




EVALUATION COMPLEXITY FOR NONLINEAR CONSTRAINED
OPTIMIZATION USING UNSCALED KKT CONDITIONS
AND HIGH-ORDER MODELS

by E. G. Birgin, J. L. Gardenghi, J. M. Martínez,
S. A. Santos and Ph. L. Toint

Report NAXYS-08-2015

10 July 2015



University of Namur, 61, rue de Bruxelles, B5000 Namur (Belgium)

<http://www.unamur.be/sciences/naxys>

Evaluation complexity for nonlinear constrained optimization using unscaled KKT conditions and high-order models*

E. G. Birgin[†] J. L. Gardenghi[†] J. M. Martínez[‡] S. A. Santos[‡]

Ph. L. Toint[§]

10 July 2015

Abstract

The evaluation complexity of general nonlinear, possibly nonconvex, constrained optimization is analyzed. It is shown that, under suitable smoothness conditions, an ϵ -approximate first-order critical point of the problem can be computed in order $O(\epsilon^{1-2(p+1)/p})$ evaluations of the problem's function and their first p derivatives. This is achieved by using a two-phases algorithm inspired by Cartis, Gould, and Toint [8, 11]. It is also shown that strong guarantees (in terms of handling degeneracies) on the possible limit points of the sequence of iterates generated by this algorithm can be obtained at the cost of increased complexity. At variance with previous results, the ϵ -approximate first-order criticality is defined by satisfying a version of the KKT conditions with an accuracy that does not depend on the size of the Lagrange multipliers.

Key words: Nonlinear programming, complexity, approximate KKT point.

1 Introduction

Complexity analysis of numerical nonlinear optimization is currently an active research area (see, for instance, [17, 18, 3, 9, 6, 8, 12, 15, 21, 14]). In this domain, the worst-case (first-order) evaluation complexity of general smooth nonlinear optimization, that is the maximal number of evaluations of the problem's objective function, constraints, and their derivatives

*This work has been partially supported by the Brazilian agencies FAPESP (grants 2010/10133-0, 2013/03447-6, 2013/05475-7, 2013/07375-0, and 2013/23494-9) and CNPq (grants 304032/2010-7, 309517/2014-1, 303750/2014-6, and 490326/2013-7) and by the Belgian National Fund for Scientific Research (FNRS).

[†]Department of Computer Science, Institute of Mathematics and Statistics, University of São Paulo, Rua do Matão, 1010, Cidade Universitária, 05508-090, São Paulo, SP, Brazil. e-mail: {egbirgin | john}@ime.usp.br

[‡]Department of Applied Mathematics, Institute of Mathematics, Statistics, and Scientific Computing, University of Campinas, Campinas, SP, Brazil. e-mail: {martinez | sandra}@ime.unicamp.br

[§]Department of Mathematics, The University of Namur, 61, rue de Bruxelles, B-5000, Namur, Belgium. e-mail: philippe.toint@unamur.be

that is needed for obtaining an approximate first-order critical point, has been the subject of recent papers by Cartis, Gould, and Toint [8, 12, 11]. In the first two of these contributions, it is shown that a two-phases trust-region based algorithm needs at most $O(\epsilon^{-2})$ evaluations of these functions (and their gradients) to compute an ϵ -approximate first-order solution, that is either an ϵ -approximate scaled first-order critical point of the problem or, as expected barring global optimization, an infeasible approximate critical point of the constraints' violation. By “ ϵ -approximate scaled first-order critical point”, we mean a point satisfying the first-order Karush-Kuhn-Tucker (KKT) condition on Lagrange multipliers up to an accuracy which is proportional to the size of the multipliers. The use of second derivatives was subsequently investigated in [11], where it was shown that a similar two-phases algorithm needs at most $O(\epsilon^{-3/2})$ evaluations of the problem functions, gradients, and Hessians to compute a point satisfying similar conditions.

The purpose of this paper is to extend these results in two different ways. The first is to consider unscaled KKT conditions (i.e. where the size of the Lagrange multipliers does not appear explicitly in the accuracy of the approximate criticality condition) and the second is to examine what can be achieved if evaluation of derivatives up to order $p > 2$ is allowed. We show below that a two-phases algorithm needs a maximum number of evaluations of the problem's functions and derivatives up to order p ranging from $O(\epsilon^{1-2(p+1)/p})$ to $O(\epsilon^{1-3(p+1)/p})$ to produce an ϵ -approximate *unscaled* first-order critical point of the problem (or an infeasible approximate critical point of the constraints' violation), depending on the identified degeneracy level. The extension of the theory to arbitrary p finds its basis in a proposal [4] by the authors of the present paper which extends to high order the available evaluation complexity results for unconstrained optimization (see [10, 17, 18]).

The paper is organised as follows. Section 2 states the problem more formally, describes a class of algorithms for its solution, and discusses the proposed termination criteria. The convergence and worst-case evaluation complexity analysis is presented in Section 3 and the complexity results are further discussed in Section 4. Conclusions are finally outlined in Section 5.

Notations: In what follows, $\|\cdot\|$ denotes the Euclidean norm and, if $v(x)$ is a vector function, $v(x)_+ \stackrel{\text{def}}{=} \max[v(x), 0]$ where the maximum is taken componentwise. $\nabla v(x)$ will denote the gradient of a function v defined on \mathbb{R}^n with respect to its variable x . The notation $[x]_j$ denotes the j -th component of a vector x whenever the simpler notation x_j might lead to confusion.

2 The problem and a class of algorithms for its solution

We consider the optimization problem given by:

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & f(x) \\ \text{s.t.} \quad & c_{\mathcal{E}}(x) = 0, \\ & c_{\mathcal{I}}(x) \leq 0, \end{aligned} \tag{2.1}$$

where f is a function from \mathbb{R}^n into \mathbb{R} and $c_{\mathcal{E}}$ and $c_{\mathcal{I}}$ are functions from \mathbb{R}^n to \mathbb{R}^m and \mathbb{R}^q , respectively. We will assume that all these functions are p times continuously differentiable. We

now define, for all $x \in \mathbb{R}^n$, the infeasibility measure

$$\theta(x) \stackrel{\text{def}}{=} \|c_{\mathcal{E}}(x)\|^2 + \|c_{\mathcal{I}}(x)_+\|^2, \quad (2.2)$$

and, for given $v > 0$, its associated level set

$$\mathcal{L}(v) \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \mid \theta(x) \leq v\}.$$

Moreover, given $t \in \mathbb{R}$, we define

$$\Phi(x, t) \stackrel{\text{def}}{=} \theta(x) + [f(x) - t]_+^2 \text{ for all } x \in \mathbb{R}^n,$$

where the scalar t is the *target*. We use the notation $\nabla\Phi(x, t)$ to denote the gradient of $\Phi(x, t)$ with respect to x .

The FTARGET (Feasibility and Target following) algorithm defined on the following page is inspired by that proposed in [11] and computes a sequence of iterates x_k by means of *outer iterations* indexed by $k = 0, 1, 2, \dots$. For obtaining the outer iterate x_{k+1} , the FTARGET algorithm uses a given (inner) Unconstrained Minimization (UM) algorithm which computes *inner iterations* indexed by $j = 0, 1, 2, \dots$ by using derivatives of its objective function up to order p . Consequently, inner iterates will be denoted by $x_{k,j}$.

It is clear that the FTARGET algorithm is in fact a class of algorithms depending on the specific choices of the derivative order $p \geq 1$ and on the UM minimizer adapted to this choice.

2.1 The meaning of the stopping criteria

We now discuss the nature of the point returned by the FTARGET algorithm as a function of the stopping criterion activated. We start by outlining the main points, leaving a detailed discussion for the following subsections.

Stopping at Step T4 means that an *approximate KKT* (AKKT) point has been found. Stopping at Step F2 has two possible meanings: (i) an infeasible stationary point of the infeasibility may be identified in the limit; or (ii) a situation analogous to the one represented by stopping at Step T3 has happened. The interpretation of stopping at Step T3 depends on a weak condition on the tolerances ϵ_p and ϵ_D in the limit and, more significantly, on the choice of the function ψ . For three different choices of ψ , it will be shown that stopping at Step T3 means that: (i) an ϵ_p -feasible point z has been found such that the gradients of active constraints at z are not uniformly linear independent (with non-negative coefficients for the inequality constraints); (ii) a feasible point which does not satisfy the Mangasarian-Fromowitz constraint qualification exists in the limit; or (iii) a feasible point which does not satisfy the Łojasiewicz inequality exists in the limit. (These claims are precisely discussed below.) We may therefore conclude globally that, for the three considered choices of ψ , FTARGET always finds an unscaled approximate KKT point under suitable “non-degeneracy” assumptions.

Algorithm 2.1: The FTARGET algorithm

Input: Let $\epsilon_P \in (0, 1]$, a primal accuracy threshold, and $\epsilon_D \in (0, 1]$, a dual one, be given. Let $\rho \in (0, 1)$, an $x_{-1} \in \mathbb{R}^n$. Let $\psi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be a continuous and non-decreasing function such that $\psi(0) = 0$. Let p be the order of available derivatives of the functions f , $c_{\mathcal{E}}$, and $c_{\mathcal{I}}$.

PHASE 1: *Computing an approximately feasible initial guess*

Step F1. Minimize $\theta(x)$ using the UM algorithm starting from $x_{-1,0} = x_{-1}$, to compute an iterate $x_{-1,j}$, $j \in \{0, 1, 2, \dots\}$, such that $\theta(x_{-1,j}) \leq \theta(x_{-1,0})$ and such that $x_{-1,j}$ satisfies at least one of the following conditions:

$$\theta(x_{-1,j}) \leq 0.99 \epsilon_P^2, \quad (2.3)$$

$$\theta(x_{-1,j}) > 0.99 \epsilon_P^2 \quad \text{and} \quad \|\nabla\theta(x_{-1,j})\| \leq \psi(\epsilon_D). \quad (2.4)$$

Define $x_0 = x_{-1,j}$.

Step F2. If $\theta(x_0) > 0.99 \epsilon_P^2$ and $\|\nabla\theta(x_0)\| \leq \psi(\epsilon_D)$, **stop** returning x_0 .

PHASE 2: *Improving dual feasibility (target following)*

Step T0. Initialize $k \leftarrow 0$.

Step T1. Compute $t_k = f(x_k) - \sqrt{\epsilon_P^2 - \theta(x_k)}$.

Step T2. Minimize $\Phi(x, t_k)$ using the UM algorithm starting from $x_{k,0} = x_k$, to compute an iterate $x_{k,j}$, $j \in \{0, 1, 2, \dots\}$, such that $\Phi(x_{k,j}, t_k) \leq \Phi(x_{k,0}, t_k) = \epsilon_P^2$ and such that $x_{k,j}$ satisfies at least one of the following conditions:

$$f(x_{k,j}) \leq t_k + \rho(f(x_k) - t_k) \quad \text{and} \quad \theta(x_{k,j}) \leq 0.99 \epsilon_P^2, \quad (2.5)$$

$$f(x_{k,j}) > t_k \quad \text{and} \quad \|\nabla\Phi(x_{k,j}, t_k)\| \leq 2\epsilon_D[f(x_{k,j}) - t_k]_+, \quad (2.6)$$

$$\theta(x_{k,j}) > 0.99 \epsilon_P^2 \quad \text{and} \quad \|\nabla\theta(x_{k,j})\| \leq \psi(\epsilon_D). \quad (2.7)$$

Define $x_{k+1} = x_{k,j}$.

Step T3. If $\theta(x_{k+1}) > 0.99 \epsilon_P^2$ and $\|\nabla\theta(x_{k+1})\| \leq \psi(\epsilon_D)$, **stop** returning x_{k+1} .

Step T4. If $f(x_{k+1}) > t_k$ and $\|\nabla\Phi(x_{k+1}, t_k)\| \leq 2\epsilon_D[f(x_{k+1}) - t_k]_+$, **stop** returning x_{k+1} .

Step T5. Set $k \leftarrow k + 1$, and go to Step T1.

2.1.1 Terminating at Step T4

When the FTARGET algorithm stops at iteration k because the stopping criterion is satisfied at Step T4, it returns x_{k+1} such that

$$\begin{aligned} \sqrt{\theta(x_{k+1})} &\leq \epsilon_P, \\ \left\| \nabla f(x_{k+1}) + \sum_{j=1}^m \lambda_j [\nabla c_{\mathcal{E}}(x_{k+1})]_j + \sum_{j=1}^q \mu_j [\nabla c_{\mathcal{I}}(x_{k+1})]_j \right\| &\leq \epsilon_D, \end{aligned} \quad (2.8)$$

where

$$\lambda_j = \frac{[c_{\mathcal{E}}(x_{k+1})]_j}{f(x_{k+1}) - t_k} \text{ for } j = 1, \dots, m \quad \text{and} \quad \mu_j = \frac{[c_{\mathcal{I}}(x_{k+1})]_j}{f(x_{k+1}) - t_k} \text{ for } j = 1, \dots, q.$$

Note also that $f(x_{k+1}) - t_k > 0$ in the two expressions above and, hence, that the multipliers λ_j and μ_j are well defined. We thus have that if $[c_{\mathcal{I}}(x_{k+1})]_j \leq 0$ then $\mu_j = 0$ and, hence,

$$\min[\mu_j, -[c_{\mathcal{I}}(x_{k+1})]_j] = \mu_j = 0.$$

On the other hand, if $[c_{\mathcal{I}}(x_{k+1})]_j > 0$ then

$$\min[\mu_j, -[c_{\mathcal{I}}(x_{k+1})]_j] = -[c_{\mathcal{I}}(x_{k+1})]_j \leq \epsilon_P. \quad (2.9)$$

We may then conclude that complementarity (as measured with the min function) is satisfied with precision ϵ_P . Note that the accuracies in the right-hand side of the second inequality in (2.8) and in (2.9) do not involve the size of the Lagrange multipliers, in contrast with the termination rule used in [11, 12].

In asymptotic terms, if we assume that the FTARGET algorithm is run infinitely many times with $\epsilon_P = \epsilon_{P,\ell} \rightarrow 0$ and $\epsilon_D = \epsilon_{D,\ell} \rightarrow 0$, and that it stops infinitely many times (ℓ_1, ℓ_2, \dots) returning a point z_{ℓ_j} that satisfies the stopping criterion at Step T4, then we have that any accumulation point z of the sequence $\{z_{\ell_j}\}$ is, by definition, an AKKT [1, 2, 19] point.

2.1.2 Terminating at Step F2 or Step T3

Some explanations regarding the choice and interpretation of the function ψ are now in order. Note that ψ is used in the algorithm to stop the execution (at Steps F2 or T3) when $\theta(x)$ is “large” (i.e. $\theta(x) \geq 0.99 \epsilon_P^2$) and $\nabla \theta(x)$ is “small” (i.e. $\|\nabla \theta(x)\| \leq \psi(\epsilon_D)$). At a first glance, these occurrences seem to indicate that x is an “approximate infeasible stationary point of $\theta(x)$ ”. However, a more careful analysis of particular cases reveals some subtleties.

When the FTARGET algorithm stops at Step F2, it returns a point z such that

$$\theta(z) > 0.99 \epsilon_P^2 \quad \text{and} \quad \|\nabla \theta(z)\| \leq \psi(\epsilon_D). \quad (2.10)$$

The analysis of these conditions is better done in asymptotic terms. Therefore, assume that the FTARGET algorithm is executed infinitely many times with $\epsilon_P = \epsilon_{P,\ell} \rightarrow 0$ and $\epsilon_D = \epsilon_{D,\ell} \rightarrow 0$.

Moreover, assume that it stops infinitely many times ℓ_1, ℓ_2, \dots at Step F2 returning z_{ℓ_j} ($j = 1, 2, \dots$) such that $\lim_{j \rightarrow \infty} z_{\ell_j} = z$. Clearly, by (2.10), we have that

$$\lim_{j \rightarrow \infty} \|\nabla\theta(z_{\ell_j})\| = 0. \quad (2.11)$$

If, in addition, $\theta(z_{\ell_j})$ is bounded away from zero, by (2.11), then stopping at Step F2 means that the infeasible point z is a stationary point of the infeasibility measure θ . If $\lim_{j \rightarrow \infty} \theta(z_{\ell_j}) = 0$ then the interpretation of this fact depends on the choice of function ψ and follows exactly the same analysis that we now investigate for the case where the FTARGET algorithm stops at Step T3.

When the FTARGET algorithm stops at Step T3, it returns a point z such that

$$0.99 \epsilon_P^2 < \theta(z) \leq \epsilon_P^2 \quad \text{and} \quad \|\nabla\theta(z)\| \leq \psi(\epsilon_D), \quad (2.12)$$

whose interpretation clearly depends on the choice of ψ . We now consider functions ψ of the form

$$\psi(\epsilon_D) = \sigma_1 \epsilon_D^{\sigma_2} \quad (2.13)$$

with $\sigma_1 > 0$ and the three possible choices: **(a)** $\sigma_2 = 1$, **(b)** $\sigma_2 \in (1, 2)$, and **(c)** $\sigma_2 = 2$.

In case **(a)**, (2.12) implies that

$$\frac{\|\nabla\theta(z)\|}{\sqrt{\theta(z)}} < \frac{\sigma_1}{\sqrt{0.99}} \frac{\epsilon_D}{\epsilon_P}. \quad (2.14)$$

Note that, by the definition (2.2) of θ ,

$$\frac{\|\nabla\theta(z)\|}{\sqrt{\theta(z)}} = 2 \left\| \sum_{i=1}^m \lambda_i [\nabla c_{\mathcal{E}}(z)]_i + \sum_{j=1}^q \mu_j [\nabla c_{\mathcal{I}}(z)]_j \right\|, \quad (2.15)$$

where

$$\lambda_i = \frac{[c_{\mathcal{E}}(z)]_i}{\left(\sum_{i=1}^m [c_{\mathcal{E}}(z)]_i^2 + \sum_{j=1}^q [c_{\mathcal{I}}(z)]_j^2 \right)^{1/2}} \quad \text{for } i = 1, \dots, m \quad (2.16)$$

and

$$\mu_j = \frac{[c_{\mathcal{I}}(z)]_j}{\left(\sum_{i=1}^m [c_{\mathcal{E}}(z)]_i^2 + \sum_{j=1}^q [c_{\mathcal{I}}(z)]_j^2 \right)^{1/2}} \quad \text{for } j = 1, \dots, q. \quad (2.17)$$

Moreover, by (2.16) and (2.17), we have that

$$\sum_{i=1}^m \lambda_i^2 + \sum_{j=1}^q \mu_j^2 = 1 \quad \text{and} \quad \mu_j = 0 \text{ whenever } [c_{\mathcal{I}}(z)]_j < 0. \quad (2.18)$$

A definition is now necessary.

Definition 1 Given $\xi > 0$, we say that x has ξ -uniformly positive linear independent gradients with respect to problem (2.1) if for all $(\lambda, \mu) \in \mathbb{R}^m \times \mathbb{R}_+^q$ such that $\sum_{i=1}^m \lambda_i^2 + \sum_{j=1}^q \mu_j^2 = 1$ and $\mu_j = 0$ whenever $[c_{\mathcal{I}}(x)]_j < 0$ we have that

$$\left\| \sum_{i=1}^m \lambda_i [\nabla c_{\mathcal{E}}(x)]_i + \sum_{j=1}^q \mu_j [\nabla c_{\mathcal{I}}(x)]_j \right\| \geq \xi.$$

This means that, by (2.14), the point z returned by FTARGET has not ξ -uniformly positive linear independent gradients of active constraints with

$$\xi = \frac{\sigma_1}{2\sqrt{0.99}} \frac{\epsilon_{\mathcal{D}}}{\epsilon_{\mathcal{P}}}.$$

If $\sigma_1 \epsilon_{\mathcal{D}} / \epsilon_{\mathcal{P}}$ is small enough, this indicates that an approximate Mangasarian-Fromovitz constraint qualification with tolerance ξ is not satisfied at z .

The analysis of cases **(b)** and **(c)** must also be done in asymptotic terms. Therefore, assume again that the FTARGET algorithm is executed infinitely many times with $\epsilon_{\mathcal{P}} = \epsilon_{\mathcal{P},\ell} \rightarrow 0$ and $\epsilon_{\mathcal{D}} = \epsilon_{\mathcal{D},\ell} \rightarrow 0$. Moreover, assume that it stops infinitely many times ℓ_1, ℓ_2, \dots at Step T3 returning z_{ℓ_j} ($j = 1, 2, \dots$) such that $\lim_{j \rightarrow \infty} z_{\ell_j} = z$. Let us assume, for the remainder of this section, that

$$\omega \stackrel{\text{def}}{=} \limsup_{j \rightarrow \infty} \frac{\epsilon_{\mathcal{D},\ell_j}}{\epsilon_{\mathcal{P},\ell_j}} < \infty. \quad (2.19)$$

In cases **(b)** and **(c)**, by (2.12), we have that

$$\frac{\|\nabla \theta(z_{\ell_j})\|}{\sqrt{\theta(z_{\ell_j})}} < \frac{\sigma_1}{\sqrt{0.99}} \frac{\epsilon_{\mathcal{D},\ell_j}^{\sigma_2}}{\epsilon_{\mathcal{P},\ell_j}} \quad \text{for } j = 1, 2, \dots \quad (2.20)$$

Taking limits in (2.20), and using (2.19) and the fact that $\sigma_2 > 1$, we have, by (2.15–2.18), that the gradients of active constraints at z are not positively linearly independent, so the Mangasarian-Fromowitz constraint qualification does not hold at the feasible point z [20].

As a consequence, if the feasible set is compact and all the feasible points satisfy the Mangasarian-Fromowitz constraint qualification then, for j large enough, stopping at T3 is impossible in view of (2.19). Therefore, Phase 2 of the FTARGET algorithm can only stop at T4 with an AKKT point.

In order to further analyze case **(c)**, we define the Łojasiewicz [16] inequality.

Definition 2 A continuously differentiable function $v : \mathbb{R}^n \rightarrow \mathbb{R}$ satisfies the Łojasiewicz inequality at \bar{x} if there exist $\delta > 0$, $\tau \in (0, 1)$, and $\kappa > 0$ such that, for all $x \in B(\bar{x}, \delta)$,

$$|v(x) - v(\bar{x})|^\tau \leq \kappa \|\nabla v(x)\|. \quad (2.21)$$

The properties of functions that satisfy the inequality (2.21) have been studied in several recent papers in connection with minimization methods, complexity theory, asymptotic analysis of partial differential equations, tame optimization, and the fulfillment of AKKT conditions in augmented Lagrangian methods [2, 5]. Smooth functions satisfy this inequality under fairly weak conditions. For example, analytic functions satisfy the Łojasiewicz inequality.

In case **(c)**, by (2.12), we have that

$$\frac{\|\nabla\theta(z_{\ell_j})\|}{\theta(z_{\ell_j})} < \frac{\sigma_1}{0.99} \left(\frac{\epsilon_{D,\ell_j}}{\epsilon_{P,\ell_j}} \right)^2 \quad \text{for } j = 1, 2, \dots$$

We therefore obtain that, for arbitrary $\tau \in (0, 1)$,

$$\theta(z_{\ell_j})^\tau = \theta(z_{\ell_j})^{\tau-1} \theta(z_{\ell_j}) > \theta(z_{\ell_j})^{\tau-1} \left[\frac{\sigma_1}{0.99} \left(\frac{\epsilon_{D,\ell_j}}{\epsilon_{P,\ell_j}} \right)^2 \right]^{-1} \|\nabla\theta(z_{\ell_j})\|.$$

Since, by (2.12), $\theta(z_{\ell_j}) \leq \epsilon_{P,\ell_j}^2$, then $\theta(z_{\ell_j}) \rightarrow 0$ when j tends to infinity. As a consequence, $\theta(z_{\ell_j})^{\tau-1}$ tends to infinity because $\tau \in (0, 1)$, and thus, using (2.19), for every $\delta > 0$ and $\kappa > 0$, there exists j sufficiently large such that

$$z_{\ell_j} \in \mathcal{B}(z, \delta) \quad \text{and} \quad \theta(z_{\ell_j})^\tau > \kappa \|\nabla\theta(z_{\ell_j})\|.$$

This implies that the function $\theta(\cdot)$ does not satisfy the Łojasiewicz inequality at any possible limit point z . Therefore, if we assume that the function $\theta(\cdot)$ satisfies the Łojasiewicz inequality at every feasible point and that j is large enough, then, in view of (2.19), Phase 2 of the FTARGET algorithm can only stop at Step T4 returning an AKKT point.

3 Finite Termination, Convergence and Complexity

In this section we will prove that the FTARGET algorithm is well defined and terminates in a finite number of iterations, provided that the UM algorithm employed to minimize $\theta(x)$ at Phase 1 and $\Phi(x, t_k)$ at Phase 2 possesses standard convergence properties. Moreover, if the UM algorithm also enjoys suitable evaluation complexity properties, we can establish complexity bounds for the FTARGET algorithm itself.

Assumption A1 *The function f is bounded below on the set $\mathcal{L}(\epsilon_P^2)$ in that there exists a constant f_{low} such that $f(x) \geq f_{\text{low}}$ for all $x \in \mathcal{L}(\epsilon_P^2)$.*

Assumption A2 *$\|\nabla f\|$ is bounded above on $\mathcal{L}(\epsilon_P^2)$ in that there exists a constant κ_∇ such that $\|\nabla f(x)\| \leq \kappa_\nabla$ for all $x \in \mathcal{L}(\epsilon_P^2)$.*

In order to simplify notation below, we also assume, without loss of generality, that

$$\kappa_\nabla \geq \max \left[1, \frac{5\sigma_1}{2\rho} \right]. \quad (3.22)$$

Assumption A1 is enough to prove a first result that is essential for further analysis of convergence and complexity. Lemma 3.1 shows that, independently of the UM algorithm used for unconstrained minimizations, the number of outer iterations computed by the FTARGET algorithm cannot exceed a bound that only depends on ϵ_P and the lower bound of f .

Lemma 3.1 Suppose that Assumption A1 holds. Then the FTARGET algorithm stops at Phase 1 or it performs, at most,

$$\left\lfloor \frac{f(x_0) - f_{\text{low}}}{0.1(1 - \rho)\epsilon_P} \right\rfloor + 1$$

outer iterations at Phase 2.

Proof. Assume that the FTARGET algorithm does not stop at Phase 1. By the definition of this algorithm, an outer iteration is launched only when $\theta(x_k) \leq 0.99\epsilon_P^2$ and is followed by other outer iteration only when

$$f(x_{k+1}) \leq t_k + \rho(f(x_k) - t_k) \quad \text{and} \quad \theta(x_{k+1}) \leq 0.99\epsilon_P^2.$$

By the definition of t_k this implies that

$$\begin{aligned} f(x_{k+1}) &\leq t_k + \rho\sqrt{\epsilon_P^2 - \theta(x_k)} \\ &= f(x_k) - \sqrt{\epsilon_P^2 - \theta(x_k)} + \rho\sqrt{\epsilon_P^2 - \theta(x_k)} \\ &= f(x_k) - (1 - \rho)\sqrt{\epsilon_P^2 - \theta(x_k)}. \end{aligned}$$

Thus, since $\theta(x_k) \leq 0.99\epsilon_P^2$,

$$f(x_{k+1}) \leq f(x_k) - (1 - \rho)\sqrt{\epsilon_P^2 - 0.99\epsilon_P^2} = f(x_k) - 0.1(1 - \rho)\epsilon_P.$$

In other words, when an outer iteration is completed satisfying (2.5) we obtain a decrease of at least $0.1(1 - \rho)\epsilon_P$ in the objective function. Then, since $f(x) \geq f_{\text{low}}$ for all $x \in \mathcal{L}(\epsilon_P^2)$ by Assumption A1, the number of outer iterations that are completed satisfying (2.5) must be smaller than or equal to

$$\left\lfloor \frac{f(x_0) - f_{\text{low}}}{0.1(1 - \rho)\epsilon_P} \right\rfloor.$$

We therefore obtain the desired result by adding a final iteration at which (2.5) may not hold. \square

Observe that, according to this lemma, the number of outer iterations performed by the FTARGET algorithm could be only one, a case that occurs if ϵ_P is sufficiently large and, consequently, the first target t_0 is very low.

Now we need to prove that, once an outer iteration is launched, it can be completed in a finite number of inner iterations. The following lemma establishes that, if $\theta(x_k) \leq 0.99\epsilon_P^2$ and the UM

algorithm does not stop at $x_{k,j}$, the gradient norm $\|\nabla\Phi(x_{k,j}, t_k)\|$ is bounded away from zero by a quantity that only depends on ϵ_P , ϵ_D , and a bound on the norm of the gradient of f on $\mathcal{L}(\epsilon_P^2)$.

Lemma 3.2 Suppose that Assumption A2 holds, that $\theta(x_k) \leq 0.99\epsilon_P^2$, and that the UM algorithm used for minimizing $\Phi(x, t_k)$ does not stop at the inner iterate $x_{k,j}$. Then,

$$\|\nabla\Phi(x_{k,j}, t_k)\| \geq \min \left[0.2\rho\epsilon_D\epsilon_P, \frac{\psi(\epsilon_D)}{2}, \frac{\psi(\epsilon_D)\epsilon_D}{2\kappa_\nabla} \right]. \quad (3.23)$$

Proof. Since the UM algorithm does not stop at $x_{k,j}$ we have that none of the conditions (2.5), (2.6), and (2.7) hold at $x_{k,j}$. We will consider two cases:

$$f(x_{k,j}) > t_k + \rho(f(x_k) - t_k) \quad (3.24)$$

and

$$f(x_{k,j}) \leq t_k + \rho(f(x_k) - t_k). \quad (3.25)$$

Consider the case (3.24) first. Then, $f(x_{k,j}) > t_k$. Since (2.6) does not hold, we have that

$$\|\nabla\Phi(x_{k,j}, t_k)\| > 2\epsilon_D(f(x_{k,j}) - t_k). \quad (3.26)$$

Now, by (3.24) and the definition of t_k , since $\theta(x_k) \leq 0.99\epsilon_P^2$,

$$f(x_{k,j}) - t_k > \rho(f(x_k) - t_k) = \rho\sqrt{\epsilon_P^2 - \theta(x_k)} \geq 0.1\rho\epsilon_P.$$

Therefore, by (3.26),

$$\|\nabla\Phi(x_{k,j}, t_k)\| > 0.2\rho\epsilon_D\epsilon_P. \quad (3.27)$$

Now consider the case in which (3.25) holds. Then, $\theta(x_{k,j}) > 0.99\epsilon_P^2$, otherwise (2.5) would have been satisfied. Thus, since (2.7) does not hold, we also have that

$$\begin{aligned} \psi(\epsilon_D) &< \|\nabla\theta(x_{k,j})\| \\ &= \|\nabla\theta(x_{k,j}) + 2(f(x_{k,j}) - t_k)_+\nabla f(x_{k,j}) - 2(f(x_{k,j}) - t_k)_+\nabla f(x_{k,j})\| \\ &\leq \|\nabla\theta(x_{k,j}) + 2(f(x_{k,j}) - t_k)_+\nabla f(x_{k,j})\| + 2\|\nabla f(x_{k,j})\|(f(x_{k,j}) - t_k)_+ \\ &\leq \|\nabla\Phi(x_{k,j}, t_k)\| + 2\kappa_\nabla(f(x_{k,j}) - t_k)_+, \end{aligned}$$

where we have used Assumption A2 to derive the last inequality, yielding that

$$\|\nabla\Phi(x_{k,j}, t_k)\| > \psi(\epsilon_D) - 2\kappa_\nabla(f(x_{k,j}) - t_k)_+. \quad (3.28)$$

We now consider two cases. In the first one,

$$2\kappa_\nabla(f(x_{k,j}) - t_k)_+ \leq \frac{\psi(\epsilon_D)}{2}. \quad (3.29)$$

Then, by (3.28),

$$\|\nabla\Phi(x_{k,j}, t_k)\| \geq \frac{\psi(\epsilon_D)}{2}.$$

In the second case, i.e. if (3.29) is not true, we have that

$$2(f(x_{k,j}) - t_k)_+ = 2(f(x_{k,j}) - t_k) \geq \frac{\psi(\epsilon_D)}{2\kappa_\nabla}.$$

Then, since (2.6) does not hold,

$$\|\nabla\Phi(x_{k,j}, t_k)\| \geq 2\epsilon_D(f(x_{k,j}) - t_k)_+ \geq \frac{\psi(\epsilon_D)\epsilon_D}{2\kappa_\nabla}.$$

This completes the proof. \square

The following assumption expresses that the UM algorithm enjoys sensible first-order convergence properties in that the sequence of iterates $\{x_j\}$ it generates when applied to the minimization of a bounded-below smooth function contains a subsequence at which the gradient of this function converges to zero. We formulate this assumption in a way suitable for its application to the convergence of the FTARGET algorithm.

Assumption A3 *For an arbitrary $\varepsilon > 0$, one has that, if the UM algorithm is applied to the minimization of $\theta(x)$ or $\Phi(x, t)$ with respect to x , starting from an arbitrary initial point x_0 , then, in a finite number of iterations, this algorithm finds an iterate x such that either $\theta(x) \leq \theta(x_0)$ and $\|\nabla\theta(x)\| \leq \varepsilon$, or $\Phi(x, t) \leq \Phi(x_0, t)$ and $\|\nabla\Phi(x, t)\| \leq \varepsilon$, respectively.*

The following lemma establishes that, under Assumptions A2–A3 and given the iterate x_k of the FTARGET algorithm, the iterate x_{k+1} is well defined and it is computed in a finite number of iterations.

Lemma 3.3 Suppose that Assumptions A2–A3 hold. Then, for all $k = 0, 1, 2, \dots$, the UM algorithm applied to the minimization of $\Phi(x, t_k)$ finds a point $x_{k,j}$ that satisfies at least one of the criteria (2.5), (2.6), and (2.7) in a finite number of iterations.

Proof. Define

$$\varepsilon = \min \left[0.2\rho\epsilon_D\epsilon_P, \frac{\psi(\epsilon_D)}{2}, \frac{\psi(\epsilon_D)\epsilon_D}{2\kappa_\nabla} \right] > 0.$$

By Assumption A3, the UM algorithm eventually finds an iteration j_0 such that x_{k,j_0} satisfies $\|\nabla\Phi(x_{k,j_0}, t_k)\| \leq \varepsilon$. By Lemma 3.2 and the definition of ε , x_{k,j_0} satisfies (2.5), (2.6), or (2.7) and the iteration k of the FTARGET algorithm terminates at some $x_{k,j}$ with $j \leq j_0$. \square

Now we are ready to prove finite termination of the FTARGET algorithm.

Theorem 3.4 Suppose that Assumptions A1–A3 hold. Then, the FTARGET algorithm stops at Step F2 of Phase 1 or after at most

$$\left\lceil \frac{f(x_0) - f_{\text{low}}}{0.1(1 - \rho)\epsilon_P} \right\rceil + 1 \quad (3.30)$$

iterations of Phase 2 with x_{k+1} satisfying T3 or T4.

Proof. By Assumption A3, the problem at Phase 1 is solved by the UM algorithm in a finite number of iterations. By Lemma 3.1 the FTARGET algorithm cannot perform more than (3.30) outer iterations. By Lemma 3.3 each outer iteration is well-defined and terminates in finite time. Therefore, the FTARGET algorithm stops at Step F2 of Phase 1 or, at the last outer iteration of Phase 2, x_{k+1} satisfies T3 or T4. \square

The following assumption prescribes an additional property of the UM algorithm. Whereas Assumption A3 says that, for any $\varepsilon > 0$, the UM algorithm finds a point that verifies $\|\nabla\theta(x)\| \leq \varepsilon$ or $\|\nabla\Phi(x, t)\| \leq \varepsilon$ in a finite number of iterations, Assumption A4 aims to quantify the number of function evaluations that are necessary to achieve a sufficiently small gradient.

Assumption A4 *There exists $\alpha \geq 0$ and a constant $\kappa_\theta > 0$ (depending on properties of the functions $c_\mathcal{E}$ and $c_\mathcal{I}$ and on parameters of the UM algorithm) such that, given $\varepsilon > 0$, if the UM algorithm is applied to the minimization of $\theta(x)$ starting from an arbitrary initial point x_{-1} , the algorithm finds an iterate x such that $\theta(x) \leq \theta(x_{-1})$ and $\|\nabla\theta(x)\| \leq \varepsilon$ employing, at most,*

$$\kappa_\theta \left\lceil \frac{\theta(x_{-1})}{\varepsilon^\alpha} \right\rceil$$

evaluations of $c_\mathcal{E}$, $c_\mathcal{I}$, and their derivatives. Moreover, if, for any $t \in \mathbb{R}$, the UM algorithm is applied to the minimization of $\Phi(x, t)$ (with respect to x) starting from an arbitrary initial point x_0 , there exists a constant $\kappa_\Phi > 0$ (depending on properties of the functions f , $c_\mathcal{E}$, and $c_\mathcal{I}$ and on parameters of the UM algorithm) such that the algorithm finds an iterate x such that $\Phi(x, t) \leq \Phi(x_0, t)$ and $\|\nabla\Phi(x, t)\| \leq \varepsilon$ employing, at most,

$$\kappa_\Phi \left\lceil \frac{\Phi(x_0, t)}{\varepsilon^\alpha} \right\rceil$$

evaluations of f , $c_\mathcal{E}$, $c_\mathcal{I}$, and their derivatives.

The constants κ_θ and κ_Φ mentioned in Assumption A4 depend on algorithmic parameters of the UM algorithm and on quantities associated with the objective function (θ or Φ) of the

unconstrained minimization problem being solved. For example, if the UM algorithm is a typical first-order linesearch method and L_Φ is a Lipschitz constant for $\nabla\Phi(x, t)$ (for all t), we have that $\alpha = 2$ and

$$\kappa_\Phi = L_\Phi \kappa_{\text{UM}},$$

where κ_{UM} depends on algorithmic parameters, such as sufficient descent tolerances and angle conditions' constants (see [13] for instance). Note that the assumption that the same L_Φ may be a Lipschitz constant independently of t is plausible, if one assumes Lipschitz continuity of ∇f , $\nabla c_\mathcal{E}$, and $\nabla c_\mathcal{I}$. Similar conclusions hold if one consider the first-order adaptive cubic regularization method ARC [7] instead of a first-order linesearch method.

The situation is however slightly more complex if one wishes to exploit derivatives of order larger than one for obtaining better worst-case complexity bounds, as we now describe. We base our argument on a recent paper by Birgin *et al.* [4] where unconstrained optimization using high-order models is considered. It results from this last paper that if one wishes to minimize $u(x) = v(x) + w(x)$ over \mathbb{R}^n (where $v(x)$ is at least continuously differentiable and $w(x)$ is p times continuously differentiable with a Lipschitz continuous p -th derivative), and if one is ready to supply derivatives of w up to order p , then a variant of the ARC method starting from \bar{x} and using high-order models can be shown to produce an approximate first-order critical point ($\|\nabla u(x)\| \leq \varepsilon$) in a number of evaluations of $w(x)$ and its derivatives at most equal to

$$\kappa_A \left[\frac{u(\bar{x}) - u_{\text{low}}}{\varepsilon^{(p+1)/p}} \right],$$

where u_{low} is a global lower bound on $u(x)$ and κ_A depends on the Lipschitz constant of the p -th derivative of w , on p , and on algorithmic parameters only. Notice that the number of evaluations of $v(x)$ might be higher (because it is explicitly included in the model and has to be evaluated, possibly with its first derivative, every time the model (and its derivative) is computed). In order to apply this technique, we now reformulate our initial problem (2.1) in the equivalent form

$$\begin{aligned} \min_{x, y, z \in \mathbb{R}^{n+q+1}} \quad & z \\ \text{s.t.} \quad & \bar{c}_\mathcal{E}(x, y, z) = 0, \quad \text{with} \quad \bar{c}_\mathcal{E}(x, y, z) \stackrel{\text{def}}{=} \begin{pmatrix} f(x) - z \\ c_\mathcal{E}(x) \\ c_\mathcal{I}(x) - y \end{pmatrix} \\ & y \leq 0, \end{aligned} \quad (3.31)$$

and construct the associated $\Phi(x, y, z, t)$ function as

$$\Phi(x, y, z, t) = \underbrace{\|\bar{c}_\mathcal{E}(x, y, z)\|^2}_{w(x, y, z)} + \underbrace{\|y_+\|^2 + [z - t]_+^2}_{v(y, z)}.$$

Note that $v(y, z)$ is continuously differentiable and that the differentiability properties of f , $c_\mathcal{E}$, and $c_\mathcal{I}$ are transferred to $\bar{c}_\mathcal{E}(x, y, z)$: the Lipschitz continuity of the p -th derivative of $\Phi(x, y, z, t)$ with respect to its first three variables being ensured if all derivatives of f , $c_\mathcal{E}$, and $c_\mathcal{I}$ are bounded up to order $p - 1$ and Lipschitz continuous up to order p . Note also that the (very simple) evaluation of $v(y, z)$ does not involve any of the problem's functions or derivatives, and thus that the number of these evaluations does not affect the evaluation complexity of the FTARGET algorithm. In these conditions, we may then conclude that the high-order ARC algorithm presented in [4] satisfies Assumption A4 with κ_θ and κ_Φ equal to κ_A and $\alpha = (p+1)/p$.

The following result is a reformulation of Lemma 3.3 in which the number of evaluations employed by the UM algorithm for minimizing $\Phi(x, t_k)$ is quantified.

Lemma 3.5 Suppose that Assumptions A1–A2 and A4 hold. Then, for all $k \geq 0$, the UM algorithm applied to the minimization of $\Phi(x, t_k)$ finds, using no more than

$$\left[\kappa_\Phi \left(\frac{2\kappa_\nabla}{\sigma_1} \right)^\alpha \epsilon_P^2 \epsilon_D^{-\alpha} \min[\epsilon_D^{\sigma_2}, \epsilon_P]^{-\alpha} \right] \quad (3.32)$$

evaluations of f , $c_\mathcal{E}$, $c_\mathcal{I}$, and their derivatives, a point $x_{k,j}$ that satisfies at least one of the criteria (2.5), (2.6), and (2.7).

Proof. Observing that $\Phi(x_{k,0}, t_k) = \Phi(x_k, t_k) = \epsilon_P^2$ and using Lemma 3.2 and Assumption A4, we deduce that the UM algorithm needs no more than

$$\kappa_\Phi \left[\frac{\epsilon_P^2}{\min \left[0.2\rho \epsilon_D \epsilon_P, \frac{\psi(\epsilon_D)}{2}, \frac{\psi(\epsilon_D) \epsilon_D}{2\kappa_\nabla} \right]^\alpha} \right]$$

evaluations to find an approximate minimizer. We obtain the desired bound using the definition of ψ , (3.22), and the inequalities $\sigma_2 \geq 1$, $\epsilon_P \leq 1$, and $\epsilon_D \leq 1$. \square

Notice that, as in Assumption A4, the constant κ_Φ in (3.32) depends only on f , $c_\mathcal{E}$, $c_\mathcal{I}$, and the parameters of the UM algorithm.

It is now possible to prove a complexity result for the FTARGET algorithm.

Theorem 3.6 Suppose that Assumptions A1–A2 and A4 hold and that $f(x) \leq f_{\text{up}}$ for all $x \in \mathcal{L}(\epsilon_{\text{P}}^2)$. Then the FTARGET algorithm stops at Step F2 of Phase 1 or there exists $k \in \{0, 1, 2, \dots\}$ such that FTARGET stops at iteration k of Phase 2 with x_{k+1} satisfying T3 or T4. Moreover, the FTARGET algorithm employs, in Phase 1, at most,

$$\left\lceil \kappa_{\theta} \left[\frac{\theta(x_{-1})}{\sigma_1^\alpha} \right] \epsilon_{\text{D}}^{-\sigma_2 \alpha} \right\rceil \quad (3.33)$$

evaluations of $c_{\mathcal{E}}$, $c_{\mathcal{I}}$, and their derivatives, and, in Phase 2, at most,

$$\left\lceil \left[\kappa_{\Phi} \left(\frac{2\kappa_{\nabla}}{\sigma_1} \right)^\alpha \left(\frac{f_{\text{up}} - f_{\text{low}}}{0.1(1-\rho)} + 1 \right) \right] \epsilon_{\text{P}} \epsilon_{\text{D}}^{-\alpha} \min[\epsilon_{\text{D}}^{\sigma_2}, \epsilon_{\text{P}}]^{-\alpha} \right\rceil \quad (3.34)$$

evaluations of f , $c_{\mathcal{E}}$, $c_{\mathcal{I}}$, and their derivatives.

Proof. The first part of the theorem follows from the definition of ψ and Theorem 3.4 since Assumption A4 implies Assumption A3. The bound on the total number of evaluations of Phase 1 follows directly from Assumption A4; while the bound

$$\kappa_{\Phi} \left[\left(\frac{2\kappa_{\nabla}}{\sigma_1} \right)^\alpha \epsilon_{\text{P}}^2 \epsilon_{\text{D}}^{-\alpha} \min[\epsilon_{\text{D}}^{\sigma_2}, \epsilon_{\text{P}}]^{-\alpha} \right] \left[\left\lceil \frac{f(x_0) - f_{\text{low}}}{0.1(1-\rho) \epsilon_{\text{P}}} \right\rceil + 1 \right]$$

on the total number of evaluations of Phase 2 follows from Lemmas 3.1 and 3.5. The bound (3.34) then follows from the assumption that $f(x) \leq f_{\text{up}}$ for all $x \in \mathcal{L}(\epsilon_{\text{P}}^2)$ and the fact that $\epsilon_{\text{P}} \leq 1$. \square

Notice that, as in Assumption A4, the constant κ_{θ} in (3.33) depends only on $c_{\mathcal{E}}$, $c_{\mathcal{I}}$, and parameters of the UM algorithm, whereas the constant κ_{Φ} in (3.34) depends only on f , $c_{\mathcal{E}}$, $c_{\mathcal{I}}$, and parameters of the UM algorithm. Note also that, by introducing the upper bound f_{up} on $f(x)$ in the neighbourhood $\mathcal{L}(\epsilon_{\text{P}}^2)$ of the feasible set, we have made the complexity bound independent of x_0 (which is not a data of the problem nor an input parameter, but the result of applying the UM algorithm to the minimization of $\theta(x)$ starting from x_{-1} at Phase 1 of the FTARGET algorithm).

4 Discussion

Some comments are now useful to make the result of Theorem 3.6 more explicit. Table 4.1 on the next page summarizes our convergence and complexity results for the case where

$$\epsilon = \epsilon_{\text{D}} = \epsilon_{\text{P}} \quad (4.1)$$

is small enough and for different choices of the UM algorithm and values of σ_2 in the definition (2.13) of ψ . Note that, from the definition in (2.19), $\omega = 1$ in this case, making our discussion

of Section 2.1 relevant. Despite its appealing symmetry, the choice (4.1) is somewhat arbitrary, and variations in the complexities displayed in the table will result from different choices. (The alternative $\epsilon_D = \epsilon_P^{2/3}$ is suggested in [11], leading to a bound in $O(\epsilon_P^{-3/2})$ for the case where $p = 2$, $\sigma_2 = 1$, and $\alpha = 3/2$, but implying that $\omega = \infty$.) In particular, (4.1) implicitly assumes that the problem's scaling is reasonable. Note that the complexity of Phase 1 mentioned in the table is only informative since the overall complexity of the FTARGET algorithm is given by the complexity of Phase 2 which dominates that of Phase 1 in all cases.

		$\psi(\epsilon) = \sigma_1 \epsilon^{\sigma_2}$		
		$\sigma_2 = 1$	$\sigma_2 \in (1, 2)$	$\sigma_2 = 2$
Convergence results		Approximate infeasible stationary point for $\epsilon \rightarrow 0$ (Phase 1), or approximate KKT point (Phase 2), or ...		
		no ξ -uniform positive linear independence of gradients of active constraints with $\xi = \sigma_1 / (2\sqrt{0.99})$	MFCQ fails for $\epsilon \rightarrow 0$	Lojasiewicz fails for $\epsilon \rightarrow 0$
Assumption A4 holds for UM with	$\alpha = 2$ ($p = 1$)	Phase 1: $O(\epsilon^{-2})$ Phase 2: $O(\epsilon^{-3})$	Phase 1: $O(\epsilon^{-2\sigma_2})$ Phase 2: $O(\epsilon^{-1-2\sigma_2})$	Phase 1: $O(\epsilon^{-4})$ Phase 2: $O(\epsilon^{-5})$
	$\alpha = 3/2$ ($p = 2$)	Phase 1: $O(\epsilon^{-1.5})$ Phase 2: $O(\epsilon^{-2})$	Phase 1: $O(\epsilon^{-1.5\sigma_2})$ Phase 2: $O(\epsilon^{-0.5-1.5\sigma_2})$	Phase 1: $O(\epsilon^{-3})$ Phase 2: $O(\epsilon^{-3.5})$
	$\alpha = (p+1)/p$ ($p \geq 1$)	Phase 1: $O(\epsilon^{-\frac{p+1}{p}})$ Phase 2: $O(\epsilon^{1-2\frac{(p+1)}{p}})$	Phase 1: $O(\epsilon^{-\sigma_2 \frac{p+1}{p}})$ Phase 2: $O(\epsilon^{1-(1+\sigma_2)\frac{(p+1)}{p}})$	Phase 1: $O(\epsilon^{-2\frac{(p+1)}{p}})$ Phase 2: $O(\epsilon^{1-3\frac{(p+1)}{p}})$

Table 4.1: Summary of convergence and complexity results for different choices of the UM algorithm and function $\psi(\epsilon) = \sigma_1 \epsilon^{\sigma_2}$ with $\sigma_1 > 0$ and different choices for $\sigma_2 \in [1, 2]$.

With $\sigma_2 = 2$, we obtain the maximal quality of the algorithmic results in the sense that the FTARGET algorithm stops when an approximate KKT point (without scaling) is found or when a very weak property (Łojasiewicz inequality) does not hold in the limit. With $\sigma_2 = 1$, the FTARGET algorithm stops when an approximate KKT point is found or when, in the limit, a relaxed Mangasarian-Fromowitz property with tolerance $\xi = \sigma_1/(2\sqrt{0.99})$ does not hold. Between those extremes, when $1 < \sigma_2 < 2$, either the final point is an approximate KKT point or the full Mangasarian-Fromowitz property fails in the limit. As expected, the complexity goes in the opposite direction. The best complexity is obtained when $\sigma_2 = 1$ and the worst one when $\sigma_2 = 2$. Also as expected, when the UM algorithm satisfies Assumption A4 with $\alpha = (p + 1)/p$ ($p > 2$), complexities are better than those with $\alpha = 2$ ($p = 1$) and $\alpha = 3/2$ ($p = 2$). The best complexity is obtained when α approaches 1 because p grows and $\sigma_2 = 1$.

We now compare the results obtained with those obtained in [8, 11, 12] for $p = 1$ and $p = 2$, focussing as above on the case where $\epsilon = \epsilon_p = \epsilon_D$. In these contributions, the complexity of achieving scaled KKT conditions is considered, at variance with the unscaled approach used in the present paper. By scaled KKT conditions, we mean that (2.8) is modified to take (for a general approximate first-order critical triple (x, λ, μ)), the form

$$\begin{aligned}
 \sqrt{\theta(x)} &\leq \epsilon, \\
 \left\| \nabla f(x) + \sum_{j=1}^m \lambda_j [\nabla c_{\mathcal{E}}(x)]_j + \sum_{j=1}^q \mu_j [\nabla c_{\mathcal{I}}(x)]_j \right\| &\leq \epsilon \|(1, \lambda, \mu)\|, \\
 \sum_{j=1}^q \mu_j [c_{\mathcal{I}}(x)]_j &\leq 2\epsilon \|(1, \lambda, \mu)\|,
 \end{aligned} \tag{4.2}$$

where $\mu_j \leq 0$ for $j = 1, \dots, q$.

It is shown in [8, 12] that, if f , $c_{\mathcal{E}}$, and $c_{\mathcal{I}}$ are continuously differentiable with Lipschitz continuous gradients, then a triple (x, λ, μ) can be found satisfying (4.2) or (2.4)/(2.7) in at most

$$\left\lceil \kappa_1 \frac{\sqrt{\theta(x_{-1})} + f_{\text{up}} - f_{\text{low}}}{\epsilon^2} + \kappa_2 |\log \epsilon| + \kappa_3 \right\rceil$$

evaluations of f , $c_{\mathcal{E}}$, and $c_{\mathcal{I}}$ (and their first derivatives), where κ_1 , κ_2 , and κ_3 are constants independent of ϵ . This result is one order better than the bound of $O(\epsilon^{-3})$ evaluations reported in Table 4.1 for the case $\alpha = 2$ and $\sigma_2 = 1$, indicating that achieving scaled KKT conditions seems easier than achieving unscaled KKT conditions when using first derivatives only. We have focussed on the case where $\sigma_2 = 1$ because the implications in terms of degeneracies for $\epsilon \rightarrow 0$ are not discussed in [8, 12], meaning that it is not clear whether the algorithms being compared declare failure in satisfying scaled or unscaled KKT conditions in the same situations.

The comparison is more difficult if we now allow the use of first and second derivatives, because the results in [11] are expressed using a different first-order criticality measure $\chi(x)$ whose value is the maximal decrease that is achievable on the linearized function under consideration in the intersection of the unit sphere and the feasible domain defined by positivity constraints on slack variables for inequalities (see [11] for details). This criticality measure is also used in

scaled form, in that it has to be below $\epsilon \|(1, \lambda, \mu)\|$ when applied on the Lagrangian (in the spirit of (4.2)), or below ϵ when applied to $\sqrt{\theta(\cdot)}$. This last condition is conceptually similar to (2.4)/(2.7), but is significantly stronger because it involves the gradient of $\sqrt{\theta(\cdot)}$ rather than that of $\theta(\cdot)$; indeed strong assumptions on the singular values of the constraints' Jacobians are needed to ensure their equivalence in order. In this context, it is shown that, if f , $c_{\mathcal{E}}$, and $c_{\mathcal{I}}$ are twice continuously differentiable with Lipschitz continuous gradients and Hessians, and if a possibly restrictive assumption on the subproblem solution holds (see [11] for details), then a triple (x, λ, μ) satisfying these alternative scaled criticality conditions can be obtained in at most

$$\left\lceil \kappa_1 \frac{\sqrt{\theta(x_{-1})} + f_{\text{up}} - f_{\text{low}}}{\epsilon^2} + \kappa_2 \right\rceil$$

evaluations of f , $c_{\mathcal{E}}$, and $c_{\mathcal{I}}$ (and their first and second derivatives), where κ_1 and κ_2 are again constants independent of ϵ . In contrast with the case where $p = 1$, this bound is now (in order) identical to the corresponding result in Table 4.1 ($\alpha = 3/2$, $\sigma_2 = 1$). Whether it could be improved to ensure that either (4.2) or (2.4)/(2.7) holds is an interesting open question.

5 Conclusions

We have presented worst-case evaluation complexity bounds for computing approximate first-order critical points of smooth constrained optimization problems. At variance with previous bounds, these involve the unscaled Karush-Kuhn-Tucker optimality conditions, and cover cases where high-order derivatives are used. As was the case in [8, 11, 12], the complexity bounds are obtained by applying a two-phase algorithm which first enforces approximate feasibility before improving optimality without deteriorating feasibility.

At this stage, the applicable nature of the two-phase algorithm is uncertain. While the specific version presented in this paper is very unlikely to be practical because it closely follows potentially nonlinear constraints, thereby enforcing possibly very short steps, the question of whether more efficient variants of the idea can be made to work in practice remains to be explored.

References

- [1] R. Andreani, G. Haeser, and J.M. Martínez. On sequential optimality conditions for smooth constrained optimization. *Optimization*, 60:627–641, 2011.
- [2] R. Andreani, J.M. Martínez, and B. F. Svaiter. A new sequential optimality condition for constrained optimization and algorithmic consequences. *SIAM Journal on Optimization*, 20:3533–3554, 2010.
- [3] S. Bellavia, C. Cartis, N. I. M. Gould, B. Morini, and Ph. L. Toint. Convergence of a regularized Euclidean residual algorithm for nonlinear least-squares. *SIAM Journal on Numerical Analysis*, 48(1):1–29, 2010.

- [4] E. G. Birgin, J. L. Gardenghi, J. M. Martínez, S. A. Santos, and Ph. L. Toint. Worst-case evaluation complexity for unconstrained nonlinear optimization using high-order regularized models. Technical Report naXys-05-2015, Namur Center for Complex Systems (naXys), University of Namur, Namur, Belgium, 2015.
- [5] J. Bolte, A. Danilidis, O. Ley, and L. Mazet. Characterizations of Łojasiewicz inequalities: Subgradient flows, talweg, convexity. *Transactions of the American Mathematical Society*, 362:3319–3363, 2010.
- [6] C. Cartis, N. I. M. Gould, and Ph. L. Toint. On the complexity of steepest descent, Newton’s and regularized Newton’s methods for nonconvex unconstrained optimization. *SIAM Journal on Optimization*, 20(6):2833–2852, 2010.
- [7] C. Cartis, N. I. M. Gould, and Ph. L. Toint. Adaptive cubic overestimation methods for unconstrained optimization. Part II: worst-case function-evaluation complexity. *Mathematical Programming, Series A*, 130(2):295–319, 2011.
- [8] 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(4):1721–1739, 2011.
- [9] C. Cartis, N. I. M. Gould, and Ph. L. Toint. An adaptive cubic regularization algorithm for nonconvex optimization with convex constraints and its function-evaluation complexity. *IMA Journal of Numerical Analysis*, 32(4):1662–1645, 2012.
- [10] C. Cartis, N. I. M. Gould, and Ph. L. Toint. On the complexity of finding first-order critical points in constrained nonlinear optimization. *Mathematical Programming, Series A*, 144(1):93–106, 2013. DOI: 10.1007/s10107-012-0617-9.
- [11] C. Cartis, N. I. M. Gould, and Ph. L. Toint. On the evaluation complexity of cubic regularization methods for potentially rank-deficient nonlinear least-squares problems and its relevance to constrained nonlinear optimization. *SIAM Journal on Optimization*, 23(3):1553–1574, 2013.
- [12] C. Cartis, N. I. M. Gould, and Ph. L. Toint. Corrigendum: On the complexity of finding first-order critical points in constrained nonlinear optimization. Technical Report RAL-P-2014-013, Rutherford Appleton Laboratory, Chilton, Oxfordshire, England, 2014.
- [13] C. Cartis, Ph. R. Sampaio, and Ph. L. Toint. Worst-case complexity of first-order non-monotone gradient-related algorithms for unconstrained optimization. Technical Report naXys-02-2013, Namur Center for Complex Systems (naXys), University of Namur, Namur, Belgium, 2013.
- [14] J. P. Dussault. Simple unified convergence proofs for the trust-region and a new ARC variant. Technical report, University of Sherbrooke, Sherbrooke, Canada, 2015.
- [15] G. N. Grapiglia, J. Yuan, and Y. Yuan. On the convergence and worst-case complexity of trust-region and regularization methods for unconstrained optimization. *Mathematical Programming, Series A*, (to appear), 2014.

- [16] S. Lojasiewicz. Ensembles semi-analytiques. Technical report, Institut des hautes Etudes Scientifiques, Bures-sur-Yvette, France, 1965. Available online at <http://perso.univ-rennes1.fr/michel.coste/Lojasiewicz.pdf>.
- [17] Yu. Nesterov. *Introductory Lectures on Convex Optimization*. Applied Optimization. Kluwer Academic Publishers, Dordrecht, The Netherlands, 2004.
- [18] Yu. Nesterov and B. T. Polyak. Cubic regularization of Newton method and its global performance. *Mathematical Programming, Series A*, 108(1):177–205, 2006.
- [19] L. Qi and Z. Wei. On the constant positive linear dependence condition and its application to SQP methods. *SIAM Journal on Optimization*, 10:963–981, 2000.
- [20] R. T. Rockafellar. Lagrange multipliers and optimality. *SIAM Review*, 35:183–238, 1993.
- [21] L. N. Vicente. Worst case complexity of direct search. *EURO Journal on Computational Optimization*, 1:143–153, 2013.