Global linear convergence
of an augmented Lagrangian algorithm
for solving convex quadratic optimization problems

Frédéric Delbos[†] and J. Charles Gilbert[‡]

February 25, 2004 (revised version)

Abstract. We consider an augmented Lagrangian algorithm for minimizing a convex quadratic function subject to linear inequality constraints. Linear optimization is an important particular instance of this problem. We show that, provided the augmentation parameter is large enough, the constraint value converges *globally* linearly to zero. This property is viewed as a consequence of the proximal interpretation of the algorithm and of the global radial Lipschitz continuity of the reciprocal of the dual function subdifferential. This Lipschitz property is itself obtained by means of a lemma of general interest, which compares the distances from a point in the positive orthant to an affine space, on the one hand, and to the polyhedron given by the intersection of this affine space and the positive orthant, on the other hand. No strict complementarity assumption is needed. The result is illustrated by numerical experiments and algorithmic implications, including complexity issues, are discussed.

Key words. Augmented Lagrangian, convex quadratic optimization, distance to a polyhedron, error bound, global linear convergence, iterative complexity, linear constraints, proximal algorithm.

AMS subject classification. 49M29, 65K05, 90C05, 90C06, 90C20.

1   Introduction

Convex quadratic programs (QP) arise in their own right and as subproblems in some numerical algorithms for solving optimization problems. On the one hand, since no strictly convex assumption is made, the important family of linear optimization problems, with a zero quadratic term in their objective, enters this framework. On the other hand, the SQP algorithm decomposes a regularized constrained least-squares problem into a sequence of strictly convex QP's (see [10, 5] for an example in reflection tomography, which partly motivates this study; see [19, 2] for recent books describing the SQP algorithm). Finding efficient algorithms for solving this basic multi-faceted problem in all possible situations is an objective that has been pursued for decades (see for example the already 20 year old survey on quadratic programming in [20] and the monographs on interior point methods in [15, 7, 18, 30, 14, 29, 31, 2]).

[†]Institut Français du Pétrole, 1 & 4 avenue de Bois-Préau, 92852 Rueil-Malmaison, France; e-mail: `Frederic.Delbos@ifp.fr`.

[‡]INRIA Rocquencourt, BP 105, 78153 Le Chesnay Cedex, France; e-mail: `Jean-Charles.Gilbert@inria.fr`.

The convex quadratic problem we consider is written

$$\begin{cases} \min_x \frac{1}{2} x^\top Q x + q^\top x \\ l \le Cx \le u, \end{cases} \tag{1.1}$$

where Q is an $n \times n$ positive semi-definite symmetric matrix, $q \in \mathbb{R}^n$, C is $m \times n$, and the m-dimensional vectors l and u satisfy $l < u$ (i.e., $l_i < u_i$ for all indices $i = 1, \ldots, m$) and may have infinite components. With lower and upper bounds, problem (1.1) is close to what is actually implemented in numerical codes (see [6] for an example). We have not included equality constraints, like $Ax = b$, in (1.1) to make the presentation simple, but such constraints can be expressed like in (1.1) by adding two inequalities $Ax \le b$ and $-Ax \le -b$. Therefore, our analysis covers problems with equality constraints. Note that, since Q may be zero, (1.1) also modelizes linear optimization.

The method that we further explore in this paper fits into the class of dual approaches, since it is essentially the augmented Lagrangian (AL) method of Hestenes [11] and Powell [22] that is applied to the convex QP (1.1). This algorithm can be implemented in such a way that it does not require any matrix factorization. It is therefore appropriate when the problem is so large that such a factorization is impracticable or too much time consuming. This is a motivation for using the AL algorithm when the optimization problem deals with systems governed by partial differential equations [8]. In that case, however, a good preconditioner for the unavoidable conjugate gradient iterations must be available.

The version of the AL algorithm we analyze is defined on an equivalent form of (1.1) obtained by introducing an auxiliary variable $y \in \mathbb{R}^m$ [10]:

$$\begin{cases} \min_x \frac{1}{2} x^\top Q x + q^\top x \\ y = Cx \\ l \le y \le u. \end{cases} \tag{1.2}$$

The algorithm generates a sequence $\{\lambda_k\} \subset \mathbb{R}^m$ converging to some optimal multiplier associated with the equality constraint of (1.2). At each iteration, an auxiliary bound constrained QP has to be solved, so that the approach can be viewed as transforming (1.1) into a sequence of bound constrained convex quadratic subproblems. Two facts contribute to the possible success of this method. First, a bound constraint QP is much easier to solve than (1.1), which has general linear constraints (see [17, 9] and the references therein). Second, because of its dual and constraint convergence, the AL algorithm usually identifies the active constraints of (1.1) in a finite number of iterations. Since often these constraints are also the active constraints of the subproblems close to the solution, the combinatorial aspect of the bound constrained QP's rapidly decreases in intensity as the convergence progresses (and usually disappears after finitely many AL iterations). This reasoning is valid, for instance, when Q is positive definite and strict complementarity holds at the solution.

The AL algorithm also generates primal iterates $(x_k, y_k) \in \mathbb{R}^n \times \mathbb{R}^m$ and is controlled by the convergence of the constraint values to zero: if $\|y_k - Cx_k\|$ is less than a given tolerance, optimality can be considered to be reached. The algorithm is also driven

by a so-called augmentation parameter r_k, whose role on the speed of convergence is major. This paper essentially shows that, provided r_k is larger than a certain positive threshold, the convergence of the constraint norm to zero is *globally* linear, meaning that *at each iteration* the constraint norm decreases by a factor uniformly less than one. This property makes predictable the number of iterations to converge to a given precision and offers a possibility to study the global iterative complexity of the algorithm.

The paper is organized as follows. In section 2, the AL algorithm under investigation is presented with the appropriate level of details. In section 3, we give the tools from convex analysis that are useful for the study of the method. We already set out some of the properties of the algorithm. This section also gives a lemma of general interest, which compares the distances from a point in the positive orthant to an affine space, on the one hand, and to the polyhedron given by the intersection of this affine space and the positive orthant, on the other hand. Section 4 deals with the global linear convergence of the AL algorithm. It starts by showing a global error bound for the dual solution set, in terms of the subgradient of the dual function. The global linear convergence is then seen as an easy consequence of this property. We conclude in section 5 by relating some numerical experiments on a seismic tomography problem and by a discussion on algorithmic consequences.

Notation

We denote the Euclidean norm by $\| \cdot \|$. The distance associated with this norm is denoted by "dist", $B := \{x : \|x\| \leq 1\}$ is the closed unit ball, and $\partial B := \{x : \|x\| = 1\}$ is the unit sphere. We note $\bar{\mathbb{R}} := \mathbb{R} \cup \{-\infty, +\infty\}$. The nonnegative orthant of $\mathbb{R}^n$ is denoted by $\mathbb{R}^n_+ := \{x \in \mathbb{R}^n : x \geq 0\}$, while $\mathbb{R}^n_{++} := \{x \in \mathbb{R}^n : x > 0\}$. The null space and range space of a matrix A are respectively denoted by $N(A)$ and $R(A)$. We write $A \succeq 0$ [resp. $A \succ 0$] to indicate that a symmetric matrix A is positive semi-definite [resp. positive definite].

Let E be a finite dimensional Euclidean space. The indicator function of a set $S \subset E$ is denoted by $\mathcal{I}_S$ (this is the function that vanishes on S and takes the value $+\infty$ outside S). The domain of a function $f : E \to \mathbb{R} \cup \{+\infty\}$ is defined by dom $f := \{x \in E : f(x) < +\infty\}$ and its epigraph by epi $f := \{(x, \alpha) \in E \times \mathbb{R} : f(x) \leq \alpha\}$. As in [12], $\overline{\text{Conv}}(E)$ is the set of functions $f : E \to \mathbb{R} \cup \{+\infty\}$ that are convex (epi f is convex), proper (epi $f \neq \emptyset$), and closed (epi f is closed). The subdifferential at $x \in E$ of a proper convex function $f : E \to \mathbb{R} \cup \{+\infty\}$ is denoted by $\partial f(x)$. We denote by $N_C(x)$ the normal cone at x to a convex set $C \subset E$. The orthogonal projection of a point x onto a nonempty closed convex set C is denoted by $P_C(x)$.

2 An AL algorithm for solving the QP

We **assume throughout that problem (1.2) has a solution** and denote by $\mathcal{S}_P$ the set of its solutions $(\bar{x}, \bar{y})$. The projections of $\mathcal{S}_P$ onto $\mathbb{R}^n$ and $\mathbb{R}^m$ are respectively denoted by $\mathcal{S}_P^x := \{\bar{x} \in \mathbb{R}^n : (\bar{x}, \bar{y}) \in \mathcal{S}_P \text{ for some } \bar{x} \in \mathbb{R}^n\}$ (this is also $\{\bar{x} \in \mathbb{R}^n : (\bar{x}, C\bar{x}) \in \mathcal{S}_P\}$, the solution set of (1.1)) and $\mathcal{S}_P^y := \{\bar{y} \in \mathbb{R}^m : (\bar{x}, \bar{y}) \in \mathcal{S}_P \text{ for some}$

$\bar{x} \in \mathbb{R}^n\}$. Since the constraints of (1.2) are qualified, there exist optimal multipliers, which certainly implies that the affine subspace

$$\Lambda := \{\lambda \in \mathbb{R}^m : C^\top \lambda \in q + R(Q)\} \tag{2.1}$$

is nonempty. Note that $\Lambda = \mathbb{R}^m$ if $Q \succ 0$.

The augmented Lagrangian is obtained by dualizing the equality constraint of (1.2). It is the function $\ell_r : (x, y, \lambda) \in \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^m \mapsto \mathbb{R}$, defined by

$$\ell_r(x, y, \lambda) = \frac{1}{2} x^\top Q x + q^\top x + \lambda^\top (y - Cx) + \frac{r}{2} \|y - Cx\|^2, \tag{2.2}$$

where $r \geq 0$ is called the *augmentation parameter* (see [1, 11, 22]).

We can now give a precise statement of the AL algorithm we study in this paper, which is basically the method of Hestenes [11] and Powell [22] applied to (1.2).

AL ALGORITHM for solving (1.1)

Initialization: choose $\lambda_0 \in \mathbb{R}^m$ and $r_0 > 0$.
Repeat for $k = 0, 1, 2, \ldots$
1. Solve:
$$\min_{\substack{x \\ l \leq y \leq u}} \ell_{r_k}(x, y, \lambda_k). \tag{2.3}$$

   Denote a solution by (x_{k+1}, y_{k+1}).
2. Update the multiplier

$$\lambda_{k+1} = \lambda_k + r_k(y_{k+1} - Cx_{k+1}). \tag{2.4}$$

3. Stop if $y_{k+1} \simeq Cx_{k+1}$.
4. Choose a new augmentation parameter: $r_{k+1} > 0$.

This algorithm deserves some comments.

1. Under the sole assumption that problem (1.1) has a solution, the QP in step (2.3) has also a solution. This fact is clarified in proposition 3.3. This solution is not necessary unique however.
2. Even though (x_{k+1}, y_{k+1}) is not uniquely determined as a solution to (2.3), $y_{k+1} - Cx_{k+1}$ is independent of that solution, so that the multipliers λ_k are unambiguously generated.
3. The augmentation parameter r_k can change from iteration to iteration, but the same value must be used in the AL minimized in step 1 and in the multiplier update in step 2. If the "step-size" in (2.4) is different from r_k (with the aim at minimizing better the dual function, as in [21, section 4.2] for example), several properties of the AL algorithm may no longer hold, such as the finite identification of active constraints (in the presence of strict complementarity) and the global linear convergence of section 4.

4. The larger are the augmentation parameters r_k, the faster is the convergence. The only limitation on a large value for r_k comes from the ill-conditioning that such a value induces in the AL and the resulting difficulty in solving (2.3). Actually, it is clear from the structure of the AL in (2.2) that a large r gives priority to the restoration of the equality constraint, leaving aside the minimization of the Lagrangian (whose role is to provide optimality).

In comparison with an interior point method, which faces the combinatorial aspect of (1.1) by transforming the problem into a sequence of linear systems, the AL algorithm goes around this difficulty by transforming a general QP into a sequence of QP's with simple bounds, which are easier to solve. Indeed, a number of efficient algorithms are available for dealing with the bound constraints on the AL in step 1. A possibility would be to minimize first analytically ℓ_r in y and then to minimize the resulting function in x. Unfortunately this function of x, which is the AL associated with the inequality constrained QP (1.1) [24, 26], has a combinatorial structure (it contains maxima) that is not easier to deal with than the direct numerical minimization of ℓ_r in (x, y) with bounds on y. In our code QPAL [6], used for the numerical experiments of section 5 and in [5], we have adapted to the (x, y) structure of problem (1.2) an active set strategy together with the gradient projection algorithm and conjugate gradient iterations on the activated faces (see [17, 9] and the references therein).

As opposed to standard (non shifted) interior point methods, whose elementary linear systems have an exploding condition number, the AL algorithm does not require the penalty parameter r_k to go to infinity. Actually, any sequence $\{r_k\}$ that remains bounded away from zero guarantees the convergence, even though large values speed it up [28]. Therefore, the bound constrained QP's in step 1 can be maintained reasonably well conditioned, keeping satisfactory the efficiency of a conjugate gradient based solver. This remark reinforces the viewpoint that considers the AL algorithm as a method suitable for large problems.

3 Convex analysis tools

3.1 Duality

As a dual function associated with problem (1.2), we use the one obtained by dualizing its equality constraints. It is the function $\delta : \mathbb{R}^m \to \mathbb{R} \cup \{+\infty\}$ defined by

$$\lambda \mapsto \delta(\lambda) := - \inf_{\substack{x \\ y \in [l,u]}} \left(\frac{1}{2} x^\top Q x + q^\top x + \lambda^\top (y - Cx) \right). \tag{3.1}$$

Clearly $\delta \in \mathrm{C\overline{onv}}(\mathbb{R}^m)$ (it takes a finite value, for instance, when λ is an optimal multiplier associated with the equality constraint of (1.2)).

For a given $\lambda \in \mathbb{R}^m$, (x_λ, y_λ) is a solution to the *Lagrange problem*, the minimization problem in (3.1), if and only if $x_\lambda \in X_\lambda$ and $y_\lambda \in Y_\lambda$, where X_λ is the affine space

$$X_\lambda := \{x \in \mathbb{R}^n : Qx = C^\top \lambda - q\} \tag{3.2}$$

and Y_λ is the Cartesian product of the following intervals

$$(Y_\lambda)_i = \begin{cases} [u_i, u_i] & \text{if } \lambda_i < 0 \\ [l_i, u_i] & \text{if } \lambda_i = 0 \\ [l_i, l_i] & \text{if } \lambda_i > 0. \end{cases} \qquad (3.3)$$

These intervals, with their possible infinite bounds, have to be understood in a broad sense: for example, $[l_i, u_i]$ is the interval $] - \infty, u_i]$ if $l_i = -\infty$ and u_i is finite, $[l_i, l_i]$ is the empty set if $l_i = -\infty$, *etc.* Since the Lagrange problem is always feasible and since a feasible convex quadratic problem has a solution if and only if it is bounded (see [2, theorem 17.1] for example), the domain of δ is the set of λ's for which the Lagrange problem has a solution. Therefore $\operatorname{dom}\delta$ can be written as the nonempty polyhedron

$$\operatorname{dom}\delta = \{\lambda \in \mathbb{R}^m : X_\lambda \neq \emptyset, \ Y_\lambda \neq \emptyset\} = \mathbb{R}^m_{l,u} \cap \Lambda,$$

where we used the fact that $X_\lambda \neq \emptyset$ if and only if $\lambda \in \Lambda$ and the notation

$$\mathbb{R}^m_{l,u} := \{\lambda \in \mathbb{R}^m : \lambda_i \leq 0 \text{ if } l_i = -\infty, \ \lambda_i \geq 0 \text{ if } u_i = +\infty\}.$$

Observe finally that the multivalued function $\lambda \mapsto -Y_\lambda$ is monotone: for λ and $\lambda' \in \mathbb{R}^m$, and for $y_\lambda \in Y_\lambda$ and $y_{\lambda'} \in Y_{\lambda'}$, there holds

$$-(y_{\lambda'} - y_\lambda)^\top(\lambda' - \lambda) \geq 0. \qquad (3.4)$$

Let $Q^\dagger$ be the pseudo-inverse of Q and take the notation

$$H := CQ^\dagger C^\top \qquad \text{and} \qquad v := CQ^\dagger q. \qquad (3.5)$$

Let $\lambda \in \operatorname{dom}\delta$. Then, the minimization in x of the Lagrangian in (3.1) is achieved by any $x_\lambda \in X_\lambda$, hence satisfying $Qx_\lambda + q = C^\top\lambda$. This minimizer is not necessarily unique and we can take the minimum norm minimizer $x_\lambda^\dagger := Q^\dagger(C^\top\lambda - q)$. After substitution in the function minimized in (3.1) and the use of $Q^\dagger Q Q^\dagger = Q^\dagger$, one gets

$$\delta(\lambda) = \sup_{y \in [l,u]} \left(\frac{1}{2}\lambda^\top H\lambda - (v + y)^\top\lambda + \frac{1}{2}q^\top Q^\dagger q \right), \quad \text{for } \lambda \in \operatorname{dom}\delta. \qquad (3.6)$$

On its domain, the dual function δ is therefore the maximum of a finite number of convex quadratic functions (those defined by the argument of the supremum in (3.6) with the components of y set to l_i or u_i; only the finite values of these bounds must be considered), which only differ by their slope at the origin (in particular, they have the same Hessian H).

Lemma 3.1 *The subdifferential of the dual function (3.1) is given at $\lambda \in \operatorname{dom}\delta$ by*

$$\partial\delta(\lambda) = H\lambda - v - Y_\lambda + C(N(Q)).$$

PROOF. We write δ as the sum of three convex functions. Let $\tilde{l}$ and $\tilde{u} \in \mathbb{R}^m$ be chosen such that $\tilde{l} < \tilde{u}$, $\tilde{l}_i = l_i$ if l_i is finite, and $\tilde{u}_i = u_i$ if u_i is finite. Define $\tilde{Y}_\lambda$ by formula (3.3) with l and u respectively replaced by $\tilde{l}$ and $\tilde{u}$. Then the following *finite* value function $\tilde{\delta} \in \mathrm{C}\overline{\mathrm{onv}}(\mathbb{R}^m)$ is identical to the right hand side of (3.6) on $\mathbb{R}^m_{l,u}$:

$$\tilde{\delta}(\lambda) = \sup_{\substack{y = (y_i)_{i=1}^m \\ y_i = \tilde{l}_i \text{ or } \tilde{u}_i}} \left(\frac{1}{2} \lambda^\top H \lambda - (v + y)^\top \lambda + \frac{1}{2} q^\top Q^\dagger q \right).$$

Clearly $\delta = \tilde{\delta} + \mathcal{I}_{\mathbb{R}^m_{l,u}} + \mathcal{I}_\Lambda$, so that theorem 23.8 in [23] implies that for $\lambda \in \mathrm{dom}\,\delta$:

$$\partial \delta(\lambda) = \partial \tilde{\delta}(\lambda) + \partial \mathcal{I}_{\mathbb{R}^m_{l,u}}(\lambda) + \partial \mathcal{I}_\Lambda(\lambda).$$

Equality holds above because $\mathcal{I}_{\mathbb{R}^m_{l,u}}$ and $\mathcal{I}_\Lambda$ are polyhedral and because $\mathrm{ri}\,\mathrm{dom}\,\tilde{\delta}$ ($= \mathbb{R}^m$, ri denotes the relative interior), $\mathrm{dom}\,\mathcal{I}_{\mathbb{R}^m_{l,u}} = \mathbb{R}^m_{l,u}$, and $\mathrm{dom}\,\mathcal{I}_\Lambda = \Lambda$ have a point in common (one in $\mathrm{dom}\,\delta \neq \emptyset$).

To compute $\partial \tilde{\delta}(\lambda)$, we use corollary VI.4.3.2 in [12] (conv denotes the convex hull):

$$\partial \tilde{\delta}(\lambda) = \mathrm{conv}\left\{ H\lambda - v - y : y \in \tilde{Y}_\lambda \text{ and } (y_i = \tilde{l}_i \text{ or } \tilde{u}_i \text{ if } \lambda_i = 0) \right\} = H\lambda - v - \tilde{Y}_\lambda.$$

On the other hand, $\partial \mathcal{I}_{\mathbb{R}^m_{l,u}}(\lambda) = N_{\mathbb{R}^m_{l,u}}(\lambda)$, which is the set of vectors $\nu \in \mathbb{R}^m$ satisfying

$$\begin{cases} \nu_i \geq 0 & \text{when } \lambda_i = 0,\ l_i = -\infty, \text{ and } u_i \text{ is finite} \\ \nu_i \leq 0 & \text{when } \lambda_i = 0,\ l_i \text{ is finite, and } u_i = +\infty \\ \nu_i \in \mathbb{R} & \text{when } \lambda_i = 0,\ l_i = -\infty, \text{ and } u_i = +\infty \\ \nu_i = 0 & \text{when } \lambda_i \neq 0. \end{cases}$$

We deduce from this computation that for $\lambda \in \mathrm{dom}\,\delta$

$$\tilde{Y}_\lambda - \partial \mathcal{I}_{\mathbb{R}^m_{l,u}}(\lambda) = Y_\lambda.$$

Finally $\partial \mathcal{I}_\Lambda(\lambda) = N_\Lambda(\lambda) = \{\mu \in \mathbb{R}^m : C^\top \mu \in R(Q)\}^\perp = C(N(Q))$. Adding the last three subdifferentials provides the formula of $\partial \delta(\lambda)$ given in the statement of the lemma. $\square$

Let us denote by $\mathcal{S}_D$ the set of dual solutions:

$$\mathcal{S}_D := \{\bar{\lambda} \in \mathbb{R}^m : 0 \in \partial \delta(\bar{\lambda})\}.$$

Not surprisingly, this is a convex polyhedron, which can be described in the standard form. It will be useful to make this form explicit and we do so in lemma 3.2 below. For this, we take a partition of $\{1, \ldots, m\}$ into the index sets

$$\begin{aligned} I_l &:= \{i : \bar{y}_i = l_i \text{ for all } (\bar{x}, \bar{y}) \in \mathcal{S}_P\}, \\ J &:= \{i : l_i < \bar{y}_i < u_i \text{ for some } (\bar{x}, \bar{y}) \in \mathcal{S}_P\}, \\ I_u &:= \{i : \bar{y}_i = u_i \text{ for all } (\bar{x}, \bar{y}) \in \mathcal{S}_P\}. \end{aligned} \tag{3.7}$$

7

We also introduce the orthant face $\mathcal{O}$ and the affine subspace $\mathcal{A}$

$$\mathcal{O} := \{\lambda \in \mathbb{R}^m : \lambda_{I_l} \geq 0,\ \lambda_J = 0,\ \lambda_{I_u} \leq 0\},$$
$$\mathcal{A} := \{\lambda \in \mathbb{R}^m : C^\top \lambda = Q\bar{x} + q\}. \tag{3.8}$$

In the definition of $\mathcal{A}$, $\bar{x}$ is an arbitrary primal solution. We have not made this dependence explicit in the symbol of the set since, as shown in the proof of the next lemma, $\mathcal{A}$ does not depend on the choice of $\bar{x} \in \mathcal{S}_P^x$.

Lemma 3.2 *The set of dual solutions $\mathcal{S}_D$ can be written as the intersection*

$$\mathcal{S}_D = \mathcal{O} \cap \mathcal{A}.$$

Furthermore, for any $\bar{\lambda} \in \mathcal{S}_D$ and any $\bar{y} \in \mathcal{S}_P^y$, we have $\bar{y} \in Y_{\bar{\lambda}}$ and $H\bar{\lambda} = v + \bar{y} + C\bar{u}$ for some $\bar{u} \in N(Q)$.

PROOF. Let $\ell(x, y, \lambda) = \frac{1}{2}x^\top Q x + q^\top x + \mathcal{I}_{[l,u]}(y) + \lambda^\top(y - Cx)$ be the Lagrangian function of the problem $\min_{(x,y)}\{\frac{1}{2}x^\top Q x + q^\top x + \mathcal{I}_{[l,u]}(y) : y = Cx\}$, which has the same dual function as problem (1.2). Since the constraint of this problem is qualified, $\bar{\lambda} \in \mathcal{S}_D$ if and only if $0 \in \partial_{(x,y)}\ell(\bar{x}, \bar{y}, \bar{\lambda})$, where $(\bar{x}, \bar{y})$ is an arbitrary primal solution. This can also be written $Q\bar{x} + q = C^\top\bar{\lambda}$ and $0 \in N_{[l,u]}(\bar{y}) + \bar{\lambda}$, which is equivalent to $\bar{\lambda} \in \mathcal{A}_{\bar{x}} \cap \mathcal{O}_{\bar{y}}$, where

$$\mathcal{A}_{\bar{x}} := \{\lambda \in \mathbb{R}^m : C^\top\lambda = Q\bar{x} + q\},$$
$$\mathcal{O}_{\bar{y}} := \{\lambda \in \mathbb{R}^m : \lambda_i \geq 0 \text{ if } \bar{y}_i = l_i,\ \lambda_i = 0 \text{ if } l_i < \bar{y}_i < u_i,\ \lambda_i \leq 0 \text{ if } \bar{y}_i = u_i\}.$$

By varying $(\bar{x}, \bar{y}) \in \mathcal{S}_P$, we see that $\mathcal{A}_{\bar{x}} = \mathcal{A}$ is independent of the chosen primal solution $\bar{x} \in \mathcal{S}_P^x$ and that $\bar{\lambda} \in \cap\{\mathcal{O}_{\bar{y}} : \bar{y} \in \mathcal{S}_P^y\} = \mathcal{O}$.

For proving the second part of the lemma, take $\bar{\lambda} \in \mathcal{S}_D$ and $(\bar{x}, \bar{y}) \in \mathcal{S}_P$. We have shown that $\bar{\lambda} \in \mathcal{A}_{\bar{x}} \cap \mathcal{O}_{\bar{y}}$. Actually, $\bar{\lambda} \in \mathcal{O}_{\bar{y}}$ is equivalent to $\bar{y} \in Y_{\bar{\lambda}}$. By $\bar{\lambda} \in \mathcal{A}_{\bar{x}}$, we have that $C^\top\bar{\lambda} = Q\bar{x} + q$. Multiplying to the left both sides of this equation by $CQ^\dagger$ provides $H\bar{\lambda} = v + \bar{y} + C\bar{u}$, where $\bar{u} := (Q^\dagger Q - I)\bar{x} \in N(Q)$. □

The fact observed in the proof above that the gradient of the criterion of the primal problem at a solution, here $Q\bar{x} + q$, is independent of the chosen solution is a property of general convex problems; see [16, 3]. This fact can also be deduced from the property that the subdifferential of a convex function (here the criterion of problem (1.1)) is constant on the relative interior of a set on which this function is constant (here the solution set).

3.2 Proximality

We will use the fundamental result of Rockafellar [25], according to which the AL algorithm of section 2 is the proximal algorithm on the dual function δ. More precisely, the multiplier λ_{k+1} computed in step 2 of the AL algorithm is also the unique solution to

$$\inf_{\lambda \in \mathbb{R}^m} \left(\delta(\lambda) + \frac{1}{2r_k}\|\lambda - \lambda_k\|^2\right). \tag{3.9}$$

The same parameter $r_k > 0$ is used above and in (2.3). In addition, the optimal value of this problem is the opposite of the optimal value of problem (2.3). The optimality conditions of problem (3.9) can be written $0 \in \partial\delta(\lambda_{k+1}) + (\lambda_{k+1} - \lambda_k)/r_k$. Using (2.4), we see that:

$$Cx_k - y_k \in \partial\delta(\lambda_k), \quad \forall k \geq 1. \tag{3.10}$$

Note that, since λ_{k+1} is uniquely determined as the solution to (3.9), this is also the case for $y_{k+1} - Cx_{k+1}$, even though x_{k+1} and y_{k+1} are not uniquely determined.

Let us now clarify the conditions ensuring that the augmented Lagrange problem (2.3) has a solution.

Proposition 3.3 *The following three properties are equivalent:*
- *(i)* $\operatorname{dom}\delta \neq \emptyset$,
- *(ii) problem (1.1), with some (or any) finite shift of its finite bounds to make it feasible, has a solution,*
- *(iii) for some (or any) $r_k > 0$ and $\lambda_k \in \mathbb{R}^m$, problem (2.3) has a solution.*

PROOF. $[(i) \Rightarrow (iii)]$ Fix $r_k > 0$ and $\lambda_k \in \mathbb{R}^m$ (not necessarily the kth iterate). Since $\operatorname{dom}\delta \neq \emptyset$, the optimal value of (3.9) is finite, so that the optimal value of problem (2.3) is also finite. As a feasible bounded convex quadratic problem, (2.3) must have a solution [2, theorem 17.1].

$[(iii) \Rightarrow (ii)]$ We proceed by contradiction. Suppose that $\tilde{l}$ and $\tilde{u} \in \mathbb{R}^m$ are such that $\tilde{l} < \tilde{u}$, $\tilde{l}_i = -\infty$ iff $l_i = -\infty$, $\tilde{u}_i = +\infty$ iff $u_i = +\infty$, and $[\tilde{l}, \tilde{u}] \cap R(C) \neq \emptyset$ (these bounds $\tilde{l}$ and $\tilde{u}$ result from a finite shift of the finite bounds of (1.1) that makes this problem feasible) and assume that the feasible problem $\min\{f(x) : \tilde{l} \leq Cx \leq \tilde{u}\}$, where $f(x) := (1/2)x^\top Q x + q^\top x$, has no solution. Then, there exists a sequence $\{x_j\}$ such that $Cx_j \in [\tilde{l}, \tilde{u}]$ and $f(x_j) \to -\infty$ when $j \to \infty$ (this is because a bounded feasible convex quadratic problem has a solution). Let $y_j := P_{[l,u]}(Cx_j)$ be the projection of Cx_j onto $[l, u]$. Then $\|y_j - Cx_j\| \leq m^{1/2}\|y_j - Cx_j\|_\infty \leq m^{1/2}\gamma$, where $\gamma := \max(\|\tilde{l} - l\|_\infty, \|\tilde{u} - u\|_\infty)$ (these norms are taken on the finite components of l and u), and (x_j, y_j) is feasible for problem (2.3). On the other hand, for an arbitrary $r_k > 0$ and $\lambda_k \in \mathbb{R}^m$, $\ell_{r_k}(x_j, y_j, \lambda_k) \leq f(x_j) + m^{1/2}\gamma\|\lambda_k\| + (r_k/2)m\gamma^2 \to -\infty$ when $j \to \infty$. Therefore problem (2.3) has no solution.

$[(ii) \Rightarrow (i)]$ Let $\tilde{l}$ and $\tilde{u}$ be some finite shifts of the finite bounds of problem (1.1), such that the problem $\min\{f(x) : \tilde{l} \leq Cx \leq \tilde{u}\}$, with f as in the previous paragraph, has a solution, $\tilde{x}$ say. Since its constraints are qualified, there exist $\tilde{\lambda}^l$ and $\tilde{\lambda}^u$ such that $Q\tilde{x} + q = C^\top(\tilde{\lambda}^l - \tilde{\lambda}^u)$, $\tilde{\lambda}^l \geq 0$, $\tilde{\lambda}^u \geq 0$, $\tilde{\lambda}_i^l = 0$ if $l_i = -\infty$, and $\tilde{\lambda}_i^u = 0$ if $u_i = +\infty$. It is easy to check that $\tilde{\lambda}^l - \tilde{\lambda}^u \in \mathbb{R}_{l,u}^m \cap \Lambda = \operatorname{dom}\delta$. $\square$

If the original quadratic problem (1.1) has a solution, condition (ii) above holds (without having to shift the bounds), so that the augmented Lagrange problem (2.3) has a solution.

3.3 Projection onto a convex polyhedron

This section gives two lemmas related to the projection onto a convex polyhedron. The first lemma has a general interest. It compares the distance from a point x in the positive orthant to a convex polyhedron $\mathcal{X}$ defined in the standard form and the distance from x to the underlying affine space $\mathcal{A}$. It is claimed that the second distance is bounded below by a positive constant (independent of $x \geq 0$) times the first one. Of course, since $\mathcal{X} \subset \mathcal{A}$, $\operatorname{dist}(x, \mathcal{A}) \leq \operatorname{dist}(x, \mathcal{X})$.

Lemma 3.4 *Let A be an $m \times n$ matrix and $b \in \mathbb{R}^m$. Consider the affine subspace $\mathcal{A}$ and the convex polyhedron $\mathcal{X}$ defined by*

$$\mathcal{A} := \{x \in \mathbb{R}^n : Ax = b\} \qquad and \qquad \mathcal{X} := \{x \in \mathbb{R}^n : Ax = b, \ x \geq 0\}.$$

These sets are supposed to be nonempty. Then, there exists a constant $\gamma > 0$ such that

$$\forall x \in \mathbb{R}_+^n, \quad \operatorname{dist}(x, \mathcal{A}) \geq \gamma \operatorname{dist}(x, \mathcal{X}).$$

PROOF. *First stage: reformulation of the statement of the lemma.* By using the triangle inequality, it is easy to see that the conclusion of the lemma is equivalent to claim that

$$\exists \gamma' > 0, \quad \forall x \in \mathbb{R}_+^n, \quad \|x - P_{\mathcal{A}}(x)\| \geq \gamma' \|P_{\mathcal{A}}(x) - P_{\mathcal{X}}(x)\|, \tag{3.11}$$

This inequality suggests that a certain function (whose value is the left hand side of the inequality in (3.11) divided by the factor of γ' in the right hand side) has a positive slope. This is the strategy we follow to establish (3.11).

Instead of testing the validity of the inequality in (3.11) for any $x \in \mathbb{R}_+^n$, we consider all the possible points x^0 that are projections onto $\mathcal{X}$ of points in the positive orthant and reformulate (3.11). For a point $x^0 \in \mathcal{X}$ and a direction $d \in N_{\mathcal{X}}(x^0) \cap N(A) \cap \partial B$, let us introduce the function

$$\varphi : \alpha \in \mathbb{R}_{++} \mapsto \varphi(\alpha) := \inf \left\{ \|A^\top y\| : y \in \mathbb{R}^m, \ x^0 + \alpha d + A^\top y \geq 0 \right\}. \tag{3.12}$$

This function depends on the choice of x^0 and d, but we do not mention this dependence to keep the notation light. Let us show that (3.11) holds if

$$\exists \gamma' > 0, \quad \forall x^0 \in \mathcal{X}, \quad \forall d \in N_{\mathcal{X}}(x^0) \cap N(A) \cap \partial B, \quad \forall \alpha > 0, \quad \varphi(\alpha) \geq \gamma' \alpha. \tag{3.13}$$

Let $x \in \mathbb{R}_+^n$, $x^1 := P_{\mathcal{A}}(x)$, and $x^0 := P_{\mathcal{X}}(x)$. One can assume that $x^1 \neq x^0$ (since otherwise the inequality in (3.11) is trivially satisfied). Set $d := (x^1 - x^0)/\|x^1 - x^0\|$. It is clear that $d \in N(A) \cap \partial B$ (note that both x^1 and $x^0 \in \mathcal{A}$). To show that $d \in N_{\mathcal{X}}(x^0)$, observe that $x^1 := P_{\mathcal{A}}(x)$ implies that $x = x^1 + A^\top y = x^0 + \alpha d + A^\top y$ for $\alpha := \|x^1 - x^0\| > 0$ and a certain $y \in \mathbb{R}^m$. Then, for all $z \in \mathcal{X}$, there holds

$$d^\top (z - x^0) = \frac{1}{\alpha}(x - x^0 - A^\top y)^\top (z - x^0) = \frac{1}{\alpha}(x - x^0)^\top (z - x^0) \leq 0.$$

We have used the fact that $z - x^0 \in N(A)$ and that $x^0 := P_{\mathcal{X}}(x)$ to get the last inequality. This shows that $d \in N_{\mathcal{X}}(x^0)$. Using (3.13) and the fact that $x = x^0 + \alpha d + A^\top y \geq 0$, we see that the inequality in (3.11) holds:

$$\|x - x^1\| = \|A^\top y\| \geq \varphi(\alpha) \geq \gamma' \alpha = \gamma' \|x^1 - x^0\|.$$

The claim (3.13) can be simplified. Observe that $\varphi(\alpha) \geq 0$, that $\varphi(0) = 0$, and that $\varphi(t\alpha) \leq t\varphi(\alpha)$ when $\alpha \geq 0$ and $t \in \,]0,1]$. To prove this last property of φ, assume that $\varphi(\alpha) < \infty$ (otherwise, there is nothing to show). Then take $\varepsilon > 0$ and $y \in \mathbb{R}^m$ such that $x^0 + \alpha d + A^\top y \geq 0$ and $\|A^\top y\| \leq \varphi(\alpha) + \varepsilon$. Since $x^0 \geq 0$, there holds $0 \leq (1-t)x^0 + t(x^0 + \alpha d + A^\top y) = x^0 + t\alpha d + A^\top(ty)$ and therefore $\varphi(t\alpha) \leq \|A^\top(ty)\| \leq t\varphi(\alpha) + \varepsilon$. Since $\varepsilon > 0$ is arbitrary, there holds $\varphi(t\alpha) \leq t\varphi(\alpha)$. Now, this property of φ implies that $\alpha \in \mathbb{R}_{++} \mapsto \varphi(\alpha)/\alpha$ is nondecreasing. Therefore, we have reduced the problem to showing that

$$\exists \gamma' > 0, \quad \forall x^0 \in \mathcal{X}, \quad \forall d \in N_{\mathcal{X}}(x^0) \cap N(A) \cap \partial B, \quad \varphi'(0;1) \geq \gamma', \qquad (3.14)$$

where $\varphi'(0;1)$ denotes the right derivative of φ at zero.

Second stage: control of the decomposition of the normal directions. Consider a point $x^0 \in \mathcal{X}$ having a unitary normal direction in the null space of A, say $d \in N_{x^0}(\mathcal{X}) \cap N(A) \cap \partial B$. Define $I := I(x^0) := \{i : x_i^0 = 0\}$ and $J := J(x^0) := \{i : x_i^0 > 0\}$. These directions d are characterized by the conditions

$$d = A^\top z - r, \quad r_I \geq 0, \quad r_J = 0, \quad Ad = 0, \quad \text{and} \quad \|d\| = 1, \qquad (3.15)$$

for some vectors $z \in \mathbb{R}^m$ and $r \in \mathbb{R}^n$. The decomposition of d in $A^\top z - r$ as above is not necessarily unique. It will be useful to identify a decomposition that provides the smallest value to $\|A^\top z\|$, which is therefore a solution to

$$\begin{cases} \min_{(z,r)} \frac{1}{2}\|A^\top z\|^2 \\ A^\top z - r = d \\ r_I \geq 0 \\ r_J = 0. \end{cases}$$

It is easy to show that this problem has a solution, which is characterized by (3.15) and

$$A(A^\top z - s) = 0, \quad s_I \geq 0, \quad \text{and} \quad s_I^\top r_I = 0, \qquad (3.16)$$

for some vector $s \in \mathbb{R}^n$.

Let us show that

$$\max_{\substack{x^0 \in \mathcal{X} \\ }} \sup_{\substack{d \in N_{\mathcal{X}}(x^0) \\ Ad = 0 \\ \|d\| = 1}} \min_{\substack{(z,r) \in \mathbb{R}^m \times \mathbb{R}^n \\ A^\top z - r = d \\ r_{I(x^0)} \geq 0 \\ r_{J(x^0)} = 0}} \|A^\top z\| \; < \; +\infty. \qquad (3.17)$$

We see on (3.15) that two points $x^0 \in \mathcal{X}$ having the same index set I have the same normal cone. Therefore, the point $x^0 \in \mathcal{X}$ intervenes in (3.17) only through its index

sets I and J. Since there is a finite number of such sets, one can fix x^0, hence I and J. Let us continue by contradiction, assuming that there exists a sequence $\{(d^k, z^k, r^k, s^k)\}$ such that $Ad^k = 0$, $\|d^k\| = 1$, $A^\top z^k - r^k = d^k$, $r_I^k \geq 0$, $r_J^k = 0$, $A(A^\top z^k - s^k) = 0$, $s_I^k \geq 0$, $(s_I^k)^\top r_I^k = 0$, and $\|A^\top z^k\| \to \infty$. Extracting a subsequence if necessary, it can be assumed that $A^\top z^k / \|A^\top z^k\| \to A^\top \bar{z}$, a vector of unit norm. Since $\{d^k\}$ is bounded, the identity $A^\top z^k - r^k = d^k$ shows that $r^k / \|A^\top z^k\|$ converges to $\bar{r} := A^\top \bar{z}$. Multiplying the identity $A(A^\top z^k - s^k) = 0$ by $\bar{z}$, one finds for sufficiently large k

$$ 0 = \bar{z}^\top A A^\top z^k - \bar{r}^\top s^k = \bar{z}^\top A A^\top z^k, $$

because, when $\bar{r}_i > 0$, then $i \in I$ and, for all sufficiently large k, $r_i^k > 0$, so that $s_i^k = 0$. Dividing the right hand side by $\|A^\top z^k\|$ and taking the limit, one would find $A^\top \bar{z} = 0$, which provides the expected contradiction.

Third stage: lower bound for $\varphi'(0; 1)$ and conclusion. Let us introduce $\upsilon : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$, the value function of the problem

$$ \begin{cases} \inf_y \|A^\top y\| \\ x^0 + A^\top y \geq 0, \end{cases} \tag{3.18} $$

which is the proper convex function defined by $\upsilon(p) := \inf\{\|A^\top y\| : x^0 + A^\top y \geq p\}$. Then, for fixed $x^0 \in \mathcal{X}$ and $d \in N_\mathcal{X}(x^0) \cap N(A) \cap \partial B$, $\varphi(\alpha)$ defined by (3.12) can be written $\varphi(\alpha) = \upsilon(-\alpha d)$. Therefore

$$ \varphi'(0; 1) = \upsilon'(0; -d) \geq -g^\top d, \quad \forall g \in \partial \upsilon(0). \tag{3.19} $$

As for the subdifferential $\partial \upsilon(0)$, it is formed of the optimal multipliers associated with the constraint of (3.18), which are the g-parts of the pairs (g, u) satisfying

$$ g \in u + N(A), \quad \|u\| \leq 1, \quad g \geq 0, \quad \text{and} \quad (x^0)^\top g = 0. \tag{3.20} $$

Let $d = A^\top z - r$ be a decomposition of d satisfying (3.15)-(3.16). If $A^\top z = 0$, then $g := -\alpha d = \alpha r$ is a subgradient of υ at zero for any $\alpha \geq 0$ (the conditions (3.20) are satisfied with $u = 0$; recall that $d \in N(A)$), so that (3.19) shows that $\varphi'(0; 1) = +\infty$. If $A^\top z \neq 0$, then $g := (A^\top z - d)/\|A^\top z\| = r/\|A^\top z\|$ is a subgradient of υ at zero (the conditions (3.20) are satisfied with $u = A^\top z / \|A^\top z\|$). Then (3.19) shows that $\varphi'(0; 1) \geq 1/\|A^\top z\|$. Since for these decompositions, stage 2 of the proof has shown that $A^\top z$ is bounded, (3.14) holds and, consequently, the result is proven. $\qquad \square$

As shown by the following example, lemma 3.4 no longer holds with all its generality when $\mathcal{X}$ is the intersection of an affine space $\mathcal{A}$ and an arbitrary closed convex cone.

Example 3.5 Let us introduce the following closed convex cone $K := \{x \in \mathbb{R}^3 : x_2 x_3 \geq x_1^2, x_2 \geq 0, x_3 \geq 0\}$, the 1×3 matrix $A := (0\ 1\ 0)$, and $b = 0 \in \mathbb{R}$. Define the affine space $\mathcal{A}$ and its intersection with K by

$$ \mathcal{A} := \{x \in \mathbb{R}^3 : Ax = b\} = \{x \in \mathbb{R}^3 : x_2 = 0\}, $$
$$ \mathcal{X} := K \cap \mathcal{A} = \{x \in \mathbb{R}^3 : x_1 = x_2 = 0, \ x_3 \geq 0\}. $$

Then the conclusion of lemma 3.4 does not hold for these sets $\mathcal{A}$ and $\mathcal{X}$. To see this, fix $x_1 > 0$ and consider the points $x^t := (x_1, x_1^2/t, t)$ for $t \uparrow +\infty$. Clearly $x^t \in K$, $P_{\mathcal{A}}(x^t) = (x_1, 0, t)$, and $P_{\mathcal{X}}(x^t) = (0, 0, t)$. Therefore $\|x^t - P_{\mathcal{A}}(x^t)\|/\|P_{\mathcal{A}}(x^t) - P_{\mathcal{X}}(x^t)\| = x_1/t$, which is not bounded away from zero. $\qquad\square$

Actually, it will be useful below to have the following relaxed version of lemma 3.4. This one allows the projected point x not to belong to $\mathbb{R}^n_+$. This point must however be sufficiently close to the positive orthant with respect to its distance to $\mathcal{X}$.

Corollary 3.6 *Assume the framework defined in the statement of lemma 3.4. Then, there exist two constants $\tau > 0$ and $\gamma > 0$ such that for all $x \in \mathbb{R}^n$,*

$$\operatorname{dist}(x, \mathbb{R}^n_+) \leq \tau \operatorname{dist}(x, \mathcal{X}) \qquad \Longrightarrow \qquad \operatorname{dist}(x, \mathcal{A}) \geq \gamma \operatorname{dist}(x, \mathcal{X}).$$

PROOF. Let γ be the constant given by lemma 3.4 and set

$$\tau := \frac{\gamma}{4(1 + \gamma)}.$$

Let x be such that $\operatorname{dist}(x, \mathbb{R}^n_+) \leq \tau \operatorname{dist}(x, \mathcal{X})$. To simplify the notation, let us define

$$x^0 := P_{\mathcal{X}}(x), \qquad x^1 := P_{\mathcal{A}}(x), \qquad \text{and} \qquad \bar{x} := P_{\mathbb{R}^n_+}(x),$$

Using several times the triangle inequality, lemma 3.4, the non-expansiveness of the projectors $P_{\mathcal{A}}$ and $P_{\mathcal{X}}$, and the definition of τ, one can write

$$
\begin{aligned}
\|x - x^1\| &\geq \|\bar{x} - P_{\mathcal{A}}(\bar{x})\| - \|P_{\mathcal{A}}(\bar{x}) - P_{\mathcal{A}}(x)\| - \|x - \bar{x}\| \\
&\geq \gamma\|\bar{x} - P_{\mathcal{X}}(\bar{x})\| - 2\|x - \bar{x}\| \\
&\geq \gamma\|\bar{x} - x^0\| - (2 + \gamma)\|x - \bar{x}\| \\
&\geq \gamma\|x - x^0\| - 2(1 + \gamma)\|x - \bar{x}\| \\
&\geq \gamma\|x - x^0\| - 2\tau(1 + \gamma)\|x - x^0\| \\
&= \frac{\gamma}{2}\|x - x^0\|.
\end{aligned}
$$

This is the expected inequality. $\qquad\square$

The following lemma will be also useful. If $I \subset \{1, \ldots, n\}$, we denote by I^c the complementary set of I in $\{1, \ldots, n\}$.

Lemma 3.7 *Let A be an $m \times n$ matrix, $b \in \mathbb{R}^m$, $I \subset \{1, \ldots, n\}$, and $\varphi : \mathbb{R}^n \to \mathbb{R}$ be a convex differentiable function. If $\bar{x}$ is a solution to the problem*

$$\min\left\{\varphi(x) : Ax = b, \ x_I \geq 0, \ x_{I^c} = 0\right\},$$

then there is a subset of indices $J \subset \{1, \ldots, n\}$, containing I, such that $\bar{x}$ is also a solution to the problem

$$\min\left\{\varphi(x) : Ax = b, \ x_J \geq 0, \ x_{J^c} \leq 0\right\}.$$

PROOF. The constraints of the first problem are affine, hence qualified. Therefore, there exist vectors $y \in \mathbb{R}^m$ and $s \in \mathbb{R}^n$ such that

$$\nabla \varphi(\bar{x}) + A^\top y + s = 0, \quad \bar{x}_I \geq 0, \quad s_I \leq 0, \quad s_I^\top \bar{x}_I = 0, \quad \bar{x}_{I^c} = 0.$$

Define

$$J := I \cup \{i \in I^c : s_i \leq 0\}.$$

Then

$$\nabla \varphi(\bar{x}) + A^\top y + s = 0, \quad \bar{x}_J \geq 0, \quad s_J \leq 0, \quad s_J^\top \bar{x}_J = 0,$$
$$\bar{x}_{J^c} \leq 0, \quad s_{J^c} \geq 0, \quad s_{J^c}^\top \bar{x}_{J^c} = 0.$$

By convexity, these conditions suffice to show that $\bar{x}$ is also a solution to the second problem. □

4 Global linear convergence of the algorithm

The global linear convergence of the AL algorithm will be shown in section 4.2 to be a consequence of the radial Lipschitz continuity of the multifunction $\partial \delta^{-1}$, the reciprocal of the subdifferential of the dual function (this argument is taken from [28]). The latter property is the subject of section 4.1.

4.1 Global dual error bound

Two normed spaces E and F being given, a multifunction $T : E \twoheadrightarrow F$ is said to be *radially Lipschitz continuous at $x_0 \in E$ with constant $L \geq 0$* if for all $x \in E$ and all $y \in T(x)$, there holds $\mathrm{dist}(y, T(x_0)) \leq L\|x - x_0\|$ ("dist" denotes here the distance associated with the norm of F). Consider the multifunction

$$\partial \delta^{-1} : g \in \mathbb{R}^m \mapsto \{\lambda \in \mathbb{R}^m : g \in \partial \delta(\lambda)\} \subset \mathbb{R}^m,$$

where δ is the dual function defined in (3.1). Clearly $\partial \delta^{-1}(0) = \mathcal{S}_D$, the set of dual solutions. Then $\partial \delta^{-1}$ is radially Lipschitz continuous at 0 with constant $L \geq 0$, if

$$\forall \lambda \in \mathbb{R}^m, \ \forall g \in \partial \delta(\lambda) : \ \mathrm{dist}(\lambda, \mathcal{S}_D) \leq L\|g\|. \tag{4.1}$$

Such a property is sometimes called a *global error bound* for the dual solution set $\mathcal{S}_D$ in terms of the dual function subgradient (see the review paper by Pang [21] and the contribution of Izmailov and Solodov [13]). In this section, we show that this property holds in a weaker form: λ has to stay at a bounded distance from $\mathcal{S}_D$ (the Lipschitz constant L depends on this distance). Nevertheless, this property still has a global nature, since λ is not required to be close to $\mathcal{S}_D$ and g is not required to be close to 0.

To show that this property is natural, consider first a quadratic problem with only equality constraints:

$$\begin{cases} \inf_x \ \frac{1}{2} x^\top Q x + q^\top x \\ Cx = b. \end{cases} \tag{4.2}$$

It is assumed that this problem is convex ($Q \succcurlyeq 0$) and has a solution. It is therefore feasible: $b \in R(C)$. Since the constraint is qualified, there exist optimal multipliers, which implies that the affine subspace Λ defined in (2.1) is nonempty.

Using the pseudo-inverse $Q^\dagger$ of Q, the symmetric matrix $H \succcurlyeq 0$, and the vector v defined in (3.5), the dual function δ associated with problem (4.2) can be written

$$\delta(\lambda) = \begin{cases} \frac{1}{2}\lambda^\top H \lambda - (v+b)^\top \lambda + \frac{1}{2}q^\top Q^\dagger q & \text{for } \lambda \in \Lambda \\ +\infty & \text{otherwise.} \end{cases} \tag{4.3}$$

A computation like in the proof of lemma 3.1 shows that

$$\partial \delta(\lambda) = H\lambda - v - b + C(N(Q)), \quad \text{for } \lambda \in \Lambda.$$

Since $\mathcal{S}_D$ is defined as the set of minimizers of δ, one finds

$$\mathcal{S}_D = \{\bar{\lambda} \in \Lambda : H\bar{\lambda} \in v + b + C(N(Q))\}.$$

It is useful to introduce

$$\sigma := \inf_{\substack{\mu \in \partial B \cap R(C) \\ C^\top \mu \in R(Q^\dagger)}} \mu^\top H \mu, \tag{4.4}$$

which, by definition, takes the value $+\infty$ when $\{\mu \in R(C) : C^\top \mu \in R(Q^\dagger)\} = \{0\}$. Below, the smallest *nonzero* eigenvalue of a *zero* matrix is defined to be $+\infty$.

Lemma 4.1 *The value in $\bar{\mathbb{R}}$ defined by (4.4) satisfies $\sigma > 0$. It is the smallest nonzero eigenvalue of H when $R(C^\top) \subset R(Q^\dagger)$.*

PROOF. We only have to consider the case when $\{\mu \in R(C) : C^\top \mu \in R(Q^\dagger)\} \neq \{0\}$. Then, $Q^\dagger \neq 0$ (because $C^\top \mu = 0$ and $\mu \in R(C)$ imply that $\mu = 0$) and $C \neq 0$. Now, when $C^\top \mu \in R(Q^\dagger) = N(Q^\dagger)^\perp$, $\mu^\top H \mu = \mu^\top C Q^\dagger C^\top \mu \geq \zeta_{\min}(Q^\dagger)\|C^\top \mu\|^2$, where $\zeta_{\min}(Q^\dagger)$ is the smallest nonzero eigenvalue of $Q^\dagger$. On the other hand, when $\mu \in R(C)$, $\|C^\top \mu\| \geq \sigma_{\min}(C)\|\mu\|$, where $\sigma_{\min}(C)$ is the smallest nonzero singular value of C. We have shown that

$$\sigma \geq \zeta_{\min}(Q^\dagger)\, \sigma_{\min}(C)^2 > 0.$$

Suppose now that $R(C^\top) \subset R(Q^\dagger)$. Then $\sigma = \inf\{\mu^\top H \mu : \mu \in \partial B \cap R(C)\}$, so that σ will be the smallest nonzero eigenvalue of H if we show that $R(C) = N(H)^\perp$ or that $N(H) = N(C^\top)$. The inclusion $N(C^\top) \subset N(H)$ is clear. Conversely, let $\nu \in N(H)$, which reads $CQ^\dagger C^\top \nu = 0$. This implies that $C^\top \nu \in N(Q^\dagger) = R(Q^\dagger)^\perp \subset N(C)$, by assumption. Then $C^\top \nu = 0$. $\qquad\square$

Note that when $R(C^\top) \not\subset R(Q^\dagger)$, σ is not the smallest nonzero eigenvalue of H. Here is an example

$$Q^\dagger = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \quad \text{and} \quad C = \begin{pmatrix} 1 & 1 \end{pmatrix}.$$

Then $H = 1$, while $\sigma = +\infty$ since there is no nonzero μ such that $C^\top \mu \in R(Q^\dagger)$.

Proposition 4.2 *Consider problem* (4.2) *with* $Q \succcurlyeq 0$ *and suppose that it has a solution. Then property* (4.1) *is satisfied by the dual function* (4.3), *with the Euclidean norm and a constant* $L = 1/\sigma$, *where* σ *is defined by* (4.4).

PROOF. Since problem (4.2) has a solution and its constraint is qualified, $\mathcal{S}_D$ is nonempty (it is identical to the set of optimal multipliers). To prove (4.1), we only have to consider the dual variables $\lambda \in \Lambda$, since otherwise $\partial \delta(\lambda)$ is empty. Note also that we only have to consider the case when $H \neq 0$ since otherwise $\mathcal{S}_D = \Lambda$ and (4.1) is trivially satisfied with $L = 0$.

Let $\lambda \in \operatorname{dom} \delta = \Lambda$, $g \in \partial \delta(\lambda)$, and $\bar{\lambda}$ be the projection of λ onto $\mathcal{S}_D$. We have for some u and $\bar{u} \in N(Q)$

$$g = H\lambda - (v + b) + Cu \qquad \text{and} \qquad 0 = H\bar{\lambda} - (v + b) + C\bar{u}.$$

Subtracting these two identities and taking the scalar product with $(\lambda - \bar{\lambda})$ yield

$$g^\top (\lambda - \bar{\lambda}) = (\lambda - \bar{\lambda})^\top H (\lambda - \bar{\lambda}) + (u - \bar{u})^\top C^\top (\lambda - \bar{\lambda}).$$

The last term vanishes, since $C^\top (\lambda - \bar{\lambda}) \in R(Q)$ and $Q(u - \bar{u}) = 0$. Now observe that $\lambda - \bar{\lambda} \in R(C)$, since $\bar{\lambda} + N(C^\top) \subset \mathcal{S}_D$ (equality holds actually), and that $C^\top (\lambda - \bar{\lambda}) \in R(Q) = R(Q^\dagger)$, since both λ and $\bar{\lambda} \in \Lambda$. Therefore

$$g^\top (\lambda - \bar{\lambda}) \geq \left(\inf_{\substack{\mu \in \partial B \cap R(C) \\ C^\top \mu \in R(Q^\dagger)}} \mu^\top H \mu \right) \|\lambda - \bar{\lambda}\|^2 = \sigma \|\lambda - \bar{\lambda}\|^2.$$

Now (4.1) with $L = 1/\sigma$ follows by using the Cauchy-Schwarz inequality on the left hand side. $\square$

When C is surjective and $Q \succ 0$, extending this result to the dual function associated with the strictly convex quadratic problem (1.2) is an easy exercise (then H is positive definite and L is the inverse of the smallest eigenvalue of H). On the other hand, when C is not surjective, property (4.1) cannot hold without being lightly weakened, as shown by the following example.

Example 4.3 Consider the special QP with a single inactive constraint ($m = 1$, $C = 0$, and $l < 0 < u$) and a zero optimal value ($q = 0$). Then δ is the function

$$\delta(\lambda) = \begin{cases} -u\lambda & \text{if } \lambda \leq 0 \\ -l\lambda & \text{if } \lambda > 0. \end{cases}$$

Clearly, (4.1) can hold only if λ is not too far from the dual solution set: $|\lambda| \leq L \min(-l, u)$. $\square$

The analysis of the inequality constrained QP is more difficult than the one of problem (4.2), since it has to cover simultaneously two different cases: the quadratic dual function of the equality constrained QP (4.2) and the sharp dual function of the previous example.

In the case of an inequality constrained QP, it will be shown that the Lipschitz constant L may also depend on the largest gap Δ between the $\bar{y} \in \mathcal{S}_P^y$ and the inactive bounds. More precisely, Δ is defined by

$$\Delta := \sup_{\bar{y} \in \mathcal{S}_P^y} \min \left(\min_{i \in I_l \cup J} (u_i - \bar{y}_i), \min_{i \in J \cup I_u} (\bar{y}_i - l_i) \right), \tag{4.5}$$

where the index sets I_l, J, and I_u are introduced in (3.7). If $J = \emptyset$, either I_l or $I_u \neq \emptyset$, and the fact that $l < u$ implies that $\Delta > 0$. If $J \neq \emptyset$, the convexity of $\mathcal{S}_P$ implies that there is a $\bar{y} \in \mathcal{S}_P^y$ such that $l_J < \bar{y}_J < u_J$, in which case also $\Delta > 0$. The dependence of L on Δ is clearly visible in example 4.3: for λ at a unit distance from the solution, we must have $L \geq 1/\min(-l, u) = 1/\Delta$. This lower bound on L goes to infinity when l or u tends to zero, and it goes to zero when $l \to -\infty$ and $u \to +\infty$.

Proposition 4.4 *Consider problem* (1.1) *with* $Q \succcurlyeq 0$ *and suppose that it has a solution. Then, for any bounded set* $\mathcal{B} \subset \mathbb{R}^m$, *there exists a constant* L, *such that*

$$\forall \lambda \in \mathcal{S}_D + \mathcal{B}, \ \forall g \in \partial\delta(\lambda): \ \text{dist}(\lambda, \mathcal{S}_D) \leq L\|g\|. \tag{4.6}$$

PROOF. *First stage: definition of* L. We know that $\mathcal{S}_D \neq \emptyset$. Let $\mathcal{B}$ be a bounded set in $\mathbb{R}^m$, i.e., $\mathcal{B} \subset \beta B$ for some $\beta > 0$. To make the proof rigorous, we now define $L > 0$, even though the motivation for its definition will not look quite clear at this point.

Let $\mathcal{K}$ be the collection of index sets $K \subset \{1, \ldots, m\}$ such that $I_l \subset K$ and $I_u \subset K^c := \{1, \ldots, m\} \backslash K$ (the index sets I_l and I_u are defined in (3.7)). With any index set $K \subset \{1, \ldots, m\}$, we associate the orthant

$$\mathcal{O}_K := \{\lambda \in \mathbb{R}^m : \lambda_K \geq 0, \lambda_{K^c} \leq 0\}.$$

Define $\mathcal{O}$ and $\mathcal{A}$ by (3.8). For any index set $K \in \mathcal{K}$, $\mathcal{O}_K \cap \mathcal{A}$ is nonempty (since it contains $\mathcal{S}_D = \mathcal{O} \cap \mathcal{A}$, see lemma 3.2). Therefore, with an index set $K \in \mathcal{K}$, corollary 3.6 associates two constants $\tau_K > 0$ and $\gamma_K > 0$ such that for any $\lambda \in \mathbb{R}^m$:

$$\text{dist}(\lambda, \mathcal{O}_K) \leq \tau_K \text{dist}(\lambda, \mathcal{O}_K \cap \mathcal{A}) \implies \text{dist}(\lambda, \mathcal{A}) \geq \gamma_K \text{dist}(\lambda, \mathcal{O}_K \cap \mathcal{A}).$$

Since $\mathcal{K}$ is finite, the constants

$$\tau := \min_{K \in \mathcal{K}} \tau_K \quad \text{and} \quad \gamma := \min_{K \in \mathcal{K}} \gamma_K$$

are positive. Therefore, we have found two constants $\tau > 0$ and $\gamma > 0$ such that, for any $K \in \mathcal{K}$, there holds

$$\lambda \in \mathcal{O}_K^\tau := \{\lambda' \in \mathbb{R}^m : \text{dist}(\lambda', \mathcal{O}_K) \leq \tau \text{dist}(\lambda', \mathcal{O}_K \cap \mathcal{A})\}$$
$$\implies \text{dist}(\lambda, \mathcal{A}) \geq \gamma \text{dist}(\lambda, \mathcal{O}_K \cap \mathcal{A}). \tag{4.7}$$

Recall the definitions (4.4) of $\sigma > 0$ and (4.5) of $\Delta > 0$. Then, the constant $L \geq 0$ is defined by

$$L := \max \left(\frac{1}{\sigma\gamma^2}, \frac{\beta}{\tau\Delta} \right). \tag{4.8}$$

17

In this formula, the constants $\sigma > 0$ and $\Delta > 0$ may take the value $+\infty$. Therefore, L is finite, but can vanish.

Second stage: proof of (4.6). Fix $\lambda \in \mathcal{S}_D + \mathcal{B}$ and $g \in \partial\delta(\lambda)$ (necessarily, $\lambda \in \operatorname{dom}\delta$). Denote the projection of λ onto $\mathcal{S}_D$ by $\bar{\lambda} := P_{\mathcal{S}_D}(\lambda)$. Observe that,

$$\|\lambda - \bar{\lambda}\| \leq \beta. \tag{4.9}$$

Let $\varepsilon \in \,]0, \Delta[$ and define L_ε by formula (4.8), but with $\Delta - \varepsilon$ in place of Δ. Observe now that showing

$$g^\top(\lambda - \bar{\lambda}) \geq \frac{1}{L_\varepsilon}\|\lambda - \bar{\lambda}\|^2 \tag{4.10}$$

suffices to conclude the proof since then the inequality in (4.6) follows from the Cauchy-Schwarz inequality applied to the left hand side of (4.10) and the fact that ε can be chosen arbitrarily close to zero.

From the form of the subdifferential $\partial\delta(\lambda)$ given by lemma 3.1, we have for some $y_\lambda \in Y_\lambda$, some $y_{\bar{\lambda}} \in Y_{\bar{\lambda}}$, and some $u, \bar{u} \in N(Q)$:

$$g = H\lambda - v - y_\lambda + Cu \qquad \text{and} \qquad 0 = H\bar{\lambda} - v - y_{\bar{\lambda}} + C\bar{u}. \tag{4.11}$$

According to lemma 3.2, $y_{\bar{\lambda}}$ can be chosen arbitrarily in $\mathcal{S}_P^y$ and we take it such that

$$\min\left(\min_{i \in I_l \cup J}(u_i - \bar{y}_i), \min_{i \in J \cup I_u}(\bar{y}_i - l_i)\right) \geq \Delta - \varepsilon. \tag{4.12}$$

As in the proof of proposition 4.2, $(u - \bar{u})^\top C^\top(\lambda - \bar{\lambda}) = 0$, because $C^\top(\lambda - \bar{\lambda}) \in R(Q)$ (both λ and $\bar{\lambda} \in \operatorname{dom}\delta \subset \Lambda$) and $Q(u - \bar{u}) = 0$. Therefore, subtracting the identities in (4.11) and taking the scalar product with $(\lambda - \bar{\lambda})$ yield

$$g^\top(\lambda - \bar{\lambda}) = (\lambda - \bar{\lambda})^\top H(\lambda - \bar{\lambda}) - (y_\lambda - y_{\bar{\lambda}})^\top(\lambda - \bar{\lambda}). \tag{4.13}$$

We will get (4.10) by finding a lower bound of the right hand side of (4.13). Note that the two terms are nonnegative (this is clear for the first one, since H is positive semi-definite; for the second one, use the monotonicity property (3.4)).

Since $\bar{\lambda} = P_{\mathcal{A} \cap \mathcal{O}}(\lambda)$, by lemma 3.7, one can find an index set $K \subset \mathcal{K}$ such that $\bar{\lambda} = P_{\mathcal{A} \cap \mathcal{O}_K}(\lambda)$. We analyze successively two complementary cases, using the set $\mathcal{O}_K^\tau$ defined in (4.7) and $\lambda^t := (1-t)\bar{\lambda} + t\lambda$ for $t \in \mathbb{R}$.

Case A: there exists a $t \in \,]0,1]$ such that $\lambda^t \in \mathcal{O}_K^\tau$. In this case, we work on the first term in the right hand side of (4.13), discarding the second one. Because $P_{\mathcal{A} \cap \mathcal{O}_K}(\lambda^t) = \bar{\lambda}$, (4.7) gives

$$\gamma\|\lambda^t - \bar{\lambda}\| \leq \|\lambda^t - P_{\mathcal{A}}(\lambda^t)\|.$$

Decompose $\lambda^t - \bar{\lambda} = \mu_0 + \mu_1$, where $\mu_0 \in N(C^\top)$ and $\mu_1 \in R(C)$, and observe that $C^\top\mu_1 = C^\top(\lambda^t - \bar{\lambda}) \in R(Q) = R(Q^\dagger)$ (since both λ^t and $\bar{\lambda} \in \operatorname{dom}\delta \subset \Lambda$). Then, using the definition (4.4) of σ, one finds

$$(\lambda^t - \bar{\lambda})^\top H(\lambda^t - \bar{\lambda}) = \mu_1^\top H\mu_1 \geq \sigma\|\mu_1\|^2.$$

From the fact that $\mu_1 \in R(C)$ and that $C^\top(\lambda^t - \mu_1) = C^\top(\bar\lambda + \mu_0) = C^\top\bar\lambda = Q\bar{x} + q$, we deduce that $\mu_1 = \lambda^t - P_{\mathcal{A}}(\lambda^t)$. Therefore

$$(\lambda^t - \bar\lambda)^\top H(\lambda^t - \bar\lambda) \geq \sigma\gamma^2\|\lambda^t - \bar\lambda\|^2.$$

Since $\lambda - \bar\lambda = (\lambda^t - \bar\lambda)/t$, we also have

$$(\lambda - \bar\lambda)^\top H(\lambda - \bar\lambda) \geq \sigma\gamma^2\|\lambda - \bar\lambda\|^2.$$

Discarding the second term in the right hand side of (4.13) (it is nonnegative) and using the definition of L in (4.8), it follows that

$$g^\top(\lambda - \bar\lambda) \geq (\lambda - \bar\lambda)^\top H(\lambda - \bar\lambda) \geq \sigma\gamma^2\|\lambda - \bar\lambda\|^2 \geq \frac{1}{L}\|\lambda - \bar\lambda\|^2 \geq \frac{1}{L_\varepsilon}\|\lambda - \bar\lambda\|^2,$$

which is the expected inequality (4.10).

   Case B: for any $t \in\ ]0,1]$, $\lambda^t \notin \mathcal{O}_K^\tau$. In this case, we work on the second term in the right hand side of (4.13), discarding the first one. Let us start by choosing $t \in\ ]0,1]$ sufficiently small such that $\lambda_i^t\bar\lambda_i > 0$ when $\bar\lambda_i \neq 0$; by assumption, this $\lambda^t \notin \mathcal{O}_K^\tau$. Let $g^t \in \partial\delta(\lambda^t)$, so that $g^t = H\lambda^t - v - y_{\lambda^t} + Cu^t$ for some $y_{\lambda^t} \in Y_{\lambda^t}$ and some $u^t \in N(Q)$ (compare with (4.11)). Proceeding as before, we get an identity like (4.13):

$$(g^t)^\top(\lambda^t - \bar\lambda) = (\lambda^t - \bar\lambda)^\top H(\lambda^t - \bar\lambda) - (y_{\lambda^t} - y_{\bar\lambda})^\top(\lambda^t - \bar\lambda). \tag{4.14}$$

Denote by $\lambda_K^t := P_{\mathcal{O}_K}(\lambda^t)$ the projection of λ^t onto $\mathcal{O}_K$ and decompose

$$-(y_{\lambda^t} - y_{\bar\lambda})^\top(\lambda^t - \bar\lambda) = -(y_{\lambda^t} - y_{\bar\lambda})^\top(\lambda^t - \lambda_K^t) - (y_{\lambda^t} - y_{\bar\lambda})^\top(\lambda_K^t - \bar\lambda).$$

The last term in the right hand side is nonnegative. Indeed, by the choice of t, if $\bar\lambda_i \neq 0$, one has $\lambda_i^t\bar\lambda_i > 0$ and therefore $(y_{\lambda^t} - y_{\bar\lambda})_i = 0$ (see the definition (3.3) of Y_λ). The only nonzero terms of the last scalar product are therefore of the form $(y_{\bar\lambda} - y_{\lambda^t})_i(\lambda_K^t)_i$. If $(\lambda_K^t)_i > 0$, one has $\lambda_i^t > 0$ (since λ_K^t is the projection of λ^t onto $\mathcal{O}_K$) and therefore $(y_{\lambda^t})_i = l_i$, so that the term can be written $((y_{\bar\lambda})_i - l_i)(\lambda_K^t)_i \geq 0$ (since $l_i \leq (y_{\bar\lambda})_i \leq u_i$). Similarly, the term is nonnegative when $(\lambda_K^t)_i < 0$. Therefore

$$-(y_{\lambda^t} - y_{\bar\lambda})^\top(\lambda^t - \bar\lambda) \geq -(y_{\lambda^t} - y_{\bar\lambda})^\top(\lambda^t - \lambda_K^t) = \sum_{i \in I_{K,\lambda^t}} (y_{\bar\lambda} - y_{\lambda^t})_i(\lambda^t - \lambda_K^t)_i,$$

where we have introduced the index set

$$I_{K,\lambda^t} := \{i : \lambda_i^t \neq (\lambda_K^t)_i\}.$$

Let us show that all the terms of the sum on the indices $i \in I_{K,\lambda^t}$ above are positive. Observe first that $(\lambda_K^t)_i = 0$ (since $\lambda_i^t \neq (\lambda_K^t)_i$ and λ_K^t is the projection of λ^t onto the orthant $\mathcal{O}_K$). On the other hand, if $\lambda_i^t > 0$, then $(y_{\lambda^t})_i = l_i$ and $l_i < (y_{\bar\lambda})_i \leq u_i$ (this is because $i \in K^c$ when $\lambda_i^t > 0$ and $(\lambda_K^t)_i = 0$, and because $I_l \subset K$); therefore $(y_{\bar\lambda} - y_{\lambda^t})_i = (y_{\bar\lambda})_i - l_i \geq \Delta - \varepsilon > 0$ (see (4.12)). Similarly, if $\lambda_i^t < 0$, then $(y_{\lambda^t})_i = u_i$,

19

$i \in K$, and $(y_{\bar\lambda} - y_{\lambda^t})_i = (y_{\bar\lambda})_i - u_i \leq -(\Delta - \varepsilon) < 0$. In particular, all the terms of the sum are positive. Therefore, we get ($\|\cdot\|_1$ denotes the ℓ_1-norm)

$$-(y_{\lambda^t} - y_{\bar\lambda})^\top (\lambda^t - \bar\lambda) \geq (\Delta - \varepsilon)\|\lambda^t - \lambda_K^t\|_1 \geq (\Delta - \varepsilon)\|\lambda^t - \lambda_K^t\| \geq \tau(\Delta - \varepsilon)\|\lambda^t - \bar\lambda\|, \quad (4.15)$$

where the last inequality comes from the fact that $\lambda^t \notin \mathcal{O}_K^\tau$ and that $P_{\mathcal{A} \cap \mathcal{O}_K}(\lambda^t) = \bar\lambda$. We can now conclude:

$$
\begin{aligned}
g^\top(\lambda - \bar\lambda) \quad &\geq \quad (g^t)^\top(\lambda - \bar\lambda) \qquad \text{[monotonicity of the subdifferential]} \\
&= \quad \frac{1}{t}(g^t)^\top(\lambda^t - \bar\lambda) \qquad \text{[definition of λ^t]} \\
&\geq \quad -\frac{1}{t}(y_{\lambda^t} - y_{\bar\lambda})^\top(\lambda^t - \bar\lambda) \qquad \text{[(4.14)]} \\
&\geq \quad \frac{\tau(\Delta - \varepsilon)}{t}\|\lambda^t - \bar\lambda\| \qquad \text{[(4.15)]} \\
&\geq \quad \tau(\Delta - \varepsilon)\|\lambda - \bar\lambda\| \qquad \text{[definition of λ^t]} \\
&\geq \quad \frac{\tau(\Delta - \varepsilon)}{\beta}\|\lambda - \bar\lambda\|^2 \qquad \text{[(4.9)]} \\
&\geq \quad \frac{1}{L_\varepsilon}\|\lambda - \bar\lambda\|^2 \qquad \text{[definition of L_ε].}
\end{aligned}
$$

This is the expected inequality (4.10). $\qquad\square$

4.2 Global linear convergence

We can now state the global linear convergence of the constraint norm towards zero in the AL algorithm of section 2. Note that the rate of convergence $\min(L/r_k, 1)$ may depend through L on the distance from the initial iterate λ_0 to the dual solution set $\mathcal{S}_D$.

Theorem 4.5 *Suppose that problem (1.1) with $Q \succcurlyeq 0$ has a solution. Consider the AL algorithm of section 2. For any $\beta > 0$, there exists an $L > 0$, such that $\mathrm{dist}(\lambda_0, \mathcal{S}_D) \leq \beta$ implies that*

$$\|y_{k+1} - Cx_{k+1}\| \leq \min\left(\frac{L}{r_k}, 1\right) \|y_k - Cx_k\|, \qquad \text{for all } k \geq 1. \qquad (4.16)$$

In particular, if $r_k \geq \bar{r}$ for all $k \geq 1$ and some $\bar{r} > L$, the constraint norm tends to zero globally linearly.

PROOF. The proof gathers known techniques (see for example [27, 28]) with the result of proposition 4.4. We give the details for completeness.

Let us note $g_{k+1} := Cx_{k+1} - y_{k+1}$. Recall from (3.10) that $g_{k+1} \in \partial\delta(\lambda_{k+1})$. Subtracting two consecutive iteration identities (2.4) provides

$$\frac{1}{r_{k+1}}(\lambda_{k+2} - \lambda_{k+1}) + (g_{k+2} - g_{k+1}) = \frac{1}{r_k}(\lambda_{k+1} - \lambda_k).$$

Taking norms, using the monotonicity of the subdifferential (which implies that $(g_{k+2} - g_{k+1})^\top(\lambda_{k+2} - \lambda_{k+1}) \geq 0$), and discarding $\|g_{k+2} - g_{k+1}\|^2 \geq 0$, we get $\|\lambda_{k+2} - \lambda_{k+1}\|^2 / r_{k+1}^2 \leq \|\lambda_{k+1} - \lambda_k\|^2 / r_k^2$ or $\|g_{k+2}\| \leq \|g_{k+1}\|$. This yields the second part of the min in (4.16).

Subtracting an arbitrary dual solution $\bar\lambda \in \mathcal{S}_D$ from both sides of the iteration identity (2.4) gives

$$\lambda_{k+1} - \bar\lambda + r_k g_{k+1} = \lambda_k - \bar\lambda, \quad \text{for } k \geq 0.$$

Taking norms and using the monotonicity of the subdifferential lead to

$$\|\lambda_{k+1} - \bar\lambda\|^2 + r_k^2 \|g_{k+1}\|^2 \leq \|\lambda_k - \bar\lambda\|^2, \quad \text{for } k \geq 0. \tag{4.17}$$

This shows in particular that the sequence $\{\|\lambda_k - \bar\lambda\|\}_{k \geq 0}$ is nonincreasing. Since $\bar\lambda$ is arbitrary in $\mathcal{S}_D$, there holds

$$\text{dist}(\lambda_{k+1}, \mathcal{S}_D) \leq \|\lambda_{k+1} - P_{\mathcal{S}_D}(\lambda_k)\| \leq \|\lambda_k - P_{\mathcal{S}_D}(\lambda_k)\| = \text{dist}(\lambda_k, \mathcal{S}_D).$$

Therefore, $\{\text{dist}(\lambda_k, \mathcal{S}_D)\}_{k \geq 0}$ is also nonincreasing, so that $\lambda_k \in \mathcal{S}_D + \beta B$ for all $k \geq 0$. Now, let $L > 0$ be the constant that proposition 4.4 associates with $\mathcal{B} := \beta B$. By this proposition, $\|\lambda_k - P_{\mathcal{S}_D}(\lambda_k)\| \leq L\|g_k\|$. Discarding the first term in the left hand side of (4.17) and using $P_{\mathcal{S}_D}(\lambda_k)$ for $\bar\lambda$, we get $\|g_{k+1}\| \leq (L/r_k)\|g_k\|$. This yields the first part of the min in (4.16). $\qquad\square$

5 Numerical experiments and discussion

The aim of this section is to illustrate by numerical experiments the global linear convergence property of the AL algorithm studied in this paper and to assess the quality of the bound given by theorem 4.5. The numerical experiments are taken from seismic reflection tomography applications. We conclude with a discussion on algorithmic implications.

5.1 A seismic reflection tomography problem

Seismic reflection tomography is a technique used to recover the geological structure of the subsoil from the measurements of the travel-times of seismic waves (see [10] for a description of the approach). From an optimization viewpoint, the problem consists in minimizing a nonlinear least-squares function subject to nonlinear constraints. In [5], a Gauss-Newton SQP method globalized by line-search is proposed and analyzed. At each iteration, a solution to a strictly convex quadratic model of the objective function subject to linear constraints is computed using our code QPAL [4, 6].

We have chosen here to present the results obtained with the problem KARINE, which is representative of those observed with our collection of 2D and 3D seismic reflection problems. These have a number of variables up to $15\ 10^3$ and a number of constraints up to 10^4. The features of the selected problem are summarized in table 1.

n	m	m^*_{act}	κ_2
442	320	108	8.4 10^5

Table 1: Description of the tomography problem KARINE

It is a 2D model depending on $n = 442$ parameters and having $m = 320$ linear inequality constraints. The matrix Q of the selected quadratic subproblem (1.1) is positive definite and has its ℓ_2 condition number equal to $\kappa_2 = 8.4 \ 10^5$. Its constraint matrix C has been balanced (the Euclidean norm of each of its rows is equal to 1). The number of active constraints at the solution is $m^*_{act} = 108$, which represents 34% of the number of constraints.

The results presented in section 5.2 have been obtained using the AL algorithm described in section 2, with a fixed augmentation parameter r. In order to study the dependence of the results on r, we have run the QP solver for 21 different values of r, ranging from 1 to 10^5. In each case, the AL algorithm is initialized with a null Lagrangian multiplier ($\lambda_0 = 0$) and is stopped when the constraint norm is sufficiently small ($\|y_k - Cx_k\| \leq 10^{-10}$).

5.2 Assessing the global linear convergence result

In this section, we illustrate the global linear convergence property of the AL algorithm established in theorem 4.5. The actual global rate of linear convergence is given by $\rho := \sup\{\|y_{k+1} - Cx_{k+1}\| / \|y_k - Cx_k\| : k \geq 1\}$ and can be estimated during a particular run by

$$\rho_{\text{est}} := \max_{1 \leq k \leq n_{AL}} \frac{\|y_{k+1} - Cx_{k+1}\|}{\|y_k - Cx_k\|} \leq \rho, \tag{5.1}$$

where n_{AL} is the number of AL iterations actually performed to reach the required accuracy of 10^{-10} on the constraint norm.

Theorem 4.5 has shown that ρ is bounded above by a function of r:

$$\rho \leq \min\left(\frac{L}{r}, 1\right) \qquad \text{or} \qquad \log \rho \leq \min(\log L - \log r, 0). \tag{5.2}$$

A natural question is to know whether this bound is tight in practice. This is difficult to say, since the value of L is generally unknown, but the appearance of ρ_{est} as a function of r in the considered problem may give a clue on this question.

The plain curve in figure 1 gives $\log \rho_{\text{est}}$ as a function of $\log r$ (double logarithmic scale). As predicted by the theory, we see that $\rho_{\text{est}} \leq 1$ for all positive r. Furthermore, the larger is the augmentation parameter r, the faster is the convergence: $\rho \simeq 1$ for $r \leq 10$ (convergence is hardly detectable) and $\rho \simeq 3.10^{-3}$ for $r = 10^5$ (convergence is obtained in very few AL iterations). We have represented by a dotted line the tangent to the ρ_{est} curve with a slope -1. This line crosses the top horizontal line of the graph at the horizontal coordinate $L_{\text{inf}} \simeq 304$. According to (5.2), L_{inf} provides a lower estimate of the value of L. Since both curves (the plain and dotted ones) are quite
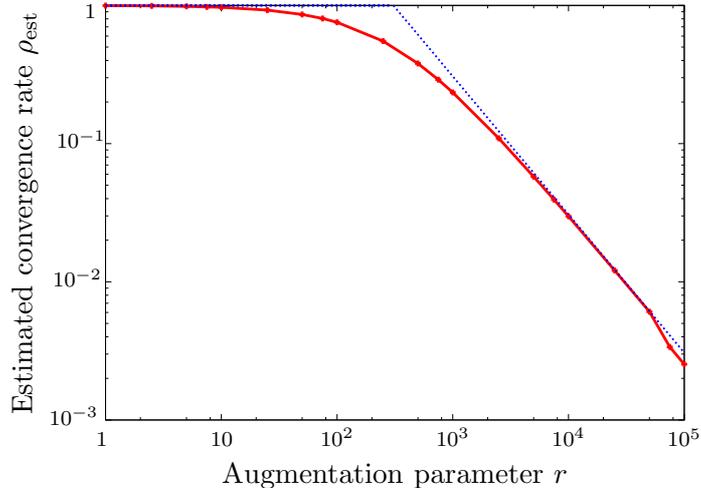
Figure 1: Global linear convergence rate of the constraint norm as a function of r

close, it is likely that the dotted curve is close to the upper bound on ρ given by (5.2). As a result, it is likely that the upper bound given by theorem 4.5 is tight. Note that the small discrepancy between both curves for large values of r ($r \geq 7.5 \, 10^4$) is due here to the fact that the AL algorithm reaches the required constraint norm accuracy in very few AL iterations ($n_{AL} \leq 3$), so that the maximum in (5.1) is taken on that few number. In other cases, such a discrepancy can come from an inexact solve of the bound constraint problem (2.3), due to a large value of r.

5.3 Discussion

As shown in this paper, the global linear rate of convergence of the AL algorithm depends on the Lipschitz constant L given by (4.8) and on the value of $\bar{r} := \inf r_k$, where r_k is the value of the augmentation parameter at iteration k. More precisely, the decrease of the constraint norm at iteration k is bounded above by L/r_k. It is usually impossible to compute L in practice, since it depends on the constants γ, σ, and Δ (see lemma 3.4, (4.4), (4.5), and finally (4.8)), which are either unknown or too expensive to compute. As a Lipschitz constant, however, L has easily computable *lower* estimates.

The estimate L_{inf} of L given in section 5.2 is not available at run time, since it requires to run the AL algorithm on a particular problem for various values of r. Nevertheless, the quantities

$$L_{\mathrm{inf},k} := \max_{1 \leq i \leq k} \left(r_i \frac{\|y_{i+1} - Cx_{i+1}\|}{\|y_i - Cx_i\|} \right)$$

satisfy $L_{\mathrm{inf},k} \leq L$ and can therefore be used as a lower estimate of L, after iteration k is completed. A given desired rate of convergence $\rho_{\mathrm{des}} \in \,]0, 1[$ is then likely to be obtained at iteration $k + 1$ by taking

$$r_{k+1} \geq \frac{L_{\mathrm{inf},k}}{\rho_{\mathrm{des}}}. \tag{5.3}$$

It is the fact that the estimate (4.16) has a *global* validity that gives sense to an update of the value of r_k in this way at *each* iteration. It should be clear at this point that the AL algorithm gains in efficiency by taking r_k as large as possible, the only limitation being that problem (2.3) needs to be numerically solvable. Since it is sometimes difficult to tell what is a large value for a particular problem, the lower bound on r_{k+1} in (5.3) may also be useful as a reference.

We conclude with a result providing an estimate of the number of iterations needed to reach a given tolerance on the constraint norm. Assume that a number $\rho_{\text{des}} \in \,]0,1[$ is given as a desired rate of convergence. Of course, since the Lipschitz constant L is unknown, this rate of convergence cannot be ensured, but the algorithm can try to approach it by updating r_k when it feels it is necessary. The next result gives an estimate of the iterative complexity of the AL algorithm with an update rule based on (5.3). More precisely, defining

$$\rho_k := \frac{\|y_{k+1} - Cx_{k+1}\|}{\|y_k - Cx_k\|},$$

the AL algorithm is supposed to update the value of r_k, for $k \geq 1$, according to:

$$\text{if } \rho_k \leq \rho_{\text{des}}, \quad \text{then } r_{k+1} = r_k, \quad \text{else } r_{k+1} = \tfrac{\rho_k}{\rho_{\text{des}}}\, r_k. \tag{5.4}$$

There is nothing magic in the update rule of r_k above. It could equally use $r_{k+1} = 10\,\rho_k r_k/\rho_{\text{des}}$ or simply $r_{k+1} = 10\,r_k$ when r_k needs to be increased.

Proposition 5.1 *Suppose that the AL algorithm of section 2 uses the rule* (5.4) *to update the augmentation parameter r_k, for $k \geq 1$. Let $\varepsilon \in \,]0,1]$ and let L be the positive constant given by theorem 4.5. Fix any $t \in \,]\rho_{\text{des}}, 1[$. Then*

$$\|y_{k+1} - Cx_{k+1}\| \leq \varepsilon \|y_1 - Cx_1\|, \tag{5.5}$$

as soon as

$$k \geq \frac{\log \varepsilon}{\log t} + \max\left(1 + \frac{\log(L/(tr_1))}{\log(t/\rho_{\text{des}})}, 0\right). \tag{5.6}$$

PROOF. Let $t \in \,]\rho_{\text{des}}, 1[$. Clearly, since $\rho_i \leq 1$,

$$\frac{\|y_{k+1} - Cx_{k+1}\|}{\|y_1 - Cx_1\|} = \prod_{1 \leq i \leq k} \rho_i \leq \prod_{\substack{1 \leq i \leq k \\ \rho_i \leq t}} \rho_i \leq t^{k_t},$$

where $k_t := |K_t|$ is the number of elements in $K_t := \{i \in \mathbb{N} : 1 \leq i \leq k, \ \rho_i \leq t\}$. Taking logarithms, we see that (5.5) holds as soon as $k_t \geq (\log \varepsilon)/(\log t)$.

If $K_t^c := \{1, \ldots, k\} \backslash K_t$ is empty, then $k = k_t$ and the result is proven.

Suppose now that $K_t^c \neq \emptyset$. Since $\rho_i \leq L/r_i$ (by theorem 4.5), $i \in K_t$ as soon as $r_i \geq L/t$. Let j be the last index in K_t^c, namely the $(k-k_t)$th one, if any. Then r_j is the result of $k - k_t - 1$ updates from r_1, using factors ρ_i/ρ_{des} that are $\geq t/\rho_{\text{des}}$ (see the

24

update rule (5.4)). Hence we must have $(t/\rho_{\text{des}})^{k-k_t-1} r_1 \leq r_j \leq L/t$. This gives an upper bound on the number of elements of K_t^c, namely

$$k - k_t \leq 1 + \frac{\log(L/(tr_1))}{\log(t/\rho_{\text{des}})}.$$

The total number of iterations to satisfy (5.5) is therefore at most this upper bound on the number of elements in K_t^c, plus the lower bound on k_t obtained above. $\qquad\square$

Roughly expressed, the number of iterations needed to reach precision $\varepsilon > 0$ on the relative constraint norm is of order $O(\log \varepsilon) + O(\log L)$. As shown in the proof of proposition 5.1, the first term of order $O(\log \varepsilon)$ is due to the linear convergence of the constraint norm towards zero, which is triggered when the augmentation parameter is large enough (a consequence of theorem 4.5). The second term of order $O(\log L)$, which is the only place where the dimension of the problem can intervene, is due to a possible too small value of r_1 and to the number of iterations that the rule (5.4) needs to make r_k large enough. This term can be made as small as desired by choosing a large value for r_1 or by adopting an update rule of r_k that increases these values more rapidly than in (5.4). As a result, the *computational* complexity of the AL algorithm of section 2 essentially rests on the one of the AL subproblems (2.3). When strict complementarity holds, the finite identification of the active constraints in (2.3) occurs and the computational complexity is then basically induced by the very first AL subproblems. Our experience with the AL algorithm, limited to the seismic reflection tomography problems described in section 5.1, supports that conclusion.

References

[1] K.J. Arrow, R.M. Solow (1958). Gradient methods for constrained maxima with weakened assumptions. In K.J. Arrow, L. Hurwicz, H. Uzawa (editors), *Studies in Linear and Nonlinear Programming*, pages 166–176. Stanford University Press, Standford, Calif.

[2] J.F. Bonnans, J.Ch. Gilbert, C. Lemaréchal, C. Sagastizábal (2003). *Numerical Optimization – Theoretical and Practical Aspects*. Springer Verlag, Berlin

[3] J.V. Burke, M.C. Ferris (1991). Characterization of solution sets of convex programs. *Operations Research Letters*, 10, 57–60

[4] F. Delbos (2004). *Problèmes d'Optimisation de Grande Taille avec Contraintes en Tomographie de Réflexion 3D*. PhD Thesis, University Pierre et Marie Curie (Paris VI), Paris. (to appear)

[5] F. Delbos, J.Ch. Gilbert, R. Glowinski, D. Sinoquet (2003). Constrained optimization in reflection tomography: an augmented Lagrangian SQP approach. Technical report, Institut Français du Pétrole. (to appear)

[6] F. Delbos, J.Ch. Gilbert, D. Sinoquet (2003). QPAL: a solver of large-scale convex quadratic optimization problems using an augmented Lagrangian approach. Technical report, INRIA, BP 105, 78153 Le Chesnay, France. (to appear)

[7] D. den Hertog (1992). *Interior Point Approach to Linear, Quadratic and Convex Programming*. Mathematics and its Applications 277. Kluwer Academic Publishers, Dordrecht

[8] M. Fortin, R. Glowinski (1982). *Méthodes de Lagrangien Augmenté – Applications à la Résolution Numérique de Problèmes aux Limites.* Méthodes Mathématiques de l'Informatique 9. Dunod, Paris

[9] A. Friedlander, J.M. Martínez (1994). On the maximization of a concave quadratic function with box constraints. *SIAM Journal on Optimization*, 4, 177–192

[10] R. Glowinski, Q.-H. Tran (1993). Constrained optimization in reflexion tomography: the augmented Lagrangian method. *East-West J. Numer. Math.*, 1(3), 213–234

[11] M.R. Hestenes (1969). Multiplier and gradient methods. *Journal of Optimization Theory and Applications*, 4, 303–320

[12] J.-B. Hiriart-Urruty, C. Lemaréchal (1993). *Convex Analysis and Minimization Algorithms.* Grundlehren der mathematischen Wissenschaften 305-306. Springer-Verlag

[13] A.F. Izmailov, M.V. Solodov (2001). Error bounds for 2-regular mappings with Lipschitzian derivatives and their applications. *Mathematical Programming*, 89, 413–435

[14] B. Jansen (1997). *Interior Point Techniques in Optimization – Complementarity, Sensitivity and Algorithms.* Applied Optimization 6. Kluwer Academic Publishers, Dordrecht

[15] M. Kojima, N. Megiddo, T. Noma, A. Yoshise (1991). *A Unified Approach to Interior Point Algorithms for Linear Complementarity Problems.* Lecture Notes in Computer Science 538. Springer-Verlag, Berlin

[16] O.L. Mangasarian (1988). A simple characterization of solution sets of convex programs. *Operations Research Letters*, 7, 21–26

[17] J.J. Moré, G. Toraldo (1991). On the solution of large quadratic programming problems with bound constraints. *SIAM Journal on Optimization*, 1, 93–113

[18] Y.E. Nesterov, A.S. Nemirovskii (1994). *Interior-Point Polynomial Algorithms in Convex Programming.* SIAM Studies in Applied Mathematics 13. SIAM, Philadelphia

[19] J. Nocedal, S.J. Wright (1999). *Numerical Optimization.* Springer Series in Operations Research. Springer, New York

[20] J.S. Pang (1983). Methods for quadratic programming: a survey. *Computers and Chemical Engineering*, 7, 583–594

[21] J.S. Pang (1997). Error bounds in mathematical programming. *Mathematical Programming*, 79, 299–332

[22] M.J.D. Powell (1969). A method for nonlinear constraints in minimization problems. In R. Fletcher (editor), *Optimization*, pages 283–298. Academic Press, London

[23] R.T. Rockafellar (1970). *Convex Analysis.* Princeton Mathematics Ser. 28. Princeton University Press, Princeton, New Jersey

[24] R.T. Rockafellar (1971). New applications of duality in convex programming. In *Proceedings of the 4th Conference of Probability, Brasov, Romania*, pages 73–81

[25] R.T. Rockafellar (1973). A dual approach to solving nonlinear programming problems by unconstrained optimization. *Mathematical Programming*, 5, 354–373

[26] R.T. Rockafellar (1974). Augmented Lagrange multiplier functions and duality in nonconvex programming. *SIAM Journal on Control*, 12, 268–285

[27] R.T. Rockafellar (1976). Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization*, 14, 877–898

[28] R.T. Rockafellar (1976). Augmented Lagrangians and applications of the proximal point algorithm in convex programming. *Mathematics of Operations Research*, 1, 97–116

[29] C. Roos, T. Terlaky, J.-Ph. Vial (1997). *Theory and Algorithms for Linear Optimization – An Interior Point Approach.* John Wiley & Sons, Chichester

[30] T. Terlaky (editor) (1996). *Interior Point Methods of Mathematical Programming.* Kluwer Academic Press, Dordrecht

[31] S.J. Wright (1997). *Primal-Dual Interior-Point Methods.* SIAM Publication, Philadelphia