Mathematical Programming 5 (1973) 235–242. North-Holland Publishing Company

COMPUTABLE ERROR BOUNDS FOR NONLINEAR PROGRAMMING*

Stephen M. ROBINSON
University of Wisconsin, Madison, Wisc., U.S.A.

Received 11 October 1972
Revised manuscript received 26 April 1973

This paper presents a method for obtaining computable bounds for the error in an approximate Kuhn–Tucker point of a nonlinear program. Techniques of interval analysis are employed to compute the error bounds.

1. Introduction

Consider a general nonlinear programming problem of the form

$$\text{minimize } \{\theta(x):\ g(x) \leqslant 0,\ h(x) = 0,\ x \in \Gamma\}, \tag{1}$$

where θ, g and h are differentiable, not necessarily convex functions from an open set $\Gamma \subset \mathbf{R}^n$ into $\mathbf{R}$, $\mathbf{R}^m$, and $\mathbf{R}^q$, respectively. The usual result of the numerical solution of (1) by an algorithm is a triple $(\tilde{x}, \tilde{u}, \tilde{v}) \in \mathbf{R}^{n+m+q}$, with $\tilde{x} \in \Gamma$, at which the (first-order) Kuhn–Tucker conditions [2, §2.1],

$$\theta'(x) + u^{\mathrm{T}} g'(x) + v^{\mathrm{T}} h'(x) = 0,$$

$$g(x) \leqslant 0, \qquad u^{\mathrm{T}} g(x) = 0, \qquad h(x) = 0, \qquad u \geqslant 0 \tag{2}$$

are approximately satisfied; we usually have $\tilde{u} \geqslant 0$, but in general none of the other conditions will be exactly satisfied. It is usually assumed that there is a Kuhn–Tucker point $(\hat{x}, \hat{u}, \hat{v})$ for (1) "close" to $(\tilde{x}, \tilde{u}, \tilde{v})$.

* Sponsored by the United States Army under Contract No. DA-31-124-ARO-D-462.

However, it is of interest to know whether this assumption is actually true, and if so how far away from $(\tilde{x}, \tilde{u}, \tilde{v})$ the point $(\hat{x}, \hat{u}, \hat{v})$ could be.

In this paper we outline a procedure applicable to a wide class of problems of the form (1), including in particular all those in which θ, g, and h are rational functions. If certain conditions are satisfied, this technique enables us to prove that the point $(\hat{x}, \hat{u}, \hat{v})$ exists and to give rigorous, computable bounds for its distance from $(\tilde{x}, \tilde{u}, \tilde{v})$. The entire procedure is implementable on a computer; it employs interval arithmetic [5] and is based on recent work of Nickel [6] which in turn rests on earlier results due to Hansen [3].

2. Notation and definitions

We introduce in this section several definitions and notational conventions, mostly related to interval arithmetic, which we shall use in what follows. A complete discussion of interval analysis can be found in [5].

A real interval is an interval $[\alpha, \beta]$ on the real line $\mathbf{R}$; an interval in $\mathbf{R}^p$ is a p-vector each of whose components is a real interval, and an interval matrix is defined in the corresponding way. All interval quantities will be written with a bar over them: e.g., $\bar{y}$ is an interval vector. We shall not distinguish between the vector $y \in \mathbf{R}^p$ and the degenerate interval vector whose components are $[y_i, y_i]$ for $1 \leqslant i \leqslant p$; this convention permits us to treat points as intervals if it is more convenient to do so.

Let $\bar{z} \subset \mathbf{R}^p$ be a fixed interval vector, and let f be a function from $\mathbf{R}^p$ into $\mathbf{R}^k$. We denote by $\bar{f}$ an *interval extension* of f on $\bar{z}$; that is, an interval function (one which operates on intervals to produce intervals) satisfying the following two properties:

(i) for each $z \in \bar{z}$, $\bar{f}(z) = f(z)$,

(ii) for each $\bar{w} \subset \bar{y} \subset \bar{z}$, $\bar{f}(\bar{w}) \subset \bar{f}(\bar{y})$.

Let $\bar{M}$ be an $n \times n$ interval matrix. If each $n \times n$ matrix $P \in \bar{M}$ (i.e., such that $P_{ij} \in \bar{M}_{ij}$ for $1 \leqslant i, j \leqslant n$) is invertible, then we say that $\bar{M}$ is *nonsingular*. We denote by $\bar{M}^{-1}$ an interval inverse of $\bar{M}$; that is, a matrix such that for each $P \in \bar{M}$ we have $P^{-1} \in \bar{M}^{-1}$. Interval inverses, like interval extensions, are not unique.

One of the reasons for employing interval analysis in computing is that it can be implemented on a computer in such a way as to take into account *automatically* the errors due to roundoff at each stage of the

computations. For details of two such implementations, see [1,4]. One can thus in many cases obtain rigorous error bounds for a computed quantity by means of interval analysis when such bounds could not have been found by conventional (analytic) methods. For examples, see [3, 5, 6].

3. Error bounds for a Kuhn–Tucker point

In this section we develop an error-bounding technique for nonlinear programming. The method is based on two observations: first, that the Kuhn–Tucker conditions for a stationary point can be written as a system of nonlinear equations augmented by certain inequality conditions; second, that techniques are known [6] for obtaining computable bounds on approximate solutions of systems of nonlinear equations by means of interval analysis. Our strategy for bounding a Kuhn–Tucker point will be to combine these two facts with a simple device to ensure that the inequalities of the Kuhn–Tucker conditions are taken into account.

The system of nonlinear equations which we shall employ is obtained by letting $p := n + m + q$, denoting by z the triple (x, u, v), and defining a function f from $\mathbf{R}^p$ into itself by

$$f(z) := \begin{bmatrix} [\theta'(x) + u^{\mathrm{T}} g'(x) + v^{\mathrm{T}} h'(x)]^{\mathrm{T}} \\ u_1 g_1(x) \\ \vdots \\ u_m g_m(x) \\ h_1(x) \\ \vdots \\ h_q(x) \end{bmatrix}, \tag{3}$$

where we have represented the derivatives by matrices of n columns. Note that (x, u, v) is a Kuhn–Tucker point of (1) if and only if $f(z) = 0$, $g(x) \leqslant 0$ and $u \geqslant 0$. The function f is basic to the analysis of optimality and sensitivity in nonlinear programming ([2; §§2.3, 5.2]; see also [7]).

To obtain the additional conditions required to deal with the inequalities in the Kuhn–Tucker conditions, we introduce two definitions: For any interval $\bar{c} \subset \mathbf{R}$, we let

$$\bar{c}_0 := \begin{cases} \text{conv}[0, \bar{c}] & \text{if } \bar{c} \subset \mathbf{R}_+, \\ \bar{c} & \text{otherwise,} \end{cases}$$

$$\bar{c}^0 := \begin{cases} \text{conv}[0, \bar{c}] & \text{if } \bar{c} \subset \mathbf{R}_-, \\ \bar{c} & \text{otherwise,} \end{cases}$$

where conv$[0, \bar{c}]$ denotes the convex hull of 0 and $\bar{c}$, and $\mathbf{R}_+$ and $\mathbf{R}_-$ are respectively the non-negative and non-positive real numbers. We also recall that *strict complementary slackness* is said to hold at a Kuhn–Tucker point $(\hat{x}, \hat{u}, \hat{v})$ if for each i, $1 \leqslant i \leqslant m$, exactly one of the quantities $\hat{u}_i$ and $g_i(\hat{x})$ is zero.

We are now ready to present the main result. To do so, we must first define two interval quantities which will be used in stating the theorem. Let $\bar{x}$, $\bar{u}$ and $\bar{v}$ be intervals in $\mathbf{R}^n$, $\mathbf{R}^m$ and $\mathbf{R}^q$, and let θ, g and h be twice differentiable on $\bar{x}$. Define $\bar{z} := (\bar{x}, \bar{u}, \bar{v})$, and denote the components of $\bar{z}$ by $\bar{z}_i$ for $1 \leqslant i \leqslant p$. For $1 \leqslant i, j \leqslant p$, we suppose that there are available interval extensions, denoted by $\partial \bar{f}_i(z)/\partial z_j$, of $\partial f_i(z)/\partial z_j$ on $\bar{z}$, where f is defined by (3). We define an interval matrix $\bar{F}'(\overline{w}, \overline{y})$ for $\overline{w}, \overline{y} \subset \bar{z}$ by

$$\bar{F}'(\overline{w}, \overline{y})_{ij} := \partial \bar{f}_i(\overline{w}_1, \ldots, \overline{w}_{j-1}, \overline{y}_j, \ldots, \overline{y}_n)/\partial z_j, \qquad 1 \leqslant i, j \leqslant p.$$

Let the interval Newton operator $\bar{N}$ be defined for $z \in \bar{z}$ by $\bar{N}(z, \bar{z}) := z - \bar{F}'(z, \bar{z})^{-1} f(z)$; note that since the last expression involves an interval inverse, $\bar{N}$ is not uniquely defined.

Theorem. Let θ, g and h be functions of class C'' from an open subset Γ of $\mathbf{R}^n$ into $\mathbf{R}$, $\mathbf{R}^m$ and $\mathbf{R}^q$, respectively. Let $\bar{x} \subset \Gamma \subset \mathbf{R}^n$, $\bar{u} \subset \mathbf{R}^m$ and $\bar{v} \subset \mathbf{R}^q$, and let $\bar{g}$ be an interval extension of g on $\bar{x}$. Suppose that $0 \notin \bar{g}_i(\bar{x})_0 \cap \bar{u}_i^0$ for $1 \leqslant i \leqslant m$, that $\bar{F}'(\bar{z}, \bar{z})$ is nonsingular, and that there is a point $\tilde{z} \in \bar{z}$ such that $\bar{N}(\tilde{z}, \bar{z}) \subset \bar{z}$. Then $\bar{N}(\tilde{z}, \bar{z})$ contains a Kuhn–Tucker triple $\hat{z} := (\hat{x}, \hat{u}, \hat{v})$ of (1) at which strict complementary slackness holds with linear independence of the gradients to the active constraints; further, $(\hat{x}, \hat{u}, \hat{v})$ is the only Kuhn–Tucker triple of (1) in $\bar{z}$.

Proof. By a theorem of Nickel [6], a point $\hat{z} \in \bar{N}(\tilde{z}, \bar{z})$ exists with $f(\hat{z}) = 0$, and $\hat{z}$ is the unique zero of f in $\bar{z}$. Since $f(\hat{z}) = 0$, the inequalities of the Kuhn–Tucker conditions are satisfied. Choose any i between 1 and m; since $\hat{u}_i g_i(\hat{x}) = 0$ we have either $\hat{u}_i = 0$ or $g_i(\hat{x}) = 0$. Suppose the former, then $0 = \hat{u}_i \in \bar{u}_i \subset \bar{u}_i^0$, so by hypothesis $0 \notin \bar{g}_i(\bar{x})_0$,

and it then follows easily that $\bar{g}_i(\bar{x}) \subset \operatorname{int} \mathbf{R}_-$, where int denotes interior. But the interval extension property implies that $\bar{g}_i(\bar{x}) \supset \bar{g}_i(\hat{x}) = g_i(\hat{x})$, so we must have $g_i(\hat{x}) < 0$. Similarly, we can show that if $g_i(\hat{x}) = 0$, then $\hat{u}_i$ must be strictly positive. Hence $\hat{u} \geqslant 0$ and $g(\hat{x}) \leqslant 0$, so $\hat{z}$ is a Kuhn–Tucker triple for (1) at which strict complementary slackness holds. Since $\bar{F}'(\bar{z}, \bar{z})$ is nonsingular, the matrix $f'(\hat{z})$ (which, by the interval extension property, equals $\bar{F}'(\hat{z}, \hat{z})$ and so is contained in $\bar{F}'(\bar{z}, \bar{z})$) is also nonsingular. Linear independence of the gradients to the active constraints follows from the nonsingularity of $f'(\hat{z})$ and the strict complementary slackness. This completes the proof.

It is of interest at this point to consider whether the hypotheses of this theorem can be satisfied for real problems. In order to apply the theorem, one needs to be able to find interval extensions for g and for all of the second partial derivatives of θ, g and h; in addition, to apply the techniques in practice one will require interval extensions also for h and for all of the first partial derivatives of θ, g and h; these are necessary because it is generally impossible to compute $f(\tilde{z})$ exactly, and so one is forced to resort to an interval extension $\bar{f}$ in computing $\bar{N}(\tilde{z}, \bar{z})$. For wide classes of problems, such extensions are available; for example, if θ, g and h are rational functions of x there are available the so-called "natural" extensions obtained by replacing the real variable x with its interval equivalent $\bar{x}$ and computing the functions in interval arithmetic. These may or may not be the best extensions to use in a given situation; for more details, see [5, §6].

Another question bearing on the practicality of this procedure is that of determining $\bar{z}$. This interval must be large enough to contain $\hat{z}$, yet for purposes of improved accuracy it is desirable to choose $\bar{z}$ to be as small as possible. If the wrong choice is made, the procedure cannot be carried out: one may have $\bar{F}'(\bar{z}, \bar{z})$ singular, or else $\bar{N}(\tilde{z}, \bar{z}) \not\subset \bar{z}$. However, in practice it should not be very hard to choose a suitable $\bar{z}$, since it is very often possible to estimate from the progress of the computation roughly how many decimals of accuracy one has in an approximate solution. In such a case it is not hard to select a $\bar{z}$ which one has good reason to think will contain $\hat{z}$. Of course, if $\bar{N}(\tilde{z}, \bar{z}) \subset \bar{z}$, then one knows with certainty that a correct choice was made: the procedure will never yield a false error bound.

It is not difficult to show that if $\hat{z}$ satisfies the second-order sufficiency conditions [2] with strict complementary slackness and linear independence of the gradients to the active constraints, if the second

derivatives of θ, g and h are continuous in a neighborhood of $\hat{x}$, and if the interval extensions employed are continuous in the usual interval metric topology [5, Ch. 4], then there is an open neighborhood of $\hat{z}$ such that any interval $\bar{z}$ contained in this neighborhood and containing $\hat{z}$ in its interior will satisfy the conditions of the theorem, provided that $\tilde{z}$ is close enough to $\hat{z}$. Thus, in theory the method can always be made to work with good enough initial approximations. In practice one would simply choose $\bar{z}$ and try the method; if it does not work, adjust $\bar{z}$ (and possibly $\tilde{z}$) and try again.

We note that if $\tilde{z}$ is close to $\hat{z}$ and the matrix $f'(\hat{z})$ is not ill conditioned, then $\bar{N}(\tilde{z}, \bar{z})$ will usually be much smaller than $\bar{z}$. Hence the method improves the approximate solution at the same time as it establishes error bounds. This phenomenon is illustrated by the example in the next section. Of course, the improvement is paid for by the fairly considerable labor involved in calculating all of the second partial derivatives of θ, g and h. For this reason, since a great many derivatives will be involved if the problem is even moderately large, it would seem advisable in such cases to employ a pre-compiler which performs automatic (analytic) differentiation on the computer. Such programs are available, and their use often saves time as well as removing a possible source of error.

4. Numerical example

We present here a simple numerical example to show how the method actually works in practice. Consider the nonlinear program in $\mathbf{R}^2$ given by

$$\begin{aligned} &\text{minimize} \quad \theta(x) := x_2, \\ &\text{subject to} \quad g(x) := \begin{bmatrix} -x_1 + x_2^2 \\ (x_1 - 1)^2 + x_2^2 - 1 \\ x_1^2 + x_2^2 - 1 \end{bmatrix} \leqslant 0. \end{aligned} \tag{4}$$

Suppose that by using some algorithm we have obtained the following approximation to the solution $\hat{x}$ and the associated vector $\hat{u}$ of multipliers:

$$\tilde{x} = \begin{pmatrix} 0.618 \\ -0.786 \end{pmatrix}, \quad \tilde{u} = \begin{pmatrix} 0.352 \\ 0 \\ 0.284 \end{pmatrix}.$$

We define $\tilde{z} := (0.618, -0.786, 0.352, 0, 0.284)^{\mathrm{T}}$, and let $\bar{z} := \tilde{z} + \bar{d}$,

where $\bar{d}$ is an interval vector each of whose components is $[-0.001, 0.001]$. Using the techniques developed above, we find by computation with interval arithmetic, using the natural interval extensions for all functions involved, that

$$\bar{N}(\tilde{z}, \bar{z}) \subset \begin{bmatrix} [\ 0.618032, & 0.618037] \\ [-0.786153, & -0.786150] \\ [\ 0.351574, & 0.351581] \\ [-0.1727 \times 10^{-5}, & 0.1727 \times 10^{-5}] \\ [\ 0.284430, & 0.284434] \end{bmatrix} \subset \bar{z},$$

while with $\bar{z} := (\bar{x}, \bar{u})$ we have

$$\bar{g}(\bar{x}) \subset \begin{bmatrix} [-0.28 \times 10^{-2}, \ 0.24 \times 10^{-2}] \\ [-0.2387, -0.2339] \\ [-0.31 \times 10^{-2}, \ 0.26 \times 10^{-2}] \end{bmatrix},$$

so that $\bar{g}_i(\bar{x})_0 \cap \bar{u}_i^0 = \emptyset$ for $1 \leqslant i \leqslant 3$. Applying the theorem, we conclude that there is a Kuhn–Tucker point in the interval $\bar{N}(\tilde{z}, \bar{z})$ at which strict complementary slackness holds, and that this is the only Kuhn–Tucker point of (4) in $\bar{z}$. Note that the method has produced a bound exact to approximately five decimal places, whereas the original $\tilde{z}$ was correct to only three.

In fact, (4) can easily be solved analytically; if we carry out the solution we obtain

$$\hat{x}_1 = \tfrac{1}{2}(-1 + \sqrt{5}) \in [0.61803398, 0.61803399],$$

$$\hat{x}_2 = -\sqrt{\hat{x}_1} \in [-0.78615138, -0.78615137],$$

while $\hat{u}_2 = 0$, and $\hat{u}_1$ and $\hat{u}_3$ are the solutions of

$$\begin{bmatrix} 0 \\ 1 \end{bmatrix} + \begin{bmatrix} -1 & 2\hat{x}_1 \\ 2\hat{x}_2 & 2\hat{x}_2 \end{bmatrix} \begin{bmatrix} \hat{u}_1 \\ \hat{u}_3 \end{bmatrix} = 0,$$

so that

$$\hat{u}_1 \in [0.35157758, 0.35157759], \qquad \hat{u}_3 \in [0.28443224, 0.28443225].$$

Hence the error bounds given above are correct.

References

[1] N. Apostolatos et al., "The algorithmic language Triplex-ALGOL 60", *Numerische Mathematik* 11 (1968) 175–180.

[2] A.V. Fiacco and G.P. McCormick, *Nonlinear programming: Sequential unconstrained minimization techniques* (Wiley, New York, 1968).

[3] E.R. Hansen, "On solving systems of equations using interval arithmetic", *Mathematics of Computation* 22 (1968) 374–384.

[4] T.D. Ladner and J.M. Yohe, "An interval arithmetic package for the UNIVAC 1108", Technical Summary Report No. 1055, Mathematics Research Center, University of Wisconsin, Madison, Wisc. (1970).

[5] R.E. Moore, *Interval analysis* (Prentice-Hall, Englewood Cliffs, N.J., 1966).

[6] K. Nickel, "On the Newton method in interval analysis", Technical Summary Report No. 1136, Mathematics Research Center, University of Wisconsin, Madison, Wisc. (1971).

[7] S.M. Robinson, "Perturbed Kuhn–Tucker points and rates of convergence for a class of nonlinear-programming algorithms", Technical Summary Report No. 1298, Mathematics Research Center, University of Wisconsin, Madison, Wisc. (1972), to appear in *Mathematical Programming*.