

Complexity and performance of an Augmented Lagrangian algorithm*

E. G. Birgin[†] J. M. Martínez[‡]

July 4, 2019

Abstract

Algencan is a well established safeguarded Augmented Lagrangian algorithm introduced in [R. Andreani, E. G. Birgin, J. M. Martínez and M. L. Schuverdt, On Augmented Lagrangian methods with general lower-level constraints, *SIAM Journal on Optimization* 18, pp. 1286-1309, 2008]. Complexity results that report its worst-case behavior in terms of iterations and evaluations of functions and derivatives that are necessary to obtain suitable stopping criteria are presented in this work. In addition, the computational performance of a new version of the method is presented, which shows that the updated software is a useful tool for solving large-scale constrained optimization problems.

Keywords: Nonlinear programming, Augmented Lagrangian methods, complexity, numerical experiments.

1 Introduction

Augmented Lagrangian methods have a long tradition in numerical optimization. The main ideas were introduced by Powell [43], Hestenes [39], and Rockafellar [45]. At each (outer) iteration of an Augmented Lagrangian method one minimizes the objective function plus a term that penalizes the non-fulfillment of the constraints with respect to suitable shifted tolerances. Whereas the classical external penalty method [34, 35] needs to employ penalty parameters that tend to infinity, the shifting technique aims to produce convergence by means of displacements of the constraints that generate approximations to a solution with moderate penalty parameters [20]. As a by-product, one obtains approximations of the Lagrange multipliers associated with the original optimization problem. The safeguarded version of the method [3] discards Lagrange multiplier approximations when they become very large. The convergence theory for safeguarded Augmented Lagrangian methods was given in [3, 20]. Recently, examples that illustrate the convenience of safeguarded Augmented Lagrangians were given in [41].

Conn, Gould, and Toint [27] produced the celebrated package Lancelot, that solves constrained optimization problems using Augmented Lagrangians in which the constraints are

*This work was supported by FAPESP (grants 2013/07375-0, 2016/01860-1, and 2018/24293-0) and CNPq (grants 309517/2014-1 and 303750/2014-6).

[†]Dept. of Computer Science, Institute of Mathematics and Statistics, University of São Paulo, Rua do Matão, 1010, Cidade Universitária, 05508-090, São Paulo, SP, Brazil. email: egbirgin@ime.usp.br

[‡]Dept. of Applied Mathematics, Institute of Mathematics, Statistics, and Scientific Computing, State University of Campinas, 13083-859, Campinas, SP, Brazil. email: martinez@ime.unicamp.br

defined by equalities and bounds. The technique was extended to the case of equality constraints plus linear constraints in [26]. Differently from Lancelot, in Algencan [3, 20] (see, also, [4, 5, 14, 15, 17, 18, 19]), the Augmented Lagrangian is defined not only with respect to equality constraints but also with respect to inequalities. The theory presented in [3] and [20] admits the presence of lower-level constraints not restricted to boxes or polytopes. However, in the practical implementations of Algencan, lower-level constraints are always boxes.

In the last 10 years, the interest in Augmented Lagrangian methods was renewed due to their ability to solve large-scale problems. Dostál and Beremlijski [31, 32] employed Augmented Lagrangian methods for solving quadratic programming problems that appear in structural optimization. Fletcher [36] applied Augmented Lagrangian ideas to the minimization of quadratics with box constraints. Armand and Omheni [12] employed an Augmented Lagrangian technique for solving equality constrained optimization problems and handled inequality constraints by means of logarithmic barriers [13]. Curtis, Gould, Jiang, and Robinson [28, 29] defined an Augmented Lagrangian algorithm in which decreasing the penalty parameters is possible following intrinsic algorithmic criteria. Local convergence results without constraint qualifications were proved in [33]. The case with (possibly complementarity) degenerate constraints was analyzed in [40]. Chatzipanagiotis and Zavlanos [25] defined and analyzed Augmented Lagrangian methods in the context of distributed computation. An Exact Penalty algorithm for constrained optimization with complexity results was introduced in [24]. Grapiglia and Yuan [38] analyzed the complexity of an Augmented Lagrangian algorithm for inequality constraints based on the approach of Sun and Yuan [46] and assuming that a feasible initial point is available.

In this paper, we report the main features of a new implementation of Algencan. The new Algencan preserves the main characteristics of the previous algorithm: constraints are considered in the form of equalities and inequalities, without slack variables; box-constrained subproblems are solved using active-set strategies; and global convergence properties are fully preserved. A new acceleration procedure is introduced by means of which an approximate KKT point may be obtained. It consists in applying a local Newton method to a semismooth KKT system [42, 44] starting from every Augmented Lagrangian iterate. Special attention is given to the box-constraint algorithm used for solving subproblems. The algorithm presented in this paper is able to handle large-scale problems but not “huge” ones. This means that we deal with number of variables and Hessian structures that make it affordable to use sparse factorizations. Larger problems need the help of iterative linear solvers which are not available in the new Algencan yet. Exhaustive numerical experimentation is given and all the software employed is available on a free basis in <http://www.ime.usp.br/~egbirgin/>, so that computational results are fully reproducible.

The paper is organized as follows. In Section 2, we recall the definition of Algencan with box lower-level constraints and we review global convergence results. In Section 3, we prove complexity properties. In Section 4, we describe the algorithm for solving box-constrained subproblems. In Section 5, we describe the computer implementation. In Section 6, we report numerical experiments. Conclusions are given in Section 7.

Notation. If $C \subseteq \mathbb{R}^n$ is a convex set, $P_C(v)$ denotes the Euclidean projection of v onto C . If $\ell, u \in \mathbb{R}^n$, $[\ell, u]$ denotes the box defined by $\{x \in \mathbb{R}^n \mid \ell \leq x \leq u\}$. $(\cdot)_+ = \max\{0, \cdot\}$. If $v \in \mathbb{R}^n$, v_+ denotes the vector with components $(v_i)_+$ for $i = 1, \dots, n$. If $v, w \in \mathbb{R}^n$, $\min\{v, w\}$ denotes the vector with components $\min\{v_i, w_i\}$ for $i = 1, \dots, n$. The symbol $\|\cdot\|$ denotes the Euclidean norm.

2 Augmented Lagrangian

In this section, we consider constrained optimization problems defined by

$$\text{Minimize } f(x) \text{ subject to } h(x) = 0, g(x) \leq 0, \text{ and } \ell \leq x \leq u, \quad (1)$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$, $h : \mathbb{R}^n \rightarrow \mathbb{R}^m$, and $g : \mathbb{R}^n \rightarrow \mathbb{R}^p$ are continuously differentiable.

We consider the Augmented Lagrangian method in the way analyzed in [3] and [20]. This method has interesting global theoretical properties. On the one hand, every limit point is a stationary point of the problem of minimizing infeasibility. On the other hand, every feasible limit point satisfies a sequential optimality condition [7, 8, 9]. This implies that every feasible limit point is KKT-stationary under very mild constraint qualifications [8, 9]. The basic definition of the method and the main theoretical results are reviewed in this section.

The Augmented Lagrangian function [39, 43, 45] associated with problem (1) is defined by

$$L_\rho(x, \lambda, \mu) = f(x) + \frac{\rho}{2} \left[\sum_{i=1}^m \left(h_i(x) + \frac{\lambda_i}{\rho} \right)^2 + \sum_{i=1}^p \left(g_i(x) + \frac{\mu_i}{\rho} \right)_+^2 \right]$$

for all $x \in [\ell, u]$, $\rho > 0$, $\lambda \in \mathbb{R}^m$, and $\mu \in \mathbb{R}_+^p$. The Augmented Lagrangian model algorithm follows.

Algorithm 2.1: Assume that $x^0 \in \mathbb{R}^n$, $\lambda_{\min} < \lambda_{\max}$, $\bar{\lambda}^1 \in [\lambda_{\min}, \lambda_{\max}]^m$, $\mu_{\max} > 0$, $\bar{\mu}^1 \in [0, \mu_{\max}]^p$, $\rho_1 > 0$, $\gamma > 1$, $0 < \tau < 1$, and $\{\varepsilon_k\}_{k=1}^\infty$ are given. Initialize $k \leftarrow 1$.

Step 1. Find $x^k \in [\ell, u]$ as an approximate solution to

$$\text{Minimize } L_{\rho_k}(x, \bar{\lambda}^k, \bar{\mu}^k) \text{ subject to } \ell \leq x \leq u \quad (2)$$

satisfying

$$\left\| P_{[\ell, u]} \left(x^k - \nabla L_{\rho_k}(x^k, \bar{\lambda}^k, \bar{\mu}^k) \right) - x^k \right\| \leq \varepsilon_k. \quad (3)$$

Step 2. Define

$$V^k = \min \left\{ -g(x^k), \frac{\bar{\mu}^k}{\rho_k} \right\}.$$

If $k = 1$ or

$$\max \left\{ \|h(x^k)\|, \|V^k\| \right\} \leq \tau \max \left\{ \|h(x^{k-1})\|, \|V^{k-1}\| \right\}, \quad (4)$$

choose $\rho_{k+1} = \rho_k$. Otherwise, define $\rho_{k+1} = \gamma \rho_k$.

Step 3. Compute

$$\lambda^{k+1} = \bar{\lambda}^k + \rho_k h(x^k) \text{ and } \mu^{k+1} = \left(\bar{\mu}^k + \rho_k g(x^k) \right)_+. \quad (5)$$

Compute $\bar{\lambda}^{k+1} \in [\lambda_{\min}, \lambda_{\max}]^m$ and $\bar{\mu}_i^{k+1} \in [0, \mu_{\max}]^p$. Set $k \leftarrow k + 1$ and go to Step 1.

The problem of finding an approximate minimizer of $L_{\rho_k}(x, \bar{\lambda}^k, \bar{\mu}^k)$ onto $[\ell, u]$ in the sense of (3) can always be solved. In fact, due to the compactness of $[\ell, u]$, a global minimizer, that obviously satisfies (3), always exists. Moreover, local minimization algorithms are able to find an approximate stationary point satisfying (3) in a finite number of iterations. Therefore, given

an iterate x^k , the iterate x^{k+1} is well defined. So, Algorithm 2.1 generates an infinite sequence $\{x^k\}$ whose properties are surveyed below. Of course, as it will be seen later, suitable stopping criteria can be defined by means of which acceptable approximate solutions to (1) are usually obtained.

Algorithm 2.1 has been presented without a “stopping criterion”. This means that, in principle, the algorithm generates an infinite sequence of primal iterates x^k and Lagrange-multiplier estimators. Complexity results presented in this work report the worst-case effort that could be necessary to obtain different properties, that may be used as stopping criteria in practical implementations or not. We believe that the interpretation of these results helps to decide which stopping criteria should be used in a practical application.

The relevant theoretical properties of this algorithm are the following:

1. Every limit point x^* of the sequence generated by the algorithm satisfies the complementarity condition

$$\mu_i^{k+1} = 0 \text{ whenever } g_i(x^*) < 0 \quad (6)$$

for k large enough. (See [20, Thm.4.1].)

2. Every limit point x^* of the sequence generated by the algorithm satisfies the first-order optimality conditions of the feasibility problem

$$\text{Minimize } \|h(x)\|^2 + \|g(x)_+\|^2 \text{ subject to } \ell \leq x \leq u. \quad (7)$$

(See [20, Thm.6.5].)

3. If, for all $k \in \{1, 2, \dots\}$, x^k is an approximate global minimizer of $L_{\rho_k}(x, \bar{\lambda}^k, \bar{\mu}^k)$ onto $[\ell, u]$ with tolerance $\eta > 0$, every limit point of $\{x^k\}$ is a global minimizer of the infeasibility function $\|h(x)\|^2 + \|g(x)_+\|^2$. Condition (3) does not need to hold in this case. (See [20, Thm.5.1].)
4. If, for all $k \in \{1, 2, \dots\}$, x^k is an approximate global minimizer of $L_{\rho_k}(x, \bar{\lambda}^k, \bar{\mu}^k)$ onto $[\ell, u]$ with tolerance $\eta_k \downarrow 0$, every feasible limit point of $\{x^k\}$ is a global minimizer of the general constrained minimization problem (1). As before, condition (3) is not necessary in this case. (See [20, Thm.5.2].)
5. If $\varepsilon_k \downarrow 0$, every feasible limit point of the sequence $\{x^k\}$ satisfies the sequential optimality condition AKKT [7] given by

$$\lim_{k \in K} \left\| P_{[\ell, u]} \left(x^k - \left(\nabla f(x^k) + \nabla h(x^k) \lambda^{k+1} + \nabla g(x^k) \mu^{k+1} \right) \right) - x^k \right\| = 0 \quad (8)$$

and

$$\lim_{k \in K} \max \{ \|h(x^k)\|_\infty, \|\min\{-g(x^k), \mu^{k+1}\}\|_\infty \} = 0, \quad (9)$$

where the sequence of indices K is such that $\lim_{k \in K} x^k = x^*$. (See [20, Thm.6.4].)

Under an additional Lojasiewicz-like condition, it is obtained that $\lim_{k \in K} \sum_{i=1}^p \mu_i^{k+1} g_i(x^k) = 0$ (see [10]). Moreover, in [6], it was proved that an even stronger sequential optimality condition is satisfied by the sequence $\{x^k\}$, perhaps associated with different Lagrange multipliers approximations than the ones generated by the Augmented Lagrangian algorithm.

These properties say that, even if ε_k does not tend to zero, Algorithm 2.1 finds stationary points of the infeasibility measure $\|h(x)\|^2 + \|g(x)_+\|^2$ and that, when ε_k tends to zero, feasible limit points satisfy a sequential optimality condition. Thus, under very weak constraint qualifications, feasible limit points satisfy Karush-Kuhn-Tucker conditions. See [8, 9]. Some of these properties, but not all, are shared by other constrained optimization algorithms. For example, the property that feasible limit points satisfy optimality KKT conditions is proved to be satisfied by other optimization algorithms only under much stronger constraint qualifications than the ones required by Algorithm 2.1. Moreover, the Newton-Lagrange method may fail to satisfy approximate KKT conditions even when it converges to the solution of rather simple constrained optimization problems [1, 2].

Augmented Lagrangian implementations have a modular structure. At each iteration, a box-constrained optimization problem is approximately solved. The efficiency of the Augmented Lagrangian algorithm is strongly linked to the efficiency of the box-constraint solver.

Algencan may be considered to be a conservative visit to the Augmented Lagrangian framework. For example, subproblems are solved with relatively high precision, instead of stopping subproblem solvers prematurely according to information related to the constrained optimization landscape. It could be argued that solving subproblems with high precision at points that may be far from the solution represents a waste of time. Nevertheless, our point of view is that saving subproblem iterations when one is close to a subproblem solution is not worthwhile because in that region Newton-like solvers tend to be very fast; and accurate subproblems' solutions help to produce better approximations of Lagrange multipliers. Algencan is also conservative when subproblems' solvers use minimal information about the structure of the Augmented Lagrangian function they minimize. The reason for this decision is connected to the modular structure of Algencan. Subproblem solvers are continuously being improved due to the continuous and fruitful activity in bound-constrained minimization. Therefore, we aim to take advantage of those improvements with minimal modifications of subproblem algorithms when applied to minimize Augmented Lagrangians.

3 Complexity

This section is devoted to worst-case complexity results related to Algorithm 2.1. Algorithm 2.1 was not devised with the aim of optimizing complexity. Nevertheless, our point of view is that the complexity analysis that follows helps to understand the actual behavior of the algorithm, filling a gap opened by the convergence theory.

By (5) and straightforward calculations, we have that, for all $k = 1, 2, 3, \dots$,

$$\nabla f(x^k) + \nabla h(x^k)\lambda^{k+1} + \nabla g(x^k)\mu^{k+1} = \nabla L_{\rho_k}(x^k, \bar{\lambda}^k, \bar{\mu}^k).$$

Therefore, the fulfillment of

$$\|P_{[\ell, u]}(x^k - \nabla L_{\rho_k}(x^k, \bar{\lambda}^k, \bar{\mu}^k)) - x^k\| \leq \varepsilon \tag{10}$$

implies that the projected gradient of the Lagrangian with multipliers λ^{k+1} and μ^{k+1} approximately vanishes with precision ε . In the next lemma, we show that the fulfillment of

$$\max\{\|h(x^k)\|_\infty, \|V_k\|_\infty\} \leq \delta \tag{11}$$

implies that feasibility and complementarity hold at x^k with precision δ . For these reasons, in the context of Algorithm 2.1, iterates that satisfy (10) and (11) are considered approximate stationary points of problem (1).

Lemma 3.1 *For all $\delta > 0$,*

$$\max\{\|h(x^k)\|_\infty, \|V_k\|_\infty\} \leq \delta \quad (12)$$

implies that

$$\|h(x^k)\|_\infty \leq \delta, \|g(x^k)_+\|_\infty \leq \delta, \text{ and, for all } j = 1, \dots, p, \mu_j^{k+1} = 0 \text{ if } g_j(x^k) < -\delta. \quad (13)$$

Proof: By (12), $\|h(x^k)\|_\infty \leq \delta$ and $|\min\{-g_j(x^k), \bar{\mu}_j^k/\rho_k\}| \leq \delta$ for all $j = 1, \dots, p$. Therefore, $-g_j(x^k) \geq -\delta$, so $g_j(x^k) \leq \delta$ for all $j = 1, \dots, p$. Moreover, by (12), if $g_j(x^k) < -\delta$, we necessarily have that $\bar{\mu}_j^k/\rho_k \leq \delta$. Adding these two inequalities, we obtain that, if $g_j(x^k) < -\delta$ then $g_j(x^k) + \bar{\mu}_j^k/\rho_k < 0$. Consequently, $\rho_k g_j(x^k) + \bar{\mu}_j^k < 0$, so $\mu_j^{k+1} = 0$. Therefore, (12) implies (13) as we wanted to prove. \square

In Theorem 3.1 below, we assume that the sequence $\{\rho_k\}$ is bounded. Sufficient conditions for this requirement, where the bound $\bar{\rho}$ only depends on algorithmic parameters and characteristics of the problem, were given in [3] and [20]. We also assume that there exists $N(\varepsilon) \in \{1, 2, 3, \dots\}$ such that $\varepsilon_k \leq \varepsilon$ for all $k \geq N(\varepsilon)$. Clearly, this condition can be enforced by the criterion used to define $\{\varepsilon_k\}$. For example, $\varepsilon_{k+1} = \frac{1}{2}\varepsilon_k$ obviously implies that $\varepsilon_k \leq \varepsilon$ if $k > N(\varepsilon) \equiv \log(\varepsilon)/\log(\varepsilon_1)$.

Lemma 3.2 *There exists $c_{\text{big}} > 0$ such that, for all $k \geq 1$,*

$$\max\{\|h(x^k)\|_\infty, \|V_k\|_\infty\} \leq c_{\text{big}}. \quad (14)$$

Proof: Since, by definition of the algorithm, $\rho_k \geq \rho_1$, the bound (14) comes from the continuity of h and g , the compactness of the domain $[\ell, u]$, and the boundedness of $\bar{\mu}^k$. \square

From now on, c_{big} will denote a positive constant satisfying (14), whose existence is guaranteed by Lemma 3.2.

Theorem 3.1 *Let $\delta > 0$ and $\varepsilon > 0$ be given. Assume that, for all $k \in \{1, 2, 3, \dots\}$, $\rho_k \leq \bar{\rho}$. Moreover, assume that, for all $k \geq N(\varepsilon)$, we have that $\varepsilon_k \leq \varepsilon$. Then, after at most*

$$N(\varepsilon) + [\log(\bar{\rho}/\rho_1)/\log(\gamma)] \times [\log(\delta/c_{\text{big}})/\log(\tau)] \quad (15)$$

iterations, we obtain $x^k \in [\ell, u]$, $\lambda^{k+1} \in \mathbb{R}^m$, and $\mu^{k+1} \in \mathbb{R}_+^p$ such that

$$\left\| P_{[\ell, u]} \left(x^k - \left(\nabla f(x^k) + \nabla h(x^k) \lambda^{k+1} + \nabla g(x^k) \mu^{k+1} \right) \right) - x^k \right\| \leq \varepsilon, \quad (16)$$

$$\|h(x^k)\|_\infty \leq \delta, \|g(x^k)_+\|_\infty \leq \delta, \quad (17)$$

and, for all $j = 1, \dots, p$,

$$\mu_j^{k+1} = 0 \text{ whenever } g_j(x^k) < -\delta. \quad (18)$$

Proof: The number of iterations such that $\rho_{k+1} = \gamma\rho_k$ is bounded above by

$$\log(\bar{\rho}/\rho_1)/\log(\gamma). \quad (19)$$

Therefore, this is also a bound for the number of iterations at which (4) does not hold.

By (14), if (4) holds during

$$\log(\delta/c_{\text{big}})/\log \tau \quad (20)$$

consecutive iterations, we get that

$$\max\{\|h(x^k)\|_\infty, \|V_k\|_\infty\} \leq \delta,$$

which, by Lemma 3.1, implies (17) and (18).

Now, by hypothesis, after $N(\varepsilon)$ iterations, we have that $\varepsilon_k \leq \varepsilon$. Therefore, by (19) and (20), after at most

$$N(\varepsilon) + [\log(\bar{\rho}/\rho_1)/\log(\gamma)] \times [\log(\delta/c_{\text{big}})/\log(\tau)] \quad (21)$$

iterations, we have that (16), (17), and (18) hold. \square

Theorem 3.1 shows that, as expected, if ρ_k is bounded, we obtain approximate feasibility and optimality. In the following theorem, we assume that the subproblems are solved by means of some method that, for obtaining precision $\varepsilon > 0$, employs at most $c\varepsilon^{-q}$ iterations and evaluations, where c only depends on characteristics of the problem, the upper bound for ρ_k , and algorithmic parameters of the method.

Theorem 3.2 *In addition to the hypotheses of Theorem 3.1, assume that there exist $c_{\text{inner}} > 0$ and $q > 0$, where c_{inner} only depends on $\bar{\rho}$, λ_{\min} , λ_{\max} , μ_{\max} , ℓ , u , and characteristics of the functions f , h , and g , such that the number of inner iterations, function and derivative evaluations that are necessary to obtain (3) is bounded above by $c_{\text{inner}}\varepsilon_k^{-q}$. Then, the number of inner iterations, function evaluations, and derivative evaluations that are necessary to obtain k such that (16), (17), and (18) hold is bounded above by*

$$c_{\text{inner}}\varepsilon_{\min}^{-q} \{N(\varepsilon) + [\log(\bar{\rho}/\log \rho_1)/\log(\gamma)] \times [\log(\delta/c_{\text{big}})/\log(\tau)]\},$$

where

$$\varepsilon_{\min} = \min\{\varepsilon_k \mid k \leq N(\varepsilon) + [\log(\bar{\rho}/\log \rho_1)/\log(\gamma)] \times [\log(\delta/c_{\text{big}})/\log(\tau)]\}. \quad (22)$$

Proof: The desired result follows from Theorem 3.1 and the assumptions of this theorem. \square

Note that, in Theorem 3.2, we admit the possibility that ε_k decrease after completing $N(\varepsilon)$ iterations. This is the reason for the definition of ε_{\min} (22). In practical implementations, it is reasonable to stop decreasing ε_k when it achieves a user-given stopping tolerance ε . According to Theorem 3.2, the complexity bounds related to approximate optimality, feasibility, and complementarity depend on the optimality tolerance ε in, essentially, the same way that the complexity of the subproblem solver depends on its stopping tolerance. In other words, under the assumption of boundedness of penalty parameters, the worst-case complexity of the Augmented Lagrangian method is essentially the same as the complexity of the subproblem solver.

In computer implementations, it is usual to employ, in addition to a (successful) stopping criterion based on (16), (17), and (18), an (unsuccessful) stopping criterion based on the size of the penalty parameter. The rationale is that if the penalty parameter grew to be very large, it is not worthwhile to expect further improvements with respect to feasibility and we are probably close to an infeasible local minimizer of infeasibility. The complexity results that correspond to this decision are given below.

Theorem 3.3 *Let $\delta > 0$, $\varepsilon > 0$, and $\rho_{\text{big}} > \rho_1$ be given. Assume that, for all $k \geq N(\varepsilon)$, we have that $\varepsilon_k \leq \varepsilon$. Then, after at most*

$$N(\varepsilon) + \lceil \log(\rho_{\text{big}}/\rho_1) / \log(\gamma) \rceil \times \lceil \log(\delta/c_{\text{big}}) / \log(\tau) \rceil \quad (23)$$

iterations, we obtain $x^k \in [\ell, u]$, $\lambda^{k+1} \in \mathbb{R}^m$, and $\mu^{k+1} \in \mathbb{R}_+^p$ such that (16), (17), and (18) hold or we obtain an iteration such that $\rho_k \geq \rho_{\text{big}}$.

Proof: If $\rho_k \leq \rho_{\text{big}}$ for all $k \leq N(\varepsilon) + \lceil \log(\rho_{\text{big}}/\rho_1) / \log(\gamma) \rceil \times \lceil \log(\delta/c_{\text{big}}) / \log(\tau) \rceil$, by the same argument used in the proof of Theorem 3.1, with ρ_{big} replacing $\bar{\rho}$, we obtain that (16), (17), and (18) hold. \square

Theorem 3.4 *In addition to the hypotheses of Theorem 3.3, assume that there exist $c_{\text{inner}} > 0$ and $q > 0$, where c_{inner} only depends on ρ_{big} , λ_{min} , λ_{max} , μ_{max} , ℓ , u , and characteristics of the functions f , h , and g , such that the number of inner iterations, function and derivative evaluations that are necessary to obtain (3) is bounded above by $c_{\text{inner}}\varepsilon_k^{-q}$. Then, the number of inner iterations, function evaluations, and derivative evaluations that are necessary to obtain k such that (16), (17), and (18) hold or such that $\rho_k > \rho_{\text{big}}$ is bounded above by*

$$c_{\text{inner}}\varepsilon_{\min,2}^{-q} \{N(\varepsilon) + \lceil \log(\rho_{\text{big}}/\rho_1) / \log(\gamma) \rceil \times \lceil \log(\delta/c_{\text{big}}) / \log(\tau) \rceil\},$$

where

$$\varepsilon_{\min,2} = \min\{\varepsilon_k \mid k \leq \{N(\varepsilon) + \lceil \log(\rho_{\text{max}}/\log \rho_1) / \log(\gamma) \rceil \times \lceil \log(\delta/c_{\text{big}}) / \log(\tau) \rceil\}\}. \quad (24)$$

Proof: The desired result follows directly from Theorem 3.3. \square

The complexity results proved up to now indicate that suitable stopping criteria for Algorithm 2.1 could be based on the fulfillment of (16), (17), and (18) or, alternatively, on the occurrence of an undesirable big penalty parameter. The advantage of these criteria is that, according to them, worst-case complexity is of the same order as the complexity of subproblem solvers. Convergence results establish that solutions obtained with very large penalty parameters are close to stationary points of the infeasibility. However, stationary points of infeasibility may be feasible points and, again, convergence theory shows that when Algorithm 2.1 converges to a feasible point, this point satisfies AKKT optimality conditions, independently of constraint qualifications. As a consequence, the danger exists of interrupting executions prematurely, in situations in which meaningful progress could be obtained admitting further increases of the penalty parameter. This state of facts leads one to analyze complexity of Algorithm 2.1 independently of penalty parameter growth and introducing a possibly more reliable criterion for

detecting infeasible stationary points of infeasibility. Roughly speaking, we will say that an iterate seems to be an infeasible stationary point of infeasibility when the projected gradient of the infeasibility measure is significantly smaller than the infeasibility value. The natural question that arises is whether the employment of this (more reliable) stopping criterion has an important effect on the complexity bounds.

Assumptions on the limitation of ρ_k are given up from now on. Note that the possibility that $\rho_k \rightarrow \infty$ needs to be considered since necessarily takes place, for example, when the feasible region is empty.

Lemma 3.3 *There exist $c_{\text{lips}}, c_f > 0$ such that, for all $x \in [\ell, u]$, $\lambda \in [\lambda_{\min}, \lambda_{\max}]^m$, and $\mu \in [0, \mu_{\max}]^p$, one has*

$$\|\nabla h(x)\| \|\lambda\| + \|\nabla g(x)\| \|\mu\| \leq c_{\text{lips}} \quad (25)$$

and

$$\|\nabla f(x)\| \leq c_f. \quad (26)$$

Proof: The desired result follows from the boundedness of the domain, the continuity of the functions, and the boundedness of λ and μ . \square

The following lemma establishes a bound for the projected gradient of the infeasibility measure in terms of the value of the displaced infeasibility and the value of the penalty parameter.

Lemma 3.4 *For all $x \in [\ell, u]$, $\lambda \in [\lambda_{\min}, \lambda_{\max}]^m$, $\mu \in [0, \mu_{\max}]^p$, and $\rho > 0$, one has that*

$$\begin{aligned} & \|P_{[\ell, u]}(x - \nabla [\|h(x)\|^2 + \|g(x)_+\|^2]) - x\| \\ & \leq \|P_{[\ell, u]}(x - \nabla [\|h(x) + \lambda/\rho\|^2 + \|(g(x) + \mu/\rho)_+\|^2]) - x\| + 2c_{\text{lips}}/\rho, \end{aligned}$$

where c_{lips} is defined in Lemma 3.3.

Proof: Note that

$$\frac{1}{2} \nabla [\|h(x) + \lambda/\rho\|^2 + \|(g(x) + \mu/\rho)_+\|^2] = h'(x)^T (h(x) + \lambda/\rho) + g'(x)^T (g(x) + \mu/\rho)_+$$

and

$$\frac{1}{2} \nabla [\|h(x)\|^2 + \|g(x)_+\|^2] = \nabla h(x)h(x) + \nabla g(x)g(x)_+.$$

Therefore,

$$\begin{aligned} & \left\| \frac{1}{2} \nabla [\|h(x) + \lambda/\rho\|^2 + \|(g(x) + \mu/\rho)_+\|^2] - \frac{1}{2} \nabla [\|h(x)\|^2 + \|g(x)_+\|^2] \right\| \\ & \leq \|\nabla h(x)\lambda/\rho + \nabla g(x) [(g(x) + \mu/\rho)_+ - g(x)_+]\| \leq \frac{1}{\rho} [\|\nabla h(x)\| \|\lambda\| + \|\nabla g(x)\| \|\mu\|]. \end{aligned}$$

Then, by (25), if $\rho > 0$, $x \in [\ell, u]$, $\lambda \in [\lambda_{\min}, \lambda_{\max}]^m$, and $\mu \in [0, \mu_{\max}]^p$,

$$\|\nabla [\|h(x)\|^2 + \|g(x)_+\|^2] - \nabla [\|h(x) + \lambda/\rho\|^2 + \|(g(x) + \mu/\rho)_+\|^2]\| \leq 2c_{\text{lips}}/\rho.$$

So, by the non-expansivity of projections,

$$\|P_{[\ell,u]}(x - \nabla [\|h(x)\|^2 + \|g(x)_+\|^2]) - P_{[\ell,u]}(x - \nabla [\|h(x) + \lambda/\rho\|^2 + \|(g(x) + \mu/\rho)_+\|^2])\| \leq 2c_{\text{lips}}/\rho.$$

Thus, the thesis is proved. \square

The following theorem establishes that, before the number of iterations given by (27), we necessarily find an approximate KKT point or we find an infeasible point that, very likely, is close to an infeasible stationary point of the infeasibility measure. The latter type of infeasible points is characterized by the fact that the projected gradient of the infeasibility is smaller than δ_{low} whereas the infeasibility value is bigger than $\delta \gg \delta_{\text{low}}$.

Theorem 3.5 *Let $\delta > 0$, $\delta_{\text{low}} \in (0, \delta)$, and $\varepsilon > 0$ be given. Assume that $N(\delta_{\text{low}}, \varepsilon)$ is such that $\varepsilon_k \leq \min\{\varepsilon, \delta_{\text{low}}\}/4$ for all $k \geq N(\delta_{\text{low}}, \varepsilon)$. Then, after at most*

$$N(\delta_{\text{low}}, \varepsilon) + \left\lceil \frac{\log(\delta/c_{\text{big}})}{\log(\tau)} \right\rceil \times \left\lceil \frac{\log(\rho_{\text{max}}/\rho_1)}{\log(\gamma)} \right\rceil \quad (27)$$

iterations, where

$$\rho_{\text{max}} = \max \left\{ 1, \frac{4c_{\text{lips}}}{\delta_{\text{low}}}, \frac{\mu_{\text{max}}}{\delta}, \frac{4c_f}{\delta_{\text{low}}} \right\}, \quad (28)$$

we obtain an iteration k such that one of the following two facts takes place:

1. The iterate $x^k \in [\ell, u]$ verifies

$$\|P_{[\ell,u]}(x^k - \nabla [\|h(x^k)\|^2 + \|g(x^k)_+\|^2]) - x^k\| \leq \delta_{\text{low}} \text{ and } \max\{\|h(x^k)\|_\infty, \|g(x^k)_+\|_\infty\} > \delta. \quad (29)$$

2. The multipliers $\lambda^{k+1} \in \mathbb{R}^m$ and $\mu^{k+1} \in \mathbb{R}_+^p$ are such that

$$\|P_{[\ell,u]}(x^k - (\nabla f(x^k) + \nabla h(x^k)\lambda^{k+1} + \nabla g(x^k)\mu^{k+1})) - x^k\| \leq \varepsilon, \quad (30)$$

$$\|h(x^k)\|_\infty \leq \delta, \quad \|g(x^k)_+\|_\infty \leq \delta, \quad (31)$$

and, for all $j = 1, \dots, p$,

$$\mu_j^{k+1} = 0 \text{ whenever } g_j(x^k) < -\delta. \quad (32)$$

Proof: Let k_{end} be such that

$$\|P_{[\ell,u]}(x^k - \nabla [\|h(x^k)\|^2 + \|g(x^k)_+\|^2]) - x^k\| \leq \delta_{\text{low}} \Rightarrow \max\{\|h(x^k)\|_\infty, \|g(x^k)_+\|_\infty\} \leq \delta \quad (33)$$

for all $k \leq k_{\text{end}}$ whereas (33) does not hold if $k = k_{\text{end}} + 1$. (With some abuse of notation, we say that $k_{\text{end}} = \infty$ when (33) holds for all k .) In other words, if $k \leq k_{\text{end}}$,

$$\|P_{[\ell,u]}(x^k - \nabla [\|h(x^k)\|^2 + \|g(x^k)_+\|^2]) - x^k\| > \delta_{\text{low}} \text{ or } \max\{\|h(x^k)\|_\infty, \|g(x^k)_+\|_\infty\} \leq \delta, \quad (34)$$

whereas (34) does not hold if $k = k_{\text{end}} + 1$.

We consider two possibilities:

$$k_{\text{end}} < N(\delta_{\text{low}}, \varepsilon) + \left\lceil \frac{\log(\delta/c_{\text{big}})}{\log(\tau)} \right\rceil \times \left\lceil \frac{\log(\rho_{\text{max}}/\rho_1)}{\log(\gamma)} \right\rceil \quad (35)$$

and

$$k_{\text{end}} \geq N(\delta_{\text{low}}, \varepsilon) + \left\lceil \frac{\log(\delta/c_{\text{big}})}{\log(\tau)} \right\rceil \times \left\lceil \frac{\log(\rho_{\text{max}}/\rho_1)}{\log(\gamma)} \right\rceil. \quad (36)$$

In the first case, since (33) does not hold for $k = k_{\text{end}} + 1$, it turns out that (29) occurs at iteration $k_{\text{end}} + 1$. It remains to analyze the case in which (36) takes place.

Suppose that

$$k \leq N(\delta_{\text{low}}, \varepsilon) + \left\lceil \frac{\log(\delta/c_{\text{big}})}{\log(\tau)} \right\rceil \times \left\lceil \frac{\log(\rho_{\text{max}}/\rho_1)}{\log(\gamma)} \right\rceil, \quad (37)$$

$$\varepsilon_k \leq \delta_{\text{low}}/4, \quad (38)$$

$$\rho_k \geq 1, \quad (39)$$

$$\rho_k \geq 4c_f/\delta_{\text{low}}, \quad (40)$$

$$\rho_k \geq 4c_{\text{lips}}/\delta_{\text{low}}, \quad (41)$$

$$\rho_k \geq \mu_{\text{max}}/\delta, \quad (42)$$

$$k \geq N(\delta_{\text{low}}, \varepsilon). \quad (43)$$

By (3), for all $k \geq 1$, we have that

$$\left\| P_{[\ell, u]} \left(x^k - \nabla f(x^k) - \frac{\rho_k}{2} \nabla \left\{ \sum_{i=1}^m \left[h_i(x^k) + \frac{\bar{\lambda}_i^k}{\rho_k} \right]^2 + \sum_{i=1}^p \left[\left(g_i(x^k) + \frac{\bar{\mu}_i^k}{\rho_k} \right)_+ \right]^2 \right\} \right) - x^k \right\| \leq \varepsilon_k.$$

Therefore, by (39),

$$\left\| P_{[\ell, u]} \left(x^k - \frac{1}{\rho_k} \nabla f(x^k) - \frac{1}{2} \nabla \left(\|h(x^k) + \bar{\lambda}^k/\rho_k\|^2 + \|(g(x^k) + \bar{\mu}^k/\rho_k)_+\|^2 \right) \right) - x^k \right\| \leq \varepsilon_k.$$

Therefore, by the non-expansivity of projections and (26), we have that

$$\left\| P_{[\ell, u]} \left(x^k - \frac{1}{2} \nabla \left(\|h(x^k) + \bar{\lambda}^k/\rho_k\|^2 + \|(g(x^k) + \bar{\mu}^k/\rho_k)_+\|^2 \right) \right) - x^k \right\| \leq \varepsilon_k + \frac{c_f}{\rho_k}. \quad (44)$$

So, by (38) and (40),

$$\left\| P_{[\ell, u]} \left(x^k - \nabla \left(\|h(x^k) + \bar{\lambda}^k/\rho_k\|^2 + \|(g(x^k) + \bar{\mu}^k/\rho_k)_+\|^2 \right) \right) - x^k \right\| \leq \delta_{\text{low}}/2. \quad (45)$$

Therefore, by Lemma 3.4 and (41),

$$\left\| P_{[\ell, u]} \left(x^k - \nabla \left(\|h(x^k)\|^2 + \|g(x^k)_+\|^2 \right) \right) - x^k \right\| \leq \delta_{\text{low}}. \quad (46)$$

By (36) and (37), we have that $k \leq k_{\text{end}}$, so, by (46),

$$\|h(x^k)\|_\infty \leq \delta \text{ and } \|g(x^k)_+\|_\infty \leq \delta. \quad (47)$$

By (47), $g_j(x^k) \leq \delta$ for all $j = 1, \dots, p$. Now, if $g_j(x^k) < -\delta$, we have that $\bar{\mu}_j^k + \rho_k g_j(x^k) < \bar{\mu}_j^k - \delta \rho_k$, which is smaller than zero because of (42), so $\mu_j^{k+1} = 0$.

Therefore, the approximate feasibility and complementarity conditions

$$\|h(x^k)\|_\infty \leq \delta, \|g(x^k)_+\| \leq \delta, \text{ and } \mu_j^k = 0 \text{ if } g_j(x^k) < -\delta \quad (48)$$

hold at x^k . Moreover, by (43) and Lemma 3.1, we have that (30) also holds. Therefore, we proved that (36), (37), (38), (39), (40), (41), (42), and (43) imply (30), (31), and (32). So, we only need to show that there exists k that satisfies (37)–(43) or satisfies (37), (30), (31), and (32). In other words, we must prove that, before completing

$$N(\delta_{\text{low}}, \varepsilon) + \left\lceil \frac{\log(\delta/c_{\text{big}})}{\log(\tau)} \right\rceil \times \left\lceil \frac{\log(\rho_{\text{max}}/\rho_1)}{\log(\gamma)} \right\rceil,$$

iterations, we get (30), (31), and (32) or we get (37)–(43).

To prove this statement, suppose that, for all k satisfying (37), at least one among the conditions (30), (31), and (32) does not hold. Since (30) necessarily holds if $k \geq N(\delta_{\text{low}}, \varepsilon)$, this implies that for all k satisfying (37) and (43) at least one among the conditions (31) and (32) does not hold. By Lemma 3.1, this implies that for all k satisfying (37) and (43),

$$\max\{\|h(x^k)\|_\infty, \|V_k\|_\infty\} > \delta.$$

Then, by (14), for $k \geq N(\delta_{\text{low}}, \varepsilon)$, the existence of more than $\log(\delta/c_{\text{big}})/\log(\tau)$ consecutive iterations $k, k+1, k+2, \dots$ satisfying (4) and (37) is impossible.

Therefore, after the first $N(\delta_{\text{low}}, \varepsilon)$ iterations, if ρ_k is increased at iterations $k_1 < k_2$, but not at any iteration $k \in (k_1, k_2)$, we have that $k_2 - k_1 \leq \log(\delta/c_{\text{big}})/\log(\tau)$. This means that, after the first $N(\delta_{\text{low}}, \varepsilon)$ iterations, the number of iterations at which ρ_k is not increased is bounded above by $\log(\delta/c_{\text{big}})/\log(\tau)$ times the number of iterations at which ρ_k is increased. Now, for obtaining (39)–(42), $\log(\rho_{\text{max}}/\rho_1)/\log(\gamma)$ iterations in which ρ_k is increased are obviously sufficient. This completes the desired result. \square

Theorem 3.6 *In addition to the hypotheses of Theorem 3.5, assume that there exist $c_{\text{inner}} > 0$, $v > 0$, and $q > 0$, where c_{inner} only depends on λ_{min} , λ_{max} , μ_{max} , ℓ , u , and characteristics of the functions f , h , and g , such that the number of inner iterations, function and derivative evaluations that are necessary to obtain (3) is bounded above by $c_{\text{inner}} \rho_k^v \varepsilon_k^{-q}$. Then, the number of inner iterations, function evaluations, and derivative evaluations that are necessary to obtain k such that (29) holds or (30), (31) and (32) hold is bounded above by*

$$c_{\text{inner}} \rho_{\text{max}}^v \varepsilon_{\text{min},3}^{-q} \left\{ N(\delta_{\text{low}}, \varepsilon) + \left\lceil \frac{\log(\delta/c_{\text{big}})}{\log(\tau)} \right\rceil \times \left\lceil \frac{\log(\rho_{\text{max}}/\rho_1)}{\log(\gamma)} \right\rceil \right\},$$

where ρ_{max} is given by (28) and

$$\varepsilon_{\text{min},3} = \min\{\varepsilon_k \mid k \leq N(\delta_{\text{low}}, \varepsilon) + \left\lceil \frac{\log(\delta/c_{\text{big}})}{\log(\tau)} \right\rceil \times \left\lceil \frac{\log(\rho_{\text{max}}/\rho_1)}{\log(\gamma)} \right\rceil\}. \quad (49)$$

Proof: The desired result follows from Theorem 3.5 and the assumptions of this theorem. \square

The comparison between Theorems 3.4 and 3.6 is interesting. This comparison seems to indicate that, if we want to be confident that the diagnostic “ x^k is an infeasible stationary point of infeasibility” is correct, we must be prepared to pay for that certainty. In fact, the bound ρ_{\max} on the penalty parameter for the algorithm is defined by (28), which not only grows with $1/\delta_{\text{low}}$, but also depends on global bounds of the problem c_{lips} and c_f . Moreover, ε_k also needs to decrease below $\delta_{\text{low}}/4$ because the decrease of the projected gradient of infeasibility is only guaranteed by a stronger decrease of the projected gradient of the Augmented Lagrangian.

4 Solving the Augmented Lagrangian subproblems

The problem considered in this section is

$$\text{Minimize } \Phi(x) \text{ subject to } x \in \Omega, \quad (50)$$

where $\Omega = \{x \in \mathbb{R}^n \mid \ell \leq x \leq u\}$. We assume that Φ has continuous first derivatives and that second derivatives exist almost everywhere. When the Hessian at a point x does not exist we call $\nabla^2\Phi(x)$ the limit of $\nabla^2\Phi(x^j)$ for a sequence x^j that converges to x . Problem (50) is of the same type of the problem that is approximately solved at Step 1 of Algorithm 2.1 and we have in mind the case $\Phi(x) \equiv L_{\rho_k}(x, \bar{\lambda}^k, \bar{\mu}^k)$.

For all $I \subseteq \{1, \dots, 2n\}$, we define the *open face*

$$F_I = \{x \in \Omega \mid x_i = \ell_i \text{ if } i \in I, x_i = u_i \text{ if } n+i \in I, \ell_i < x_i < u_i \text{ otherwise}\}.$$

By definition, Ω is the union of its open faces and the open faces are disjoint. This means that every $x \in \Omega$ belongs to exactly one face F_I . The variables x_i such that $\ell_i < x_i < u_i$ are called *free variables*. For every $x \in \Omega$, we also define the continuous projected gradient of Φ given by

$$g_P(x) = P_\Omega(x - \nabla\Phi(x)) - x \quad (51)$$

and, if F_I is the open face to which x belongs, the continuous projected internal gradient $g_I(x)$ given by

$$[g_I(x)]_i = \begin{cases} [g_P(x)]_i, & \text{if } x_i \text{ is a free variable,} \\ 0, & \text{otherwise.} \end{cases}$$

Note that, sometimes, we write $g_I(x)$ omitting the fact that the subindex I refers the face F_I to which the argument $x \in \Omega$ belongs.

The bound-constrained minimization method described in the current section can be seen as a second-order counterpart of the method introduced in [16]. (See, also, [11].) The iterates visit the different faces of the box Ω preserving the current face while the quotient $\|g_I(x)\|/\|g_P(x)\|$ is big enough or the new iterate does not hit the boundary. When this quotient reveals that few progress can be expected from staying in the current face, the face is abandoned by means of a spectral projected gradient [21, 22, 23] iteration. Within each face, iterations obey a safeguarded Newton scheme with line searches. The employment of this method is coherent with the conservative point of view of Algencan. For example, we do not aim to predict the active constraints at the solution and the inactive bounds have no influence in the iterations independently of the

distance of the current iterate to a bound. Moreover, we do not try to use second-order information for leaving the faces. Of course, we do not deny the efficiency of methods that employ such procedures, but we feel comfortable with the conservative strategy because the number of algorithmic parameters can be reduced to a minimum.

Algorithm 4.1: Assume that $x^0 \in \Omega$, $\Phi_{\text{target}} \in \mathbb{R}$, $r \in (0, 1]$, $0 < \tau_1 \leq \tau_2 < 1$, $\gamma \in (0, 1)$, $\beta \in (0, 1)$, $0 < \eta$, $0 < \lambda_{\min}^{\text{SPG}} < \lambda_{\max}^{\text{SPG}}$, $0 < \sigma_{\text{small}}$, $0 < \sigma_{\min} \leq \sigma_{\max}$, $0 < \underline{h} < \bar{h}$, $t_{\max}^{\text{ext}} \in \mathbb{N}_{\geq 0}$ are given. Initialize $k \leftarrow 0$.

Step 1. If $\Phi(x^k) \leq \Phi_{\text{target}}$ then stop. Otherwise, if $\|g_I(x^k)\|_{\infty} \geq r\|g_P(x^k)\|_{\infty}$ then go to Step 2 to perform an *inner-to-the-face iteration using Newton with line search* else go to Step 5 to perform a *leaving-face iteration using spectral projected gradients (SPG)*.

Step 2. Let \bar{n} be the number of free variables and let $\bar{H}_k \in \mathbb{R}^{\bar{n} \times \bar{n}}$ be the Hessian $\nabla^2 \Phi(x^k)$ in which rows and columns associated with *non* free variables were removed.

Step 2.1. If \bar{H}_k is positive definite then set $\sigma \leftarrow 0$ and compute $\bar{d}^k \in \mathbb{R}^{\bar{n}}$ as the solution of $\bar{H}_k \bar{d} = -\bar{g}^k$, where $\bar{g}^k \in \mathbb{R}^{\bar{n}}$ corresponds to $\nabla \Phi(x^k)$ with the components associated with the *non* free variables removed, and go to Step 2.3.

Step 2.2. Inertia correction

Step 2.2.1. If σ^{ini} is undefined then set $\sigma^{\text{ini}} \leftarrow P_{[\sigma_{\min}, \sigma_{\max}]}(\sigma_{\text{small}} h)$, where $h = P_{[\underline{h}, \bar{h}]}(\max_{\{i=1, \dots, \bar{n}\}} \{ |[\bar{H}_k]_{ii}| \})$.

Step 2.2.2. Set $\sigma \leftarrow \sigma^{\text{ini}}$ and while $\bar{H}_k + \sigma I$ is not positive definite do $\sigma \leftarrow 10\sigma$.

Step 2.2.3. Compute $\bar{d}^k \in \mathbb{R}^{\bar{n}}$ as a solution of $(\bar{H}_k + \sigma I)\bar{d} = -\bar{g}^k$, where $\bar{g}^k \in \mathbb{R}^{\bar{n}}$ corresponds to $\nabla \Phi(x^k)$ with the components associated with the *non* free variables removed.

Step 2.2.4. While $\|\bar{d}^k\|_2 > \eta \max\{1, \|\bar{x}^k\|_2\}$ do $\sigma \leftarrow 10\sigma$ and redefine \bar{d}^k as the solution of $(\bar{H}_k + \sigma I)\bar{d} = -\bar{g}^k$. ($\bar{x}^k \in \mathbb{R}^{\bar{n}}$ corresponds to the free components of x^k .)

Step 2.3. Let $d^k \in \mathbb{R}^n$ be the “expansion” of \bar{d}^k with $[d^k]_i = 0$ if i is a *non*-free-variable index.

Step 2.4. If $\sigma > 0$ then set $\sigma^{\text{ini}} \leftarrow P_{[\sigma_{\min}, \sigma_{\max}]}(\frac{1}{2}\sigma)$. Otherwise, set $\sigma^{\text{ini}} \leftarrow P_{[\sigma_{\min}, \sigma_{\max}]}(\frac{1}{2}\sigma^{\text{ini}})$.

Step 3. Line search plus possible projection

Step 3.1. Compute α_{\max} as the largest $\alpha > 0$ such that $x^k + \alpha d^k \in \Omega$. If $\alpha_{\max} \geq 1$ (i.e. $x^k + d^k \in \Omega$) then skip Step 3.2 below (i.e. go to Step 3.3).

Step 3.2. Set $x_{\text{trial}} \leftarrow P_{\Omega}(x^k + d^k)$. If $\Phi(x_{\text{trial}}) \leq \Phi_{\text{target}}$ or $\Phi(x_{\text{trial}}) \leq \Phi(x^k)$ then set $x^{k+1} = x_{\text{trial}}$ and go to Step 4.

Step 3.3. Set $t \leftarrow \min\{1, \alpha_{\max}\}$ and $x_{\text{trial}} \leftarrow x^k + td^k$. While $\Phi(x_{\text{trial}}) > \Phi_{\text{target}}$ and $\Phi(x_{\text{trial}}) > \Phi(x^k) + t\gamma \langle \nabla \Phi(x^k), d^k \rangle$, choose $t_{\text{new}} \in [\tau_1 t, \tau_2 t]$ and set $t \leftarrow t_{\text{new}}$ and $x_{\text{trial}} \leftarrow x^k + td^k$.

Step 3.4. Set $x^{k+1} = x_{\text{trial}}$. If $t = \alpha_{\max}$ or ($t = 1$ and $\langle \nabla \Phi(x_{\text{trial}}), d^k \rangle > \beta \langle \nabla \Phi(x^k), d^k \rangle$) then go to Step 4. Otherwise, go to Step 6.

Step 4. Extrapolation

Step 4.1. Set $t \leftarrow 1$, $x_{\text{ref}} \leftarrow x_{\text{trial}}$, and $x_{\text{ext}} \leftarrow P_{\Omega}(x^k + 2^t(x_{\text{trial}} - x^k))$.

Step 4.2. While $t \leq t_{\text{max}}^{\text{ext}}$ and $\Phi(x_{\text{ext}}) < f(x_{\text{ref}})$ and $\Phi(x_{\text{ref}}) > \Phi_{\text{target}}$ do
 $t \leftarrow t + 1$, $x_{\text{ref}} \leftarrow x_{\text{ext}}$, and $x_{\text{ext}} \leftarrow P_{\Omega}(x^k + 2^t(x_{\text{trial}} - x^k))$.

Step 4.3. Reset $x^{k+1} = x_{\text{ref}}$ and go to Step 6.

Step 5. *Leaving-face SPG iteration*

Step 5.1. If $k = 0$ or $\langle x^k - x^{k-1}, \nabla f(x^k) - \nabla f(x^{k-1}) \rangle \leq 0$ then set

$$\lambda_k^{\text{SPG}} = \max \left\{ 1, \|x^k\|_2 / \|g_P(x^k)\|_2 \right\}.$$

Otherwise, set $\lambda_k^{\text{SPG}} = \|x^k - x^{k-1}\|_2^2 / \langle x^k - x^{k-1}, \nabla \Phi(x^k) - \nabla \Phi(x^{k-1}) \rangle$.

In any case, redefine λ_k^{SPG} as $\max\{\lambda_{\text{min}}^{\text{SPG}}, \min\{\lambda_k^{\text{SPG}}, \lambda_{\text{max}}^{\text{SPG}}\}\}$.

Step 5.2. Set $t \leftarrow 1$, $x_{\text{trial}} \leftarrow P_{\Omega}(x^k - \lambda_k^{\text{SPG}} \nabla \Phi(x^k))$, and $d^k = x_{\text{trial}} - x^k$.

Step 5.3. While $\Phi(x_{\text{trial}}) > \Phi_{\text{target}}$ and $\Phi(x_{\text{trial}}) > \Phi(x^k) + t\gamma \langle \nabla \Phi(x^k), d^k \rangle$,
choose $t_{\text{new}} \in [\tau_1 t, \tau_2 t]$ and set $t \leftarrow t_{\text{new}}$ and $x_{\text{trial}} \leftarrow x^k + t d^k$.

Step 5.4. Set $x^{k+1} = x_{\text{trial}}$.

Step 6. In the current iteration, the definition of x^{k+1} implied in the evaluation of Φ at several points named x_{trial} . If, for any of them, we have that $\Phi(x_{\text{trial}}) < \Phi(x^{k+1})$ then reset $x^{k+1} = x_{\text{trial}}$. In any case, set $k \leftarrow k + 1$ and go to Step 1.

Remark 1. At Steps 3.3 and 5.3, interpolation is done with safeguarded quadratic interpolation. This means that, given

$$t_{\text{temp}} = - \frac{\langle \nabla \Phi(x^k), d^k \rangle t^2}{2(\Phi(x_{\text{trial}}) - \Phi(x^k)) - t \langle \nabla \Phi(x^k), d^k \rangle},$$

if $t_{\text{temp}} \in [\tau_1 t, \tau_2 t]$ then $t_{\text{new}} = t_{\text{temp}}$. Otherwise, $t_{\text{new}} = \frac{1}{2}t$. This choice requires $0 < \tau_1 \leq \frac{1}{2} \leq \tau_2 < 1$ instead of simply $0 < \tau_1 \leq \tau_2 < 1$.

5 Implementation details and parameters

We implemented Algorithms 2.1 and 4.1 in Fortran 90. Implementation is freely available at <http://www.ime.usp.br/~egbirgin/>. Interfaces for solving user-defined problems coded in Fortran 90 as well as problems from the CUTEst [37] collection are available. All tests reported below were conducted on a computer with 3.5 GHz Intel Core i7 processor and 16GB 1600 MHz DDR3 RAM memory, running OS X High Sierra (version 10.13.6). Codes were compiled by the GFortran compiler of GCC (version 8.2.0) with the -O3 optimization directive enabled.

5.1 Augmented Lagrangian method

Algorithm 2.1 was devised to be applied to a scaled version of problem (50). Following the Ipopt strategy described in [47, p.46], in the scaled problem, the objective function f is multiplied by

$$s_f = \max \left\{ 10^{-8}, \frac{100}{\max\{1, \|\nabla f(x^0)\|_\infty\}} \right\},$$

each constraint h_j ($j = 1, \dots, m$) is multiplied by

$$s_{h_j} = \max \left\{ 10^{-8}, \frac{100}{\max\{1, \|\nabla h_j(x^0)\|_\infty\}} \right\},$$

and each constraint g_j ($j = 1, \dots, p$) is multiplied by

$$s_{g_j} = \max \left\{ 10^{-8}, \frac{100}{\max\{1, \|\nabla g_j(x^0)\|_\infty\}} \right\},$$

where $x^0 \in \mathbb{R}^n$ is the given initial guess. The scaling is optional and it is used when the input parameter “scale” is set to “true”. If the parameter is set to “false”, the original problem, that corresponds to considering all scaling factors equal to one, is solved.

As stopping criterion, we say that an iterate $x^k \in [\ell, u]$ with its associated Lagrange multipliers λ^{k+1} and μ^{k+1} satisfies the main stopping criterion when

$$\max \left\{ \|h(x^k)\|_\infty, \|g(x^k)_+\|_\infty \right\} \leq \varepsilon_{\text{feas}}, \quad (52)$$

$$\left\| P_{[\ell, u]} \left(x^k - \left[s_f \nabla f(x^k) + \sum_{j=1}^m \lambda_j^{k+1} s_{h_j} \nabla h_j(x^k) + \sum_{j=1}^p \mu_j^{k+1} s_{g_j} \nabla g_j(x^k) \right] \right) - x^k \right\|_\infty \leq \varepsilon_{\text{opt}}, \quad (53)$$

$$\max_{j=1, \dots, p} \left\{ \min \left\{ -s_{g_j} g_j(x^k), \mu_j^{k+1} \right\} \right\} \leq \varepsilon_{\text{compl}}, \quad (54)$$

where $\varepsilon_{\text{feas}} > 0$, $\varepsilon_{\text{opt}} > 0$, and $\varepsilon_{\text{compl}} > 0$ are given constants. This means that the stopping criterion requires *unscaled* feasibility with tolerance $\varepsilon_{\text{feas}}$ plus *scaled* optimality with tolerance ε_{opt} and *scaled* complementarity (measured with the min function) with tolerance $\varepsilon_{\text{compl}}$. Note that $x^k \in [\ell, u]$, i.e. it satisfies the bound-constraints with zero tolerance. In addition to this stopping criterion, Algorithm 2.1 also stops if the penalty parameter ρ_k reaches the value ρ_{big} or if, in three consecutive iterations, the inner solver that is used at Step 1 fails at finding a point $x^k \in [\ell, u]$ that satisfies (3).

In (3) and (4), we consider $\|\cdot\| = \|\cdot\|_\infty$. At Step 2, we consider $\varepsilon_1 = \sqrt{\varepsilon_{\text{opt}}}$ and $\varepsilon_k = \max\{\varepsilon_{\text{opt}}, 0.1\varepsilon_{k-1}\}$ for $k > 1$; and, at Step 3, if $\lambda^{k+1} \in [\lambda_{\min}, \lambda_{\max}]^m$ and $\mu_i^{k+1} \in [0, \mu_{\max}]^p$ then we set $\bar{\lambda}^{k+1} = \lambda^{k+1}$ and $\bar{\mu}^{k+1} = \mu^{k+1}$. Otherwise, we set $\bar{\lambda}^{k+1} = 0$ and $\bar{\mu}^{k+1} = 0$. In the numerical experiments, we set $\varepsilon_{\text{feas}} = \varepsilon_{\text{opt}} = \varepsilon_{\text{compl}} = 10^{-8}$, $\rho_{\text{big}} = 10^{20}$, $\lambda_{\min} = -10^{16}$, $\lambda_{\max} = 10^{16}$, $\mu_{\max} = 10^{16}$, $\gamma = 10$, $\tau = 0.5$, $\bar{\lambda}^1 = 0$, $\bar{\mu}^1 = 0$, and

$$\rho_1 = 10 \max \left\{ 1, \frac{|f(x^0)|}{\max\{\|h(x^0)\|_2^2 + \|g(x^0)_+\|_2^2\}} \right\}.$$

Two additional strategies complete the implementation of Algorithm 2.1. On the one hand, if Algorithm 2.1 fails at finding a point that satisfies (52), the feasibility problem (7) is tackled with Algorithm 4.1 with the purpose of, at least, finding a feasible point to the original NLP problem (1). On the other hand, at every iteration k , prior to the subproblem minimization at Step 1, $(x^{k-1}, \lambda^k, \mu^k)$ is used as initial guess to perform ten iterations of the “pure” Newton method (no line search, no inertia correction) applied to the semismooth KKT system [42, 44] associated with problem (50), with dimension $3n + m + p$, given by

$$\begin{pmatrix} \nabla f(x) + \sum_{j=1}^m \lambda_j \nabla h_j(x) + \sum_{j=1}^p \mu_j \nabla g_j(x) - \nu^\ell + \nu^u \\ h(x) \\ \min\{-g(x), \mu\} \\ \min\{x - \ell, \nu^\ell\} \\ \min\{u - x, \nu^u\} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix},$$

where $\nu^\ell, \nu^u \in \mathbb{R}^n$ are the Lagrange multipliers associated with the bound constraints $\ell \leq x$ and $x \leq u$, respectively. This process is related to the so-called acceleration process described in [18] in which a different KKT system was considered. (See [18] for details.) The stopping criteria for the acceleration process are (i) “the Jacobian of the KKT system has the ‘wrong’ inertia”, (ii) “a maximum of 10 iterations was reached”, and (iii)

$$\max \{ \|h(x)\|_\infty, \|g(x)_+\|_\infty, \|(\ell - x)_+\|_\infty, \|(x - u)_+\|_\infty \} \leq \varepsilon_{\text{feas}}, \quad (55)$$

$$\left\| \nabla f(x) + \sum_{j=1}^m \lambda_j \nabla h_j(x) + \sum_{j=1}^p \mu_j \nabla g_j(x) - \nu^\ell + \nu^u \right\|_\infty \leq \varepsilon_{\text{opt}}, \quad (56)$$

$$\max \left\{ \max_{j=1, \dots, p} \{[\min\{-g(x), \mu\}]_j\}, \max_{i=1, \dots, n} \{[\min\{x - \ell, \nu^\ell\}]_i\}, \max_{i=1, \dots, n} \{[\min\{u - x, \nu^u\}]_i\} \right\} \leq \varepsilon_{\text{compl}}. \quad (57)$$

Note that criterion (iii) corresponds to satisfying approximate KKT conditions for the *unscaled* original problem (1). On the other hand, differently from an iterate $x^k \in [\ell, u]$ of Algorithm 2.1 that satisfies (52,53,54), a point that satisfies criterion (iii) may violate the bound constraints with tolerance $\varepsilon_{\text{feas}}$.

If the acceleration process stops satisfying criterion (i) or (ii), everything it was done in the acceleration is discarded and the iterations of Algorithm 2.1 continue. On the other hand, assume that a point satisfying criterion (iii) was found by the acceleration process. If $(x^{k-1}, \lambda^k, \mu^k)$ satisfies (52,53,54) with half the precision, i.e. with $\varepsilon_{\text{feas}}$, ε_{opt} , and $\varepsilon_{\text{compl}}$ substituted by $\varepsilon_{\text{feas}}^{1/2}$, $\varepsilon_{\text{opt}}^{1/2}$, and $\varepsilon_{\text{compl}}^{1/2}$, respectively, then we say the acceleration was successful, the point found by the acceleration process is returned, and the optimization process stops. On the other hand, if $(x^{k-1}, \lambda^k, \mu^k)$ is far from satisfying (52,53,54), we believe the approximate KKT point the acceleration found may be an undesirable point. The point is saved for further references, but the optimization process continues; and the next Augmented Lagrangian subproblem is tackled by Algorithm 4.1 starting from x^{k-1} and ignoring the point found by the acceleration process.

5.2 Bound-constrained minimization method

As main stopping criterion of Algorithm 4.1, we considered the condition

$$\|g_P(x^k)\|_\infty \leq \varepsilon \quad (58)$$

where $g_P(x^k) = P_{[\ell, u]}(x^k - \nabla\Phi(x^k)) - x^k$ as defined in (51). When an unconstrained or bound-constrained problem is being solved, in (58) and in the alternative stopping criteria described below, we use $\varepsilon = \varepsilon_{\text{opt}} = 10^{-8}$. When the problem being tackled by Algorithm 4.1 is a subproblem of Algorithm 2.1, the value of ε in (58) and in the alternative stopping criteria described below is the one described in Section 5.1 (that we cannot mention here since we are using k to denote iterations of both Algorithms 2.1 and 4.1). In addition, Algorithm 4.1 may also stop at iteration k by any of the following alternative stopping criteria: **(a)** $\|g_P(x^{k-\ell})\|_\infty < \sqrt{\varepsilon}$ for all $0 \leq \ell < 100$; **(b)** $\|g_P(x^{k-\ell})\|_\infty < \varepsilon^{1/4}$ for all $0 \leq \ell < 5,000$; **(c)** $\|g_P(x^{k-\ell})\|_\infty < \varepsilon^{1/8}$ for all $0 \leq \ell < 10,000$; **(d)** $\Phi(x^k) \leq \Phi_{\text{target}}$; **(e)** $k \geq k_{\text{max}} = 50,000$; and **(f)** k_{best} is the smallest index such that $\Phi(x^{k_{\text{best}}}) = \min\{\Phi(x^0), \Phi(x^1), \dots, \Phi(x^k)\}$ and $k - k_{\text{best}} > 3$, i.e. the best functional value so far obtained is not updated in three consecutive iterations.

In Algorithm 4.1, although the theory allows us a wide range of possibilities, in practice we consider $H_k = \nabla^2\Phi(x^k)$ for all k . The linear systems at Step 2.1 and 2.2.2 are solved with subroutine MA57 from HSL [48] (using all its default parameters). In the experiments, we set $\Phi_{\text{target}} = -10^{12}$, $r = 0.1$, $\tau_1 = 0.1$, $\tau_2 = 0.9$, $\gamma = 10^{-4}$, $\beta = 0.5$, $\eta = 10^4$, $\lambda_{\text{min}}^{\text{SPG}} = 10^{-16}$, $\lambda_{\text{max}}^{\text{SPG}} = 10^{16}$, $\sigma_{\text{small}} = 10^{-8}$, $\sigma_{\text{min}} = 10^{-8}$, $\sigma_{\text{max}} = 10^{16}$, $\underline{h} = 10^{-8}$, $\bar{h} = 10^8$, and $t_{\text{max}}^{\text{ext}} = 20$.

When Algorithm 4.1 is used to solve a subproblem of Algorithm 2.1, we have that $\nabla^2\Phi(x) = \nabla^2 L_{\rho_k}(x, \bar{\lambda}_k, \bar{\mu}_k)$, i.e. $\nabla^2\Phi(x)$ is the Hessian of the augmented Lagrangian associated with the scaled version of problem (50) given by

$$s_f \nabla^2 f(x) + \sum_{j=1}^m \left\{ \bar{\lambda}_j^k s_{h_j} \nabla^2 h_j(x) + \rho_k s_{\bar{h}_j}^2 \nabla h_j(x) \nabla h_j(x)^T \right\} + \sum_{j \in I_k} \left\{ \bar{\mu}_j^k s_{g_j} \nabla^2 g_j(x) + \rho_k s_{\bar{g}_j}^2 \nabla g_j(x) \nabla g_j(x)^T \right\}, \quad (59)$$

where $I_k = I_{\rho_k}(x^k, \bar{\mu}^k) = \{j = 1, \dots, p \mid \bar{\mu}^k + \rho_k s_{g_j} g_j(x^k) > 0\}$. A relevant issue from the practical point of view is that, despite the sparsity of the Hessian of the Lagrangian and the sparsity of the Jacobian of the constraints, this matrix may be dense. Thus, factorizing, or even building it, may be prohibitive. As an alternative, instead of building and factorizing the Hessian above, it can be solved an augmented linear system with the coefficients' matrix given by

$$\left(\begin{array}{c|c} s_f \nabla^2 f(x) + \sum_{j=1}^m \left\{ \bar{\lambda}_j^k s_{h_j} \nabla^2 h_j(x) \right\} + \sum_{j \in I_k} \left\{ \bar{\mu}_j^k s_{g_j} \nabla^2 g_j(x) \right\} & J(x)^T \\ \hline J(x) & -\frac{1}{\rho_k} I \end{array} \right), \quad (60)$$

where $J(x)$ is a matrix whose columns are $\nabla h_1(x), \dots, \nabla h_m(x)$ plus the gradients $\nabla g_j(x)$ such that $j \in I_k$. This matrix preserves the sparsity of the Hessian of the Lagrangian and of the Jacobian of the constraints. The implementation of Algorithms 4.1 dynamically selects one of the two approaches.

Another relevant fact from the practical point of view, related to matrices (59) and (60), is that the current tools available in CUTEst compute the full Jacobian of the constraints and $\sum_{j=1}^p \bar{\mu}_j^k s_{g_j} \nabla^2 g_j(x)$ with $\bar{\mu}_j^k = 0$ if $j \notin I_k$ instead of $J(x)$ and $\sum_{j \in I_k} \bar{\mu}_j^k s_{g_j} \nabla^2 g_j(x)$, respectively. On the one hand, this feature preserves the Jacobian's and the Hessian-of-the-Lagrangian's sparsity structures independently of $\bar{\mu}^k$ and x , as required by some solvers. On the other hand, it impairs Algorithm 2.1, when applied to problems from the CUTEst collection, of fully exploiting

the potential advantage of dealing with inequality constraints without adding slack variables. In summary, Algorithm 2.1–4.1 is prepared to deal with matrices with different sparsity structures at every iteration and, for that reason, it performs the analysis step of the factorization at every iteration. This is the price to pay for exploiting inequality constraints without adding slack variables. However, the CUTEst subroutines are not prepared to exploit this feature and Algorithm 2.1–4.1, when solving problems from the CUTEst collection, pays the price without enjoying the advantages. Of course, this CUTEst inconvenient influences negatively the comparison of Algencan with other solvers if the CPU time is used as a performance measure.

6 Numerical experiments

In this section, we aim to evaluate the performance of Algorithm 2.1–4.1 (referred as Algencan from now on) for solving unconstrained, bound-constrained, feasibility, and nonlinear programming problems. The performance of Ipopt [47] (version 3.12.12) is also exhibited. Both methods were run in the same computational environment, compiled with the same BLAS routines, and also using the same subroutine MA57 from HSL for solving the linear systems. All Ipopt default parameters were used¹. A CPU time limit of 10 minutes per problem was imposed. In the numerical experiments, we considered all 1,258 problems from the CUTEst collection [37] with their default dimensions. In the collection, there are 217 unconstrained problems, 144 bound-constrained problems, 157 feasibility problems, and 740 nonlinear programming problems. A hint on the number of variables in each family is given in Table 1.

Problem type	# of problems	n_{\max}	# of problems with $n \geq \omega n_{\max}$		
			$\omega = 0.1$	$\omega = 0.01$	$\omega = 0.001$
unconstrained	217	100,000	15	87	97
bound-constrained	144	149,624	5	60	72
feasibility	156	123,200	5	40	55
NLP	740	250,997	67	263	379

Table 1: Distribution of the number of variables n in the CUTEst collection test problems.

Large tables with a detailed description of the output of each method in the 1,258 problems can be found in <http://www.ime.usp.br/~egbirgin/>. A brief overview follows. Note that, since the methods differ in the stopping criteria, arbitrary decisions will be made. A point in common is that both methods seek satisfying the (sup-norm of the) violation of the unscaled equality and inequality constraints with precision $\varepsilon_{\text{feas}} = 10^{-8}$. However, as described in [47, §3.5], Ipopt considers a *relative* initial relaxation of the bound constraints (whose default value is 10^{-8}); and it may apply repeated additional relaxations during the optimization process. Table 2 shows the number of problems in which each method found a point satisfying

$$\max\{\|h(x)\|_{\infty}, \| [g(x)]_+ \|_{\infty}\} \leq \varepsilon_{\text{feas}} \quad (61)$$

¹Option 'honor_original_bounds no', that does not affect Ipopt's optimization process, was used. Ipopt might relax the bounds during the optimization beyond its initial *relative* relaxation factor whose default value is 10^{-8} . Option 'honor_original_bounds no' simply avoids the final iterate to be projected back onto the box defined by the bound constraints. So, the actual absolute violation of the bound constraints at the final iterate can be measured.

plus

$$\max\{\|(\ell - x)_+\|_\infty, \|(x - u)_+\|_\infty\} \leq \bar{\varepsilon}_{\text{feas}} \quad (62)$$

with $\varepsilon_{\text{feas}} = 10^{-8}$ and $\bar{\varepsilon}_{\text{feas}} \in \{0.1, 10^{-2}, \dots, 10^{-16}, 0\}$. Figures in the table show that, in most cases, Algencan satisfies the bound constraints with zero tolerance and that the violation of the bound constraints rarely exceeds the tolerance 10^{-8} . This is an expected result, since the method satisfies these requirements by definition. Regarding Ipopt, the table shows in which way the amount of problems in which (62) holds varies as a function of the tolerance $\bar{\varepsilon}_{\text{feas}}$.

	$\bar{\varepsilon}_{\text{feas}}$																
	0.1	10^{-2}	10^{-3}	10^{-4}	10^{-5}	10^{-6}	10^{-7}	10^{-8}	10^{-9}	10^{-10}	10^{-11}	10^{-12}	10^{-13}	10^{-14}	10^{-15}	10^{-16}	0
Algencan	1,132	1,132	1,131	1,131	1,131	1,130	1,130	1,130	1,121	1,115	1,112	1,105	1,092	1,082	1,077	1,069	1,058
Ipopt	1,073	1,072	1,070	1,068	1,056	1,044	1,016	970	794	793	793	793	793	792	792	792	791

Table 2: Number of problems in which a point satisfying (61,62) was found by Algencan and Ipopt with $\varepsilon_{\text{feas}} = 10^{-8}$ and $\bar{\varepsilon}_{\text{feas}} \in \{0.1, 10^{-2}, \dots, 10^{-16}, 0\}$.

If the violation of the bound constraints is disregarded, Table 2 shows that Algencan found points satisfying (61,62) with $\varepsilon_{\text{feas}} = 10^{-8}$ and $\bar{\varepsilon}_{\text{feas}} = 0.1$ in 1,132 problems; while Ipopt found the same in 1,073. There are in the CUTEst collection 85 problems (62 feasibility problems and 23 nonlinear programming problems) in which the number of equality constraints is larger than the number of variables. Ipopt does not apply to these problems and, thus, of course, it does not find a point satisfying (61,62). Algencan *did* find a point satisfying (61,62) in 28 out of the 85 problems to which Ipopt does not apply; and this explains half of the difference between the methods. In any case, it can be said that, over a universe of 1,258 problems, both methods found “feasible points” in a large fraction of the problems; recalling that the collection contains infeasible problems.

We now consider the set of 757 problems in which both methods found a point satisfying (61) with $\varepsilon_{\text{feas}} = 10^{-8}$ and (62) with $\bar{\varepsilon}_{\text{feas}} = 0$. For a given problem, let f_1 be the value of the objective function at the point found by Algencan; let f_2 be the value of the objective function at the point found by Ipopt; and let $f^{\min} = \min\{f_1, f_2\}$. Table 3 shows in how many problems it holds

$$f_i \leq f^{\min} + f_{\text{tol}} \max\{1, |f^{\min}|\} \text{ for } i = 1, 2 \quad (63)$$

and $f_{\text{tol}} \in \{0.1, 10^{-2}, \dots, 10^{-8}, 0\}$.

	f_{tol}									
	0.1	10^{-2}	10^{-3}	10^{-4}	10^{-5}	10^{-6}	10^{-7}	10^{-8}	0	
Algencan	722	715	706	694	691	678	675	663	498	
Ipopt	723	708	699	694	683	653	623	592	383	

Table 3: Number of problems in which a point satisfying (61) with $\varepsilon_{\text{feas}} = 10^{-8}$, (62) with $\bar{\varepsilon}_{\text{feas}} = 0$, and (63) with $f_{\text{tol}} \in \{0.1, 10^{-2}, \dots, 10^{-8}, 0\}$ was found by Algencan and Ipopt.

Finally, we consider the set of 688 problems in which both, Algencan and Ipopt, found a point that satisfies (61) with $\varepsilon_{\text{feas}} = 10^{-8}$, (62) with $\bar{\varepsilon}_{\text{feas}} = 0$, and (63) with $f_{\text{tol}} = 0.1$. For this set of problems, Figure 1 shows the performance profile [30] that considers, as performance

measure, the CPU time spent by each method. In the figure, for $i \in M \equiv \{\text{Algencan}, \text{Ipopt}\}$,

$$\Gamma_i(\tau) = \frac{\#\{j \in \{1, \dots, q\} \mid t_{ij} \leq \tau \min_{s \in M} \{t_{sj}\}\}}{q},$$

where $\#\mathcal{S}$ denotes the cardinality of set \mathcal{S} , $q = 688$ is the number of considered problems, and t_{ij} is the performance measure (CPU time) of method i applied to problem j . Thus, $\Gamma_{\text{Algencan}}(1) = 0.48$ and $\Gamma_{\text{Ipopt}}(1) = 0.53$ says that Algencan was faster than Ipopt in 48% of the problems and Ipopt was faster than Algencan in 53% of the problems. Complementing the performance profile, we can report that there are 9 problems in which both methods spent at least a second of CPU time and one of the methods is at least ten times faster than the other. Among these 9 problems, Ipopt is faster in 5 and Algencan is faster in the other 4.

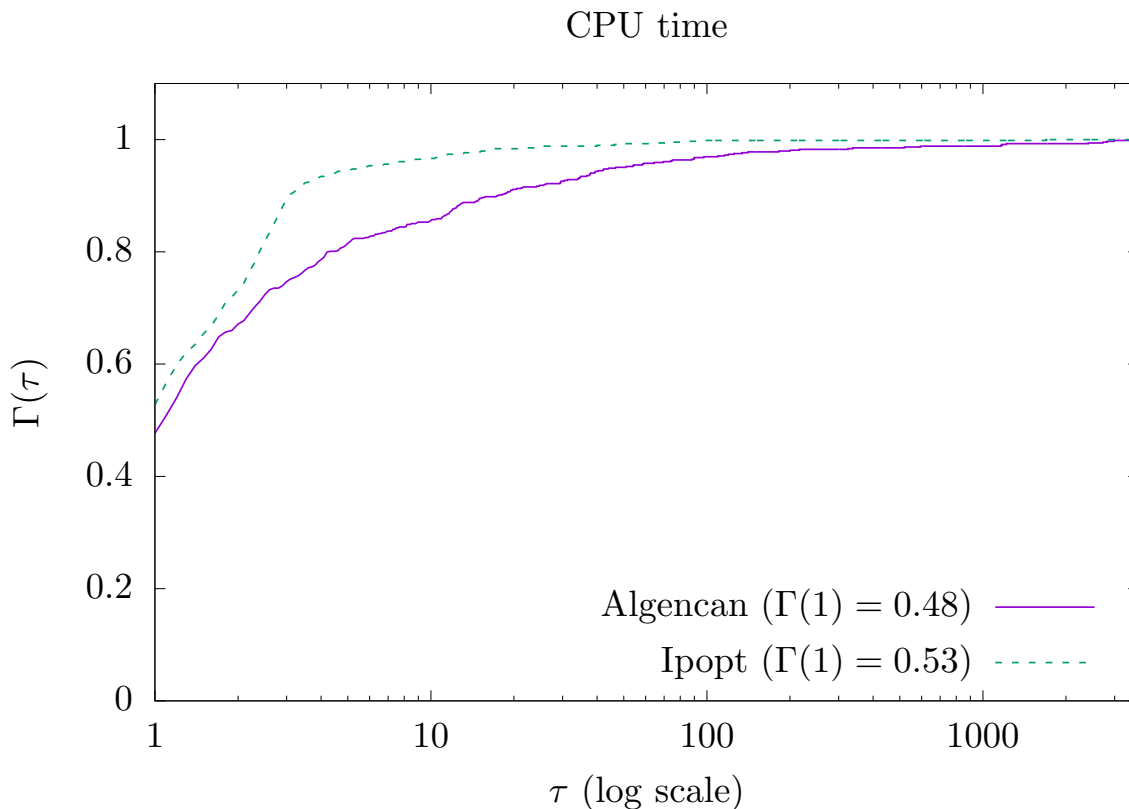


Figure 1: Performance profiles comparing the CPU time spent by Algencan and Ipopt in the 688 problems in which both methods found a point that satisfies (61) with $\varepsilon_{\text{feas}} = 10^{-8}$, (62) with $\bar{\varepsilon}_{\text{feas}} = 0$, and (63) with $f_{\text{tol}} = 0.1$.

7 Conclusions

In this work, a version of the (safeguarded) Augmented Lagrangian algorithm Algencan [3, 20] that possesses iteration and evaluation complexity was described, implemented, and evaluated.

Moreover, the convergence theory of Algencan was complemented with new complexity results. The way in which an Augmented Lagrangian method was able to inherit the complexity properties from a method for bound-constrained minimization is a nice example of the advantages of the modularity feature that Augmented Lagrangian methods usually possess.

As a byproduct of this development, a new version of Algencan that uses a Newtonian method with line search to solve the subproblems was developed from scratch. Moreover, the acceleration process described in [18] was revisited. In particular, the KKT system with complementarity modelled with the product between constraints and multipliers was replaced with the KKT system that models the complementarity constraints with the semismooth min function.

We provided a fully reproducible comparison with Ipopt, which is, probably, the most effective and best known free software for constrained optimization. The main feature we want to stress is that there exist a significant number of problems that Algencan solves satisfactorily whereas Ipopt does not, and vice versa. This is not surprising because the way in which Augmented Lagrangians and Interior Point Newtonian methods handle problems are qualitatively different. Constrained Optimization is an extremely heterogeneous family. Therefore, we believe that what justifies the existence of new algorithms or the survival of traditional ones is not their capacity of solving a large number of problems using slightly smaller computer time than “competitors”, but the potentiality of solving some problems that other algorithms fail to solve. Engineers and practitioners should not care about the choice between algorithm A or B according to subtle efficiency criteria. The best strategy is to contemplate both, using one or the other according to their behavior on the family of problems that they need to solve in practice. As in many aspects of life, competition should give place to cooperation.

Acknowledgements. The authors are indebted to Iain Duff, Nick Gould, Dominique Orban, and Tyrone Rees for their help in issues related to the usage of MA57 from HSL and the CUTEst collection.

References

- [1] R. Andreani, J. M. Martínez, and L. T. Santos, Newton’s method may fail to recognize proximity to optimal points in constrained optimization, *Mathematical Programming* 160, pp. 547–555, 2016.
- [2] R. Andreani, J. M. Martínez, L. T. Santos, and B. F. Svaiter, On the behavior of constrained optimization methods when Lagrange multipliers do not exist, *Optimization Methods and Software* 29, pp. 646–657, 2014.
- [3] R. Andreani, E. G. Birgin, J. M. Martínez, and M. L. Schuverdt, On augmented Lagrangian methods with general lower-level constraints, *SIAM Journal on Optimization* 18, pp. 1286–1309, 2008.
- [4] R. Andreani, E. G. Birgin, J. M. Martínez and M. L. Schuverdt, Augmented Lagrangian methods under the Constant Positive Linear Dependence constraint qualification, *Mathematical Programming* 111, pp. 5–32, 2008.

- [5] R. Andreani, E. G. Birgin, J. M. Martínez, and M. L. Schuverdt, Second-order negative-curvature methods for box-constrained and general constrained optimization, *Computational Optimization and Applications* 45, pp. 209–236, 2010.
- [6] R. Andreani, N. S. Fazzio, M. L. Schuverdt, and L. D. Secchin, A Sequential Optimality Condition Related to the Quasi-normality Constraint Qualification and its Algorithmic Consequences, *SIAM Journal on Optimization* 29, pp. 743–766, 2019.
- [7] R. Andreani, G. Haeser, and J. M. Martínez, On sequential optimality conditions for smooth constrained optimization, *Optimization* 60, pp. 627–641, 2011.
- [8] R. Andreani, J. M. Martínez, A. Ramos, and P. J. S. Silva, A Cone-Continuity Constraint Qualification and algorithmic consequences, *SIAM Journal on Optimization* 26, pp. 96–110, 2016.
- [9] R. Andreani, J. M. Martínez, A. Ramos, and P. J. S. Silva, Strict constraint qualifications and sequential optimality conditions for constrained optimization, *Mathematics of Operations Research*, to appear. doi:10.1287/moor.2017.0879
- [10] R. Andreani, J. M. Martínez, and B. F. Svaiter, A new sequential optimality condition for constrained optimization and algorithmic consequences, *SIAM Journal on Optimization* 20, pp. 3533–3554, 2010.
- [11] M. Andretta, E. G. Birgin and J. M. Martínez, Practical active-set Euclidian trust-region method with spectral projected gradients for bound-constrained minimization, *Optimization* 54, pp. 305–325, 2005.
- [12] P. Armand and R. Omhenni, A globally and quadratically convergent primal-dual augmented Lagrangian algorithm for equality constrained optimization, *Optimization Methods and Software* 32, pp. 1–21, 2017.
- [13] P. Armand and R. Omhenni, A mixed logarithmic barrier-augmented Lagrangian method for nonlinear optimization, *Journal of Optimization Theory and Applications* 173, pp. 523–547, 2017.
- [14] E. G. Birgin, D. Fernández, and J. M. Martínez, On the boundedness of penalty parameters in an Augmented Lagrangian method with lower level constraints, *Optimization Methods and Software* 27, pp. 1001–1024, 2012.
- [15] E. G. Birgin, C. A. Floudas, and J. M. Martínez, Global minimization using an Augmented Lagrangian method with variable lower-level constraints, *Mathematical Programming* 125, pp. 139–162, 2010.
- [16] E. G. Birgin and J. M. Martínez, Large-scale active-set box-constrained optimization method with spectral projected gradients, *Computational Optimization and Applications* 23, pp. 101–125, 2002.
- [17] E. G. Birgin and J. M. Martínez, Structured minimal-memory inexact quasi-Newton method and secant preconditioners for Augmented Lagrangian Optimization, *Computational Optimization and Applications* 39, pp. 1–16, 2008.

- [18] E. G. Birgin and J. M. Martínez, Improving ultimate convergence of an Augmented Lagrangian method, *Optimization Methods and Software* 23, pp. 177–195, 2008.
- [19] E. G. Birgin and J. M. Martínez, Augmented Lagrangian method with nonmonotone penalty parameters for constrained optimization, *Computational Optimization and Applications* 51, pp. 941–965, 2012.
- [20] E. G. Birgin and J. M. Martínez, *Practical Augmented Lagrangian Methods for Constrained Optimization*, vol. 10 of Fundamentals of Algorithms, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2014. doi:10.1137/1.9781611973365.
- [21] E. G. Birgin, J. M. Martínez and M. Raydan, Nonmonotone spectral projected gradient methods on convex sets, *SIAM Journal on Optimization* 10, pp. 1196–1211, 2000.
- [22] E. G. Birgin, J. M. Martínez and M. Raydan, Algorithm 813: SPG - software for convex-constrained optimization, *ACM Transactions on Mathematical Software* 27, pp. 340–349, 2001.
- [23] E. G. Birgin, J. M. Martínez and M. Raydan, Spectral Projected Gradient methods: Review and Perspectives, *Journal of Statistical Software* 60, issue 3, 2014. doi:10.18637/jss.v060.i03.
- [24] C. Cartis, N. I. M. Gould, and Ph. L. Toint, On the evaluation complexity of composite function minimization with applications to nonconvex nonlinear programming, *SIAM Journal on Optimization* 21, pp. 1721–1739, 2011.
- [25] N. Chatzipanagiotis and M. M. Zavlanos, On the convergence of a distributed Augmented Lagrangian method for nonconvex optimization, *IEEE Transactions on Automatic Control* 62, pp. 4405–4420, 2017.
- [26] A. R. Conn, N. I. M. Gould, A. Sartenaer and Ph. 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, pp. 674–703, 1996.
- [27] A. R. Conn, N. I. M. Gould and Ph. L. Toint, *Lancelot: A Fortran package for large scale nonlinear optimization*, Springer-Verlag, Berlin, 1992.
- [28] F. E. Curtis, H. Jiang, and D. P. Robinson, An adaptive Augmented Lagrangian method for large-scale constrained optimization, *Mathematical Programming* 152, pp. 201–245, 2015.
- [29] F. E. Curtis, N. I. M. Gould, H. Jiang, and D. P. Robinson, Adaptive Augmented Lagrangian methods: Algorithms and practical numerical experience, *Optimization Methods & Software* 31, pp. 157–186, 2016.
- [30] E. D. Dolan and J. J. Moré, Benchmarking optimization software with performance profiles, *Mathematical Programming* 91, pp. 201–213, 2002.
- [31] Z. Dostál, Optimal Quadratic Programming Algorithms, volume 23 of *Optimization and its Applications*, Springer, New York, 2009.
- [32] Z. Dostál and P. Beremlijski, On convergence of inexact Augmented Lagrangians for separable and equality convex QCQP problems without constraint qualification, *Advances in Electrical and Electronic Engineering* 15, pp. 215–222, 2017.

- [33] D. Fernández and M. V. Solodov, Local convergence of exact and inexact augmented Lagrangian methods under the second-order sufficient optimality condition, *SIAM Journal on Optimization* 22, pp. 384–407, 2012.
- [34] A. V. Fiacco and G.P. McCormick, *Nonlinear Programming: Sequential Unconstrained Minimization Techniques*, John Wiley & Sons, New York, 1968.
- [35] R. Fletcher, *Practical Methods of Optimization*, Academic Press, London, 1987.
- [36] R. Fletcher, Augmented Lagrangians, box constrained QP and extensions, *IMA Journal of Numerical Analysis* 37, pp. 1635–1656, 2017.
- [37] N. I. M. Gould, D. Orban, and Ph. L. Toint, CUTEst: a constrained and unconstrained testing environment with safe threads for mathematical optimization, *Computational Optimization and Applications* 60, pp. 545–557, 2014.
- [38] G. N. Grapiglia and Y. Yuan, On the complexity of an Augmented Lagrangian method for nonconvex optimization, arXiv:1906.05622v1.
- [39] M. R. Hestenes, Multiplier and gradient methods, *Journal of Optimization Theory and Applications* 4, pp. 303–320, 1969.
- [40] A. F. Izmailov, M. V. Solodov, and E. I. Uskov, Global convergence of augmented Lagrangian methods applied to optimization problems with degenerate constraints, including problems with complementarity constraints, *SIAM Journal on Optimization* 22, pp. 1579–1606, 2012.
- [41] C. Kanzow and D. Steck, An example comparing the standard and safeguarded augmented Lagrangian methods, *Operations Research Letters* 45, pp. 598–603, 2017.
- [42] J. M. Martínez and L. Qi, Inexact Newton methods for solving nonsmooth equation, *Journal of Computational and Applied Mathematics* 60, pp. 127–145, 1995.
- [43] M. J. D. Powell, A method for nonlinear constraints in minimization problems, in *Optimization*, R. Fletcher (ed.), Academic Press, New York, NY, pp. 283–298, 1969.
- [44] L. Qi and J. Sun, A nonsmooth version of a Newton’s method, *Mathematical Programming* 58, pp. 353–367, 1993.
- [45] R. T. Rockafellar, Augmented Lagrange multiplier functions and duality in nonconvex programming, *SIAM Journal on Control and Optimization* 12, pp. 268–285, 1974.
- [46] W. Sun and Y. Yuan, *Optimization Theory and Methods: Nonlinear Programming*, Springer, Berlin, 2006.
- [47] A. Wächter and L. T. Biegler, On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming, *Mathematical Programming* 106, pp. 25–57, 2006.
- [48] *HSL. A collection of fortran codes for large scale scientific computation*, <http://www.hsl.rl.ac.uk/>.