteration.

sol_prt – Boolean

will be **TRUE** when the print function is called with the final result.

new_lm – Boolean

will be **TRUE** when the Lagrange multipliers have been updated.

user – double *

iuser – Integer *

p – Pointer

pointers for communication of user information. If used they must be allocated memory by the user either before entry to **nag_opt_lin_lsq** or during a call to **print_fun**. The type Pointer will be **void *** with a C compiler that defines **void *** and **char *** otherwise.

9. Error Indications and Warnings

NE_INT_ARG_LT

On entry, **m** must not be less than 1: **m** = *<value>*.

On entry, **n** must not be less than 1: **n** = *<value>*.

On entry, **nclin** must not be less than 0: **nclin** = *<value>*.

NE_2_INT_ARG_LT

On entry, **tda** = *<value>* while **n** = *<value>*. These parameters must satisfy **tda** ≥ **n**.

On entry, **tdh** = *<value>* while **n** = *<value>*. These parameters must satisfy **tdh** ≥ **n**.

NE_ARRAY_CONS

The contents of array **kx** are not valid.

Constraint: must contain a permutation of integers 1, 2, ..., **n**.

NE_OPT_NOT_INIT

Options structure not initialized.

NE_BAD_PARAM

On entry, parameter **options.print_level** had an illegal value.

On entry, parameter **options.prob** had an illegal value.

On entry, parameter **options.start** had an illegal value.

NE_INVALID_INT_RANGE_1

Value *<value>* given to **options.max_iter** is not valid. Correct range is **max_iter** ≥ 0.

Value *<value>* given to **options.fmax_iter** is not valid. Correct range is **fmax_iter** ≥ 0.

NE_INVALID_REAL_RANGE_FF

Value *<value>* given to **options.crash_tol** is not valid. Correct range is $0.0 \leq \text{crash_tol} \leq 1.0$.

Value *<value>* given to **options.rank_tol** is not valid. Correct range is $0.0 < \text{rank_tol} < 1.0$.

NE_INVALID_REAL_RANGE_F

Value *<value>* given to **options.ftol** is not valid. Correct range is **ftol** > 0.0.

Value *<value>* given to **options.inf_bound** is not valid. Correct range is **inf_bound** > 0.0.

Value *<value>* given to **options.inf_step** is not valid. Correct range is **inf_step** > 0.0.

NE_CVEC_NULL

options.prob = *<value>* but argument **cvec** = NULL.

NE_B_NULL

options.prob = $\langle value \rangle$ but argument **b** = NULL.

NE_H_NULL_QP

options.prob = $\langle value \rangle$ but argument **h** = NULL. This problem type requires an array to be supplied in parameter **h**.

NE_WARM_START

options.start = **Nag_Warm** but pointer **options.state** = NULL.

NE_BOUND

The lower bound for variable $\langle value \rangle$ (array element **bl**[$\langle value \rangle$]) is greater than the upper bound.

NE_BOUND_LCON

The lower bound for linear constraint $\langle value \rangle$ (array element **bl**[$\langle value \rangle$]) is greater than the upper bound.

NE_STATE_VAL

options.state[$\langle value \rangle$] is out of range. **state**[$\langle value \rangle$] = $\langle value \rangle$.

NE_ALLOC_FAIL

Memory allocation failed.

NW_SOLN_NOT_UNIQUE

Optimal solution is not unique.

The point in **x** is a weak local minimum, i.e., the projected gradient is negligible, the Lagrange multipliers are optimal, but either R_z (see Section 7.3) is singular or there is a small multiplier. This means that x is not unique.

NE_UNBOUNDED

Solution appears to be unbounded.

This error indicator implies that a step as large as optional parameter **inf_step** (default value 10^{20} ; see Section 8.2) would have to be taken in order to continue the algorithm. This situation can occur only when H is singular, there is an explicit linear term, and at least one variable has no upper or lower bound.

NW_NOT_FEASIBLE

No feasible point was found for the linear constraints.

It was not possible to satisfy all the constraints to within the feasibility tolerance. In this case, the constraint violations at the final x will reveal a value of the tolerance for which a feasible point will exist – for example, if the feasibility tolerance for each violated constraint exceeds its **Residual** (see Section 4.1) at the final point. The modified problem (with an altered value of the optional feasibility tolerance, **ftol**) may then be solved using optional parameter **start** = **Nag_Warm** (see Section 8.2). The user should check that there are no constraint redundancies. If the data for the constraints are accurate only to the absolute precision σ , the user should ensure that the value of **ftol** is *greater* than σ . For example, if all elements of A are of order unity and are accurate only to three decimal places, **ftol** should be at least 10^{-3} .

NW_TOO_MANY_ITER

The maximum number of iterations, $\langle value \rangle$, have been performed.

The limiting number of iterations (determined by the optional parameters **max_iter** and **fmax_iter**, see Section 8.2) was reached before normal termination occurred. If the method appears to be making progress (e.g., the objective function is being satisfactorily reduced), either increase the iteration limits or, alternatively, rerun nag_opt_lin_lsq using the optional parameter **start** = **Nag_Warm** to specify the initial working set. If the iteration limit is already large, but some of the constraints could be nearly linearly dependent, check the extended iteration printout (see Section 8.3) for a repeated pattern of constraints entering and leaving the working set. (Near-dependencies are often indicated by wide variations in size in the diagonal elements of the matrix T (see Section 7.2), which will be printed if optional parameter **print_level** = **Nag_Soln_Iter_Full** (default value = **Nag_Soln_Iter**; see Section 8.2). In this case, the algorithm could be cycling (see the comments below for **fail.code** = **NE_CYCLING**).

NE_CYCLING

The algorithm could be cycling, since a total of 50 changes were made to the working set without altering x . Check the detailed iteration printout for a repeated pattern of constraint deletions and additions.

If a sequence of constraint changes is being repeated, the iterates are probably cycling. (`nag_opt_lin_lsq` does not contain a method that is guaranteed to avoid cycling; such a method would be combinatorial in nature.) Cycling may occur in two circumstances: at a constrained stationary point where there are some small or zero Lagrange multipliers; or at a point (usually a vertex) where the constraints that are satisfied exactly are nearly linearly dependent. In the latter case, the user has the option of identifying the offending dependent constraints and removing them from the problem, or restarting the run with a larger value of the optional parameter **ftol** (default value = $\sqrt{\epsilon}$, where ϵ is the *machine precision*; see Section 8.2). If this error exit occurs but no suspicious pattern of constraint changes can be observed, it may be worthwhile to restart with the final x (with optional parameter **start** = **Nag.Cold** or **Nag.Warm**).

NW_OVERFLOW_WARN

Serious ill conditioning in the working set after adding constraint $\langle value \rangle$. Overflow may occur in subsequent iterations.

If overflow occurs preceded by this warning then serious ill conditioning has probably occurred in the working set when adding a constraint. It may be possible to avoid the difficulty by increasing the magnitude of the optional parameter **ftol** and re-running the program. If the message recurs even after this change, the offending linearly dependent constraint j must be removed from the problem.

NE_NOT_APPEND_FILE

Cannot open file $\langle string \rangle$ for appending.

NE_WRITE_ERROR

Error occurred when writing to file $\langle string \rangle$.

NE_NOT_CLOSE_FILE

Cannot close file $\langle string \rangle$.

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.

10. Further Comments**10.1. Termination Criteria**

`nag_opt_lin_lsq` exits with **fail.code** = **NE.NOERROR** if x is a strong local minimizer, i.e., the reduced gradient is negligible, the Lagrange multipliers are optimal (see Section 4.1) and R_z (see Section 7.3) is non-singular.

10.2. Scaling

Sensible scaling of the problem is likely to reduce the number of iterations required and make the problem less sensitive to perturbations in the data, thus improving the condition of the problem. In the absence of better information it is usually sensible to make the Euclidean lengths of each constraint of comparable magnitude. See the Chapter Introduction and Gill *et al* (1981) for further information and advice.

10.3. Accuracy

`nag_opt_lin_lsq` implements a numerically stable active set strategy and returns solutions that are as accurate as the condition of the problem warrants on the machine.

11. References

Gill P E Hammarling S Murray W Saunders M A and Wright M H (1986) User's Guide for LSSOL (Version 1.0) *Report SOL 86-1*. Department of Operations Research, Stanford University.

Gill P E Murray W Saunders M A and Wright M H (1984) Procedures for optimization problems with a mixture of bounds and general linear constraints *ACM Trans. Math. Softw.* **10** 282–298.
 Gill P E Murray W and Wright M H (1981) *Practical Optimization*. Academic Press.
 Stoer J (1971) On the numerical solution of constrained least-squares problems. *SIAM J. Numer. Anal.* **8** 382–411.

12. See Also

nag_opt_lp (e04mfc)
 nag_opt_init (e04xxc)
 nag_opt_read (e04xyc)
 nag_opt_free (e04xzc)

13. Example 2

To minimize the quadratic function $c^T x + \frac{1}{2}x^T Hx$, where

$$c = (-4.0, -1.0, -1.0, -1.0, -1.0, -1.0, -1.0, -1.0, -0.3)^T,$$

$$H = \begin{pmatrix} 2 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 2 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 2 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 2 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

subject to the bounds

$$\begin{aligned} -2 &\leq x_1 \leq 2 \\ -2 &\leq x_2 \leq 2 \\ -2 &\leq x_3 \leq 2 \\ -2 &\leq x_4 \leq 2 \\ -2 &\leq x_5 \leq 2 \\ -2 &\leq x_6 \leq 2 \\ -2 &\leq x_7 \leq 2 \\ -2 &\leq x_8 \leq 2 \\ -2 &\leq x_9 \leq 2 \end{aligned}$$

and to the general constraints

$$\begin{aligned} -2.0 &\leq x_1 + x_2 + x_3 + x_4 + x_5 + x_6 + x_7 + x_8 + 4x_9 \leq 1.5 \\ -2.0 &\leq x_1 + 2x_2 + 3x_3 + 4x_4 - 2x_5 + x_6 + x_7 + x_8 + x_9 \leq 1.5 \\ -2.0 &\leq x_1 - x_2 + x_3 - x_4 + x_5 + x_6 + x_7 + x_8 + x_9 \leq 4.0 \end{aligned}$$

The initial point, which is feasible, is

$$x_0 = (0, 0, 0, 0, 0, 0, 0, 0, 0)^T,$$

and $F(x_0) = 0$.

The optimal solution (to five figures) is

$$x^* = (2.0, -0.23333, -0.26667, -0.3, -0.1, 2.0, -1.7777, -0.45555)^T,$$

and $F(x^*) = -8.0678$. Three bound constraints and two general constraints are active at the solution. Note that, although the Hessian matrix is positive semi-definite, the point x^* is unique.

This example illustrates the use of the **options** structure. Since the problem is of type QP2 (as described in Section 3) and the default value of the optional parameter **prob** is **Nag_LS1**, it is necessary to reset this parameter to **Nag_QP2** in order to solve the problem. This is achieved by declaring the **options** structure and initialising it by calling `nag_opt_init (e04xxc)`. Then **options.prob** is assigned directly, before calling `nag_opt_lin_lsq`. Note that the **cvec** parameter to `nag_opt_lin_lsq`, which was NULL in Example 1, needs to be supplied here, whereas **b**, which was supplied in Example 1 (and is required for all LS type problems), does not appear in QP problems and is NULL here. On return from `nag_opt_lin_lsq`, `nag_opt_free (e04xzc)` is used to free the memory assigned to the pointers in the options structure. Users should **not** use the standard C function `free()` for this purpose.

13.1. Program Text

```
static void ex2(void)
#else
static void ex2()
#endif
{
#define MMAX 10
#define NMAX 10
#define NCLIN 10
#define MAXBND NMAX+NCLIN

/* Local variables */
double a[NCLIN][NMAX], bl[MAXBND], bu[MAXBND];
double cvec[NMAX], h[NMAX][NMAX], x[NMAX];
double objf;
Integer kx[NMAX];
Integer i, j;
Integer m, n, nclin;
Integer tda, tdh;
Nag_E04_Opt options;
static NagError fail1, fail2;

fail1.print = TRUE;
fail2.print = TRUE;

Vprintf("\nExample 2: some options are set\n");
Vscanf(" %*[^\\n]"); /* Skip heading in data file */

/* Read problem dimensions */
Vscanf(" %*[^\\n]");
Vscanf("%ld%ld%ld%*[^\\n]", &m, &n, &nclin);

if (m <= MMAX && n <= NMAX && nclin <= NCLIN)
{
    tda = NMAX;
    tdh = NMAX;

    /* We solve a QP2 type problem in this example */

    /* Read cvec, h, a, bl, bu and x from data file */

    Vscanf(" %*[^\\n]");
    for (i = 0; i < m; ++i)
        Vscanf("%lf",&cvec[i]);

    Vscanf(" %*[^\\n]");
    for (i = 0; i < m; ++i)
```

```

    for (j = 0; j < n; ++j)
        Vscanf("%lf",&h[i][j]);

    if (nclin > 0)
    {
        Vscanf(" %*[^\\n]");
        for (i = 0; i < nclin; ++i)
            for (j = 0; j < n; ++j)
                Vscanf("%lf",&a[i][j]);
    }

    /* Read lower bounds */
    Vscanf(" %*[^\\n]");
    for (i = 0; i < n + nclin; ++i)
        Vscanf("%lf",&bl[i]);

    /* Read upper bounds */
    Vscanf(" %*[^\\n]");
    for (i = 0; i < n + nclin; ++i)
        Vscanf("%lf",&bu[i]);

    /* Read the initial point x */
    Vscanf(" %*[^\\n]");
    for (i = 0; i < n; ++i)
        Vscanf("%lf",&x[i]);

    /* Change the problem type */
    e04xxc(&options);
    options.prob = Nag_QP2;

    e04ncc(m, n, nclin, (double*)a, tda, bl, bu, cvec, (double*)0,
          (double*)h, tdh, kx, x, &objf,
          &options, NAGCOMM_NULL, &fail1);

    /* Free options memory */
    e04xzc(&options, "all", &fail2);
}
} /* ex2 */

```

13.2. Program Data

Values of m, n, nclin
 9 9 3

Objective function vector cvec
 -4.0 -1.0 -1.0 -1.0 -1.0 -1.0 -1.0 -0.1 -0.3

Objective function matrix H

2.0	1.0	1.0	1.0	1.0	0.0	0.0	0.0	0.0
1.0	2.0	1.0	1.0	1.0	0.0	0.0	0.0	0.0
1.0	1.0	2.0	1.0	1.0	0.0	0.0	0.0	0.0
1.0	1.0	1.0	2.0	1.0	0.0	0.0	0.0	0.0
1.0	1.0	1.0	1.0	2.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Linear constraint matrix A

1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	4.0
1.0	2.0	3.0	4.0	-2.0	1.0	1.0	1.0	1.0
1.0	-1.0	1.0	-1.0	1.0	1.0	1.0	1.0	1.0

Lower bounds
 -2.0 -2.0 -2.0 -2.0 -2.0 -2.0 -2.0 -2.0 -2.0 -2.0 -2.0 -2.0

Upper bounds
 2.0 2.0 2.0 2.0 2.0 2.0 2.0 2.0 2.0 1.5 1.5 4.0

Initial estimate of x
 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0

13.3. Program Results

Parameters to e04ncc

Linear constraints..... 3 Number of variables..... 9
 Objective matrix rows..... 9

prob..... Nag_QP2 start..... Nag_Cold
 ftol..... 1.05e-08 rank_tol..... 1.05e-07
 crash_tol..... 1.00e-02 hessian..... FALSE
 inf_bound..... 1.00e+20 inf_step..... 1.00e+20
 fmax_iter..... 60 max_iter..... 60
 machine precision..... 1.11e-16
 print_level..... Nag_Soln_Iter
 outfile..... stdout

Memory allocation:

state..... Nag
 ax..... Nag lambda..... Nag

Rank of the objective function data matrix = 5

Itn	Step	Ninf	Sinf/Objective	Norm Gz
0	0.0e+00	0	0.000000e+00	4.5e+00
1	7.5e-01	0	-4.375000e+00	5.0e-01
2	1.0e+00	0	-4.400000e+00	2.8e-17
3	3.0e-01	0	-4.700000e+00	8.9e-01
4	1.0e+00	0	-5.100000e+00	2.4e-17
5	5.4e-01	0	-6.055714e+00	1.7e+00
6	1.1e-02	0	-6.113326e+00	1.6e+00
7	1.1e-01	0	-6.215049e+00	1.2e+00
8	1.0e+00	0	-6.538008e+00	1.8e-17
9	6.5e-01	0	-7.428704e+00	7.2e-02
10	1.0e+00	0	-7.429717e+00	1.8e-17
11	1.0e+00	0	-8.067718e+00	1.8e-17
12	1.0e+00	0	-8.067778e+00	1.8e-17

Exit from QP problem after 12 iterations.

Varbl	State	Value	Lower Bound	Upper Bound	Lagr Mult	Residual
V 1	UL	2.00000e+00	-2.00000e+00	2.00000e+00	-8.0000e-01	0.0000e+00
V 2	FR	-2.33333e-01	-2.00000e+00	2.00000e+00	0.0000e+00	1.7667e+00
V 3	FR	-2.66667e-01	-2.00000e+00	2.00000e+00	0.0000e+00	1.7333e+00
V 4	FR	-3.00000e-01	-2.00000e+00	2.00000e+00	0.0000e+00	1.7000e+00
V 5	FR	-1.00000e-01	-2.00000e+00	2.00000e+00	0.0000e+00	1.9000e+00
V 6	UL	2.00000e+00	-2.00000e+00	2.00000e+00	-9.0000e-01	0.0000e+00
V 7	UL	2.00000e+00	-2.00000e+00	2.00000e+00	-9.0000e-01	0.0000e+00
V 8	FR	-1.77778e+00	-2.00000e+00	2.00000e+00	0.0000e+00	2.2222e-01
V 9	FR	-4.55556e-01	-2.00000e+00	2.00000e+00	0.0000e+00	1.5444e+00

L Con	State	Value	Lower Bound	Upper Bound	Lagr Mult	Residual
L 1	UL	1.50000e+00	-2.00000e+00	1.50000e+00	-6.6667e-02	1.1102e-15
L 2	UL	1.50000e+00	-2.00000e+00	1.50000e+00	-3.3333e-02	-4.4409e-16
L 3	FR	3.93333e+00	-2.00000e+00	4.00000e+00	0.0000e+00	6.6667e-02

Exit after 12 iterations.

Optimal QP solution found.

Final QP objective value = -8.0677778e+00