

NAG Toolbox for MATLAB

e04nq

1 Purpose

e04nq solves sparse linear programming or convex quadratic programming problems. The initialization function e04np **must** have been called before calling e04nq.

2 Syntax

```
[hs, x, pi, rc, ns, ninf, sinf, obj, user, cw, iw, rw, ifail] =
e04nq(start, qphx, m, n, lenc, ncolh, iobj, objadd, prob, acol, inda,
loca, bl, bu, c, names, helast, hs, x, ns, cw, iw, rw, 'ne', ne,
'nname', nname, 'user', user, 'lencw', lencw, 'leniw', leniw, 'lenrw',
lenrw)
```

Before calling e04nq or one of the option setting functions e04ns, e04nt or e04nu, e04np **must** be called.

3 Description

e04nq is designed to solve large-scale *linear* or *quadratic programming problems* of the form:

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

where x is an n -vector of variables, l and u are constant lower and upper bounds, A is an m by n 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** 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). e04nq is suitable for all *convex* quadratic programs. The defining feature of a *convex* QP is that the matrix H must be *positive semidefinite*, i.e., it must satisfy $x^T H x \geq 0$ for all x . If not, $f(x)$ is nonconvex and e04nq will terminate with the error indicator **ifail** = 11. If $f(x)$ is nonconvex it may be more appropriate to call e04vh 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 semidefinite

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 description of the optional parameter **Elastic Mode**).

e04nq 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 e04nq by means of the three array parameters **acol**, **inda** and **loca**. This allows you to specify the pattern of nonzero elements in A .

e04nq exploits structure in H by requiring H to be defined *implicitly* in a (sub)program 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 e04mf, e04nc or e04nf.

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, e04nq 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:

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

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 a two-phase iterative procedure (in which the general constraints are satisfied throughout): a *feasibility phase* (Phase 1), in which the sum of infeasibilities with respect to the bounds on x and s is minimized to find a feasible point that satisfies all constraints within a specified feasibility tolerance; and an *optimality phase* (Phase 2), in which $f(x)$ is minimized (or maximized) by constructing a sequence of iterates that lies within the feasible region.

Phase 1 involves solving a linear program of the form

Phase 1

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

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. Then 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 you have the option of either terminating or continuing in so-called *elastic mode* (see the discussion of the

optional parameter **Elastic Mode**). 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. You are 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, e04nq 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) $\underset{x,s,v,w}{\text{minimize}} f(x) + \gamma \sum_{j=1}^{n+m} (v_j + w_j) \quad \text{subject to } Ax - s = 0, \quad l \leq \begin{pmatrix} x \\ s \end{pmatrix} - v + w \leq u, \quad v \geq 0, \quad w \geq 0,$

where γ is a nonnegative parameter known as the *elastic weight* (see the description of the optional parameter **Elastic Weight**), 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 the original objective $f(x)$, the sum of infeasibilities, or anything in-between. If $\gamma = 0$, e04nq will attempt to minimize f subject to the (true) upper and lower bounds on the non-elastic variables (and declare the problem infeasible if the non-elastic 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 nonsingular. 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**).

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 $(A \ -I)$. (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.

e04nq is based on SQOPT, which is part of the SNOPT package described in Gill *et al.* 2005. 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 2005 Users’ guide for SQOPT 7: a Fortran package for large-scale linear and quadratic programming *Report NA 05-1* Department of Mathematics, University of California, San Diego URL: <ftp://www.cam.ucsd.edu/pub/peg/reports/sqdoc7.pdf>
- Gill P E, Murray W and Saunders M A 2005 Users’ guide for SNOPT 7.1: a Fortran package for large-scale linear nonlinear programming *Report NA 05-2* Department of Mathematics, University of California, San Diego URL: <ftp://www.cam.ucsd.edu/pub/peg/reports/sndoc7.pdf>
- 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 .

5.1 Compulsory Input Parameters

1: **start – string**

Indicates how a starting basis (and certain other items) will be obtained.

start = 'C'

Requests that an internal Crash procedure be used to choose an initial basis, unless a Basis file is provided via optional parameters **Old Basis File**, **Insert File** or **Load File**.

start = 'B'

Is the same as **start** = 'C' but is more meaningful when a Basis file is given.

start = 'W'

Means that a basis is already defined in **hs** and a start point is already defined in **x** (probably from an earlier call).

Constraint: **start** = 'B', 'C' or 'W'.

2: **qphx** – string containing name of m-file

For QP problems, you must supply a version of **qphx** to compute the matrix product Hx for a given vector x . If H has rows and columns of zeros, it is most efficient to order x so that the nonlinear variables appear first. For example, if $x = (y, z)^T$ and only y enters the objective quadratically then

$$Hx = \begin{pmatrix} H_1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} y \\ z \end{pmatrix} = \begin{pmatrix} H_1 y \\ 0 \end{pmatrix}. \quad (2)$$

In this case, **ncolh** should be the dimension of y , and **qphx** should compute $H_1 y$. For FP and LP problems, **qphx** will never be called by e04nq and hence **qphx** may be the string 'e04nsh'.

Its specification is:

```
[hx, user] = qphx(ncolh, x, nstate, user)
```

Input Parameters

1: **ncolh** – int32 scalar

This is the same parameter **ncolh** as supplied to e04nq.

2: **x(ncolh)** – double array

The first **ncolh** elements of the vector x .

3: **nstate** – int32 scalar

Allows you to save computation time if certain data must be read or calculated only once. To preserve this data for a subsequent calculation place it in one of **user**, **user** or **user**.

nstate = 1

e04nq is calling **qphx** for the first time.

nstate = 0

There is nothing special about the current call of **qphx**.

nstate ≥ 2

e04nq is calling **qphx** for the last time. This parameter setting allows you to perform some additional computation on the final solution.

nstate = 2

The current x is optimal.

nstate = 3

The problem appears to be infeasible.

	nstate = 4
	The problem appears to be unbounded.
	nstate = 5
	The iterations limit was reached.
4:	user – Any MATLAB object
	Output Parameters
1:	hx(ncolh) – double array
	The product Hx . If ncolh is less than the input parameter n , Hx is really the product H_1y in (2).
2:	user – Any MATLAB object

3: **m** – int32 scalar

m , the number of general linear constraints (or slacks). This is the number of rows in the linear constraint matrix A , including the free row (if any; see **iobj**). Note that A must have at least one row. If your problem has no constraints, or only upper or lower bounds on the variables, then you must include a dummy row with sufficiently wide upper and lower bounds (see also **acol**, **inda** and **loca**).

Constraint: $m \geq 1$.

4: **n** – int32 scalar

n , the number of variables (excluding slacks). This is the number of columns in the linear constraint matrix A .

Constraint: $n \geq 1$.

5: **lenc** – int32 scalar

the number of elements in the constant objective vector c .

If **lenc** > 0, the first **lenc** elements of x belong to variables corresponding to the constant objective term c .

Constraint: $0 \leq \text{lenc} \leq n$.

6: **ncolh** – int32 scalar

n_H , the number of leading nonzero columns of the Hessian matrix H . For FP and LP problems, **ncolh** must be set to zero.

The first **ncolh** elements of x belong to variables corresponding to the nonzero block of the QP Hessian.

Constraint: $0 \leq \text{ncolh} \leq n$.

7: **iobj** – int32 scalar

If **iobj** > 0, row **iobj** of A is a free row containing the nonzero 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: $0 \leq \text{iobj} \leq m$.

8: **objadd – double scalar**

The constant q , to be added to the objective for printing purposes. Typically **objadd** = 0.0D0.

9: **prob – string**

The name for the problem. It is used in the printed solution and in some functions that output Basis files. A blank name may be used.

10: **acol(ne) – double array**

The nonzero elements of A , ordered by increasing column index. Note that all elements must be assigned a value in the calling program.

11: **inda(ne) – int32 array**

inda(i) must contain the row index of the nonzero element stored in **acol**(i), for $i = 1, 2, \dots, \mathbf{ne}$. Thus a pair of values (**acol**(i), **inda**(i)) contains a matrix element and its corresponding row index.

Note that the row indices for a column may be supplied in any order.

Constraint: $1 \leq \mathbf{inda}(i) \leq \mathbf{m}$, for $i = 1, 2, \dots, \mathbf{ne}$.

12: **loca(n + 1) – int32 array**

loca(j) must contain the index in **acol** and **inda** of the start of the j th column, for $j = 1, 2, \dots, \mathbf{n}$. Thus for $j = 1 : \mathbf{n}$, the entries of column j are held in **acol**($k : l$) and their corresponding row indices are in **inda**($k : l$), where $k = \mathbf{loca}(j)$ and $l = \mathbf{loca}(j + 1) - 1$. 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 **loca**(1) = 1 and **loca**($\mathbf{n} + 1$) = $\mathbf{ne} + 1$. 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 $\mathbf{ne} = 1$, **acol**(1) = 0.0, **inda**(1) = 1, **loca**(1) = 1, and **loca**(j) = 2, for $j = 2 : \mathbf{n} + 1$. This row is made ‘free’ by setting its bounds to be **bl**($\mathbf{n} + 1$) = $-\mathbf{infbnd}$ and **bu**($\mathbf{n} + 1$) = \mathbf{infbnd} , where \mathbf{infbnd} is the value of the optional parameter **Infinite Bound Size**.

Constraints:

$$\begin{aligned} \mathbf{loca}(1) &= 1; \\ \mathbf{loca}(j) &\geq 1, \text{ for } j = 2, 3, \dots, \mathbf{n}; \\ \mathbf{loca}(\mathbf{n} + 1) &= \mathbf{ne} + 1; \\ 0 &\leq \mathbf{loca}(j + 1) - \mathbf{loca}(j) \leq \mathbf{m}, \text{ for } j = 1, 2, \dots, \mathbf{n}. \end{aligned}$$

13: **bl(n + m) – double array**

l , the lower bounds for all the variables and general constraints, in the following order. The first \mathbf{n} elements of **bl** must contain the bounds on the variables x , and the next \mathbf{m} elements the bounds for the general linear constraints Ax (which, equivalently, are the bounds for the slacks, s) and the free row (if any). To fix the j th variable, set **bl**(j) = **bu**(j) = β , say, where $|\beta| < \mathbf{infbnd}$. To specify a nonexistent lower bound (i.e., $l_j = -\infty$), set **bl**(j) $\leq -\mathbf{infbnd}$. Here, \mathbf{infbnd} is the value of the optional parameter **Infinite Bound Size**. To specify the j th constraint as an *equality*, set **bl**($\mathbf{n} + j$) = **bu**($\mathbf{n} + j$) = β , say, where $|\beta| < \mathbf{infbnd}$. Note that the lower bound corresponding to the free row must be set to $-\infty$ and stored in **bl**($\mathbf{n} + \mathbf{iobj}$).

Constraint: if $\mathbf{iobj} > 0$, **bl**($\mathbf{n} + \mathbf{iobj}$) $\leq -\mathbf{infbnd}$.

(See also the description for **bu**.)

14: **bu(n + m) – double array**

u , the upper bounds for all the variables and general constraints, in the following order. The first \mathbf{n} elements of **bu** must contain the bounds on the variables x , and the next \mathbf{m} elements the bounds for the general linear constraints Ax (which, equivalently, are the bounds for the slacks, s) and the free row (if any). To specify a nonexistent upper bound (i.e., $u_j = +\infty$), set **bu**(j) $\geq \mathbf{infbnd}$. Note that the upper bound corresponding to the free row must be set to $+\infty$ and stored in **bu**($\mathbf{n} + \mathbf{iobj}$).

Constraints:

if $\mathbf{iobj} > 0$, $\mathbf{bu}(\mathbf{n} + \mathbf{iobj}) \geq \mathbf{infbnd}$;
 $\mathbf{bl}(i) \leq \mathbf{bu}(i)$ otherwise.

15: **c(*)** – double array

Note: the dimension of the array **c** must be at least $\max(1, \mathbf{lenc})$ if $\mathbf{iobj} \neq 0$, and at least 1 otherwise.

Contains the explicit objective vector *c* (if any). If the problem is of type FP, or if $\mathbf{lenc} = 0$, then **c** is not referenced. (In that case, **c** may be dimensioned (1), or it could be any convenient array.)

16: **names(nname)** – string array

The optional column and row names, respectively.

If $\mathbf{nname} = 1$, **names** is not referenced and the printed output will use default names for the columns and rows.

If $\mathbf{nname} = \mathbf{n} + \mathbf{m}$, 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 **names(n + iobj)**.

17: **helast(n + m)** – int32 array

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 <i>j</i> is non-elastic and cannot be infeasible
1	Variable <i>j</i> can violate its lower bound
2	Variable <i>j</i> can violate its upper bound
3	Variable <i>j</i> 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** \neq 0, for $j = 1, 2, \dots, \mathbf{n} + \mathbf{m}$.

18: **hs(n + m)** – int32 array

If **start** = 'C' or 'B', and a Basis file of some sort is to be input (see the description of the optional parameters **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 <i>x(j)</i> 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 $j = 1, 2, \dots, \mathbf{n}$. 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)** \leq **x(j)** \leq **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 e04nq has been called previously with the same values of **n** and **m**, **hs** already contains satisfactory information.

Constraints:

if **start** = 'C', $0 \leq \mathbf{hs}(j) \leq 5$, for $j = 1, 2, \dots, \mathbf{n}$;
 if **start** = 'W', $0 \leq \mathbf{hs}(j) \leq 3$, for $j = 1, 2, \dots, \mathbf{n} + \mathbf{m}$.

19: **x(n + m) – double array**

The initial values of the variables x , and, if **start** = 'W', the slacks s , i.e., (x, s) . (See the description for parameter **hs**.)

20: **ns – int32 scalar**

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.

21: **cw(len cw) – string array**

Constraint: **len cw** ≥ 600 .

22: **iw(len iw) – int32 array**

Constraint: **len iw** ≥ 600 .

23: **rw(len rw) – double array**

Constraint: **len rw** ≥ 600 .

5.2 Optional Input Parameters

1: **ne – int32 scalar**

Default: The dimension of the arrays **acol**, **inda**. (An error is raised if these dimensions are not equal.)

the number of nonzero elements in A .

Constraint: $1 \leq \mathbf{ne} \leq \mathbf{n} \times \mathbf{m}$.

2: **nname – int32 scalar**

Default: The dimension of the array **names**.

the number of column (i.e., variable) and row names supplied in the array **names**.

nname = 1

There are no names. Default names will be used in the printed output.

nname = **n** + **m**

All names must be supplied.

Constraint: **nname** = 1 or **n** + **m**.

3: **user – Any MATLAB object**

user is not used by e04nq, but is passed to **qphx**. Note that for large objects it may be more efficient to use a global variable which is accessible from the m-files than to use **user**.

4: **lencw – int32 scalar**

Default: The dimension of the array **cw**.

Constraint: **lencw** \geq 600.

5: **leniw – int32 scalar**

Default: The dimension of the array **iw**.

Constraint: **leniw** \geq 600.

6: **lenrw – int32 scalar**

Default: The dimension of the array **rw**.

Constraint: **lenrw** \geq 600.

5.3 Input Parameters Omitted from the MATLAB Interface

None.

5.4 Output Parameters1: **hs(n + m) – int32 array**

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 $\mathbf{x}(j)$
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**. Note that unless the optional parameter **Scale Option** = 0 is specified, the optional parameter **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 optional parameter **Feasibility Tolerance**, and there may be some nonbasic variables for which $\mathbf{x}(j)$ 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).

2: **x(n + m) – double array**

The final values of the variables and slacks (x, s) .

3: **pi(m) – double array**

Contains the dual variables π (a set of Lagrange multipliers (shadow prices) for the general constraints).

4: **rc(n + m) – double array**

Contains the reduced costs, $g - (A \quad -I)^T \pi$. The vector g is the gradient of the objective if \mathbf{x} is feasible, otherwise it is the gradient of the Phase 1 objective. In the former case, $g(i) = 0$, for $i = \mathbf{n} + 1 : \mathbf{m}$, hence **rc**($\mathbf{n} + 1 : \mathbf{m}$) = π .

5: **ns – int32 scalar**

The final number of superbasics. This will be zero for FP and LP problems.

6: **ninf – int32 scalar**

The number of infeasibilities.

7: **sinf – double scalar**

The sum of the scaled infeasibilities. This will be zero if **ninf** = 0, and is most meaningful when **Scale Option** = 0.

8: **obj – double scalar**

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.

Note that **obj** does not include contributions from the constant term **objadd** or the objective row, if any.

9: **user – Any MATLAB object**

user is not used by e04nq, but is passed to **qphx**. Note that for large objects it may be more efficient to use a global variable which is accessible from the m-files than to use **user**.

10: **cw(lenw) – string array**11: **iw(leniw) – int32 array**12: **rw(lenrw) – double array**13: **ifail – int32 scalar**

0 unless the function detects an error (see Section 6).

6 Error Indicators and Warnings

Note: e04nq may return useful information for one or more of the following detected errors or warnings.

ifail = 1

The initialization function e04np has not been called or at least one of **lenw**, **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 QP solution found. The final x is not unique.

This exit will occur when the following are true:

- (i) the problem is feasible;
- (ii) the reduced gradient is negligible;
- (iii) the Lagrange multipliers are optimal; and
- (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**.

Feasibility is measured with respect to the upper and lower bounds on the variables and slacks. 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** 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.

If **Elastic Mode** $\neq 0$, e04nq will optimize the QP objective and the sum of infeasibilities, suitably weighted using the optional parameter **Elastic Mode**. The function will tend to determine a ‘good’ infeasible point if the elastic weight is sufficiently large.

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 be increased or decreased by an arbitrary amount without causing a basic variable to violate a bound. Consider adding an upper or lower bound to the variable. Also, examine the constraints that have nonzeros 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 optional parameter **Scale Option**.

ifail = 7

Too many iterations. The value of the optional parameter **Iterations Limit** 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 value of the optional parameter **Superbasics Limit** is too small. The current set of basic and superbasic variables have been optimized as much as possible and a pricing operation is necessary to continue, but there are already **Superbasics Limit** superbasics (and no room for any more).

In general, raise the **Superbasics Limit** s by a reasonable amount, bearing in mind the storage needed for reduced Hessian (see Section 10.2). (The **Reduced Hessian Dimension** h will also increase to s unless specified otherwise, and the associated storage will be about $\frac{1}{2}s^2$ words.) In some cases you may have to set $h < s$ to conserve storage, but beware that the rate of convergence will probably fall off severely.

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** 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 x_B , 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** 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 optional parameter **LU Factor Tolerance** to 2.0 (or possibly even smaller, but not less than 1.0).

ifail = 11

An indefinite matrix was detected during the computation of the reduced Hessian factor R (see Section 10.2). This may be caused by H being indefinite. Check also that (sub)program **qphx** has been coded correctly and that all relevant elements of Hx have been assigned their correct values. If **qphx** is coded correctly and H is positive semidefinite, the failure may be caused by ill conditioning. Try reducing the values of the optional parameters **LU Factor Tolerance** and **LU Update Tolerance**. If there are very large values in H , check the scaling of the variables and constraints.

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 parameter **Print File** and examine the output carefully for further information.

ifail = 15

An unexpected error has occurred. Set the optional parameter **Print File** and examine the output carefully for further information.

7 Accuracy

e04nq 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 j_q . If the problem is purely

linear, variable j_q will usually become basic immediately (unless it should happen to reach its opposite bound and return to the nonbasic set).

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

Label	Description
Itn	is the iteration count.
pp	is the partial-price 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 j th column of the constraint matrix $(A \ -I)$. 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 nonzero, otherwise it becomes nonbasic.
-BS	is the variable removed from the basis to become nonbasic.
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 to exclude very small elements of y from consideration during the computation of $Step$.
nInf	is the number of violated constraints (infeasibilities) before the present iteration. This number will not increase unless iterations are in elastic mode.
sInf	is the sum of infeasibilities before the present iteration. It will usually decrease at each nonzero step, but if $nInf$ decreases by 2 or more, $sInf$ may occasionally increase. However, in elastic mode it will decrease monotonically.
Objective	is the value of the current objective function after the present iteration. Note, if Elastic Mode is 2, the heading is Composite Obj .
L+U	L is the number of nonzeros in the basis factor L . Immediately after a basis factorization $B = LU$, L contains $lenL$ (see Section 12). Further nonzeros are added to L when various columns of B are later replaced. (Thus, L increases monotonically.) U is the number of nonzeros in the basis factor U . Immediately after a basis factorization $B = LU$, U contains $lenU$ (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 superbasic set is non-empty.

Label	Description
rgNorm	is the largest reduced-gradient among the superbasic variables after the current iteration. During the optimality phase, this will be approximately zero after a unit step.
nS	is the current number of superbasic variables.
condHz	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 condHz.

9 Example

e04nq_qphx.m

```
function [hx, user] = qphx(ncolh, x, nstate, user)
    hx = zeros(ncolh, 1);

    hx(1) = 2*x(1);
    hx(2) = 2*x(2);
    hx(3) = 2*(x(3)+x(4));
    hx(4) = hx(3);
    hx(5) = 2*x(5);
    hx(6) = 2*(x(6)+x(7));
    hx(7) = hx(6);
```

```
start = 'C';
m = int32(8);
n = int32(7);
lenc = int32(0);
ncolh = int32(7);
iobj = int32(8);
objadd = 0;
prob = '          ';
acol = [0.02;
        0.02;
        0.03;
        1;
        0.7;
        0.02;
        0.15;
        -200;
        0.06;
        0.75;
        0.03;
        0.04;
        0.05;
        0.04;
        1;
        -2000;
        0.02;
        1;
        0.01;
        0.08;
        0.08;
        0.8;
        -2000;
        1;
        0.12;
        0.02;
        0.02;
```

```

0.75;
0.04;
-2000;
0.01;
0.8;
0.02;
1;
0.02;
0.06;
0.02;
-2000;
1;
0.01;
0.01;
0.97;
0.01;
400;
0.97;
0.03;
1;
400];
inda = [int32(7);
int32(5);
int32(3);
int32(1);
int32(6);
int32(4);
int32(2);
int32(8);
int32(7);
int32(6);
int32(5);
int32(4);
int32(3);
int32(2);
int32(1);
int32(8);
int32(2);
int32(1);
int32(4);
int32(3);
int32(7);
int32(6);
int32(8);
int32(1);
int32(7);
int32(3);
int32(4);
int32(6);
int32(2);
int32(8);
int32(5);
int32(6);
int32(7);
int32(1);
int32(2);
int32(3);
int32(4);
int32(8);
int32(1);
int32(2);
int32(3);
int32(6);
int32(7);
int32(8);
int32(7);
int32(2);
int32(1);
int32(8)];
loca = [int32(1);

```

[illegible]

```

    int32(0)];
x = [0;
     0;
     0;
     0;
     0;
     0;
     0;
     0;
     0;
     0;
     0;
     0;
     0;
     0;
     0;
     0];
ns = int32(0);
[cw, iw, rw, ifail] = e04np();
[hsOut, xOut, pi, rc, nsOut, ninf, sinf, obj, cuser, cwOut, iwOut, rwOut,
 ifail] = ...
    e04nq(start, 'e04nq_qphx', m, n, lenc, ncolh, iobj, objadd, prob,
acol, ...
    inda, loca, bl, bu, c, names, helast, hs, x, ns, cw, iw, rw)

hsOut =
    0
    3
    2
    2
    3
    3
    3
    0
    3
    1
    3
    3
    0
    0
    3

xOut =
    1.0e+06 *
    0
    0.0003
    0.0006
    0.0002
    0.0004
    0.0003
    0.0002
    0.0020
    0.0000
    0.0001
    0.0000
    0.0000
    0.0015
    0.0002
    -2.9887

pi =
    1.0e+04 *
    -1.2901
    0.0000
    -0.2325
    0
    0
    1.4455
    1.4581
    -0.0001

rc =
    1.0e+04 *
    0.2361

```

```

-0.0000
-0.0000
 0.0000
 0.0000
 0.0000
 0.0000
 0.0000
-1.2901
 0.0000
-0.2325
      0
      0
    1.4455
    1.4581
    -0.0001
nsOut =
      2
ninf =
      0
sinf =
      0
obj =
-1.8478e+06
cuser =
      0
cwOut =
    array elided
iwOut =
    array elided
rwOut =
    array elided
ifail =
      0

```

Note: the remainder of this document is intended for more advanced users. Section 10 contains a detailed description of the algorithm which may be needed in order to understand Sections 11 and 12. Section 11 describes the optional parameters which may be set by calls to e04ns, e04nt and/or e04nu. Section 12 describes the quantities which can be requested to monitor the course of the computation.

10 Algorithmic Details

This section contains a detailed description of the method used by e04nq.

10.1 Overview

e04nq 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 function 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 (sub)programs. 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 8.1) to the quadratic objective function (the printed quantity `Objective`; see Section 8.1).

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

$$\begin{pmatrix} \bar{x} \\ \bar{s} \end{pmatrix} = \begin{pmatrix} x \\ s \end{pmatrix} + \alpha p, \quad (3)$$

where the *step length* α is a nonnegative 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**). 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

$$Wp = 0. \quad (4)$$

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 (7) and (8)). The direction p will satisfy (4) if

$$p = Zp_Z, \quad (5)$$

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 $(A \ -I)$ 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 $(A \ -I)$ by (conceptually) partitioning the constraints $Ax - s = 0$ so that

$$Bx_B + Sx_S + Nx_N = 0, \quad (6)$$

where B is a square nonsingular basis and x_B , x_S and x_N are the basic, superbasic 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 (6).

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

$$WP = \begin{pmatrix} B & S & N \\ 0 & 0 & I_N \end{pmatrix}, \quad (7)$$

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 as the *BSN* partition changes. Given the permutation P , the null space basis is given by

$$Z = P \begin{pmatrix} -B^{-1}S \\ I \\ 0 \end{pmatrix}. \quad (8)$$

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:

$$g_Z = Z^T g \quad \text{and} \quad H_Z = Z^T H Z, \quad (9)$$

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 condHz 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 semidefinite 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 nonsingular 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 nonsingular).

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

$$W^T \lambda = g(x). \quad (10)$$

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**. If a multiplier is nonoptimal, 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 nonzero, there is no feasible point and the function terminates immediately with **ifail** = 3 (see Section 6).

The special form (7) of the working set allows the multiplier vector λ , the solution of (10), to be written in terms of the vector

$$d = \begin{pmatrix} g \\ 0 \end{pmatrix} - \begin{pmatrix} A^T \\ -I \end{pmatrix} \pi = \begin{pmatrix} g - A^T \pi \\ \pi \end{pmatrix}, \quad (11)$$

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 (9). 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 **rgNorm**, **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 (8) and (12)). 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 nonsingular, R is nonsingular and p_Z in (5) is computed from the equations

$$R^T R p_Z = -g_Z, \quad (12)$$

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_N , 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

$$p_Z^T H_Z p_Z = 0 \quad \text{and} \quad g_Z^T p_Z \leq 0, \quad (13)$$

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 function 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.

e04nq makes explicit allowance for infeasible constraints. Infeasible linear constraints are detected first by solving a problem of the form

$$\underset{x, v, w}{\text{minimize}} \quad e^T(v + w) \quad \text{subject to} \quad l \leq \begin{pmatrix} x \\ Gx - v + w \end{pmatrix} \leq u, \quad v \geq 0, \quad w \geq 0, \quad (14)$$

where $e^T = (1, 1, \dots, 1)$. 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 (8) could be arbitrarily high. To guard against this, the function implements a ‘basis repair’ feature in which the LUSOL package (see Gill *et al.* 1991) is used to compute the rectangular factorization

$$(B \ S)^T = LU, \quad (15)$$

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 (7). 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** is δ . Over a period of K iterations (where K is the value of the optional parameter **Expand Frequency**), the feasibility

tolerance actually used by the function (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 nontrivial adjustments made. If the count is nonzero, 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 function reaches an apparently optimal, infeasible or unbounded solution, unless this situation has already occurred twice. If any nontrivial 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 e04nq define choices in the problem specification or the algorithm logic. In order to reduce the number of formal parameters of e04nq these optional parameters have associated *default values* that are appropriate for most problems. Therefore, you need only specify those optional parameters whose values are to be different from their default values.

The remainder of this section can be skipped if you 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 functions e04ns, e04nt and e04nu before a call to e04nq, but after a call to e04np.

e04ns, e04nt or e04nu can be called to supply options directly, one call being necessary for each optional parameter. e04ns, e04nt or e04nu should be consulted for a full description of this method of supplying optional parameters.

All optional parameters not specified by you are set to their default values. Optional parameters specified by you are unaltered by e04nq (unless they define invalid values) and so remain in effect for subsequent calls unless altered by you.

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 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 x02aj). The variable *infbnd* holds the value of **Infinite Bound Size**.

Optional Parameters	Default Values
<u>Backup Basis File</u>	Default = 0. See <u>New Basis File</u> .
<u>Check Frequency</u>	Default = 60
<u>Crash Option</u>	Default = 3
<u>Crash Tolerance</u>	Default = 0.1. See <u>Crash Option</u> .
<u>Defaults</u>	
<u>Dump File</u>	Default = 0
<u>Elastic Mode</u>	Default = 1
<u>Elastic Objective</u>	Default = 1

<u>Elastic Weight</u>	Default = 1.0
<u>Expand Frequency</u>	Default = 10000
<u>Factorization Frequency</u>	Default = 100(LP) or 50(QP)
<u>Feasibility Tolerance</u>	Default = $\max\{10^{-6}, \sqrt{\epsilon}\}$
<u>Feasible Point</u>	See <u>Minimize</u>
<u>Infinite Bound Size</u>	Default = 10^{20}
<u>Insert File</u>	Default = 0. See <u>Punch File</u> .
<u>Iterations Limit</u>	Default = $\max\{10000, 10 \max\{m, n\}\}$
<u>LU Complete Pivoting</u>	See <u>LU Partial Pivoting</u>
<u>LU Density Tolerance</u>	Default = 0.6
<u>LU Factor Tolerance</u>	Default = 100.0
<u>LU Partial Pivoting</u>	Default
<u>LU Rook Pivoting</u>	See <u>LU Partial Pivoting</u>
<u>LU Singularity Tolerance</u>	Default = $\epsilon^{\frac{2}{3}}$. See <u>LU Density Tolerance</u> .
<u>LU Update Tolerance</u>	Default = 10.0. See <u>LU Factor Tolerance</u> .
<u>List</u>	See <u>Nolist</u>
<u>Load File</u>	Default = 0. See <u>Dump File</u> .
<u>Maximize</u>	See <u>Minimize</u>
<u>Minimize</u>	Default
<u>New Basis File</u>	Default = 0
<u>Nolist</u>	Default
<u>Old Basis File</u>	Default = 0
<u>Optimality Tolerance</u>	Default = $\max\{10^{-6}, \sqrt{\epsilon}\}$
<u>Partial Price</u>	Default = 10(LP) or 1(QP)
<u>Pivot Tolerance</u>	Default = $\epsilon^{\frac{2}{3}}$
<u>Print File</u>	Default = 0
<u>Print Frequency</u>	Default = 100
<u>Print Level</u>	Default = 1
<u>Punch File</u>	Default = 0
<u>QPSolver CG</u>	See <u>QPSolver Cholesky</u>
<u>QPSolver Cholesky</u>	Default
<u>QPSolver QN</u>	See <u>QPSolver Cholesky</u>
<u>Reduced Hessian Dimension</u>	Default = 1(LP) or $\min(2000, n_H + 1, n)$ (QP)
<u>Save Frequency</u>	Default = 100. See <u>New Basis File</u> .
<u>Scale Option</u>	Default = 2
<u>Scale Print</u>	See <u>Scale Option</u>
<u>Scale Tolerance</u>	Default = 0.9. See <u>Scale Option</u> .
<u>Solution File</u>	Default = 0
<u>Summary File</u>	Default = 0
<u>Summary Frequency</u>	Default = 100. See <u>Summary File</u> .
<u>Superbasics Limit</u>	Default = 1(LP) or $\min\{n_H + 1, n\}$ (QP)
<u>Suppress Parameters</u>	Default
<u>System Information No</u>	See <u>System Information No</u>
<u>System Information Yes</u>	Default = 0
<u>Timing Level</u>	Default = 0
<u>Unbounded Step Size</u>	Default = <i>infbnd</i>

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 refactorized and the basic variables recomputed to satisfy the constraints more accurately. If $i \leq 0$, 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 these options do 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 $(A \ -I)$. 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, treating 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.

The **Crash Tolerance** allows the Crash procedure to ignore certain 'small' nonzero elements in each column of A . If a_{\max} is the largest element in column j , other nonzeros 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 a **Crash Tolerance** of r for some value of $r > 0.5$.

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

Optional parameters **Dump File** and **Load File** are similar to optional parameters **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 optional parameters **Dump File** and **Load File** is given in Gill *et al.* 2005.

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 optional parameters **Old Basis File** or **Insert File** are 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>i</i> | Meaning |
|----------|---|
| 0 | Elastic mode is never invoked. e04nq 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 the optional parameters 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 if a feasible point exists.

Elastic Objective *i* Default = 1

This determines the form of the composite objective $f(x) + \gamma \sum_j (v_j + w_j)$ in Phase 2 (γ). Three types of composite objectives are available.

- | <i>i</i> | Meaning |
|----------|---|
| 0 | Include only the true objective $f(x)$ in the composite objective. This option sets $\gamma = 0$ in the composite objective and allows e04nq to ignore the elastic bounds and find a solution that minimizes $f(x)$ subject to the non-elastic 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 specified by the input array helast . |

Elastic Weight *r* Default = 1.0

This defines the value of γ in the composite objective in Phase 2 (γ).

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 $f(x)$ is the quadratic objective.

Note that the effect of γ is *not* disabled once a feasible point 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 e04nq (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 the description of the optional parameter **Pivot Tolerance**).

If $i \leq 0$, the value $i = 9999999$ is used and effectively no anti-cycling procedure is invoked.

Factorization 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 optional parameter **Check Frequency** to ensure that the linear constraints $Ax - s = 0$ are satisfied. Occasionally, 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 = $\max\{10^{-6}, \sqrt{\epsilon}\}$

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 .)

e04nq 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 e04nq 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*. e04nq will not attempt to find the solution that minimizes the sum unless **Elastic Objective** = 2.

If the constraints and variables have been scaled (see the description of the optional parameter **Scale Option**), 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 $r \geq 0$, r defines the ‘infinite’ bound *infbnd* in the definition of the problem constraints. Any upper bound greater than or equal to *infbnd* will be regarded as plus infinity (and similarly any lower bound less than or equal to $-infbnd$ will be regarded as minus infinity). If $r < 0$, the default value is used.

Iterations Limit i Default = $\max\{10000, 10 \max\{m, n\}\}$

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; and feasibility and optimality will be checked. No iterations will be performed. If $i < 0$, the default value is used.

LU Density Tolerance r_1 Default = 0.6

LU Singularity Tolerance r_2 Default = ϵ^2

The density tolerance r_1 is used during *LU* factorization of the basis matrix. Columns of L and rows of U are formed one at a time, and the remaining rows and columns of the basis are altered appropriately. At any stage, if the density of the remaining matrix exceeds r_1 , the Markowitz strategy for choosing pivots is terminated. The remaining matrix is factored by a dense *LU* procedure. Raising the density tolerance towards 1.0 may give slightly sparser *LU* factors, with a slight increase in factorization time.

If $r_2 > 0$, r_2 defines the singularity tolerance used to guard against ill-conditioned basis matrices. After B is refactorized, the diagonal elements of U are tested as follows. If $|u_{jj}| \leq r_2$ or $|u_{jj}| < r_2 \max_i |u_{ij}|$, the j th column of the basis is replaced by the corresponding slack variable. If $r_2 \leq 0$, the default value is used.

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 sub-matrix of the form

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

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

LU Partial Pivoting	Default
LU Complete Pivoting	
LU Rook Pivoting	

The LU factorization implements a Markowitz-type search for pivots that locally minimize the fill-in subject to a threshold pivoting stability criterion. The default option is to use threshold partial pivoting. The options **LU Complete Pivoting** and **LU Rook Pivoting** are more expensive but more stable and better at revealing rank, as long as the **LU Factor Tolerance** is not too large (say < 2.0).

Minimize	Default
Maximize	
Feasible Point	

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.

The option **Feasible Point** means ‘ignore the objective function, while finding a feasible point for the linear constraints’. It can be used to check that the constraints are feasible without altering the call to e04nq.

New Basis File	i_1	Default = 0
Backup Basis File	i_2	Default = 0
Save Frequency	i_3	Default = 100

Optional parameters **New Basis File** and **Backup Basis File** are 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 nontrivial 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.

Use of $i_2 > 0$ 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 e04nq 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.* 2005.

Nolist **List**

Default

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

Old Basis File

 i

Default = 0

If $i > 0$, the basis maps information will be obtained from this file. The file will usually have been output previously as a **New Basis File** or **Backup Basis File**. A full description of information recorded in **New Basis File** and **Backup Basis File** is given in Gill *et al.* 2005.

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

Optimality Tolerance

 r Default = $\max\{10^{-6}, \sqrt{\epsilon}\}$

This is used to judge the size of the reduced gradients $d_j = g_j - a_j^T \pi$, 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

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

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, I_j , for $j = 1, 2, \dots, i$ (modulo i). If the previous pricing search was successful on A_{j-1}, I_{j-1} , the next search begins on the segments A_j and I_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}, I_{j+1} , and so on. If $i \leq 0$, the default value is used.

Pivot Tolerance r Default = ϵ^2

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’ determines which component of x reaches an upper or lower bound first. The corresponding element of p is called the pivot element. 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 optional parameter **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 optional parameter **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. If $i \leq 0$, the value of $i = 9999999$ is used and effectively no checks are made.

Print Level i Default = 1

This controls the amount of printing produced by e04nq as follows.

i	Meaning
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 the optional parameter **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.* 2005.

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

If $i_2 > 0$, 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.

QPSolver Cholesky	Default
QPSolver CG	
QPSolver QN	

Specifies the active-set algorithm used to solve the quadratic program in Phase 2 (γ). **QPSolver Cholesky** holds the full Cholesky factor R of the reduced Hessian $Z^T H Z$. As the QP iterations proceed, the dimension of R changes with the number of superbasic variables. If the number of superbasic variables needs to increase beyond the value of **Reduced Hessian Dimension**, the reduced Hessian cannot be stored and the solver switches to **QPSolver CG**. The Cholesky solver is reactivated if the number of superbasics stabilizes at a value less than **Reduced Hessian Dimension**.

QPSolver QN solves the QP using a quasi-Newton method. In this case, R is the factor of a quasi-Newton approximate Hessian.

QPSolver CG uses an active-set method similar to **QPSolver QN**, but uses the conjugate-gradient method to solve all systems involving the reduced Hessian.

The Cholesky QP solver is the most robust, but may require a significant amount of computation if there are many superbasics.

The quasi-Newton QP solver does not require computation of the exact R at the start of Phase 2 (γ). It may be appropriate when the number of superbasics is large but relatively few iterations are needed to reach a solution (e.g., if e04nq is called with a Warm Start).

The conjugate-gradient QP solver is appropriate for problems with many degrees of freedom (say, more than 2000 superbasics).

Reduced Hessian Dimension	i	Default = 1(LP) or $\min(2000, n_H + 1, n)$ (QP)
----------------------------------	-----	--

This specifies that an i by i triangular matrix R (to define the reduced Hessian according to $R^T R = Z^T H Z$). is to be available for use by the Cholesky QP solver.

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

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 100).

- 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.

Optional parameter **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 nonzero coefficients in each column:

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

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. The value of r should lie in the range $0 < r < 1$.

Scale Print causes the row scales $r(i)$ and column scales $c(j)$ to be printed to **Print File**, if **System Information Yes** has been specified. The scaled matrix coefficients are $\bar{a}_{ij} = a_{ij}c(j)/r(i)$, and the scaled bounds on the variables and slacks are $\bar{l}_j = l_j/c(j)$, $\bar{u}_j = u_j/c(j)$, where $c(j) = r(j-n)$ if $j > n$.

Solution File i Default = 0

If $i > 0$, the final solution will be output to file i (whether optimal or not).

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 $i_1 > 0$, the Summary file is output to file i_1 , including a line of the iteration log 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. If $i_2 \leq 0$, the value of $i_2 = 99999999$ is used and effectively no checks are made.

Superbasics Limit i Default = 1(LP) or $\min\{n_H + 1, n\}$ (QP)

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 $n_H + 1$, where n_H is the number of leading nonzero columns of H . For many problems, i may be considerably smaller than n_H . This will save storage if n_H is very large.

Suppress Parameters

Normally e04nq prints the options file as it is being read, and then prints a complete list of the available keywords and their final values. The optional parameter **Suppress Parameters** tells e04nq not to print the full list.

System Information No **System Information Yes**

Default

This option prints additional information on the progress of major and minor iterations, and Crash statistics. See Section 12.

Timing Level

 i

Default = 0

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

Unbounded Step Size

 r Default = *infbnd*

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. See **Infinite Bound Size** for the definition of *infbnd*.

12 Description of Monitoring Information

This section describes the intermediate printout and final printout which constitutes the monitoring information produced by e04nq. (See also the description of the optional parameters **Print File** and **Print Level**.) You can control the level of printed output.

12.1 Crash Statistics

When **Print Level** ≥ 10 , **Print File** > 0 and **System Information Yes** has been specified, the following lines of intermediate printout (less than 120 characters) are produced on the unit number specified by optional parameter **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., $\mathbf{hs}(j) = 3$ for some $j \leq n$). It will be a subset of the columns for which $\mathbf{hs}(j) = 3$ 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 nonsingular triangle).

12.2 Basis Factorization Statistics

When **Print Level** ≥ 10 and **Print File** > 0 , the first seven items of intermediate printout in the list below are produced on the unit number specified by optional parameter **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**. In addition, if **System Information Yes** has been specified, the entries from **Elms** onwards are also output.

Label	Description
Factor	the number of factorizations since the start of the run.

Demand	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 <i>LU</i> factorization.</td></tr> <tr> <td>1</td><td>The number of updates reached the Factorization Frequency.</td></tr> <tr> <td>2</td><td>The nonzeros in the updated factors have increased significantly.</td></tr> <tr> <td>7</td><td>Not enough storage to update factors.</td></tr> <tr> <td>10</td><td>Row residuals are too large (see the description of 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 <i>LU</i> factorization.	1	The number of updates reached the Factorization Frequency .	2	The nonzeros in the updated factors have increased significantly.	7	Not enough storage to update factors.	10	Row residuals are too large (see the description of the optional parameter Check Frequency).	11	Ill-conditioning has caused inconsistent results.
Code	Meaning														
0	First <i>LU</i> factorization.														
1	The number of updates reached the Factorization Frequency .														
2	The nonzeros in the updated factors have increased significantly.														
7	Not enough storage to update factors.														
10	Row residuals are too large (see the description of the optional parameter Check Frequency).														
11	Ill-conditioning has caused inconsistent results.														
Itn	is the current minor iteration number.														
Nonlin	is the number of nonlinear variables in the current basis <i>B</i> .														
Linear	is the number of linear variables in <i>B</i> .														
Slacks	is the number of slack variables in <i>B</i> .														
B, BR, BS or BT factorize	is the type of <i>LU</i> factorization.														
	<table> <tr> <td>B</td><td>periodic factorization of the basis <i>B</i>.</td></tr> <tr> <td>BR</td><td>more careful rank-revealing factorization of <i>B</i> using threshold rook pivoting. This occurs mainly at the start, if the first basis factors seem singular or ill-conditioned. Followed by a normal B factorize.</td></tr> <tr> <td>BS</td><td>B_S is factorized to choose a well-conditioned <i>B</i> from the current (<i>B S</i>). Followed by a normal B factorize.</td></tr> <tr> <td>BT</td><td>same as BS except the current <i>B</i> is tried first and accepted if it appears to be not much more ill-conditioned than after the previous BS factorize.</td></tr> </table>	B	periodic factorization of the basis <i>B</i> .	BR	more careful rank-revealing factorization of <i>B</i> using threshold rook pivoting. This occurs mainly at the start, if the first basis factors seem singular or ill-conditioned. Followed by a normal B factorize.	BS	B_S is factorized to choose a well-conditioned <i>B</i> from the current (<i>B S</i>). Followed by a normal B factorize.	BT	same as BS except the current <i>B</i> is tried first and accepted if it appears to be not much more ill-conditioned than after the previous BS factorize.						
B	periodic factorization of the basis <i>B</i> .														
BR	more careful rank-revealing factorization of <i>B</i> using threshold rook pivoting. This occurs mainly at the start, if the first basis factors seem singular or ill-conditioned. Followed by a normal B factorize.														
BS	B_S is factorized to choose a well-conditioned <i>B</i> from the current (<i>B S</i>). Followed by a normal B factorize.														
BT	same as BS except the current <i>B</i> is tried first and accepted if it appears to be not much more ill-conditioned than after the previous BS factorize.														
m	is the number of rows in <i>B</i> or B_S .														
n	is the number of columns in <i>B</i> or B_S . Preceded by '=' or '>' respectively.														
Elms	is the number of nonzero elements in <i>B</i> or B_S .														
Amax	is the largest nonzero in <i>B</i> or B_S .														
Density	is the percentage nonzero density of <i>B</i> or B_S .														
Merit/MerRP/MerCP	Merit is the average Markowitz merit count for the elements chosen to be the diagonals of <i>PUQ</i> . Each merit count is defined to be $(c-1)(r-1)$ where <i>c</i> and <i>r</i> are the number of nonzeros in the column and row containing the element at the time it is selected to be the next diagonal. Merit is the average of <i>n</i> such quantities. It gives an indication of how much work was required to preserve sparsity during the factorization. If LU Complete Pivoting or LU Rook Pivoting has been selected, this heading is changed to MerCP, respectively MerRP.														
lenL	is the number of nonzeros in <i>L</i> .														
L+U	is the number of nonzeros representing the basis factors <i>L</i> and <i>U</i> . Immediately after a basis factorization $B = LU$, this is lenL+lenU, the number of subdiagonal elements in the columns of a lower triangular matrix and the number of diagonal and superdiagonal elements in the rows of an upper-triangular matrix. Further nonzeros are added to <i>L</i> when various columns of <i>B</i> are later replaced. As columns of <i>B</i> are replaced, the matrix <i>U</i> is maintained explicitly (in sparse form). The value of <i>L</i> will steadily increase, whereas the value of <i>U</i> may fluctuate up or down. Thus the value of <i>L+U</i> may fluctuate up or down (in general, it will tend to increase).														
Cmpressns	is the number of times the data structure holding the partially factored matrix needed to be compressed to recover unused storage. Ideally this number should be zero. If it is more than 3 or 4, the amount of workspace available to e04nq should be increased for efficiency.														

Incres	is the percentage increase in the number of nonzeros in L and U relative to the number of nonzeros in B or B_S .
Utri	is the number of triangular rows of B or B_S at the top of U .
lenU	the number of nonzeros in U , including its diagonals.
Ltol	is the largest subdiagonal element allowed in L . This is the specified LU Factor Tolerance or a smaller value that is currently being used for greater stability.
Umax	the maximum nonzero element in U .
Ugrwth	is the ratio U_{\max}/A_{\max} , which ideally should not be substantially larger than 10.0 or 100.0. If it is orders of magnitude larger, it may be advisable to reduce the LU Factor Tolerance to 5.0, 4.0, 3.0 or 2.0, say (but bigger than 1.0). As long as L_{\max} is not large (say 5.0 or less), $\max(A_{\max}, U_{\max})/D_{\min}$ gives an estimate of the condition number B . If this is extremely large, the basis is nearly singular. Slacks are used to replace suspect columns of B and the modified basis is refactored.
Ltri	is the number of triangular columns of B or B_S at the left of L .
dense1	is the number of columns remaining when the density of the basis matrix being factorized reached 0.3.
Lmax	is the actual maximum subdiagonal element in L (bounded by L_{tol}).
Akmax	is the largest nonzero generated at any stage of the LU factorization. (Values much larger than A_{\max} indicate instability.) Ak_{\max} is not printed if LU Partial Pivoting is selected.
Agrwth	is the ratio Ak_{\max}/A_{\max} . Values much larger than 100 (say) indicate instability. $Agrwth$ is not printed if LU Partial Pivoting is selected.
bump	is the size of the block to be factorized nontrivially after the triangular rows and columns of B or B_S have been removed.
dense2	is the number of columns remaining when the density of the basis matrix being factorized reached 0.6. (The Markowitz pivot strategy searches fewer columns at that stage.)
DUmax	is the largest diagonal of PUQ .
DUmin	is the smallest diagonal of PUQ .
condU	the ratio DU_{\max}/DU_{\min} , which estimates the condition number of U (and of B if L_{tol} is less than 5.0, say).

12.3 Basis Map

When **Print Level** ≥ 10 and **Print File** > 0 , the following lines of intermediate printout (less than 80 characters) are produced on the unit number specified by optional parameter **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 problem unnamed).
Infeasibilities	gives the number of infeasibilities. Printed only if the final point is infeasible.
Objective Value	gives the objective value at the final point (or the value of the sum of infeasibilities). Printed only if the final point is feasible.
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)).
Iteration	gives the iteration number when the file was created.
Superbasics	gives the number of superbasic variables.
Objective	gives the name of the free row for the problem (blank if objective unnamed).
RHS	gives the name of the constraint right-hand side for the problem (blank if objective unnamed).
Ranges	gives the name of the ranges for the problem (blank if objective unnamed).
Bounds	gives the name of the bounds for the problem (blank if objective unnamed).

12.4 Solution Output

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.

12.4.1 The ROWS section

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

$$\alpha \leq \nu_i x \leq \beta,$$

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, the slack variable s_i is directly available, and it is sometimes convenient to refer to its state. It should satisfy $\alpha \leq s_i \leq \beta$. 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: <table> <tr> <td>LL</td><td>s_i is nonbasic at its lower limit, α.</td></tr> <tr> <td>UL</td><td>s_i is nonbasic at its upper limit, β.</td></tr> <tr> <td>EQ</td><td>s_i is nonbasic and fixed at the value $\alpha = \beta$.</td></tr> <tr> <td>FR</td><td>s_i is nonbasic and currently zero, even though it is free to take any value between its bounds α and β.</td></tr> <tr> <td>BS</td><td>s_i is basic.</td></tr> <tr> <td>SBS</td><td>s_i is superbasic.</td></tr> </table>	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.
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. Note that unless the optional parameter **Scale Option** = 0 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 <i>might</i> also change.
D	<i>Degenerate.</i> The variable is basic or superbasic, but it is equal (or very close) to one of its bounds.
I	<i>Infeasible.</i> The variable is basic or superbasic and is currently violating one of its bounds by more than the value of the Feasibility Tolerance .
N	<i>Not precisely optimal.</i> If the slack is superbasic, the dual variable π_i is not sufficiently small, as measured by the Optimality Tolerance . If the slack is nonbasic, π_i is not sufficiently positive or negative. If a loose Optimality Tolerance has been used, or if iterations were terminated before optimality, this key might be helpful in deciding whether or not to restart the run.
Activity	is the value of $\nu_j x$ at the final iterate.
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 $\mathbf{bl}(j) \leq -\mathit{infbnd}$.
Upper Limit	is β , the upper bound specified for the variable s_i . None indicates that $\mathbf{bu}(j) \geq \mathit{infbnd}$.
Dual Activity	is the value of the dual variable π_i (the Lagrange multiplier for ν_i ; see Section 10.3). For FP problems, π_i is set to zero.
i	gives the index i of the i th row.

12.4.2 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. Note that unless the optional parameter **Scale Option** = 0 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 (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 **Feasibility Tolerance**.
- N *Not precisely optimal.* If the slack is superbasic, the dual variable π_i is not sufficiently small, as measured by the **Optimality Tolerance**. If the slack is nonbasic, π_i is not sufficiently positive or negative. If a loose **Optimality Tolerance** has been used, or if iterations were terminated before optimality, this key might be helpful in deciding whether or not to restart the run.

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 Limit	is the lower bound specified for the variable. None indicates that $\text{bl}(j) \leq -\text{infbnd}$.
Upper Limit	is the upper bound specified for the variable. None indicates that $\text{bu}(j) \geq \text{infbnd}$.
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.

12.5 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 $\pm 10^{20}$ rather than None. Other real values are output with format 1pe16.6. The maximum line 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 lines would be ignored. Each subsequent line 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 line that starts with a 1 and is otherwise blank. If this and the next 4 lines are skipped, the COLUMNS section can then be read under the same format. (There should be no need to use any BACKSPACE statements.)

12.6 The Summary File

If **Summary File** f is specified with $f > 0$, certain brief information will be output to unit f . When e04nq is run interactively, unit 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 line length is 72 characters, including a carriage-control character in column 1.

The following information is included:

1. The optional parameters supplied via the option setting functions, if any;
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**;
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 , **names**(j) replaced by **names**($n + j$), **bl**(j) and **bu**(j) are replaced by **bl**($n + j$) and **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.
