

NAG Fortran Library Routine Document

E04NQF

Note: before using this routine, please read the Users' Note for your implementation to check the interpretation of ***bold italicised*** terms and other implementation-dependent details.

Note: *this routine uses **optional parameters** to define choices in the problem specification and in the details of the algorithm. If you wish to use default settings for all of the optional parameters, you need only read Sections 1 to 9 of this document. Refer to the additional Sections 10, 11 and 12 for a detailed description of the algorithm, the specification of the optional parameters and a description of the monitoring information produced by the routine.*

1 Purpose

E04NQF solves sparse linear programming or convex quadratic programming problems. The initialization routine E04NPF **must** have been called prior to calling E04NQF.

2 Specification

```

SUBROUTINE E04NQF (START, QPHX, M, N, NE, NNAME, LENC, NCOLH, IOBJ,
1      OBJADD, PROB, ACOL, INDA, LOCA, BL, BU, C, NAMES,
2      HELAST, HS, X, PI, RC, NS, NINF, SINP, OBJ, CW,
3      LENCW, IW, LENIW, RW, LENRW, CUSER, IUSER, RUSER,
4      IFAIL)

    INTEGER          M, N, NE, NNAME, LENC, NCOLH, IOBJ, INDA(NE),
1      LOCA(N+1), HELAST(N+M), HS(N+M), NS, NINF, LENCW,
2      IW(LENIW), LENIW, LENRW, IUSER(*), IFAIL
    double precision OBJADD, ACOL(NE), BL(N+M), BU(N+M), C(*), X(N+M),
1      PI(M), RC(N+M), SINP, OBJ, RW(LENRW), RUSER(*)
    CHARACTER*1      START
    CHARACTER*8      PROB, NAMES(NNAME), CW(LENCW), CUSER(*)
    EXTERNAL         QPHX

```

Before calling E04NQF or one of the option setting routines E04NRF, E04NSF, E04NTF or E04NUF, routine E04NPF **must** be called. The specification for E04NPF is:

```

SUBROUTINE E04NPF (CW, LENCW, IW, LENIW, RW, LENRW, IFAIL)

    INTEGER          LENCW, IW(LENIW), LENIW, LENRW, IFAIL
    double precision RW(LENRW)
    CHARACTER*8      CW(LENCW)

```

LENCW, LENIW and LENRW, the declared lengths of CW, IW and RW respectively, must satisfy:

$$\text{LENCW} \geq 600$$

$$\text{LENIW} \geq 600$$

$$\text{LENRW} \geq 600$$

The contents of the arrays CW, IW and RW **must not** be altered between calling routines E04NPF, E04NQF, E04NRF, E04NSF, E04NTF and E04NUF.

After calling E04NQF you can call one or both of the routines E04NXF or E04NYF to obtain the current value of an optional parameter.

3 Description

E04NQF is designed to solve large scale *linear or quadratic programming problems* that are assumed to be stated in the following general form:

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

where x is a set of n variables, l and u are constant lower and upper bounds, and A is a sparse matrix and $f(x)$ is a linear or quadratic objective function that may be specified in a variety of ways, depending upon the particular problem being solved. The optional parameter **Maximize** (see Section 11.2) may be used to specify a problem in which $f(x)$ is maximized instead of minimized.

Upper and lower bounds are specified for all variables and constraints. This form allows full generality in specifying various types of constraint. In particular, the j th constraint may be defined as an equality by setting $l_j = u_j$. If certain bounds are not present, the associated elements of l or u may be set to special values that are treated as $-\infty$ or $+\infty$.

The possible forms for the function $f(x)$ are summarized in Table 1. The most general form for $f(x)$ is

$$f(x) = q + c^T x + \frac{1}{2} x^T H x = q + \sum_{j=1}^n c_j x_j + \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n x_i H_{ij} x_j$$

where q is a constant, c is a constant n vector and H is a constant symmetric n by n matrix with elements $\{H_{ij}\}$. In this form, f is a quadratic function of x and (1) is known as a *quadratic program* (QP). E04NQF is suitable for all *convex* quadratic programs. The defining feature of a *convex* QP is that the matrix H must be *positive semi-definite*, i.e., it must satisfy $x^T H x \geq 0$ for all x . If not, $f(x)$ is nonconvex and E04NQF will terminate with the error indicator IFAIL = 11. If $f(x)$ is nonconvex it may be more appropriate to call E04VHF instead.

Problem type	Objective function $f(x)$	Hessian matrix H
FP	Not applicable	$q = c = H = 0$
LP	$q + c^T x$	$H = 0$
QP	$q + c^T x + \frac{1}{2} x^T H x$	Symmetric positive semi-definite

Table 1
Choices for the objective function $f(x)$

If $H = 0$, then $f(x) = q + c^T x$ and the problem is known as a *linear program* (LP). In this case, rather than defining an H with zero elements, you can define H to have no columns by setting NCOLH = 0 (see Section 5).

If $H = 0$, $q = 0$, and $c = 0$, there is no objective function and the problem is a *feasible point problem* (FP), which is equivalent to finding a point that satisfies the constraints on x . In the situation where no feasible point exists, several options are available for finding a point that minimizes the constraint violations (see the optional parameter **Elastic Mode** in Section 11.2).

E04NQF is suitable for large LPs and QPs in which the matrix A is *sparse*, i.e., when there are sufficiently many zero elements in A to justify storing them implicitly. The matrix A is input to E04NQF by means of the three array parameters ACOL, IND A and LOCA. This allows the user to specify the pattern of non-zero elements in A .

E04NQF exploits structure or sparsity in H by requiring H to be defined *implicitly* in a subroutine that computes the product Hx for any given vector x . In many cases, the product Hx can be computed very efficiently for any given x , e.g., H may be a sparse matrix, or a sum of matrices of rank-one.

For problems in which A can be treated as a *dense* matrix, it is usually more efficient to use E04MFF/E04MFA, E04NCF/E04NCA or E04NFF/E04NFA.

There is considerable flexibility allowed in the definition of $f(x)$ in Table 1. The vector c defining the linear term $c^T x$ can be input in three ways: as a sparse row of A ; as an explicit dense vector c ; or as both a sparse row and an explicit vector (in which case, $c^T x$ will be the sum of two linear terms). When stored in A , c is the IOBJth row of A , which is known as the *objective row*. The objective row must always be a *free* row of A in the sense that its lower and upper bounds must be $-\infty$ and $+\infty$. Storing c as part of A is recommended if c is a sparse vector. Storing c as an explicit vector is recommended for a sequence of problems, each with a different objective (see parameters C and LENC).

The upper and lower bounds on the m elements of Ax are said to define the *general constraints* of the problem. Internally, E04NQF converts the general constraints to equalities by introducing a set of *slack variables* s , where $s = (s_1, s_2, \dots, s_m)^T$. For example, the linear constraint $5 \leq 2x_1 + 3x_2 \leq +\infty$ is replaced by $2x_1 + 3x_2 - s_1 = 0$, together with the bounded slack $5 \leq s_1 \leq +\infty$. The problem defined by (1) can therefore be re-written in the following equivalent form:

$$\underset{x \in R^n, s \in R^m}{\text{minimize}} f(x) \quad \text{subject to } Ax - s = 0, \quad \ell \leq \begin{pmatrix} x \\ s \end{pmatrix} \leq u.$$

Since the slack variables s are subject to the same upper and lower bounds as the elements of Ax , the bounds on x and Ax can simply be thought of as bounds on the combined vector (x, s) . (In order to indicate their special role in QP problems, the original variables x are sometimes known as ‘column variables’, and the slack variables s are known as ‘row variables’.)

Each LP or QP problem is solved using an *active-set* method. This is an iterative procedure with two phases: a *feasibility phase* (Phase 1), in which the sum of infeasibilities is minimized to find a feasible point; and an *optimality phase* (Phase 2), in which $f(x)$ is minimized (or maximized) by constructing a sequence of iterations that lies within the feasible region.

Phase 1 involves solving a linear program of the form

Phase 1

$$\underset{x, v, w}{\text{minimize}} \sum_{j=1}^{n+m} (v_j + w_j) \quad \text{subject to } Ax - s = 0, \quad \ell \leq \begin{pmatrix} x \\ s \end{pmatrix} - v + w \leq u, \quad v \geq 0, \quad w \geq 0$$

which is equivalent to minimizing the sum of the constraint violations. If the constraints are feasible (i.e., at least one feasible point exists), eventually a point will be found at which both v and w are zero. The associated value of (x, s) satisfies the original constraints and is used as the starting point for the Phase 2 iterations for minimizing $f(x)$.

If the constraints are infeasible (i.e., $v \neq 0$ or $w \neq 0$ at the end of Phase 1), no solution exists for (1) and the user has the option of either terminating or continuing in so-called *Elastic mode* (see the discussion of the optional parameter **Elastic Mode** in Section 11.2). In elastic mode, a ‘relaxed’ or ‘perturbed’ problem is solved in which $f(x)$ is minimized while allowing some of the bounds to become ‘elastic’, i.e., to change from their specified values. Variables subject to elastic bounds are known as *elastic variables*. An elastic variable is free to violate one or both of its original upper or lower bounds. The user is able to assign which bounds will become elastic if elastic mode is ever started (see the parameter HELAST in Section 5).

To make the relaxed problem meaningful, E04NQF minimizes $f(x)$ while (in some sense) finding the ‘smallest’ violation of the elastic variables. In the situation where all the variables are elastic, the relaxed problem has the form

Phase 2 (γ)

$$\underset{x, v, w}{\text{minimize}} f(x) + \gamma \sum_{j=1}^{n+m} (v_j + w_j) \quad \text{subject to } Ax - s = 0, \quad \ell \leq \begin{pmatrix} x \\ s \end{pmatrix} - v + w \leq u, \quad v \geq 0, \quad w \geq 0,$$

where γ is a nonnegative parameter known as the elastic weight (see the optional parameter **Elastic Weight** in Section 11.2), and $f(x) + \gamma \sum_j (v_j + w_j)$ is called the *composite objective*. In the more general situation where only a subset of the bounds are elastic, the v ’s and w ’s for the non-elastic bounds are fixed at zero.

The *elastic weight* can be chosen to make the composite objective behave like either the original objective $f(x)$ or the sum of infeasibilities. If $\gamma = 0$, E04NQF will attempt to minimize f subject to the (true) upper and lower bounds on the nonelastic variables (and declare the problem infeasible if the nonelastic variables cannot be made feasible).

At the other extreme, choosing γ sufficiently large, will have the effect of minimizing the sum of the violations of the elastic variables subject to the original constraints on the non-elastic variables. Choosing a large value of the elastic weight is useful for defining a ‘least-infeasible’ point for an infeasible problem.

In Phase 1 and elastic mode, all calculations involving v and w are done implicitly in the sense that an elastic variable x_j is allowed to violate its lower bound (say) and an explicit value of v can be recovered as $v_j = l_j - x_j$.

A constraint is said to be *active* or *binding* at x if the associated element of either x or Ax is equal to one of its upper or lower bounds. Since an active constraint in Ax has its associated slack variable at a bound, the status of both simple and general upper and lower bounds can be conveniently described in terms of the status of the variables (x, s) . A variable is said to be *nonbasic* if it is temporarily fixed at its upper or lower bound. It follows that regarding a general constraint as being *active* is equivalent to thinking of its associated slack as being *nonbasic*.

At each iteration of an active-set method, the constraints $Ax - s = 0$ are (conceptually) partitioned into the form

$$Bx_B + Sx_S + Nx_N = 0,$$

where x_N consists of the nonbasic elements of (x, s) and the *basis matrix* B is square and non-singular. The elements of x_B and x_S are called the *basic* and *superbasic* variables respectively; with x_N they are a permutation of the elements of x and s . At a QP solution, the basic and superbasic variables will lie somewhere between their upper or lower bounds, while the nonbasic variables will be equal to one of their bounds. At each iteration, x_S is regarded as a set of independent variables that are free to move in any desired direction, namely one that will improve the value of the objective function (or sum of infeasibilities). The basic variables are then adjusted in order to ensure that (x, s) continues to satisfy $Ax - s = 0$. The number of superbasic variables (n_S say) therefore indicates the number of degrees of freedom remaining after the constraints have been satisfied. In broad terms, n_S is a measure of *how nonlinear* the problem is. In particular, n_S will always be zero for FP and LP problems.

If it appears that no improvement can be made with the current definition of B , S and N , a nonbasic variable is selected to be added to S , and the process is repeated with the value of n_S increased by one. At all stages, if a basic or superbasic variable encounters one of its bounds, the variable is made nonbasic and the value of n_S is decreased by one.

Associated with each of the m equality constraints $Ax - s = 0$ is a *dual variable* π_i . Similarly, each variable in (x, s) has an associated *reduced gradient* d_j (also known as a *reduced cost*). The reduced gradients for the variables x are the quantities $g - A^T\pi$, where g is the gradient of the QP objective function; and the reduced gradients for the slack variables s are the dual variables π . The QP subproblem is optimal if $d_j \geq 0$ for all nonbasic variables at their lower bounds, $d_j \leq 0$ for all nonbasic variables at their upper bounds and $d_j = 0$ for all superbasic variables. In practice, an *approximate* QP solution is found by slightly relaxing these conditions on d_j (see the description of the optional parameter **Optimality Tolerance** in Section 11.2).

The process of computing and comparing reduced gradients is known as *pricing* (a term first introduced in the context of the simplex method for linear programming). To ‘price’ a nonbasic variable x_j means that the reduced gradient d_j associated with the relevant active upper or lower bound on x_j is computed via the formula $d_j = g_j - a_j^T\pi$, where a_j is the j th column of $(A \mid -I)$. (The variable selected by such a process and the corresponding value of d_j (i.e., its reduced gradient) are the quantities +SBS and dj in the monitoring file output; see Section 12.) If A has significantly more columns than rows (i.e., $n \gg m$), pricing can be computationally expensive. In this case, a strategy known as *partial pricing* can be used to compute and compare only a subset of the d_j ’s.

E04NQF is based on SQOPT, which is part of the SNOPT package described in Gill *et al.* (1999). It uses stable numerical methods throughout and includes a reliable basis package (for maintaining sparse LU factors of the basis matrix B), a practical anti-degeneracy procedure, efficient handling of linear constraints

and bounds on the variables (by an active-set strategy), as well as automatic scaling of the constraints. Further details can be found in Section 10.

4 References

- Fourer R (1982) Solving staircase linear programs by the simplex method *Math. Programming* **23** 274–313
- Gill P E and Murray W (1978) Numerically stable methods for quadratic programming *Math. Programming* **14** 349–372
- Gill P E, Murray W and Saunders M A (1995) User's guide for QPOPT 1.0: a Fortran package for quadratic programming *Report SOL 95-4* Department of Operations Research, Stanford University
- Gill P E, Murray W and Saunders M A (1999) Users' guide for SQOPT 5.3: a Fortran package for large-scale linear and quadratic programming *Report SOL 99* Department of Operations Research, Stanford University
- Gill P E, Murray W, Saunders M A and Wright M H (1987) Maintaining *LU* factors of a general sparse matrix *Linear Algebra and its Applics.* **88/89** 239–270
- Gill P E, Murray W, Saunders M A and Wright M H (1989) A practical anti-cycling procedure for linearly constrained optimization *Math. Programming* **45** 437–474
- Gill P E, Murray W, Saunders M A and Wright M H (1991) Inertia-controlling methods for general quadratic programming *SIAM Rev.* **33** 1–36
- Hall J A J and McKinnon K I M (1996) The Simplest Examples where the Simplex Method Cycles and Conditions where EXPAND Fails to Prevent Cycling *Report MS 96–100* Department of Mathematics and Statistics, University of Edinburgh

5 Parameters

The first n entries of the parameters BL, BU, HS and X refer to the variables x . The last m entries refer to the slacks s .

- 1: START – CHARACTER*1

Input

On entry: indicates how a starting basis (and certain other items) are to be obtained.

If START = 'C' (Cold Start), then an internal Crash procedure will be used to choose an initial Basis matrix, unless a basis file is provided via **Old Basis File**, **Load File** or **Insert File** (see Section 11.2);
 if START = 'B' (Basis file), is the same as START = 'C' but is more meaningful when a basis file is given (see **Old Basis File** in Section 11.2);
 if START = 'W' (Warm Start), then a basis is already defined in HS (probably from a previous call).

Constraint: START = 'C', 'B' or 'W'.

- 2: QPHX – SUBROUTINE, supplied by the NAG Fortran Library or the user. *External Procedure*

For QP problems, you must supply a version of QPHX to compute the matrix product Hx for the given vector x . If H has rows and columns of zeros, it is most efficient to order the variables $x = (y \ z)^T$ so that

$$Hx = \begin{pmatrix} H_1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} y \\ z \end{pmatrix} = \begin{pmatrix} H_1 y \\ 0 \end{pmatrix},$$

where the nonlinear variables y appear first as shown. The number of columns of H_1 is specified in NCOLH. For FP and LP problems, QPHX will never be called by E04NQF and hence QPHX may be the dummy routine E04NSH.

Its specification is:

<pre> SUBROUTINE QPHX (NCOLH, X, HX, NSTATE, CUSER, IUSER, RUSER) INTEGER NCOLH, NSTATE, IUSER(*) double precision X(NCOLH), HX(NCOLH), RUSER(*) CHARACTER*8 CUSER(*) </pre>		
1:	NCOLH – INTEGER	<i>Input</i>
	<i>On entry:</i> this is the same parameter NCOLH as supplied to E04NQF (see above).	
2:	X(NCOLH) – double precision array	<i>Input</i>
	<i>On entry:</i> the first NCOLH elements of the vector x .	
3:	HX(NCOLH) – double precision array	<i>Output</i>
	<i>On exit:</i> the product Hx .	
4:	NSTATE – INTEGER	<i>Input</i>
	<i>On entry:</i> if NSTATE = 1, then E04NQF is calling QPHX for the first time. This parameter setting allows you to save computation time if certain data must be read or calculated only once. To preserve this data for subsequent calculation place it in one of CUSER, RUSER or IUSER. If NSTATE = 0 there is nothing special about the current call of QPHX. If NSTATE ≥ 2 , then E04NQF is calling QPHX for the last time. This parameter setting allows you to perform some additional computation on the final solution. On the last call of QPHX, if NSTATE = 2, the current x is optimal; if NSTATE = 3, the problem appears to be infeasible; if NSTATE = 4, the problem appears to be unbounded; and if NSTATE = 5, the iterations limit was reached.	
5:	CUSER(*) – CHARACTER*8 array	<i>User Workspace</i>
6:	IUSER(*) – INTEGER array	<i>User Workspace</i>
7:	RUSER(*) – double precision array	<i>User Workspace</i>
	QPHX is called from E04NQF with the parameters CUSER, IUSER and RUSER as supplied to E04NQF. These parameters are not touched by E04NQF. The array parameters can be used; as an alternative to using COMMON.	

QPHX must be declared as EXTERNAL in the (sub)program from which E04NQF is called. Parameters denoted as *Input* must **not** be changed by this procedure.

- | | | |
|----|---|--------------|
| 3: | M – INTEGER | <i>Input</i> |
| | <i>On entry:</i> m , the number of general linear constraints (or slacks). This is the number of rows in A , including the free row (if any; see IOBJ below). | |
| | <i>Constraint:</i> $M \geq 1$. | |
| 4: | N – INTEGER | <i>Input</i> |
| | <i>On entry:</i> n , the number of variables (excluding slacks). This is the number of columns in the linear constraint matrix A . | |
| | <i>Constraint:</i> $N \geq 1$. | |
| 5: | NE – INTEGER | <i>Input</i> |
| | <i>On entry:</i> the number of non-zero elements in A . | |
| | <i>Constraint:</i> $1 \leq NE \leq N \times M$. | |

- 6: NNAME – INTEGER *Input*
On entry: the number of column (i.e., variable) and row names supplied in the array NAMES.
 If NNAME = 1, there are no names. Default names will be used in the printed output;
 if NNAME = N + M, all names must be supplied.
Constraint: NNAME = 1 or N + M.
- 7: LENC – INTEGER *Input*
On entry: the number of elements in the constant objective vector c .
Constraint: $0 \leq \text{LENC} \leq N$.
- 8: NCOLH – INTEGER *Input*
On entry: n_H , the number of leading non-zero columns of the Hessian matrix H . For FP and LP problems, NCOLH must be set to zero.
Constraint: $0 \leq \text{NCOLH} \leq N$.
- 9: IOBJ – INTEGER *Input*
On entry: if IOBJ > 0, row IOBJ of A is a free row containing the non-zero elements of the vector c appearing in the linear objective term $c^T x$. If IOBJ = 0, there is no free row, i.e., the problem is either an FP problem, or a QP problem with $c = 0$.
Constraint: $0 \leq \text{IOBJ} \leq M$.
- 10: OBJADD – *double precision* *Input*
On entry: the constant q , to be added to the objective for printing purposes. Typically OBJADD = 0.0D0.
- 11: PROB – CHARACTER*8 *Input*
On entry: the name for the problem. It is used in the printed solution and in some routines that output Basis files. A blank name may be used.
- 12: ACOL(NE) – *double precision* array *Input*
On entry: the non-zero elements of A , ordered by increasing column index. Note that all elements must be assigned a value in the calling program.
- 13: INDA(NE) – INTEGER array *Input*
On entry: INDA(i) must contain the row index of the non-zero element stored in ACOL(i), for $i = 1, 2, \dots, \text{NE}$. Thus a pair of values (ACOL(k), INDA(k)) contains a matrix element and its corresponding row index.
 If LENC > 0, the first LENC elements of ACOL and INDA belong to variables corresponding to the constant objective term c .
 If the problem has a quadratic objective, the first NCOLH columns of ACOL and INDA belong to variables corresponding to the non-zero block of the QP Hessian. Subroutine QPHX knows about these variables.
 Note that the row indices for a column may be supplied in any order.
Constraint: $1 \leq \text{INDA}(i) \leq M$, for $i = 1, 2, \dots, \text{NE}$.
- 14: LOCA(N + 1) – INTEGER array *Input*
On entry: LOCA(j) must contain the index in ACOL and INDA of the start of the j th column, for $j = 1, 2, \dots, N$. Thus for $j = 1 : n$, the entries of column j are held in ACOL($k:l$) and their corresponding row indices are in INDA($k:l$), where $k = \text{LOCA}(j)$ and $l = \text{LOCA}(j + 1) - 1$. To specify the j th column as empty, set LOCA(j) = LOCA($j + 1$). Note that the first and last

elements of LOCA must be such that $\text{LOCA}(1) = 1$ and $\text{LOCA}(N + 1) = \text{NE} + 1$. If your problem has no constraints, or just bounds on the variables, you may include a dummy ‘free’ row with a single (zero) element by setting $\text{ACOL}(1) = 0.0$, $\text{INDA}(1) = 1$, $\text{LOCA}(1) = 1$, and $\text{LOCA}(j) = 2$ for $j = 2 : n + 1$. This row is made ‘free’ by setting its bounds to be $\text{BL}(n + 1) = -\text{bigbnd}$ and $\text{BU}(n + 1) = \text{bigbnd}$.

Constraints:

$$\begin{aligned} \text{LOCA}(1) &= 1; \\ \text{LOCA}(j) &\geq 1, \text{ for } j = 2, 3, \dots, N; \\ \text{LOCA}(N + 1) &= \text{NE} + 1; \\ 0 &\leq \text{LOCA}(j + 1) - \text{LOCA}(j) \leq M, \text{ for } j = 1, 2, \dots, N. \end{aligned}$$

- 15: $\text{BL}(N + M)$ – **double precision** array *Input/Output*

On entry: l , the lower bounds for all the variables and general constraints, in the following order. The first N elements of BL must contain the bounds on the variables x , and the next M elements the bounds for the general linear constraints Ax (or slacks s) and the free row (if any). To fix the j th variable, set $\text{BL}(j) = \text{BU}(j) = \beta$, say, where $|\beta| < \text{bigbnd}$. To specify a non-existent lower bound (i.e., $l_j = -\infty$), set $\text{BL}(j) \leq -\text{bigbnd}$, where bigbnd is the value of the optional parameter **Infinite Bound Size** (see Section 11.2). To specify the j th constraint as an *equality*, set $\text{BL}(n + j) = \text{BU}(n + j) = \beta$, say, where $|\beta| < \text{bigbnd}$. Note that the lower bound corresponding to the free row must be set to $-\infty$ and stored in $\text{BL}(N + \text{IOBJ})$.

On exit: used as internal workspace prior to being restored and hence is unchanged.

Constraint: if $\text{IOBJ} > 0$, $\text{BL}(N + \text{IOBJ}) \leq -\text{bigbnd}$

(See also the description for BU below.).

- 16: $\text{BU}(N + M)$ – **double precision** array *Input/Output*

On entry: u , the upper bounds for all the variables and general constraints, in the following order. The first N elements of BU must contain the bounds on the variables x , and the next M elements the bounds for the general linear constraints Ax (or slacks s) and the free row (if any). To specify a non-existent upper bound (i.e., $u_j = +\infty$), set $\text{BU}(j) \geq \text{bigbnd}$. Note that the upper bound corresponding to the free row must be set to $+\infty$ and stored in $\text{BU}(N + \text{IOBJ})$.

On exit: used as internal workspace prior to being restored and hence is unchanged.

Constraints:

$$\begin{aligned} \text{if } \text{IOBJ} > 0, & \text{BU}(N + \text{IOBJ}) \geq \text{bigbnd}; \\ & \text{BL}(i) \leq \text{BU}(i) \text{ otherwise.} \end{aligned}$$

- 17: $\text{C}(*)$ – **double precision** array *Input*

On entry: contains the explicit objective vector c (if any). If the problem is of type *FP*, or if $\text{LENC} = 0$, then C is not referenced. (In that case, C may be dimensioned (1), or it could be any convenient array.)

- 18: $\text{NAMES}(\text{NNAME})$ – CHARACTER*8 array *Input*

On entry: the optional column and row names, respectively.

If $\text{NNAME} = 1$, NAMES is not referenced and the printed output will use default names for the columns and rows;

if $\text{NNAME} = N + M$, the first N elements must contain the names for the columns and the next M elements must contain the names for the rows. Note that the name for the free row (if any) must be stored in $\text{NAMES}(N + \text{IOBJ})$.

- 19: $\text{HELAST}(N + M)$ – INTEGER array *Input*

On entry: defines which variables are to be treated as being elastic in elastic mode. The allowed values of HELAST are:

HELAST(*j*) Status in elastic mode

- 0 variable *j* is non-elastic and cannot be infeasible
- 1 variable *j* can violate its lower bound
- 2 variable *j* can violate its upper bound
- 3 variable *j* can violate either its lower or upper bound

HELAST need not be assigned if optional parameter **Elastic Mode** = 0.

Constraint: HELAST(*j*) = 0, 1, 2, 3 if **Elastic Mode** ≠ 0, for *j* = 1, 2, ..., N + M.

20: HS(N + M) – INTEGER array

Input/Output

On entry: if START = 'C' or 'B', and a Basis file of some sort is to be input (an Old Basis file, Insert file or Load file), then HS and X need not be set at all.

If START = 'C' and there is no basis file, the first N elements of HS and X must specify the initial states and values, respectively, of the variables *x*. (The slacks *s* need not be initialized.) An internal Crash procedure is then used to select an initial basis matrix *B*. The initial basis matrix will be triangular (neglecting certain small elements in each column). It is chosen from various rows and columns of (*A* −*I*). Possible values for HS(*j*) are as follows:

HS(*j*) State of X(*j*) during Crash procedure

- 0 or 1 Eligible for the basis
- 2 Ignored
- 3 Eligible for the basis (given preference over 0 or 1)
- 4 or 5 Ignored

If nothing special is known about the problem, or there is no wish to provide special information, you may set HS(*j*) = 0 and X(*j*) = 0.0, for *j* = 1, 2, ..., N. All variables will then be eligible for the initial basis. Less trivially, to say that the *j*th variable will probably be equal to one of its bounds, set HS(*j*) = 4 and X(*j*) = BL(*j*) or HS(*j*) = 5 and X(*j*) = BU(*j*) as appropriate.

Following the Crash procedure, variables for which HS(*j*) = 2 are made superbasic. Other variables not selected for the basis are then made nonbasic at the value X(*j*) if BL(*j*) ≤ X(*j*) ≤ BU(*j*), or at the value BL(*j*) or BU(*j*) closest to X(*j*).

If START = 'W', HS and X must specify the initial states and values, respectively, of the variables and slacks (*x*, *s*). If E04NQF has been called previously with the same values of N and M, HS already contains satisfactory information.

Constraints:

- if START = 'C', 0 ≤ HS(*j*) ≤ 5 for *j* = 1, 2, ..., N;
- if START = 'W', 0 ≤ HS(*j*) ≤ 3 for *j* = 1, 2, ..., N + M.

On exit: the final states of the variables and slacks (*x*, *s*). The significance of each possible value of HS(*j*) is as follows:

HS(*j*) State of variable *j* Normal value of X(*j*)

- 0 Nonbasic BL(*j*)
- 1 Nonbasic BU(*j*)
- 2 Superbasic Between BL(*j*) and BU(*j*)
- 3 Basic Between BL(*j*) and BU(*j*)

If NINF = 0, basic and superbasic variables may be outside their bounds by as much as the value of the optional parameter **Feasibility Tolerance** (see Section 11.2). Note that unless the optional parameter **Scale Option** = 0 (see Section 11.2) is specified, the **Feasibility Tolerance** applies to the variables of the scaled problem. In this case, the variables of the original problem may be as much as 0.1 outside their bounds, but this is unlikely unless the problem is very badly scaled.

Very occasionally some nonbasic variables may be outside their bounds by as much as the **Feasibility Tolerance**, and there may be some nonbasic variables for which X(*j*) lies strictly between its bounds.

If $NINF > 0$, some basic and superbasic variables may be outside their bounds by an arbitrary amount (bounded by $SINF$ if **Scale Option** = 0).

21: $X(N + M)$ – **double precision** array *Input/Output*

On entry: the initial values of the variables x , if $START = 'W'$ and slacks s , i.e., (x, s) . (See the description for HS above.)

On exit: the final values of the variables and slacks (x, s) .

22: $PI(M)$ – **double precision** array *Output*

On exit: contains the dual variables π (a set of Lagrange multipliers (shadow prices) for the general constraints).

23: $RC(N + M)$ – **double precision** array *Output*

On exit: the first N elements contain the reduced costs, $g - (A \ -I)^T \pi$, where g is the gradient of the objective if X is feasible (or the gradient of the Phase 1 objective otherwise). The last M entries are π .

24: NS – INTEGER *Input/Output*

On entry: n_S , the number of superbasics. For QP problems, NS need not be specified if $START = 'C'$, but must retain its value from a previous call when $START = 'W'$. For FP and LP problems, NS need not be initialized.

On exit: the final number of superbasics. This will be zero for FP and LP problems.

25: $NINF$ – INTEGER *Output*

On exit: the number of infeasibilities.

26: $SINF$ – **double precision** *Output*

On exit: the sum of the scaled infeasibilities. This will be zero if $NINF = 0$, and is most meaningful when **Scale Option** = 0.

27: OBJ – **double precision** *Output*

On exit: the value of the objective function.

If $NINF = 0$, OBJ includes the quadratic objective term $\frac{1}{2}x^T Hx$ (if any);

if $NINF > 0$, OBJ is just the linear objective term $c^T x$ (if any).

For FP problems, OBJ is set to zero.

28: $CW(LENCW)$ – CHARACTER*8 array *Communication Array*

29: $LENCW$ – INTEGER *Input*

On entry: the dimension of the array CW as declared in the (sub)program from which E04NQF is called.

Constraint: $LENCW \geq 600$.

30: $IW(LENIW)$ – INTEGER array *Communication Array*

31: $LENIW$ – INTEGER *Input*

On entry: the dimension of the array IW as declared in the (sub)program from which E04NQF is called.

Constraint: $LENIW \geq 600$.

- 32: RW(LENRW) – **double precision** array Communication Array
 33: LENRW – INTEGER Input

On entry: the dimension of the array RW as declared in the (sub)program from which E04NQF is called.

Constraint: $\text{LENRW} \geq 600$.

- 34: CUSER(*) – CHARACTER*8 array User Workspace

Note: the dimension of the array CUSER must be at least 1.

CUSER is not used by E04NQF, but is passed directly to the external procedure QPHX and may be used to pass information to that routine.

- 35: IUSER(*) – INTEGER array User Workspace

Note: the dimension of the array IUSER must be at least 1.

IUSER is not used by E04NQF, but is passed directly to the external procedure QPHX and may be used to pass information to that routine.

- 36: RUSER(*) – **double precision** array User Workspace

Note: the dimension of the array RUSER must be at least 1.

RUSER is not used by E04NQF, but is passed directly to the external procedure QPHX and may be used to pass information to that routine.

- 37: IFAIL – INTEGER Input/Output

On initial entry: IFAIL must be set to 0, –1 or 1. Users who are unfamiliar with this parameter should refer to Chapter P01 for details.

On final exit: IFAIL = 0 unless the routine detects an error (see Section 6). An exit value of 0 does not imply that the objective function was convex.

For environments where it might be inappropriate to halt program execution when an error is detected, the value –1 or 1 is recommended. If the output of error messages is undesirable, then the value 1 is recommended. Otherwise, because for this routine the values of the output parameters may be useful even if IFAIL $\neq$ 0 on exit, the recommended value is –1. **When the value –1 or 1 is used it is essential to test the value of IFAIL on exit.**

E04NQF returns with IFAIL = 0 if the reduced gradient (Norm rg; see Section 8.1) is negligible, the Lagrange multipliers (Lagr Mult; see Section 8.1) are optimal and x satisfies the constraints to the accuracy requested by the value of the optional parameter **Feasibility Tolerance** (see Section 11.2).

6 Error Indicators and Warnings

If on entry IFAIL = 0 or –1, explanatory error messages are output on the current error message unit (as defined by X04AAF).

Errors or warnings detected by the routine:

IFAIL = 1

The initialization routine E04NPF has not been called or at least one of LENCW, LENIW and LENRW is less than 600.

IFAIL = 2

An input parameter is invalid.

IFAIL = 3

The requested accuracy could not be achieved.

IFAIL = 4

Weak solution found. The final x is not unique.

This exit will occur when

- (i) the problem is feasible;
- (ii) the reduced gradient is negligible;
- (iii) the Lagrange multipliers are optimal;
- (iv) the reduced Hessian is singular or there are some very small multipliers.

This exit cannot occur if H is positive-definite (i.e., $f(x)$ is strictly convex).

IFAIL = 5

The problem is infeasible. The general constraints cannot all be satisfied simultaneously to within the value of the optional parameter **Feasibility Tolerance** (see Section 11.2).

Feasibility is measured with respect to the upper and lower bounds on the variables. The message tells us that among all the points satisfying the general constraints $Ax - s = 0$, there is apparently no point that satisfies the bounds on x and s . Violations as small as the **Feasibility Tolerance** (see Section 11.2) are ignored, but at least one component of x or s violates a bound by more than the tolerance.

Note: although the objective function is the sum of infeasibilities (when $NINF > 0$), this sum will not necessarily have been minimized when **Elastic Mode** = 1.

IFAIL = 6

The problem is unbounded (or badly scaled). For a minimization problem, the objective function is not bounded below in the feasible region.

For linear problems, unboundedness is detected by the simplex method when a nonbasic variable can apparently be increased or decreased by an arbitrary amount without causing a basic variable to violate a bound. A message will give the index of the nonbasic variable. Consider adding an upper or lower bound to the variable. Also, examine the constraints that have non-zeros in the associated column, to see if they have been formulated as intended.

Very rarely, the scaling of the problem could be so poor that numerical error will give an erroneous indication of unboundedness. Consider using the **Scale Option** (see Section 11.2).

IFAIL = 7

Too many iterations. The value of the optional parameter **Iteration Limit** (see Section 11.2) is too small.

The Iterations limit was exceeded before the required solution could be found. Check the iteration log to be sure that progress was being made. If so, restart the run using a basis file that was saved at the end of the run.

IFAIL = 8

The reduced Hessian matrix $Z^T H Z$ (see Section 10.2) exceeds its assigned dimension. The value of the optional parameter **Superbasics Limit** (see Section 11.2) is too small.

In general, raise the **Superbasics Limit** (see Section 11.2) s by a reasonable amount, bearing in mind the storage needed for the reduced Hessian is $\frac{1}{2}s^2$.

IFAIL = 9

The basis is singular after several attempts to factorize it (adding slacks where necessary). Either the problem is badly scaled or the value of the optional parameter **LU Factor Tolerance** (see Section 11.2) is too large.

IFAIL = 10

Numerical error in trying to satisfy the general constraints. The basis is very ill-conditioned.

An *LU* factorization of the basis has just been obtained and used to recompute the basic variables xB , given the present values of the superbasic and nonbasic variables. However, a row check has revealed that the resulting solution does not satisfy the current constraints $Ax - s = 0$ sufficiently well.

This probably means that the current basis is very ill-conditioned. Request the **Scale Option** (see Section 11.2) if there are any linear constraints and variables.

For certain highly structured basis matrices (notably those with band structure), a systematic growth may occur in the factor U . Consult the description of **Umax**, **Umin** and **Growth** in Section 12, and set the **LU Factor Tolerance** to 2.0 (or possibly even smaller, but not less than 1.0).

IFAIL = 11

The Hessian matrix H appears to be indefinite. This sometimes occurs because the values of the optional parameters **LU Factor Tolerance** and **LU Update Tolerance** (see Section 11.2) are too large. Check also that subroutine QPHX has been coded correctly and that all relevant elements of Hx have been assigned their correct values.

IFAIL = 12

Internal memory allocation failed when attempting to obtain the required workspace. Please contact NAG.

IFAIL = 13

Internal memory allocation was insufficient. Please contact NAG.

IFAIL = 14

An error has occurred in the basis package, perhaps indicating incorrect setup of arrays **INDA** and **LOCA**. Set the optional argument **Print File** (see Section 11.2) and examine the output carefully for further information.

IFAIL = 15

An unexpected error has occurred. Set the optional argument **Print File** (see Section 11.2) and examine the output carefully for further information.

7 Accuracy

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

8 Further Comments

This section contains a description of the printed output.

8.1 Description of the Printed Output

If **Print Level** > 0, one line of information is output to the **Print File** every k th iteration, where k is the specified **Print Frequency**. A heading is printed before the first such line following a basis factorization. The heading contains the items described below. In this description, a pricing operation is defined to be the process by which one or more nonbasic variables are selected to become superbasic (in addition to those already in the superbasic set). The variable selected will be denoted by jq . If the problem is purely linear, variable jq will usually become basic immediately (unless it should happen to reach its opposite bound and return to the nonbasic set).

If **Partial Price** is in effect, variable j_q is selected from A_{pp} or I_{pp} , the pp th segments of the constraint matrix $(A \mid -I)$.

Label	Description
Itn	is the iteration count.
pp	is the optional indicator. The variable selected by the last pricing operation came from the pp th partition of A and $-I$. Note that pp is reset to zero whenever the basis is refactorized.
dj	is the value of the reduced gradient (or reduced cost) for the variable selected by the pricing operation at the start of the current iteration. Algebraically, d_j is $d_j = g_j - \pi^T a_j$, for $j = j_q$ where g_j is the gradient of the current objective function, π is the vector of dual variables, and a_j is the jth column of the constraint matrix $(A \mid -I)$. Note that d_j is the norm of the reduced-gradient vector at the start of the iteration, just after the pricing operation.
+SBS	is the variable j_q selected by the pricing operation to be added to the superbasic set.
-SBS	is the variable chosen to leave the superbasic set. It has become basic if the entry under $-B$ is non-zero, otherwise it becomes nonbasic.
-BS	is the variable removed from the basis (if any) to become nonbasic.
-B	is the variable chosen to leave the set of basics (if any) in a special basic $\leftrightarrow$ superbasic swap. The entry under $-SBS$ has become basic if this entry is non-zero, and nonbasic otherwise. The swap is done to ensure that there are no superbasic slacks.
Step	is the value of the step length α taken along the current search direction p . The variables x have just been changed to $x + \alpha p$. If a variable is made superbasic during the current iteration (i.e., $+SBS$ is positive), $Step$ will be the step to the nearest bound. During the optimality phase, the step can be greater than unity only if the reduced Hessian is not positive-definite.
Pivot	is the r th element of a vector y satisfying $By = a_q$ whenever a_q (the q th column of the constraint matrix $(A - I)$) replaces the r th column of the basis matrix B . Wherever possible, $Step$ is chosen so as to avoid extremely small values of $Pivot$ (since they may cause the basis to be nearly singular). In extreme cases, it may be necessary to increase the value of the optional parameter Pivot Tolerance (see Section 11.2) to exclude very small elements of y from consideration during the computation of $Step$.
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 the 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.
L	is the number of non-zeros in the basis factor L . Immediately after a basis factorization $B = LU$, this entry contains $1enL$ (see Section 12). Further non-zeros are added to L when various columns of B are later replaced. (Thus, L increases monotonically.)

U	is the number of non-zeros in the basis factor U . Immediately after a basis factorization $B = LU$, this entry contains <code>lenU</code> (see Section 12). As columns of B are replaced, the matrix U is maintained explicitly (in sparse form). The value of U may fluctuate up or down; in general, it will tend to increase.
Ncp	is the number of compressions required to recover workspace in the data structure for U . This includes the number of compressions needed during the previous basis factorization. Normally, Ncp should increase very slowly.

The following will be output if the problem is QP or if the superbasis is non-empty (i.e., if the current solution is nonbasic).

Label	Description
Norm rg	is $\ d_S\ $, the Euclidean norm of the reduced gradient (see Section 10.3). During the optimality phase, this norm will be approximately zero after a unit step.
Ns	is the current number of superbasis variables.
Cond Hz	is a lower bound on the condition number of the reduced Hessian (see Section 10.2). The larger this number, the more difficult the problem. Attention should be given to the scaling of the variables and the constraints to guard against high values of Cond Hz.

9 Example

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

$$c = (-200.0, -2000.0, -2000.0, -2000.0, -2000.0, 400.0, 400.0)^T$$

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

subject to the bounds

$$\begin{aligned} 0 &\leq x_1 \leq 200 \\ 0 &\leq x_2 \leq 2500 \\ 400 &\leq x_3 \leq 800 \\ 100 &\leq x_4 \leq 700 \\ 0 &\leq x_5 \leq 1500 \\ 0 &\leq x_6 \\ 0 &\leq x_7 \end{aligned}$$

and to the linear constraints

$$\begin{aligned} x_1 + x_2 + x_3 + x_4 + x_5 + x_6 + x_7 &= 2000 \\ 0.15x_1 + 0.04x_2 + 0.02x_3 + 0.04x_4 + 0.02x_5 + 0.01x_6 + 0.03x_7 &\leq 60 \\ 0.03x_1 + 0.05x_2 + 0.08x_3 + 0.02x_4 + 0.06x_5 + 0.01x_6 &\leq 100 \\ 0.02x_1 + 0.04x_2 + 0.01x_3 + 0.02x_4 + 0.02x_5 &\leq 40 \\ 0.02x_1 + 0.03x_2 &+ 0.01x_5 \leq 30 \\ 1500 &\leq 0.70x_1 + 0.75x_2 + 0.80x_3 + 0.75x_4 + 0.80x_5 + 0.97x_6 \\ 250 &\leq 0.02x_1 + 0.06x_2 + 0.08x_3 + 0.12x_4 + 0.02x_5 + 0.01x_6 + 0.97x_7 \leq 300 \end{aligned}$$

The initial point, which is infeasible, is

$$x_0 = (0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0)^T.$$

The optimal solution (to five figures) is

$$x^* = (0.0, 349.40, 648.85, 172.85, 407.52, 271.36, 150.02)^T.$$

One bound constraint and four linear constraints are active at the solution. Note that the Hessian matrix H is positive semi-definite.

9.1 Program Text

Note: the listing of the example program presented below uses **bold italicised** terms to denote precision-dependent details. Please read the Users' Note for your implementation to check the interpretation of these terms. As explained in the Essential Introduction to this manual, the results produced may not be identical for all implementations.

```
*      E04NQF Example Program Text
*      Mark 21 Release. NAG Copyright 2004.
      IMPLICIT      NONE
*
*      .. Parameters ..
      INTEGER       NIN, NOUT
      PARAMETER     (NIN=5,NOUT=6)
      INTEGER       NMAX, MMAX, NEMAX
      PARAMETER     (NMAX=100,MMAX=100,NEMAX=100)
      INTEGER       LENCW, LENIW, LENRW
      PARAMETER     (LENCW=600,LENIW=600,LENRW=600)
*
*      .. Local Scalars ..
      DOUBLE PRECISION OBJ, OBJADD, SINF
      INTEGER       I, ICOL, IFAIL, IOBJ, J, JCOL, LENC, M, N, NCOLH,
+
      CHARACTER     NE, NINF, NNAME, NS
      CHARACTER     START
      CHARACTER*8    PROB
*
*      .. Local Arrays ..
      DOUBLE PRECISION ACOL(NEMAX), BL(NMAX+MMAX), BU(NMAX+MMAX), C(1),
+
      PI(MMAX), RC(NMAX+MMAX), RUSER(1), RW(LENRW),
+
      X(NMAX+MMAX)
      INTEGER       HELAST(NMAX+MMAX), HS(NMAX+MMAX), INDA(NEMAX),
+
      IUSER(1), IW(LENIW), LOCA(NMAX+1)
      CHARACTER*8    CUSER(1), CW(LENCW), NAMES(NMAX+MMAX)
*
*      .. External Subroutines ..
      EXTERNAL      E04NPF, E04NQF, E04NSF, E04NTF, QPHX
*
*      .. Executable Statements ..
      WRITE (NOUT,*) 'E04NQF Example Program Results'
*
*      Skip heading in data file.
      READ (NIN,*)
      READ (NIN,*) N, M
      IF (N.LE.NMAX .AND. M.LE.MMAX) THEN
*
*          Read NE, IOBJ, NCOLH, START and NNAME from data file.
      READ (NIN,*) NE, IOBJ, NCOLH, START, NNAME
*
*          Read NAMES from data file.
      READ (NIN,*) (NAMES(I),I=1,NNAME)
*
*          Read the matrix ACOL from data file. Set up LOCA.
      JCOL = 1
      LOCA(JCOL) = 1
      DO 40 I = 1, NE
*
*          Element ( INDA( I ), ICOL ) is stored in ACOL( I ).
      READ (NIN,*) ACOL(I), INDA(I), ICOL
*
*          IF (ICOL.LT.JCOL) THEN
*              Elements not ordered by increasing column index.
      WRITE (NOUT,99999) 'Element in column', ICOL,
+
      ' found after element in column', JCOL, '. Problem',
+
      ' abandoned.'
      STOP
      ELSE IF (ICOL.EQ.JCOL+1) THEN
```

```

*           Index in ACOL of the start of the ICOL-th column equals I.
              LOCA(ICOL) = I
              JCOL = ICOL
        ELSE IF (ICOL.GT.JCOL+1) THEN
*           Index in ACOL of the start of the ICOL-th column equals I,
*           but columns JCOL+1,JCOL+2,...,ICOL-1 are empty. Set the
*           corresponding elements of LOCA to I.
              DO 20 J = JCOL + 1, ICOL - 1
                  LOCA(J) = I
20          CONTINUE
              LOCA(ICOL) = I
              JCOL = ICOL
        END IF
40    CONTINUE

*           LOCA(N+1) = NE + 1

*
*           IF (N.GT.ICOL) THEN
*           Columns N,N-1,...,ICOL+1 are empty. Set the corresponding
*           elements of LOCA accordingly.
              DO 60 I = N, ICOL + 1, -1
                  LOCA(I) = LOCA(I+1)
60          CONTINUE
        END IF

*
*           Read BL, BU, HS and X from data file.
              READ (NIN,*) (BL(I),I=1,N+M)
              READ (NIN,*) (BU(I),I=1,N+M)
              IF (START.EQ.'C') THEN
                  READ (NIN,*) (HS(I),I=1,N)
              ELSE IF (START.EQ.'W') THEN
                  READ (NIN,*) (HS(I),I=1,N+M)
              END IF
              READ (NIN,*) (X(I),I=1,N)

*
*           Call E04NPF to initialise E04NQF.
              IFAIL = -1
              CALL E04NPF(CW,LENCW,IW,LENIW,RW,LENRW,IFAIL)

*
*           By default E04NQF does not print monitoring
*           information. Set the print file unit or the summary
*           file unit to get information.
              CALL E04NTF('Print file',NOUT,CW,IW,RW,IFAIL)

*
*           We have no explicit objective vector so set LENC = 0; the
*           objective vector is stored in row IOBJ of ACOL.
              LENC = 0
              OBJADD = 0.0DO
              PROB = ' '

*
*           Do not allow any elastic variables (i.e. they cannot be
*           infeasible). If we'd set optional argument "Elastic mode" to 0,
*           we wouldn't need to set the individual elements of array HELAST.
              DO 80 I = 1, N + M
                  HELAST(I) = 0
80          CONTINUE

*
*           Solve the QP problem.
              IFAIL = -1
              CALL E04NQF(START,QPHX,M,N,NE,NNAME,LENC,NCOLH,IOBJ,OBJADD,
+                 PROB,ACOL,INDA,LOCA,BL,BU,C,NAMES,HELAST,HS,X,PI,
+                 RC,NS,NINF,SINF,OBJ,CW,LENCW,IW,LENIW,RW,LENRW,
+                 CUSER,IUSER,RUSER,IFAIL)

*
              WRITE (NOUT,*)
              WRITE (NOUT,99998) IFAIL
              IF (IFAIL.EQ.0) THEN
                  WRITE (NOUT,99997) OBJ
                  WRITE (NOUT,99996) (X(I),I=1,N)
              END IF

*

```

```

      END IF
      STOP
*
99999 FORMAT (1X,A,I5,A,I5,A,A)
99998 FORMAT (1X,'On exit from E04NQF, IFAIL = ',I5)
99997 FORMAT (1X,'Final objective value = ',1P,E11.3)
99996 FORMAT (1X,'Optimal X = ',7F9.2)
      END
*
      SUBROUTINE QPHX(NCOLH,X,HX,NSTATE,CUSER,IUSER,RUSER)
*      Routine to compute H*x. (In this version of QPHX, the Hessian
*      matrix H is not referenced explicitly.)
*      .. Parameters ..
      DOUBLE PRECISION TWO
      PARAMETER      (TWO=2.0D+0)
*      .. Scalar Arguments ..
      INTEGER        NCOLH, NSTATE
*      .. Array Arguments ..
      DOUBLE PRECISION HX(NCOLH), RUSER(*), X(NCOLH)
      INTEGER        IUSER(*)
      CHARACTER*8     CUSER(*)
*      .. Executable Statements ..
      HX(1) = TWO*X(1)
      HX(2) = TWO*X(2)
      HX(3) = TWO*(X(3)+X(4))
      HX(4) = HX(3)
      HX(5) = TWO*X(5)
      HX(6) = TWO*(X(6)+X(7))
      HX(7) = HX(6)
      RETURN
      END

```

9.2 Program Data

E04NQF Example Program Data

```

  7  8      : Values of N and M
48  8  7  'C' 15      : Values of NNZ, IOBJ, NCOLH, START and NNAME

'...X1...' '...X2...' '...X3...' '...X4...' '...X5...'
'...X6...' '...X7...' '..ROW1..' '..ROW2..' '..ROW3..'
'..ROW4..' '..ROW5..' '..ROW6..' '..ROW7..' '..COST..' : End of array NAMES

    0.02    7    1      : Sparse matrix A, ordered by increasing column index;
    0.02    5    1      : each row contains ACOL(i), INDA(i), ICOL (= column index)
    0.03    3    1      : The row indices may be in any order. In this example
    1.00    1    1      : row 8 defines the linear objective term transpose(C)*X.
    0.70    6    1
    0.02    4    1
    0.15    2    1
-200.00    8    1
    0.06    7    2
    0.75    6    2
    0.03    5    2
    0.04    4    2
    0.05    3    2
    0.04    2    2
    1.00    1    2
-2000.00   8    2
    0.02    2    3
    1.00    1    3
    0.01    4    3
    0.08    3    3
    0.08    7    3
    0.80    6    3
-2000.00   8    3
    1.00    1    4
    0.12    7    4
    0.02    3    4
    0.02    4    4
    0.75    6    4

```

```

0.04    2    4
-2000.00 8    4
0.01    5    5
0.80    6    5
0.02    7    5
1.00    1    5
0.02    2    5
0.06    3    5
0.02    4    5
-2000.00 8    5
1.00    1    6
0.01    2    6
0.01    3    6
0.97    6    6
0.01    7    6
400.00   8    6
0.97    7    7
0.03    2    7
1.00    1    7
400.00   8    7      : End of matrix A

0.0      0.0      4.0E+02  1.0E+02  0.0      0.0
0.0      2.0E+03 -1.0E+25 -1.0E+25 -1.0E+25 -1.0E+25
1.5E+03  2.5E+02 -1.0E+25      : End of lower bounds array BL

2.0E+02  2.5E+03  8.0E+02  7.0E+02  1.5E+03  1.0E+25
1.0E+25  2.0E+03  6.0E+01  1.0E+02  4.0E+01  3.0E+01
1.0E+25  3.0E+02  1.0E+25      : End of upper bounds array BU

0  0  0  0  0  0  0      : Initial array HS
0.0 0.0 0.0 0.0 0.0 0.0 0.0 : Initial vector X

```

9.3 Program Results

E04NQF Example Program Results

Parameters
=====

Files

Solution file.....	0	Old basis file	0	(Print file).....	6
Insert file.....	0	New basis file	0	(Summary file).....	0
Punch file.....	0	Backup basis file.....	0		
Load file.....	0	Dump file.....	0		

Frequencies

Print frequency.....	100	Check frequency.....	60	Save new basis map....	100
Summary frequency.....	100	Factorization frequency	50	Expand frequency.....	10000

LP/QP Parameters

Minimize.....		QPsolver Cholesky.....		Cold start.....	
Scale tolerance.....	0.900	Feasibility tolerance..	1.00E-06	Iteration limit.....	10000
Scale option.....	2	Optimality tolerance...	1.00E-06	Print level.....	1
Crash tolerance.....	0.100	Pivot tolerance.....	2.04E-11	Partial price.....	1
Crash option.....	3	Elastic weight.....	1.00E+00	Prtl price section (A)	7
Elastic mode.....	1	Elastic objective.....	1	Prtl price section (-I)	8

QP objective

Objective variables....	7	Hessian columns.....	7	Superbasics limit.....	7
Nonlin Objective vars..	7	Unbounded step size....	1.00E+20		
Linear Objective vars..	0				

Miscellaneous

LU factor tolerance....	100.00	LU singularity tol.....	2.04E-11	Timing level.....	0
LU update tolerance....	10.00	LU swap tolerance.....	1.03E-04	Debug level.....	0

```

LU partial pivoting...          eps (machine precision) 1.11E-16      System information.....      No

Nonlinear constraints      0      Linear constraints      8
Nonlinear variables       7      Linear variables      0
Jacobian variables       0      Objective variables   7
Total constraints         8      Total variables      7

Itn      1: Feasible constraints

E04NQF EXIT  0 -- finished successfully
E04NQF INFO  1 -- optimality conditions satisfied

Problem name
No. of iterations          9      Objective value      -1.8477846771E+06
No. of Hessian products   16      Linear objective      -2.9886903537E+06
                                   Quadratic objective  1.1409056766E+06
No. of superbasics        2      No. of basic nonlinears  4
No. of degenerate steps   0      Percentage           0.00
Max x      (scaled)      3 2.4E-01      Max pi      (scaled)      6 4.7E+07
Max x      3 6.5E+02      Max pi      7 1.5E+04
Max Prim inf(scaled)      0 0.0E+00      Max Dual inf(scaled)      6 1.3E-08
Max Primal infeas        0 0.0E+00      Max Dual infeas          9 3.8E-11

Name                                     Objective Value      -1.8477846771E+06

Status      Optimal Soln      Iteration      9      Superbasics      2

Section 1 - Rows

Number  ...Row.. State  ...Activity... Slack Activity  ..Lower Limit.  ..Upper Limit.  ..Dual Activity  ..i

      8  ..ROW1..  EQ      2000.00000      .      2000.00000      2000.00000      -12900.76766      1
      9  ..ROW2..  BS      49.23160      -10.76840      None      60.00000      0.00000      2
     10  ..ROW3..  UL      100.00000      .      None      100.00000      -2324.86620      3
     11  ..ROW4..  BS      32.07187      -7.92813      None      40.00000      .      4
     12  ..ROW5..  BS      14.55719      -15.44281      None      30.00000      .      5
     13  ..ROW6..  LL      1500.00000      .      1500.00000      None      14454.60290      6
     14  ..ROW7..  LL      250.00000      .      250.00000      300.00000      14580.95432      7
     15  ..COST..  BS     -2988690.35370     -2988690.35370      None      None      -1.0      8

Section 2 - Columns

Number  .Column. State  ...Activity...  .Obj Gradient.  ..Lower Limit.  ..Upper Limit.  Reduced Gradnt  m+j

      1  ...X1...  LL      .      -200.00000      .      200.00000      2360.67253      9
      2  ...X2...  BS      349.39923     -1301.20153      .      2500.00000      0.00000      10
      3  ...X3...  SBS      648.85342     -356.59829      400.00000      800.00000      0.00000      11
      4  ...X4...  SBS      172.84743     -356.59829      100.00000      700.00000      0.00000      12
      5  ...X5...  BS      407.52089     -1184.95822      .      1500.00000      0.00000      13
      6  ...X6...  BS      271.35624      1242.75804      .      None      0.00000      14
      7  ...X7...  BS      150.02278      1242.75804      .      None      0.00000      15

On exit from E04NQF, IFAIL =      0
Final objective value = -1.848E+06
Optimal x =      0.00      349.40      648.85      172.85      407.52      271.36      150.02

```

Note: the remainder of this document is intended for more advanced users. Section 10 contains a detailed algorithm description that may be needed in order to understand Sections 11 and 12. Section 11 describes the optional parameters that may be set by calls to E04NRF, E04NSF, E04NTF and/or E04NUF. Section 12 describes the quantities that can be requested to monitor the course of the computation.

10 Algorithmic Details

This section contains a description of the method used by E04NQF.

10.1 Overview

E04NQF is based on an inertia-controlling method that maintains a Cholesky factorization of the reduced Hessian (see below). The method is similar to that of Gill and Murray (1978), and is described in detail by Gill *et al.* (1991). Here we briefly summarize the main features of the method. Where possible, explicit reference is made to the names of variables that are parameters of the routine or appear in the printed output.

The method used has two distinct 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 computations in both phases are performed by the same subroutines. The two-phase nature of the algorithm is reflected by changing the function being minimized from the sum of infeasibilities (the printed quantity `Sinf`; see Section 12) to the quadratic objective function (the printed quantity `Objective`; see Section 12).

In general, an iterative process is required to solve a quadratic program. Given an iterate (x, s) in both the original variables x and the slack variables s , a new iterate $(\bar{x}, \bar{s})$ is defined by

$$\begin{pmatrix} \bar{x} \\ \bar{s} \end{pmatrix} = \begin{pmatrix} x \\ s \end{pmatrix} + \alpha p, \quad (2)$$

where the *step length* α is a non-negative scalar (the printed quantity `Step`; see Section 12), and p is called the *search direction*. (For simplicity, we shall consider a typical iteration and avoid reference to the index of the iteration.) Once an iterate is feasible (i.e., satisfies the constraints), all subsequent iterates remain feasible.

10.2 Definition of the Working Set and Search Direction

At each iterate (x, s) , a *working set* of constraints is defined to be a linearly independent subset of the constraints that are satisfied ‘exactly’ (to within the value of the optional parameter **Feasibility Tolerance**; see Section 11.2). The working set is the current prediction of the constraints that hold with equality at a solution of the LP or QP problem. Let m_W denote the number of constraints in the working set (including bounds), and let W denote the associated m_W by $(n + m)$ *working set matrix* consisting of the m_W gradients of the working set constraints.

The search direction is defined so that constraints in the working set remain *unaltered* for any value of the step length. It follows that p must satisfy the identity

$$Wp = 0. \quad (3)$$

This characterization allows p to be computed using any n by n_Z full-rank matrix Z that spans the null space of W . (Thus, $n_Z = n - m_W$ and $WZ = 0$.) The null space matrix Z is defined from a sparse LU factorization of part of W (see (6) and (7) below). The direction p will satisfy (3) if

$$p = Zp_Z, \quad (4)$$

where p_Z is any n_Z -vector.

The working set contains the constraints $Ax - s = 0$ and a subset of the upper and lower bounds on the variables (x, s) . Since the gradient of a bound constraint $x_j \geq l_j$ or $x_j \leq u_j$ is a vector of all zeros except for ± 1 in position j , it follows that the working set matrix contains the rows of $(A \mid -I)$ and the unit rows associated with the upper and lower bounds in the working set.

The working set matrix W can be represented in terms of a certain column partition of the matrix $(A \mid -I)$ by (conceptually) partitioning the constraints $Ax - s = 0$ so that

$$Bx_B + Sx_S + Nx_N = 0, \quad (5)$$

where B is a square non-singular basis and x_B , x_S and x_N are the basic, superbasis and nonbasic variables

respectively. The nonbasic variables are equal to their upper or lower bounds at (x, s) , and the superbasic variables are independent variables that are chosen to improve the value of the current objective function. The number of superbasic variables is n_S (the printed quantity Ns ; see Section 12). Given values of x_N and x_S , the basic variables x_B are adjusted so that (x, s) satisfies (5).

If P is a permutation matrix such that $(A \mid -I)P = (B \mid S \mid N)$, then W satisfies

$$WP = \begin{pmatrix} B & S & N \\ 0 & 0 & I_N \end{pmatrix}, \quad (6)$$

where I_N is the identity matrix with the same number of columns as N .

The null space matrix Z is defined from a sparse LU factorization of part of W . In particular, Z is maintained in ‘reduced gradient’ form, using the LUSOL package (see Gill *et al.* (1991)) to maintain sparse LU factors of the basis matrix B that alters as the working set W changes. Given the permutation P , the null space basis is given by

$$Z = P \begin{pmatrix} -B^{-1}S \\ I \\ 0 \end{pmatrix}. \quad (7)$$

This matrix is used only as an operator, i.e., it is never computed explicitly. Products of the form Zv and $Z^T g$ are obtained by solving with B or B^T . This choice of Z implies that n_Z , the number of ‘degrees of freedom’ at (x, s) , is the same as n_S , the number of superbasic variables.

Let g_Z and H_Z denote the *reduced gradient* and *reduced Hessian* of the objective function:

$$g_Z = Z^T g \quad \text{and} \quad H_Z = Z^T H Z, \quad (8)$$

where g is the objective gradient at (x, s) . Roughly speaking, g_Z and H_Z describe the first and second derivatives of an n_S -dimensional *unconstrained* problem for the calculation of p_Z . (The condition estimator of H_Z is the quantity Cond Hz in the monitoring file output; see Section 12.)

At each iteration, an upper triangular factor R is available such that $H_Z = R^T R$. Normally, R is computed from $R^T R = Z^T H Z$ at the start of the optimality phase and then updated as the QP working set changes. For efficiency, the dimension of R should not be excessive (say, $n_S \leq 1000$). This is guaranteed if the number of nonlinear variables is ‘moderate’.

If the QP problem contains linear variables, H is positive semi-definite and R may be singular with at least one zero diagonal element. However, an inertia-controlling strategy is used to ensure that only the last diagonal element of R can be zero. (See Gill *et al.* (1991) for a discussion of a similar strategy for indefinite quadratic programming.)

If the initial R is singular, enough variables are fixed at their current value to give a non-singular R . This is equivalent to including temporary bound constraints in the working set. Thereafter, R can become singular only when a constraint is deleted from the working set (in which case no further constraints are deleted until R becomes non-singular).

10.3 Main Iteration

If the reduced gradient is zero, (x, s) is a constrained stationary point on the working set. During the feasibility phase, the reduced gradient will usually be zero only at a vertex (although it may be zero elsewhere in the presence of constraint dependencies). During the optimality phase, a zero reduced gradient implies that x minimizes the quadratic objective function when the constraints in the working set are treated as equalities. At a constrained stationary point, Lagrange multipliers λ are defined from the equations

$$W^T \lambda = g(x). \quad (9)$$

A Lagrange multiplier λ_j corresponding to an inequality constraint in the working set is said to be *optimal* if $\lambda_j \leq \sigma$ when the associated constraint is at its *upper bound*, or if $\lambda_j \geq -\sigma$ when the associated constraint is at its *lower bound*, where σ depends on the value of the optional parameter **Optimality Tolerance** (see Section 11.2). If a multiplier is non-optimal, the objective function (either the true objective or the sum of infeasibilities) can be reduced by continuing the minimization with the

corresponding constraint excluded from the working set. (This step is sometimes referred to as ‘deleting’ a constraint from the working set.) If optimal multipliers occur during the feasibility phase but the sum of infeasibilities is non-zero, there is no feasible point and the routine terminates immediately with IFAIL = 3 (see Section 6).

The special form (6) of the working set allows the multiplier vector λ , the solution of (9), to be written in terms of the vector

$$d = \begin{pmatrix} g \\ 0 \end{pmatrix} - (A_B^T - I)^T \pi = \begin{pmatrix} g - A_B^T \pi \\ \pi \end{pmatrix}, \quad (10)$$

where π satisfies the equations $B^T \pi = g_B$, and g_B denotes the basic elements of g . The elements of π are the Lagrange multipliers λ_j associated with the equality constraints $Ax - s = 0$. The vector d_N of nonbasic elements of d consists of the Lagrange multipliers λ_j associated with the upper and lower bound constraints in the working set. The vector d_S of superbasic elements of d is the reduced gradient g_Z in (8). The vector d_B of basic elements of d is zero, by construction. (The Euclidean norm of d_S and the final values of d_S , g and π are the quantities Norm rg, Reduced Gradnt, Obj Gradient and Dual Activity in the monitoring file output; see Section 12.)

If the reduced gradient is not zero, Lagrange multipliers need not be computed and the search direction is given by $p = Zp_Z$ (see (7) and (11)). The step length is chosen to maintain feasibility with respect to the satisfied constraints.

There are two possible choices for p_Z , depending on whether or not H_Z is singular. If H_Z is non-singular, R is non-singular and p_Z in (4) is computed from the equations

$$R^T R p_Z = -g_Z, \quad (11)$$

where g_Z is the reduced gradient at x . In this case, $(x, s) + p$ is the minimizer of the objective function subject to the working set constraints being treated as equalities. If $(x, s) + p$ is feasible, α is defined to be unity. In this case, the reduced gradient at $(\bar{x}, \bar{s})$ will be zero, and Lagrange multipliers are computed at the next iteration. Otherwise, α is set to α_M , the step to the ‘nearest’ constraint along p . This constraint is then added to the working set at the next iteration.

If H_Z is singular, then R must also be singular, and an inertia-controlling strategy is used to ensure that only the last diagonal element of R is zero. (See Gill *et al.* (1991) for a discussion of a similar strategy for indefinite quadratic programming.) In this case, p_Z satisfies

$$p_Z^T H_Z p_Z = 0 \quad \text{and} \quad g_Z^T p_Z \leq 0, \quad (12)$$

which allows the objective function to be reduced by any step of the form $(x, s) + \alpha p$, where $\alpha > 0$. The vector $p = Zp_Z$ is a direction of unbounded descent for the QP problem in the sense that the QP objective is linear and decreases without bound along p . If no finite step of the form $(x, s) + \alpha p$ (where $\alpha > 0$) reaches a constraint not in the working set, the QP problem is unbounded and the routine terminates immediately with IFAIL = 6 (see Section 6). Otherwise, α is defined as the maximum feasible step along p and a constraint active at $(x, s) + \alpha p$ is added to the working set for the next iteration.

E04NQF makes explicit allowance for infeasible constraints. Infeasible linear constraints are detected first by solving a problem of the form

$$\underset{x,v,w}{\text{minimize}} \quad e^T(v+w) \quad \text{subject to} \quad l \leq \begin{Bmatrix} x \\ Gx - v + w \end{Bmatrix} \leq u, \quad v \geq 0, \quad w \geq 0, \quad (13)$$

where $e = (1, 1, \dots, 1)^T$. This is equivalent to minimizing the sum of the general linear constraint violations subject to the simple bounds. (In the linear programming literature, the approach is often called *elastic programming*.)

10.4 Miscellaneous

If the basis matrix is not chosen carefully, the condition of the null space matrix Z in (7) could be arbitrarily high. To guard against this, the routine implements a ‘basis repair’ feature in which the LUSOL package (see Gill *et al.* (1991)) is used to compute the rectangular factorization

$$(B \ S)^T = LU, \quad (14)$$

returning just the permutation P that makes PLP^T unit lower triangular. The pivot tolerance is set to require $|PLP^T|_{ij} \leq 2$, and the permutation is used to define P in (6). It can be shown that $\|Z\|$ is likely to be little more than unity. Hence, Z should be well-conditioned *regardless of the condition of W* . This feature is applied at the beginning of the optimality phase if a potential $B - S$ ordering is known.

The EXPAND procedure (see Gill *et al.* (1989)) is used to reduce the possibility of cycling at a point where the active constraints are nearly linearly dependent. Although there is no absolute guarantee that cycling will not occur, the probability of cycling is extremely small (see Hall and McKinnon (1996)). The main feature of EXPAND is that the feasibility tolerance is increased at the start of every iteration. This allows a positive step to be taken at every iteration, perhaps at the expense of violating the bounds on (x, s) by a small amount.

Suppose that the value of the optional parameter **Feasibility Tolerance** (see Section 11.2) is δ . Over a period of K iterations (where K is the value of the optional parameter **Expand Frequency**; see Section 11.2), the feasibility tolerance actually used by the routine (i.e., the *working* feasibility tolerance) increases from 0.5δ to δ (in steps of $0.5\delta/K$).

At certain stages the following ‘resetting procedure’ is used to remove small constraint infeasibilities. First, all nonbasic variables are moved exactly onto their bounds. A count is kept of the number of non-trivial adjustments made. If the count is non-zero, the basic variables are recomputed. Finally, the working feasibility tolerance is reinitialized to 0.5δ .

If a problem requires more than K iterations, the resetting procedure is invoked and a new cycle of iterations is started. (The decision to resume the feasibility phase or optimality phase is based on comparing any constraint infeasibilities with δ .)

The resetting procedure is also invoked when the routine reaches an apparently optimal, infeasible or unbounded solution, unless this situation has already occurred twice. If any non-trivial adjustments are made, iterations are continued.

The EXPAND procedure not only allows a positive step to be taken at every iteration, but also provides a potential *choice* of constraints to be added to the working set. All constraints at a distance α (where $\alpha \leq \alpha_N$) 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. This strategy helps keep the basis matrix B well-conditioned.

11 Optional Parameters

Several optional parameters in E04NQF define choices in the problem specification or the algorithm logic. In order to reduce the number of formal parameters of E04NQF these optional parameters have associated *default values* that are appropriate for most problems. Therefore, the user need only specify those optional parameters whose values are to be different from their default values.

The remainder of this section can be skipped by users who wish to use the default values for *all* optional parameters. A complete list of optional parameters and their default values is given in Section 11.1.

Optional parameters may be specified by calling one, or any, of the routines E04NRF, E04NSF, E04NTF and E04NUF prior to a call to E04NQF, but after a call to E04NPF.

E04NRF reads options from an external options file, with *Begin* and *End* as the first and last lines respectively and each intermediate line defining a single optional parameter. For example,

```
Begin
  Print Level = 5
End
```

The call

```
CALL E04NRF (IOPTNS, CW, IW, RW, INFORM)
```

can then be used to read the file on unit IOPTNS. INFORM will be zero on successful exit. E04NRF should be consulted for a full description of this method of supplying optional parameters.

E04NSF, E04NTF or E04NUF can be called to supply options directly, one call being necessary for each optional parameter. E04NSF, E04NTF or E04NUF should be consulted for a full description of this method of supplying optional parameters.

All optional parameters not specified by the user are set to their default values. Optional parameters specified by the user are unaltered by E04NQF (unless they define invalid values) and so remain in effect for subsequent calls unless altered by the user.

11.1 Optional parameter checklist and default values

The following list gives the valid options. For each option, we give the keyword, any essential optional qualifiers and the default value. A definition for each option can be found in Section 11.2. The minimum abbreviation of each keyword is underlined. The qualifier may be omitted. The letters *i* and *r* denote INTEGER and *double precision* values required with certain options. The default value of an option is used whenever the condition $|i| \geq 100000000$ is satisfied. The number ϵ is a generic notation for *machine precision* (see X02AJF).

Optional Parameters	Default Values
<u>Backup</u> <u>Basis</u> <u>File</u>	Default = 0
<u>Check</u> <u>Frequency</u>	Default = 60
<u>Crash</u> <u>Option</u>	Default = 3
<u>Crash</u> <u>Tolerance</u>	Default = 0.1
<u>Defaults</u>	
<u>Dump</u> <u>File</u>	Default = 0
<u>Elastic</u> <u>Mode</u>	Default = 1
<u>Elastic</u> <u>Objective</u>	Default = 1
<u>Elastic</u> <u>Weight</u>	Default = 1.0
<u>Expand</u> <u>Frequency</u>	Default = 10000
<u>Factorisation</u> <u>Frequency</u>	Default = 100(LP) or 50(QP)
<u>Feasibility</u> <u>Tolerance</u>	Default = 10^{-6}
<u>Infinite</u> <u>Bound</u> <u>Size</u>	Default = 10^{20}
<u>Insert</u> <u>File</u>	Default = 0
<u>Iteration</u> <u>Limit</u>	Default = $\max(10000, m)$
<u>Iters</u>	
<u>Itns</u>	
<u>List</u>	Default = Nolist
<u>Load</u> <u>File</u>	Default = 0
<u>LU</u> <u>Factor</u> <u>Tolerance</u>	Default = 100.0
<u>LU</u> <u>Singularity</u> <u>Tolerance</u>	Default = 2.010^{-6}
<u>LU</u> <u>Update</u> <u>Tolerance</u>	Default = 10.0
<u>Maximize</u>	Default = Minimize
<u>Minimize</u>	
<u>New</u> <u>Basis</u> <u>File</u>	Default = 0
<u>Nolist</u>	
<u>Optimality</u> <u>Tolerance</u>	Default = 10^{-6}
<u>Old</u> <u>Basis</u> <u>File</u>	Default = 0
<u>Partial</u> <u>Price</u>	Default = 10(LP) or 1(QP)
<u>Pivot</u> <u>Tolerance</u>	Default = $\epsilon^{0.67}$
<u>Print</u> <u>File</u>	Default = 0
<u>Print</u> <u>Frequency</u>	Default = 100
<u>Print</u> <u>Level</u>	Default = 1
<u>Punch</u> <u>File</u>	Default = 0
<u>Save</u> <u>Frequency</u>	Default = 100
<u>Scale</u> <u>Option</u>	Default = 2
<u>Scale</u> <u>Tolerance</u>	Default = 0.9
<u>Solution</u> <u>File</u>	Default = 0
<u>Summary</u> <u>File</u>	Default = 0
<u>Summary</u> <u>Frequency</u>	Default = 100
<u>Superbasics</u> <u>Limit</u>	Default = $\min(500, n_H + 1, n)$

Suppress Parameters**Timing Level**

Default = 0

Unbounded Step SizeDefault = $\max(\text{bigbnd}, 10^{20})$ **11.2 Description of the Optional Parameters****Check Frequency** i

Default = 60

Every i th iteration after the most recent basis factorization, a numerical test is made to see if the current solution (x, s) satisfies the linear constraints $Ax - s = 0$. If the largest element of the residual vector $r = Ax - s$ is judged to be too large, the current basis is refactored and the basic variables recomputed to satisfy the constraints more accurately. If $i < 0$, the default value is used. If $i = 0$, the value $i = 99999999$ is used and effectively no checks are made.

Check Frequency = 1 is useful for debugging purposes, but otherwise this option should not be needed.

Crash Option i

Default = 3

Crash Tolerance r

Default = 0.1

Note that this option does not apply when **START** = 'W' (see Section 5).

If **START** = 'C', an internal Crash procedure is used to select an initial basis from various rows and columns of the constraint matrix $(A \mid -I)$. The value of i determines which rows and columns of A are initially eligible for the basis, and how many times the Crash procedure is called. Columns of $-I$ are used to pad the basis where necessary.

i	Meaning
0	The initial basis contains only slack variables: $B = I$.
1	The Crash procedure is called once, looking for a triangular basis in all rows and columns of the matrix A .
2	The Crash procedure is called once, looking for a triangular basis in rows.
3	The Crash procedure is called twice. The two calls treat linear equalities and linear inequalities separately.

If $i \geq 1$, certain slacks on inequality rows are selected for the basis first. (If $i \geq 2$, numerical values are used to exclude slacks that are close to a bound.) The Crash procedure then makes several passes through the columns of A , searching for a basis matrix that is essentially triangular. A column is assigned to 'pivot' on a particular row if the column contains a suitably large element in a row that has not yet been assigned. (The pivot elements ultimately form the diagonals of the triangular basis.) For remaining unassigned rows, slack variables are inserted to complete the basis.

This value allows the Crash procedure to ignore certain 'small' non-zero elements in each column of A . If $a_{\max}$ is the largest element in column j , other non-zeros a_{ij} in the column are ignored if $|a_{ij}| \leq a_{\max} \times r$. (To be meaningful, r should be in the range $0 \leq r < 1$.)

When $r > 0.0$, the basis obtained by the Crash procedure may not be strictly triangular, but it is likely to be nonsingular and almost triangular. The intention is to obtain a starting basis containing more columns of A and fewer (arbitrary) slacks. A feasible solution may be reached sooner on some problems.

For example, suppose the first m columns of A form the matrix shown under **LU** factor tolerance; i.e., a tridiagonal matrix with entries $-1, 4, -1$. To help the Crash procedure choose all m columns for the initial basis, we would specify Crash tolerance r for some value of $r > \frac{1}{4}$.

Defaults

This special keyword may be used to reset all optional parameters to their default values.

Dump File i_1

Default = 0

Load File i_2

Default = 0

Dump File and **Load File** are similar to **Punch File** and **Insert File**, but they record solution information in a manner that is more direct and more easily modified. A full description of information recorded in **Dump File** and **Load File** is given in Gill *et al.* (1999).

If $i_1 > 0$, the last solution obtained will be output to the file with unit number i .

If $i_2 > 0$, the **Load File** containing basis information will be read. The file will usually have been output previously as a **Dump File**. The file will not be accessed if an **Old Basis File** or an **Insert File** is specified.

Elastic Mode

 i

Default = 1

This parameter determines if (and when) elastic mode is to be started. Three elastic modes are available as follows:

- | i | Meaning |
|-----|---|
| 0 | Elastic mode is never invoked. E04NQF will terminate as soon as infeasibility is detected. There may be other points with significantly smaller sums of infeasibilities. |
| 1 | Elastic mode is invoked only if the constraints are found to be infeasible (the default). If the constraints are infeasible, continue in elastic mode with the composite objective determined by the values of Elastic Objective and Elastic Weight . |
| 2 | The iterations start and remain in elastic mode. This option allows you to minimize the composite objective function directly without first performing Phase 1 iterations. |

The success of this option will depend critically on your choice of **Elastic Weight**. If **Elastic Weight** is sufficiently large and the constraints are feasible, the minimizer of the composite objective and the solution of the original problem are identical. However, if the **Elastic Weight** is not sufficiently large, the minimizer of the composite function may be infeasible, even though a feasible point for the constraints may exist.

Elastic Objective

 i

Default = 1

This option determines the form of the composite objective. Three types of composite objectives are available.

- | i | Meaning |
|-----|--|
| 0 | Include only the true objective $f(x)$ in the composite objective. This option sets $\gamma = 0$ in the composite objective and will allow E04NQF to ignore the elastic bounds and find a solution that minimizes f subject to the nonelastic constraints. This option is useful if there are some ‘soft’ constraints that you would like to ignore if the constraints are infeasible. |
| 1 | Use a composite objective defined with γ determined by the value of Elastic Weight . This value is intended to be used in conjunction with Elastic Mode = 2. |
| 2 | Include only the elastic variables in the composite objective. The elastics are weighted by $\gamma = 1$. This choice minimizes the violations of the elastic variables at the expense of possibly increasing the true objective. This option can be used to find a point that minimizes the sum of the violations of a subset of constraints determined by the parameter HELAST. |

Elastic Weight

 r

Default = 1.0

This keyword defines the value of γ in the composite objective.

At each iteration of elastic mode, the composite objective is defined to be

$$\text{minimize } \sigma f(x) + \gamma (\text{sum of infeasibilities});$$

where $\sigma = 1$ for **Minimize**, $\sigma = -1$ for **Maximize**, and f is the current objective value.

Note that the effect of γ is *not* disabled once a feasible iterate is obtained.

Expand Frequency

 i

Default = 10000

This option is part of an anti-cycling procedure (see Section 10.4) designed to allow progress even on highly degenerate problems.

The strategy is to force a positive step at every iteration, at the expense of violating the constraints by a small amount. Suppose that the value of the optional parameter **Feasibility Tolerance** is δ . Over a period of i iterations, the feasibility tolerance actually used by E04NQF (i.e., the *working* feasibility tolerance) increases from 0.5δ to δ (in steps of $0.5\delta/i$).

Increasing the value of i helps reduce the number of slightly infeasible nonbasic variables (most of which are eliminated during the resetting procedure). However, it also diminishes the freedom to choose a large pivot element (see **Pivot Tolerance** below).

If $i < 0$, the default value is used. If $i = 0$, the value $i = 99999999$ is used and effectively no anti-cycling procedure is invoked.

Factorisation Frequency i Default = 100(LP) or 50(QP)

If $i > 0$, at most i basis changes will occur between factorizations of the basis matrix. For LP problems, the basis factors are usually updated at every iteration. Higher values of i may be more efficient on problems that are extremely sparse and well scaled. For QP problems, fewer basis updates will occur as the solution is approached. The number of iterations between basis factorizations will therefore increase. During these iterations a test is made regularly according to the value of **Check Frequency** (see above) to ensure that the linear constraints $Ax - s = 0$ are satisfied. If necessary, the basis will be refactorized before the limit of i updates is reached. If $i \leq 0$, the default value is used.

Feasibility Tolerance r Default = 10^{-6}

A *feasible problem* is one in which all variables satisfy their upper and lower bounds to within the absolute tolerance r . (This includes slack variables. Hence, the general constraints are also satisfied to within r .)

E04NQF attempts to find a feasible solution before optimizing the objective function. If the sum of infeasibilities cannot be reduced to zero, the problem is assumed to be *infeasible*. Let $Sinf$ be the corresponding sum of infeasibilities. If $Sinf$ is quite small, it may be appropriate to raise r by a factor of 10 or 100. Otherwise, some error in the data should be suspected.

Note that if $SINF$ is not small and you have not asked E04NQF to minimize the violations of the elastic variables (i.e., you have not specified **Elastic Objective** = 2, there may be other points that have a *significantly smaller sum of infeasibilities*. E04NQF will not attempt to find the solution that minimizes the sum unless **Elastic Objective** = 2.

If the constraints and variables have been scaled (see **Scale Option** below), then feasibility is defined in terms of the scaled problem (since it is more likely to be meaningful).

Infinite Bound Size r Default = 10^{20}

If $r > 0$, r defines the ‘infinite’ bound $bigbnd$ in the definition of the problem constraints. Any upper bound greater than or equal to $bigbnd$ will be regarded as plus infinity (and similarly any lower bound less than or equal to $-bigbnd$ will be regarded as minus infinity). If $r \leq 0$, the default value is used.

Iteration Limit i Default = $\max(10000, m)$
Iters
Itns

The value of i specifies the maximum number of iterations allowed before termination. Setting $i = 0$ and **Print Level** > 0 means that the workspace needed to start solving the problem will be computed and printed, but no iterations will be performed. If $i < 0$, the default value is used.

List Default = **Nolist**
Nolist

Normally each optional parameter specification is printed as it is supplied. **Nolist** may be used to suppress the printing and **List** may be used to restore printing.

LU Factor Tolerance r_1 Default = 100.0
LU Update Tolerance r_2 Default = 10.0

The values of r_1 and r_2 affect the stability and sparsity of the basis factorization $B = LU$, during refactorization and updates respectively. The lower triangular matrix L is a product of matrices of the form

$$\begin{pmatrix} 1 & \\ \mu & 1 \end{pmatrix}$$

where the multipliers μ will satisfy $|\mu| \leq r_i$. The default values of r_1 and r_2 usually strike a good compromise between stability and sparsity. They must satisfy $r_1, r_2 \geq 1.0$.

For large and relatively dense problems, $r_1 = 10.0$ or 5.0 (say) may give a useful improvement in stability without impairing sparsity to a serious degree.

For certain very regular structures (e.g., band matrices) it may be necessary to reduce r_1 and/or r_2 in order to achieve stability. For example, if the columns of A include a submatrix of the form

$$\begin{pmatrix} 4 & -1 & & & \\ -1 & 4 & -1 & & \\ & -1 & 4 & -1 & \\ & & \cdot & \cdot & \cdot \\ & & & -1 & 4 & -1 \\ & & & & -1 & 4 \end{pmatrix},$$

one should set both r_1 and r_2 to values in the range $1.0 \leq r_i < 4.0$.

LU Singularity Tolerance

r

Default = 2.010^{-6}

If $r > 0$, r defines the singularity tolerance used to guard against ill-conditioned basis matrices. Whenever the basis is refactorized, the diagonal elements of U are tested as follows. If $|u_{jj}| \leq r$ or $|u_{jj}| < r \times \max_i |u_{ij}|$, the j th column of the basis is replaced by the corresponding slack variable. If $r \leq 0$, the default value is used.

Maximize **Minimize**

Default = **Minimize**

This option specifies the required direction of the optimization. It applies to both linear and nonlinear terms (if any) in the objective function. Note that if two problems are the same except that one minimizes $f(x)$ and the other maximizes $-f(x)$, their solutions will be the same but the signs of the dual variables π_i and the reduced gradients d_j (see Section 10.3) will be reversed.

New Basis File

i_1

Default = 0

Backup Basis File

i_2

Default = 0

Save Frequency

i_3

Default = 100

New Basis File and **Backup Basis File** sometimes referred to as basis maps. They contain the most compact representation of the state of each variable. They are intended for restarting the solution of a problem at a point that was reached by an earlier run. For non-trivial problems, it is advisable to save basis maps at the end of a run, in order to restart the run if necessary.

If $i_1 > 0$, a basis map will be saved on file i_1 every i_3 th iteration, where i_3 is the **Save Frequency**. The first record of the file will contain the word PROCEEDING if the run is still in progress. A basis map will also be saved at the end of a run, with some other word indicating the final solution status.

Using $i_2 > 0$, is intended as a safeguard against losing the results of a long run. Suppose that a **New Basis File** is being saved every 100 (**Save Frequency**) iterations, and that E04NQF is about to save such a basis at iteration 2000. It is conceivable that the run may be interrupted during the next few milliseconds (in the middle of the save). In this case the basis file will be corrupted and the run will have been essentially wasted.

To eliminate this risk, both a **New Basis File** and a **Backup Basis File** may be specified. The following would be suitable for the above example:

```
Backup Basis File 11
New Basis File 12
```

The current basis will then be saved every 100 iterations, first on file 12 and then immediately on file 11. If the run is interrupted at iteration 2000 during the save on file 12, there will still be a usable basis on file 11 (corresponding to iteration 1900).

Note that a new basis will be saved in **New Basis File** at the end of a run if it terminates normally, but it will not be saved in **Backup Basis File**. In the above example, if an optimum solution is found at iteration

2050 (or if the iteration limit is 2050), the final basis on file 12 will correspond to iteration 2050, but the last basis saved on file 11 will be the one for iteration 2000.

A full description of information recorded in **New Basis File** and **Backup Basis File** is given in Gill *et al.* (1999).

Old Basis File

 i

Default = 0

If $i > 0$, the basis maps information will be obtained from this file. A full description of information recorded in **New Basis File** and **Backup Basis File** is given in Gill *et al.* (1999). The file will usually have been output previously as a **New Basis File** or **Backup Basis File**.

The file will not be acceptable if the number of rows or columns in the problem has been altered.

Optimality Tolerance

 r Default = 10^{-6}

This is used to judge the size of the reduced gradients $d_j = g_j - \pi a_j$, where g_j is the j th component of the gradient, a_j is the associated column of the constraint matrix $(A \ -I)$, and π is the set of dual variables.

By construction, the reduced gradients for basic variables are always zero. The problem will be declared optimal if the reduced gradients for nonbasic variables at their lower or upper bounds satisfy

$$d_j/\|\pi\| \geq -r \quad \text{or} \quad d_j/\|\pi\| \leq r$$

respectively, and if $|d_j|/\|\pi\| \leq r$ for superbasic variables.

In the above tests, $\|\pi\|$ is a measure of the size of the dual variables. It is included to make the tests independent of a scale factor on the objective function.

The quantity $\|\pi\|$ actually used is defined by

$$\|\pi\| = \max\{\sigma\sqrt{m}, 1\}, \quad \text{where } \sigma = \sum_{i=1}^m |\pi_i|j$$

so that only large scale factors are allowed for.

If the objective is scaled down to be very *small*, the optimality test reduces to comparing d_j against $0.01r$.

Partial Price

 i Default = $10(LP)$ or $1(QP)$

This option is recommended for large FP or LP problems that have significantly more variables than constraints (i.e., $n \gg m$). It reduces the work required for each pricing operation (i.e., when a nonbasic variable is selected to enter the basis). If $i = 1$, all columns of the constraint matrix $(A \ -I)$ are searched. If $i > 1$, A and I are partitioned to give i roughly equal segments A_j, K_j , for $j = 1, 2, \dots, p$ (modulo p). If the previous pricing search was successful on A_{j-1}, K_{j-1} , the next search begins on the segments A_j, K_j . If a reduced gradient is found that is larger than some dynamic tolerance, the variable with the largest such reduced gradient (of appropriate sign) is selected to enter the basis. If nothing is found, the search continues on the next segments A_{j+1}, K_{j+1} , and so on. If $i \leq 0$, the default value is used.

Pivot Tolerance

 r Default = $\epsilon^{0.67}$

Broadly speaking, the pivot tolerance is used to prevent columns entering the basis if they would cause the basis to become almost singular.

When x changes to $x + \alpha p$ for some search direction p , a ‘ratio test’ is used to determine which component of x reaches an upper or lower bound first. The corresponding element of p is called the pivot element.

For linear problems, elements of p are ignored (and therefore cannot be pivot elements) if they are smaller than the pivot tolerance r .

It is common for two or more variables to reach a bound at essentially the same time. In such cases, the **Feasibility Tolerance** (say t) provides some freedom to maximize the pivot element and thereby improve numerical stability. Excessively small values of t should therefore not be specified.

To a lesser extent, the **Expand Frequency** (say f) also provides some freedom to maximize the pivot element. Excessively *large* values of f should therefore not be specified.

Print File i Default = 0

If $i > 0$, the following information is output to i during the solution of each problem:

- a listing of the optional parameters;
- some statistics about the problem;
- the amount of storage available for the LU factorization of the basis matrix;
- notes about the initial basis resulting from a crash procedure or a Basis File;
- the iteration log;
- basis factorization statistics;
- the exit IFAIL condition and some statistics about the solution obtained;
- the printed solution, if requested.

The last four items are described in Sections 8 and 12. Further brief output may be directed to the **Summary File**.

Print Frequency i Default = 100

If $i > 0$, one line of the iteration log will be printed every i th iteration. A value such as $i = 10$ is suggested for those interested only in the final solution.

Print Level i Default = 1

This controls the amount of printing produced by E04NQF as follows.

- 0 No output except error messages. If you want to suppress all output, set **Print File** = 0.
- = 1 The set of selected options, problem statistics, summary of the scaling procedure, information about the initial basis resulting from a crash or a basis file. a single line of output at each iteration (controlled by **Print Frequency**), and the exit condition with a summary of the final solution.
- ≥ 10 Basis factorization statistics.

Punch File i_1 Default = 0
Insert File i_2 Default = 0

These files provide compatibility with commercial mathematical programming systems. The **Punch File** from a previous run may be used as an **Insert File** for a later run on the same problem. A full description of information recorded in **Insert File** and **Punch File** is given in Gill *et al.* (1999).

If $i_2 > 0$, the final solution obtained will be output to file i_1 . For linear programs, this format is compatible with various commercial systems.

If $i_1 > 0$, the **Insert File** containing basis information will be read. The file will usually have been output previously as a **Punch File**. The file will not be accessed if **Old Basis File** is specified.

Scale Option i Default = 2
Scale Tolerance r Default = 0.9

Three scale options are available as follows:

- | i | Meaning |
|-----|---|
| 0 | No scaling. This is recommended if it is known that x and the constraint matrix never have very large elements (say, larger than 1000). |
| 1 | The constraints and variables are scaled by an iterative procedure that attempts to make the matrix coefficients as close as possible to 1.0 (see Fourer (1982)). This will sometimes improve the performance of the solution procedures. |

- 2 The constraints and variables are scaled by the iterative procedure. Also, a certain additional scaling is performed that may be helpful if the right-hand side b or the solution x is large. This takes into account columns of $(A \ -I)$ that are fixed or have positive lower bounds or negative upper bounds.

Scale Tolerance affects how many passes might be needed through the constraint matrix. On each pass, the scaling procedure computes the ratio of the largest and smallest non-zero coefficients in each column:

$$\rho_j = \max_j |a_{ij}| / \min_i |a_{ij}| \quad (a_{ij} \neq 0).$$

If $\max_j \rho_j$ is less than r times its previous value, another scaling pass is performed to adjust the row and column scales. Raising r from 0.9 to 0.99 (say) usually increases the number of scaling passes through A . At most 10 passes are made.

Solution File i Default = 0

If $i > 0$, the final solution will be output to file i (whether optimal or not). All numbers are printed in 1pe16.6 format.

To see more significant digits in the printed solution, it will sometimes be useful to make i refer to the system **Print File**.

Summary File i_1 Default = 0
Summary Frequency i_2 Default = 100

If $i_1 > 0$, a brief log will be output to file i_1 , including one line of information every i_2 th iteration. In an interactive environment, it is useful to direct this output to the terminal, to allow a run to be monitored on-line. (If something looks wrong, the run can be manually terminated.) Further details are given in Section 12.

Superbasics Limit i Default = $\min(500, n_H + 1, n)$

This places a limit on the storage allocated for superbasic variables. Ideally, i should be set slightly larger than the ‘number of degrees of freedom’ expected at an optimal solution.

For linear programs, an optimum is normally a basic solution with no degrees of freedom. (The number of variables lying strictly between their bounds is no more than m , the number of general constraints.) The default value of i is therefore 1.

For quadratic problems, the number of degrees of freedom is often called the ‘number of independent variables’.

Normally, i need not be greater than $\text{NCOLH} + 1$, where NCOLH is the number of leading non-zero columns of H , n_H .

For many problems, i may be considerably smaller than NCOLH . This will save storage if NCOLH is very large.

Suppress Parameters

Normally E04NQF prints the optional file as it is being read, and then prints a complete list of the available keywords and their final values. The **Suppress Parameters** option tells E04NQF not to print the full list.

Timing Level i Default = 0

If $i > 0$, some timing information will be output to the **Print File**, if it is > 0 .

Unbounded Step Size r Default = $\max(\text{bigbnd}, 10^{20})$

If $r > 0$, r 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 not positive-definite.) If the change in x during an iteration would exceed the value of r , the objective function is considered to be unbounded below in the feasible region. If $r \leq 0$, the default value is used.

12 Description of Monitoring Information

This section describes the intermediate printout and final printout which constitutes the monitoring information produced by E04NQF. (See also the description of the optional parameters **Print File** and **Print Level** in Section 11.2.) The level of printed output can be controlled by the user.

When **Print Level** > 20 and **Print File** > 0, the following lines of intermediate printout (< 120 characters) are produced on the unit number specified by **Print File** whenever the matrix B or $B_S = (B \ S)^T$ is factorized. Gaussian elimination is used to compute an LU factorization of B or B_S , where PLP^T is a lower triangular matrix and PUQ is an upper triangular matrix for some permutation matrices P and Q . The factorization is stabilized in the manner described under the optional parameter **LU Factor Tolerance** (see Section 11.2).

Label	Description														
Factorize	is the factorization count.														
Demand	is a code giving the reason for the present factorization as follows:														
	<table> <tr> <th>Code</th><th>Meaning</th></tr> <tr> <td>0</td><td>First LU factorization.</td></tr> <tr> <td>1</td><td>The number of updates reached the value of the optional parameter Factorization Frequency (see Section 11.2).</td></tr> <tr> <td>2</td><td>The number of non-zeros in the updated factors has increased significantly.</td></tr> <tr> <td>7</td><td>Not enough storage to update factors.</td></tr> <tr> <td>10</td><td>Row residuals too large (see the description for the optional parameter Check Frequency).</td></tr> <tr> <td>11</td><td>Ill-conditioning has caused inconsistent results.</td></tr> </table>	Code	Meaning	0	First LU factorization.	1	The number of updates reached the value of the optional parameter Factorization Frequency (see Section 11.2).	2	The number of non-zeros in the updated factors has increased significantly.	7	Not enough storage to update factors.	10	Row residuals too large (see the description for the optional parameter Check Frequency).	11	Ill-conditioning has caused inconsistent results.
Code	Meaning														
0	First LU factorization.														
1	The number of updates reached the value of the optional parameter Factorization Frequency (see Section 11.2).														
2	The number of non-zeros in the updated factors has increased significantly.														
7	Not enough storage to update factors.														
10	Row residuals too large (see the description for the optional parameter Check Frequency).														
11	Ill-conditioning has caused inconsistent results.														
Iteration	is the iteration count.														
Infeas	the number of infeasibilities at the <i>start</i> of the previous iteration.														
Objective	if Infeas > 0, this is the Sum of Infeasibilities at the <i>start</i> of the previous iteration.														
Nonlinear	is the number of nonlinear variables in the current basis B (not printed if B_S is factorized). If Infeas = 0, this is the value of the objective function after the previous iteration.														
Linear	is the number of linear variables in B (not printed if B_S is factorized).														
Slacks	is the number of slack variables in B (not printed if B_S is factorized).														
Elms	is the number of non-zeros in B (not printed if B_S is factorized)														
Density	is the percentage non-zero density of B (not printed if B_S is factorized). More precisely, $\text{Density} = 100 \times \text{Elms} / (m \times m)$, where m is the number of rows in the problem ($m = \text{Linear} + \text{Slacks}$).														
Compressns	is the number of times the data structure holding the partially factorized matrix needed to be compressed, in order to recover unused workspace. Ideally, it should be zero.														
Merit	is the average Markowitz merit count for the elements chosen to be the diagonals of PUQ . Each merit count is defined to be $(c - 1)(r - 1)$, where c and r are the number of non-zeros in the column and row containing the element at the time it is selected to be the next diagonal. Merit is the average of m such quantities. It gives an indication of how much work was required to preserve sparsity during the factorization.														
lenL	is the number of non-zeros in L .														
lenU	is the number of non-zeros in U .														
Increase	is the percentage increase in the number of non-zeros in L and U relative to the number of non-zeros in B . More precisely, $\text{Increase} = 100 \times (\text{lenL} + \text{lenU} - \text{Elms}) / \text{Elms}$.														

m	is the number of rows in the problem. Note that $m = U_t + L_t + b_p$.
U_t	is the number of triangular rows of B at the top of U .
d1	is the number of columns remaining when the density of the basis matrix being factorized reached 0.3.
Lmax	is the maximum sub-diagonal element in the columns of L . This will not exceed the value of the optional parameter LU Factor Tolerance (see Section 11.2).
Bmax	is the maximum non-zero element in B (not printed if B_S is factorized).
BSmax	is the maximum non-zero element in B_S (not printed if B is factorized).
Umax	is the maximum non-zero element in U , excluding elements of B that remain in U unchanged. (For example, if a slack variable is in the basis, the corresponding row of B will become a row of U without modification. Elements in such rows will not contribute to Umax. If the basis is strictly triangular then <i>none</i> of the elements of B will contribute and Umax will be zero.) Ideally, Umax should not be significantly larger than Bmax. If it is several orders of magnitude larger, it may be advisable to reset the LU Factor Tolerance to some value nearer unity. Umax is not printed if B_S is factorized.
Umin	is the magnitude of the smallest diagonal element of PUQ (not printed if B_S is factorized).
Growth	is the value of the ratio $U_{\max}/B_{\max}$, which should not be too large. Providing Lmax is not large (say < 10.0), the ratio $\max(B_{\max}, U_{\max})/U_{\min}$ is an estimate of the condition number of B . If this number is extremely large, the basis is nearly singular and some numerical difficulties might occur. (However, an effort is made to avoid near-singularity by using slacks to replace columns of B that would have made Umin extremely small and the modified basis is refactorized.) Growth is not printed if B_S is factorized.
L_t	is the number of triangular columns of B at the left of L .
b_p	is the size of the ‘bump’ or block to be factorized nontrivially after the triangular rows and columns of B have been removed.
d2	is the number of columns remaining when the density of the basis matrix being factorized has reached 0.6.

When **Print Level** > 20 and **Print File** > 0 , the following lines of intermediate printout (< 120 characters) are produced on the unit number specified by **Print File** whenever **START** = 'C' (see Section 5). They refer to the number of columns selected by the Crash procedure during each of several passes through A , whilst searching for a triangular basis matrix.

Label	Description
Slacks	is the number of slacks selected initially.
Free cols	is the number of free columns in the basis, including those whose bounds are rather far apart.
Preferred	is the number of ‘preferred’ columns in the basis (i.e., $HS(j) = 3$ for some $j \leq n$). It will be a subset of the columns for which $HS(j) = 3$ was specified.
Unit	is the number of unit columns in the basis.
Double	is the number of double columns in the basis.
Triangle	is the number of triangular columns in the basis.
Pad	is the number of slacks used to pad the basis (to make it a non-singular triangle).

When **Print Level** > 20 and **Print File** > 0, the following lines of intermediate printout (< 80 characters) are produced on the unit number specified by **Print File**. They refer to the elements of the NAMES array (see Section 5).

Label	Description
Name	gives the name for the problem (blank if none).
Status	gives the exit status for the problem (i.e., Optimal soln, Weak soln, Unbounded, Infeasible, Excess itns, Error condn or Feasble soln) followed by details of the direction of the optimization (i.e., (Min) or (Max)).
Objective	gives the name of the free row for the problem (blank if none).
RHS	gives the name of the constraint right-hand side for the problem (blank if none).
Ranges	gives the name of the ranges for the problem (blank if none).
Bounds	gives the name of the bounds for the problem (blank if none).

At the end of a run, the final solution will be output to the **Print File**. Some header information appears first to identify the problem and the final state of the optimization procedure. A ROWS section and a COLUMNS section then follow, giving one line of information for each row and column.

The ROWS section

The general constraints take the form $l \leq Ax \leq u$. The i th constraint is therefore of the

$$\infty \leq \nu_i^T x \leq \beta,$$

where ν_i is the i th row of A .

Internally, the constraints take the form $Ax - s = 0$, where s is the set of slack variables (which happen to satisfy the bounds $l \leq s \leq u$). For the i th constraint it is the slack variable s_i that is directly available, and it is sometimes convenient to refer to its state. A ‘.’ is printed for any numerical value that is exactly zero.

Label	Description
Number	is the value of $n + i$. (This is used internally to refer to s_i in the intermediate output.)
Row	gives the name of ν_i .
State	the state of ν_i (the state of s_i relative to the bounds α and β . The various states possible are as follows: LL s_i is nonbasic at its lower limit, α . UL s_i is nonbasic at its upper limit, β . EQ s_i is nonbasic and fixed at the value $\alpha = \beta$. FR s_i is nonbasic and currently zero, even though it is free to take any value between its bounds α and β . BS s_i is basic. SBS s_i is superbasic. A key is sometimes printed before State to give some additional information about the state of a variable. Note that unless the optional parameter Scale Option = 0 (see Section 11.2) is specified, the tests for assigning a key are applied to the variables of the scaled problem. A <i>Alternative optimum possible</i> . The variable is nonbasic, but its reduced gradient is essentially zero. This means that if the variable were allowed to start moving away from its bound, there would be no change in the value of the objective function. The values of the other free variables <i>might</i> 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 basic or superbasic, but it is equal to (or very close to) one of its bounds.
- I *Infeasible.* The variable is basic or superbasic and is currently violating one of its bounds by more than the value of the optional parameter **Feasibility Tolerance** (see Section 11.2).
- N *Not precisely optimal.* The variable is nonbasic or superbasic. If the value of the reduced gradient for the variable exceeds the value of the optional parameter **Optimality Tolerance** (see Section 11.2), the solution would not be declared optimal because the reduced gradient for the variable would not be considered negligible.

Activity	is the value of v_i at the final iterate (the i th element of $A^T x$).
Slack Activity	is the value by which the row differs from its nearest bound. (For the free row (if any), it is set to Activity.)
Lower Limit	is α , the lower bound specified for the variable s_i . None indicates that $BL(j) \leq -bigbnd$.
Upper Bound	is β , the upper bound specified for the variable s_i . None indicates that $BU(j) \geq bigbnd$.
Dual Activity	is the value of the dual variable π_i (the Lagrange multiplier for v_i ; see Section 10.3). For FP problems, π_i is set to zero.
i	gives the index i of the i th row.

The COLUMNS section

Let the j th component of x be the variable x_j and assume that it satisfies the bounds $\alpha \leq x_j \leq \beta$. A ‘.’ is printed for any numerical value that is exactly zero.

Label	Description
Number	is the column number j . (This is used internally to refer to x_j in the intermediate output.)
Column	gives the name of x_j .
State	the state of x_j relative to the bounds α and β . The various states possible are as follows:
LL	x_j is nonbasic at its lower limit, α .
UL	x_j is nonbasic at its upper limit, β .
EQ	x_j is nonbasic and fixed at the value $\alpha = \beta$.
FR	x_j is nonbasic and currently zero, even though it is free to take any value between its bounds α and β .
BS	x_j is basic.
SBS	x_j is superbasic.

A key is sometimes printed before State to give some additional information about the state of a variable. Note that unless the optional parameter **Scale Option** = 0 (see Section 11.2) is specified, the tests for assigning a key are applied to the variables of the scaled problem.

- A *Alternative optimum possible.* The variable is nonbasic, but its reduced gradient is essentially zero. This means that if the variable were allowed to start moving away from its bound, there would be no change in the value of 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 basic or superbasic, but it is equal to (or very close to) one of its bounds.
- I *Infeasible.* The variable is basic or superbasic and is currently violating one of its bounds by more than the value of the optional parameter **Feasibility Tolerance** (see Section 11.2).
- N *Not precisely optimal.* The variable is nonbasic or superbasic. If the value of the reduced gradient for the variable exceeds the value of the optional parameter **Optimality Tolerance** (see Section 11.2), the solution would not be declared optimal because the reduced gradient for the variable would not be considered negligible.

Activity	is the value of x_j at the final iterate.
Obj Gradient	is the value of g_j at the final iterate. For FP problems, g_j is set to zero.
Lower Bound	is the lower bound specified for the variable. None indicates that $BL(j) \leq -bigbnd$.
Upper Bound	is the upper bound specified for the variable. None indicates that $BU(j) \geq bigbnd$.
Reduced Gradient	is the value of d_j at the final iterate (see Section 10.3). For FP problems, d_j is set to zero.
$m + j$	is the value of $m + j$.

Note: if two problems are the same except that one minimizes $f(x)$ and the other maximizes $-f(x)$, their solutions will be the same but the signs of the dual variables π_i and the reduced gradients d_j will be reversed.

The SOLUTION file

If a positive **Solution File** is specified, the information contained in a printed solution may also be output to the relevant file (which may be the PRINT file if so desired). Infinite Upper and Lower limits appear as 10^{20} rather than None. Other real values are output with format 1pe16.6. Again, the maximum record length is 111 characters, including what would be the carriage-control character if the file were printed.

A SOLUTION file is intended to be read from disk by a self-contained program that extracts and saves certain values as required for possible further computation. Typically the first 14 records would be ignored. Each subsequent record may be read using

```
FORMAT (i8, 2x, 2a4, 1x, a1, 1x, a3, 5e16.6, i7)
```

adapted to suit the occasion. The end of the ROWS section is marked by a record that starts with a 1 and is otherwise blank. If this and the next 4 records are skipped, the COLUMNS section can then be read under the same format. (There should be no need to use any BACKSPACE statements.)

The SUMMARY file

If **Summary File** f is specified with $f > 0$, certain brief information will be output to file f . When E04NQF is run interactively, file f will usually be the terminal. For batch jobs a disk file should be used, to retain a concise log of each run if desired. (A **Summary File** is more easily perused than the associated **Print File**).

A **Summary File** (like the **Print File**) is not rewound after a problem has been processed. The maximum record length is 72 characters, including a carriage-control character in column 1.

The following information is included:

1. The Begin card from the optional parameters file, if used (see ISPECS);
2. The basis file loaded, if any;

3. The status of the solution after each basis factorization (whether feasible; the objective value; the number of function calls so far);
4. The same information every k th iteration, where k is the specified **Summary Frequency** (see Section 11.2);
5. Warnings and error messages;
6. The exit condition and a summary of the final solution.

Item 4 is preceded by a blank line, but item 5 is not.

The meaning of the printout for linear constraints is the same as that given above for variables, with ‘variable’ replaced by ‘constraint’, n replaced by m , $\text{NAMES}(j)$ replaced by $\text{NAMES}(n + j)$, $\text{BL}(j)$ and $\text{BU}(j)$ are replaced by $\text{BL}(n + j)$ and $\text{BU}(n + j)$ respectively, and with the following change in the heading:

`Constrnt` gives the name 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.
