

SANDIA REPORT

SAND2003-8550

Unlimited Release

Printed October 2003

Stationarity Results for Generating Set Search for Linearly Constrained Optimization

Tamara G. Kolda, Robert Michael Lewis, and Virginia Torczon

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865)576-8401
Facsimile: (865)576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.doe.gov/bridge>

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800)553-6847
Facsimile: (703)605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>



SAND2003-8550
Unlimited Release
Printed October 2003

Stationarity Results for Generating Set Search for Linearly Constrained Optimization

Tamara G. Kolda*
Computational Sciences and Mathematics Research Department
Sandia National Laboratories
Livermore, CA 94551-9217

Robert Michael Lewis†
Department of Mathematics
College of William & Mary
P.O. Box 8795
Williamsburg, VA 23187-8795

Virginia J. Torczon‡
Department of Computer Science
College of William & Mary
P.O. Box 8795
Williamsburg, VA 23187-8795

ABSTRACT

We derive new stationarity results for derivative-free, generating set search methods for linearly constrained optimization. We show that a particular measure of stationarity is of the same order as the step length at an identifiable subset of the iterations. Thus, even in the absence of explicit knowledge of the derivatives of the objective function, we still have information about stationarity. These results help both unify the convergence analysis of several classes of direct search algorithms and clarify the fundamental geometrical ideas that underlie them. In addition, these results validate a practical stopping criterion for such algorithms.

Keywords: constrained optimization, direct search, generating set search, global convergence analysis, nonlinear programming, derivative-free methods, pattern search

*Corresponding author. Email: tgkolda@sandia.gov. This work was supported by the Mathematical, Information, and Computational Sciences Program of the U.S. Department of Energy, under contract DE-AC04-94AL85000 with Sandia Corporation.

†Email: buckaroo@math.wm.edu. This research was supported by the Computer Science Research Institute at Sandia National Laboratories and by the National Science Foundation under Grant DMS-0215444.

‡Email: va@cs.wm.edu. This research was funded by the Computer Science Research Institute at Sandia National Laboratories.

This page intentionally left blank.

1. Introduction. In this paper, we consider the convergence properties of generating set search methods for linearly constrained optimization:

$$(1.1) \quad \begin{array}{ll} \text{minimize} & f(x) \\ \text{subject to} & Ax \leq b. \end{array}$$

Here A is an $m \times n$ matrix, and Ω denotes the feasible set: $\Omega = \{ x \mid Ax \leq b \}$.

In this paper, we restrict attention to *feasible iterates* approaches. Thus, we require the initial iterate x_0 to be feasible and arrange for all subsequent iterates x_k to satisfy $x_k \in \Omega$. Derivative-free algorithms of the class we term *generating set search* (GSS) [5] were introduced for bound constraints in [6, 9] and for general linear constraints in [7, 10] ([10] also allows for nonlinear constraints whose derivatives are known). It was shown that under appropriate hypotheses, the limit points of the sequence of iterates produced by these algorithms are Karush–Kuhn–Tucker (KKT) points of (1.1).

In this paper, we give a different approach to the convergence analysis of GSS methods for linearly constrained optimization. The approach is based on new stationarity results for these algorithms. Specifically, we show that at an identifiable subsequence of iterations, there is a linear relationship between a particular measure of stationarity and a parameter Δ_k that controls the length of the step admissible at iteration k of a GSS algorithm. This relationship is useful since Δ_k is known to the user, while any direct measure of stationarity would involve $\nabla f(x)$, which is presumed to be unavailable or unreliable in the derivative-free context. Such a relationship exists for the unconstrained case [3] and the bound constrained case [6, 8]; however, the relationship for the linearly constrained case in [7], using a measure of stationarity different from the one used here, is less satisfying. The results presented here rectify this shortcoming.

Our results also help both unify the convergence analysis of several classes of direct search algorithms and clarify the fundamental geometrical ideas that underlie them. In addition, they validate a practical stopping criterion for GSS algorithms. Finally, these stationarity results are needed to extend GSS methods [4] to the augmented Lagrangian approach of Conn, Gould, Sartenaer, and Toint [1].

In §2 we introduce some notation and terminology. We outline the class of algorithms under consideration in §3. The stationarity results and their significance are discussed in §4; preliminary statements of these results appeared in [5]. The geometrical facts that underlie the stationarity results follow in Appendix A. Variations of the stationarity results in §4 are developed in Appendix B.

2. Notation and geometrical quantities of interest. For GSS methods, the set of search directions \mathcal{D}_k for linearly constrained optimization must reflect the geometry of the boundary of the feasible region near x_k . This requires that we be able to identify the active, or nearly active, constraints so that we can specify an appropriate set of generators to serve as search directions.

Let a_i^T be the i th row of the constraint matrix A in (1.1). Let

$$\mathcal{C}_i = \{ y \mid a_i^T y = b_i \}$$

denote the set where the i th constraint is binding. Given $x \in \Omega$ and $\varepsilon \geq 0$, define the index set $I(x, \varepsilon)$ to be the indices i of those constraints for which x is within distance ε of \mathcal{C}_i :

$$(2.1) \quad i \in I(x, \varepsilon) \text{ if and only if } \text{dist}(x, \mathcal{C}_i) \leq \varepsilon.$$

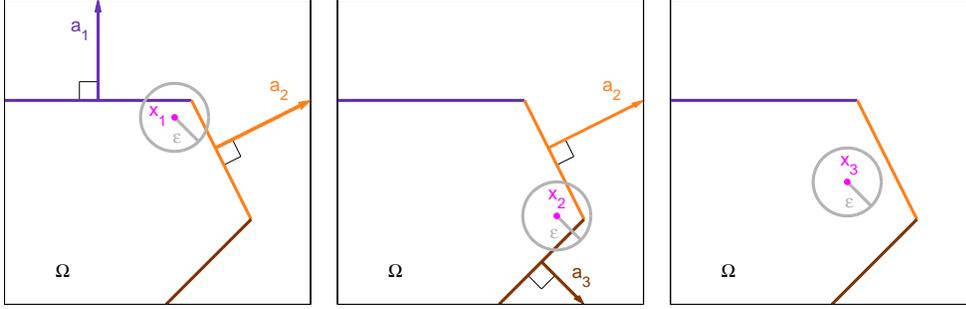


FIG. 2.1. The outward-pointing normals a_i for the index set $I(x_1, \varepsilon) = \{1, 2\}$ and a_i for the index set $I(x_2, \varepsilon) = \{2, 3\}$. Since the distance from x_3 to $\partial\Omega$ is greater than ε , $I(x_3, \varepsilon) = \emptyset$.

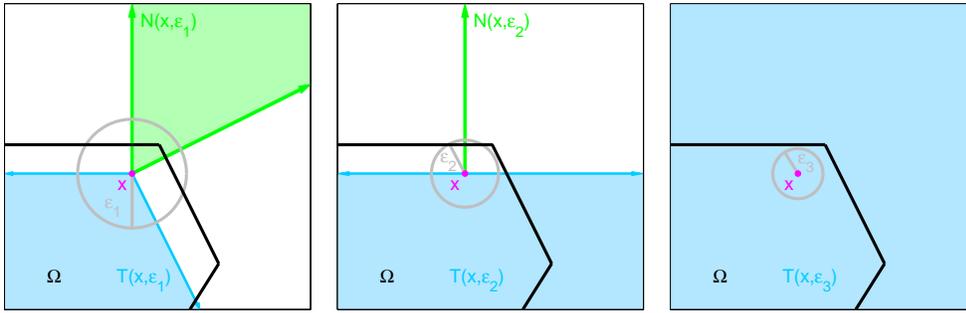


FIG. 2.2. The cones $N(x, \varepsilon)$ and $T(x, \varepsilon)$ for the values ε_1 , ε_2 , and ε_3 . Note that for this example, as ε varies from ε_1 to 0, there are only the three distinct pairs of cones illustrated ($N(x, \varepsilon_3) = \{0\}$).

The vectors a_i for $i \in I(x, \varepsilon)$ are the outward-pointing normals to the faces of the boundary of Ω within distance ε of x . Examples are shown in Figure 2.1 for three choices of $x \in \Omega$.

Given $x \in \Omega$, we denote by $N(x, \varepsilon)$ the cone generated by $\{0\}$ and the vectors a_i for $i \in I(x, \varepsilon)$. Its polar cone is denoted by $T(x, \varepsilon)$:

$$T(x, \varepsilon) = \{ v \mid w^T v \leq 0 \text{ for all } w \in N(x, \varepsilon) \}.$$

If $N(x, \varepsilon) = \{0\}$, which is the case when $I(x, \varepsilon) = \emptyset$, then $T(x, \varepsilon) = \mathbb{R}^n$ (in other words, if the boundary is more than distance ε away, and one looks only within distance ε of x , then the problem looks unconstrained). Several examples are illustrated in Figure 2.2. Observe that $N(x, 0)$ is the tangent cone of Ω at x , while $T(x, 0)$ is the tangent cone of Ω at x ,

The significance of $N(x, \varepsilon)$ is that for suitable choices of ε , its polar $T(x, \varepsilon)$ approximates the polyhedron Ω near x , where “near” is in terms of ε . (More precisely, $x + T(x, \varepsilon)$ approximates the feasible region near x .) The polar cone $T(x, \varepsilon)$ is important because if $T(x, \varepsilon) \neq \{0\}$, then one can proceed from x along all directions in $T(x, \varepsilon)$ for a distance of at least ε , and still remain inside the feasible region (see Proposition A.6). Again, Figure 2.2 illustrates this point. The same is not true for directions in the tangent cone, since the tangent cone does not reflect the proximity of the boundary for points close to, but not on, the boundary.

Finally, given a vector v , let $v_{N(x, \varepsilon)}$ and $v_{T(x, \varepsilon)}$ denote the projections (in the

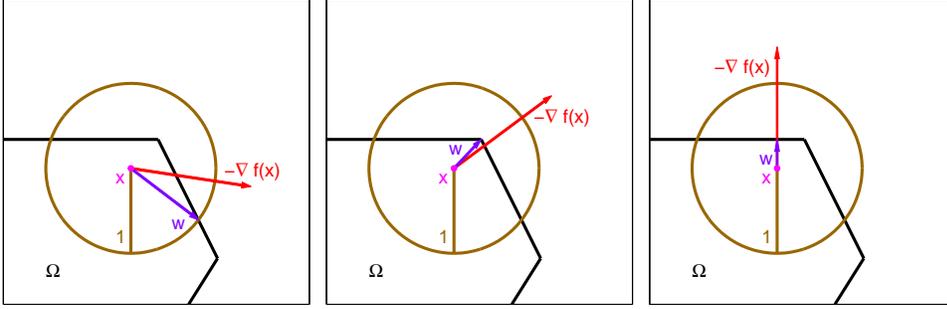


FIG. 2.3. How the w in (2.2) varies with $-\nabla f(x)$ when $x - \nabla f(x) \notin \Omega$.

Euclidean norm) of v onto $N(x, \varepsilon)$ and its polar $T(x, \varepsilon)$, respectively. We frequently abbreviate this notation as v_N and v_T . The polar decomposition [12, 13, 14] says that any vector v can be written as the sum of its projection onto a cone and its polar and the projections are orthogonal. We make use of the fact that $v = v_N + v_T$, where $v_N^T v_T = 0$.

We use the following quantity to measure progress toward a KKT point of (1.1). For $x \in \Omega$, let

$$(2.2) \quad \chi(x) = \max_{\substack{x+w \in \Omega \\ \|w\| \leq 1}} -\nabla f(x)^T w.$$

The function χ has the following properties [2]:

1. $\chi(x)$ is continuous.
2. $\chi(x) \geq 0$.
3. $\chi(x) = 0$ if and only if x is a KKT point for (1.1).

Showing that $\chi(x_k) \rightarrow 0$ as $k \rightarrow \infty$ then constitutes a global first-order stationarity result. The w 's that define $\chi(x)$ in (2.2) are illustrated in Figure 2.3 for three choices of $-\nabla f(x)$.

3. GSS algorithms for linearly constrained problems. The general form of the algorithm for linearly-constrained generating set search is given in Figure 3.1. There are four places where the description in Figure 3.1 is intentionally vague; all are flagged by the word ‘‘admissible.’’ In the subsections that follow, we make clear what is meant by ‘‘admissible.’’ We begin with some terminology.

At iteration k , x_k is always the best feasible point discovered thus far; i.e., $f(x_k) \leq f(x_\ell)$ for all $\ell \leq k$. The scalar Δ_k denotes the *step-length control parameter*; i.e., it controls the maximum length of the admissible step along any search direction $d \in \mathcal{D}_k$. The scalar function $\rho(\cdot)$ is called the *forcing function*. The requirements on $\rho(\cdot)$ are given in §3.2.

An iteration is called *successful* if the decrease condition in Step 2 is satisfied; i.e., there exists $d_k \in \mathcal{D}_k$ and an admissible $\tilde{\Delta}_k \in [0, \Delta_k]$ such that

$$(3.1) \quad f(x_k + \tilde{\Delta}_k d_k) < f(x_k) - \rho(\Delta_k).$$

The set of indices of all successful iterations is denoted by \mathcal{S} : $k \in \mathcal{S}$ if and only if iteration k is successful. Note that d_k , $\tilde{\Delta}_k$, and ϕ_k are defined only if $k \in \mathcal{S}$.

If the iteration is not successful, then we deem it *unsuccessful*. It must be the case then that for each $d \in \mathcal{G}_k$, either $x_k + \Delta_k d \notin \Omega$ or

$$f(x_k + \Delta_k d) \geq f(x_k) - \rho(\Delta_k).$$

ALGORITHM 3.1 (Linearly constrained generating set search)

INITIALIZATION.

Let $x_0 \in \Omega$ be the initial guess.

Let $\Delta_{\text{tol}} > 0$ be the tolerance used to test for convergence.

Let $\Delta_0 > \Delta_{\text{tol}}$ be the initial value of the step-length control parameter.

Let $\varepsilon_\star > 0$ be the maximum distance used to identify nearby constraints.

ALGORITHM. For each iteration $k = 0, 1, 2, \dots$

STEP 1. Choose an admissible set of search directions $\mathcal{D}_k = \mathcal{G}_k \cup \mathcal{H}_k$.

STEP 2. If there exists $d_k \in \mathcal{D}_k$ and a corresponding admissible $\tilde{\Delta}_k \in [0, \Delta_k]$ such that $x_k + \tilde{\Delta}_k d_k \in \Omega$ and

$$f(x_k + \tilde{\Delta}_k d_k) < f(x_k) - \rho(\Delta_k),$$

then:

- Set $x_{k+1} = x_k + \tilde{\Delta}_k d_k$.
- Set $\Delta_{k+1} = \phi_k \Delta_k$ for an admissible choice of $\phi_k \geq 1$.

STEP 3. Otherwise, for every $d \in \mathcal{G}_k$, either $x_k + \Delta_k d \notin \Omega$ or

$$f(x_k + \Delta_k d) \geq f(x_k) - \rho(\Delta_k).$$

In this case:

- Set $x_{k+1} = x_k$ (no change).
- Set $\Delta_{k+1} = \theta_k \Delta_k$ for an admissible choice of $\theta_k \in (0, 1)$.

If $\Delta_{k+1} < \Delta_{\text{tol}}$, then terminate.

FIG. 3.1. *Linearly constrained GSS*

The set of indices of all unsuccessful iterations is denoted by \mathcal{U} : $k \in \mathcal{U}$ if and only if iteration k is unsuccessful. Note that $\mathcal{S} \cap \mathcal{U} = \emptyset$. Also, θ_k is defined only if $k \in \mathcal{U}$.

The scalar constant Δ_{tol} denotes the *convergence tolerance* on the step-length control parameter. Once $\Delta_{k+1} < \Delta_{\text{tol}}$, the algorithm terminates. However, Δ_k is only reduced for $k \in \mathcal{U}$, so the algorithm can only terminate at the end of an unsuccessful iteration. The stationarity results in §4 make clear that this stopping criterion is appropriate.

3.1. Choosing an admissible set of search directions. The set of search directions at iteration k is $\mathcal{D}_k = \mathcal{G}_k \cup \mathcal{H}_k$. For now, we assume the following two conditions are satisfied by the \mathcal{G}_k used at each iteration of the algorithm in Figure 3.1. Conditions on \mathcal{H}_k are discussed in §3.2. A variant of Condition 1 is discussed in Appendix B.

CONDITION 1. There exists a constant $c_{\min} > 0$, independent of k , such that for all k the following holds. For every $\varepsilon \in [0, \varepsilon_\star]$, there exists $\mathcal{G} \subseteq \mathcal{G}_k$ such that \mathcal{G} generates $T(x_k, \varepsilon)$ and, furthermore, $c_{A.1}(\mathcal{G}) > c_{\min}$, where $c_{A.1}(\cdot)$ is the quantity from Proposition A.1.

Condition 1 says that to be admissible, \mathcal{G}_k must contain generators for all $T(x_k, \varepsilon)$ as ε varies from 0 to ε_\star . Observe, however, that as ε varies from 0 to ε_\star there is only a finite number of distinct cones $N(x_k, \varepsilon)$ (and correspondingly $T(x_k, \varepsilon)$) since there is only a finite number of faces of Ω . For instance, in Figure 2.2, if $\varepsilon_\star = \varepsilon_1$, then for $0 \leq \varepsilon \leq \varepsilon_\star$ there are only three distinct cones $N(x_k, \varepsilon)$ since there are only two faces of Ω within distance ε_\star of x_k . The notion of requiring some of the search directions to generate $T(x_k, \varepsilon)$ for every $\varepsilon \in [0, \varepsilon_\star]$ originally appeared in a slightly restricted form in [11], and later in [7, 10].

Condition 1 also imposes a uniformity condition which ensures that if $-\nabla f(x)$ is not normal to Ω at x , then at least one search direction in \mathcal{G}_k lies safely within 90° of $-\nabla f(x)$. To understand the necessity of this condition, suppose the problem were unconstrained. Then $T(x_k, \varepsilon) = \mathbb{R}^n$ for all ε . If we apply Proposition A.1 with $K = \mathbb{R}^n$ and $v = -\nabla f(x_k)$, then $v_K = -\nabla f(x_k)$ and (A.1) reduces to

$$c_{A.1}(\mathcal{G}) \|\nabla f(x)\| \|d\| \leq -\nabla f(x)^T d.$$

This we recognize as a condition which says that the angle made by at least one of the search directions in \mathcal{G}_k with the direction of steepest descent remains bounded away from 90° . This is a familiar condition in line search methods. See Section 3.4.1 of [5] for a further discussion of this point.

Condition 2 imposes bounds on the norm of all the search directions in \mathcal{G}_k .

CONDITION 2. There exist $\beta_{\min} > 0$ and $\beta_{\max} > 0$, independent of k , such that for all k the following holds.

$$\beta_{\min} \leq \|d\| \leq \beta_{\max} \quad \text{for all } d \in \mathcal{G}_k.$$

Since the lengths of the steps that can be taken depend both on Δ_k (the step-length control parameter) and on $\|d\|$, for all $d \in \mathcal{G}_k$ we require bounds on the lengths of the search directions in \mathcal{G}_k so that the length of steps of the form $\tilde{\Delta}_k d$, for $\tilde{\Delta}_k \in [0, \Delta_k]$ and $d \in \mathcal{G}_k$, can be monitored using Δ_k .

3.2. Globalization strategies. In Steps 2–3 of the linearly constrained GSS algorithm given in Figure 3.1, the choice of an admissible step, as well as admissible parameters ϕ_k and θ_k , depends on the type of globalization strategy that is in effect.

The globalization of GSS is discussed in depth in [5]. The basic idea is to enforce conditions which ensure that at least

$$\liminf_{k \rightarrow \infty} \Delta_k = 0.$$

We consider two possibilities here. We first specify the requirements on the forcing function $\rho(\cdot)$.

CONDITION 3. (Requirements on the forcing function)

1. The function $\rho(\cdot)$ is a nonnegative continuous function on $[0, +\infty)$.
2. The function $\rho(\cdot)$ is $o(t)$ as $t \downarrow 0$; i.e., $\lim_{t \downarrow 0} \rho(t) / t = 0$.
3. The function $\rho(\cdot)$ is nondecreasing; i.e., $\rho(t_1) \leq \rho(t_2)$ if $t_1 \leq t_2$.

Note that $\rho(\Delta) \equiv 0$ satisfies these requirements, as does $\rho(\Delta) = \alpha\Delta^p$ where α is some positive scalar and $p > 1$.

3.2.1. Globalization via a sufficient decrease condition. One possibility is to use a sufficient decrease condition to ensure global convergence. This globalization strategy requires the following of the forcing function $\rho(\cdot)$ and the choice of θ_k .

CONDITION 4. (Sufficient decrease)

1. Condition 3 is satisfied.
2. The forcing function $\rho(\cdot)$ is such that $\rho(t) > 0$ for $t > 0$.
3. A constant $\theta_{\max} < 1$ exists such that $\theta_k < \theta_{\max}$ for all k .

Full details are discussed in Section 3.7.1 of [5]. The requirements of Condition 4 are easily satisfied by choosing, say, $\rho(\Delta) = 10^{-2}\Delta^2$ and $\theta_k = \frac{1}{2}$. Note that this condition does not make any assumptions about the generating sets. The only other assumption necessary to yield a globalization result is that f be bounded below, as follows.

THEOREM 3.1 (Theorem 3.4 of [5]). *Assume that the linearly constrained GSS method in Figure 3.1 satisfies Condition 4. Furthermore, assume the function f is bounded below. Then $\liminf_{k \rightarrow \infty} \Delta_k = 0$.*

3.2.2. Globalization via a rational lattice. Another possibility requires all iterates to lie on a rational lattice. In this case, restrictions are required both on the types of steps and on the types of search directions that are allowed.

CONDITION 5. (Rational lattice)

1. Condition 3 is satisfied.
2. The set $\mathbf{G} = \{d^{(1)}, \dots, d^{(p)}\}$ is a finite set of search directions, and every vector $d \in \mathbf{G}$ is of the form $d = Bc$ where $B \in \mathbb{R}^{n \times n}$ is a nonsingular matrix and $c \in \mathbb{Q}^n$.
3. All generators are chosen from \mathbf{G} ; i.e., $\mathcal{G}_k \subseteq \mathbf{G}$ for all k .
4. All extra directions are integer combinations of the elements of \mathbf{G} ; i.e., $\mathcal{H}_k \subset \{\sum_{i=0}^p \xi^{(i)} d^{(i)} \mid \xi^{(i)} \in \{0, 1, 2, \dots\}\}$.
5. The scalar Λ is a fixed positive integer.
6. For all $k \in \mathcal{S}$, ϕ_k is of the form $\phi_k = \Lambda^{\ell_k}$ where $\ell_k \in \{0, 1, 2, \dots\}$.
7. For all $k \in \mathcal{U}$, θ_k is of the form $\theta_k = \Lambda^{m_k}$ where $m_k \in \{-1, -2, \dots\}$.
8. All steps $\tilde{\Delta}_k \in [0, \Delta_k]$ satisfy either $\tilde{\Delta}_k = 0$ or $\tilde{\Delta}_k = \Lambda^m \Delta_k > \Delta_{\text{tol}}$, where $m \in \mathbb{Z}$, $m \leq 0$.

While the list of requirements in Condition 5 looks onerous, in fact most can be satisfied in a straightforward fashion. A full discussion of the reasons for these requirements can be found in Section 3.7.2 of [5]; here we limit ourselves to a few observations.

First, the choice $\rho(\Delta) \equiv 0$ is standard for the rational lattice globalization strategy and means that Condition 3 is satisfied automatically.

Note also that the set \mathbf{G} is a conceptual construct that describes the set of all *possible* search directions, rather than necessarily being a set of directions that are actually formed and used in the search. For instance, \mathbf{G} may be viewed as containing all the generators for all possible cones $T(x, \varepsilon)$, for all $x \in \Omega$ and $\varepsilon \in [0, \varepsilon_*]$. Furthermore, since the number of faces in the polyhedron Ω is finite, the set \mathbf{G} may be chosen to be finite. In this way the third requirement in Condition 5, necessary for globalization, is satisfied. In fact, it is not necessary to construct the set \mathbf{G} of all potential search directions. At any iteration, a smaller set of search directions particular to x_k is all that is needed.

A standard assumption [7] is that the linear constraints are rational, i.e., $A \in \mathbb{Q}^{m \times n}$. Then, if the generators are constructed with care (see Section 8 in [7]), every generator will be rational. Choosing B to be a diagonal matrix with nonzero entries along the diagonal then ensures that the vectors d in \mathbf{G} are of the form $d = Bc$ where $B \in \mathbb{R}^{n \times n}$ is nonsingular and $c \in \mathbb{Q}^n$, while remaining generators for all possible cones $T(x, \varepsilon)$. Note that Condition 2 is satisfied automatically because \mathbf{G} is finite.

Finally, consider the requirements on the scaling of the step. The usual choice of Λ is 2. In the unconstrained case, ϕ_k typically is chosen from the set $\{1, 2\}$ so that $\ell_k \in \{0, 1\}$ for all k while θ_k usually is chosen to be $\frac{1}{2}$ so that $m_k = -1$ for all k . In the case of linear constraints, it is convenient to make use of the added flexibility in choosing m_k so that it is possible to take feasible steps that move the iterate close to or onto a face of Ω , as we discuss next. That is the import of the final requirement in Condition 5.

As before, note that Condition 5 does not make any assumption about the generating sets. The only other assumption necessary to yield a globalization result is that the set $\mathcal{F} = \Omega \cap \mathcal{L}_f(x_0)$ is bounded, where $\mathcal{L}_f(x_0) = \{x \mid f(x) \leq f(x_0)\}$ denotes the level set of f at x_0 , as follows.

THEOREM 3.2 (Theorem 3.8 of [5]). *Assume that the linearly constrained GSS method in Figure 3.1 satisfies Condition 5, and also that the set $\mathcal{F} = \Omega \cap \mathcal{L}_f(x_0)$ is bounded. Then $\liminf_{k \rightarrow \infty} \Delta_k = 0$.*

3.2.3. Choosing an admissible step. In this section, it is helpful to enumerate the search directions in \mathcal{D}_k . At iteration k , suppose

$$\mathcal{D}_k = \left\{ d_k^{(1)}, d_k^{(2)}, \dots, d_k^{(p_k)} \right\}.$$

In the unconstrained case, the trial points at iteration k are specified by

$$\left\{ x_k + \Delta_k d_k^{(i)} \mid i = 1, \dots, p_k \right\}.$$

In the linearly constrained case, however, we may choose to take shorter steps, depending on whether or not the full step (Δ_k) yields a feasible trial point.

Since we are restricting our attention to feasible iterates methods for solving (1.1), we need to choose a scaling factor, which we denote $\tilde{\Delta}_k^{(i)}$, to handle those situations where $x_k + \Delta_k d_k^{(i)} \notin \Omega$. Thus, the set of trial points in the constrained case is defined by

$$\left\{ x_k + \tilde{\Delta}_k^{(i)} d_k^{(i)} \mid i = 1, \dots, p_k \right\}.$$

The choice of $\tilde{\Delta}_k^{(i)}$ must be done in a way that ensures feasible iterates and yet ensures global convergence. If the i th trial point yields success, we use the notation $\hat{\Delta}_k = \tilde{\Delta}_k^{(i)}$ and $d_k = d_k^{(i)}$; i.e., we drop the superscript.

To ensure convergence to KKT points, we want to take a full step if it yields a feasible trial point. This leads to the following condition.

CONDITION 6. If $x_k + \Delta_k d_k^{(i)} \in \Omega$, then $\tilde{\Delta}_k^{(i)} = \Delta_k$.

On the other hand, if $x_k + \Delta_k d_k^{(i)} \notin \Omega$, then there are at least two ways to choose $\tilde{\Delta}_k^{(i)}$ in Step 2. The first approach is straightforward. For a given $d_k^{(i)} \in \mathcal{D}_k$,

$$(3.2) \quad \tilde{\Delta}_k^{(i)} = \begin{cases} \Delta_k & \text{if } x_k + \Delta_k d_k^{(i)} \in \Omega \\ 0 & \text{otherwise.} \end{cases}$$

This corresponds to exact penalization (see the discussion in Section 8.1 of [5]) since the effect of (3.2) is to reject any step $\Delta_k d_k^{(i)}$ (by setting $\tilde{\Delta}_k^{(i)} = 0$) that would take the search outside the feasible region Ω .

The second possibility is to take the longest possible step along $d_k^{(i)}$ that keeps the trial point feasible while satisfying the requirements of a given globalization strategy. If a sufficient decrease condition is being employed (Condition 4), then an admissible choice of $\tilde{\Delta}_k^{(i)}$ is the solution to

$$(3.3) \quad \begin{aligned} & \text{maximize} && \Delta \\ & \text{subject to} && 0 \leq \Delta \leq \Delta_k, \\ & && x_k + \Delta d_k^{(i)} \in \Omega. \end{aligned}$$

If the rational lattice strategy is being used (Condition 5), then an admissible choice of $\tilde{\Delta}_k^{(i)}$ is the solution to

$$\begin{aligned} & \text{maximize} && \Delta \\ & \text{subject to} && \Delta = 0 \text{ or } \Delta = \Lambda^m \Delta_k, \text{ with } \Delta_{\text{tol}} < \Lambda^m \Delta_k, \ m \in \{-1, -2, \dots\}, \\ & && x_k + \Delta d_k^{(i)} \in \Omega. \end{aligned}$$

In other words, the goal is to take the longest possible step that keeps the trial point feasible while remaining on the rational lattice that underlies the search. The condition $\Delta_{\text{tol}} < \Lambda^m \Delta_k, m \in \{-1, -2, \dots\}$, serves to ensure that m is bounded below, independent of k .

4. Stationarity results. At *unsuccessful* iterations of the GSS method outlined in Figure 3.1, we can bound the measure of stationarity $\chi(x_k)$ in terms of Δ_k . To do so, we make the following assumption. Let $\mathcal{F} = \{x \in \Omega \mid f(x) \leq f(x_0)\}$.

ASSUMPTION 4.1. *The set \mathcal{F} is bounded.*

ASSUMPTION 4.2. *The gradient of f is Lipschitz continuous with constant M on \mathcal{F} .*

If both Assumption 4.1 and Assumption 4.2 hold, then there exists $\gamma > 0$ such that for all $x \in \mathcal{F}$,

$$(4.1) \quad \|\nabla f(x)\| < \gamma.$$

The Lipschitz assumption is purely for convenience; a version of our results can be proved assuming only continuity of $\nabla f(x)$. We then have the following results regarding GSS methods for linearly constrained problems.

THEOREM 4.3. *Assume that the linearly constrained GSS method in Figure 3.1 satisfies Conditions 1, 2, and 6. Let $\varepsilon_\star > 0$ be given. Suppose also that Assumption 4.2 holds. If k is an unsuccessful iteration and $\Delta_k \beta_{\max} \leq \varepsilon_\star$, then*

$$(4.2) \quad \left\| [-\nabla f(x_k)]_{T(x_k, \varepsilon)} \right\| \leq \frac{1}{c_{\min}} \left(M \Delta_k \beta_{\max} + \frac{\rho(\Delta_k)}{\Delta_k \beta_{\min}} \right)$$

for any ε satisfying $\Delta_k \beta_{\max} \leq \varepsilon \leq \varepsilon_\star$.

Proof. Consider any ε for which $\Delta_k \beta_{\max} \leq \varepsilon \leq \varepsilon_\star$. Clearly, we need only consider the case when $[-\nabla f(x_k)]_{T(x_k, \varepsilon)} \neq 0$.

Condition 1 guarantees a set $\mathcal{G} \subseteq \mathcal{G}_k$ that generates $T(x_k, \varepsilon)$. Then we can apply Proposition A.1 with $K = T(x_k, \varepsilon)$ and $v = -\nabla f(x_k)$ to conclude that there exists some $\hat{d} \in \mathcal{G}$ such that

$$(4.3) \quad c_{A.1}(\mathcal{G}) \left\| [-\nabla f(x_k)]_{T(x_k, \varepsilon)} \right\| \|\hat{d}\| \leq -\nabla f(x_k)^T \hat{d}.$$

Condition 6 and the fact that iteration k is unsuccessful tell us that

$$f(x_k + \Delta_k d) \geq f(x_k) - \rho(\Delta_k) \quad \text{for all } d \in \mathcal{G}_k \text{ for which } x_k + \Delta_k d \in \Omega.$$

Condition 2 ensures that for all $d \in \mathcal{G}$, $\|\Delta_k d\| \leq \Delta_k \beta_{\max}$ and, by assumption, $\Delta_k \beta_{\max} \leq \varepsilon$, so we have $\|\Delta_k d\| \leq \varepsilon$ for all $d \in \mathcal{G}$. Proposition A.6 assures us that $x_k + \Delta_k d \in \Omega$ for all $d \in \mathcal{G}$. Thus,

$$(4.4) \quad f(x_k + \Delta_k d) - f(x_k) + \rho(\Delta_k) \geq 0 \quad \text{for all } d \in \mathcal{G}.$$

Meanwhile, since the gradient of f is assumed to be continuous (Assumption 4.2), we can apply the mean value theorem to obtain, for some $\alpha_k \in (0, 1)$,

$$f(x_k + \Delta_k d) - f(x_k) = \Delta_k \nabla f(x_k + \alpha_k \Delta_k d)^T d \quad \text{for all } d \in \mathcal{G}.$$

Putting this together with (4.4), we obtain

$$0 \leq \Delta_k \nabla f(x_k + \alpha_k \Delta_k d)^T d + \rho(\Delta_k) \quad \text{for all } d \in \mathcal{G}.$$

Dividing through by Δ_k and subtracting $\nabla f(x_k)^T d$ from both sides yields

$$-\nabla f(x_k)^T d \leq (\nabla f(x_k + \alpha_k \Delta_k d) - \nabla f(x_k))^T d + \rho(\Delta_k)/\Delta_k \quad \text{for all } d \in \mathcal{G}.$$

Since $\nabla f(x)$ is Lipschitz continuous (Assumption 4.2) and $0 < \alpha_k < 1$, we obtain

$$(4.5) \quad -\nabla f(x_k)^T d \leq M \Delta_k \|d\|^2 + \rho(\Delta_k)/\Delta_k \quad \text{for all } d \in \mathcal{G}.$$

Since (4.5) holds for all $d \in \mathcal{G}$, (4.3) tells us that for some $\hat{d} \in \mathcal{G}$

$$c_{A.1}(\mathcal{G}) \left\| [-\nabla f(x_k)]_{T(x_k, \varepsilon)} \right\| \leq M \Delta_k \|\hat{d}\| + \frac{\rho(\Delta_k)}{\Delta_k \|\hat{d}\|}.$$

Using the bounds on $\|\hat{d}\|$ in Condition 2,

$$\left\| [-\nabla f(x_k)]_{T(x_k, \varepsilon)} \right\| \leq \frac{1}{c_{A.1}(\mathcal{G})} \left(M \Delta_k \beta_{\max} + \frac{\rho(\Delta_k)}{\Delta_k \beta_{\min}} \right).$$

The theorem then follows from the fact that $c_{A.1}(\mathcal{G}) > c_{\min}$ (Condition 1). \square

Theorem 4.3 makes it possible to use linearly constrained GSS methods in connection with the augmented Lagrangian framework presented in [1]. The approach in [1] proceeds by successive approximate minimization of the augmented Lagrangian. The stopping criterion in the subproblems involves the norm of the projection onto $T(x_k, \omega_k)$ of the negative gradient of the augmented Lagrangian, for a parameter $\omega_k \downarrow 0$. In the derivative-free setting the gradient is unavailable; however, Theorem 4.3 enables us to use Δ_k as an alternative measure of stationarity in the subproblems. Details will appear in [4].

THEOREM 4.4. *Assume that the linearly constrained GSS method in Figure 3.1 satisfies Conditions 1, 2, and 6. Let $\varepsilon_\star > 0$ be given. Suppose also that Assumptions 4.1–4.2 hold. Then there exists $c_{4.4} > 0$, independent of k , but depending on A , c_{\min} , the γ from (4.1), and M , such that if k is an unsuccessful iteration and $\Delta_k \beta_{\max} \leq \varepsilon_\star$, then*

$$\chi(x_k) \leq c_{4.4} \left(M \Delta_k \beta_{\max} + \frac{\rho(\Delta_k)}{\Delta_k \beta_{\min}} \right).$$

Proof. Clearly, we need only consider the case when $\chi(x_k) \neq 0$.

Case I. First, suppose

$$\Delta_k \beta_{\max} \leq r_{A.4}(\gamma, A) \chi(x_k),$$

where $r_{A.4}(\gamma, A)$ is from Proposition A.4 and β_{\max} is the upper bound from Condition 2 on the norms of the search directions. Let

$$\varepsilon = \min\{r_{A.4}(\gamma, A)\chi(x_k), \varepsilon_\star\}.$$

Then $\Delta_k \beta_{\max} \leq \varepsilon \leq \varepsilon_\star$, so we can apply Theorem 4.3 to obtain

$$\|[-\nabla f(x_k)]_{T(x_k, \varepsilon)}\| \leq \frac{1}{c_{\min}} \left(M \Delta_k \beta_{\max} + \frac{\rho(\Delta_k)}{\Delta_k \beta_{\min}} \right).$$

Moreover, $\varepsilon \leq r_{A.4}(\gamma, A)\chi(x_k)$, so we can apply Proposition A.4 with $v = -\nabla f(x_k)$ to obtain

$$\|[-\nabla f(x_k)]_{T(x_k, \varepsilon)}\| \geq \frac{1}{2} \chi(x_k).$$

The two preceding relations yield

$$(4.6) \quad \chi(x_k) \leq \frac{2}{c_{\min}} \left(M \Delta_k \beta_{\max} + \frac{\rho(\Delta_k)}{\Delta_k \beta_{\min}} \right).$$

Case II. The second case to consider is

$$\Delta_k \beta_{\max} > r_{A.4}(\gamma, A) \chi(x_k).$$

This can be rewritten as

$$(4.7) \quad \chi(x_k) < \frac{1}{r_{A.4}(\gamma, A) M} M \Delta_k \beta_{\max}.$$

From (4.6) and (4.7), choosing

$$c_{4.4} = \min \left\{ \frac{2}{c_{\min}}, \frac{1}{r_{A.4}(\gamma, A) M} \right\}$$

yields the desired result. \square

Theorem 4.4 suggests that at unsuccessful iterations, the continuous measure of stationarity $\chi(\cdot)$ tends to decrease as Δ_k is decreased. Since Δ_k is reduced only at unsuccessful iterations, where the result of Theorem 4.4 holds, it is reasonable to terminate the algorithm when Δ_k is reduced to below some tolerance.

As an immediate corollary of Theorem 4.4, we obtain a first-order convergence result for the GSS algorithm in Figure 3.1.

THEOREM 4.5. *Assume that the linearly constrained GSS method in Figure 3.1 satisfies Conditions 1, 2, 6, and either Condition 4 or 5. Let $\varepsilon_\star > 0$ be given. Suppose also that Assumptions 4.1–4.2 hold. Then*

$$\liminf_{k \rightarrow 0} \chi(x_k) = 0.$$

In Appendix B we prove results similar to Theorems 4.3–4.4 when the set of search directions at iteration k is allowed to be smaller than that considered here.

5. Conclusions. In [7], stationarity was measured using the quantity

$$q(x) = P_\Omega(x - \nabla f(x)),$$

where P_Ω is the projection onto Ω . In this case, $q(x)$ is continuous and $\|q(x)\| = 0$ if and only if x is a KKT point of (1.1). In [7], the authors showed that the bound

$$\|q(x)\| = O\left(\Delta_k^{\frac{1}{2}}\right)$$

held at unsuccessful iterations. They conjectured that this could be sharpened to $O(\Delta_k)$, but encountered obstacles to proving such a result.

On the other hand, Theorem 4.4 shows that at unsuccessful iterations, $\chi(x)$ is $O(\Delta_k)$. The results we have proved have at least two benefits. The first is that Theorem 4.4 provides a justification for the stopping criterion prescribed for linearly constrained GSS in Figure 3.1. The situation is much the same as that discussed for unconstrained and bound constrained problems in [3, 8]. The step-length control parameter Δ_k , which appears in the definition of GSS methods, provides a reliable asymptotic measure of first-order stationarity. A second consequence, which results from Theorem 4.3, is that it now will be possible to extend the results in [1] to direct search methods [4]. The situation is much the same as that discussed in [8]. The stopping criterion proposed by Conn, Gould, Sartenaer, and Toint for the solution of the subproblem requires explicit knowledge of derivatives. Such information is presumed absent in direct search methods, but we can replace this with a stopping criterion based on the size of Δ_k in a way that preserves the convergence properties of the algorithm in [1].

Theorem 4.3 and Theorem 4.4 also bring out some of the elements common to the approaches described in [6, 7] and [9, 10]. In particular, we see how the convergence analysis may be viewed as comprising two parts: showing that at least a subsequence of $\{\Delta_k\}$ converges to 0 and showing that a stationarity bound like that in Theorem 4.4 holds at the unsuccessful iterations. For the first part, the algorithms in [6, 7] rely on simple decrease in the objective and the fact that the iterates lie on a successively refined lattice, while in [9, 10] the derivative-free sufficient decrease condition is used. On the other hand, for both classes of algorithms the situation at unsuccessful iterations is the same, as described in Theorem 4.3 and Theorem 4.4. These bounds

are consequences of the choice of search directions. The stationarity results in §4 do not hold for all the algorithmic possibilities discussed in [10], so in Appendix B we investigate what can be said using a variant of Condition 1 that allows \mathcal{G}_k to contain generators for a single cone $T(x_k, \varepsilon_k)$, where ε_k is updated at each iteration.

Since the choice of search directions for a linearly constrained GSS method depends on Ω , we have attempted to clarify some of the fundamental geometrical ideas, pertaining to the fact that Ω is a polyhedron, that are necessary for the stationarity results in §4 and Appendix B to hold. These geometrical results, which do not depend on any assumptions regarding f and its derivatives, are developed in the appendix that follows.

Appendix A. Geometric results on cones and polyhedra. Here we present geometrical results having to do with our use of $\chi(\cdot)$ as a measure of stationarity. Since these are largely technical in nature, we have relegated them to an appendix.

The first result says that if a vector v is not in the polar K° of a finitely generated cone K and $v_K \neq 0$, then v must be within 90° of at least one of the generators of K . The proof follows that of Corollary 10.4 in [7].

PROPOSITION A.1. *Let K be a convex cone in \mathbb{R}^n generated by the finite set \mathcal{G} . Then there exists $c_{A.1}(\mathcal{G}) > 0$, depending only on \mathcal{G} , for which the following holds. For any v for which $v_K \neq 0$,*

$$(A.1) \quad \max_{d \in \mathcal{G}} \frac{v^T d}{\|d\|} \geq c_{A.1}(\mathcal{G}) \|v_K\|.$$

The next proposition says that if one can move from x to $x+v$ and remain feasible, then v cannot be too outward-pointing with respect to the constraints near x .

PROPOSITION A.2. *There exists $c_{A.2}(A) > 0$, depending only on A , for which the following holds. If $x \in \Omega$ and $x+v \in \Omega$, then for any $\varepsilon \geq 0$, $\|v_{N(x,\varepsilon)}\| \leq c_{A.2}(A) \varepsilon$.*

Proof. Let $N = N(x, \varepsilon)$. The result is immediate if $v_N = 0$, so we need only consider the case when $v_N \neq 0$. Recall that N is generated by the outward-pointing normals to the binding constraints within distance ε of x ; thus, the set $\mathcal{G} = \{a_i \mid i \in I(x, \varepsilon)\}$ generates N . A simple calculation shows that the distance from x to $\{y \mid a_i^T y = b_i\}$ is $(b_i - a_i^T x) / \|a_i\|$, so it follows that

$$\frac{b_i - a_i^T x}{\|a_i\|} \leq \varepsilon \quad \text{for all } i \in I(x, \varepsilon).$$

Meanwhile, since $x+v \in \Omega$, we have

$$a_i^T x + a_i^T v \leq b_i \quad \text{for all } i.$$

The preceding two relations then lead to

$$(A.2) \quad a_i^T v \leq b_i - a_i^T x \leq \varepsilon \|a_i\| \quad \text{for all } i \in I(x, \varepsilon).$$

Since N is generated by $\mathcal{G} = \{a_i \mid i \in I(x, \varepsilon)\}$, and $v_N \neq 0$, we may apply Proposition A.1 and (A.2) to obtain

$$c_{A.1}(\mathcal{G}) \|v_N\| \leq \max_{i \in I(x, \varepsilon)} \frac{v^T a_i}{\|a_i\|} \leq \max_{i \in I(x, \varepsilon)} \frac{\varepsilon \|a_i\|}{\|a_i\|} = \varepsilon.$$

Any \mathcal{G} used here consists of combinations of columns of A , and the number of such combinations is finite. Thus there is a lower bound $c_* > 0$ for $c_{A.1}(\mathcal{G})$ that depends only on A . The result follows with $c_{A.2}(A) = 1/c_*$. \square

For $x \in \Omega$ and $v \in \mathbb{R}^n$, define

$$(A.3) \quad \hat{\chi}(x; v) = \max_{\substack{x+w \in \Omega \\ \|w\| \leq 1}} w^T v.$$

Note from (2.2) that $\chi(x) = \hat{\chi}(x; -\nabla f(x))$. We use v in (A.3) to emphasize that the following results are purely geometric facts about cones and polyhedra.

The following proposition relates $\hat{\chi}(x; v)$ to the cones $T(x, \varepsilon)$ and $N(x, \varepsilon)$.

PROPOSITION A.3. *There exists $c_{A.3}(A) > 0$, depending only on A , such that if $x \in \Omega$, then for all $\varepsilon \geq 0$,*

$$\hat{\chi}(x; v) \leq \|v_{T(x, \varepsilon)}\| + c_{A.3}(A) \|v_{N(x, \varepsilon)}\| \varepsilon.$$

Proof. Let $N = N(x, \varepsilon)$ and $T = T(x, \varepsilon)$. Writing v in terms of its polar decomposition, $v = v_N + v_T$, we obtain

$$\hat{\chi}(x; v) = \max_{\substack{x+w \in \Omega \\ \|w\| \leq 1}} w^T v \leq \max_{\substack{x+w \in \Omega \\ \|w\| \leq 1}} w^T v_T + \max_{\substack{x+w \in \Omega \\ \|w\| \leq 1}} w^T v_N.$$

For the first term on the right-hand side we have

$$\max_{\substack{x+w \in \Omega \\ \|w\| \leq 1}} w^T v_T \leq \|v_T\|.$$

Meanwhile, for any w we have

$$w^T v_N = (w_T + w_N)^T v_N \leq w_N^T v_N$$

since $w_T^T v_N \leq 0$. Thus,

$$\max_{\substack{x+w \in \Omega \\ \|w\| \leq 1}} w^T v_N \leq \max_{\substack{x+w \in \Omega \\ \|w\| \leq 1}} \|w_N\| \|v_N\|.$$

However, since $x + w \in \Omega$, Proposition A.2 tells us that

$$\|w_N\| \leq c_{A.2}(A) \varepsilon.$$

Therefore,

$$\hat{\chi}(x; v) \leq \|v_T\| + c_{A.2}(A) \|v_N\| \varepsilon.$$

Setting $c_{A.3}(A) = c_{A.2}(A)$ completes the proof. \square

The following corollary says that if $\hat{\chi}(x; v) \neq 0$ and ε is sufficiently small (relative to $\hat{\chi}(x; v)$), then $\hat{\chi}(x; v)$ cannot be very large unless $v_{T(x, \varepsilon)}$, the tangential part of v , is also at least of the same order of magnitude.

PROPOSITION A.4. *Let $\gamma > 0$ be given, and let $v \in \mathbb{R}^n$ satisfy $\|v\| \leq \gamma$. Suppose $x \in \Omega$ is such that $\hat{\chi}(x; v) > 0$. Then there exists $r_{A.4}(\gamma, A) > 0$, depending only on γ and A , such that if $\varepsilon \leq r_{A.4}(\gamma, A) \hat{\chi}(x; v)$, then*

$$(A.4) \quad \|v_{T(x, \varepsilon)}\| \geq \frac{1}{2} \hat{\chi}(x; v).$$

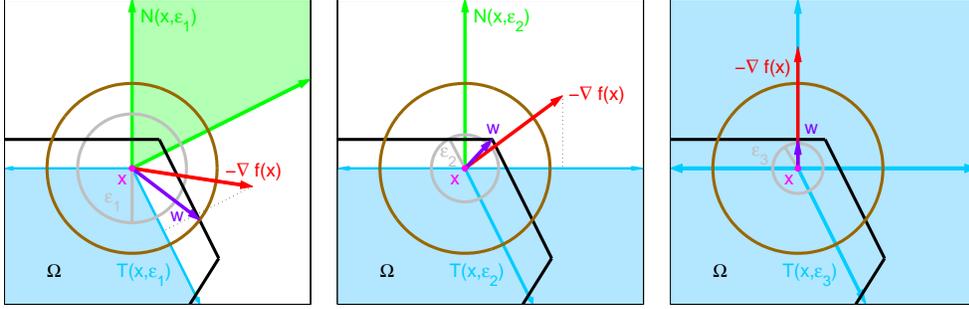


FIG. A.1. The cones $N(x, \varepsilon)$ and $T(x, \varepsilon)$ for the values ε_1 , ε_2 , and ε_3 , and three examples (from Figure 2.3) showing the projection of the negative gradient onto a generator of each of the cones $T(x, \varepsilon_1)$, $T(x, \varepsilon_2)$, and $T(x, \varepsilon_3)$.

Proof. Let $N = N(x, \varepsilon)$ and $T = T(x, \varepsilon)$. By Proposition A.3, for any $\varepsilon \geq 0$,

$$\hat{\chi}(x; v) \leq \|v_T\| + c_{A.3}(A) \|v_N\| \varepsilon.$$

Since $\|v_N\| \leq \|v\|$ (because the projection onto convex sets is contractive), we have

$$\hat{\chi}(x; v) \leq \|v_T\| + c_{A.3}(A)\gamma\varepsilon.$$

Set

$$r_{A.4}(\gamma, A) = \frac{1}{2c_{A.3}(A)\gamma};$$

then, if $\varepsilon \leq r_{A.4}(\gamma, A)\hat{\chi}(x; v)$, we have

$$\frac{1}{2} \hat{\chi}(x; v) \leq \|v_T\|.$$

□

The following proposition ensures that at least one of the directions in \mathcal{D}_k , a valid set of directions for linearly constrained GSS methods (see Figure 3.1), is a descent direction. Illustrations can be seen in Figure A.1 for the three examples from Figure 2.3 as ε is reduced to satisfy Condition 1.

PROPOSITION A.5. *Let $\gamma > 0$ be given, and let $v \in \mathbb{R}^n$ satisfy $\|v\| \leq \gamma$. Suppose $x \in \Omega$ is such that $\hat{\chi}(x; v) > 0$. Given $\varepsilon \geq 0$, let \mathcal{G} generate $T(x, \varepsilon)$. Then there exists $c_{A.5}(\mathcal{G}) > 0$, depending only on \mathcal{G} , such that if $\varepsilon \leq r_{A.4}(\gamma, A)\hat{\chi}(x; v)$, then there exists $d \in \mathcal{G}$ such that*

$$(A.5) \quad c_{A.5}(\mathcal{G}) \hat{\chi}(x; v) \|d\| \leq v^T d.$$

Proof. Let $\varepsilon \leq r_{A.4}(\gamma, A)\hat{\chi}(x; v)$. Proposition A.4 tells us that

$$\|v_{T(x, \varepsilon)}\| \geq \frac{1}{2} \hat{\chi}(x; v).$$

Since, by hypothesis, $\hat{\chi}(x; v) > 0$, it follows that $v_{T(x, \varepsilon)} \neq 0$. Proposition A.1 then says that there exists $d \in \mathcal{G}$ such that

$$v^T d \geq c_{A.1}(\mathcal{G}) \|v_{T(x, \varepsilon)}\| \|d\|.$$

Combining the two previous inequalities yields

$$v^T d \geq \frac{1}{2} c_{A.1}(\mathcal{G}) \hat{\chi}(x; v) \|d\|.$$

Let $c_{A.5}(\mathcal{G}) = \frac{1}{2} c_{A.1}(\mathcal{G})$. \square

Another point needed for the stationarity analysis is that it is possible to take a suitably long step along the descent direction in \mathcal{G} promised by Proposition A.5. The following result ensures that it is possible to take steps of at least length ε along directions in $T(x, \varepsilon)$ and remain feasible, as in the examples in Figure A.1.

PROPOSITION A.6. *If $x \in \Omega$, and $v \in T(x, \varepsilon)$ satisfies $\|v\| \leq \varepsilon$, then $x + v \in \Omega$.*

Proof. Suppose not; i.e., $v \in T(x, \varepsilon)$ and $\|v\| \leq \varepsilon$, but $x + v \notin \Omega$. Then there exists i such that $a_i^T(x + v) > b_i$. Using the fact that $x \in \Omega$, so $a_i^T x \leq b_i$, we have

$$(A.6) \quad a_i^T v > b_i - a_i^T x \geq 0.$$

Define

$$t = \frac{b_i - a_i^T x}{a_i^T v}.$$

Note that $t < 1$ by (A.6). Let $y = x + tv$. Then $a_i^T y = b_i$ and $\|x - y\| = \|tv\| < \varepsilon$. Thus, $i \in I(x, \varepsilon)$ and $a_i \in N(x, \varepsilon)$. Since, by hypothesis, $v \in T(x, \varepsilon)$, we must have $a_i^T v \leq 0$. However, this contradicts (A.6). \square

Appendix B. Results for a more limited set of search directions. Results similar to Theorems 4.3–4.4 can be proved when, rather than requiring \mathcal{G}_k to contain generators for $T(x_k, \varepsilon)$ for all $\varepsilon \in [0, \varepsilon_*]$ as in Condition 1, \mathcal{G}_k need only contain generators for a single cone $T(x_k, \varepsilon_k)$, where ε_k is updated at each iteration. Algorithm 1 in [10] is based on such a set of search directions. This leads to the following relaxation of Condition 1.

CONDITION 7. There exists a constant $c_{\min} > 0$, independent of k , such that for all k the following holds. There exists $\mathcal{G} \subseteq \mathcal{G}_k$ such that \mathcal{G} generates $T(x_k, \varepsilon_k)$ and, furthermore, $c_{A.1}(\mathcal{G}) > c_{\min}$, where $c_{A.1}(\cdot)$ is the quantity from Proposition A.1.

Note, however, that the uniformity condition is unchanged.

To ensure global convergence, at each iteration the parameter ε_k must be updated in such a way that $\lim_{k \in \mathcal{U}} \varepsilon_k = 0$ (or at least $\liminf_{k \in \mathcal{U}} \varepsilon_k = 0$). In Algorithm 1 in [10], this is done by setting $\varepsilon_{k+1} = \varepsilon_k$ for $k \in \mathcal{S}$ and $\varepsilon_{k+1} = \eta \varepsilon_k$ for $k \in \mathcal{U}$ where η is some fixed constant with $\eta \in (0, 1)$. Unlike in [10], we require the following relation between ε_k and Δ_k which restricts ε_k from becoming too small relative to Δ_k .

CONDITION 8. There exists a constant $c_\varepsilon > 0$, independent of k , such that $c_\varepsilon \Delta_k \leq \varepsilon_k$ for all k .

This condition can be ensured if ε_k is reduced by the same factor θ_k as Δ_k at unsuccessful iterations; i.e., $\varepsilon_{k+1} = \theta_k \varepsilon_k$ for all $k \in \mathcal{U}$.

The resulting algorithm is similar to Algorithm 1 in [10], though our approach differs in several ways. The most material difference is that we impose a more stringent step acceptance criterion, requiring the improvement in f to be $\rho(\Delta_k)$, rather than

$\rho(\tilde{\Delta}_k)$. Another difference, as noted previously, is that we require a relationship between ε_k and Δ_k . We then have the following analog of Theorem 4.3.

THEOREM B.1. *Assume that the linearly constrained GSS method in Figure 3.1 satisfies Conditions 2, 6, 7, and 8 and that $\tilde{\Delta}_k^{(i)}$ is chosen according to (3.3). Suppose also that Assumption 4.2 holds. If k is an unsuccessful iteration, then*

$$(B.1) \quad \left\| [-\nabla f(x_k)]_{T(x_k, \varepsilon_k)} \right\| \leq \frac{1}{c_{\min}} \left(M\Delta_k\beta_{\max} + \frac{\rho(\Delta_k)}{\mu\Delta_k\beta_{\min}} \right),$$

where $\mu > 0$ is independent of k and depends only on c_ε and β_{\max} .

Proof. The proof closely resembles that of Theorem 4.3. By Condition 7, there is a subset $\mathcal{G} \subseteq \mathcal{G}_k$ that generates $T(x_k, \varepsilon_k)$. Applying Proposition A.1 with $K = T(x_k, \varepsilon_k)$ and $v = -\nabla f(x_k)$, we are assured that there exists some $d_k^{(i)} \in \mathcal{G}$ for which

$$(B.2) \quad c_{A.1}(\mathcal{G}) \left\| [-\nabla f(x_k)]_{T(x_k, \varepsilon_k)} \right\| \left\| d_k^{(i)} \right\| \leq -\nabla f(x_k)^T d_k^{(i)}.$$

Now, the algorithm always tries a step $\tilde{\Delta}_k^{(i)} d_k^{(i)}$ along each generator $d_k^{(i)}$. Since $\tilde{\Delta}_k^{(i)}$ is chosen as the solution to (3.3), either $\tilde{\Delta}_k^{(i)} = \Delta_k$ or $\tilde{\Delta}_k^{(i)}$ is as long as possible while still having $x + \tilde{\Delta}_k^{(i)} d_k^{(i)} \in \Omega$. In the latter case, Proposition A.6 says we can move a distance of at least ε_k along the generators of $T(x_k, \varepsilon_k)$ and remain feasible. Therefore, if $\tilde{\Delta}_k^{(i)} < \Delta_k$, we still have at least $\tilde{\Delta}_k^{(i)} \left\| d_k^{(i)} \right\| \geq \varepsilon_k$. From Condition 8 it then follows that

$$\tilde{\Delta}_k^{(i)} \geq \frac{\varepsilon_k}{\left\| d_k^{(i)} \right\|} \geq \frac{c_\varepsilon \Delta_k}{\beta_{\max}}.$$

Thus, in either case we have

$$(B.3) \quad \tilde{\Delta}_k^{(i)} \geq \min \left(1, \frac{c_\varepsilon}{\beta_{\max}} \right) \Delta_k.$$

Let

$$\mu = \min \left(1, \frac{c_\varepsilon}{\beta_{\max}} \right).$$

If iteration k is unsuccessful, then

$$(B.4) \quad f(x_k + \tilde{\Delta}_k^{(i)} d_k^{(i)}) - f(x_k) + \rho(\Delta_k) \geq 0 \quad \text{for all } d_k^{(i)} \in \mathcal{G}_k.$$

Then, as in the proof of Theorem 4.3, using the mean-value theorem and the Lipschitz continuity of $\nabla f(x)$, we find that

$$-\nabla f(x_k)^T d_k^{(i)} \leq M\tilde{\Delta}_k^{(i)} \left\| d_k^{(i)} \right\|^2 + \rho(\Delta_k)/\tilde{\Delta}_k^{(i)} \quad \text{for all } d_k^{(i)} \in \mathcal{G}_k.$$

From (B.3) we then have

$$(B.5) \quad -\nabla f(x_k)^T d_k^{(i)} \leq M\Delta_k \left\| d_k^{(i)} \right\|^2 + \rho(\Delta_k)/(\mu\Delta_k) \quad \text{for all } d_k^{(i)} \in \mathcal{G}_k.$$

Since (B.5) holds for all $d_k^{(i)} \in \mathcal{G}_k$, (B.2) tells us that

$$c_{A.1}(\mathcal{G}) \left\| [-\nabla f(x_k)]_{T(x_k, \varepsilon_k)} \right\| \leq M\Delta_k \left\| \hat{d}_k^{(i)} \right\| + \frac{\rho(\Delta_k)}{\mu\Delta_k \left\| \hat{d}_k^{(i)} \right\|}.$$

Because Condition 2 is satisfied, and

$$\|[-\nabla f(x_k)]_{T(x_k, \varepsilon_k)}\| \leq \frac{1}{c_{A.1}(\mathcal{G})} \left(M \Delta_k \beta_{\max} + \frac{\rho(\Delta_k)}{\mu \Delta_k \beta_{\min}} \right)$$

Appealing to Condition 7 and choosing $c_{4.3} = 1/c_{\min}$ yields the result. \square

Note that this theorem differs from Theorem 4.3 because it has the factor μ in the denominator of the second term of the right hand side. Since $\mu \leq 1$, this means that the bound in (B.1) may not be as tight as the bound (4.2) obtained when \mathcal{G}_k contains generators for multiple cones (as in Condition 1).

We also have the following analog of Theorem 4.4.

THEOREM B.2. *Assume that the linearly constrained GSS method in Figure 3.1 satisfies Conditions 2, 6, 7, and 8 and that $\tilde{\Delta}_k^{(i)}$ is chosen according to (3.3). Suppose also that Assumptions 4.1–4.2 hold. If k is an unsuccessful iteration, then*

$$\chi(x_k) \leq \frac{1}{c_{\min}} \left(M \Delta_k \beta_{\max} + \frac{\rho(\Delta_k)}{\mu \Delta_k \beta_{\min}} \right) + c_{A.3}(A) \gamma \varepsilon_k,$$

where γ is from (4.1) and μ is as in Theorem B.1.

Proof. From Proposition A.3 we have

$$\chi(x) \leq \|[-\nabla f(x_k)]_{T(x, \varepsilon_k)}\| + c_{A.3}(A) \|[-\nabla f(x_k)]_{N(x, \varepsilon_k)}\| \varepsilon_k.$$

The fact that $\|[-\nabla f(x_k)]_{N(x, \varepsilon_k)}\| \leq \gamma$ and the bound on $\|[-\nabla f(x_k)]_{T(x, \varepsilon_k)}\|$ from Theorem B.1 then yield the theorem. \square

Finally, we have a first-order convergence result analogous to Theorem 4.5.

THEOREM B.3. *Assume that the linearly constrained GSS method in Figure 3.1 satisfies Conditions 2, 6, 7, and 8. In addition, assume that either Condition 4 or 5 holds. Further, assume $\liminf_{k \in \mathcal{U}} \varepsilon_k = 0$. If Assumptions 4.1–4.2 also hold, then*

$$\liminf_{k \rightarrow 0} \chi(x_k) = 0.$$

REFERENCES

- [1] A. R. CONN, N. GOULD, A. SARTENAER, AND P. L. TOINT, *Convergence properties of an augmented Lagrangian algorithm for optimization with a combination of general equality and linear constraints*, SIAM Journal on Optimization, 6 (1996), pp. 674–703.
- [2] A. R. CONN, N. I. M. GOULD, AND P. L. TOINT, *Trust-Region Methods*, vol. 1 of MPS/SIAM Series on Optimization, SIAM, Philadelphia, 2000.
- [3] E. D. DOLAN, R. M. LEWIS, AND V. J. TORCZON, *On the local convergence properties of pattern search*, SIAM Journal on Optimization, (to appear).
- [4] T. G. KOLDA, R. M. LEWIS, AND V. TORCZON, *Convergence properties of an augmented Lagrangian direct search algorithm for optimization with a combination of general equality and linear constraints*. in preparation, 2003.
- [5] ———, *Optimization by direct search: New perspectives on some classical and modern methods*, SIAM Review, 45 (2003), pp. 385–482.
- [6] R. M. LEWIS AND V. TORCZON, *Pattern search algorithms for bound constrained minimization*, SIAM Journal on Optimization, 9 (1999), pp. 1082–1099.
- [7] ———, *Pattern search methods for linearly constrained minimization*, SIAM Journal on Optimization, 10 (2000), pp. 917–941.
- [8] ———, *A globally convergent augmented Lagrangian pattern search algorithm for optimization with general constraints and simple bounds*, SIAM Journal on Optimization, 12 (2002), pp. 1075–1089.

- [9] S. LUCIDI AND M. SCIANDRONE, *On the global convergence of derivative-free methods for unconstrained optimization*, SIAM Journal on Optimization, 13 (2002), pp. 97–116.
- [10] S. LUCIDI, M. SCIANDRONE, AND P. TSENG, *Objective-derivative-free methods for constrained optimization*, Mathematical Programming, 92 (2002), pp. 37–59.
- [11] J. H. MAY, *Linearly Constrained Nonlinear Programming: A Solution Method That Does Not Require Analytic Derivatives*, PhD thesis, Yale University, New Haven, Connecticut, December 1974.
- [12] J. J. MOREAU, *Décomposition orthogonale d'un espace hilbertien selon deux cônes mutuellement polaires*, Comptes Rendus de l'Académie des Sciences de Paris, 255 (1962), pp. 238–240.
- [13] R. T. ROCKAFELLAR, *Convex Analysis*, Princeton University Press, Princeton, NJ, 1970.
- [14] E. H. ZARANTONELLO, *Projections on convex sets in Hilbert space and spectral theory*, in Contributions to Nonlinear Functional Analysis, E. H. Zarantonello, ed., Academic Press, London and New York, 1971, pp. 237–424.