

THEORY OF GLOBALLY CONVERGENT PROBABILITY-ONE HOMOTOPIES FOR NONLINEAR PROGRAMMING*

LAYNE T. WATSON[†]

Abstract. For many years globally convergent probability-one homotopy methods have been remarkably successful on difficult realistic engineering optimization problems, most of which were attacked by homotopy methods because other optimization algorithms failed or were ineffective. Convergence theory has been derived for a few particular problems, and considerable fixed point theory exists, but generally convergence theory for the homotopy maps used in practice for nonlinear constrained optimization has been lacking. This paper derives some probability-one homotopy convergence theorems for unconstrained and inequality constrained optimization, for linear and nonlinear inequality constraints, and with and without convexity. Some insight is provided into why the homotopies used in engineering practice are so successful, and why this success is more than dumb luck. By presenting the theory as variations on a prototype probability-one homotopy convergence theorem, the essence of such convergence theory is elucidated.

Key words. constrained optimization, globally convergent, homotopy algorithm, nonlinear equations, probability-one homotopy

AMS(MOS) subject classifications. 65F10, 65F50, 65H10, 65K10

1. Introduction. Continuation methods for optimization, as for nonlinear systems of equations, have been around for a long time and studied extensively. This paper concerns only a recent variant known as *globally convergent probability-one homotopy* methods. The words “continuation” and “homotopy” are often used interchangeably, but subtle and fundamental distinctions can be drawn between continuation, homotopy, and probability-one homotopy methods. These distinctions have been discussed numerous times in the literature [5], [8], [25], [28], [33]. The purpose of this paper is to help close a gap in the convergence theory for globally convergent probability-one homotopy methods applied to nonlinear programming, and to offer some theoretical justification for the observed success of homotopies in engineering practice.

From a high level perspective, all the fundamental convergence theory was done by Chow, Mallet-Paret, Yorke [5], and Watson [23], and all that remains is to verify that a particular homotopy map has the right properties. Alas, the devil is in the details, which are indeed often nontrivial. It is appropriate to sketch out here what is well understood and where gaps remain.

Much of the early work on computational homotopy algorithms was motivated by Brouwer fixed point problems: given a continuous function f from a compact, convex subset of finite dimensional Euclidean space into itself, find a fixed point $x = f(x)$. The algorithms and theory are elegant and well understood for both simplicial [3], [6], [7] and continuous [3], [5], [20] approaches.

* Received by the editors ; accepted for publication (in revised form) .

[†] Departments of Computer Science and Mathematics, Virginia Polytechnic Institute and State University, Blacksburg, Virginia 24061-0106. This work was supported in part by National Science Foundation Grant DMS-9625968, and Air Force Office of Scientific Research Grant F49620-96-1-0104.

For nonlinear systems of equations $F(x) = 0$ not derived from Brouwer fixed point problems, the theory [3], [5], [25], [26] and algorithms [28], [33] are well developed in terms of properties of F . Special cases, such as when F is a polynomial system, have a deep and rich supporting theory [13], and special, highly sophisticated algorithms have been devised to exploit the structure of F [14], [15], [33]. However, except in rare instances that usually result in polynomial systems, a physical model does not directly result in a finite dimensional nonlinear system of equations $F(x) = 0$. Rather, $F(x) = 0$ results from a discretization, approximation, or iteration step of another mathematical model of the physical phenomenon. The catch is that abstract conditions on F (for a homotopy algorithm to converge) do not easily translate into meaningful or verifiable conditions on the physical model or on the discretization/approximation/iteration process. The gap here is considerable: not many homotopy convergence theorems are stated at the level of physical modelling or the high level processes that spawn the nonlinear systems $F(x) = 0$ to be solved.

One notable exception is the solution of nonlinear two-point boundary value problems (BVPs). Conditions on the original two-point boundary value problem itself for which an approximation $F(x) = 0$ is solvable by a globally convergent homotopy algorithm have been derived in a series of papers. Convergence theorems directly addressing the nonlinear two-point boundary value problem exist for approximation processes based on shooting [21], finite differences [23], collocation [31], and finite elements [32]. This is significant because many physical models reduce to two-point boundary value problems, and thus convergence theory exists for a large class of problems of interest.

For nonlinear constrained optimization the gap has been large. Global convergence theorems, stated in terms of conditions on the objective function and constraints, for homotopy algorithms have been an elusive quarry. Some attempts include [16], [22], [24], [27], [29], [30]. Recently Lin et al. [9], [10], using a particular classical homotopy map (not a probability-one map) and constraint aggregation, have obtained convergence results for general nonlinear programming problems with a strong “normal cone condition” assumption. The convergence theory presented here has comparatively weak assumptions, applies to homotopy maps actually used in practice, and does not use constraint aggregation, which is numerically ill-conditioned in practice [30]. Probability-one homotopy algorithms have been enormously successful in engineering practice, notwithstanding the lack of theory. The goal of this work is to narrow the gap by providing such theorems for inequality constraints, and to help explain theoretically the observed success in practice [26]. Extending the theory for nonlinear equality constraints seems to require a homotopy existence theory for underdetermined nonlinear systems, and would at least involve a nontrivial extension of the proofs here. Nonlinear equality constraints are undeniably important, which mandates future work on homotopy theory for them.

There is a variant of probability-one homotopy theory for piecewise smooth functions [1], [2], and this might seem like a more natural tool for constrained optimization. Recent work along these lines includes [4], [18], and [19]. Despite the appeal of these nonsmooth formulations, they are not yet seriously competitive with the existing sophisticated numerical implementations for smooth formulations on realistic large scale problems [4].

After some background in §2, the theory is presented as a series of refinements applied to successively more general optimization problems. Sections 3 and 4 summarize some known results, but all the results after Theorem 4.1 are new. The progression is from unconstrained (§3) to nonnegative constraints (§4) to linear constraints (§5) to nonlinear convex constraints (§6) to general nonlinear constraints (§7). First convexity is assumed and then finally dropped in §7. Section 7 also provides some insight into why the homotopy maps used in engineering practice might (or might not) work.

2. Background and notation. Let E^n denote n -dimensional Euclidean space, $E^{m \times n}$ the set of real $m \times n$ matrices, and the i th component of a vector $x \in E^n$ by x_i . The i th row of a matrix $A \in E^{m \times n}$ is denoted by $A_{i.}$, and the j th column by $A_{.j}$. For sets of indices M and N , A_{MN} is the submatrix of A with rows indexed by M and columns indexed by N . Similarly x_M is the subvector of the vector x corresponding to the indices in M . No distinction is made between row vectors and column vectors, except when matrix arithmetic is involved. Following Mangasarian's notation for $x \in E^n$, $x > 0$ means all $x_i > 0$, $x \geq 0$ means all $x_i \geq 0$, and $x \geq 0$ means $x \geq 0$ but $x \neq 0$. $\|\cdot\|$ is the 2-norm unless otherwise indicated.

The *gradient* of a differentiable function $f : E^n \rightarrow E$ is the row vector $\nabla f(x) = \left(\frac{\partial f}{\partial x_1}(x), \dots, \frac{\partial f}{\partial x_n}(x) \right)$. The *Jacobian matrix* of $F : E^n \rightarrow E^m$ is

$$DF(x) = \nabla F(x) = \begin{pmatrix} \nabla F_1(x) \\ \vdots \\ \nabla F_m(x) \end{pmatrix}.$$

The *Hessian matrix* of the C^2 function $f : E^n \rightarrow E$ is

$$\nabla^2 f(x) = D(\nabla f(x)) = \begin{pmatrix} \frac{\partial^2 f}{\partial x_1^2}(x) & \cdots & \frac{\partial^2 f}{\partial x_n \partial x_1}(x) \\ \vdots & & \vdots \\ \frac{\partial^2 f}{\partial x_1 \partial x_n}(x) & \cdots & \frac{\partial^2 f}{\partial x_n^2}(x) \end{pmatrix}.$$

For open $U \subset E^n$, open $V \subset E^m$, $n > m$, a C^2 map $\rho : U \rightarrow V$ is said to be *transversal to zero* if $D\rho$ has full rank on $\rho^{-1}(0)$. Note that in the trivial case where $\rho^{-1}(0)$ is empty, ρ is trivially transversal to zero.

The theoretical foundation of probability-one homotopies (referred to in early work as the Chow-Yorke algorithm) was laid by Chow, Mallet-Paret, and Yorke [5], and the algorithm was immediately recast as a practical computational procedure by Watson [20]. The intent here is not to summarize or survey probability-one homotopy developments—see the survey papers [25] (early history), [26] (applications), and [33] for the latest numerical algorithms.

Depending on the context and intended use, the supporting theory is presented differently. The best formulation for the work here is contained in Lemmas 2.1–2.3 from [22], which are restated here for convenience.

LEMMA 2.1. *Let $\rho : E^m \times [0, 1) \times E^n \rightarrow E^n$ be a C^2 map which is transversal to zero, and define*

$$\rho_a(\lambda, z) = \rho(a, \lambda, z).$$

Then for almost all $a \in E^m$, the map ρ_a is also transversal to zero.

Lemma 2.1 is known as a parametrized Sard's theorem, and its significance is partially given by:

LEMMA 2.2. *In addition to the hypotheses of Lemma 2.1, suppose that for each $a \in E^m$ the system $\rho_a(0, z) = 0$ has a unique nonsingular solution $z^{(0)}$. Then for almost all $a \in E^m$ there is a smooth zero curve $\gamma \subset [0, 1) \times E^n$ of $\rho_a(\lambda, z)$, emanating from $(0, z^{(0)})$, along which the Jacobian matrix $D\rho_a(\lambda, z)$ has rank n . γ does not intersect itself or any other zero curves of ρ_a , does not bifurcate, has finite arc length in any compact subset of $[0, 1) \times E^n$, and either goes to infinity or reaches (has an accumulation point in) the hyperplane $\lambda = 1$.*

LEMMA 2.3. *Under the hypotheses of Lemma 2.2, if the zero curve γ is bounded, then it has an accumulation point $(1, \bar{z})$. Furthermore, if $\text{rank } D\rho_a(1, \bar{z}) = n$, then γ has finite arc length.*

Conceptually, how all this relates to optimization is as follows: (1) convert an optimization problem to a nonlinear system of n equations in n unknowns, $F(x) = 0$. (2) Construct a homotopy map $\rho_a(\lambda, x)$ satisfying the hypotheses of the above lemmas, and with $\rho_a(1, x) = F(x)$. (3) Track the zero curve γ of ρ_a from the known point $(0, z^{(0)})$ to a point $(1, \bar{x})$. \bar{x} then solves the original optimization problem. Each of these steps can be fraught with theoretical and computational difficulties, and homotopy algorithms are often considered (with some truth) more art than science. The third step, homotopy zero curve tracking, is close to routine, with robust, numerically stable mathematical software [33] being available. The homotopy construction step is definitely an art, but good maps ρ_a are known for large classes of problems, and several books exist on the topic [13], [3]. The first step, conversion of an optimization problem to a nonlinear system, is perhaps the least understood and most debatable. Why convert a difficult optimization problem into a (possibly even more) difficult nonlinear system? There are enough examples of such counterintuitive conversions being successful (e.g., Karmarkar's algorithm converts a linear program into a series of nonlinear programs) to keep the question open.

3. Unconstrained convex optimization. The simplest possible case, convex unconstrained optimization, is worth mentioning because it shows how everything should work in the ideal case. While a homotopy algorithm is not advocated for convex unconstrained optimization, it is nevertheless reassuring that the theory does cover this case elegantly.

THEOREM 3.1. *Let $f : E^n \rightarrow E$ be a C^3 convex map with a minimum at \tilde{x} , $\|\tilde{x}\| \leq M$. Then for almost all a , $\|a\| < M$, there is a zero curve γ of the homotopy map*

$$\rho_a(\lambda, x) = \lambda \nabla f(x) + (1 - \lambda)(x - a),$$

along which the Jacobian matrix $D\rho_a(\lambda, x)$ has full rank, emanating from $(0, a)$ and having an accumulation point $(1, \bar{x})$, where \bar{x} solves

$$\min_x f(x).$$

If the Hessian matrix $\nabla^2 f(\bar{x})$ is nonsingular, then γ has finite arc length.

Theorem 3.1 is proved in [24], but a sketch of the proof is repeated here for several reasons. First, it illustrates that a simple proof suffices for the unconstrained case. Second, this proof is a prototype for many homotopy convergence proofs. Often, the essence of a homotopy convergence theorem proof is to construct a map $\rho_a(\lambda, x)$ to which this proof applies, or to generalize the prototype proof to apply to a particular ρ_a .

Let (λ, x) be any point with $0 \leq \lambda < 1$ and $\|x\| = 3M$. Now $\|a\| < M$ and $\|\tilde{x}\| \leq M$ give

$$(x - \tilde{x})(x - a) > 0,$$

and the convexity of f at the minimum \tilde{x} gives

$$(x - \tilde{x})\nabla f(x) = (x - \tilde{x})(\nabla f(x) - \nabla f(\tilde{x})) \geq 0.$$

Combining these inequalities yields

$$(x - \tilde{x})[\lambda \nabla f(x) + (1 - \lambda)(x - a)] > 0,$$

which means that $\rho_a(\lambda, x) \neq 0$ for $0 \leq \lambda < 1$ and $\|x\| = 3M$. Hence γ is bounded, being contained in the solid cylinder $[0, 1] \times \{x \mid \|x\| \leq 3M\}$. The conclusion follows from Lemma 2.3. \square

The essence of the above proof is that the zero curve γ of $\rho_a(\lambda, x)$ emanating from the trivially found start point $(0, a)$ does not pierce the surface

$$[0, 1] \times \{x \mid \|x\| = r\}$$

of some sufficiently large (solid) cylinder containing $(0, a)$. Then γ must be contained inside the solid cylinder, hence bounded, and must therefore pierce (or at least accumulate at) the hyperplane $\lambda = 1$ at a point $(1, \bar{x})$. This prototype convergence proof reveals a fundamental difference between continuation, homotopy, and probability-one homotopy algorithms. For the former two, a convergence theorem would have to address the *existence* and *connectivity* of γ for $0 \leq \lambda \leq 1$, requiring assumptions beyond the mere boundedness of γ . In contrast, a probability-one homotopy convergence proof essentially amounts to proving the connected component of $\rho_a^{-1}(0)$ containing $(0, z^{(0)})$ is bounded. The other requirements—transversality of ρ , $\rho_a(0, z)$ being a trivial map, $\rho_a(1, z) = F(z)$ —are normally trivially satisfied by the construction of ρ . Finally, note that continuation and homotopy algorithms must typically explicitly deal with singularities along γ , whereas a well constructed probability-one zero curve γ has no singularities, theoretical or numerical.

4. Nonnegatively constrained convex optimization. Let $f : E^n \rightarrow E$ be a C^3 convex function, and say that f is *uniformly convex* if there exists $\nu > 0$ such that $x[\nabla^2 f(z)]x \geq \nu \|x\|^2$ for all $x, z \in E^n$. Consider next the constrained optimization problem

$$\min f(x) \quad \text{such that} \quad x \geq 0. \quad (4.1)$$

Since f is convex and Slater's constraint qualification is satisfied, the Kuhn-Tucker optimality conditions are both necessary and sufficient. Hence (4.1) is equivalent to the *nonlinear complementarity problem*

$$x \geq 0, \quad F(x) \geq 0, \quad x F(x) = 0, \quad (4.2)$$

where $F(x) = \nabla f(x)$. There are numerous ways to rewrite (4.2) as a nonsmooth [4] or smooth nonlinear system of equations, but the simplest way (meeting the C^2 requirement for smooth probability-one homotopies), due to Mangasarian [12], is as

$$K(x) = 0, \quad (4.3)$$

where

$$K_i(x) = -|F_i(x) - x_i|^3 + (F_i(x))^3 + x_i^3. \quad (4.4)$$

This choice for $K(x)$ permits the use of the canonical homotopy map

$$\rho_a(\lambda, x) = \lambda K(x) + (1 - \lambda)(x - a). \quad (4.5)$$

Since (4.1), (4.2), and (4.3) are equivalent, $\rho_a(1, \bar{x}) = 0$ gives a solution \bar{x} to (4.1).

A convergence theorem for (4.5) uses a general existence result for the nonlinear complementarity problem from [24]:

LEMMA 4.1. *Suppose every zero of $K(x)$ lies in the ball $\|x\| < r$, where r is such that $x \geq 0$ and $\|x\| \geq r$ imply $x_k > 0$ and $F_k(x) \geq 0$ for some index k . Then there exists $\delta > 0$ such that for almost all $a \geq 0$ with $\|a\| < \delta$ there is a zero curve γ of $\rho_a(\lambda, x)$, along which $D\rho_a(\lambda, x)$ has full rank, connecting $(0, a)$ to $(1, \bar{x})$, where \bar{x} is a zero of $K(x)$.*

This lemma directly gives the result (from [24]):

THEOREM 4.1. *Let $f : E^n \rightarrow E$ be a C^3 uniformly convex map. Then there exists $\delta > 0$ such that for almost all $a \geq 0$ with $\|a\| < \delta$ there is a zero curve γ of*

$$\rho_a(\lambda, x) = \lambda K(x) + (1 - \lambda)(x - a),$$

where

$$K_i(x) = -\left|\frac{\partial f(x)}{\partial x_i} - x_i\right|^3 + \left(\frac{\partial f(x)}{\partial x_i}\right)^3 + x_i^3,$$

along which $D\rho_a(\lambda, x)$ has full rank, connecting $(0, a)$ to a point $(1, \bar{x})$, where \bar{x} solves the constrained optimization problem (4.1).

Note that homotopy convergence theorems are often also *existence* theorems, as is the case with Theorem 4.1, and consequently the assumptions certainly cannot be weaker than are required for existence of a solution. The uniform convexity assumption

of Theorem 4.1 is one way to guarantee the existence of a solution to (4.1). If one *assumes* that (4.1) has a solution, then a theorem like the following is possible.

THEOREM 4.2. *Let $f : E^n \rightarrow E$ be a C^3 convex map, and assume that (4.1) has a solution \tilde{x} , and that the level sets of f are bounded. Then there exists $\delta > 0$ such that for almost all $a \geq 0$ with $\|a\| < \delta$ there is a zero curve γ of*

$$\rho_a(\lambda, x) = \lambda K(x) + (1 - \lambda)(x - a),$$

where

$$K_i(x) = - \left| \frac{\partial f(x)}{\partial x_i} - x_i \right|^3 + \left(\frac{\partial f(x)}{\partial x_i} \right)^3 + x_i^3,$$

along which $D\rho_a(\lambda, x)$ has full rank, emanating from $(0, a)$ and reaching a point $(1, \bar{x})$, where \bar{x} solves (4.1).

Proof. Since $K(x) = 0$ is equivalent to (4.2), which is equivalent to (4.1), it suffices to verify the hypotheses of Lemma 4.1 for the nonlinear complementarity problem with $F(x) = \nabla f(x)$. First note that by assumption the solutions of (4.1) are bounded, and therefore all the zeros of $K(x)$ lie in open some ball $B(M) = \{x \in E^n \mid \|x\| < M\}$. That is, every solution \tilde{x} of (4.1) satisfies $\|\tilde{x}\| < M$.

Observe that it suffices to consider only points (λ, x) with $0 \leq \lambda < 1$ and $x \geq 0$, since $x_i < 0, a_i \geq 0$ imply $K_i(x) < 0$ and $x_i - a_i < 0$, which then imply $[\rho_a(\lambda, x)]_i < 0$; hence $x \geq 0$ along the zero curve γ of ρ_a . $f(x)$ has a maximum at some point \hat{x} on the compact set

$$S_1 = \{x \in E^n \mid x \geq 0, \quad \|x\| = M\}.$$

By assumption, the level set

$$S_2 = \{y \in E^n \mid y \geq 0, \quad f(y) \leq f(\hat{x})\}$$

is contained in some closed ball $\{x \in E^n \mid \|x\| \leq r/2\}$. Since $\hat{x} \in S_1 \cap S_2$, $0 < M \leq r/2 < r$. Now consider any $z \geq 0$ with $\|z\| = r$. It follows that

$$f(z) > f(\hat{x}) \geq f\left(\frac{M}{r}z\right) > f(\tilde{x}),$$

and from the convexity of f ,

$$\begin{aligned} f\left(\frac{M}{r}z\right) &\geq f(z) + \nabla f(z) \left(\frac{M}{r}z - z\right) \Rightarrow \\ \left(1 - \frac{M}{r}\right) z \nabla f(z) &\geq f(z) - f\left(\frac{M}{r}z\right) > 0 \Rightarrow \\ z \nabla f(z) &> 0 \Rightarrow \end{aligned}$$

$z_k > 0$ and $(\nabla f(z))_k > 0$ for some index k , the requirement of Lemma 4.1. \square

Using concepts like recession cones and indicator functions from convex analysis [17], very short proofs can be given for the next two theorems. The essential fact from [17] is that if one nonempty level set is bounded, then all the level sets are bounded. In the interest of maintaining an elementary exposition, short direct proofs are given

here. A variant of Theorem 4.2 can be obtained without reference to level sets. One such possibility is the next theorem.

THEOREM 4.3. *Let $f : E^n \rightarrow E$ be a C^3 convex map, and assume that (4.1) has a solution \tilde{x} at which f is strictly convex. Then the conclusion of Theorem 4.2 holds.*

Proof. \tilde{x} is the unique minimum point from strict convexity, and hence the zeros of $K(x)$ are bounded. The proof of Theorem 4.2 applies if it can be shown that the level set S_2 is bounded. Suppose not. Then there exists a sequence $y^{(k)} \in S_2$ with $\|y^{(k)}\| \rightarrow \infty$. The vectors $y^{(k)}/\|y^{(k)}\|$ lie on the compact unit sphere, and therefore have a convergent subsequence $y^{(k_i)}/\|y^{(k_i)}\| \rightarrow y \geq 0$. Reduce to this subsequence. For each k , choose $0 < t_k < 1$ such that

$$\left\| (1 - t_k)\tilde{x} + t_k y^{(k)} \right\| = M. \quad (4.6)$$

Now by the strict convexity of f at \tilde{x} ,

$$f\left((1 - t_k)\tilde{x} + t_k y^{(k)}\right) < (1 - t_k)f(\tilde{x}) + t_k f\left(y^{(k)}\right) \leq (1 - t_k)f(\tilde{x}) + t_k f(\hat{x}). \quad (4.7)$$

Taking the limit as $k \rightarrow \infty$ ($\|y^{(k)}\| \rightarrow \infty$ and (4.6) give $t_k \rightarrow 0$) yields

$$f(\tilde{x} + \alpha y) \leq f(\tilde{x}), \quad (4.8)$$

where $\tilde{x} + \alpha y \geq 0$, $\|\tilde{x} + \alpha y\| = M > \|\tilde{x}\|$, which contradicts the strict convexity of f at the minimum point \tilde{x} . \square

The most general version of the homotopy convergence theorem for (4.1), whose proof is a refinement of the previous proof, is given last. Theorems 4.2 and 4.3 could have been dispensed with, but presenting the proofs as a series of refinements is instructive.

THEOREM 4.4. *Let $f : E^n \rightarrow E$ be a C^3 convex map, assume that (4.1) has a solution \tilde{x} , and that every solution \tilde{x} satisfies $\|\tilde{x}\| < M$. Then the conclusion of Theorem 4.2 holds.*

Proof. By assumption the zeros of $K(x)$ are bounded. The proof of Theorem 4.3, after the first sentence, applies verbatim, with the following changes. Without strict convexity, the strict inequality in (4.7) becomes inequality (\leq), but this doesn't matter in the limit. (4.8) still obtains, but now the contradiction is that $\tilde{x} + \alpha y$ is also a solution of (4.1), which does not satisfy $\|\tilde{x} + \alpha y\| < M$. \square

5. Linearly constrained convex optimization. Let $f : E^n \rightarrow E$ be a C^3 convex function, let $A \in E^{m \times n}$, $b \in E^m$. First consider the problem

$$\min f(x) \quad \text{subject to} \quad g(x) = Ax - b \leq 0. \quad (5.1)$$

Since both f and g are convex and g satisfies the Arrow-Hurwicz-Uzawa constraint qualification, (5.1) is equivalent to the Kuhn-Tucker problem

$$(\nabla f(x))^t + A^t u = 0, \quad (5.2)$$

$$Ax - b \leq 0, \quad (5.3)$$

$$u \geq 0, \quad (5.4)$$

$$u^t(Ax - b) = 0. \quad (5.5)$$

As before, the complementarity conditions (5.3)–(5.5) can be replaced by a nonlinear system $K(x, u) = 0$, defined by

$$K_i(x, u) = -|b_i - A_i \cdot x - u_i|^3 + (b_i - A_i \cdot x)^3 + u_i^3.$$

One possible homotopy map, that has been successful in practice on some difficult engineering optimization problems [30], is

$$\rho_a(\lambda, x, u) = \lambda \begin{pmatrix} (\nabla f(x))^t + A^t u \\ K(x, u) \end{pmatrix} + (1 - \lambda) \begin{pmatrix} x - x^0 \\ u - u^0 \end{pmatrix}, \quad (5.6)$$

where $a = (x^0, u^0)$ is the random probability-one homotopy parameter vector, and $u^0 > 0$. This is the direct generalization of what was done for simple nonnegativity constraints $x \geq 0$, and one would expect it to work. For instance, it is known that a quadratic programming problem with general linear inequality constraints is equivalent to a quadratic programming problem with only nonnegativity constraints. Unfortunately, the homotopy map (5.6) does not suffice. Qualitatively, it worked before because (with convex f) $K(x, u)$ and $u - u^0$ had the same sign for large arguments, thus $\rho_a(\lambda, x, u)$ could not be zero outside some large ball. This meant the zero curve γ of ρ_a could not penetrate the surface of that ball, and hence had to reach a solution of the original problem. In (5.6), x and u can play off against other permitting $\rho_a^{-1}(0)$ to be unbounded.

As a simple example, consider $f(x) = (1/2)x^2$ and $g(x) = 1 - x \leq 0$, and take $x^0 = -1$, $u^0 = 0.1$. The zero curve γ of the homotopy map ρ_a in (5.6) is unbounded, as shown in Fig. 1. Of course a lucky guess for (x^0, u^0) may still work. For $(x^0, u^0) = (0, -1)$, γ has several turning points but still reaches $\lambda = 1$ in finite arc length.




FIG. 1. Example of unbounded homotopy zero curve.

The idea behind the repair of (5.6) is to replace (5.3) by

$$Ax - b - (1 - \lambda)b^0 \leq 0, \quad (5.7)$$

and because of the technical necessity to preserve transversality, replace $\lambda K(x, u) + (1 - \lambda)(u - u^0)$ by

$$K(\lambda, x, u) - (1 - \lambda)c^0. \quad (5.8)$$

Assume that the feasible set $\{x \mid Ax - b \leq 0\}$ is nonempty and bounded (this is not an obstacle in practice, because variable bounds can always be added). For some arbitrary initial guess $x^0 \in E^n$, choose $b^0 \in E^m$, $b^0 > 0$ such that $Ax^0 - b - b^0 < 0$. Also choose $c^0 \in E^m$ such that $c^0 > 0$. Define $S_\lambda = \{x \mid Ax - b - (1 - \lambda)b^0 \leq 0\}$, and observe that

$$S_0 \supset S_{\lambda_1} \supset S_{\lambda_2} \supset S_1 \neq \emptyset \text{ for } 0 \leq \lambda_1 < \lambda_2 \leq 1. \quad (5.9)$$

It would be desirable if complementarity could be automatically enforced by defining K with something like

$$\begin{aligned} K_i(\lambda, x, u, b^0, c^0) = & -|(1 - \lambda)b_i^0 + b_i - A_i \cdot x - (u_i - (1 - \lambda)c_i^0)|^3 \\ & + ((1 - \lambda)b_i^0 + b_i - A_i \cdot x)^3 + (u_i - (1 - \lambda)c_i^0)^3, \quad i = 1, \dots, m, \end{aligned} \quad (5.10)$$

and then

$$\rho(x^0, b^0, c^0, \lambda, x, u) = \begin{pmatrix} \lambda [(\nabla f(x))^t + A^t u] + (1 - \lambda)(x - x^0) \\ K(\lambda, x, u, b^0, c^0) \end{pmatrix}. \quad (5.11)$$

As always for probability-one homotopies, technically $0 \leq \lambda < 1$. Unfortunately this K results in ρ in (5.11) not being transversal to zero, due to inherent cancellation in the structure of K in (5.10). Something a bit more complicated is required, such as

$$\begin{aligned} K_i(\lambda, x, u, b^0, c^0) = & -|(1 - \lambda)b_i^0 + b_i - A_i \cdot x - u_i|^3 + ((1 - \lambda)b_i^0 + b_i - A_i \cdot x)^3 \\ & + u_i^3 - (1 - \lambda)c_i^0, \quad i = 1, \dots, m. \end{aligned} \quad (5.12)$$

The complication is that given $c^0 > 0$ and $(1 - \lambda)b^0 + b - Ax^0 > 0$, some work is required to find the starting point at $\lambda = 0$: the value u^0 for u such that $K(0, x^0, u^0, b^0, c^0) = 0$. However, it is easily verified that K_i is a strictly monotone increasing function of u_i , and thus u^0 can always be uniquely determined.

Let $a = (x^0, b^0, c^0)$, and define $\rho_a(\lambda, x, u) = \rho(x^0, b^0, c^0, \lambda, x, u)$, using (5.11) and (5.12).

THEOREM 5.1. *Let $f : E^n \rightarrow E$ be a C^3 convex function, let $A \in E^{m \times n}$, $b \in E^m$, and assume that $S_1 = \{x \mid Ax - b \leq 0\}$ is nonempty and bounded. Let $\rho_a(\lambda, x, u) = \rho(x^0, b^0, c^0, \lambda, x, u)$ be defined from (5.11) and (5.12). Then for almost all $x^0 \in E^n$, almost all $b^0 \in E^m$ such that $b^0 > 0$ and $Ax^0 - b - b^0 < 0$, and almost all $c^0 \in E^m$ with $c^0 > 0$, there exists a zero curve γ of $\rho_a(\lambda, x, u)$ emanating from $(0, x^0, u^0)$, along which the Jacobian matrix $D\rho_a(\lambda, x, u)$ has rank $n + m$, and reaching a point $(1, \bar{x}, \bar{u})$, where \bar{x} solves $\min_{x \in S_1} f(x)$. If $\text{rank } D\rho_a(1, \bar{x}, \bar{u}) = n + m$, then γ has finite arc length.*

Proof. Several facts need to be verified first.

(i) S_λ is nonempty and bounded for $0 \leq \lambda \leq 1$. This follows from the assumption that $S_1 \neq \emptyset$, (5.9), and the fact that boundedness is unrelated to the constant term $b + (1 - \lambda)b^0$: S_λ is bounded if and only if $\{x \mid Ax \leq 0\} = \{0\}$. Furthermore, observe that $\text{int } S_\lambda \neq \emptyset$ for $0 \leq \lambda < 1$.

(ii) ρ defined from (5.11) and (5.12) is transversal to zero, for $0 \leq \lambda < 1$, and $c^0 > 0$. It is easily verified that full rank $(n + m)$ comes from the D_{x^0} and D_{c^0} columns. This is a good illustration of the fact that the dimension of the probability-one homotopy parameter vector $a = (x^0, b^0, c^0)$ need not equal the dimension of the homotopy map ρ , and how this flexibility can be used to advantage.

(iii) There is a unique point (x^0, u^0) such that $\rho_a(0, x, u) = 0$. For $\lambda = 0$, clearly $x = x^0$ from (5.11). Now given x^0, b^0 such that $b^0 + b - Ax^0 > 0$, and $c^0 > 0$, it can be verified that $K = 0$ from (5.12) has a unique solution $u^0 > 0$.

By Lemma 2.1, $\rho_a(\lambda, x, u)$ is also transversal to zero for almost all $a = (x^0, b^0, c^0) \in E^n \times E^m \times (0, \infty)^m$. The statement of the theorem restricts b^0 to depend on x^0 , but this is immaterial to the transversality of ρ_a , since the full rank of $D\rho$ does not depend on D_{b^0} . The set of all (x^0, b^0, c^0) described in the theorem is open, and ρ_a is transversal to zero for almost all points a in this set.

From (iii), there is exactly one solution (x^0, u^0) to $\rho_a = 0$ at $\lambda = 0$. Therefore Lemma 2.2 applies, and the existence (for almost all the prescribed points (x^0, b^0, c^0)) of a zero curve γ and the full rank of $D\rho_a$ along γ follow. γ emanates from the point $(0, x^0, u^0)$, and either reaches a solution point $(1, \bar{x}, \bar{u})$, or wanders off to infinity. By Lemma 2.3, it suffices to prove that γ is bounded. The finite arc length statement about γ also follows from Lemma 2.3.

(iv) Consider an arbitrary point (λ, x, u) on γ for $0 < \lambda < 1$. A careful examination of the signs of the terms in K_i in (5.12) reveals that $K_i < 0$ if $u_i < 0$ or $(1 - \lambda)b_i^0 + b_i - A_i \cdot x < 0$. Therefore everywhere along γ , $u \geq 0$ and $x \in S_\lambda \subset S_0$ is bounded from (i) and (5.9). Furthermore, $u_i > 0$ and $(1 - \lambda)b_i^0 + b_i - A_i \cdot x > 0$ along γ .

Suppose that γ is not bounded, and let $(\lambda^{(k)}, x^{(k)}, u^{(k)}) \rightarrow \infty$ be a sequence of points on γ . Since $[0, 1] \times S_0$ is compact, $\{(\lambda^{(k)}, x^{(k)})\}_{k=1}^\infty$ has a convergent subsequence $(\lambda^{(k_i)}, x^{(k_i)}) \rightarrow (\hat{\lambda}, \hat{x})$. Now from (5.11), this means that $u^{(k_i)} \geq 0$, $\|u^{(k_i)}\| \rightarrow \infty$, and

$$\hat{\lambda} \left[(\nabla f(\hat{x}))^t + A^t u^{(k_i)} \right] + (1 - \hat{\lambda})(\hat{x} - x^0) \rightarrow 0, \quad (5.13)$$

$$K(\hat{\lambda}, \hat{x}, u^{(k_i)}, b^0, c^0) \rightarrow 0. \quad (5.14)$$

If $\hat{\lambda} = 1$, then \hat{x} is a solution to (5.1), and \bar{u} corresponding to $\bar{x} = \hat{x}$ can be constructed from $u^{(k_i)}$. In a degenerate case, γ converges to a (possibly unbounded) manifold of solution points $(1, \bar{x}, \bar{u})$. So now consider $0 \leq \hat{\lambda} < 1$ and two cases.

Case 1: For some j , $(1 - \hat{\lambda})b_j^0 + b_j - A_j \hat{x} > 0$ and $\limsup_{k_i \rightarrow \infty} u_j^{(k_i)} = \infty$. As observed earlier, K_j is a strictly monotone increasing function of u_j . Therefore $\|u^{(k_i)}\| \rightarrow \infty$ implies $\left\| K(\hat{\lambda}, \hat{x}, u^{(k_i)}, b^0, c^0) \right\|_\infty$ is increasing, which contradicts (5.14). Therefore γ is bounded and the theorem follows.

Case 2: $(1 - \hat{\lambda})b_j^0 + b_j - A_j \hat{x} = 0$ for every j with $\limsup_{k_i \rightarrow \infty} u_j^{(k_i)} = \infty$; denote this set of indices by J . Suppose first that $\hat{\lambda} = 0$. Then from (5.13), $A^t \lambda^{(k_i)} u^{(k_i)} \rightarrow x^0 - \hat{x}$. A subsequence argument yields a vector w such that $(A_J)^t w = x^0 - \hat{x}$, $w \geq 0$. Combining the relations $Ax^0 - b - b^0 < 0$ and $(b^0 + b - A\hat{x})_J = 0$ gives $A_J \cdot (x^0 - \hat{x}) < 0$. Now all these relations result in

$$0 \geq (x^0 - \hat{x})^t (A_J)^t w = (x^0 - \hat{x})^t (x^0 - \hat{x}) > 0,$$

a contradiction. Therefore $\hat{\lambda} \neq 0$. As observed in item (i), $\{x \mid Ax \leq 0\} = \{0\}$, which is equivalent to the positive cone $\mathcal{C}(A^t) = \{A^t y \mid y \geq 0\} = E^n$. Therefore there exists w such that

$$A^t w = -(\nabla f(\hat{x}))^t - (1 - \hat{\lambda})(\hat{x} - x^0)/\hat{\lambda}, \quad w \geq 0.$$

Writing $u^{(k_i)} = w + v^{(k_i)}$ then gives $A^t v^{(k_i)} \rightarrow 0$ and $\|v^{(k_i)}\| \rightarrow \infty$. Recall that $u^{(k_i)} = w + v^{(k_i)} \geq 0$, which means that any negative components of $v^{(k_i)}$ must be bounded by $\|w\|_\infty$, and therefore negative components of $v^{(k_i)} / \|v^{(k_i)}\|_\infty$ converge to zero as $\|v^{(k_i)}\| \rightarrow \infty$. The bounded sequence $\{v^{(k_i)} / \|v^{(k_i)}\|_\infty\}_{i=1}^\infty$ has a subsequence converging to some point $v \in E^m$ with $\|v\|_\infty = 1$, and $v \geq 0$ by the preceding remark. Finally $A^t v^{(k_i)} \rightarrow 0$ and $\|v^{(k_i)}\| \rightarrow \infty$ imply $A^t (v^{(k_i)} / \|v^{(k_i)}\|_\infty) \rightarrow 0$ yielding $A^t v = 0$, $v \geq 0$, $\|v\|_\infty = 1$, or $(A_J)^t v_J = 0$, $v_J \geq 0$. By Gordan's Theorem of the Alternative [11], $A_J z > 0$ has no solution. However, since $\text{int } S_{\hat{\lambda}} \neq \emptyset$, there is an interior feasible point x so that combining the relations $Ax - b - (1 - \hat{\lambda})b^0 < 0$ and $((1 - \hat{\lambda})b^0 + b - A\hat{x})_J = 0$ yields $A_J \cdot (x - \hat{x}) < 0$. This contradiction proves that γ is bounded, and the theorem follows. \square

COROLLARY 5.1. *Suppose that the zero curve γ of $\rho_a(\lambda, x, u)$ defined from (5.11)–(5.12) has the property that $\|x\|$ is bounded but $\|u\| \rightarrow \infty$ along γ , for $0 \leq \lambda < 1$. Then there exists $v \in E^m$ such that $A^t v = 0$, $v \geq 0$, and every index i for which $v_i \neq 0$ has $|u_i| \rightarrow \infty$.*

It is instructive to consider why the proof of Theorem 5.1 worked, and how it could have gone wrong. After all, homotopy methods don't always work, and the conclusion of Corollary 5.1 is a plausible situation. The structure of ρ in (5.11)–(5.12) is important for transversality (Lemma 2.2) and the fact that γ (for this particular ρ) cannot return to $\lambda = 0$. Yet many other choices for ρ could possess these properties equally well. The boundedness of x along γ is especially opportune, being a direct consequence of the relaxation (5.7) and the ‘‘interior’’ map (5.8), which forces points (λ, x, u) along γ to be strictly feasible (interior) for *different* constraints (from the original ones). Effectively λ and x are under control (not so if S_λ were unbounded),

and $\|u\| \rightarrow \infty$ is the only potential problem. Controlling $\|u\|$ is delicate, and was achieved here by having $\text{int } S_\lambda \neq \emptyset$ for $0 \leq \lambda < 1$ (trivially true for the $g(x)$ in (5.1) but not so trivial in more general contexts). $\|u\|$ was controlled in earlier sections by a global monotonicity property, a much stronger condition than S_λ bounded with nonempty interior for $0 \leq \lambda < 1$.

6. General nonlinear convex optimization. The optimization problem (5.1) can be generalized in several different directions. If the convexity assumption is dropped for either f or g , then (5.2)–(5.5) become only necessary conditions. Certainly many optimization algorithms are based on necessary optimality conditions, in which case only convergence to a stationary point is guaranteed. This particular direction of generalization will be pursued in a later section. Here in this section the intent is to preserve the optimality conditions being both necessary and sufficient, and thus $g(x)$ in (5.1) will be generalized.

Let $f : E^n \rightarrow E$ and $g : E^n \rightarrow E^m$ be C^3 convex functions, and assume that g satisfies the Arrow-Hurwicz-Uzawa constraint qualification at every solution of

$$\min f(x) \quad \text{subject to} \quad g(x) \leq 0. \quad (6.1)$$

Under these assumptions (6.1) is equivalent to the Kuhn-Tucker problem

$$(\nabla f(x))^t + (\nabla g(x))^t u = 0, \quad (6.2)$$

$$g(x) \leq 0, \quad (6.3)$$

$$u \geq 0, \quad (6.4)$$

$$u^t g(x) = 0. \quad (6.5)$$

Given the discussion in the last section, it seems reasonable to try the direct generalization of (5.11) for the homotopy map

$$\rho(x^0, b^0, c^0, \lambda, x, u) = \begin{pmatrix} \lambda \left[(\nabla f(x))^t + (\nabla g(x))^t u \right] + (1 - \lambda)(x - x^0) \\ K(\lambda, x, u, b^0, c^0) \end{pmatrix}, \quad (6.6)$$

where

$$K_i(\lambda, x, u, b^0, c^0) = -|(1 - \lambda)b_i^0 - g_i(x) - u_i|^3 + ((1 - \lambda)b_i^0 - g_i(x))^3 + u_i^3 - (1 - \lambda)c_i^0, \quad i = 1, \dots, m, \quad (6.7)$$

is the direct generalization of (5.12). The question is how changing from linear constraints $Ax - b \leq 0$ to nonlinear convex constraints $g(x) \leq 0$ affects the proof of Theorem 5.1. The crux of the proof seems to be the sets

$$S_\lambda = \{x \in E^n \mid g(x) - (1 - \lambda)b^0 \leq 0\}, \quad (6.8)$$

which need to satisfy (5.9), $\text{int } S_\lambda \neq \emptyset$ for $0 \leq \lambda < 1$, and be bounded for $0 \leq \lambda \leq 1$. Given some arbitrary initial guess $x^0 \in E^n$, choose $b^0 \in E^m$, $b^0 > 0$ such that $g(x^0) - b^0 < 0$. As before choose $c^0 \in E^m$ such that $c^0 > 0$. Assuming that the feasible set $S_1 = \{x \in E^n \mid g(x) \leq 0\}$ is nonempty and bounded is not a severe

restriction, since for any practical problem variable bounds can always be added. The boundedness of the sets S_λ for $0 \leq \lambda \leq 1$ follows from Corollary 8.3.3, Theorem 8.4, and Theorem 8.7 of [17]. For completeness a short direct proof follows.

LEMMA 6.1. *Let $g : E^n \rightarrow E^m$ be a C^3 convex function, and let $x^0 \in E^n$, $b^0 \in E^m$, $\delta \in E$ be such that $b^0 \geq \delta e > 0$ and $g(x^0) - b^0 < 0$. Define*

$$S_\lambda = \{x \in E^n \mid g(x) - (1 - \lambda)b^0 \leq 0\}.$$

If S_1 is nonempty and bounded, then S_λ is nonempty and bounded for $0 \leq \lambda \leq 1$, and $\text{int } S_\lambda \neq \emptyset$ for $0 \leq \lambda < 1$.

Proof. Since $S_1 \neq \emptyset$,

$$S_0 \supset S_{\lambda_1} \supset S_{\lambda_2} \supset S_1 \neq \emptyset \text{ for } 0 \leq \lambda_1 < \lambda_2 \leq 1.$$

Suppose that S_{λ_1} is unbounded while S_{λ_2} is bounded, say $S_{\lambda_2} \subset B(r/2) = \{x \in E^n \mid \|x\| < r/2\}$. Pick any point $\tilde{x} \in S_{\lambda_2}$ (hence $\|\tilde{x}\| < r/2$). Now there exists a sequence $y^{(k)} \in S_{\lambda_1}$ with $r < \|y^{(k)}\| \rightarrow \infty$. Reduce to a convergent subsequence $y^{(k_i)} / \|y^{(k_i)}\| \rightarrow y$, and for each subsequence index k choose $0 < t_k < 1$ such that

$$\left\| (1 - t_k)\tilde{x} + t_k y^{(k)} \right\| = r. \quad (6.9)$$

Now by the convexity of g ,

$$\begin{aligned} g\left((1 - t_k)\tilde{x} + t_k y^{(k)}\right) &\leq (1 - t_k)g(\tilde{x}) + t_k g\left(y^{(k)}\right) \\ &\leq (1 - t_k)(1 - \lambda_2)b^0 + t_k(1 - \lambda_1)b^0. \end{aligned}$$

Taking the limit as $k \rightarrow \infty$ ($\|y^{(k)}\| \rightarrow \infty$ and (6.9) give $t_k \rightarrow 0$) yields

$$g(\tilde{x} + \alpha y) \leq (1 - \lambda_2)b^0 \quad (6.10)$$

for $\|\tilde{x} + \alpha y\| = r$. Now (6.10) $\implies \tilde{x} + \alpha y \in S_{\lambda_2} \implies r = \|\tilde{x} + \alpha y\| < r/2$, a contradiction. Therefore S_λ must be nonempty and bounded for all $0 \leq \lambda \leq 1$. The statement about $\text{int } S_\lambda$ follows easily by continuity. \square

Let $a = (x^0, b^0, c^0)$, and define $\rho_a(\lambda, x, u) = \rho(x^0, b^0, c^0, \lambda, x, u)$, according to (6.6) and (6.7). As before, u^0 is uniquely defined by $K(0, x^0, u^0, b^0, c^0) = 0$.

THEOREM 6.1. *Let $f : E^n \rightarrow E$ and $g : E^n \rightarrow E^m$ be C^3 convex functions, let g satisfy the Arrow-Hurwicz-Uzawa constraint qualification at every solution of (6.1), and assume that $S_1 = \{x \in E^n \mid g(x) \leq 0\}$ is nonempty and bounded. Let $\rho_a(\lambda, x, u) = \rho(x^0, b^0, c^0, \lambda, x, u)$ be defined from (6.6) and (6.7). Then for almost all $x^0 \in E^n$, almost all $b^0 \in E^m$ such that $b^0 > 0$ and $g(x^0) - b^0 < 0$, and almost all $c^0 \in E^m$ with $c^0 > 0$, there exists a zero curve γ of $\rho_a(\lambda, x, u)$ emanating from $(0, x^0, u^0)$, along which the Jacobian matrix $D\rho_a(\lambda, x, u)$ has rank $n + m$, reaching a point $(1, \bar{x}, \bar{u})$, where \bar{x} solves $\min_{x \in S_1} f(x)$. If $\text{rank } D\rho_a(1, \bar{x}, \bar{u}) = n + m$, then γ has finite arc length.*

Proof. By the convexity and constraint qualification assumptions, (6.1) is equivalent to (6.2)–(6.5), which are equivalent to $\rho(x^0, b^0, c^0, 1, x, u) = 0$. A careful examination of the proof of Theorem 5.1 reveals that it is valid if (a) $Ax - b$ is replaced by $g(x)$, (b) ρ and K from (5.11)–(5.12) are replaced by ρ and K from (6.6)–(6.7), (c) the sets S_λ from (6.8) are nonempty and bounded for $0 \leq \lambda \leq 1$ and have nonempty interiors for $0 \leq \lambda < 1$, and (d) the conclusion of Corollary 5.1 also leads to a contradiction in the present more general situation. Take each item in turn.

(a) Replacing the function $Ax - b$ by the function $g(x)$ affects nothing in the proof of Theorem 5.1. The appearance of $\nabla g(x)$ rather than (constant) A in ρ does have some effect. The argument ruling out the possibility $\hat{\lambda} = 0$ becomes $\nabla g_J(\hat{x})(x^0 - \hat{x}) \leq g_J(x^0) - g_J(\hat{x}) < 0$, using the convexity of g . Arguments involving (5.13) and $\mathcal{C}(A^t)$ are valid with $\nabla g(\hat{x})$ replacing A . The final contradiction in Case 2 is addressed below in item (d).

(b) The transversality and other fundamental properties of ρ and K from (6.6)–(6.7) are easily verified.

(c) Lemma 6.1 provides these crucial facts about the sets S_λ .

(d) The question here is given $(\nabla g_J(\hat{x}))^t v_J = 0$, $v_J \geq 0$, and $\text{int } S_{\hat{\lambda}} \neq \emptyset$, does the same contradiction ensue by finding a vector z such that $(\nabla g_J(\hat{x}))z < 0$. The answer is yes since $x \in \text{int } S_{\hat{\lambda}}$ gives $g(x) - (1 - \hat{\lambda})b^0 < 0$, and then subtracting $(g(\hat{x}) - (1 - \hat{\lambda})b^0)_J = 0$ gives

$$\nabla g_J(\hat{x})(x - \hat{x}) \leq (g(x) - g(\hat{x}))_J < 0. \quad (6.11)$$

□

A result similar to Corollary 5.1 could be derived, but the conclusion — $(\nabla g(\hat{x}))^t v = 0$, $v \geq 0$ has a solution — is not interesting since the point \hat{x} has no special significance.

The post-mortem comments on the proof of Theorem 5.1 apply also to the proof of Theorem 6.1, with the latter being technically (but not conceptually) more difficult. Properties of the sets S_λ required proof, and γ again could not return to $\lambda = 0$. Convexity of g was crucial in converting constraint values into a gradient inequality as in (6.11), both for obtaining $\hat{\lambda} \neq 0$ and for the final contradiction (6.11). Neither quasiconvexity nor pseudoconvexity suffice for g .

7. Nonconvex programs. The convergence theory in the preceding sections might, at first glance, seem trivial and contrived (with the assumptions dictated more by the exigencies of the proof rather than by practical applications), and not to address the homotopy maps actually used on practical engineering problem. Indeed, homotopy maps like (4.5), (5.11), and (6.6), although they work, are rarely used in practice. There are two significant questions to be answered: (1) how important is convexity, which has figured prominently in the discussion so far? (2) How important is it that homotopy maps as in (4.5), (5.11), or (6.6) be used? To both questions, the answer turns out to be: not very!

Convexity simplifies proofs, but is really only needed to make the Kuhn-Tucker conditions sufficient for optimality. Without convexity, convergence only to a stationary point can be guaranteed. The proofs in §§ 3 and 4 used convexity, but only because those results were for the canonical map $\lambda F(x) + (1 - \lambda)(x - a)$, and were done using a nonlinear complementarity result that depended on convexity (pseudoconvexity). Note, for instance, that the homotopy map (6.6) is not of the canonical form $\lambda F(x) + (1 - \lambda)(x - a)$, which is the map used for Brouwer fixed point problems $x = f(x)$ (where $F(x) = x - f(x)$) [5]. This particular (canonical) map, which unfortunately is often thought of as “the” homotopy map, is only appropriate when F comes from a fixed point map, or has some sort of global monotonicity property like $x F(x) \geq 0$ for all $\|x\| \geq r$ for some sufficiently large $r > 0$. Convexity (pseudoconvexity) is sufficient, but not necessary, for such global monotonicity, and hence is a natural assumption when using the map $\lambda F(x) + (1 - \lambda)(x - a)$. The theory in §§ 3 and 4 could be generalized to assume something like “ $f(x)$ acts like a pseudoconvex function for $\|x\| \geq r$ sufficiently large,” but it hardly seems worth the trouble, since (as will be shown) homotopy maps like (6.6) obviate the need for convexity (when abandoning *sufficient* conditions for optimality!). §§ 5 and 6 used convexity to derive properties like γ can not return to $\lambda = 0$, $\nabla g_J(\hat{x})z < 0$ has a solution z , and the boundedness of the sets S_λ in (6.8) for $0 \leq \lambda \leq 1$. Convexity is overkill, though, and these much weaker properties can be explicitly assumed.

Let $f : E^n \rightarrow E$ and $g : E^n \rightarrow E^m$ be C^3 functions, and assume that g satisfies the Arrow-Hurwicz-Uzawa constraint qualification at every local solution of

$$\min f(x) \quad \text{subject to} \quad g(x) \leq 0. \quad (7.1)$$

If \bar{x} solves (7.1) locally, then there exists $\bar{u} \in E^m$ such that (\bar{x}, \bar{u}) solves the Kuhn-Tucker problem

$$(\nabla f(x))^t + (\nabla g(x))^t u = 0, \quad (7.2)$$

$$g(x) \leq 0, \quad (7.3)$$

$$u \geq 0, \quad (7.4)$$

$$u^t g(x) = 0. \quad (7.5)$$

Let $F : E^n \times [0, 1] \rightarrow E$ and $G : E^n \times [0, 1] \rightarrow E^m$ be C^3 functions such that

$$F(x, 1) = f(x), \quad G(x, 1) = g(x), \quad (7.6)$$

and the optimization problem

$$\min F(x, 0) \quad \text{subject to} \quad G(x, 0) \leq 0 \quad (7.7)$$

has an easily obtained (local) solution x^0 . In practice $F(x, \lambda)$, $G(x, \lambda)$ represent a family of optimization problems

$$\min F(x, \lambda) \quad \text{subject to} \quad G(x, \lambda) \leq 0, \quad (7.8)$$

where λ is embedded deeply and nonlinearly in the objective function $F(x, \lambda)$ and constraints $G(x, \lambda)$. This embedding often embodies considerable physical insight into the problem (7.1), and (7.7) is a version of (7.1) with simplified physics and/or geometry. A good choice for (7.8) may take years to develop, and generally requires considerable problem specific knowledge and the intimate involvement of an engineer or scientist. The payoff will be a robust, globally convergent algorithm that is more efficient than applying an “off-the-shelf” algorithm, and avoids spurious solutions (e.g., unstable equilibria in mechanics or unstable circuit operating points can be expressly avoided).

One could naively solve (7.8) with continuation varying λ from 0 to 1, but this is precisely the point at which the probability-one theory can make a significant improvement over simple continuation in λ (and also over arc length continuation). A probability-one homotopy for (7.8) guarantees the existence of a zero curve γ with good numerical properties, the importance of which for practical computation cannot be overstated. The homotopy map (6.6) is generalized to

$$\rho(x^0, b^0, c^0, \lambda, x, u) = \begin{pmatrix} \lambda \left[(\nabla_x F(x, \lambda))^t + (\nabla_x G(x, \lambda))^t u \right] + (1 - \lambda)(x - x^0) \\ K(\lambda, x, u, b^0, c^0) \end{pmatrix}, \quad (7.9)$$

where

$$K_i(\lambda, x, u, b^0, c^0) = -|(1 - \lambda)b_i^0 - G_i(x, \lambda) - u_i|^3 + ((1 - \lambda)b_i^0 - G_i(x, \lambda))^3 + u_i^3 - (1 - \lambda)c_i^0, \quad i = 1, \dots, m, \quad (7.10)$$

is the direct generalization of (6.7). The map (7.9), or some minor variation thereof, is what is typically used in practice, and has been extremely successful on industrial optimization problems.

The key observation in the proofs in §§ 5 and 6 is that *what matters most is not the structure of the homotopy map ρ , but the nature of the sets S_λ* . (Of course, ρ still has to satisfy the hypotheses of Lemma 2.3, and some technical conditions along γ are required.) The general convergence theory is now developed.

Given some arbitrary initial guess $x^0 \in E^n$, choose $b^0 \in E^m$ such that $b^0 > 0$ and $G(x^0, 0) - b^0 < 0$. Choose $c^0 \in E^m$ such that $c^0 > 0$. Consider the sets

$$S_\lambda = \{x \in E^n \mid G(x, \lambda) - (1 - \lambda)b^0 \leq 0\}, \quad 0 \leq \lambda \leq 1. \quad (7.11)$$

Note that $S_\lambda \neq \emptyset$ for small λ since $x^0 \in \text{int } S_0$. However, since the constraints $G(x, \lambda)$ now can change with λ , (5.9) need not hold, i.e., the sets S_λ do not necessarily form

a chain $S_0 \supset S_{\lambda_1} \supset S_{\lambda_2}$ for $0 < \lambda_1 < \lambda_2$. The question is exactly what properties must S_λ have in order for the proofs of Theorems 5.1 and 6.1 to extend to the general nonconvex problem (7.1)? It is not necessary for the sets S_λ to form a chain as in (5.9), or even to satisfy $\bigcap_{0 \leq \lambda \leq 1} S_\lambda \neq \emptyset$. Certainly each S_λ must be nonempty, otherwise K

from (7.10) cannot possibly be zero: $S_\lambda = \emptyset$ implies for each $x \in E^n$ there is an index i such that $(1 - \lambda)b_i^0 - G_i(x, \lambda) < 0$, which means for all x some $K_i(\lambda, x, u, b^0, c^0) < 0$ for all u , and thus $\rho_a(\lambda, x, u) \neq 0$ for any x, u .

A point (λ, x, u) on the zero curve γ of $\rho_a(\lambda, x, u)$ must have $x \in S_\lambda$ and $u \geq 0$ (otherwise $K \neq 0$), but S_λ bounded for $0 \leq \lambda \leq 1$ (the conclusion of Lemma 6.1) does not imply x along γ is bounded. The weakest assumption to keep x along γ bounded would then seem to be: $\bigcup_{0 \leq \lambda \leq 1} S_\lambda$ is bounded. This condition is a bit subtle, though,

as S_λ depends indirectly on x^0 , and x^0 is supposed to be generic. Precisely, the requirement is as follows. Let $X^0 \subset E^n$, $B^0 \subset E^m$ be open nonempty sets such that for each point $x^0 \in X^0$, there exists $b^0 \in B^0$ such that $b^0 > 0$, $G(x^0, 0) - b^0 < 0$. Then $\bigcup_{0 \leq \lambda \leq 1} S_\lambda$ must be bounded for each $x^0 \in X^0$, $b^0 \in B^0$ satisfying $b^0 > 0$, $G(x^0, 0) - b^0 < 0$.

The above discussion is summarized in the hypotheses of the following theorem. Let $a = (x^0, b^0, c^0)$, and define $\rho_a(\lambda, x, u) = \rho(x^0, b^0, c^0, \lambda, x, u)$, according to (7.9) and (7.10). As always, u^0 is uniquely defined by $K(0, x^0, u^0, b^0, c^0) = 0$.

THEOREM 7.1. *Let $f : E^n \rightarrow E$ and $g : E^n \rightarrow E^m$ be C^3 functions, let g satisfy the Arrow-Hurwicz-Uzawa constraint qualification at every local solution of (7.1), let $X^0 \subset E^n$ and $B^0 \subset \{b \in E^m \mid b > 0\}$ be open and nonempty, and for $b^0 \in B^0$ and $0 \leq \lambda \leq 1$ define*

$$S_\lambda(b^0) = \{x \in E^n \mid G(x, \lambda) - (1 - \lambda)b^0 \leq 0\}.$$

For each $x^0 \in X^0$ assume there exists $b^0 \in B^0$ such that $G(x^0, 0) - b^0 < 0$. For each $x^0 \in X^0$ and $b^0 \in B^0$ satisfying $G(x^0, 0) - b^0 < 0$, further assume that $S_\lambda(b^0)$ is nonempty for $0 \leq \lambda \leq 1$, and that $\bigcup_{0 \leq \lambda \leq 1} S_\lambda(b^0)$ is bounded. Let $\rho_a(\lambda, x, u) =$

$\rho(x^0, b^0, c^0, \lambda, x, u)$ be defined from (7.9) and (7.10). Then for almost all $x^0 \in X^0$, almost all $b^0 \in B^0$ such that $G(x^0, 0) - b^0 < 0$, and almost all $c^0 \in E^m$ with $c^0 > 0$, there exists a zero curve γ of $\rho_a(\lambda, x, u)$ emanating from $(0, x^0, u^0)$, along which the Jacobian matrix $D\rho_a(\lambda, x, u)$ has rank $n + m$. If in addition there exists $\kappa > 0$ such that for any point (λ, x, u) on γ ,

$$\|(\lambda, x, u) - (0, x^0, u^0)\| > 1 \implies \lambda \geq \kappa,$$

and for any accumulation point $(\hat{\lambda}, \hat{x})$ of (λ, x) along γ

$$\left[\nabla_x G_J(\hat{x}, \hat{\lambda}) \right] z > 0 \text{ has a solution } z,$$

where $J = \{j \mid G_j(\hat{x}, \hat{\lambda}) - (1 - \hat{\lambda})b_j^0 = 0\}$, then γ reaches a point $(1, \bar{x}, \bar{u})$, where (\bar{x}, \bar{u}) solves the Kuhn-Tucker problem (7.2)–(7.5). If rank $D\rho_a(1, \bar{x}, \bar{u}) = n + m$, then γ has finite arc length.

Proof. The homotopy map ρ defined from (7.9)–(7.10), similar to ρ and K from (5.11)–(5.12) and (6.6)–(6.7), satisfies the hypotheses of Lemma 2.2. Thus a homotopy zero curve γ exists and it only remains to show γ is bounded. If γ reaches a point $(1, \bar{x}, \bar{u})$, since $\rho_a(1, \bar{x}, \bar{u}) = 0$ is equivalent to the necessary optimality conditions (7.2)–(7.5), (\bar{x}, \bar{u}) will be a stationary point for the original optimization problem (7.1).

As before, by the nature of K and the boundedness of $\bigcup S_\lambda$, (λ, x) is bounded along γ . Suppose that γ is unbounded, and let $(\lambda^{(k)}, x^{(k)}, u^{(k)}) \rightarrow \infty$ be a sequence of points along γ . As before, there is a subsequence $(\lambda^{(k_i)}, x^{(k_i)}) \rightarrow (\hat{\lambda}, \hat{x})$, with $u^{(k_i)} \geq 0$ and $\|u^{(k_i)}\| \rightarrow \infty$. By assumption, $\hat{\lambda} \geq \kappa > 0$ and thus $\hat{\lambda} \neq 0$. The argument for the case $\hat{\lambda} = 1$ is identical to that in the proof of Theorem 5.1. Now consider only $0 < \hat{\lambda} < 1$. The argument for Case 1, where for some j ,

$$(1 - \hat{\lambda})b_j^0 - G_j(\hat{x}, \hat{\lambda}) > 0 \quad \text{and} \quad \overline{\lim}_{k_i \rightarrow \infty} u_j^{(k_i)} = \infty,$$

is identical to that for Theorem 5.1.

Case 2, where for every j ,

$$\overline{\lim}_{k_i \rightarrow \infty} u_j^{(k_i)} = \infty \implies (1 - \hat{\lambda})b_j^0 - G_j(\hat{x}, \hat{\lambda}) = 0,$$

leads to the system

$$(\nabla_x G_J(\hat{x}, \hat{\lambda}))^t v_J = 0, \quad v_J \geq 0,$$

having a solution v_J , where $J \subset \{j \mid G_j(\hat{x}, \hat{\lambda}) - (1 - \hat{\lambda})b_j^0 = 0\}$. By Gordan's Theorem of the Alternative,

$$(\nabla_x G_J(\hat{x}, \hat{\lambda}))z > 0$$

has no solution z . This contradicts the explicit hypothesis about $\nabla_x G_J$ along γ . Therefore γ is bounded, and the theorem follows. \square

COROLLARY 7.1. *If the assumption in Theorem 7.1 about $\lambda \geq \kappa > 0$ for points on γ far from $(0, x^0, u^0)$ is replaced by*

$$\text{rank } D_{(x,u)}\rho_a(\lambda, x, u) = n + m$$

along γ , then the conclusions of Theorem 7.1 hold.

Proof. The rank assumption implies that γ has no turning points, a much stronger assumption that simply $\lambda \geq \kappa > 0$ eventually, i.e., as arc length s increases, γ does not asymptotically approach the hyperplane $\lambda = 0$. \square

On many realistic engineering applications, γ does in fact have several turning points, and if the convergence theory could not accommodate turning points, it would not accurately reflect practice. Theorems 5.1, 6.1, and 7.1 have been presented as a series of generalizations with proof refinements, and so the post-mortem comments on Theorems 5.1 and 6.1 in essence apply here also. Proving that γ is bounded amounts to controlling, in some fashion, each of λ , x , u along γ .

λ is controlled by preventing $\liminf_{s \rightarrow \infty} \lambda(s) = 0$. This frequently happens when the homotopy map is poorly chosen. For instance, $\lambda(s) \rightarrow 0$ as $u(s) \rightarrow \infty$ in Fig. 1 for

the homotopy map (5.6). The theory here shows that with the right homotopy map, $0 < \liminf_{s \rightarrow \infty} \lambda(s) = \hat{\lambda} < 1$ cannot happen except in rare degenerate cases involving the active constraint gradients $\nabla_x G_J(\hat{x}, \hat{\lambda})$.

x is controlled by the property that $x \in S_\lambda(b^0)$, and by assumption $\bigcup_{0 \leq \lambda \leq 1} S_\lambda(b^0)$

is bounded. What happens if $S_\lambda(b^0) = \emptyset$ for some $0 < \lambda < 1$? For complicated problems, it is easy to unwittingly construct a family (7.8) for which some $S_\lambda(b^0)$ is empty. Consider the problem

$$\min_x F(x, \lambda) = x \quad \text{subject to} \quad G(x, \lambda) = x^2 - 1 + 2\lambda \leq 0$$

and take $x^0 = 0$, $b^0 = 1$, $c^0 = 1$. $S_{2/3}(1) = \{0\}$, $S_\lambda(1) = \emptyset$ for $\lambda > 2/3$ (i.e., there is no solution at $\lambda = 1$) so something has to fail. Figure 2 shows what happens to γ for the map (7.9).




FIG. 2. Example of homotopy zero curve when some S_λ is empty.

u is controlled by the property (in the convex case) or the assumption (in the nonconvex case) that $(\nabla_x G_J(\hat{x}, \hat{\lambda}))z > 0$ has a solution z , where J is related to active constraints at an accumulation point $(\hat{\lambda}, \hat{x})$ of $((\lambda, x)$ along) γ . This condition can be interpreted as a “constraint qualification for homotopy maps.” Since its failure to hold represents a degenerate situation, it can be achieved (in principle) by generically perturbing the map $G(x, \lambda)$.

As mentioned earlier, the map (7.9) closely resembles those used in practice, and thus Theorem 7.1 reflects practice. Generally, for each λ , the problem (7.8) is physically meaningful with S_λ being nonempty and bounded. The constraint qualification involving $\nabla_x G_J(x, \lambda)$ holds generically, and thus is not normally a concern. $\liminf_{s \rightarrow \infty} \lambda(s) > 0$ must be assumed, and this is the fly in the ointment. This condition is achieved by some sort of global monotonicity property (which often *does* hold for practical problems, related to energy considerations), the rank condition of Corollary 7.1 (extremely hard to verify for a complicated problem), or by the clever construction of (7.8). There is no silver bullet!

8. Acknowledgment. The author is indebted to Gene Allgower, Kurt Georg, Alexander Morgan, and many mentors, colleagues, and engineers who have shared their problems and insights. Thanks are due also to a referee whose insightful observations helped strengthen the theorems and proofs.

REFERENCES

- [1] J. C. ALEXANDER, *The topological theory of an imbedding method*, in Continuation Methods (H. G. Wacker, ed.), Academic Press, New York, 1978, pp. 37–68.
- [2] J. C. ALEXANDER, T.-Y. LI, AND J. A. YORKE, *Piecewise smooth homotopies*, in Homotopy Methods and Global Convergence, B. C. Eaves, F. J. Gould, H.-O. Peitgen, and M. J. Todd, eds., Plenum, New York, 1983, pp. 1–14.
- [3] E. L. ALLGOWER AND K. GEORG, *Numerical Continuation Methods*, Springer-Verlag, Berlin, (1990).
- [4] S. C. BILLUPS, A. L. SPEIGHT, AND L. T. WATSON, *Nonmonotone path following methods for nonsmooth equations and complementarity problems*, in Applications and Algorithms of Complementarity, M. C. Ferris, O. L. Mangasarian, and J.-S. Pang, (Eds.), Kluwer, Norwell, MA, to appear.
- [5] S. N. CHOW, J. MALLET-PARET, AND J. A. YORKE, *Finding zeros of maps: Homotopy methods that are constructive with probability one*, Math. Comp., 32 (1978), pp. 887–899.
- [6] B. C. EAVES, *Homotopies for computation of fixed points*, Math. Programming, 3 (1972), pp. 1–22.
- [7] B. C. EAVES AND R. SAIGAL, *Homotopies for computation of fixed points on unbounded regions*, Math. Programming, 3 (1972), pp. 225–237.
- [8] Y. GE, L. T. WATSON, E. G. COLLINS, JR., AND D. S. BERNSTEIN, *Probability-one homotopy algorithms for full and reduced order H^2/H^∞ controller synthesis*, Optimal Control Appl. Methods, 17 (1996), pp. 187–208.
- [9] Z. LIN, Y. LI, AND B. YU, *A combined homotopy interior point method for general nonlinear programming problems*, Appl. Math. Comp., 80 (1996), pp. 209–224.
- [10] Z. LIN, B. YU, AND G. FENG, *A combined homotopy interior point method for convex nonlinear programming*, Appl. Math. Comp., 84 (1997), pp. 193–211.
- [11] O. L. MANGASARIAN, *Nonlinear Programming*, McGraw-Hill, New York, 1969.
- [12] ———, *Equivalence of the complementarity problem to a system of nonlinear equations*, SIAM J. Appl. Math., 31 (1976), pp. 89–92.
- [13] A. P. MORGAN, *Solving Polynomial Systems Using Continuation for Scientific and Engineering Problems*, Prentice-Hall, Englewood Cliffs, NJ, (1987).
- [14] A. P. MORGAN AND A. J. SOMMESE, *A homotopy for solving general polynomial systems that respects m -homogeneous structures*, Appl. Math. Comput., 24 (1987), pp. 101–113.
- [15] ———, *Computing all solutions to polynomial systems using homotopy continuation*, Appl. Math. Comput., 24 (1987), pp. 115–138.
- [16] A. B. POORE AND Q. AL-HASSAN, *The expanded Lagrangian system for constrained optimization problems*, SIAM J. Control Optim., 26 (1988), pp. 417–427.
- [17] R. T. ROCKAFELLAR, *Convex Analysis*, Princeton Univ. Press, Princeton, NJ, 1970.
- [18] H. SELLAMI, *A homotopy continuation method for solving normal equations*, Math. Prog., 82 (1998), pp. 317–337.
- [19] H. SELLAMI AND S. M. ROBINSON, *Implementation of a continuation method for normal maps*, Math. Prog., 76 (1997), pp. 563–578.
- [20] L. T. WATSON, *A globally convergent algorithm for computing fixed points of C^2 maps*, Appl. Math. Comput., 5 (1979), pp. 297–311.
- [21] ———, *An algorithm that is globally convergent with probability one for a class of nonlinear two-point boundary value problems*, SIAM J. Numer. Anal., 16 (1979), pp. 394–401.
- [22] ———, *Solving the nonlinear complementarity problem by a homotopy method*, SIAM J. Control Optim., 17 (1979), pp. 36–46.
- [23] ———, *Solving finite difference approximations to nonlinear two-point boundary value problems by a homotopy method*, SIAM J. Sci. Stat. Comput., 1 (1980), pp. 467–480.
- [24] ———, *Computational experience with the Chow-Yorke algorithm*, Math. Programming, 19 (1980), pp. 92–101.

- [25] ———, *Numerical linear algebra aspects of globally convergent homotopy methods*, SIAM Rev., 28 (1986), pp. 529–545.
- [26] ———, *Globally convergent homotopy algorithms for nonlinear systems of equations*, Nonlinear Dynamics, 1 (1990) 143–191.
- [27] ———, *A survey of probability-one homotopy methods for engineering optimization*, Arabian J. Sci. Engrg., 16 (1991) 297–323.
- [28] L. T. WATSON, S. C. BILLUPS, AND A. P. MORGAN, *HOMPACK: A suite of codes for globally convergent homotopy algorithms*, ACM Trans. Math. Software, 13 (1987), pp. 281–310.
- [29] L. T. WATSON, J. P. BIXLER, AND A. B. POORE, *Continuous homotopies for the linear complementarity problem*, SIAM J. Matrix Anal. Appl., 10 (1989), pp. 259–277.
- [30] L. T. WATSON AND R. T. HAFTKA, *Modern homotopy methods in optimization*, Comput. Methods Appl. Mech. Engrg., 74 (1989) 289–305.
- [31] L. T. WATSON AND M. R. SCOTT, *Solving spline collocation approximations to nonlinear two-point boundary value problems by a homotopy method*, Appl. Math. Comput., 24 (1987), pp. 333–357.
- [32] L. T. WATSON AND L. R. SCOTT, *Solving Galerkin approximations to nonlinear two-point boundary value problems by a globally convergent homotopy method*, SIAM J. Sci. Stat. Comput., 8 (1987), pp. 768–789.
- [33] L. T. WATSON, M. SOSONKINA, R. C. MELVILLE, A. P. MORGAN, AND H. F. WALKER, *Algorithm 777: HOMPACK90: A suite of FORTRAN 90 codes for globally convergent homotopy algorithms*, ACM Trans. Math. Software, 23 (1997), pp. 514–549.