

ON THE EVALUATION COMPLEXITY OF CUBIC REGULARIZATION METHODS FOR POTENTIALLY RANK-DEFICIENT NONLINEAR LEAST-SQUARES PROBLEMS AND ITS RELEVANCE TO CONSTRAINED NONLINEAR OPTIMIZATION*

CORALIA CARTIS[†], NICHOLAS I. M. GOULD[‡], AND PHILIPPE L. TOINT[§]

Abstract. We propose a new termination criterion suitable for potentially singular, zero or nonzero residual, least-squares problems, with which cubic regularization variants take at most $\mathcal{O}(\epsilon^{-3/2})$ residual- and Jacobian-evaluations to drive either the Euclidean norm of the residual or its gradient below ϵ ; this is the best known bound for potentially rank-deficient nonlinear least-squares problems. We then apply the new optimality measure and cubic regularization steps to a family of least-squares merit functions in the context of a target-following algorithm for nonlinear equality-constrained problems; this approach yields the first evaluation complexity bound of order $\epsilon^{-3/2}$ for nonconvexly constrained problems when higher accuracy is required for primal feasibility than for dual first-order criticality.

Key words. evaluation complexity, worst-case analysis, least-squares problems, constrained nonlinear optimization, cubic regularization methods

AMS subject classifications. 68Q25, 90C60, 90C30, 90C26, 65K05

DOI. 10.1137/120869687

1. Introduction. A ubiquitous challenge in scientific computing is the minimization of an appropriate norm of a given, sufficiently smooth, vector-valued function $r : \mathbb{R}^n \rightarrow \mathbb{R}^m$. This problem formulation arises in numerous real-life applications requiring data fitting, parameter estimation, image reconstruction, weather forecasting, and so forth [23, Chapter 10]. Crucially, it is often an essential building block when solving constrained nonlinear programming problems, being used, for example, to reduce the constraint violation in various sequential programming [3, 13, 24, 25, 26], filter [17], funnel [18], and iteratively reweighted least-squares approaches [2, section 4.5.2]. Nonlinear least-squares problems are also at the heart of the path-following method for constrained problems, which we propose and analyze here as well.

Here we focus on the Euclidean-norm case that gives rise to the equivalent nonlinear least-squares problem,

$$(1.1) \quad \min_{x \in \mathbb{R}^n} \Phi(x) \stackrel{\text{def}}{=} \frac{1}{2} \|r(x)\|^2,$$

now involving the smooth function $\Phi(x)$; other norms may be of interest and some are equally acceptable in this framework. We allow arbitrary values for m and n , and so both over- and underdetermined residuals $r(x)$ are allowed in (1.1), as well as square

*Received by the editors March 12, 2012; accepted for publication (in revised form) April 9, 2013; published electronically August 1, 2013. This work was supported by the Royal Society through International Joint Project 14265.

<http://www.siam.org/journals/siopt/23-3/86968.html>

[†]School of Mathematics, University of Edinburgh, Edinburgh, EH9 3JZ, Scotland, UK (coralia.cartis@ed.ac.uk). The work of this author was supported by EPSRC grant EP/I028854/1.

[‡]Computational Science and Engineering Department, Rutherford Appleton Laboratory, Chilton, Oxfordshire, OX11 0QX, UK (nick.gould@stfc.ac.uk). The work of this author was funded by EPSRC grant EP/E053351/1.

[§]Namur Center for Complex Systems (NAXyS), University of Namur, B-5000 Namur, Belgium (philippe.toint@fundp.ac.be).

nonlinear systems of equations; in the latter two cases, one may wish to reduce $\Phi(x)$ in (1.1) to zero so as to find the zeros of the system $r(x) = 0$.

Methods for solving (1.1) differ not only in their practical performance, but also in the theoretical bounds known on their worst-case efficiency, which is the focus of this paper. Of the various methods proposed, Gauss–Newton techniques are the most popular and well researched [16, 23]. Rather than tackling the smooth formulation (1.1), recent algorithmic variants [21, 1, 9] attempt to minimize the unsquared and hence nonsmooth norm of $r(x)$ instead, in an attempt to improve the conditioning of the system that defines the change to the iterates. Using only first-order information—namely, values of the residual $r(x)$ and its Jacobian $J(x)$ at given x , obtained from a so-called black-box/oracle—both classical and modern variants can be made/shown to be globally convergent to stationary points of (1.1), namely, to points satisfying

$$(1.2) \quad \nabla_x \Phi(x) \stackrel{\text{def}}{=} J(x)^T r(x) = 0;$$

furthermore, the number of residual and Jacobian evaluations required to bring the norm of (1.2) or some (nonsmooth) first-order optimality measure within some tolerance ϵ is $\mathcal{O}(\epsilon^{-2})$, provided $J(x)$ and $r(x)$ are Lipschitz continuous [21, 1, 9, 23, 14]. Another possibility is to apply Newton-type methods to the unconstrained problem (1.1), which can ensure, for example, fast local convergence for nonzero residual problems and, most importantly here, improved global efficiency for both zero and nonzero residual problems. In particular, cubic regularization methods [19, 22, 27, 10] applied to (1.1) take $\mathcal{O}(\epsilon^{-3/2})$ residual evaluations to ensure (1.2) is within ϵ , provided $r(x)$, $J(x)$, and the Hessians $\nabla_{xx} r_i(x)$, $i = 1, \dots, m$, are Lipschitz continuous; this bound is sharp for nonlinear least squares [12], is optimal from a worst-case complexity point of view for a wide class of second-order methods and nonconvex unconstrained problems [5], and is the best known complexity for second-order methods. This bound can be further improved for gradient-dominated residuals (such as when the singular values of the Jacobian are uniformly bounded away from, or converge to, zero at the same rate as the residual) [22].

The (natural) approximate satisfaction of (1.2) as termination criterion for the cubic regularization and other methods suffers from the disadvantage that an approximate zero of $r(x)$ is guaranteed only when $J(x)$ is uniformly full-rank, with a known lower bound on its smallest singular value—this is a strong assumption. In this paper, we introduce a termination condition that can distinguish between the zero and nonzero residual cases automatically/implicitly. Namely, we argue for the use of a scaled variant of (1.2), which is precisely the gradient of $\|r(x)\|$ whenever $r(x) \neq 0$, as well as the inclusion of the size of the residual in the termination condition. Indeed, irrespective of the algorithm employed and the Jacobian’s rank properties, one may regard the proposed termination as the appropriate way of handling accuracy for nonlinear least-squares problems. Without requiring a nondegenerate Jacobian, we then show that cubic regularization methods can generate *either* an approximate scaled gradient *or* residual value within ϵ in at most $\mathcal{O}(\epsilon^{-3/2})$ residual evaluations, thus preserving the (optimal) order of the bound for cubic regularization.

Consider now the evaluation complexity of minimizing a smooth but potentially nonconvex objective $f(x) \in \mathbb{R}$ for $x \in \mathcal{C}$. When \mathcal{C} is described by finitely many smooth (but potentially nonconvex) equality and inequality constraints, we have shown that a first-order exact penalty method with bounded penalty parameters [9], as well as a short-step target-following algorithm with steepest-descent-like steps [4], take $\mathcal{O}(\epsilon^{-2})$ objective and constraint evaluations to generate an approximate KKT

point of the problem or an infeasible point of the feasibility measure with respect to the constraints. Thus adding constraints does not deteriorate the order of the worst-case evaluation complexity bounds achieved in the unconstrained case when steepest-descent-like methods are employed. A natural question arises as to whether an improved evaluation complexity bound, of the order of cubic regularization, can be shown for constrained problems. In the case when \mathcal{C} is given by convex constraints, projected cubic regularization variants can be shown to satisfy the $\mathcal{O}(\epsilon^{-3/2})$ evaluation bound [7]. In this paper, in a similar vein to [4], we propose a short-step target-following algorithm for problems with nonconvex equality constraints

$$\text{minimize } f(x) \text{ such that } c(x) = 0$$

that takes cubic regularization steps for a sequence of shifting least-squares merit functions. The evaluation complexity of the resulting algorithm is better than that for steepest descent, and can even achieve $\mathcal{O}(\epsilon^{-3/2})$, provided the (dual) KKT conditions are satisfied with lower accuracy than the (primal) feasibility with respect to the constraints.

The structure of the paper is as follows. Section 2 summarizes adaptive cubic regularization methods [10] and relevant complexity results. Section 3.1 presents the new termination criteria for (1.1) based on the scaled gradient, while section 3.2 gives the complexity result for cubic regularization applied to (1.1) with the new termination criteria. Sections 4 and 5 present the short-step target-following cubic regularization algorithm for the equality-constrained problem and its complexity analysis, respectively. Section 6 summarizes our contributions and discusses possible extensions of this work.

2. Previous cubic regularization construction and results.

2.1. Description of adaptive cubic regularization algorithm. We consider applying the Adaptive Regularization with Cubics (ARC) algorithm [11, 10] to (1.1); here, we focus on the ARC variant that has the best known and optimal worst-case evaluation complexity, so-called $\text{ARC}_{(S)}$. At each iterate x_k , $k \geq 0$, a step s_k is computed that approximately minimizes the local cubic model

$$(2.1) \quad m_k(s) = \frac{1}{2} \|r(x_k)\|^2 + s^T J(x_k)^T r(x_k) + \frac{1}{2} s^T B_k s + \frac{1}{3} \sigma_k \|s\|^3$$

of $\Phi(x_k + s)$ with respect to s , where B_k is an approximation to the Hessian of Φ at x_k and $\sigma_k > 0$ is a regularization parameter. In this method, the step s_k is computed to satisfy

$$(2.2) \quad s_k^T J(x_k)^T r(x_k) + s_k^T B_k s_k + \sigma_k \|s_k\|^3 = 0$$

and

$$(2.3) \quad s_k^T B_k s_k + \sigma_k \|s_k\|^3 \geq 0.$$

Conditions (2.2) and (2.3) are achieved whenever s_k is a global minimizer of the model m_k along the direction s_k , namely, $\arg \min_{\alpha \in \mathbb{R}} m_k(\alpha s_k) = 1$; in particular, they are satisfied whenever s_k is a global minimizer of the model m_k over a(ny) subspace [11, Theorem 3.1, Lemma 3.2]. Note that if s_k is chosen as the global minimizer of m_k over the entire space, σ_k is maintained at a sufficiently large value, and B_k is the true Hessian, then $\text{ARC}_{(S)}$ is similar to the cubic regularization technique proposed in [22].

To ensure ARC's fast local convergence, we need to go beyond unidimensional minimization, and so we terminate the inner model minimization when

$$(2.4) \quad \boxed{\text{TC.s}} \quad \|\nabla_s m_k(s_k)\| \leq \kappa_\theta \min\{1, \|s_k\|\} \|J(x_k)^T r(x_k)\|,$$

where κ_θ is any constant in $(0, 1)$; see [11, section 3.2] for a detailed description of this and other possible termination conditions. Note that $\nabla_s m_k(0) = \nabla_x \Phi(x_k) = J(x_k)^T r(x_k)$ so that (2.4) is a relative error condition, which is clearly satisfied at any minimizer s_k of m_k since then $\nabla_s m_k(s_k) = 0$. Generally, we hope that the inner minimization will be terminated before this inevitable outcome. Note that when s_k is computed by minimizing m_k over a subspace, we may increase the subspace of minimization until TC.s is satisfied. In particular, one may use a Lanczos-based approach where the subspace is the Krylov one generated by $\{\nabla_x \Phi(x_k), B_k \nabla_x \Phi(x_k), B_k^2 \nabla_x \Phi(x_k), \dots\}$. In this case, conditions (2.2) and (2.3) are also achieved [11, sections 3.2, 6, and 7].

It remains to describe the iterate updating and model improvement technique in ARC. The step s_k is accepted and the new iterate x_{k+1} set to $x_k + s_k$ whenever (a reasonable fraction of) the predicted model decrease $\Phi(x_k) - m_k(s_k)$ is realized by the actual decrease in the objective, $\Phi(x_k) - \Phi(x_k + s_k)$. This is measured by computing the ratio ρ_k in (2.5) and requiring ρ_k to be greater than a prescribed positive constant η_1 (for example, $\eta_1 = 0.1$); it can be shown that ρ_k is well-defined whenever $\nabla_x \Phi(x_k) \neq 0$ [11, Lemma 2.1]. Since the current weight σ_k has resulted in a successful step, there is no pressing reason to increase it, and indeed there may be benefits in decreasing it if the model overestimates the function locally. By contrast, if ρ_k is smaller than η_1 , we judge that the improvement in objective is insufficient—indeed there is no improvement if $\rho_k \leq 0$. If this happens, the step will be rejected and x_{k+1} left as x_k . Under these circumstances, the only recourse available is to increase the weight σ_k prior to the next iteration with the implicit intention of reducing the size of the step.

Algorithm 2.1 gives a summary of the $\text{ARC}_{(S)}$ algorithm applied to (1.1).

Algorithm 2.1: Adaptive regularization using cubics ($\text{ARC}_{(S)}$) [11, 10] applied to (1.1).

A starting point x_0 , an initial and a minimal regularization parameter $\sigma_0 \geq \sigma_{\min} > 0$, and algorithmic parameters $\gamma_2 \geq \gamma_1 > 1$ and $1 > \eta_2 \geq \eta_1 > 0$ are given.

For $k = 0, 1, \dots$, until **termination**, do:

1. Compute a step s_k that satisfies (2.2)–(2.4).
2. Compute $r(x_k + s_k)$ and

$$(2.5) \quad \rho_k = \frac{\frac{1}{2}\|r(x_k)\|^2 - \frac{1}{2}\|r(x_k + s_k)\|^2}{\frac{1}{2}\|r(x_k)\|^2 - m_k(s_k)}.$$

3. Set

$$x_{k+1} = \begin{cases} x_k + s_k & \text{if } \rho_k \geq \eta_1, \\ x_k & \text{otherwise.} \end{cases}$$

4. Set

$$(2.6) \quad \sigma_{k+1} \in \begin{cases} [\sigma_{\min}, \sigma_k] & \text{if } \rho_k > \eta_2, & \text{[very successful iteration]} \\ [\sigma_k, \gamma_1 \sigma_k] & \text{if } \eta_1 \leq \rho_k \leq \eta_2, & \text{[successful iteration]} \\ [\gamma_1 \sigma_k, \gamma_2 \sigma_k] & \text{otherwise.} & \text{[unsuccessful iteration]} \end{cases}$$

Note that we have not yet defined the condition required for $\text{ARC}_{(S)}$ to terminate. In [11, 10], we terminate ARC when $\|\nabla_x \Phi(x_k)\| \leq \epsilon$, and possibly also $\lambda_{\min}(\nabla_{xx} \Phi(x_k)) \geq -\epsilon$, for a user-specified tolerance $\epsilon \in (0, 1)$. Here, we will require that either some scaled gradient or the residual is within ϵ ; this novel termination condition, specific to (1.1), is described in section 3.1.

Note that in the important special case of (1.1) when $r(x) := \nabla f(x)$ for some sufficiently smooth objective $f(x)$, $\text{ARC}_{(S)}$ is a third-order scheme for the minimization of $f(x)$. As the main issue concerning the development of third-order schemes for nonlinear optimization is the NP-hardness of the majority of auxiliary optimization subproblems related to multivariate polynomials of degree three, the construction (2.1) offers a tractable, even computationally inexpensive, way of incorporating third-order information in the optimization technique.

2.2. Assumptions and useful results. The following assumptions are chosen to ensure that those in [11, 10] are satisfied when $\text{ARC}_{(S)}$ is applied to (1.1), which allows us to employ some crucial ARC results from [11, 10] to (1.1).

We assume that

(2.7)

$$\boxed{\text{AR.1}} \quad r_i \text{ is twice continuously differentiable on } \mathbb{R}^n \text{ for all } i \in \{1, \dots, m\}.$$

We also assume that the Jacobian J of r , the residuals r_i , and the Hessian $\nabla_{xx} r_i$ for each $i \in \{1, \dots, m\}$ are globally Lipschitz continuous on the path of all generated iterates and trial points, namely,

(2.8)

$$\boxed{\text{AR.2}} \quad |r_i(x) - r_i(x_k)| \leq \kappa_{r_i} \|x - x_k\| \text{ for all } x \in [x_k, x_k + s_k] \text{ and } k \geq 0$$

for some $\kappa_{r_i} \geq 1$ for each $i \in \{1, \dots, m\}$, and

(2.9)

$$\boxed{\text{AR.3}} \quad \|J(x) - J(x_k)\| \leq \kappa_J \|x - x_k\| \text{ for all } x \in [x_k, x_k + s_k] \text{ and } k \geq 0$$

for some $\kappa_J > 0$, and finally, for each $i \in \{1, \dots, m\}$, there exists $L_i > 0$ such that

(2.10)

$$\boxed{\text{AR.4}} \quad \|\nabla_{xx} r_i(x) - \nabla_{xx} r_i(x_k)\| \leq L_i \|x - x_k\| \text{ for all } x \in [x_k, x_k + s_k]$$

and for all $k \geq 0$. Assumption AR.2 implies that r is globally Lipschitz continuous on the path of all generated iterates and trial points, with Lipschitz constant $\kappa_r \stackrel{\text{def}}{=} \|(\kappa_{r_1}, \dots, \kappa_{r_m})\| \geq 1$. Furthermore, AR.1–AR.4 imply that the gradient $\nabla_x \Phi$ given in (1.2) and the Hessian of Φ

$$(2.11) \quad \nabla_{xx} \Phi(x) = J(x)^T J(x) + \sum_{i=1}^n r_i(x) \nabla_{xx} r_i(x)$$

are globally Lipschitz continuous on the path of all generated iterates $[x_k, x_k + s_k]$, $k \geq 0$, with Lipschitz constants

$$(2.12) \quad L_g \stackrel{\text{def}}{=} \kappa_r^2 + \|r(x_0)\| \kappa_J \geq 1$$

and

$$(2.13) \quad L \stackrel{\text{def}}{=} 2\kappa_J \kappa_r + \kappa_J \sum_{i=1}^n \kappa_{r_i} + \|r(x_0)\| \sum_{i=1}^n L_i,$$

respectively, where we also used that ARC generates monotonically decreasing function values so that $\|r(x_k)\| \leq \|r(x_0)\|$. (These are variants of assumptions AF.4 and AF.6, respectively, in [11, 10], suitable for our current purposes.)

Clearly, the values of the residual $r(x_k)$ and its Jacobian $J(x_k)$ are required to form the model (2.1) and estimate (2.5). Thus, as B_k is an approximation to the Hessian of Φ in (2.11) at x_k , only the Hessian of each r_i needs to be approximated in B_k and so it is natural to consider B_k to be of the form

$$(2.14) \quad B_k = J(x_k)^T J(x_k) + M_k,$$

where

$$(2.15) \quad M_k \approx H_\Phi(x_k) \stackrel{\text{def}}{=} \sum_{i=1}^n r_i(x_k) \nabla_{xx} r_i(x_k).$$

We require that M_k and $H_\Phi(x_k)$ in (2.15) agree along s_k in the sense that there exists some constant $C > 0$ such that

$$(2.16) \quad \boxed{\text{AM.4}} \quad \|(H_\Phi(x_k) - M_k)s_k\| \leq C\|s_k\|^2 \text{ for all } k \geq 0.$$

This, (2.11), and (2.14) imply that

$$(2.17) \quad \|\nabla_{xx}\Phi(x_k) - B_k\|s_k\| \leq C\|s_k\|^2 \text{ for all } k \geq 0,$$

which is assumption AM.4 in [11, 10]. The condition AM.4 is trivially satisfied with $C = 0$ when we set $M_k = H_\Phi(x_k)$, i.e., $B_k = \nabla_{xx}\Phi(x_k)$, for all $k \geq 0$ in the ARC algorithm. The requirement (2.16) or (2.17) is stronger than the Dennis–Moré condition [15]. The latter is achieved by some quasi-Newton updates provided further assumptions hold (see the discussion following [11, (4.6)]). Quasi-Newton methods may still satisfy AM.4 in practice, though we are not aware if this can be ensured theoretically. We have shown in [8] that AM.4 can be achieved when B_k is approximated by (forward) finite differences of gradient values, without changing the order of the worst-case evaluation complexity bound as a function of the accuracy ϵ .

The first lemma recalls some useful ARC properties, crucial to the complexity bound in section 3.2.

LEMMA 2.1. *Let AR.1–AR.4 and AM.4 hold, and apply Algorithm ARC_(S) to (1.1). Then*

$$(2.18) \quad \sigma_k \geq \frac{3}{2}(L + C) \implies k \text{ is very successful},$$

and so

$$(2.19) \quad \sigma_k \leq \max(\sigma_0, \frac{3}{2}\gamma_2(L + C)) \stackrel{\text{def}}{=} \bar{\sigma} \text{ for all } k \geq 0,$$

where L and C are defined in (2.13) and (2.16), respectively. Also, we have the function decrease

$$(2.20) \quad \frac{1}{2}\|r(x_k)\|^2 - \frac{1}{2}\|r(x_{k+1})\|^2 \geq \alpha \|J(x_{k+1})^T r(x_{k+1})\|^{3/2} \text{ for all successful } k,$$

where $\alpha \stackrel{\text{def}}{=} \eta_1 \sigma_{\min} \kappa_g^3 / 6$ and where κ_g is the positive constant

$$(2.21) \quad \kappa_g \stackrel{\text{def}}{=} \sqrt{(1 - \kappa_\theta) / (\frac{1}{2}L + C + \bar{\sigma} + \kappa_\theta L_g)},$$

with κ_θ , $\bar{\sigma}$, and L_g defined in (2.4), (2.19), and (2.12), respectively.

Proof. The relation (2.18) and the bound (2.19) both follow from [11, Lemma 5.2], and (2.20) from (2.5), $\sigma_k \geq \sigma_{\min}$ (due to (2.6)), [11, Lemma 3.3], and [10, Lemma 5.2]. \square

Relating successful and total iteration counts. The total number of (major) ARC iterations is the same as the number of residual/function evaluations (as we also need to evaluate r on unsuccessful iterations in order to be able to compute ρ_k in (2.5)), while the number of successful ARC iterations is the same as that of Jacobian/gradient evaluations.

Let us introduce some useful notation. Throughout, denote the index set

$$(2.22) \quad \mathcal{S} \stackrel{\text{def}}{=} \{k \geq 0 : k \text{ successful or very successful in the sense of (2.6)}\},$$

and, given any $j \geq 0$, let

$$(2.23) \quad \mathcal{S}_j \stackrel{\text{def}}{=} \{k \leq j : k \in \mathcal{S}\},$$

with $|\mathcal{S}_j|$ denoting the cardinality of the latter.

The lower bound on σ_k and the construction of steps 2–4 of $\text{ARC}_{(\mathcal{S})}$ allow us to quantify the total iteration count as a function of the successful ones.

THEOREM 2.2. *For any fixed $j \geq 0$, let \mathcal{S}_j be defined in (2.23). Assume that there exists $\bar{\sigma} > 0$ be such that*

$$(2.24) \quad \sigma_k \leq \bar{\sigma} \text{ for all } k \leq j.$$

Then

$$(2.25) \quad j \leq \left\lceil 1 + \frac{2}{\log \gamma_1} \log \left(\frac{\bar{\sigma}}{\sigma_{\min}} \right) \right\rceil \cdot |\mathcal{S}_j|.$$

Proof. The updates (2.6) imply that $\sigma_k \geq \sigma_{\min}$ for all k . Now apply [10, Theorem 2.1], namely, the bound [10, (2.14)] on the number of unsuccessful iterations up to j , and use the fact that the unsuccessful iterations up to j together with \mathcal{S}_j form a partition of $\{0, \dots, j\}$. \square

Values for $\bar{\sigma}$ in (2.24) are provided in (2.19), under the assumptions of Lemma 2.1. Thus, based on Theorem 2.2, it remains to bound the successful iteration count $|\mathcal{S}_j|$ since the total iteration count up to j is of the same order in ϵ as $|\mathcal{S}_j|$.

3. Evaluation complexity of cubic regularization for potentially rank-deficient nonlinear least-squares problems.

3.1. A suitable termination condition for $\text{ARC}_{(\mathcal{S})}$. Here, we depart from the standard choice of termination criterion for derivative-based optimization algorithms such as $\text{ARC}_{(\mathcal{S})}$ when applied to (1.1), namely, requiring a sufficiently small gradient $\|\nabla \Phi_x(x_k)\| = \|J(x_k)^T r(x_k)\| \leq \epsilon$, where $\epsilon > 0$ is the user-specified accuracy tolerance. Such a condition is only guaranteed to provide an approximate zero of

the residual r when $J(x)$ is uniformly full-rank and a lower bound on its smallest singular values is known, which are limiting assumptions. Such assumptions are not required for steepest-descent-like methods if appropriate optimality measures are employed [9, 4], but the complexity of such methods is worse than the best second-order methods [9, 12]. Thus, we introduce a termination condition that can distinguish between the zero and nonzero residual cases automatically/implicitly. We propose the following termination for $\text{ARC}_{(\mathcal{S})}$:

$$(3.1) \quad \textbf{termination} : \quad \|r(x_k)\| \leq \epsilon_p \quad \text{or} \quad \|g_r(x_k)\| \leq \epsilon_d,$$

where $\epsilon_p > 0$ and $\epsilon_d > 0$ are the required accuracy tolerances and where

$$(3.2) \quad g_r(x) \stackrel{\text{def}}{=} \begin{cases} \frac{J(x)^T r(x)}{\|r(x)\|} & \text{whenever } r(x) \neq 0; \\ 0 & \text{otherwise.} \end{cases}$$

Note that the scaled gradient $g_r(x)$ in (3.2) is precisely the gradient of $\|r(x)\|$ whenever $r(x) \neq 0$. If $r(x) = 0$, we are at the global minimum of r and so $g_r(x) = 0 \in \partial(\|r(x)\|)$ [20, section VI.3].

In the termination condition (3.1), the scaled gradient $g_r(x_k)$ may be bounded away from zero—for instance, when the singular values of the Jacobian are uniformly bounded away from zero—then, as we show in the next subsection, the residual values converge to zero, and so (3.1) can be achieved. When the iterates approach a nonzero residual value, then g_r converges to zero, and so again (3.1) can be satisfied. (Another suitable termination condition with similar properties is given after the main result in section 3.2.)

In the next subsection, we show that $\text{ARC}_{(\mathcal{S})}$ can generate either an approximate scaled gradient or residual value within ϵ in at most $\mathcal{O}(\epsilon^{-3/2})$ residual evaluations, thus preserving the (optimal) order of the bound for cubic regularization.

3.2. Evaluation complexity of $\text{ARC}_{(\mathcal{S})}$ with termination condition (3.1).

The first lemma exploits (2.20) to give new lower bounds on the function decrease that depend on the residual and the scaled gradient (3.2); the bounds below will also be used for the constrained case.

LEMMA 3.1. *Let AR.1–AR.4 and AM.4 hold, and apply the $\text{ARC}_{(\mathcal{S})}$ algorithm to (1.1). Then, for all successful iterations k for which $r(x_k) \neq 0$, we have*

$$(3.3) \quad \|r(x_k)\| - \|r(x_{k+1})\| \geq \min \left\{ \alpha \beta^{\frac{3}{2}} \|g_r(x_{k+1})\|^{\frac{3}{2}} \cdot \|r(x_k)\|^{\frac{1}{2}}, (1 - \beta) \|r(x_k)\| \right\}$$

and

$$(3.4) \quad \|r(x_k)\|^{\frac{1}{2}} - \|r(x_{k+1})\|^{\frac{1}{2}} \geq \min \left\{ \frac{1}{2} \alpha \beta^{\frac{3}{2}} \|g_r(x_{k+1})\|^{\frac{3}{2}}, (\beta^{-\frac{1}{2}} - 1) \|r(x_{k+1})\|^{\frac{1}{2}} \right\},$$

where α is defined just after (2.20) and $\beta \in (0, 1)$ is any fixed problem-independent constant.

Proof. Suppose that $r(x_k) \neq 0$, let $\beta \in (0, 1)$, and denote

$$(3.5) \quad \mathcal{S}_\beta \stackrel{\text{def}}{=} \{k \in \mathcal{S} : \|r(x_{k+1})\| > \beta \|r(x_k)\|\},$$

where \mathcal{S} is defined in (2.22). We first analyze the function decrease for iterations $k \in \mathcal{S}_\beta$ and then for the ones in $\mathcal{S} \setminus \mathcal{S}_\beta$. Let $k \in \mathcal{S}_\beta$; then $r(x_{k+1}) \neq 0$ since $r(x_k) \neq 0$.

From (2.20), (3.2), and (3.5), we deduce

$$\begin{aligned}
 \|r(x_k)\|^2 - \|r(x_{k+1})\|^2 &\geq 2\alpha \|J(x_{k+1})^T r(x_{k+1})\|^{3/2} \\
 &= 2\alpha \left(\frac{\|J(x_{k+1})^T r(x_{k+1})\|}{\|r(x_{k+1})\|} \right)^{3/2} \|r(x_{k+1})\|^{3/2} \\
 (3.6) \qquad &= 2\alpha \|g_r(x_{k+1})\|^{3/2} \cdot \|r(x_{k+1})\|^{3/2} \\
 &\geq 2\alpha\beta^{3/2} \|g_r(x_{k+1})\|^{3/2} \cdot \|r(x_k)\|^{3/2}.
 \end{aligned}$$

Conjugacy properties and the monotonicity relation $\|r(x_k)\| \geq \|r(x_{k+1})\|$ give

$$(3.7) \quad \|r(x_k)\| - \|r(x_{k+1})\| = \frac{\|r(x_k)\|^2 - \|r(x_{k+1})\|^2}{\|r(x_k)\| + \|r(x_{k+1})\|} \geq \frac{\|r(x_k)\|^2 - \|r(x_{k+1})\|^2}{2\|r(x_k)\|},$$

and furthermore

$$\begin{aligned}
 (3.8) \quad \sqrt{\|r(x_k)\|} - \sqrt{\|r(x_{k+1})\|} &= \frac{\|r(x_k)\| - \|r(x_{k+1})\|}{\sqrt{\|r(x_k)\|} + \sqrt{\|r(x_{k+1})\|}} \\
 &\geq \frac{\|r(x_k)\|^2 - \|r(x_{k+1})\|^2}{4\|r(x_k)\|^{3/2}}.
 \end{aligned}$$

Employing the last inequality in (3.6) into (3.7) and (3.8), respectively, we obtain

$$(3.9) \quad \|r(x_k)\| - \|r(x_{k+1})\| \geq \alpha\beta^{3/2} \|g_r(x_{k+1})\|^{3/2} \cdot \|r(x_k)\|^{1/2} \text{ for all } k \in \mathcal{S}_\beta$$

and

$$(3.10) \quad \|r(x_k)\|^{1/2} - \|r(x_{k+1})\|^{1/2} \geq \frac{\alpha\beta^{3/2}}{2} \|g_r(x_{k+1})\|^{3/2} \text{ for all } k \in \mathcal{S}_\beta.$$

Conversely, let $k \in \mathcal{S} \setminus \mathcal{S}_\beta$, which gives

$$(3.11) \quad \|r(x_{k+1})\| \leq \beta \|r(x_k)\|,$$

and so the residual values decrease linearly on such iterations. It follows from (3.11) that on such iterations we have the following function decrease:

$$(3.12) \quad \|r(x_k)\| - \|r(x_{k+1})\| \geq (1 - \beta) \|r(x_k)\| \text{ for all } k \in \mathcal{S} \setminus \mathcal{S}_\beta$$

and

$$\begin{aligned}
 (3.13) \quad \|r(x_k)\|^{1/2} - \|r(x_{k+1})\|^{1/2} &\geq (1 - \sqrt{\beta}) \|r(x_k)\|^{1/2} \\
 &\geq \frac{1 - \sqrt{\beta}}{\sqrt{\beta}} \|r(x_{k+1})\|^{1/2} \text{ for all } k \in \mathcal{S} \setminus \mathcal{S}_\beta.
 \end{aligned}$$

(Note that (3.12) and (3.13) continue to hold if $r(x_{k+1}) = 0$.) The bound (3.3) now follows from (3.9) and (3.12), and (3.4) from (3.10) and (3.13). \square

The next theorem gives a general evaluation complexity result for $\text{ARC}_{(\mathcal{S})}$ applied to (1.1) when the termination condition (3.1) is employed.

THEOREM 3.2. *Let AR.1–AR.4 and AM.4 hold, and let $\epsilon_p, \epsilon_d \in (0, 1)$. Consider applying the $\text{ARC}_{(\mathcal{S})}$ algorithm with the termination condition (3.1) to minimize (1.1). Then $\text{ARC}_{(\mathcal{S})}$ terminates after at most*

$$(3.14) \quad \left\lceil \max\{\kappa_1 \epsilon_d^{-3/2}, \kappa_2 \epsilon_p^{-1/2}\} \right\rceil + 1$$

successful iterations—or, equivalently, Jacobian evaluations—and at most

$$(3.15) \quad \left\lceil \kappa_S \max\{\kappa_1 \epsilon_d^{-3/2}, \kappa_2 \epsilon_p^{-1/2}\} \right\rceil + 1$$

total (successful and unsuccessful) iterations—or, equivalently, residual evaluations, where

$$(3.16) \quad \kappa_1 \stackrel{\text{def}}{=} 2\|r(x_0)\|^{1/2}\alpha^{-1}\beta^{-3/2}, \quad \kappa_2 \stackrel{\text{def}}{=} \|r(x_0)\|^{1/2}(\beta^{-1/2} - 1)^{-1},$$

$$(3.17) \quad \kappa_S \stackrel{\text{def}}{=} 2(1 + \kappa_S^u), \quad \text{and} \quad \kappa_S^u \stackrel{\text{def}}{=} 2\log(\bar{\sigma}/\sigma_{\min})/\log \gamma_1,$$

with α defined just after (2.20), $\bar{\sigma}$ defined as in (2.19), and $\beta \in (0, 1)$ a fixed problem-independent constant.

Proof. Clearly, if (3.1) is satisfied at the starting point, there is nothing left to prove. Assume now that (3.1) fails at $k = 0$. For any iteration $(k + 1)$ at which $\text{ARC}_{(S)}$ does not terminate, it follows from (3.1) that we have

$$(3.18) \quad \|r(x_{k+1})\| > \epsilon_p \quad \text{and} \quad \|g_r(x_{k+1})\| > \epsilon_d.$$

From (3.4) and (3.18), we deduce that, for all $k \in \mathcal{S}$ for which (3.18) holds,

$$(3.19) \quad \|r(x_k)\|^{1/2} - \|r(x_{k+1})\|^{1/2} \geq \min\{\frac{1}{2}\alpha\beta^{3/2}\epsilon_d^{3/2}, (\beta^{-1/2} - 1)\epsilon_p^{1/2}\}.$$

Summing up (3.19) over all iterations $k \in \mathcal{S}$ for which (3.18) holds, with say $j_\epsilon \leq \infty$ as the largest index, and using that the $\text{ARC}_{(S)}$ iterates remain unchanged over unsuccessful iterations, we obtain

$$(3.20) \quad \begin{aligned} \|r(x_0)\|^{1/2} - \|r(x_{j_\epsilon})\|^{1/2} &= \sum_{k=0, k \in \mathcal{S}}^{j_\epsilon-1} \left[\|r(x_k)\|^{1/2} - \|r(x_{k+1})\|^{1/2} \right] \\ &\geq |\mathcal{S}_\epsilon| \min\left\{ \frac{1}{2}\alpha\beta^{3/2}\epsilon_d^{3/2}, (\beta^{-1/2} - 1)\epsilon_p^{1/2} \right\}, \end{aligned}$$

where $|\mathcal{S}_\epsilon|$ denotes the number of successful iterations up to iteration j_ϵ . Using that $\|r(x_{j_\epsilon})\|^{1/2} \geq 0$, we further obtain from (3.20) that $j_\epsilon < \infty$ and that

$$|\mathcal{S}_\epsilon| \leq \frac{\|r(x_0)\|^{1/2}}{\min\left\{ \frac{1}{2}\alpha\beta^{3/2}\epsilon_d^{3/2}, (\beta^{-1/2} - 1)\epsilon_p^{1/2} \right\}},$$

which gives (3.14) since $|\mathcal{S}_\epsilon|$ must be an integer and since the termination condition is checked at the next iteration; see [10, (5.21), (5.22)] for full details. To derive (3.15), apply Theorem 2.2 with $j = j_\epsilon$, with $\bar{\sigma}$ defined as in (2.19), and use also that $\epsilon_p, \epsilon_d \in (0, 1)$. \square

The next corollary gives the main complexity result of this section, whose proof follows immediately from Theorem 3.2. It shows that the evaluation complexity of $\text{ARC}_{(S)}$ driving either $\|r(x)\|$ or its gradient below ϵ is $\mathcal{O}(\epsilon^{-3/2})$, an improvement of existing $\text{ARC}_{(S)}$ results which can only ensure that the gradient of $\|r(x)\|^2$ goes below ϵ in that same-order number of evaluations.

COROLLARY 3.3. *Let AR.1–AR.4 and AM.4 hold, and let $\epsilon \stackrel{\text{def}}{=} \min\{\epsilon_p, \epsilon_d\} \in (0, 1)$. Consider applying the $\text{ARC}_{(S)}$ algorithm with the termination condition (3.1) to minimize (1.1). Then $\text{ARC}_{(S)}$ terminates after at most*

$$(3.21) \quad \left\lceil \kappa_S^s \epsilon^{-3/2} \right\rceil + 1$$

successful iterations—or, equivalently, Jacobian evaluations—and at most

$$(3.22) \quad \left\lceil \kappa_S \kappa_S^s \epsilon^{-3/2} \right\rceil + 1$$

total (successful and unsuccessful) iterations—or, equivalently, residual evaluations, where

$$(3.23) \quad \kappa_S^s \stackrel{\text{def}}{=} \|r(x_0)\|^{1/2} / \min\{\frac{1}{2}\alpha\beta^{3/2}, \beta^{-1/2} - 1\},$$

with α defined just after (2.20), κ_S , as in (3.17), and $\beta \in (0, 1)$ a fixed problem-independent constant.

Some remarks on the above theorem/corollary and its proof follow:

- Note that in the nonzero residual case, namely, when $\{\|r(x_k)\|\}$ converges to some $r_* > 0$, the monotonicity of this sequence implies that $\|r(x_{k+1})\| \geq \beta\|r(x_k)\|$ for all k , with $\beta \stackrel{\text{def}}{=} r_*/\|r(x_0)\| \in (0, 1)$. Thus in this case, there is no need to consider the iterations (3.11) of faster linear convergence.
- The function decrease in (3.4) implies that instead of (3.1), we could have used the condition

$$(3.24) \quad \textbf{termination 2: } \|r(x_k)\|^{1/3} \leq \epsilon_p \text{ or } \|g_r(x_k)\| \leq \epsilon_d$$

as termination for the $\text{ARC}_{(S)}$ algorithm, without changing the order of the complexity bound as a function of (ϵ_p, ϵ_d) or even of $\epsilon = \min\{\epsilon_p, \epsilon_d\}$. In fact, using the condition (3.24) improves the bound/accuracy for the residual values reaching within ϵ_p .

- Note that the bound (3.14) is a bound on the total number of successful iterations for which (3.18) holds. Thus despite the measure (3.1) being non-monotonic, after (3.14) iterations are taken, this measure would remain below (ϵ_p, ϵ_d) for the remaining $\text{ARC}_{(S)}$ iterations, if any are taken.
- The use of conjugacy in the above proof is reminiscent of the reweighted least-squares techniques [23]. However, our attempts at applying (some modified) ARC to such variants of (1.1) have not been successful.

3.3. Is the bound (3.15) sharp for the nonlinear least-squares problem (1.1)? Recall the example in [12, section 5] that shows that $\text{ARC}_{(S)}$ takes essentially $\epsilon^{-3/2}$ iterations/evaluations to ensure that the norm of the gradient is less than ϵ . The univariate function $f : \mathbb{R} \rightarrow \mathbb{R}$ in question is positive for all $x \geq 0$ and at the iterates, and it is zero at infinity, minimum to which $\text{ARC}_{(S)}$ converges. Thus this example can be viewed as a least-squares, zero-residual problem, with r in (1.1) defined as $r \stackrel{\text{def}}{=} \sqrt{f}$. It shows that $\text{ARC}_{(S)}$ with the termination condition that the absolute value of (1.2)—which in this case is precisely the first derivative of f —is less than ϵ takes essentially $\epsilon^{-3/2}$ iterations/evaluations, and so the $\text{ARC}_{(S)}$ complexity bound is sharp for nonlinear least squares. (Note that although $\sqrt{f(x)}$ and its derivatives may not be globally Lipschitz continuous as $x \rightarrow \infty$, the first and second derivatives of $|r|^2 = f$ have this property, as we have shown in [12, section 5]. The latter conditions are sufficient for the $\mathcal{O}(\epsilon^{-3/2})$ bound to hold for $\text{ARC}_{(S)}$.) It is unclear whether the bound (3.15) for $\text{ARC}_{(S)}$ with the termination condition (3.1) is also sharp.

3.4. Further improving the evaluation complexity of cubic regularization for nonlinear least-squares with special structure. Suppose that $r(x)$ in (1.1) is gradient-dominated of degree 2 [22], namely,

$$(3.25) \quad \frac{\|J(x)^T r(x)\|}{\|r(x)\|} \geq \sigma_{\min}(J(x)) \geq \tau_2 > 0, \quad x \in \mathbb{R}^n,$$

where $\sigma_{\min}(J(x))$ denotes the smallest singular value of $J(x)$; this implies that g_r in (3.1) is bounded away from zero for all $r(x) \neq 0$. Then under the conditions of Theorem 3.2, one can deduce from (3.4) and (3.19) that $r(x_k)$ must converge to zero as $k \rightarrow \infty$, and that the asymptotic rate of this convergence is superlinear (i.e., linear with any convergence factor $\beta \in (0, 1)$); also, the algorithm takes a (problem-dependent) constant number of steps to enter this region of superlinear convergence. We do not give the details of this result here as a (slightly stronger) result of this form—where the size of the neighborhood of fast local convergence does not depend on β and $r(x_0)$ enters the complexity bound in a polynomial way—was given in [22, Theorem 7] for cubic regularization; the latter result continues to hold here for $\text{ARC}_{(S)}$ when applied to problems which we know a priori satisfy (3.25) since then (3.1) is no longer required explicitly. An advantage of our (slightly weaker) approach here is that the termination condition (3.1) “senses” naturally when (3.25) holds and ensures $\text{ARC}_{(S)}$ behaves accordingly.

Similarly, assume now that the smallest singular value of the Jacobian of $r(x)$ converges to zero at the same rate as $r(x)$, or that there exists $\tau_1 > 0$ such that $\|J(x)^T r(x)\|/\|r(x)\| \geq \tau_1 \|r(x)\|$ for all x , which is the same as $r(x)$ being gradient-dominated of degree 1 [22]. Then again we can deduce improved complexity bounds from (3.4) in the same vein as [22, Theorem 6], giving that $\text{ARC}_{(S)}$ requires at most $\mathcal{O}(\epsilon^{-1})$ evaluations to ensure $\|r(x_k)\| \leq \epsilon$. (Note the understandably weaker bound in this case since we minimize the square of the residual, when compared to the ARC bound of order $\mathcal{O}(\epsilon^{-1/2})$ for minimizing general unconstrained gradient-dominated functions of degree 1 [22, 6].) The cases of gradient-dominated residuals of some intermediate degree with value between 1 and 2 can be similarly analyzed, yielding improvement over the bound (3.15).

4. The ShS-ARC algorithm for equality-constrained problems. Consider now the equality-constrained problem

$$(4.1) \quad \text{minimize } f(x) \text{ such that } c(x) = 0,$$

where $f: \mathbb{R}^n \rightarrow \mathbb{R}$ and $c: \mathbb{R}^n \rightarrow \mathbb{R}^m$ with $m \leq n$. We define a slightly larger set than the set of approximately feasible points, namely,

$$(4.2) \quad \mathcal{C}_1 = \{x \in \mathbb{R}^n \mid \|c(x)\| < \kappa_c\},$$

where $\kappa_c > \epsilon_p$ is small constant independent of ϵ_p and where $\epsilon_p \in (0, 1)$ is the accuracy we aim to achieve in satisfying the constraints of (4.1). We assume the following.

AC.1 The function c is twice continuously differentiable on \mathbb{R}^n and f is twice continuously differentiable in a (sufficiently large) open set containing \mathcal{C}_1 .

The algorithm we now describe consists of two phases; see Figure 4.1(a). In the first, $\text{ARC}_{(S)}$ with termination condition (3.1) is applied to (1.1) with $r = c$, so as to minimize $\frac{1}{2}\|c(x)\|^2$ (independently of the objective function f), resulting in a point which is either (approximately) feasible or is an approximate infeasible stationary point of $\|c(x)\|$. The latter outcome is not desirable if one wishes to solve (4.1), but cannot be avoided by any algorithm not relying on global minimization or if \mathcal{C}_1 is empty. If an (approximate) feasible point has been found, Phase 2 of the algorithm then performs short cubic regularization steps for a parametrized family of least-squares functions so long as first-order criticality is not attained. These steps are computed by attempting to preserve approximate feasibility of the iterates while

producing values of the objective function that are close to a sequence of decreasing “targets.” To be specific, one or more $\text{ARC}_{(S)}$ iterations are applied to minimize the least-squares function $\Phi(x, t) \stackrel{\text{def}}{=} \frac{1}{2} \|r(x, t)\|^2$ with respect to x , where

$$(4.3) \quad r(x, t) \stackrel{\text{def}}{=} \begin{pmatrix} c(x) \\ f(x) - t \end{pmatrix}$$

and where t is a “target” value for $f(x)$. Clearly, the Jacobian $A(x, t)$ of the residual function $r(x, t)$ in (4.3) satisfies

$$(4.4) \quad A(x, t) \stackrel{\text{def}}{=} A(x) = \begin{pmatrix} J(x) \\ g(x) \end{pmatrix},$$

where $J(x)$ is the Jacobian of the constraint function $c(x)$ and $g(x)$ is the gradient of $f(x)$. Thus $\nabla_x \Phi(x, t) = A(x, t)^T r(x, t)$ and the scaled gradient (3.2) has the expression

$$(4.5) \quad g_r(x, t) \stackrel{\text{def}}{=} \begin{cases} \frac{A(x, t)^T r(x, t)}{\|r(x, t)\|} = \frac{J(x)^T c(x) + (f(x) - t)g(x)}{\|r(x, t)\|} & \text{whenever } r(x, t) \neq 0, \\ 0 & \text{otherwise.} \end{cases}$$

We are now ready to summarize our Short-Step ARC (ShS-ARC) algorithm.

Algorithm 4.1: The short-step ARC (ShS-ARC) algorithm for (4.1).

A starting point x_0 , initial regularization parameters σ_0 and σ_1 and a minimal one σ_{\min} such that $\min\{\sigma_0, \sigma_1\} \geq \sigma_{\min} > 0$, and algorithmic parameters $\gamma_2 \geq \gamma_1 > 1$ and $1 > \eta_2 \geq \eta_1 > 0$, as well as the tolerances $\epsilon_p \in (0, 1)$ and $\epsilon_d \in (0, 1)$, are given.

Phase 1:

Starting from x_0 , apply $\text{ARC}_{(S)}$ to minimize $\frac{1}{2} \|c(x)\|^2$ until a point x_1 is found such that (3.1) is satisfied, namely,

$$(4.6) \quad \|c(x_1)\| \leq \epsilon_p \quad \text{or} \quad \frac{\|J(x_1)^T c(x_1)\|}{\|c(x_1)\|} \leq \epsilon_d.$$

If $\|c(x_1)\| > \epsilon_p$, terminate [locally infeasible].

Phase 2:

1. Set $t_1 = f(x_1) - \sqrt{\epsilon_p^2 - \|c(x_1)\|^2}$ and $k = 1$.
2. For $k = 1, 2, \dots$, do:

2a. Starting from x_k , apply one iteration of $\text{ARC}_{(S)}$ to approximately minimize $\frac{1}{2} \|r(x, t_k)\|^2$ in (4.3).

2b. If $\rho_k \geq \eta_1$, do:

- If $\|g_r(x_{k+1}, t_k)\| \leq \epsilon_d$ and $r(x_{k+1}, t_k) \neq 0$, terminate.
- Else, set

$$(4.7) \quad t_{k+1} = f(x_{k+1}) - \sqrt{\|r(x_k, t_k)\|^2 - \|r(x_{k+1}, t_k)\|^2 + (f(x_{k+1}) - t_k)^2}.$$

Otherwise, set $t_{k+1} = t_k$.

Note that the monotonicity property of the $\text{ARC}_{(S)}$ iterates [11, (2.5), (3.19)] generated in step 2a of Phase 2 of ShS-ARC provides

$$(4.8) \quad \|r(x_k, t_k)\| \geq \|r(x_{k+1}, t_k)\| \quad \text{for all } k \geq 1,$$

and so the updating procedure for t_k in (4.7) is well-defined. Furthermore, (4.7) gives

$$(4.9) \quad t_k - t_{k+1} = -(f(x_{k+1}) - t_k) + \sqrt{\|r(x_k, t_k)\|^2 - \|r(x_{k+1}, t_k)\|^2 + (f(x_{k+1}) - t_k)^2}$$

for any successful $k \geq 1$, which we use to show next that the target values t_k decrease monotonically.

If we enter Phase 2 of ShS-ARC, we have $\|c(x_1)\| \leq \epsilon_p$. The next lemma proves that we remain approximately feasible for the constraints for all subsequent Phase 2 iterations.

LEMMA 4.1. *In every Phase 2 iteration $k \geq 1$ of the ShS-ARC algorithm, we have that*

$$(4.10) \quad t_k \geq t_{k+1},$$

$$(4.11) \quad f(x_k) - t_k \geq 0,$$

$$(4.12) \quad \|r(x_k, t_k)\| = \epsilon_p,$$

$$(4.13) \quad \|c(x_k)\| \leq \epsilon_p \quad \text{and} \quad |f(x_k) - t_k| \leq \epsilon_p,$$

and so $x_k \in \mathcal{C}_1$.

Proof. Due to (4.9), (4.10) follows immediately in the case when $f(x_{k+1}) \leq t_k$. Otherwise, when $f(x_{k+1}) > t_k$, conjugacy properties and (4.9) give

$$t_k - t_{k+1} = \frac{\|r(x_k, t_k)\|^2 - \|r(x_{k+1}, t_k)\|^2}{f(x_{k+1}) - t_k + \sqrt{\|r(x_k, t_k)\|^2 - \|r(x_{k+1}, t_k)\|^2 + (f(x_{k+1}) - t_k)^2}} \geq 0,$$

where in the last inequality, we also used (4.8).

Note that (4.11) holds at $k = 1$ due to the particular choice of t_1 and at $k > 1$, due to (4.7) and (4.8). Also, (4.13) follows straightforwardly from (4.12), which also provides that $x_k \in \mathcal{C}_1$ due to (4.2). It remains to prove (4.12), by induction on k . Again, the particular choice of t_1 gives (4.12) at $k = 1$. Assume now that (4.12) holds at $k > 1$, namely,

$$(4.14) \quad \|r(x_k, t_k)\| = \epsilon_p.$$

If k is an unsuccessful iteration, then $x_{k+1} = x_k$ and $t_{k+1} = t_k$ and so (4.12) is satisfied at $k + 1$. Otherwise, we have

$$\begin{aligned} (f(x_{k+1}) - t_{k+1})^2 &= \|r(x_k, t_k)\|^2 - \|r(x_{k+1}, t_k)\|^2 + (f(x_{k+1}) - t_k)^2 \\ &= \|r(x_k, t_k)\|^2 - \|c(x_{k+1})\|^2, \end{aligned}$$

where (4.11) and (4.7) give the first identity, while the second equality follows from (4.3). Thus we deduce, also using (4.3), that

$$\|r(x_{k+1}, t_{k+1})\|^2 = \|r(x_k, t_k)\|^2,$$

which concludes our induction step due to (4.14). \square

Phase 2 of the ShS-ARC terminates when

$$(4.15) \quad \|g_r(x_{k+1}, t_k)\| \leq \epsilon_d \text{ and } r(x_{k+1}, t_k) \neq 0,$$

where g_r is defined in (4.5) and $\epsilon_d \in (0, 1)$ is fixed at the start of the algorithm. Allowing different primal and dual accuracy tolerances makes sense if one considers the possibly different scalings of the (primal) residuals and (dual) gradients. The latter may occur, for instance, when the Jacobian $A(x)$ in (4.4) is not full rank, which is the case at KKT points of (4.1). The next lemma connects (4.15) to relative KKT points of (4.1) and to approximate critical points of the feasibility measure $\|c(x)\|$.

LEMMA 4.2. *For some (x, t) , assume that the scaled gradient (4.5) of $r(x, t)$ in (4.3) satisfies*

$$(4.16) \quad \|g_r(x, t)\| = \frac{\|J(x)^T c(x) + (f(x) - t)g(x)\|}{\|r(x, t)\|} \leq \epsilon_d.$$

Then either

$$(4.17) \quad \frac{\|J(x)^T c(x)\|}{\|c(x)\|} \leq \epsilon_d$$

or

$$(4.18) \quad \frac{\|J(x)^T y(x, t) + g(x)\|}{\|(y(x, t), 1)\|} \leq \epsilon_d, \text{ with } y(x, t) \stackrel{\text{def}}{=} \frac{c(x)}{|f(x) - t|}.$$

Proof. We distinguish two possible cases. First, assume that $f(x) = t$. Then (4.3) and (4.16) straightforwardly imply (4.17). Alternatively, we must have that $f(x) \neq t$. This allows us to divide in the numerator of (4.16) by $|f(x) - t|$, which then provides, also using $\|r(x, t)\| = \|(c(x), f(x) - t)\|$,

$$\left\| \frac{J(x)^T c(x)}{|f(x) - t|} + g(x) \right\| \leq \frac{\|r(x, t)\|}{|f(x) - t|} \epsilon_d = \left\| \left(\frac{c(x)}{|f(x) - t|}, 1 \right) \right\| \epsilon_d,$$

which gives (4.18). \square

The condition (4.18) is an instance of the *relative dual KKT stopping criterion*

$$(4.19) \quad \frac{\|J(x)^T y + g(x)\|}{\|(y, 1)\|} \leq \epsilon_d$$

for some multiplier $y \in \mathbb{R}^m$. The relative error condition (4.19) can be justified by means of a perturbation argument. Namely, considering the perturbations $x = x^* + \delta x$ and $y = y^* + \delta y$ to some KKT point x^* and to a corresponding multiplier y^* , a Taylor expansion and the KKT condition $J(x^*)^T y^* + g(x^*) = 0$ give that the perturbed dual feasibility residual $J(x)^T y + g(x)$ is to first order $[H(x^*) + \sum_{i=1}^m y_i^* \nabla_{xx} c_i(x^*)] \delta x + J(x^*)^T \delta y$. The presence of the multiplier y^* in the latter remainder illustrates that the size of the multiplier should not be ignored when measuring KKT equation residuals.

By Lemma 4.1, if we enter Phase 2 of ShS-ARC, we remain sufficiently close to the constraints for all subsequent iterations so that $\|c(x_k)\| \leq \epsilon_p$. This and Lemma 4.2 imply that when the ShS-ARC algorithm terminates with (4.15), then either we are close to a feasible critical point of the feasibility measure $\|c(x)\|$ (when $f(x) = t$)

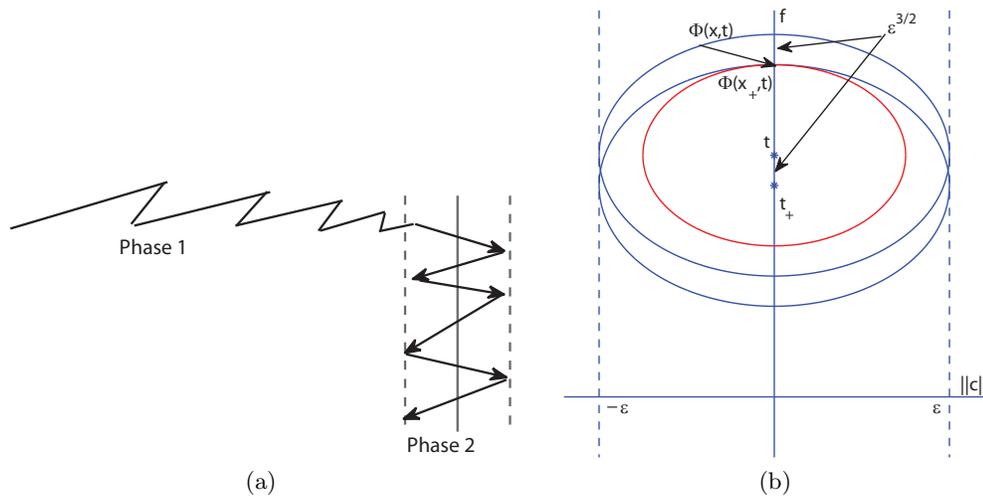


FIG. 4.1. (a) Illustration of ShS-ARC Phases 1 and 2. (b) A successful iteration of ShS-ARC's Phase 2 in the case where $\epsilon_p = \epsilon\sqrt{2}$ and $\epsilon_d = \mathcal{O}(\epsilon^{2/3})$.

or we are close to a (relative) KKT point of (4.1) (when $f(x) \neq t$). In an attempt to encourage termination with the latter rather than the former condition, one may consider allowing ShS-ARC to continue iterating when $f(x_{k+1}) = t_k$ and (4.15) hold, provided there is sufficient residual and hence target decrease (of the same order in the primal and dual tolerances as (5.7)).

In the next section, we establish that the target values t_k decrease by a fixed amount in each iteration. Thus either (4.15) holds for some k —and so we are approximately critical for (4.1) or for the constraints—or the targets reach f_* , the global minimum of f over the set of constraints, in which case again (4.15) must hold. Thus ShS-ARC will terminate; furthermore, when $\epsilon_p = \epsilon$ and $\epsilon_d = \epsilon^{2/3}$, its worst-case evaluation complexity is $\mathcal{O}(\epsilon^{-3/2})$, just like in the unconstrained case.

5. Complexity of the ShS-ARC algorithm for the equality constrained problem. Before analyzing the complexity of Algorithm ShS-ARC, we state our assumptions formally (in addition to AC.1).

AC.2 The Jacobian $J(x)$ of $c(x)$, the components $c_i(x)$ and $\nabla_{xx}c_i(x)$ are globally Lipschitz continuous on the path of all Phase 1 and Phase 2 iterates and trial points with Lipschitz constants $L_J > 0$, $L_{c_i} \geq 1$, and $L_{H,c_i} > 0$ for $i \in \{1, \dots, m\}$.

AC.3 $f(x)$, $g(x)$, and $\nabla_{xx}f(x)$ are globally Lipschitz continuous on the path of all Phase 2 iterates and trial points with (positive) Lipschitz constants L_f , $L_{g,f}$, and $L_{H,f}$, respectively.¹

AC.4 The objective $f(x)$ is bounded above and below in \mathcal{C}_1 , which is defined in (4.2); that is, there exist constants f_{low} and $f_{\text{up}} \geq f_{\text{low}} + 1$ such that

$$f_{\text{low}} \leq f(x) \leq f_{\text{up}} \quad \text{for all } x \in \mathcal{C}_1.$$

¹Note that even though $x_k \in \mathcal{C}_1$ for all Phase 2 iterations $k \geq 1$ (due to Lemma 4.1), the path of the iterates and trial points may not be included in \mathcal{C}_1 . Furthermore, upper bounds on the length of the step s_k [11, Lemma 2.2] depend on the Lipschitz constant of $f(x)$. Thus we seem unable to characterize the set that contains the Phase 2 iterates and trial points as a superset of \mathcal{C}_1 .

The assumptions AC.1–AC.4, the construction of ShS-ARC, Phase 2 iterates $x_k \in \mathcal{C}_1$ (due to Lemma 4.1), and (4.2) imply that AR.1–AR.4 hold for each of the least-squares functions that we employ in ShS-ARC, namely, $\frac{1}{2}\|c(x)\|^2$ and $\frac{1}{2}\|r(x, t_k)\|^2$ for $k \geq 1$; furthermore, the resulting constants are independent of k . In particular, the corresponding values of L_g in (2.12) for $\frac{1}{2}\|c(x)\|^2$ and $\frac{1}{2}\|r(x, t_k)\|^2$ are, respectively,

$$(5.1) \quad L_{g,1} \stackrel{\text{def}}{=} L_c^2 + \|c(x_0)\|L_J \quad \text{and} \quad L_{g,2} \stackrel{\text{def}}{=} \|(L_c, L_f)\|^2 + \kappa_c\|(L_J, L_{g,f})\|,$$

where $L_c \stackrel{\text{def}}{=} \|(L_{c_1}, \dots, L_{c_m})\| \geq 1$ is the Lipschitz constant of c , while the corresponding values of L in (2.13) for $\frac{1}{2}\|c(x)\|^2$ and $\frac{1}{2}\|r(x, t_k)\|^2$ are, respectively,

$$(5.2) \quad L_1 \stackrel{\text{def}}{=} 2L_JL_c + L_J \sum_{i=1}^m L_{c_i} + \|c(x_0)\| \sum_{i=1}^m L_{H,c_i} \quad \text{and}$$

$$(5.3) \quad L_2 \stackrel{\text{def}}{=} \|(L_J, L_{g,f})\| \left(2\|(L_c, L_f)\| + L_f + \sum_{i=1}^m L_{c_i} \right) + \kappa_c \left(L_{H,f} + \sum_{i=1}^m L_{H,c_i} \right).$$

The next lemma shows that Phase 2 of ShS-ARC consists of (at most) a constant number of unsuccessful $\text{ARC}_{(S)}$ steps followed by a successful one for minimizing $\frac{1}{2}\|r(x, t_k)\|^2$ for fixed t_k , after which t_k is decreased according to (4.7).

LEMMA 5.1. *Let AC.1–AC.4 hold, as well as AM.4 for the Hessian of $\frac{1}{2}\|r(x, t_k)\|^2$ and its approximation. Then the Phase 2 iterations of the ShS-ARC algorithm satisfy*

$$(5.4) \quad \sigma_k \leq \max(\sigma_1, \frac{3}{2}\gamma_2(L_2 + C)) \stackrel{\text{def}}{=} \bar{\sigma}_{\text{sh}} \quad \text{for all } k \geq 1,$$

where L_2 is defined in (5.3). Also, at most

$$(5.5) \quad L_{\text{sh}} \stackrel{\text{def}}{=} \left\lceil 1 + \frac{2}{\log \gamma_1} \log \left(\frac{\bar{\sigma}_{\text{sh}}}{\sigma_{\min}} \right) \right\rceil$$

ShS-ARC/ $\text{ARC}_{(S)}$ iterations are performed for each distinct target value t_k .

Proof. The implication (2.18) in Lemma 2.1 directly applies to the Phase 2 iterations of ShS-ARC, with constants $L = L_2$ defined in (5.3) and C given in AM.4, independent of k . The construction of a Phase 2 iteration of ShS-ARC and (2.6) imply that whenever σ_k is large in the sense of (2.18), we have $\sigma_{k+1} \leq \sigma_k$. Thus (5.4) follows, noting that the factor γ_2 in $\bar{\sigma}_{\text{sh}}$ is allowed for the case when σ_k is only slightly less than $3(L_2 + C)/2$ and k is not very successful, while the term σ_1 in (5.4) accounts for choices at the start of Phase 2.

Note that Theorem 2.2 directly applies to the Phase 2 iterations of ShS-ARC that employ the same target value t_k . Thus the bound (5.5) follows directly from (2.25), (5.4), the use of parameters γ_1 and σ_{\min} in Phase 2 of ShS-ARC, as well as the fact that we only take one successful ShS-ARC/ $\text{ARC}_{(S)}$ iteration for each fixed t_k (and so, here, $|\mathcal{S}_j| = 1$ in (2.25)). \square

The next lemma gives an auxiliary result to be used in Lemma 5.3.

LEMMA 5.2. *Consider the following optimization problem in two variables:*

$$(5.6) \quad \min_{(f,c) \in \mathbb{R}^2} F(f,c) \stackrel{\text{def}}{=} -f + \sqrt{\epsilon^2 - c^2} \quad \text{subject to } f^2 + c^2 \leq \tau^2,$$

where $0 < \tau < \epsilon$. The global minimum of (5.6) is attained at $(f_*, c_*) = (\tau, 0)$ and it is given by $F(f_*, c_*) = -\tau + \epsilon$.

Proof. As $F(f, c)$ is separable, linear and decreasing in f , and concave in c , the (global) solution of (5.6) is attained on the boundary of the feasible region, namely, $f_*^2 + c_*^2 = \tau^2$, and since we are minimizing, we must have $f_* = \sqrt{\tau^2 - c_*^2}$. Now,

$$F(\sqrt{\tau^2 - c^2}, c) = -\sqrt{\tau^2 - c^2} + \sqrt{\epsilon^2 - c^2} = \frac{\epsilon^2 - \tau^2}{\sqrt{\tau^2 - c^2} + \sqrt{\epsilon^2 - c^2}}$$

is strictly increasing in $|c| \in [0, \tau]$ and so its minimum is attained at $c_* = 0$. Thus $f_* = \pm\tau$, and since we are minimizing, the smallest value of $F(f, c_*)$ is at $f_* = \tau$. \square

The next lemma proves the crucial result that the targets t_k decrease by a quantity bounded below by a multiple of $\epsilon_d^{3/2} \epsilon_p^{1/2}$ at every successful Phase 2 iteration k until termination.

LEMMA 5.3. *Suppose that AC.1–AC.4 hold, as well as AM.4 for the Hessian of $\frac{1}{2}\|r(x, t_k)\|^2$ and its approximation. Set the primal and dual tolerances in that ShS-ARC algorithm such that $\epsilon_d \leq \epsilon_p^{1/3}$. Then, for every successful Phase 2 iteration $k \geq 1$ for which (4.15) fails, we have that*

$$(5.7) \quad t_k - t_{k+1} \geq \kappa_t \epsilon_d^{3/2} \epsilon_p^{1/2}$$

for some problem-dependent constant $\kappa_t \stackrel{\text{def}}{=} \min\{\alpha\beta^{3/2}, 1 - \beta\}$, where $\beta \in (0, 1)$ is any fixed problem-independent constant, $\alpha \stackrel{\text{def}}{=} \eta_1 \sigma_{\min} \kappa_{g,r}^3 / 6$ and

$$(5.8) \quad \kappa_{g,r} \stackrel{\text{def}}{=} \sqrt{(1 - \kappa_\theta) / (\frac{1}{2}L_2 + C + \bar{\sigma}_{\text{sh}} + \kappa_\theta L_{g,2})},$$

with κ_θ , L_2 , C , $\bar{\sigma}_{\text{sh}}$, and $L_{g,2}$ defined in (2.4), (5.3), (2.16), (5.4), and (5.1), respectively.

Proof. Lemma 3.1 applies to minimizing $\frac{1}{2}\|r(x, t_k)\|^2$, and so (3.3) implies that for any successful $k \geq 1$, we have

$$(5.9) \quad \|r(x_k, t_k)\| - \|r(x_{k+1}, t_k)\| \geq \kappa_t \min \left\{ \|g_r(x_{k+1}, t_k)\|^{\frac{3}{2}} \|r(x_k, t_k)\|^{\frac{1}{2}}, \|r(x_k, t_k)\| \right\},$$

where κ_t is defined below (5.7). Thus for any successful $k \geq 1$ for which (4.15) fails with $\|g_r(x_{k+1}, t_k)\| > \epsilon_d$, (5.9) becomes

$$(5.10) \quad \|r(x_k, t_k)\| - \|r(x_{k+1}, t_k)\| \geq \kappa_t \min \left\{ \epsilon_d^{3/2} \epsilon_p^{1/2}, \epsilon_p \right\} = \kappa_t \epsilon_d^{3/2} \epsilon_p^{1/2},$$

and when (4.15) fails with $r(x_{k+1}, t_k) = 0$, we trivially have that

$$(5.11) \quad \|r(x_k, t_k)\| - \|r(x_{k+1}, t_k)\| = \epsilon_p \geq \kappa_t \epsilon_d^{3/2} \epsilon_p^{1/2},$$

where in (5.10) and (5.11), we also used (4.12), $\kappa_t \in (0, 1)$, and $\epsilon_d \leq \epsilon_p^{1/3}$. Using (4.3) and the properties of the l_2 -norm, (4.9) becomes

$$(5.12) \quad \begin{aligned} t_k - t_{k+1} &= -(f(x_{k+1}) - t_k) + \sqrt{\|r(x_k, t_k)\|^2 - \|c(x_{k+1})\|^2} \\ &= -(f(x_{k+1}) - t_k) + \sqrt{\epsilon_p^2 - \|c(x_{k+1})\|^2}, \end{aligned}$$

where we used (4.12) in the second equality. It follows from (4.3) that

$$(5.13) \quad \begin{aligned} (f(x_{k+1}) - t_k)^2 + \|c(x_{k+1})\|^2 &= \|r(x_{k+1}, t_k)\|^2 \leq \left(\|r(x_k, t_k)\| - \kappa_t \epsilon_d^{3/2} \epsilon_p^{1/2} \right)^2 \\ &= \left(\epsilon_p - \kappa_t \epsilon_d^{3/2} \epsilon_p^{1/2} \right)^2, \end{aligned}$$

where in the first inequality we used (5.10) and (5.11), and in the second equality, (4.12). We now apply Lemma 5.2 to the third right-hand side of (5.12), letting $f = f(x_{k+1}) - t_k$, $c = \|c(x_{k+1})\|$, $\epsilon = \epsilon_p$, and $\tau = \epsilon_p - \kappa_t \epsilon_d^{3/2} \epsilon_p^{1/2}$. We deduce from this lemma, (5.12), and (5.13) that

$$t_k - t_{k+1} \geq -\tau + \epsilon_p = -(\epsilon_p - \kappa_t \epsilon_d^{3/2} \epsilon_p^{1/2}) + \epsilon_p = \kappa_t \epsilon_d^{3/2} \epsilon_p^{1/2},$$

which proves (5.7). \square

Figure 4.1(b) illustrates the workings of one successful Phase 2 iteration for $\epsilon \stackrel{\text{def}}{=} \epsilon_p/\sqrt{2}$ and $\epsilon_d \stackrel{\text{def}}{=} \mathcal{O}(\epsilon^{2/3})$, the case of most interest to us as it coincides with the evaluation complexity of ARC for the unconstrained case. The figure exemplifies that the amount of decrease in the target values is inherited from the merit function decrease (5.10).

Note that the ShS-ARC algorithm requires one evaluation of the objective function, its gradient (and possibly Hessian), and one evaluation of the vector of constraint functions, its Jacobian (and possibly Hessians) per iteration. We are now ready to give the main complexity result for ShS-ARC applied to (4.1).

THEOREM 5.4. *Suppose that AC.1–AC.4 hold, and that ShS-ARC is applied to minimizing (4.1) with $\epsilon_d \leq \epsilon_p^{1/3}$. Assume also that AM.4 holds for the Hessians of $\frac{1}{2}\|c(x)\|^2$ and $\frac{1}{2}\|r(x, t_k)\|^2$ and its approximations. Then the ShS-ARC algorithm generates an iterate x_k satisfying either a relative KKT condition for (4.1), namely,*

$$(5.14) \quad \|c(x_k)\| \leq \epsilon_p \quad \text{and} \quad \frac{\|J(x_k)^T y_k + g(x_k)\|}{\|(y_k, 1)\|} \leq \epsilon_d$$

for some $y_k \in \mathbb{R}^m$, or an approximate first-order criticality condition for the feasibility measure $\|c(x)\|$, namely,

$$(5.15) \quad \frac{\|J(x_k)^T c(x_k)\|}{\|c(x_k)\|} \leq \epsilon_d$$

in at most

$$(5.16) \quad \lceil \kappa_{f,c} \epsilon_d^{-3/2} \epsilon_p^{-1/2} \rceil$$

evaluations of c and f (and their derivatives), where $\kappa_{f,c} > 0$ is a problem-dependent constant, independent of ϵ_p , ϵ_d , and x_0 .

Proof. The evaluation complexity of Phase 1 follows directly from Theorem 3.2 with $\Phi(x) \stackrel{\text{def}}{=} \frac{1}{2}\|c(x)\|^2$. In particular, the evaluation complexity of obtaining x_1 is bounded above by

$$(5.17) \quad \lceil \kappa_S \max\{\kappa_1, \kappa_2\} \max\{\epsilon_d^{-3/2}, \epsilon_p^{-1/2}\} \rceil,$$

where κ_1 , κ_2 , and κ_S are defined in (3.16) and (3.17) with $r(x_0) = c(x_0)$, $L = L_1$ given in (5.2), and $L_g = L_{g,1}$ given in (5.1). If the ShS-ARC algorithm terminates at this stage, then (4.6) implies that (5.15) holds with $k = 1$ and $\|c(x_1)\| > \epsilon_p$. Assume now that Phase 2 of the ShS-ARC algorithm is entered. From AC.4 and (4.13), we have

$$f_{\text{low}} \leq f(x_k) \leq t_k + \epsilon_p \leq t_1 - i_k \kappa_t \epsilon_d^{3/2} \epsilon_p^{1/2} + \epsilon_p \leq f(x_1) - i_k \kappa_t \epsilon_d^{3/2} \epsilon_p^{1/2} + \epsilon_p,$$

where i_k is the number of successful ShS-ARC iterations from 1 to k for which (4.15) fails, and where we have also used (5.7) and the definition of t_1 in the ShS-ARC algorithm. Hence, we obtain from the inequality $f(x_1) \leq f_{\text{up}}$ (itself implied by AC.4 again) and $\epsilon_p \in (0, 1)$ that

$$(5.18) \quad i_k \leq \left\lceil \frac{f_{\text{up}} - f_{\text{low}} + 1}{\kappa_t \epsilon_d^{3/2} \epsilon_p^{1/2}} \right\rceil \stackrel{\text{def}}{=} L_{\text{sh}}^s.$$

Since for each distinct value of t_k we have one successful iteration, (5.5) in Lemma 5.1 implies that the total number of Phase 2 iterations for which (4.15) fails is bounded above by $L_{\text{sh}}^s \cdot L_{\text{sh}}$, where L_{sh} is defined as in (5.5) and L_{sh}^s is defined as in (5.18). Thus the ShS-ARC algorithm must terminate after this many iterations at most, yielding, because of Lemma 4.2, an iterate satisfying $\|c(x_k)\| \leq \epsilon_p$ and either (4.18) or (4.17); thus either (5.14) or (5.15) will hold in this case. Recalling that only one evaluation of c and f (and their derivatives, if successful) occurs per iteration, the bound (5.16) now follows by summing up the Phase 1 and Phase 2 iteration bounds, and using that $\epsilon_p \in (0, 1)$, which gives that the Phase 2 bound dominates in the order of (ϵ_p, ϵ_d) . \square

If $\epsilon_d \stackrel{\text{def}}{=} \epsilon_p^{2/3}$, then Theorem 5.4 implies that the evaluation complexity of ShS-ARC is at most $\mathcal{O}(\epsilon_p^{-3/2})$, the same as for the unconstrained case. However, if $\epsilon_d \stackrel{\text{def}}{=} \epsilon_p$, then this complexity bound worsens to $\mathcal{O}(\epsilon_p^{-2})$, the same in order as for steepest-descent-type methods for both constrained and unconstrained problems [12, 4].

6. Conclusions. We have shown that with an appropriate and practical termination condition, the (optimal) cubic regularization variant $\text{ARC}_{(\text{S})}$ takes at most $\mathcal{O}(\epsilon^{-3/2})$ evaluations to drive the residual or the scaled gradient of the potentially singular least-squares problem (1.1) below ϵ . Our analysis has focused on the Euclidean norm case, but it can be easily extended to general inner products and induced norms, and to smooth l_p -norms for $p > 2$. Though the order $\epsilon^{-3/2}$ of the ARC bound is optimal for unconstrained optimization when second-order methods are employed [5], and it is sharp for nonlinear least-squares when ensuring (1.2) is sufficiently small, it is unclear whether it is optimal or even sharp for $\text{ARC}_{(\text{S})}$ with the novel termination condition (3.1).

For the equality-constrained potentially nonconvex programming problem (4.1), we presented a target-following technique ShS-ARC that applies $\text{ARC}_{(\text{S})}$ to target-dependent least-squares merit functions tracking a path of approximately feasible points (if an initial such point can be found). Furthermore, in order to ensure approximate first-order conditions for (4.1) or for a feasibility measure—within ϵ_p for the constraint feasibility and within ϵ_d for dual (first-order) feasibility—ShS-ARC requires at most $\mathcal{O}(\epsilon_d^{-3/2} \epsilon_p^{-1/2})$ problem evaluations, which depending on the choice of tolerances ϵ_p and ϵ_d can take any value between the complexity $\mathcal{O}(\epsilon_p^{-3/2})$ of ARC (namely, when $\epsilon_d = \epsilon_p^{2/3}$) and $\mathcal{O}(\epsilon_p^{-2})$ of steepest descent (when $\epsilon_d = \epsilon_p$). Though it is natural for the primal and dual feasibility residuals to vary at different rates, and hence require different optimality tolerances (with higher accuracy for primal feasibility than for dual being common), it is an open question at the moment whether an algorithm for nonconvexly constrained problems can be devised that has worst-case evaluation complexity of order $\epsilon^{-3/2}$, where $\epsilon = \epsilon_p = \epsilon_d$. Also, extending ShS-ARC or other cubic regularization approaches to problems with nonconvex inequality constraints remains to be considered.

Acknowledgment. We thank two anonymous referees for their useful comments and for suggesting the succinct proof of Lemma 5.2.

REFERENCES

- [1] 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 J. Numer. Anal., 48 (2010), pp. 1–29.
- [2] Å. BJÖRCK, *Numerical Methods for Least Squares Problems*, SIAM, Philadelphia, 1996.
- [3] R. H. BYRD, R. B. SCHNABEL, AND G. A. SHULTZ, *A trust region algorithm for nonlinearly constrained optimization*, SIAM J. Numer. Anal., 24 (1987), pp. 1152–1170.
- [4] C. CARTIS, N. I. M. GOULD, AND PH. L. TOINT, *On the complexity of finding first-order critical points in constrained nonlinear programming*, Math. Program., DOI 10.1007/S10107-012-0617-9, 2012 (online).
- [5] C. CARTIS, N. I. M. GOULD, AND PH. L. TOINT, *Optimal Newton-Type Methods for Nonconvex Smooth Optimization Problems*, ERGO Technical Report 11-009, School of Mathematics, University of Edinburgh, 2011.
- [6] C. CARTIS, N. I. M. GOULD, AND PH. L. TOINT, *Evaluation complexity of adaptive cubic regularization methods for convex unconstrained optimization*, Optim. Methods Software, 27 (2012), pp. 197–219.
- [7] 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 J. Numer. Anal., 32 (2012), pp. 1662–1695.
- [8] C. CARTIS, N. I. M. GOULD, AND PH. L. TOINT, *On the oracle complexity of first-order and derivative-free algorithms for smooth nonconvex minimization*, SIAM J. Optim., 22 (2012), pp. 66–86.
- [9] 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 J. Optim., 21 (2011), pp. 1721–1739.
- [10] C. CARTIS, N. I. M. GOULD, AND PH. L. TOINT, *Adaptive cubic regularisation methods for unconstrained optimization. Part II: Worst-case function- and derivative-evaluation complexity*, Math. Program., 130 (2011), pp. 295–319.
- [11] C. CARTIS, N. I. M. GOULD, AND PH. L. TOINT, *Adaptive cubic regularisation methods for unconstrained optimization. Part I: Motivation, convergence and numerical results*, Math. Program., 127 (2011), pp. 245–295.
- [12] 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 problems*, SIAM J. Optim., 20 (2010), pp. 2833–2852.
- [13] M. R. CELIS, *A Trust Region Strategy for Nonlinear Equality Constrained Optimization*, Technical Report 85-4, Department of Computational and Applied Mathematics, Rice University, Houston, TX, 1985.
- [14] A. R. CONN, N. I. M. GOULD, AND PH. L. TOINT, *Trust-Region Methods*, SIAM, Philadelphia, 2000.
- [15] J. E. DENNIS AND J. J. MORÉ, *A characterization of superlinear convergence and its application to quasi-Newton methods*, Math. Comp., 28 (1974), pp. 549–560.
- [16] J. E. DENNIS AND R. B. SCHNABEL, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, Prentice-Hall, Englewood Cliffs, NJ, 1983. Reprinted as Classics Appl. Math. 16, SIAM, Philadelphia, 1996.
- [17] R. FLETCHER AND S. LEYFFER, *Nonlinear programming without a penalty function*, Math. Programming, 91 (2002), pp. 239–270.
- [18] N. I. M. GOULD AND PH. L. TOINT, *Nonlinear programming without a penalty function or a filter*, Math. Program., 122 (2010), pp. 155–196.
- [19] A. GRIEWANK, *The Modification of Newton’s Method for Unconstrained Optimization by Bounding Cubic Terms*, Technical Report NA/12 (1981), Department of Applied Mathematics and Theoretical Physics, University of Cambridge, UK, 1981.
- [20] J.-B. HIRIART-URRUTY AND C. LEMARÉCHAL, *Convex Analysis and Minimization Algorithms I*, Springer-Verlag, Berlin, Heidelberg, 1993.
- [21] YU. NESTEROV, *Modified Gauss-Newton scheme with worst case guarantees for global performance*, Optim. Methods Software, 22 (2007), pp. 469–483.
- [22] YU. NESTEROV AND B. T. POLYAK, *Cubic regularization of Newton’s method and its global performance*, Math. Program., 108 (2006), pp. 177–205.

- [23] J. NOCEDAL AND S. J. WRIGHT, *Numerical Optimization*, 2nd ed., Springer-Verlag, New York, 2006.
- [24] E. O. OMOJOKUN, *Trust Region Algorithms for Optimization with Nonlinear Equality and Inequality Constraints*, Ph.D. Thesis, University of Colorado, Boulder, CO, 1989.
- [25] M. J. D. POWELL AND Y. YUAN, *A trust region algorithm for equality constrained optimization*, *Math. Programming*, 49 (1990), pp. 189–213.
- [26] A. VARDI, *A trust region algorithm for equality constrained minimization: Convergence properties and implementation*, *SIAM J. Numer. Anal.*, 22 (1985), pp. 575–591.
- [27] M. WEISER, P. DEUFLHARD, AND B. ERDMANN, *Affine conjugate adaptive Newton methods for nonlinear elastomechanics*, *Optim. Methods Softw.*, 22 (2007), pp. 413–431.