A GLOBALLY CONVERGENT STABILIZED SQP METHOD: SUPERLINEAR CONVERGENCE

Philip E. Gill* Vyacheslav Kungurtsev† Daniel P. Robinson‡

UCSD Center for Computational Mathematics
Technical Report CCoM-14-1
June 30, 2014

Abstract

Regularized and stabilized sequential quadratic programming (SQP) methods are two classes of methods designed to resolve the numerical and theoretical difficulties associated with ill-posed or degenerate nonlinear optimization problems. Recently, a regularized SQP method has been proposed that allows convergence to points satisfying certain second-order KKT conditions (*SIAM J. Optim.*, 23(4):1983–2010, 2013). The method is formulated as a regularized SQP method with an implicit safeguarding strategy based on minimizing a bound-constrained primal-dual augmented Lagrangian. The method involves a flexible line search along a direction formed from the solution of a regularized quadratic programming subproblem and, when one exists, a direction of negative curvature for the primal-dual augmented Lagrangian. With an appropriate choice of termination condition, the method terminates in a finite number of iterations under weak assumptions on the problem functions. Safeguarding becomes relevant only when the iterates are converging to an infeasible stationary point of the norm of the constraint violations. Otherwise, the method terminates with a point that either satisfies the second-order necessary conditions for optimality, or fails to satisfy a weak second-order constraint qualification. The purpose of this paper is to establish the conditions under which this second-order regularized SQP algorithm is equivalent to the stabilized SQP method. It is shown that under conditions that are no stronger than those required by conventional stabilized SQP methods, the regularized SQP method has superlinear local convergence. The required convergence properties are obtained by allowing a small relaxation of the optimality conditions for the quadratic programming subproblem in the neighborhood of a solution.

Key words. Nonlinear programming, nonlinear constraints, augmented Lagrangian, sequential quadratic programming, SQP methods, stabilized SQP, regularized methods, primal-dual methods.

AMS subject classifications. 49J20, 49J15, 49M37, 49D37, 65F05, 65K05, 90C30

*Department of Mathematics, University of California, San Diego, La Jolla, CA 92093-0112 (pgill@ucsd.edu). Research supported by National Science Foundation grant DMS-0915220 and Department of Energy grant DE-SC0002349.

†Agent Technology Center, Department of Computer Science, Faculty of Electrical Engineering, Czech Technical University in Prague. (vyacheslav.kungurtsev@fel.cvut.cz) Research supported by the European social fund within the framework of realizing the project "Support of inter-sectoral mobility and quality enhancement of research teams at the Czech Technical University in Prague", CZ.1.07/2.3.00/30.0034.

‡Department of Applied Mathematics and Statistics, Johns Hopkins University, Baltimore, MD 21218-2682 (daniel.p.robinson@jhu.edu). Research supported by National Science Foundation grant DMS-1217153.

1. Introduction

This paper is concerned with computing solutions to the nonlinear optimization problem

$$\text{(NP)} \qquad \underset{x\in\mathbb{R}^n}{\text{minimize}} \;\; f(x) \quad \text{subject to} \;\; c(x) = 0, \quad x \geq 0,$$

where $c\colon \mathbb{R}^n \mapsto \mathbb{R}^m$ and $f\colon \mathbb{R}^n \mapsto \mathbb{R}$ are twice-continuously differentiable. This problem format assumes that all general inequality constraints have been converted to equalities by the use of slack variables. Methods for solving problem (NP) are easily extended to the more general setting with $l \leq x \leq u$.

This paper concerns the local convergence properties of the algorithm `pdSQP2` proposed by Gill, Kungurtsev and Robinson [13]. Algorithm `pdSQP2` is an extension of the first-order regularized SQP algorithm (`pdSQP`) of Gill and Robinson [15] that is designed to encourage convergence to points satisfying the second-order conditions for optimality. The method is based on the properties of the bound-constrained optimization problem:

$$\underset{x,y}{\text{minimize}} \;\; M(x, y\,; y^E, \mu) \quad \text{subject to} \;\; x \geq 0, \tag{1.1}$$

where $M(x, y\,; y^E, \mu)$ is the primal-dual augmented Lagrangian function:

$$M(x, y\,; y^E, \mu) = f(x) - c(x)^T y^E + \frac{1}{2\mu}\|c(x)\|_2^2 + \frac{\nu}{2\mu}\|c(x) + \mu(y - y^E)\|_2^2, \tag{1.2}$$

with ν a fixed nonnegative scalar, μ a positive penalty parameter, and y^E an estimate of a Lagrange multiplier vector y^*. Gill, Kungurtsev and Robinson [13, Theorem 1.3] show that if μ is sufficiently small, then a primal-dual solution (x^*, y^*) of (NP) satisfying the second-order sufficient conditions is a minimizer of $M(x, y\,; y^*, \mu)$. (The value of ν is fixed throughout the computation, and is not included as an argument of M. The choice of ν and its effect on the properties of M are discussed in [15,31] (in particular, see Gill and Robinson [15, Table 1]).)

Algorithm `pdSQP2` has an inner/outer iteration structure, with the inner iterations being those of an active-set method for solving a quadratic programming (QP) subproblem based on minimizing a local quadratic model of the primal-dual function (1.2) subject to the nonnegativity constraints on x. The outer iteration defines the QP subproblem and performs a flexible line search along a direction given by an approximate solution of the QP subproblem and, if needed, an approximate direction of negative curvature for the primal-dual function M. At the kth primal-dual iterate (x_k, y_k), the outer iteration involves the definition of a vector y_k^E that estimates the vector of Lagrange multipliers, and two penalty parameters μ_k^R and μ_k such that $\mu_k^R \leq \mu_k$, with $\mu_k^R \ll \mu_k$ in the neighborhood of a solution. The parameter μ_k^R plays the role of a small *regularization* parameter and is used in the definition of a local quadratic model of $M(x, y\,; y_k^E, \mu_k^R)$. The penalty parameter μ_k is used in the line search, where the function $M(x, y\,; y_k^E, \mu_k)$ is used as a line-search merit function for obtaining improved estimates of both the primal and dual variables.

With an appropriate choice of termination criteria, algorithm `pdSQP2` terminates in a finite number of iterations under weak assumptions on the problem functions. As in the first-order case, the method is formulated as a regularized SQP method with an augmented Lagrangian safeguarding strategy. The safeguarding becomes relevant only when the iterates are converging to an infeasible stationary point of the norm of the constraint violations. Otherwise, the method terminates with a point that either satisfies the second-order necessary conditions for optimality, or fails to satisfy a weak second-order constraint qualification.

An important property of the local quadratic model of $M(x, y\,; y_k^E, \mu_k^R)$ is that the associated bound-constrained QP subproblem is equivalent to the general QP:

$$\begin{array}{ll} \underset{x,y}{\text{minimize}} & g(x_k)^T(x - x_k) + \frac{1}{2}(x - x_k)^T H(x_k, y_k)(x - x_k) + \frac{1}{2}\mu_k^R\|y\|_2^2 \\ \text{subject to} & c(x_k) + J(x_k)(x - x_k) + \mu_k^R(y - y_k^E) = 0, \qquad x \geq 0, \end{array} \tag{1.3}$$

where g is the gradient of f, J is the Jacobian of c, and H is the Hessian of the Lagrangian function. (The definition of the local quadratic model is discussed in Section 2.2.) If $y_k^E = y_k$, the QP (1.3) is typical of those used in stabilized SQP methods (see, e.g., Hager [19], Wright [32, 33, 35], Oberlin and Wright [30], Fernández and Solodov [9], Izmailov and Solodov [23]). Moreover, the definition of μ_k^R implies that the direction provided by the solution of the QP (1.3) is an $O(\mu_k^R)$ estimate of the conventional SQP direction. Stabilized SQP has been shown to have desirable local convergence properties. In particular, superlinear convergence to primal-dual KKT points has been shown to occur without the need for a constraint qualification. The purpose of this paper is to establish the conditions under which the primal-dual regularized SQP algorithm is equivalent to the stabilized SQP method. The principal contributions are the following.

- A method is proposed for general inequality constrained optimization that has provable global convergence to either a second-order point or an infeasible stationary point. In the case of convergence to a second-order point, a local convergence analysis is given that does not require the assumption of a constraint qualification or strict complementarity condition.

- Although exact second-derivatives are used, the method does not require the solution of an indefinite quadratic programming subproblem—a process that is known to be NP-hard. In addition, the local convergence theory makes no assumptions about which local solution of the QP subproblem is computed (see Kungurtsev [26, Chapter 5] for a discussion of these issues).

- Close to the solution, the method defines iterates that are equivalent to a conventional stabilized SQP method. This equivalence holds under conditions that are no stronger than those required for the superlinear local convergence of a conventional stabilized SQP method.

- Preliminary numerical results indicate that the method has strong global and local convergence properties for both degenerate and nondegenerate problems under weak regularity assumptions.

The method has a number of features that are not shared by conventional stabilized methods: (a) either a first- or second-order model function is used in the line search, as needed; (b) a flexible line search is used; (c) the regularization and penalty parameters may increase or decrease at any iteration; (d) criteria are specified that allow the acceptance of certain inexact solutions to the QP subproblem; and (e) a local descent step may be generated based on allowing a small relaxation of the optimality conditions for the QP subproblem.

The remainder of the paper is organized as follows. This section concludes with a summary of the notation and terminology. Section 2 provides details of the second-order primal-dual regularized SQP method. The global and local convergence properties of the method are discussed in Sections 3 and 4 respectively.

1.1. Notation and terminology

Unless explicitly indicated otherwise, $\|\cdot\|$ denotes the vector two-norm or its induced matrix norm. The inertia of a real symmetric matrix A, denoted by $\operatorname{In}(A)$, is the integer triple (a_+, a_-, a_0) giving the number of positive, negative and zero eigenvalues of A. The least eigenvalue of a symmetric matrix A will be denoted by $\lambda_{\min}(A)$. Given vectors a and b with the same dimension, the vector with ith component $a_i b_i$ is denoted by $a \cdot b$. Similarly, $\min(a, b)$ is a vector with components $\min(a_i, b_i)$. The vectors e and e_j denote, respectively, the column vector of ones and the jth column of the identity matrix I. The dimensions of e, e_i and I are defined by the context. Given vectors x and y, the vector consisting of the elements of x augmented by elements of y is denoted by (x, y). The ith component of a vector labeled with a subscript will be denoted by $[\,\cdot\,]_i$, e.g., $[\,v\,]_i$ is the ith component of the vector v. For a given ℓ-vector u and index set $\mathcal{S}$, the quantity $[\,u\,]_{\mathcal{S}}$ denotes the subvector of components u_j such that $j \in \mathcal{S} \cap \{1, 2, \ldots, \ell\}$. Similarly, if M is a symmetric $\ell \times \ell$ matrix, then $[\,M\,]_{\mathcal{S}}$ denotes the symmetric matrix with elements m_{ij} for $i, j \in \mathcal{S} \cap \{1, 2, \ldots, \ell\}$. A local solution of problem (NP) is denoted by x^*. The vector $g(x)$ is used to denote $\nabla f(x)$, the gradient of $f(x)$. The matrix $J(x)$ denotes the $m \times n$ constraint Jacobian, which has ith row $\nabla c_i(x)^T$, the gradient of the ith constraint function $c_i(x)$. The Lagrangian function associated with (NP) is $L(x, y, z) = f(x) - c(x)^T y - z^T x$, where y and z are m- and n-vectors of dual variables associated with the equality constraints and bounds, respectively. The Hessian of the Lagrangian with respect to x is denoted by $H(x, y) = \nabla^2 f(x) - \sum_{i=1}^m y_i \nabla^2 c_i(x)$.

Let $\{\alpha_j\}_{j\ge 0}$ be a sequence of scalars, vectors or matrices and let $\{\beta_j\}_{j\ge 0}$ be a sequence of positive scalars. If there exists a positive constant γ such that $\|\alpha_j\| \le \gamma\beta_j$, we write $\alpha_j = O\big(\beta_j\big)$. If there exists a sequence $\{\gamma_j\} \to 0$ such that $\|\alpha_j\| \le \gamma_j\beta_j$, we say that $\alpha_j = o(\beta_j)$. If there exist positive constants γ_1 and γ_2 such that $\gamma_1\beta_j \le \|\alpha_j\| \le \gamma_2\beta_j$, we write $\alpha_j = \Theta\big(\beta_j\big)$.

2. The Primal-Dual SQP Method

The proposed method is based on the properties of a local quadratic model of the primal-dual function (1.2). At the kth iterate $v_k = (x_k, y_k)$, the model is

$$\mathcal{Q}_k(v\,; \mu) = (v - v_k)^T \nabla M(v_k\,; y^E_k, \mu) + \tfrac{1}{2}(v - v_k)^T B(v_k\,; \mu)(v - v_k), \tag{2.1}$$

where $B(v\,; \mu) \equiv B(x, y\,; \mu)$ approximates the Hessian of $M(x, y\,; y^E, \mu)$. The vector $\nabla M(v_k\,; y^E_k, \mu)$ and matrix $\nabla^2 M(x, y\,; y^E, \mu)$ are written in the form

$$\nabla M(x, y\,; y^E, \mu) = \begin{pmatrix} g(x) - J(x)^T\big(\pi(x) + \nu(\pi(x) - y)\big) \\ \nu\mu\big(y - \pi(x)\big) \end{pmatrix} \tag{2.2}$$

and

$$\nabla^2 M(x, y\,; y^E, \mu) = \begin{pmatrix} H\big(x, \pi(x) + \nu(\pi(x) - y)\big) + \frac{1}{\mu}(1 + \nu)J(x)^T J(x) & \nu J(x)^T \\ \nu J(x) & \nu\mu I \end{pmatrix}, \tag{2.3}$$

where $\pi(x) = \pi(x\,; y^E, \mu)$ denotes the first-order multiplier estimate

$$\pi(x\,; y^E, \mu) = y^E - \frac{1}{\mu}c(x). \tag{2.4}$$

The quadratic model (2.1) is based the approximation $B(v\,;\mu) \equiv B(x,y\,;\mu) \approx \nabla^2 M(x,y\,;y^E,\mu)$ defined by replacing $H\big(x,\pi(x)+\nu(\pi(x)-y)\big)$ by $H(x,y)$ in the leading block of $\nabla^2 M$. This gives the matrix

$$B(x,y\,;\mu) = \begin{pmatrix} H(x,y) + \dfrac{1}{\mu}(1+\nu)J(x)^T J(x) & \nu J(x)^T \\ \nu J(x) & \nu\mu I \end{pmatrix}, \tag{2.5}$$

which is independent of the multiplier estimate y^E.

From a numerical stability perspective, it is important that every computation be performed without forming the matrix $B(v_k\,;\mu)$ explicitly. All the relevant properties of a matrix B of the form (2.5) may be determined from either of the matrices

$$H(x,y) + \frac{1}{\mu}(1+\nu)J(x)^T J(x) \quad \text{or} \quad \begin{pmatrix} H(x,y) & J(x)^T \\ J(x) & -\mu I \end{pmatrix}.$$

These matrices are said to have "augmented Lagrangian form" and "regularized KKT form," respectively. To simplify the analysis, the formulation and analysis of the main algorithm is given in terms of matrices in augmented Lagrangian form. However, all practical computations are performed with the matrix in regularized KKT form. In particular, each iteration involves the factorization of a matrix of the form

$$K_{\mathcal{F}} = \begin{pmatrix} H_{\mathcal{F}}(x,y) & J_{\mathcal{F}}(x)^T \\ J_{\mathcal{F}}(x) & -\mu_k^R I \end{pmatrix}, \tag{2.6}$$

where quantities with the suffix "F" are defined in terms of index sets that estimate the active and free sets of primal-dual variables at a solution of problem (NP). At any nonnegative x, the active set and free set are defined as

$$\mathcal{A}(x) = \{ i : [x]_i = 0 \} \quad \text{and} \quad \mathcal{F}(x) = \{1,2,\dots,n+m\} \setminus \mathcal{A}(x), \tag{2.7}$$

respectively. (The algorithm considered in this paper does not impose bounds on the dual variables. Some implications of enforcing bounds on both the primal and dual variables are discussed by Gill and Robinson [15].) The active set at x is estimated by the "ϵ-active" set

$$\mathcal{A}_\epsilon(x,y,\mu) = \big\{ i : [x]_i \le \epsilon, \;\; \text{with} \;\; \epsilon = \min\big(\epsilon_a,\, \max\big(\mu, r(x,y)^\gamma\big)\big) \big\}, \tag{2.8}$$

where μ is a positive parameter, γ $(0 < \gamma < 1)$ and ϵ_a $(0 < \epsilon_a < 1)$ are fixed scalars, and the quantity

$$r(x,y) = \| \big(c(x),\, \min\big(x,\, g(x) - J(x)^T y\big)\big) \| \tag{2.9}$$

is a practical estimate of the distance of (x,y) to a first-order KKT point of problem (NP) (see Definition 3.1). Analogous to $\mathcal{F}(x)$, the "ϵ-free" set is the complement of $\mathcal{A}_\epsilon$ in the set of indices of the primal-dual variables, i.e.,

$$\mathcal{F}_\epsilon(x,y,\mu) = \{1,2,\dots,n+m\} \setminus \mathcal{A}_\epsilon(x,y,\mu). \tag{2.10}$$

The algorithm involves three main calculations: (i) the computation of quantities needed for the definition of the regularization parameter μ_k^R and multiplier estimate y_k^E; (ii) the computation of a feasible descent direction based on an approximate solution of a QP subproblem; and (iii) the calculation of the line-search direction and the definition of an appropriate step length using a flexible line search. A description of each of these calculations is given in the following three sections.

2.1. The regularization parameter and multiplier estimate

The definition of the local quadratic model for the QP subproblem depends on the values of the regularization parameter μ_k^R and the vector of multiplier estimates y_k^E. These values are defined in terms of the values μ_{k-1}^R and y_{k-1}^E from the previous iteration, and a scalar-vector pair $(\xi_k^{(1)}, u_{\mathcal{F}_\epsilon})$ that estimates the least negative eigenvalue of $H_{\mathcal{F}_\epsilon} + (1/\mu_{k-1}^R) J_{\mathcal{F}_\epsilon}^T J_{\mathcal{F}_\epsilon}$, where $H_{\mathcal{F}_\epsilon}$ and $J_{\mathcal{F}_\epsilon}$ are defined in terms of the primal variables with indices in the ϵ-free set, i.e., $H_{\mathcal{F}_\epsilon}$ is the matrix of ϵ-free rows and columns of $H_k = H(x_k, y_k)$, and $J_{\mathcal{F}_\epsilon}$ is the matrix of ϵ-free columns of $J_k = J(x_k)$.

The principal calculations for $(\xi_k^{(1)}, u_{\mathcal{F}_\epsilon})$ are defined in Algorithm 1, which computes a feasible direction of negative curvature (if one exists) for the local quadratic model $\mathcal{Q}_k(v\,; \mu_{k-1}^R)$. If $H_{\mathcal{F}_\epsilon} + (1/\mu_{k-1}^R) J_{\mathcal{F}_\epsilon}^T J_{\mathcal{F}_\epsilon}$ is not positive semidefinite, Algorithm 1 computes a nonzero direction $u_{\mathcal{F}_\epsilon}$ that satisfies

$$u_{\mathcal{F}_\epsilon}^T \Big(H_{\mathcal{F}_\epsilon} + \frac{1}{\mu_{k-1}^R} J_{\mathcal{F}_\epsilon}^T J_{\mathcal{F}_\epsilon} \Big) u_{\mathcal{F}_\epsilon} \le \theta \lambda_{\min} \Big(H_{\mathcal{F}_\epsilon} + \frac{1}{\mu_{k-1}^R} J_{\mathcal{F}_\epsilon}^T J_{\mathcal{F}_\epsilon} \Big) \|u_{\mathcal{F}_\epsilon}\|^2 < 0, \tag{2.11}$$

where θ is a positive scalar independent of x_k. The inequality (2.11) is used to characterize the properties of a direction of negative curvature for the quadratic model $\mathcal{Q}_k(v_k\,; \mu_{k-1}^R)$ (2.1) of the primal-dual function M, when $H_{\mathcal{F}_\epsilon} + (1/\mu_{k-1}^R) J_{\mathcal{F}_\epsilon}^T J_{\mathcal{F}_\epsilon}$ is not positive semidefinite (see [13, Lemma 2.1(ii)]). The quantity $u_{\mathcal{F}_\epsilon}$ may be computed from several available matrix factorizations as described in [16, Section 4]. For a theoretical analysis, it is necessary to assume only that the $u_{\mathcal{F}_\epsilon}$ associated with a nonzero value of $\xi_k^{(1)}$ computed in Algorithm 1 satisfies (2.11).

Algorithm 1 Curvature estimate.

1: **procedure** `CURVATURE ESTIMATE`$(x_k, y_k, \mu_{k-1}^R, J_k, H_k)$
2: Compute $H_{\mathcal{F}_\epsilon}$ and $J_{\mathcal{F}_\epsilon}$ as submatrices of H_k and J_k associated with $\mathcal{F}_\epsilon(x_k, y_k, \mu_{k-1}^R)$;
3: **if** $(H_{\mathcal{F}_\epsilon} + (1/\mu_{k-1}^R) J_{\mathcal{F}_\epsilon}^T J_{\mathcal{F}_\epsilon}$ is positive semidefinite$)$ **then**
4: $\xi_k^{(1)} = 0$; $u_k^{(1)} = 0$; $w_k^{(1)} = 0$;
5: **else**
6: Compute $u_{\mathcal{F}_\epsilon} \ne 0$ such that (2.11) holds for some positive θ independent of x_k;
7: $\xi_k^{(1)} = -u_{\mathcal{F}_\epsilon}^T \big(H_{\mathcal{F}_\epsilon} + (1/\mu_{k-1}^R) J_{\mathcal{F}_\epsilon}^T J_{\mathcal{F}_\epsilon}\big) u_{\mathcal{F}_\epsilon} / \|u_{\mathcal{F}_\epsilon}\|^2 > 0$;
8: $u_k^{(1)} = 0$; $[\,u_k^{(1)}\,]_{\mathcal{F}_\epsilon} = u_{\mathcal{F}_\epsilon}$;
9: $w_k^{(1)} = -(1/\mu_{k-1}^R) J_k u_k^{(1)}$;
10: **end if**
11: $s_k^{(1)} = (u_k^{(1)}, w_k^{(1)})$;
12: **return** $(s_k^{(1)}, \xi_k^{(1)})$;
13: **end procedure**

In Step 9 of Algorithm 1 the change $w_k^{(1)}$ in the dual variables is computed as a function of the change $u_k^{(1)}$ in the primal variables, which ensures that the equality constraints associated with the stabilized QP subproblem (1.3) are satisfied. This definition is sufficient to provide a feasible direction of negative curvature $s_k^{(1)}$ for the quadratic model $\mathcal{Q}_k(v\,; \mu_{k-1}^R)$ (2.1) when $\xi_k^{(1)} > 0$ (see [13, Lemma 2.1(ii)]).

The multiplier estimate y_k^E is updated to be the dual iterate y_k if there is improvement in a measure of the distance to a primal-dual second-order solution (x^*, y^*). Algorithm 5 uses feasibility and optimality measures $\eta(x_k)$ and $\omega(x_k, y_k)$ such that

$$\eta(x_k) = \|c(x_k)\| \;\text{ and }\; \omega(x_k, y_k) = \max\Big(\big\|\min(x_k,\, g(x_k) - J(x_k)^T y_k)\big\|\,,\; \xi_k^{(1)}\Big), \tag{2.12}$$

where the quantity $\xi_k^{(1)}$ is defined by Algorithm 1. Given $\eta(x_k)$ and $\omega(x_k, y_k)$, weighted combinations of the feasibility and optimality measures are defined as

$$\phi_V(x_k, y_k) = \eta(x_k) + \beta\omega(x_k, y_k) \qquad \text{and} \qquad \phi_O(x_k, y_k) = \beta\eta(x_k) + \omega(x_k, y_k),$$

where β is a fixed scalar such that $0 < \beta \ll 1$. The update $y_k^E = y_k$ is performed if

$$\phi_V(v_k) \le \tfrac{1}{2}\phi_V^{\max} \qquad \text{or} \qquad \phi_O(v_k) \le \tfrac{1}{2}\phi_O^{\max}, \tag{2.13}$$

where $\phi_V^{\max}$ and $\phi_O^{\max}$ are positive bounds that are reduced as the iterations proceed. The point (x_k, y_k) is called a "V-iterate" or an "O-iterate" if it satisfies the bound on $\phi_V(v_k)$ or $\phi_O(v_k)$. A "V-O iterate" is a point at which one or both of these conditions holds, and the associated iteration (or iteration index) is called a "V-O iteration." At a V-O iteration, the regularization parameter is updated as

$$\mu_k^R = \begin{cases} \min\big(\,\mu_0^R,\; \max\big(r_k, \xi_k^{(1)}\big)^\gamma\,\big) & \text{if } \max\big(r_k, \xi_k^{(1)}\big) > 0; \\ \tfrac{1}{2}\mu_{k-1}^R & \text{otherwise,} \end{cases} \tag{2.14}$$

where $r_k = r(x_k, y_k)$ is defined by (2.9). It must be emphasized that the sequence $\{\mu_k^R\}$ is not necessarily monotonic.

If the conditions for a V-O iteration do not hold, a test is made to determine if (x_k, y_k) is an approximate second-order solution of the bound-constrained problem

$$\underset{x,y}{\text{minimize}}\;\; M(x, y\,; y_{k-1}^E, \mu_{k-1}^R) \quad \text{subject to} \quad x \ge 0. \tag{2.15}$$

In particular, (x_k, y_k) is tested using the conditions:

$$\|\min\big(x_k, \nabla_{\!x} M(x_k, y_k\,; y_{k-1}^E, \mu_{k-1}^R)\big)\| \le \tau_{k-1}, \tag{2.16a}$$

$$\|\nabla_{\!y} M(x_k, y_k\,; y_{k-1}^E, \mu_{k-1}^R)\| \le \tau_{k-1}\mu_{k-1}^R, \;\text{ and} \tag{2.16b}$$

$$\xi_k^{(1)} \le \tau_{k-1}, \tag{2.16c}$$

where τ_{k-1} is a positive tolerance. If these conditions are satisfied, then (x_k, y_k) is called an "M-iterate" and the parameters are updated as in a typical conventional augmented Lagrangian method, with the multiplier estimate y_{k-1}^E replaced by the safeguarded value

$$y_k^E = \max\big(-y_{\max}e,\; \min(\,y_k,\; y_{\max}e\,)\big) \tag{2.17}$$

for some large positive scalar constant $y_{\max}$, and the new regularization parameter is given by

$$\mu_k^R = \begin{cases} \min\big(\,\tfrac{1}{2}\mu_{k-1}^R,\; \max\big(r_k, \xi_k^{(1)}\big)^\gamma\,\big), & \text{if } \max(r_k, \xi_k^{(1)}) > 0; \\ \tfrac{1}{2}\mu_{k-1}^R, & \text{otherwise.} \end{cases} \tag{2.18}$$

In addition, the tolerance τ_{k-1} is decreased by a constant factor. Conditions analogous to (2.16) are used in the first-order method of Gill and Robinson [16], in which case numerical experiments indicate that M-iterations occur infrequently relative to the total number of iterations.

Finally, if neither (2.13) nor (2.16) are satisfied, then $y_k^E = y_{k-1}^E$ and $\mu_k^R = \mu_{k-1}^R$. As the multiplier estimates and regularization parameter are *fixed* at their current values in this case, (x_k, y_k) is called an "F-iterate".

2.2. Computation of the descent direction

Given a primal-dual iterate $v_k = (x_k, y_k)$, and values of the regularization parameter μ_k^R and multiplier estimate y_k^E, the main calculation involves finding a primal-dual direction d_k such that $\nabla M_k^T d_k < 0$. The vector d_k is defined as either a "local" descent direction or "global" descent direction. In general terms, a local descent direction is the solution of a system of linear equations, whereas a global descent direction is the solution of a convex QP subproblem. If certain conditions hold at the start of the step computation, the local descent direction is computed. Additional conditions are then used to decide if the computed direction should be used for d_k. If the local descent direction is not computed, or not selected, then the QP subproblem is solved for the global descent direction. However, if a local descent direction is computed, the matrix factors used to define it are used to compute the first iterate of the QP method. Both the global descent direction and local descent direction may be regarded as being approximate solutions of a quadratic program based on minimizing the quadratic model (2.1) subject to the bounds on x.

The global descent direction. The underlying quadratic model (2.1) is not suitable as the QP objective function because $B(v_k\,; \mu_k^R)$ is not positive definite in general. An indefinite QP subproblem can have many local minima and may be unbounded. In addition, the certification of a second-order solution of an indefinite QP is computationally intractable in certain cases. In order to avoid these difficulties, a strictly convex QP subproblem is defined that incorporates the most important properties of the nonconvex quadratic model (2.1). The subproblem is

$$\begin{aligned} \underset{v}{\text{minimize}} \quad & \widehat{\mathcal{Q}}_k(v) = (v - v_k)^T \nabla M(v_k\,; y_k^E, \mu_k^R) + \tfrac{1}{2}(v - v_k)^T \widehat{B}(v_k\,; \mu_k^R)(v - v_k) \\ \text{subject to} \quad & [\,v\,]_i \geq 0, \quad i = 1, 2, \dots, n, \end{aligned} \tag{2.19}$$

where $\widehat{B}(v_k\,; \mu_k^R)$ is the matrix

$$\widehat{B}(v_k\,; \mu_k^R) = \begin{pmatrix} \widehat{H}(x_k, y_k) + \dfrac{1}{\mu_k^R}(1+\nu) J(x_k)^T J(x_k) & \nu J(x_k)^T \\ \nu J(x_k) & \nu \mu_k^R I \end{pmatrix}, \tag{2.20}$$

with $\widehat{H}(x_k, y_k)$ a modification of $H(x_k, y_k)$ chosen to make $\widehat{B}(v_k\,; \mu_k^R)$ positive definite. The matrix $\widehat{B}$ is defined by a process known as "convexification" (see [16, Section 4] for details). As discussed in the introduction to this section, in practice, the matrix $\widehat{B}(v_k\,; \mu_k^R)$ is never computed explicitly.

Solution of the QP subproblem. Given an initial feasible point $\widehat{v}_k^{(0)}$ for problem (2.19), active-set methods generate a feasible sequence $\{\widehat{v}_k^{(j)}\}_{j>0}$ such that $\widehat{\mathcal{Q}}_k(\widehat{v}_k^{(j)}) \leq \widehat{\mathcal{Q}}_k(\widehat{v}_k^{(j-1)})$

and $\widehat{v}_k^{(j)}$ minimizes $\widehat{\mathcal{Q}}_k(v)$ on a "working set" $\mathcal{W}_j$ of bound constraints. An iterate $\widehat{v}_k^{(j)}$ is optimal for (2.19) if the Lagrange multipliers for the bound constraints in the working set are nonnegative, i.e.,

$$[\,\nabla\widehat{\mathcal{Q}}_k(\widehat{v}_k^{(j)})\,]_W = [\,\nabla M(v_k\,;y_k^E,\mu_k^R) + \widehat{B}(v_k\,;\mu_k^R)(\widehat{v}_k^{(j)} - v_k)\,]_W \ge 0, \tag{2.21}$$

where the suffix "W" denotes the components of $\widehat{v}_k^{(j)}$ corresponding to the bound constraints in the working set. The first working set $\mathcal{W}_0$ is defined as the ϵ-active set $\mathcal{A}_\epsilon(x_k,y_k,\mu_k^R)$, which defines the initial feasible point $\widehat{v}_k^{(0)}$ as

$$[\,\widehat{v}_k^{(0)}\,]_{\mathcal{A}_\epsilon} = 0, \quad \text{and} \quad [\,\widehat{v}_k^{(0)}\,]_{\mathcal{F}_\epsilon} = [\,v_k\,]_{\mathcal{F}_\epsilon}, \tag{2.22}$$

where the suffices "A_ϵ" and "F_ϵ" refer to the components associated with the ϵ-active and ϵ-free sets at (x_k,y_k). In general, $\widehat{v}_k^{(0)}$ will not minimize $\widehat{\mathcal{Q}}_k(v)$ on $\mathcal{W}_0$, and an estimate of $\widehat{v}_k^{(1)}$, is defined by solving the subproblem:

$$\underset{v}{\text{minimize}} \;\; \widehat{\mathcal{Q}}_k(v) \quad \text{subject to} \;\; [\,v\,]_W = 0. \tag{2.23}$$

This estimate, if feasible, is used to define $\widehat{v}_k^{(1)}$, otherwise one of the violated bounds is added to the working set and the iteration is repeated. Eventually, the working set will include enough bounds to define an appropriate minimizer $\widehat{v}_k^{(1)}$.

The local descent direction. A minimizer (when it exists) of the general quadratic model $\mathcal{Q}_k(v\,;\mu_k^R)$ (2.1) on the ϵ-active set $\mathcal{A}_\epsilon(x_k,y_k,\mu_k^R)$ plays a vital role in the formulation of a method with superlinear local convergence. If the matrix $B_{\mathcal{F}_\epsilon}$ of ϵ-free rows and columns of $B(v_k\,;\mu_k^R)$ is positive definite, then the QP subproblem

$$\underset{v}{\text{minimize}} \;\; \mathcal{Q}_k(v\,;\mu_k^R) \quad \text{subject to} \;\; [\,v\,]_{\mathcal{A}_\epsilon} = 0 \tag{2.24}$$

has a unique solution. Given the initial QP iterate $\widehat{v}_k^{(0)}$ defined in (2.22), the solution of subproblem (2.24) may be computed as $\widehat{v}_k^{(0)} + \Delta\widehat{v}_k^{(0)}$, where $[\,\Delta\widehat{v}_k^{(0)}\,]_{\mathcal{A}_\epsilon} = 0$ and $[\,\Delta\widehat{v}_k^{(0)}\,]_{\mathcal{F}_\epsilon}$ satisfies the equations

$$B_{\mathcal{F}_\epsilon}[\,\Delta\widehat{v}_k^{(0)}\,]_{\mathcal{F}_\epsilon} = -[\,\nabla\mathcal{Q}_k(\widehat{v}_k^{(0)}\,;\mu_k^R)\,]_{\mathcal{F}_\epsilon}. \tag{2.25}$$

If d_k denotes the step from v_k to the solution of (2.24), then $d_k = \widehat{v}_k^{(0)} + \Delta\widehat{v}_k^{(0)} - v_k$. Forming the ϵ-free and ϵ-active components of d_k gives

$$[\,d_k\,]_{\mathcal{F}_\epsilon} = [\,\Delta\widehat{v}_k^{(0)}\,]_{\mathcal{F}_\epsilon} \;\text{ and }\; [\,d_k\,]_{\mathcal{A}_\epsilon} = -[\,v_k\,]_{\mathcal{A}_\epsilon} = -[\,x_k\,]_{\mathcal{A}_\epsilon} \le 0, \tag{2.26}$$

where the last inequality follows from the feasibility of x_k with respect to the bounds. Overall, the components of d_k satisfy $[\,\nabla\mathcal{Q}_k(v_k + d_k\,;\mu_k^R)\,]_{\mathcal{F}_\epsilon} = [\,\nabla M(v_k\,;y_k^E,\mu_k^R) + B_k d_k\,]_{\mathcal{F}_\epsilon} = 0$.

Consider the matrix

$$U_{\mathcal{F}_\epsilon} = \begin{pmatrix} I & -\frac{(1+\nu)}{\nu\mu_k^R} J_{\mathcal{F}_\epsilon}(x_k)^T \\ 0 & \frac{1}{\nu} I \end{pmatrix},$$

where $J_{\mathcal{F}_\epsilon}(x_k)$ denotes the matrix of ϵ-free columns of $J(x_k)$. The matrix $U_{\mathcal{F}_\epsilon}$ is nonsingular and can be applied to both sides of (2.25) without changing the solution. Using the definitions

(2.26) and performing some simplification yields

$$\begin{pmatrix} H_{\mathcal{F}_\epsilon}(x_k, y_k) & J_{\mathcal{F}_\epsilon}(x_k)^T \\ J_{\mathcal{F}_\epsilon}(x_k) & -\mu_k^R I \end{pmatrix} \begin{pmatrix} [\, p_k \,]_{\mathcal{F}_\epsilon} \\ -q_k \end{pmatrix} = - \begin{pmatrix} [\, g(x_k) + H(x_k, y_k)(\widehat{x}_k^{(0)} - x_k) - J(x_k)^T y_k \,]_{\mathcal{F}_\epsilon} \\ c(x_k) + J(x_k)(\widehat{x}_k^{(0)} - x_k) + \mu_k^R (y_k - y_k^E) \end{pmatrix}, \tag{2.27}$$

where p_k and q_k are the vectors of primal and dual components of d_k, and $H_{\mathcal{F}_\epsilon}(x_k, y_k)$ is the matrix of ϵ-free rows and columns of $H(x_k, y_k)$. If $\mathcal{A}_\epsilon(x_k, y_k, \mu_k^R) = \mathcal{A}(x^*)$, $[\, x_k \,]_{\mathcal{A}_\epsilon} = 0$, and $y_k^E = y_k$, then $\widehat{x}_k^{(0)} = x_k$ and these equations become

$$\begin{pmatrix} H_{\mathcal{F}_\epsilon}(x_k, y_k) & J_{\mathcal{F}_\epsilon}(x_k)^T \\ J_{\mathcal{F}_\epsilon}(x_k) & -\mu_k^R I \end{pmatrix} \begin{pmatrix} [\, p_k \,]_{\mathcal{F}_\epsilon} \\ -q_k \end{pmatrix} = - \begin{pmatrix} [\, g(x_k) - J(x_k)^T y_k \,]_{\mathcal{F}_\epsilon} \\ c(x_k) \end{pmatrix}, \tag{2.28}$$

which represent the dual-regularized Newton equations for minimizing M on the ϵ-active set. This property provides the motivation for the definition of a "local descent direction" in the neighborhood of a solution. If $B_{\mathcal{F}_\epsilon}$ is positive definite and v_k is a V-O iterate (so that $y_k^E = y_k$), the solution of (2.24) is considered as a potential approximate solution of a QP subproblem with objective $\mathcal{Q}_k(v\,; \mu_k^R)$. Given a small positive scalar t_k defined below in (2.30), the local descent direction d_k is used if the conditions

$$v_k + d_k \ \text{ feasible}, \qquad [\, \nabla \mathcal{Q}_k(v_k + d_k\,; \mu_k^R) \,]_{\mathcal{A}_\epsilon} \ge -t_k e, \quad \text{ and } \quad \nabla M_k^T d_k < 0 \tag{2.29}$$

are satisfied. The QP gradient condition of (2.29) relaxes the QP optimality condition (2.21) and reduces the possibility of unnecessary inner iterations in the neighborhood of a solution. As the outer iterates converge, $\mathcal{A}_\epsilon(x_k, y_k, \mu_k^R)$ provides an increasingly accurate estimate of the optimal active set. Nevertheless, close to a solution, the QP multipliers associated with variables with zero multipliers at x^* may have small negative values. The associated constraints will be removed by the QP algorithm, only to be added again at the start of the next QP subproblem. This inefficiency is prevented by the condition (2.29), which defines any small negative multiplier as being optimal.

The magnitude of the feasibility parameter t_k is based on the proximity measure $r(x_k, y_k)$ of (2.9), in particular,

$$t_k = r(x_k, y_k)^\lambda, \quad \text{where} \quad 0 < \lambda < \min\{\gamma, 1 - \gamma\} < 1, \tag{2.30}$$

with γ the parameter used in the definition (2.8) of the regularization parameter.

2.3. Definition of the line-search direction and step length

The line search defines a nonnegative step length α_k for a direction formed from the sum of two vectors d_k and s_k. The vector $d_k = (p_k, q_k) = (\widehat{x}_k - x_k, \widehat{y}_k - y_k)$ is a descent direction for $M(v\,; y_k^E, \mu_k^R)$ at v_k. The vector s_k, if nonzero, is a direction of negative curvature for the quadratic model $\mathcal{Q}_k(v\,; \mu_{k-1}^R)$. (If α_k is zero, then $(x_{k+1}, y_{k+1}) = (x_k, y_k)$ and at least one of the parameters μ_k^R or y_k^E are changed before the next line search.) If $\mu_k^R \ne \mu_{k-1}^R$, then s_k is not guaranteed to be a negative curvature direction for $\mathcal{Q}_k(v\,; \mu_k^R)$, which must be taken into account in the line-search procedure. (In particular, see the use of the quantity ξ_k^R defined in Step 12 of Algorithm 4.)

The vector $s_k^{(1)} = (u_k^{(1)}, w_k^{(1)})$, which is computed as a by-product of the computation of $\xi_k^{(1)}$ in Algorithm 1, is a direction of negative curvature for the approximate Hessian

Algorithm 2 Computation of the primal-dual descent direction d_k.

procedure DESCENT DIRECTION$(x_k, y_k, y_k^E, \mu_k^R)$
 Constants: $0 < \lambda < \min\{\gamma, 1-\gamma\} < 1$;
 $B = B(x_k, y_k\,; \mu_k^R)$; Compute a positive-definite $\widehat{B}$ from B;
 $\nabla M_k = \nabla M(x_k, y_k\,; y_k^E, \mu_k^R)$; $t_k = r(x_k, y_k)^\lambda$;
 $[\,\widehat{v}_k^{(0)}\,]_{\mathcal{A}_\epsilon} = 0$; $[\,\widehat{v}_k^{(0)}\,]_{\mathcal{F}_\epsilon} = [\,v_k\,]_{\mathcal{F}_\epsilon}$;
 if ($B_{\mathcal{F}_\epsilon}$ is positive definite **and** v_k is a V-O iterate) **then**
  $[\,\Delta\widehat{v}_k^{(0)}\,]_{\mathcal{A}_\epsilon} = 0$; Solve $B_{\mathcal{F}_\epsilon}[\,\Delta\widehat{v}_k^{(0)}\,]_{\mathcal{F}_\epsilon} = -[\,\nabla\mathcal{Q}_k(\widehat{v}_k^{(0)}\,; \mu_k^R)\,]_{\mathcal{F}_\epsilon}$; $\widehat{v}_k = \widehat{v}_k^{(0)} + \Delta\widehat{v}_k^{(0)}$;
  $d_k = \widehat{v}_k - v_k$;
  if ($v_k + d_k$ is feasible **and** $\nabla M_k^T d_k < 0$ **and** $[\,\nabla\mathcal{Q}_k(v_k + d_k\,; \mu_k^R)\,]_{\mathcal{A}_\epsilon} \ge -t_k e$) **then**
   return d_k; [local descent direction]
  end if
 end if
 $[\,\Delta\widehat{v}_k^{(0)}\,]_{\mathcal{A}_\epsilon} = 0$; Solve $\widehat{B}_{\mathcal{F}_\epsilon}[\,\Delta\widehat{v}_k^{(0)}\,]_{\mathcal{F}_\epsilon} = -[\,\nabla\widehat{\mathcal{Q}}_k(\widehat{v}_k^{(0)})\,]_{\mathcal{F}_\epsilon}$;
 Compute $\widehat{v}_k^{(1)}$ such that $\widehat{v}_k^{(1)} = \widehat{v}_k^{(0)} + \widehat{\alpha}_0\Delta\widehat{v}_k^{(0)}$ is feasible for some $\widehat{\alpha}_0 \ge 0$;
 Solve the convex QP (2.19) for $\widehat{v}_k$, starting at $\widehat{v}_k^{(1)}$;
 $d_k = \widehat{v}_k - v_k$; [global descent direction]
 return d_k;
end procedure

$B(x_k, y_k\,; \mu_{k-1}^R)$. The line-search direction of negative curvature s_k is a rescaled version of $s_k^{(1)}$. The purpose of the rescaling is to give an s_k that: (i) satisfies a normalization condition based on the approximate curvature; and (ii) gives a combined vector $d_k + s_k$ that is a feasible non-ascent direction for the merit function. The direction s_k is zero if no negative curvature is detected, but s_k must be nonzero if $\xi_k^{(1)} > 0$ and $d_k = 0$ (see [13, Lemma 2.4]), which ensures that the line-search direction is nonzero at a first-order stationary point v_k at which $B(x_k, y_k\,; \mu_{k-1}^R)$ is not positive semidefinite.

Algorithm 3 Feasible direction of negative curvature.

procedure CURVATURE DIRECTION$(s_k^{(1)},\ \xi_k^{(1)},\ x_k,\ d_k)$
 $s_k^{(2)} = \begin{cases} -s_k^{(1)}, & \text{if } \nabla M(v_k\,; y_k^E, \mu_k^R)^T s_k^{(1)} > 0; \\ \ \ s_k^{(1)}, & \text{otherwise;} \end{cases}$
 $p_k = d_k(1:n)$; $u_k^{(2)} = s_k^{(2)}(1:n)$;
 $\sigma_k = \operatorname{argmax}_{\sigma\ge 0}\ \big\{\sigma : x_k + p_k + \sigma u_k^{(2)} \ge 0,\ \|\sigma u_k^{(2)}\| \le \max(\xi_k^{(1)}, \|p_k\|)\big\}$;
 $s_k = \sigma_k s_k^{(2)}$; [scaled curvature direction]
 return s_k;
end procedure

Algorithm 4 uses the flexible line search of Gill and Robinson [13], which is an augmented Lagrangian version of the flexible line search proposed by Curtis and Nocedal [3]. Given a primal-dual search direction $\Delta v_k = d_k + s_k$, and a line-search penalty parameter μ, an Armijo condition is used to define a reduction in the function $\Psi_k(\alpha\,; \mu) = M(v_k + \alpha\Delta v_k\,; y_k^E, \mu)$ that is at least as good as the reduction in the line-search model function

$$\psi_k(\alpha\,;\mu,\iota) = \Psi_k(0\,;\mu) + \alpha\Psi_k'(0\,;\mu) + \tfrac{1}{2}(\iota-1)\alpha^2 \min\big(0, \Delta v_k^T B(x_k, y_k\,;\mu_{k-1}^R)\Delta v_k\big), \quad (2.31)$$

where Ψ_k' denotes the derivative with respect to α. The scalar ι_k is either 1 or 2, depending on the order of the line-search model function. The value $\iota_k = 1$ implies that ψ_k is an affine function, which gives a first-order line search model. The value $\iota_k = 2$ defines a quadratic ψ_k and gives a second-order line search model. The first-order line search model is used when $d_k \neq 0$, $s_k = 0$, and (x_k, y_k) is a V-O iterate. This is crucial for the proof that the line-search algorithm returns the unit step length in the neighborhood of a second-order solution (see Theorem 4.2).

Just prior to the line search, the line-search penalty parameter is increased if necessary to ensure that $\mu_k \geq \mu_k^R$, i.e.,

$$\mu_k = \max(\mu_k^R, \mu_k). \quad (2.32)$$

Given a fixed parameter $\gamma_s \in (0, \frac{1}{2})$, the flexible line search attempts to compute an α_k that satisfies the modified Armijo condition

$$\Psi_k(0\,;\mu_k^F) - \Psi_k(\alpha_k\,;\mu_k^F) \geq \gamma_s\big(\,\psi_k(0\,;\mu_k^R,\iota_k) - \psi_k(\alpha_k\,;\mu_k^R,\iota_k)\,\big) \quad (2.33)$$

for some $\mu_k^F \in [\mu_k^R, \mu_k]$. The required step is found by repeatedly reducing α_k by a constant factor until either $\rho_k(\alpha_k\,;\mu_k,\iota_k) \geq \gamma_s$ or $\rho_k(\alpha_k\,;\mu_k^R,\iota_k) \geq \gamma_s$, where

$$\rho_k(\alpha\,;\mu,\iota) = \big(\Psi_k(0\,;\mu) - \Psi_k(\alpha\,;\mu)\big)/\big(\,\psi_k(0\,;\mu_k^R,\iota) - \psi_k(\alpha\,;\mu_k^R,\iota)\,\big).$$

The Armijo procedure is not executed if either d_k and s_k are zero (see Step 4 of Algorithm 4), or if $d_k = 0$, $s_k \neq 0$, and $\xi_k^R \leq \gamma_s\xi_k^{(1)}$ (see Steps 12 and 13 of Algorithm 4). The requirement that $\xi_k^R \leq \gamma_s\xi_k^{(1)}$ ensures that any nonzero direction of negative curvature for the approximate Hessian $B(v_k\,;\mu_{k-1}^R)$ is a direction of negative curvature for the exact Hessian $\nabla^2 M(v_k\,;y_k^E,\mu_k^R)$.

Once α_k has been found, the penalty parameter for the next iteration is updated as

$$\mu_{k+1} = \begin{cases} \mu_k, & \text{if } \rho_k(\alpha_k\,;\mu_k,\iota_k) \geq \gamma_s, \text{ or } d_k = s_k = 0, \text{ or } \alpha_k = 0; \\ \max\big(\tfrac{1}{2}\mu_k, \mu_k^R\big), & \text{otherwise.} \end{cases} \quad (2.34)$$

The aim is to decrease the penalty parameter only when the merit function defined with μ_k is not sufficiently reduced by the trial step.

2.4. Algorithm summary

The regularized second-order SQP algorithm is given in Algorithm 5. Two termination criteria are defined in terms of a positive stopping tolerance τ_{stop}. The first test is applied at the beginning of an iteration and determines if the primal-dual pair (x_k, y_k) is an approximate second-order KKT point. In this case, the algorithm is terminated if

$$r(x_k, y_k) \leq \tau_{\text{stop}}, \quad \xi_k^{(1)} \leq \tau_{\text{stop}}, \quad \text{and} \quad \mu_{k-1}^R \leq \tau_{\text{stop}}. \quad (2.35)$$

The second test is applied after the definition of the parameters $(y_k^E,\, \mu_k^R)$ and determines if an infeasible x_k approximately minimizes $\|c(x)\|$ subject to $x \geq 0$. If the following conditions hold:

$$\min\big(\|c(x_k)\|, \tau_{\text{stop}}\big) > \mu_k^R, \quad \|\min\big(x_k, J(x_k)^T c(x_k)\big)\| \leq \tau_{\text{stop}}, \quad \text{and } x_k \text{ an M-iterate}, \quad (2.36)$$

Algorithm 4 Flexible line search.

procedure SEARCH($d_k, s_k, y_k^E, \mu_k, \mu_k^R, \mu_{k-1}^R, \iota_k$)
  Constant: $\gamma_s \in (0, \frac{1}{2})$;
  Compute $\nabla M = \nabla M(x_k, y_k\,; y_k^E, \mu_k^R)$.
  if $s_k = 0$ **and** $d_k = 0$ **then**
    $\alpha_k = 1$;
  else if ($d_k \neq 0$ **or** $\nabla M^T s_k < 0$ **or** $\mu_k^R = \mu_{k-1}^R$) **then**
    $\alpha_k = 1$;
    while $\rho_k(\alpha_k\,; \mu_k^R, \iota_k) < \gamma_s$ **and** $\rho_k(\alpha_k\,; \mu_k, \iota_k) < \gamma_s$ **do**
      $\alpha_k = \frac{1}{2}\alpha_k$;
    end while
  else [$d_k = 0,\ s_k \neq 0,\ \xi_k^{(1)} > 0$]
    $\xi_k^R = -\, s_k^T \nabla^2 M(v_k\,; y_k^E, \mu_k^R) s_k / \|u_k\|^2$; [by definition, $s_k = (u_k, w_k)$]
    if $\xi_k^R > \gamma_s \xi_k^{(1)}$ **then**
      $\alpha_k = 1$;
      while $\rho_k(\alpha_k\,; \mu_k^R, \iota_k) < \gamma_s$ **and** $\rho_k(\alpha_k\,; \mu_k, \iota_k) < \gamma_s$ **do**
        $\alpha_k = \frac{1}{2}\alpha$;
      end while
    else
      $\alpha_k = 0$;
    end if
  end if
  return $\alpha_k \geq 0$
end procedure

then the iterations are terminated and x_k is regarded as an infeasible stationary point of the constraint violations.

3. Global Convergence

This section concerns the global convergence of the sequence $\{(x_k, y_k)\}$ generated by Algorithm 5. The purpose is to establish that the use of a local descent direction in the neighborhood of a solution does not invalidate the global convergence results of Gill, Kungurtsev and Robinson [13].

Definition 3.1. (First- and second-order KKT points for problem (NP)**)** *A vector x^* is a* first-order KKT point *if there exists a vector y^* such that $r(x^*, y^*) = \|\big(c(x^*), \min\big(x^*, g(x^*) - J(x^*)^T y^*\big)\big)\| = 0$. If, in addition, (x^*, y^*) satisfies the condition $d^T H(x^*, y^*) d \geq 0$ for all d such that $J(x^*)d = 0$, with $d_i = 0$ for $i \in \mathcal{A}(x^*)$, then (x^*, y^*) is a* second-order KKT point.

In general, the Lagrange multiplier at a KKT point is not unique, and the set of Lagrange multiplier vectors is given by

$$\mathcal{Y}(x^*) = \{y \in \mathbb{R}^m : (x^*, y) \text{ satisfies } r(x^*, y) = 0\}. \tag{3.1}$$

As in [13], the proof of global convergence requires the following assumptions.

Algorithm 5 Second-order primal-dual SQP algorithm.

procedure PDSQP2(x_1, y_1)
  Constants: $\{\tau_{\text{stop}}, \gamma_s\} \subset (0, \frac{1}{2})$, $0 < \gamma < 1$, and $\{k_{\max}, \nu\} \subset (0, \infty)$;
  Choose $y_0^E \in \mathbb{R}^m$, $\{\tau_0, \phi_V^{\max}, \phi_O^{\max}\} \subset (0, \infty)$, and $0 < \mu_0^R \le \mu_1 < \infty$;
  for $k = 1 : k_{\max}$ **do**
    Compute the ϵ-free set $\mathcal{F}_\epsilon(x_k, y_k, \mu_{k-1}^R)$ given by (2.10);
    $J_k = J(x_k)$; $H_k = H(x_k, y_k)$;
    $\big(s_k^{(1)}, \xi_k^{(1)}\big) =$ CURVATURE$(x_k, y_k, \mu_{k-1}^R, J_k, H_k)$;   [Algorithm 1]
    Compute $r(x_k, y_k)$ from (2.9);
    if (termination criterion (2.35) holds) **then**
      return the approximate second-order KKT point (x_k, y_k);
    end if
    if $(\phi_V(x_k, y_k) \le \frac{1}{2}\phi_V^{\max})$ **then**   [V-iterate]
      $\phi_V^{\max} = \frac{1}{2}\phi_V^{\max}$; $y_k^E = y_k$; $\tau_k = \tau_{k-1}$;
      Set μ_k^R as in (2.14);
    else if $(\phi_O(x_k, y_k) \le \frac{1}{2}\phi_O^{\max})$ **then**   [O-iterate]
      $\phi_O^{\max} = \frac{1}{2}\phi_O^{\max}$; $y_k^E = y_k$; $\tau_k = \tau_{k-1}$;
      Set μ_k^R as in (2.14);
    else if $((x_k, y_k)$ satisfies (2.16a)–(2.16c)$)$ **then**   [M-iterate]
      Set y_k^E as in (2.17); $\tau_k = \frac{1}{2}\tau_{k-1}$;
      Set μ_k^R as in (2.18);
    else   [F-iterate]
      $y_k^E = y_{k-1}^E$; $\tau_k = \tau_{k-1}$; $\mu_k^R = \mu_{k-1}^R$;
    end if
    if (termination criterion (2.36) holds) **then**
      exit with the approximate infeasible stationary point x_k.
    end if
    $d_k =$ DESCENT DIRECTION$(x_k, y_k, y_k^E, \mu_k^R)$;   [Algorithm 2]
    $s_k =$ CURVATURE DIRECTION$(s_k^{(1)}, \xi_k^{(1)}, x_k, d_k)$;   [Algorithm 3]
    if $(d_k \ne 0$ **and** $s_k = 0$ **and** (x_k, y_k) is a V-O iterate$)$ **then**
      $\iota_k = 1$;
    else
      $\iota_k = 2$;
    end if
    $\mu_k = \max(\mu_k^R, \mu_k)$; $\alpha_k =$ SEARCH$(d_k, s_k, y_k^E, \mu_k, \mu_k^R, \mu_{k-1}^R, \iota_k)$;   [Algorithm 4]
    Set μ_{k+1} as in (2.34);
    $v_{k+1} = (x_{k+1}, y_{k+1}) = v_k + \alpha_k d_k + \alpha_k s_k$;
  end for
end procedure

Assumption 3.1. *The sequence of matrices* $\{\widehat{H}(x_k, y_k)\}_{k\geq 0}$ *is chosen to satisfy*

$$\|\widehat{H}(x_k, y_k)\| \leq \widehat{H}_{\max} \quad \text{and} \quad \lambda_{\min}\Big(\widehat{H}(x_k, y_k) + \frac{1}{\mu_k^R} J(x_k)^T J(x_k)\Big) \geq \underline{\lambda}_{\min},$$

for some positive $\widehat{H}_{\max}$ *and* $\underline{\lambda}_{\min}$*, and all* $k \geq 0$.

Assumption 3.2. *The functions* f *and* c *are twice continuously differentiable.*

Assumption 3.3. *The sequence* $\{x_k\}_{k\geq 0}$ *is contained in a compact set.*

The statement of each of the results established below refers to the original proof given by Gill, Kungurtsev and Robinson [13].

Lemma 3.1. ([13, Lemma 2.5]) *Let* d_k *and* s_k *be the vectors computed at the kth iteration of Algorithm 5.*

(i) *If* $d_k = 0$*, then*

  (a) $\min(x_k, g(x_k) - J(x_k)^T y_k) = 0$ *and* $\pi(x_k, y_k^E, \mu_k^R) = y_k$*;*

  (b) *if* (x_k, y_k) *is a V-O iterate, then* $r(x_k, y_k) = 0$*;*

  (c) *if* (x_k, y_k) *is an M-iterate such that* $\|y_k\|_\infty \leq y_{\max}$*, then* $r(x_k, y_k) = 0$.

(ii) *If* $d_k = s_k = 0$*, then* $\xi_k^{(1)} = 0$*,* (x_k, y_k) *is not an F-iterate, and* $\mu_k^R < \mu_{k-1}^R$.

Proof. If $d_k = 0$ then $\nabla M(v_k\,; y_k^E, \mu_k^R)^T d_k = 0$ so that the conditions of Step 9 in Algorithm 2 for the early termination of the strictly convex QP (2.19) are not satisfied. This implies that d_k is computed from Step 16, which makes it identical to the step used in [13, Lemma 2.5]. ∎

Gill, Kungurtsev and Robinson [13, Lemma 2.6] establishes the main properties of the flexible line search and line-search penalty parameter update. In particular, it is shown that the line search is guaranteed to terminate finitely. The next result indicates that the same properties hold for Algorithm 5.

Lemma 3.2. ([13, Lemma 2.6]) *If Assumption 3.1 is satisfied, then the following hold:*

(i) *The while-loops given by Steps 8 and 15 of Algorithm 4 terminate with a positive* α_k.

(ii) *If* $\mu_k < \mu_{k-1}$ *for some* $k \geq 1$*, then either the while loop given by Step 8 or the while loop given by Step 15 of Algorithm 4 was executed.*

(iii) *If* $\alpha_k = 0$*, then* (x_k, y_k) *is not an F-iterate.*

Proof. Part (ii) follows from (2.34), and part (iii) concerns the properties of a global descent direction and thus follows directly from [13, Lemma 2.6]. It remains to establish part (i). As the step computation in Algorithm 5 is the same as that in [13] when a global descent direction is computed, it may be assumed for the remainder of the proof that d_k is a local descent direction. It then follows that Step 9 of Algorithm 2 gives $\nabla M(v_k\,; y_k^E, \mu_k^R)^T d_k < 0$, and Δv_k is a descent direction for $M(v_k\,; y_k^E, \mu_k^R)$. In particular, Δv_k is nonzero and the proof of part (i) follows exactly as in [13, Subcase 1 of Case 1 in Lemma 2.6]. ∎

The following simple argument may be used to show that the proof of [13, Theorem 3.1] applies for Algorithm 5. The theorem establishes a number of results that hold under the assumption that (x_k, y_k) is an F-iterate for all k sufficiently large. However, in this situation, (x_k, y_k) cannot be a V-O iterate and d_k must be a global descent direction. Since the definition of a global descent direction is identical to the descent direction used in [13], it follows that [13, Theorem 3.1] holds for Algorithm 5.

All the auxiliary results needed for the global convergence proof given in [13] have now been established. The global convergence result, which is a combination of [13, Theorems 3.10 and 3.11] is stated below. The statement involves three constraint qualifications: the constant positive generator constraint qualification (CPGCQ); the Mangasarian-Fromovitz constraint qualification (MFCQ); and the weak constant rank condition (WCRC). Definitions of these constraint qualifications may be found, for example, in [13, Definitions 3.4–3.6].

Theorem 3.1. *Let Assumptions 3.1–3.3 hold. For the sequence $\{(x_k, y_k)\}$ of primal-dual iterates generated by Algorithm 5, define the set*

$$\mathcal{S} = \{\, k : \ (x_k, y_k) \ \textit{is a V-O iterate} \,\}. \tag{3.2}$$

Then, one of the following cases (A) *or* (B) *holds.*

(A): *The set $\mathcal{S}$ is infinite, in which case the following results hold.*

(1) *There exists a subsequence $\mathcal{S}_1 \subseteq \mathcal{S}$ and a vector x^* such that $\lim_{k\in\mathcal{S}_1} x_k = x^*$.*

(2) *Either x^* fails to satisfy the CPGCQ, or x^* is a first-order KKT point for problem* (NP)*.*

(3) *If x^* is a first-order KKT point for problem* (NP)*, then the following hold.*

(a) *If the MFCQ holds at x^*, then the sequence $\{y_k\}_{k\in\mathcal{S}_1}$ is bounded, and every limit point y^* defines a first-order KKT pair (x^*, y^*) for problem* (NP)*.*

(b) *If the MFCQ and WCRC hold at x^*, then (x^*, y^*) is a second-order KKT point.*

(B): *The set $\mathcal{S}$ is finite, in which case the set of M-iterate indices is infinite, and every limit point x^* of $\{x_k\}_{k\in\mathcal{M}}$ satisfies $c(x^*) \neq 0$ and is a KKT-point for the feasibility problem of minimizing $\|c(x)\|$ subject to $x \geq 0$.*

The next section focuses on the local convergence properties of sequences that converge to first- or second-order KKT points of problem (NP), i.e., to points considered in Case (A) of Theorem 3.1. An analysis of the rate of convergence associated with sequences converging to locally infeasible points (i.e., sequences considered in Case (B)) is beyond the scope of this paper (see, e.g., [2, 10, 27]).

4. Local Convergence

The local convergence analysis involves second-order sufficient conditions defined in terms of the set of strongly-active variables

$$\mathcal{A}_+(x, y) = \{i : [\, g(x) - J(x)^T y \,]_i > 0\}, \tag{4.1}$$

and the set of weakly-active variables

$$\mathcal{A}_0(x,y) = \{i : [\,g(x) - J(x)^T y\,]_i = 0\}. \tag{4.2}$$

Definition 4.1. (Second-order sufficient conditions (SOSC)) *A primal-dual pair* (x^*, y^*) *satisfies the* second-order sufficient optimality conditions *for problem* (NP) *if it is a first-order KKT pair, i.e.,* $r(x^*, y^*) = 0$*, and*

$$d^T H(x^*, y^*) d > 0 \;\; \textit{for all} \;\; d \in \mathcal{C}(x^*, y^*) \setminus \{0\}, \tag{4.3}$$

where $\mathcal{C}(x^*, y^*)$ *is the critical cone* $\mathcal{C}(x^*, y^*) = \operatorname{null}\big(J(x^*)\big) \cap \mathcal{C}_M(x^*, y^*)$*, with*

$$\mathcal{C}_M(x^*, y^*) = \{d : d_i = 0 \;\; \textit{for} \;\; i \in \mathcal{A}_+(x^*, y^*), \; d_i \geq 0 \;\; \textit{for} \;\; i \in \mathcal{A}_0(x^*, y^*)\,\}.$$

The local convergence analysis requires the next assumption, which strengthens Assumption 3.2.

Assumption 4.1. *The functions* $f(x)$ *and* $c(x)$ *are twice Lipschitz-continuously differentiable.*

The following assumption is made concerning the iterates generated by the algorithm.

Assumption 4.2.

(i) *The index set* $\mathcal{S}$ *of V-O iterates defined in* (3.2) *is infinite, and there exists a subsequence* $\mathcal{S}_* \subseteq \mathcal{S}$*, such that* $\lim_{k \in \mathcal{S}_*} (x_k, y_k) = (x^*, y^*)$*, with* (x^*, y^*) *a first-order KKT pair;*

(ii) *there exists a compact set* $\Lambda(x^*) \subseteq \mathcal{Y}(x^*)$ *such that* y^* *belongs to the (nonempty) interior of* $\Lambda(x^*)$ *relative to* $\mathcal{Y}(x^*)$*; and*

(iii) (x^*, y) *satisfies the SOSC of Definition* 4.1 *for every* $y \in \Lambda(x^*)$.

It follows from Theorem 3.1 that if infeasible stationary points are avoided, then the feasible case (A) applies. In this situation, the set $\mathcal{S}$ of V-O iterates is infinite and there exists a limit point x^*. Moreover, if the MFCQ and WCRC hold at x^* (which, together, imply the CPGCQ), then there exists a limit point y^* of $\{y_k\}_{k \in \mathcal{S}}$ such that the pair (x^*, y^*) is a second-order KKT point. Consequently, the key part of Assumption 4.2 is the existence of the compact set $\Lambda(x^*)$, which guarantees that the closest point in $\mathcal{Y}(x^*)$ to every element y_k of the subsequence $\{y_k\}$ satisfying $\lim_{k \to \infty} y_k = y^*$ is also in $\Lambda(x^*)$ for k sufficiently large. This is equivalent to there being a set $\mathcal{K}$, open relative to $\mathcal{Y}(x^*)$, such that $y^* \in \mathcal{K}$ and $\mathcal{K} \subset \Lambda(x^*)$. This, in turn, is equivalent to the assumption that the affine hulls of $\Lambda(x^*)$ and $\mathcal{Y}(x^*)$ are identical with y^* in the relative interior of $\Lambda(x^*)$.

The compactness of the set $\Lambda(x^*)$ in Assumption 4.2 implies the existence of a vector $y_P^*(y)$ that minimizes the distance from y to the set $\Lambda(x^*)$, i.e.,

$$y_P^*(y) \in \operatorname*{Argmin}_{\bar{y} \in \Lambda(x^*)} \|y - \bar{y}\|. \tag{4.4}$$

The existence of a vector $y_P^*(y)$ implies that the distance $\delta(x, y)$ of any primal-dual point (x, y) to the primal-dual solution set $\mathcal{V}(x^*) = \{x^*\} \times \Lambda(x^*)$ associated with x^*, may be written in the form

$$\delta(x, y) = \min_{(\bar{x}, \bar{y}) \in \mathcal{V}(x^*)} \|(x - \bar{x}, y - \bar{y})\| = \|(x - x^*, y - y_P^*(y))\|. \tag{4.5}$$

The pair $\big(x^*, y_P^*(y)\big)$ satisfies the second-order sufficient conditions as a result of Assumption 4.2.

Lemma 4.1. ([35, **Theorem 3.2**]) *If Assumptions 4.1 and 4.2 hold, then there exists a constant* $\kappa \equiv \kappa(\Lambda(x^*))$ *such that*

$$r(x_k, y_k) \in \left[\, \delta(x_k, y_k)/\kappa,\ \delta(x_k, y_k)\kappa \,\right]$$

for all $k \in \mathcal{S}_*$ *sufficiently large.*

Proof. Under the assumptions used here, the result follows from Theorem 3.2 of Wright [35], where Results A.2 and A.1 of the Appendix are used to establish that the exact and estimated distance of (x_k, y_k) to the primal-dual solution set used in [35] are equivalent up to a scalar to the definitions of $\delta(x_k, y_k)$ and $r(x_k, y_k)$ defined here. ∎

The principal steps of the local convergence analysis are summarized as follows. First, the properties of iterates with indices $k \in \mathcal{S}_* \subseteq \mathcal{S}$ are considered. It is shown that for some $k \in \mathcal{S}_*$ sufficiently large, the following results hold.

(a) The active set at x^* is identified correctly by the ϵ-active set, and the direction s_k of negative curvature is zero.

(b) A local descent direction d_k is computed, and the conditions in Step 9 of Algorithm 2 are satisfied, i.e., the local descent direction is selected for the line search.

(c) The unit step is accepted by the flexible line-search Algorithm 4, and the variables active at x^* are the same as those active at x_{k+1}.

The next step is to show that if (a)–(c) hold, then (x_{k+1}, y_{k+1}) is a V-iterate. This implies that the arguments may be repeated at x_{k+1}, and all iterates must be in $\mathcal{S}_*$ for k sufficiently large. The final step is to show that the iterates of Algorithm 5 are identical to those generated by a stabilized SQP method for which superlinear convergence has been established.

The first result shows that for $k \in \mathcal{S}_*$ sufficiently large, the set $\mathcal{A}_\epsilon$ correctly estimates the active set at the solution x^*. Moreover, for these iterations, the search direction does not include a contribution from the direction of negative curvature.

Lemma 4.2. *If Assumption 4.2 applies, then the following results hold for all* $k \in \mathcal{S}_* \subseteq \mathcal{S}$ *sufficiently large.*

(i) *The first-order proximity measure* $r(x, y)$ *converges to zero, i.e.,* $\lim_{k \in \mathcal{S}} r(x_k, y_k) = 0$.

(ii) *The* ϵ*-active sets satisfy* $\mathcal{A}_\epsilon(x_k, y_k, \mu^R_{k-1}) = \mathcal{A}_\epsilon(x_k, y_k, \mu^R_k) = \mathcal{A}(x^*)$.

(iii) *The* ϵ*-free sets satisfy* $\mathcal{F}_\epsilon(x_k, y_k, \mu^R_{k-1}) = \mathcal{F}_\epsilon(x_k, y_k, \mu^R_k) = \mathcal{F}(x^*)$.

(iv) *The matrix* $H_{\mathcal{F}}(x_k, y_k) + (1/\mu^R_{k-1}) J_{\mathcal{F}}(x_k)^T J_{\mathcal{F}}(x_k)$ *is positive definite, where the suffix "*$\mathcal{F}$*" denotes the components corresponding to the index set* $\mathcal{F}(x^*)$. *Moreover, Algorithm 1 gives* $s^{(1)}_k = 0$ *and* $\xi^{(1)}_k = 0$.

(v) *In Algorithm 2, the matrix* $B_{\mathcal{F}_\epsilon}(x_k, y_k\,; \mu^R_k)$ *is positive definite, and a local descent direction is computed.*

(vi) *The feasible direction of negative curvature* s_k *computed in Algorithm 3 is zero.*

Proof. A point (x_k, y_k) is designated as a V-O iterate if the optimality and feasibility measures satisfy condition (2.13). In this case y_k is set to y_k^E, and $\phi_V^{\max}$ or $\phi_O^{\max}$ are decreased by a fixed factor. If follows that on the infinite set $\mathcal{S}$ of V-O iterates, the condition (2.13) must hold infinitely often and at least one of the functions $\phi_V(v_k)$ or $\phi_O(v_k)$ must go to zero. The definitions of $\phi_V(v_k)$ and $\phi_O(v_k)$ in terms of the feasibility and optimality measures $\eta(x_k)$ and $\omega(x_k, y_k)$ implies that $\lim_{k\in\mathcal{S}} \eta(x_k) = 0$ and $\lim_{k\in\mathcal{S}} \omega(x_k, y_k) = 0$. The definition (2.9) of $r(x_k, y_k)$ implies that $\lim_{k\in\mathcal{S}} r(x_k, y_k) = 0$, which proves part (i). Since $r(x_k, y_k)$ goes to zero, part 2 of [13, Theorem 3.2] gives

$$\lim_{k\in\mathcal{S}} \max\left(\mu_{k-1}^R, r(x_k, y_k)^\gamma\right) = \lim_{k\in\mathcal{S}} \max\left(\mu_k^R, r(x_k, y_k)^\gamma\right) = 0.$$

Combining these limits with the definition (2.8) of the ϵ-active set gives $\mathcal{A}_\epsilon(x_k, y_k, \mu_{k-1}^R) \subseteq \mathcal{A}(x^*)$ and $\mathcal{A}_\epsilon(x_k, y_k, \mu_k^R) \subseteq \mathcal{A}(x^*)$ for $k \in \mathcal{S}$ sufficiently large.

The reverse inclusion is established by using the definition of the ϵ-active set (2.8) and the inequalities

$$\max\left(\mu_{k-1}^R, r(x_k, y_k)^\gamma\right) \geq r(x_k, y_k)^\gamma \quad \text{and} \quad \max\left(\mu_k^R, r(x_k, y_k)^\gamma\right) \geq r(x_k, y_k)^\gamma,$$

to imply that the set $\mathcal{A}_\gamma(x_k, y_k) = \left\{i : x_i \leq r(x_k, y_k)^\gamma\right\}$ satisfies $\mathcal{A}_\gamma(x_k, y_k) \subseteq \mathcal{A}_\epsilon(x_k, y_k, \mu_{k-1}^R)$ and $\mathcal{A}_\gamma(x_k, y_k) \subseteq \mathcal{A}_\epsilon(x_k, y_k, \mu_k^R)$ for $k \in \mathcal{S}$ sufficiently large. The set $\mathcal{A}_\gamma(x_k, y_k)$ is an active-set estimator that is equivalent (in the sense of Result A.2) to the active-set estimator used by Wright [35], and Facchinei, Fischer, and Kanzow [8]. This equivalence allows the application of Theorem 3.3 of [35] to obtain the inclusions

$$\mathcal{A}(x^*) \subseteq \mathcal{A}_\gamma(x_k, y_k) \subseteq \mathcal{A}_\epsilon(x_k, y_k, \mu_{k-1}^R) \quad \text{and} \quad \mathcal{A}(x^*) \subseteq \mathcal{A}_\gamma(x_k, y_k) \subseteq \mathcal{A}_\epsilon(x_k, y_k, \mu_k^R),$$

which completes the proof of part (ii).

Part (iii) follows directly from (ii) and the definition of the ϵ-free set in (2.10). For the proof of (iv) it is assumed that $k \in \mathcal{S}_* \subseteq \mathcal{S}$ is sufficiently large that (ii) and (iii) hold. From Assumption 4.2, (x^*, y^*) satisfies the SOSC and consequently, $d^T H(x^*, y^*) d > 0$ for all $d \neq 0$ such that $J(x^*) d = 0$ and $d_i = 0$ for every $i \in \mathcal{A}(x^*)$, i.e., $d_{\mathcal{F}}^T H_{\mathcal{F}}(x^*, y^*) d_{\mathcal{F}} > 0$ for all $d_{\mathcal{F}} \neq 0$ satisfying $J_{\mathcal{F}}(x^*) d_{\mathcal{F}} = 0$, where the suffix "$\mathcal{F}$" denotes quantities associated with indices in $\mathcal{F}(x^*)$. Under this assumption, [19, Lemma 3] and [13, part (2) of Theorem 3.2] imply that $H_{\mathcal{F}_\epsilon}(x_k, y_k) + (1/\mu_{k-1}^R) J_{\mathcal{F}_\epsilon}(x_k)^T J_{\mathcal{F}_\epsilon}(x_k)$ is positive definite for all $k \in \mathcal{S}_*$ sufficiently large. If this matrix is positive definite, the test of Step 3 of Algorithm 1 is satisfied, giving $s_k^{(1)} = 0$ and $\xi_k^{(1)} = 0$, as required.

As $\{\mu_k^R\} \to 0$ (see [13, part (2) of Theorem 3.2]), a similar argument to that used to establish (iv) may be used to show that $H_{\mathcal{F}_\epsilon}(x_k, y_k) + (1/\mu_k^R) J_{\mathcal{F}_\epsilon}(x_k)^T J_{\mathcal{F}_\epsilon}(x_k)$ is positive definite for all $k \in \mathcal{S}_*$ sufficiently large. This is equivalent to the matrix $B_{\mathcal{F}_\epsilon}(x_k, y_k\,; \mu_k^R)$ being positive definite for the same values of k (see Lemma 2.2 of Gill and Robinson [16]). As $B_{\mathcal{F}_\epsilon}(x_k, y_k\,; \mu_k^R)$ is positive definite and $k \in \mathcal{S}_* \subseteq \mathcal{S}$, the conditions of Step 6 of Algorithm 2 are satisfied, and a local descent direction is computed, which proves part (v).

Finally, part (iv) implies that $s_k^{(1)} = 0$, and the definition of Steps 2–5 of Algorithm 3 gives $s_k = 0$, which proves part (vi). ∎

The next result shows that the search direction d_k is nonzero for every $k \in \mathcal{S}_* \subseteq \mathcal{S}$ sufficiently large.

Lemma 4.3. *For all $k \in \mathcal{S}_* \subseteq \mathcal{S}$ sufficiently large, either $d_k \neq 0$ or $(x_k, y_k) = (x^*, y^*)$.*

Proof. The result holds trivially if $d_k \neq 0$ for all $k \in \mathcal{S}_*$ sufficiently large. Assume without loss of generality that there exists an infinite sequence $\mathcal{S}_2 \subseteq \mathcal{S}_*$ such that $d_k = 0$ for all $k \in \mathcal{S}_2$. Parts (ii) and (vi) of Lemma 4.2 imply that $\mathcal{A}_\epsilon(x_k, y_k, \mu_k^R) = \mathcal{A}(x^*)$ and $s_k = 0$ for all $k \in \mathcal{S}_2$ sufficiently large. Every $k \in \mathcal{S}_2$ is a V-O-iterate and there must exist an index $k_2 \in \mathcal{S}_2$ sufficiently large that

$$d_{k_2} = s_{k_2} = 0, \quad (x_{k_2+1}, y_{k_2+1}) = (x_{k_2}, y_{k_2}), \quad y_{k_2}^E = y_{k_2}, \quad \text{and} \quad \mathcal{A}_\epsilon(x_{k_2}, y_{k_2}, \mu_{k_2}^R) = \mathcal{A}(x^*). \quad (4.6)$$

As $d_{k_2} = 0$, parts (ia) and (ib) of Lemma 3.1 give $r(x_{k_2}, y_{k_2}) = 0$ and (x_{k_2}, y_{k_2}) is a first-order KKT point for problem (NP) and for the problem of minimizing $M(x, y\,; y_{k_2}^E, \mu_{k_2}^R)$ subject to $x \geq 0$. From (4.6) it must hold that $r(x_{k_2+1}, y_{k_2+1}) = 0$, and parts (iii) and (iv) of Lemma 4.2 and the definition of $(s_{k_2+1}^{(1)}, \xi_{k_2+1}^{(1)})$ in Algorithm 1 give $s_{k_2+1}^{(1)} = 0$ and $\xi_{k_2+1}^{(1)} = 0$. It follows that $\phi_V(x_{k_2+1}, y_{k_2+1}) = 0$, and $k_2 + 1$ is a V-iterate from condition (2.13). As a result, $y_{k_2+1}^E = y_{k_2}^E$ and $\mu_{k_2+1}^R = \frac{1}{2}\mu_{k_2}^R$, which implies that $(x_{k_2+1}, y_{k_2+1}) = (x_{k_2}, y_{k_2})$ is not only a first-order KKT point for problem (NP), but also a first-order solution of the problem of minimizing $M(x, y\,; y_{k_2+1}^E, \mu_{k_2+1}^R)$ subject to $x \geq 0$. In particular, it must hold that $d_{k_2+1} = 0$, and $s_{k_2+1} = 0$ because $\xi_{k_2+1}^{(1)} = 0$ (see Algorithm 2). Similarly, it must hold that $\mathcal{A}_\epsilon(x_{k_2+1}, y_{k_2+1}, \mu_{k_2+1}^R) = \mathcal{A}(x^*)$.

This argument may be repeated at every (x_k, y_k) such that $k \geq k_2 + 1$, and it must hold that $(x_k, y_k) = (\bar{x}, \bar{y})$ for some $(\bar{x}, \bar{y})$, and that $\mathcal{A}_\epsilon(x_k, y_k, \mu_k^R) = \mathcal{A}(x^*)$ for every $k \geq k_2$. It then follows from Assumption 4.2 that $(\bar{x}, \bar{y}) = (x^*, y^*)$, which completes the proof. ∎

For a local convergence analysis, Lemma 4.3 implies that there is no loss of generality in making the following assumption.

Assumption 4.3. *The direction d_k is nonzero for all $k \in \mathcal{S}_* \subseteq \mathcal{S}$ sufficiently large.*

Lemma 4.4. *If Assumption 4.3 holds, then $\mu_k^R = r(x_k, y_k)^\gamma > 0$ for all $k \in \mathcal{S}_* \subseteq \mathcal{S}$ sufficiently large.*

Proof. Part (iv) of Lemma 4.2 gives $\xi_k^{(1)} = 0$ for all $k \in \mathcal{S}_* \subseteq \mathcal{S}$ sufficiently large. In addition, $r(x_k, y_k)$ must be nonzero, otherwise the definition of $r(x_k, y_k)$ would imply that $c(x_k) = 0$, $y_k^E = y_k$ (because $k \in \mathcal{S}$), $\pi(x_k, y_k^E, \mu_k^R) = y_k$, $\nabla_y M(x_k, y_k\,; y_k^E, \mu_k^R) = 0$, and $\min\big(x_k, \nabla_x M(x_k, y_k\,; y_k^E, \mu_k^R)\big) = 0$. In other words, if $r(x_k, y_k)$ is zero, then (x_k, y_k) satisfies the first-order conditions for a minimizer of $M(x, y\,; y_k^E, \mu_k^R)$ subject to $x \geq 0$. This implies that there is no nonzero descent direction at (x_k, y_k), which contradicts Assumption 4.3. It follows that $r(x_k, y_k)$ is nonzero. The values $\xi_k^{(1)} = 0$ and $r(x_k, y_k) > 0$ in the definition of μ_k^R in (2.14), and part (i) of Lemma 4.2 imply that $\mu_k^R = r(x_k, y_k)^\gamma$ for $\gamma \in (0, 1)$ and $k \in \mathcal{S}_* \subseteq \mathcal{S}$ sufficiently large. ∎

Much of the local convergence analysis involves establishing that, in the limit, Algorithm 5 computes and accepts the local descent direction at every iteration. The next lemma concerns the properties of the equality-constrained subproblem for the local descent direction.

Lemma 4.5. *If $v_k = (x_k, y_k)$ is a point at which the conditions for the calculation of a local descent direction are satisfied, then the following results hold.*

(i) *The bound-constrained problem (2.24) for the local descent direction is equivalent to the stabilized QP subproblem*

$$\begin{array}{ll} \underset{x,y}{\text{minimize}} & g(x_k)^T(x-x_k) + \tfrac{1}{2}(x-x_k)^T H(x_k,y_k)(x-x_k) + \tfrac{1}{2}\mu_k^R\|y\|^2 \\ \text{subject to} & c(x_k) + J(x_k)(x-x_k) + \mu_k^R(y-y_k) = 0, \quad E_{\mathcal{A}_\epsilon}^T x = 0, \end{array} \tag{4.7}$$

where $E_{\mathcal{A}_\epsilon}$ is the matrix of columns of the identity matrix with indices in $\mathcal{A}_\epsilon$.

(ii) *If $d_k = (p_k, q_k)$ denotes the local descent direction, and $z_k = g(x_k) - J(x_k)^T y_k$, then the optimal solution to (4.7) may be written as $(x_k + p_k, y_k + q_k, [\, z_k \,]_{\mathcal{A}_\epsilon} + w_k)$, where (p_k, q_k, w_k) satisfy the nonsingular equations*

$$\begin{pmatrix} H(x_k,y_k) & J(x_k)^T & E_{\mathcal{A}_\epsilon} \\ J(x_k) & -\mu_k^R I & 0 \\ E_{\mathcal{A}_\epsilon}^T & 0 & 0 \end{pmatrix} \begin{pmatrix} p_k \\ -q_k \\ -w_k \end{pmatrix} = - \begin{pmatrix} g(x_k) - J(x_k)^T y_k - z_k^p \\ c(x_k) \\ [\, x_k \,]_{\mathcal{A}_\epsilon} \end{pmatrix}, \tag{4.8}$$

with $z_k^p = E_{\mathcal{A}_\epsilon} E_{\mathcal{A}_\epsilon}^T z_k$, i.e., the projection of z_k onto the null space of $E_{\mathcal{A}_\epsilon}^T$.

Proof. Part (i) follows from the specialization of Result 2.1 of Gill and Robinson [15] to the equality-constraint case. The equations of part (ii) are then the optimality conditions associated with (4.7). It remains to show that the equations are nonsingular. The vector (p_k, q_k) is the unique solution of (4.7) if the primal-dual Hessian of problem (4.7) is positive definite on the null-space of the constraints, which in this case is the set of vectors satisfying $J(x_k)p + \mu_k^R q = 0$ and $E_{\mathcal{A}_\epsilon}^T p = 0$. This corresponds to the requirement that

$$\begin{pmatrix} p_{\mathcal{F}_\epsilon} \\ q \end{pmatrix}^T \begin{pmatrix} H_{\mathcal{F}_\epsilon}(x_k,y_k) & 0 \\ 0 & \mu_k^R I \end{pmatrix} \begin{pmatrix} p_{\mathcal{F}_\epsilon} \\ q \end{pmatrix} = p_{\mathcal{F}_\epsilon}^T H_{\mathcal{F}_\epsilon}(x_k,y_k) p_{\mathcal{F}_\epsilon} + \frac{1}{\mu_k^R} p_{\mathcal{F}_\epsilon}^T J_{\mathcal{F}_\epsilon}(x_k)^T J_{\mathcal{F}_\epsilon}(x_k) p_{\mathcal{F}_\epsilon} > 0.$$

Lemma 2.2 of Gill and Robinson [15] establishes that $H_{\mathcal{F}_\epsilon}(x_k,y_k) + (1/\mu_k^R) J_{\mathcal{F}}(x_k)^T J_{\mathcal{F}_\epsilon}(x_k)$ is positive definite if $B_{\mathcal{F}_\epsilon}$ is positive definite, which is one of the conditions that must be satisfied for a local descent direction to be computed. ∎

The next lemma establishes that two of the three conditions for the acceptance of the local descent direction are satisfied for all $k \in \mathcal{S}_*$ sufficiently large (see Step 9 of Algorithm 2).

Lemma 4.6. *Let Assumptions 3.1, 4.1–4.3 hold. For all $k \in \mathcal{S}_* \subseteq \mathcal{S}$ sufficiently large, a local descent direction $d_k = (p_k, q_k)$ is computed that satisfies the following inequalities:*

(i) $\max\{\|p_k\|, \|q_k\|\} = O\big(\delta(x_k, y_k)\big)$; *and*

(ii) $x_k + p_k \ge 0$, $[\, \nabla \mathcal{Q}_k(v_k + d_k\, ; \mu_k^R) \,]_{\mathcal{A}_\epsilon} \ge -t_k e$, *where t_k is the positive feasibility parameter (2.30), and "$\mathcal{A}_\epsilon$" denotes the vector of components with indices in the ϵ-active set $\mathcal{A}_\epsilon(x_k, y_k, \mu_k^R)$.*

Proof. Lemma 4.5 implies that the local descent direction (p_k, q_k) satisfies the equations

$$\begin{pmatrix} H(x_k,y_k) & J(x_k)^T & E_{\mathcal{A}_\epsilon} \\ J(x_k) & -\mu_k^R I & 0 \\ E_{\mathcal{A}_\epsilon}^T & 0 & 0 \end{pmatrix} \begin{pmatrix} p_k \\ -q_k \\ -w_k \end{pmatrix} = - \begin{pmatrix} g(x_k) - J(x_k)^T y_k - z_k^p \\ c(x_k) \\ [\, x_k \,]_{\mathcal{A}_\epsilon} \end{pmatrix}, \tag{4.9}$$

where $[\,z_k\,]_{\mathcal{A}_\epsilon}+w_k$ is the vector of multipliers for the constraints $E_{\mathcal{A}_\epsilon}^T x = 0$ of problem (4.7). Let $\widetilde{\mu}_k$ denote the scalar $\widetilde{\mu}(x_k,y_k,z_k) = \|(g(x_k) - J(x_k)^T y_k - z_k^p, c(x_k), [\,x_k\,]_{\mathcal{A}_\epsilon})\|_1$. The equations (4.9) constitute a perturbation of the linear system

$$\begin{pmatrix} H(x_k,y_k) & J(x_k)^T & E_{\mathcal{A}_\epsilon} \\ J(x_k) & -\widetilde{\mu}_k I & 0 \\ E_{\mathcal{A}_\epsilon}^T & 0 & -\widetilde{\mu}_k I \end{pmatrix} \begin{pmatrix} \widetilde{p}_k \\ -\widetilde{q}_k \\ -\widetilde{w}_k \end{pmatrix} = - \begin{pmatrix} g(x_k) - J(x_k)^T y_k - z_k^p \\ c(x_k) \\ [\,x_k\,]_{\mathcal{A}_\epsilon} \end{pmatrix}, \tag{4.10}$$

which characterize the optimality conditions for the stabilized SQP subproblem associated with the equality constrained problem

$$\underset{x}{\text{minimize}}\; f(x) \;\;\text{subject to}\;\; c(x) = 0, \;\;\text{and}\;\; [\,x\,]_{\mathcal{A}_\epsilon} = E_{\mathcal{A}_\epsilon}^T x = 0. \tag{4.11}$$

The matrix of (4.10) is nonsingular and the equations have a unique solution (see Izmailov and Solodov [24, Lemma 2]). In addition, it follows from Wright [35, Lemma 4.1], Result A.3 and Lemma 4.1 that the unique solution of (4.10) satisfies

$$\|(\widetilde{p}_k,\widetilde{q}_k)\| \le \|(\widetilde{p}_k,\widetilde{q}_k,\widetilde{w}_k)\| = O(\widetilde{\mu}_k) = O\big(\delta(x_k,y_k)\big) = O\big(r(x_k,y_k)\big). \tag{4.12}$$

The underlying quadratic program associated with (4.9) satisfies the second-order sufficient conditions for optimality. Under this condition, Izmailov [21, Theorem 2.3]) establishes the Lipschitz error bound for the perturbed solutions as

$$\|(p_k - \widetilde{p}_k, q_k - \widetilde{q}_k)\| \le \|(p_k - \widetilde{p}_k, q_k - \widetilde{q}_k, w_k - \widetilde{w}_k)\| = O(\|\widetilde{\mu}_k \widetilde{w}_k + \big(\mu_k^R - \widetilde{\mu}_k\big)(q_k - \widetilde{q}_k)\|).$$

Lemma 4.4 gives $\mu_k^R = r(x_k,y_k)^\gamma$ for $\gamma \in (0,1)$. It then follows from Result A.3 the bound (4.12) and Lemma 4.1 that

$$\|(p_k - \widetilde{p}_k, q_k - \widetilde{q}_k)\| = O\big(\delta(x_k,y_k) + r(x_k,y_k)^\gamma \|q_k - \widetilde{q}_k\|\big). \tag{4.13}$$

The triangle inequality and the bounds (4.13) and (4.12) imply the existence of constants κ_1 and κ_2 that satisfy

$$\begin{aligned} \|p_k\| + \|q_k\| &\le \|p_k - \widetilde{p}_k\| + \|q_k - \widetilde{q}_k\| + \|\widetilde{p}_k\| + \|\widetilde{q}_k\| & (4.14)\\ &\le \kappa_1 \delta(x_k,y_k) + \kappa_2 r(x_k,y_k)^\gamma \|q_k - \widetilde{q}_k\|. & (4.15) \end{aligned}$$

Part (i) of Lemma 4.2 implies that $1 - \kappa_2 r(x_k,y_k)^\gamma \ge \frac{1}{2}$ for $k \in \mathcal{S}_*$ sufficiently large. This inequality may be used to derive the bound

$$\begin{aligned} &\|p_k - \widetilde{p}_k\| + \tfrac{1}{2}\|q_k - \widetilde{q}_k\| + \|\widetilde{p}_k\| + \|\widetilde{q}_k\| \\ &\qquad\qquad \le \|p_k - \widetilde{p}_k\| + \big(1 - \kappa_2 r(x_k,y_k)^\gamma\big)\|q_k - \widetilde{q}_k\| + \|\widetilde{p}_k\| + \|\widetilde{q}_k\|. \end{aligned}$$

This upper bound may be simplified using the bound on $\|p_k - \widetilde{p}_k\| + \|q_k - \widetilde{q}_k\| + \|\widetilde{p}_k\| + \|\widetilde{q}_k\|$ from (4.14)–(4.15), giving

$$\|p_k - \widetilde{p}_k\| + \tfrac{1}{2}\|q_k - \widetilde{q}_k\| + \|\widetilde{p}_k\| + \|\widetilde{q}_k\| \le \kappa_1 \delta(x_k,y_k).$$

The quantity $\frac{1}{2}(\|p_k\| + \|q_k\|)$ may be bounded using similar arguments used to give (4.14). In this case,

$$\tfrac{1}{2}(\|p_k\| + \|q_k\|) \le \|p_k - \widetilde{p}_k\| + \tfrac{1}{2}\|q_k - \widetilde{q}_k\| + \|\widetilde{p}_k\| + \|\widetilde{q}_k\| \le \kappa_1 \delta(x_k,y_k),$$

which implies that $\max\{\|p_k\|, \|q_k\|\} = O\big(\delta(x_k, y_k)\big)$, which proves part (i).

The second inequality to be established for part (ii) may be written in the equivalent form $[\,\nabla M_k + B_k d_k\,]_{\mathcal{A}_\epsilon} \ge -t_k e$, where $\nabla M_k = \nabla M(v_k\,; y^E_k, \mu^R_k)$ and $B_k = B(v_k, \mu^R_k)$. The proof requires estimates of the components of the vector $[\,\nabla M_k + B_k d_k\,]_{\mathcal{A}_\epsilon}$. After some simplification, the substitution of the quantities B_k, ∇M_k and $d_k = (p_k, q_k)$, together with the identity $J(x_k)p_k + \mu^R_k q_k = -c(x_k)$ from (4.9) give

$$[\,\nabla M_k + B_k d_k\,]_{\mathcal{A}_\epsilon} = \Big[z_k + \frac{1}{\mu^R_k} J(x_k)^T c(x_k) + H(x_k, y_k)p_k + \frac{1}{\mu^R_k} J(x_k)^T J(x_k)p_k\Big]_{\mathcal{A}_\epsilon}, \qquad (4.16)$$

where $z_k = g(x_k) - J(x_k)^T y_k$. The first part of the proof involves the estimation of a lower bound on the vector $z_k + (1/\mu^R_k)J(x_k)^T c(x_k)$. The definition of $y^*_P(\cdot)$ and the fact that (x^*, y^*) is a first-order KKT pair for problem (NP) implies that the vector $g(x^*) - J(x^*)^T y^*_P(y_k)$ is nonnegative, with

$$\begin{aligned} -[\,z_k\,]_i &= -[\,g(x_k) - J(x_k)^T y_k\,]_i \le -\big[g(x_k) - J(x_k)^T y_k - \big(g(x^*) - J(x^*)^T y^*_P(y_k)\big)\big]_i \\ &\le -[\,g(x_k) - J(x_k)^T y_k + J(x_k)^T y^*_P(y_k) - J(x_k)^T y^*_P(y_k) - g(x^*) + J(x^*)^T y^*_P(y_k)\,]_i. \end{aligned}$$

From Assumptions 4.1–4.2, $\|J(x_k)\|$ is bounded independently of k and the functions g and J are Lipschitz continuous. It follows that there exist positive constants κ_3, κ_4, and κ_5 such that

$$-[\,z_k\,]_i \le \kappa_3\|x_k - x^*\| + \kappa_4\|y_k - y^*_P(y_k)\| \le \kappa_5\delta(x_k, y_k), \qquad (4.17)$$

where the last inequality follows from the definition (4.5) of $\delta(x_k, y_k)$. As the sequence of iterates satisfies $\lim_{k\in\mathcal{S}_*}(x_k.y_k) = (x^*, y^*)$ and $\lim_{k\in\mathcal{S}_*} y^*_P(y_k) = y^*$, for $k \in \mathcal{S}_*$ sufficiently large, the assumptions needed for Lemma 4.1 apply, and

$$-[\,z_k\,]_i \le \kappa_5\delta(x_k, y_k) \le \kappa_6 r(x_k, y_k) \qquad (4.18)$$

for some positive constant κ_6. The inequality (4.18), the definition of $r(x_k, y_k)$, and the result $\mu^R_k = r(x_k, y_k)^\gamma$ of Lemma 4.4 imply that there exists a positive constant κ_7 such that

$$\begin{aligned} \Big[z_k + \frac{1}{\mu^R_k} J(x_k)^T c(x_k)\Big]_i &\ge -\kappa_6 r(x_k, y_k) - \frac{\|J(x_k)\|_1 r(x_k, y_k)}{r(x_k, y_k)^\gamma} \\ &= -\kappa_6 r(x_k, y_k) - \|J(x_k)\|_1 r(x_k, y_k)^{1-\gamma} \\ &\ge -\kappa_7 r(x_k, y_k)^{1-\gamma} \ge -\tfrac{1}{2} r(x_k, y_k)^\lambda, \end{aligned} \qquad (4.19)$$

for all i, and every $k \in \mathcal{S}_*$ sufficiently large, where the last inequality follows from the assumption $0 < \lambda < \min\{\gamma, 1-\gamma\} < 1$.

The second term of (4.16) may be bounded in a similar way using the definition $\mu^R_k = r(x_k, y_k)^\gamma$ and the bound on $\|p_k\|$ from part (i). The assumption that $H(x_k, y_k)$ and $J(x_k)$ are bounded, the estimate $\delta(x_k, y_k) = O(r(x_k, y_k))$ of Lemma 4.1, and the definition of $\mathcal{A}_\epsilon(x_k, y_k, \mu^R_k)$ give

$$\big[H(x_k, y_k)p_k + (1/\mu^R_k)J(x_k)^T J(x_k)p_k\big]_i = O\big(r(x_k, y_k)^{1-\gamma}\big) \le \tfrac{1}{2} r(x_k, y_k)^\lambda,$$

for all $k \in \mathcal{S}_*$ sufficiently large. A combination of (4.16), (4.19) and (4.20) yields

$$\begin{aligned} [\,\nabla M_k + B_k d_k\,]_{\mathcal{A}_\epsilon} &\ge \Big[z_k + \frac{1}{\mu^R_k} J(x_k)^T c(x_k)\Big]_{\mathcal{A}_\epsilon} - \Big\|\Big[H(x_k, y_k)p_k + \frac{1}{\mu^R_k} J(x_k)^T J(x_k)p_k\Big]_{\mathcal{A}_\epsilon}\Big\|_\infty e \\ &\ge -r(x_k, y_k)^\lambda e = -t_k e, \end{aligned}$$

for all $k \in \mathcal{S}_*$ sufficiently large, which proves the second result of part (ii).

The first result of Lemma 4.2(iii) implies that $\mathcal{F}(x_k, y_k, \mu_k^R) = \mathcal{F}(x^*)$ for $k \in \mathcal{S}_*$ sufficiently large. If the limit $\lim_{k \in \mathcal{S}_*} [\, x_k \,]_{\mathcal{F}_\epsilon} = [\, x^* \,]_{\mathcal{F}} > 0$ is used in conjunction with the definition $[\, x_k + p_k \,]_{\mathcal{A}_\epsilon} = 0$, and the estimate $\|[\, p_k \,]_{\mathcal{F}_\epsilon}\| = \|[\, p_k \,]_{\mathcal{F}}\| = O\big(\delta(x_k, y_k)\big)$ of part (i), it follows that $x_k + p_k \geq 0$ for $k \in \mathcal{S}_*$ sufficiently large, as required. ∎

Part (ii) of Lemma 4.6 implies that two of the three conditions needed for the acceptance of the local descent direction are satisfied. It remains to show that the third condition $\nabla M_k^T d_k < 0$ holds. The next technical result establishes some properties of the local descent direction.

Lemma 4.7. *Let Assumptions 3.1, 4.1–4.3 hold. For all $k \in \mathcal{S}_* \subseteq \mathcal{S}$ sufficiently large, a local descent direction $d_k = (p_k, q_k)$ is computed such that $(\widehat{x}_k, \widehat{y}_k) = (x_k + p_k, y_k + q_k)$ satisfies*

$$\delta(\widehat{x}_k, \widehat{y}_k) = \|\widehat{x}_k - x^*\| + \|\widehat{y}_k - y_P^*(\widehat{y}_k)\| = O\big(\delta(x_k, y_k)^{1+\gamma}\big), \tag{4.20}$$

with $y_P^*(\cdot)$ defined in (4.4).

Proof. The proof utilizes a result of Izmailov [21, Theorem 2.3] that provides a bound on the change in the solution of a problem perturbed by a quantity ε. If the second-order sufficient conditions hold at a primal-dual solution (x^*, y^*) of a problem P, then the primal-dual solution $(\widetilde{x}, \widetilde{y})$ of a perturbed problem $P(\varepsilon)$ satisfies

$$\|\widetilde{x} - x_*\| + \inf_{y \in \mathcal{Y}(x^*)} \|\widetilde{y} - y\| = O(\|\varepsilon\|). \tag{4.21}$$

For the purposes of this theorem, the unperturbed problem is an equality-constrained variant of problem (NP) for which the optimal active set has been identified. Parts (ii) and (iii) of Lemma 4.2 imply that $\mathcal{A}(x^*) = \mathcal{A}_\epsilon(x_k, y_k, \mu_k^R)$, and $\mathcal{F}(x^*) = \mathcal{F}_\epsilon(x_k, y_k, \mu_k^R)$ for $k \in \mathcal{S}_*$ sufficiently large. Let $E_\mathcal{A}$ denote the matrix of columns of the identity matrix with indices in $\mathcal{A}(x^*)$. At any iteration with $k \in \mathcal{S}_*$, consider the perturbed problem

$$\underset{x}{\text{minimize}} \;\; f(x) + x^T \varepsilon_k^{(1)} \quad \text{subject to} \;\; c(x) + \varepsilon_k^{(2)} = 0, \quad E_\mathcal{A}^T x = 0, \tag{4.22}$$

where $\varepsilon_k^{(1)}$ and $\varepsilon_k^{(2)}$ are perturbation vectors such that $\varepsilon_k = \big(\varepsilon_k^{(1)}, \varepsilon_k^{(2)}\big)$ with

$$\varepsilon_k = \begin{pmatrix} \varepsilon_k^{(1)} \\ \varepsilon_k^{(2)} \end{pmatrix} = \begin{pmatrix} g(x_k) - J(x_k)^T \widehat{y}_k - (g(\widehat{x}_k) - J(\widehat{x}_k)^T \widehat{y}_k) + H(x_k, y_k)(\widehat{x}_k - x_k) \\ c(x_k) + J(x_k)(\widehat{x}_k - x_k) - c(\widehat{x}_k) + \mu_k^R(\widehat{y}_k - y_k^E) \end{pmatrix}. \tag{4.23}$$

The following simple argument may be used to show that the perturbations go to zero as $k \to \infty$ for $k \in \mathcal{S}_*$. Part (i) of Lemma 4.6 implies that $\lim_{k \in \mathcal{S}_*} (\widehat{x}_k - x_k, \widehat{y}_k - y_k) = \lim_{k \in \mathcal{S}_*} (p_k, q_k) = 0$ for $k \in \mathcal{S}_*$ sufficiently large. Moreover, as $\lim_{k \in \mathcal{S}_*} (x_k, y_k) = (x^*, y^*)$ and $y_k^E = y_k$ for $k \in \mathcal{S}_*$, it must hold that $\lim_{k \in \mathcal{S}_*} \varepsilon_k = 0$.

The proof of (4.20) is based on applying the bound (4.21) for the values $(\widetilde{x}, \widetilde{y}) = (\widehat{x}_k, \widehat{y}_k)$. In this case, under Assumption 4.2, it holds that

$$\delta(\widehat{x}_k, \widehat{y}_k) = \|\widehat{x}_k - x^*\| + \|\widehat{y}_k - y_P^*(\widehat{y}_k)\| = \|\widehat{x}_k - x^*\| + \inf_{y \in \Lambda(x^*)} \|\widehat{y}_k - y\| = O(\|\varepsilon_k\|).$$

Three results must be established in order to apply this result. First, (x^*, y^*) must satisfy the second-order sufficient conditions for the equality-constrained problem (4.22) with $\varepsilon_k = 0$.

Second, $(\widehat{x}_k, \widehat{y}_k)$ must be an optimal solution for the perturbed problem (4.22) with perturbation (4.23). Third, the perturbation (4.23) must be bounded in terms of $\delta(x_k, y_k)$.

For the first part it must be shown that (x^*, y^*) satisfies the second-order sufficient conditions for problem (4.22) with no perturbation. The first-order KKT conditions for (4.22) are

$$g(x) - J(x)^T y + \varepsilon_k^{(1)} - E_{\mathcal{A}} z_{\mathcal{A}} = 0, \quad c(x) + \varepsilon_k^{(2)} = 0, \quad \text{and} \quad E_{\mathcal{A}}^T x = 0. \tag{4.24}$$

If $\varepsilon_k = 0$ then (x^*, y^*) satisfies these conditions, which implies that the primal-dual pair (x^*, y^*) is a first-order KKT point. The second-order conditions for problem (NP) imply that $p^T H(x^*, y^*) p > 0$ for all p such that $J(x^*)p = 0$ and $p_i = 0$ for every $i \in \mathcal{A}(x^*)$. These conditions also apply for problem (4.22) when $\varepsilon_k = 0$, which imply that (x^*, y^*) satisfies the second-order sufficient conditions for the unperturbed problem.

Next, it must be shown that $(\widehat{x}_k, \widehat{y}_k)$ is an optimal solution for the problem (4.22) with perturbation (4.23). By definition, the point $(\widehat{x}_k, \widehat{y}_k)$ satisfies the optimality conditions for the equality-constrained problem (2.24). If $y_k^E = y_k$, then these conditions are

$$\begin{aligned} g(x_k) + H(x_k, y_k)(\widehat{x}_k - x_k) - J(x_k)^T y_k - E_{\mathcal{A}} z_{\mathcal{A}} &= 0, \\ c(x_k) + J(x_k)(\widehat{x}_k - x_k) + \mu_k^R(\widehat{y}_k - y_k) &= 0, \quad \text{and} \quad E_{\mathcal{A}}^T \widehat{x}_k = 0, \end{aligned} \tag{4.25}$$

where $z_{\mathcal{A}} = [\, z_k \,]_{\mathcal{A}}$ with $z_k = g(x_k) - J(x_k)^T y_k$ (cf. (4.9)). These identities may be used to show that $(\widehat{x}_k, \widehat{y}_k)$ satisfies the optimality conditions (4.24) with ε_k defined as in (4.23).

It remains to bound the perturbation norm $\|\varepsilon_k\|$ from (4.23). The Taylor-series expansions of $g(\widehat{x}_k) = g(x_k + p_k)$ and $J(\widehat{x}_k) = J(x_k + p_k)$, together with the assumption that $\{\nabla^2 c_i(x_k)\}_{k \in \mathcal{S}_*}$ is bounded, give

$$\begin{aligned} &g(x_k) - g(x_k + p_k) + H(x_k, y_k) p_k - (J(x_k) - J(x_k + p_k))^T \widehat{y}_k \\ &\qquad = \sum_{i=1}^{m} [\, \widehat{y}_k - y_k \,]_i \nabla^2 c_i(x_k) p_k + O(\|p_k\|^2) = O\big(\|p_k\| \|\widehat{y}_k - y_k\|\big) + O\big(\|p_k\|^2\big), \end{aligned} \tag{4.26}$$

which bounds the norm of the first block of (4.23).

Three properties of the iterates are needed to bound the norm of the second block. First, a Taylor-series expansion of $c(x_k + p_k)$ gives $c(x_k) - c(x_k + p_k) + J(x_k) p_k = O(\|p_k\|^2)$. Second, as $\mathcal{S}_*$ contains only V-O iteration indices, the updating rule for y_k^E in Algorithm 5 gives $y_k^E = y_k$ for all $k \in \mathcal{S}_*$. Third, Lemma 4.4 gives $\mu_k^R = r(x_k, y_k)^\gamma$, which implies that $\mu_k^R \|\widehat{y}_k - y_k\| = r(x_k, y_k)^\gamma \|\widehat{y}_k - y_k\|$. The combination of these results gives

$$\|\varepsilon_k\| = O(\|p_k\|^2) + O(\|p_k\| \|\widehat{y}_k - y_k\|) + O(r(x_k, y_k)^\gamma \|\widehat{y}_k - y_k\|).$$

Writing $q_k = \widehat{y}_k - y_k$ and using the results $r(x_k, y_k) = O(\delta(x_k, y_k))$ (from Lemma 4.1), and $\max\{\|p_k\|, \|q_k\|\} = O\big(\delta(x_k, y_k)\big)$ (from Lemma 4.6(i)), and the definition $0 < \gamma < 1$, gives

$$\|\varepsilon_k\| = O\big(\delta(x_k, y_k)^2 + \delta(x_k, y_k)^{1+\gamma}\big) = O\big(\delta(x_k, y_k)^{1+\gamma}\big),$$

which gives the required bound (4.20). ∎

The next step is to show that the final condition required for the acceptance of the local descent direction computed in Step 8 of Algorithm 2 is satisfied, i.e., that the local descent direction is a descent direction for the merit function. The proof of this result requires the following lemma.

Lemma 4.8. *For every $k \in \mathcal{S}_* \subseteq \mathcal{S}$ it holds that*

(i) $\|y_k - \pi_k\| = O\big(\|c(x_k)\|/\mu_k^R\big)$; *and*

(ii) $\|\nabla^2 M(v_k\,; y_k^E, \mu_k^R) - B_k\| = O\big(\|c(x_k)\|/\mu_k^R\big)$,

where $\pi_k = \pi(x_k\,; y_k^E, \mu_k^R)$. Moreover, if Assumption 3.3 holds, then $\lim_{k\in\mathcal{S}_*} \|y_k - \pi_k\| = 0$ and $\lim_{k\in\mathcal{S}_*} \|\nabla^2 M(v_k\,; y_k^E, \mu_k^R) - B_k\| = 0$.

Proof. As $y_k = y_k^E$ for all $k \in \mathcal{S}_* \subseteq \mathcal{S}$, the definition of π_k gives $\|y_k - \pi_k\| = \|c(x_k)\|/\mu_k^R$. This estimate in conjunction with the definitions of $\nabla^2 M$ and B imply that part (ii) also holds.

Lemma 4.4 and part (i) of Lemma 4.2 give $\lim_{k\in\mathcal{S}_*} r(x_k, y_k) = 0$, with $\mu_k^R = r(x_k, y_k)^\gamma$ and $1 - \gamma > 0$ for all $k \in \mathcal{S}_* \subseteq \mathcal{S}$ sufficiently large. These results may be combined to give

$$0 \le \lim_{k\in\mathcal{S}_*} \frac{\|c(x_k)\|}{\mu_k^R} \le \lim_{k\in\mathcal{S}_*} \frac{r(x_k, y_k)}{\mu_k^R} = \lim_{k\in\mathcal{S}_*} \frac{r(x_k, y_k)}{r(x_k, y_k)^\gamma} = \lim_{k\in\mathcal{S}_*} r(x_k, y_k)^{1-\gamma} = 0.$$

It follows from (i) that $\lim_{k\in\mathcal{S}_*} \|y_k - \pi_k\| = 0$. Moreover, the assumption that $\{\nabla^2 c_i(x_k)\}_{k\in\mathcal{S}_*}$ is bounded gives $\lim_{k\in\mathcal{S}_*} \|\nabla^2 M(v_k\,; y_k^E, \mu_k^R) - B_k\| = 0$, as required. ∎

It remains to show that the final acceptance condition is satisfied.

Lemma 4.9. *Let Assumptions 3.1, 4.1–4.3 hold. For any $\bar\sigma$ satisfying $0 < \bar\sigma < 1$, and all $k \in \mathcal{S}_* \subseteq \mathcal{S}$ sufficiently large, a local descent direction $d_k = (p_k, q_k)$ is computed that satisfies*

$$\nabla M(v_k\,; y_k^E, \mu_k^R)^T d_k \le -\bar\sigma d_k^T B_k d_k - \bar c \|d_k\|^2 \quad \text{and} \quad \nabla M(v_k\,; y_k^E, \mu_k^R)^T d_k < 0, \tag{4.27}$$

for some positive constant $\bar c$. In particular, d_k is a strict descent direction for $M(v\,; y_k^E, \mu_k^R)$ at v_k.

Proof. Throughout the proof, the quantities $\nabla M(x_k, y_k\,; y_k^E, \mu_k^R)$ and $B(x_k, y_k\,; \mu_k^R)$ are denoted by ∇M_k and B_k, respectively. In addition, it is assumed that $k \in \mathcal{S}_* \subseteq \mathcal{S}$ is sufficiently large that parts (ii) and (iii) of Lemma 4.2 hold; i.e., $\mathcal{A}_\epsilon(x_k, y_k, \mu_k^R) = \mathcal{A}(x^*)$, and $\mathcal{F}_\epsilon(x_k, y_k, \mu_k^R) = \mathcal{F}(x^*)$. With this assumption, $[\,B_k\,]_\mathcal{A}$, $[\,B_k\,]_\mathcal{F}$ and $[\,B_k\,]_{\mathcal{A},\mathcal{F}}$ denote the rows and columns of the matrix B_k associated with the index sets $\mathcal{A}(x^*)$ and $\mathcal{F}(x^*)$.

The definition of d_k from (2.26) gives the identity $[\,\nabla M_k + B_k d_k\,]_\mathcal{F} = 0$, which, in partitioned form, is

$$[\,B_k\,]_\mathcal{F} [\,d_k\,]_\mathcal{F} + [\,B_k\,]_{\mathcal{A},\mathcal{F}}^T [\,d_k\,]_\mathcal{A} = -[\,\nabla M_k\,]_\mathcal{F}. \tag{4.28}$$

Similarly, the scalar $d_k^T B_k d_k$ may be written in the form

$$d_k^T B_k d_k = [\,d_k\,]_\mathcal{F}^T [\,B_k\,]_\mathcal{F} [\,d_k\,]_\mathcal{F} + (2[\,B_k\,]_{\mathcal{A},\mathcal{F}} [\,d_k\,]_\mathcal{F} + [\,B_k\,]_\mathcal{A} [\,d_k\,]_\mathcal{A})^T [\,d_k\,]_\mathcal{A}. \tag{4.29}$$

Combining (4.28) and (4.29) yields

$$\begin{aligned} -[\,\nabla M_k\,]_\mathcal{F}^T [\,d_k\,]_\mathcal{F} &= d_k^T B_k d_k - ([\,B_k\,]_{\mathcal{A},\mathcal{F}} [\,d_k\,]_\mathcal{F} + [\,B_k\,]_\mathcal{A} [\,d_k\,]_\mathcal{A})^T [\,d_k\,]_\mathcal{A} \\ &= d_k^T B_k d_k - [\,B_k d_k\,]_\mathcal{A}^T [\,d_k\,]_\mathcal{A}, \end{aligned} \tag{4.30}$$

which implies that, for any $\bar\sigma$ satisfying $0 < \bar\sigma < 1$, it must hold that

$$\nabla M_k^T d_k + \bar\sigma d_k^T B_k d_k = (\bar\sigma - 1) d_k^T B_k d_k + [\,B_k d_k\,]_\mathcal{A}^T [\,d_k\,]_\mathcal{A} + [\,\nabla M_k\,]_\mathcal{A}^T [\,d_k\,]_\mathcal{A}. \tag{4.31}$$

The proof involves constructing a bound on each of the terms of the right-hand side of this identity. These bounds are characterized in terms of the index sets $\mathcal{A}_+(x^*,y^*)$ and $\mathcal{A}_0(x^*,y^*)$ defined in (4.1) and (4.2), together with the set $\mathcal{F}_0(x^*,y^*) = \mathcal{A}_0(x^*,y^*) \cup \mathcal{F}(x^*,y^*)$. In what follows, $[\,B_k\,]_{\mathcal{A}_+}$ and $[\,B_k\,]_{\mathcal{F}_0}$ denote the matrices of rows and columns of B_k associated with the index sets $\mathcal{A}_+$ and $\mathcal{F}_0$, with similar definitions for $[\,B_k\,]_{\mathcal{A}_0}$ and $[\,B_k\,]_{\mathcal{A}_+,\mathcal{F}_0}$, etc. The index sets $\mathcal{F}_0$ and $\mathcal{A}_+$ define a partition of $\{1,\,2,\,\ldots,\,n+m\}$, and $d_k^TB_kd_k$ may be partitioned analogous to (4.29) as

$$d_k^TB_kd_k = [\,d_k\,]_{\mathcal{F}_0}^T[\,B_k\,]_{\mathcal{F}_0}[\,d_k\,]_{\mathcal{F}_0} + ([\,B_k\,]_{\mathcal{A}_+}[\,d_k\,]_{\mathcal{A}_+} + 2[\,B_k\,]_{\mathcal{A}_+,\mathcal{F}_0}[\,d_k\,]_{\mathcal{F}_0})^T[\,d_k\,]_{\mathcal{A}_+}. \tag{4.32}$$

The second-order sufficient conditions given in Definition 4.1, [13, Theorem 1.3 and part 2 of Theorem 3.2], together with a continuity argument imply that, for all $k\in\mathcal{S}_*$ sufficiently large, B_k is uniformly positive definite when restricted to the set $\mathcal{C} = \{(p,q)\in\mathbb{R}^{n+m} : p_{\mathcal{A}_+} = 0 \text{ and } p_{\mathcal{A}_0} \ge 0\}$. The relation $(-d)^TB_k(-d) = d^TB_kd$ implies that if d satisfies $d_{\mathcal{A}_0}\le 0$ and $d_{\mathcal{A}_+} = 0$, then $d^TB_kd > 0$. For the particular vector $d = (0,[\,d_k\,]_{\mathcal{A}_0},[\,d_k\,]_{\mathcal{F}}) = (0,[\,d_k\,]_{\mathcal{F}_0})$, for which $[\,d_k\,]_{\mathcal{A}_0}\le 0$, it must be the case that

$$[\,d_k\,]_{\mathcal{F}_0}^T[\,B_k\,]_{\mathcal{F}_0}[\,d_k\,]_{\mathcal{F}_0} \ge \kappa_1\|[\,d_k\,]_{\mathcal{F}_0}\|^2, \;\text{ for some } \kappa_1\in(0,1), \tag{4.33}$$

and all $k\in\mathcal{S}_*$ sufficiently large. This inequality provides a bound on the first term on the right-hand side of (4.32). An estimate of the second and third terms may be determined using a bound on the magnitude of the components of $[\,B_kd_k\,]_{\mathcal{A}}$, where, by definition,

$$[\,B_kd_k\,]_{\mathcal{A}_+} = \Big[\Big(H(x_k,y_k) + \frac{1}{\mu_k^R}(1+\nu)J(x_k)^TJ(x_k)\Big)p_k + \nu J(x_k)^Tq_k\Big]_{\mathcal{A}_+}.$$

For $k\in\mathcal{S}_*$ sufficiently large, Lemma 4.4 implies that $\mu_k^R = r(x_k,y_k)^\gamma$. In addition, as $\|H(x_k,y_k)\|$ and $\|J(x_k)\|$ are bounded on $\mathcal{S}$, it follows from the bounds on $\|p_k\|$ and $\|q_k\|$ given by Lemma 4.6(i), and the equivalence $r(x_k,y_k) = \Theta\big(\delta(x_k,y_k)\big)$ of Lemma 4.1, that the magnitude of the components of $[\,B_kd_k\,]_{\mathcal{A}_+}$ are estimated by

$$\|[\,B_kd_k\,]_{\mathcal{A}_+}\| = O(r(x_k,y_k)^{1-\gamma}) = O(\delta(x_k,y_k)^{1-\gamma}). \tag{4.34}$$

A similar argument gives the bound

$$\big|\big([\,B_k\,]_{\mathcal{A}_+}[\,d_k\,]_{\mathcal{A}_+} + 2[\,B_k\,]_{\mathcal{A}_+,\mathcal{F}_0}[\,d_k\,]_{\mathcal{F}_0}\big)^T[\,d_k\,]_{\mathcal{A}_+}\big| = O(\delta(x_k,y_k)^{1-\gamma}\,\|[\,d_k\,]_{\mathcal{A}_+}\|). \tag{4.35}$$

The application of the bound (4.33) and estimate (4.35) to (4.32) gives

$$-d_k^TB_kd_k \le -\kappa_1\|[\,d_k\,]_{\mathcal{F}_0}\|^2 + \kappa_2\delta(x_k,y_k)^{1-\gamma}\|[\,d_k\,]_{\mathcal{A}_+}\|, \tag{4.36}$$

for some positive κ_2 independent of k, which serves to bound $(\bar\sigma-1)d_k^TB_kd_k$, the first term of the right-hand side of (4.31).

The second and third terms of (4.31) are estimated by bounding components from the index set $\mathcal{A}_+$. The estimate (4.34) gives

$$[\,B_kd_k\,]_{\mathcal{A}_+}^T[\,d_k\,]_{\mathcal{A}_+} \le \kappa_3\delta(x_k,y_k)^{1-\gamma}\|[\,d_k\,]_{\mathcal{A}_+}\|, \;\text{ for some } \kappa_3\in(0,1). \tag{4.37}$$

A Taylor-series expansion of $\nabla M(v_k\,;y^E,\mu_k^R)$ at $y^E = y_k^E\;(=y_k)$ gives

$$\nabla M_k = \nabla M(v_k\,;y^* + (y_k-y^*),\mu_k^R) = \nabla M(v_k\,;y^*,\mu_k^R) + O(\|y_k-y^*\|). \tag{4.38}$$

A Taylor-series expansion of the inner product $[\nabla M(v\,;y^*,\mu_k^R)]_{\mathcal{A}_+}^T[d_k]_{\mathcal{A}_+}$ at $v=v^*$ gives

$$
\begin{aligned}
[\nabla M(v_k\,;y^*,\mu_k^R)]_{\mathcal{A}_+}^T[d_k]_{\mathcal{A}_+} &= [d_k]_{\mathcal{A}_+}^T[\nabla M(v^*+(v_k-v^*)\,;y^*,\mu_k^R)]_{\mathcal{A}_+}\\
&= [d_k]_{\mathcal{A}_+}^T[\nabla M(v^*\,;y^*,\mu_k^R)]_{\mathcal{A}_+} + O\Big(\frac{1}{\mu_k^R}\|[d_k]_{\mathcal{A}_+}\|\,\|v_k-v^*\|\Big).
\end{aligned}
$$

In order to bound the last term on the right-hand side, we substitute the value $\mu_k^R = r(x_k,y_k)^\gamma$ implied by Lemma 4.4, and apply the estimate $r(x_k,y_k)=\Theta\big(\delta(x_k,y_k)\big)$ from Lemma 4.1. If the resulting value is used with the value $\|[d_k]_{\mathcal{A}_+}\| = O(\|d_k\|) = O(\delta(x_k,y_k))$ of Lemma 4.6(i), then

$$
[\nabla M(v_k\,;y^*,\mu_k^R)]_{\mathcal{A}_+}^T[d_k]_{\mathcal{A}_+} = [d_k]_{\mathcal{A}_+}^T[\nabla M(v^*\,;y^*,\mu_k^R)]_{\mathcal{A}_+} + O\big(\delta(x_k,y_k)^{1-\gamma}\|v_k-v^*\|\big).
$$

This estimate can be combined with (4.38) to obtain

$$
\begin{aligned}
[\nabla M_k]_{\mathcal{A}_+}^T[d_k]_{\mathcal{A}_+} &= [d_k]_{\mathcal{A}_+}^T[\nabla M(v^*\,;y^*,\mu_k^R)]_{\mathcal{A}_+}\\
&\quad + O(\delta(x_k,y_k)^{1-\gamma}\|v_k-v^*\|) + O(\|[d_k]_{\mathcal{A}_+}\|\,\|y_k-y^*\|).
\end{aligned}\tag{4.39}
$$

As $v^*=(x^*,y^*)$ is a primal-dual KKT pair for problem (NP), it follows from the definition of $\mathcal{A}_+$ that $[\nabla M(v^*\,;y^*,\mu_k^R)]_{\mathcal{A}_+} = [g(x^*)-J(x^*)^Ty^*]_{\mathcal{A}_+} > 0$. Combining this with $[d_k]_{\mathcal{A}_+}\le 0$ from (2.26) yields

$$
[\nabla M(v^*\,;y^*,\mu_k^R)]_{\mathcal{A}_+}^T[d_k]_{\mathcal{A}_+} \le -\kappa_4\|[d_k]_{\mathcal{A}_+}\| \text{ for some positive } \kappa_4. \tag{4.40}
$$

As $\gamma<1$, the limit $\delta(x_k,y_k)\to 0$ and estimates (4.39)–(4.40) imply that

$$
[\nabla M_k]_{\mathcal{A}_+}^T[d_k]_{\mathcal{A}_+} \le -\tfrac{1}{2}\kappa_4\|[d_k]_{\mathcal{A}_+}\| \text{ for } k\in\mathcal{S}_* \text{ sufficiently large.}
$$

The combination of this inequality with (4.37) gives

$$
[B_kd_k]_{\mathcal{A}_+}^T[d_k]_{\mathcal{A}_+} + [\nabla M_k]_{\mathcal{A}_+}^T[d_k]_{\mathcal{A}_+} \le \kappa_3\delta(x_k,y_k)^{1-\gamma}\|[d_k]_{\mathcal{A}_+}\| - \tfrac{1}{2}\kappa_4\|[d_k]_{\mathcal{A}_+}\|, \tag{4.41}
$$

for all $k\in\mathcal{S}_*$ sufficiently large.

Finally, consider the last two terms of (4.31) associated with the set $\mathcal{A}_0$. As $k\in\mathcal{S}$, it holds that $y_k^E = y_k$ and $\pi_k = \pi(x_k\,;y_k^E,\mu_k^R) = y_k - c(x_k)/\mu_k^R$. Let $\widetilde{y}_k$ denote the vector $\widetilde{y}_k = \pi_k + \nu(\pi_k - y_k) = y_k - (1+\nu)c(x_k)/\mu_k^R$. The definitions of ∇M_k and B_k, together with the primal-dual partition of d_k give

$$
\begin{aligned}
&[\nabla M_k + B_kd_k]_{\mathcal{A}_0}\\
&\quad= [g(x_k) - J(x_k)^T\widetilde{y}_k + H(x_k,y_k)p_k + \frac{1}{\mu_k^R}(1+\nu)J(x_k)^TJ(x_k)p_k + \nu J(x_k)^Tq_k]_{\mathcal{A}_0}\\
&\quad= [g(x_k) - J(x_k)^T\widetilde{y}_k + H(x_k,y_k)p_k - \frac{1}{\mu_k^R}(1+\nu)J(x_k)^Tc(x_k) - J(x_k)^Tq_k]_{\mathcal{A}_0}\\
&\quad= [g(x_k) - J(x_k)^Ty_k + H(x_k,y_k)p_k - J(x_k)^Tq_k]_{\mathcal{A}_0}.
\end{aligned}\tag{4.42}
$$

It follows from (4.42) and a Taylor-series expansion with respect to x of $g(x)-J(x)^T(y_k+q_k)$ that

$$
\begin{aligned}
[\nabla M_k + B_kd_k]_{\mathcal{A}_0} &= \big[g(x_k+p_k) - J(x_k+p_k)^T(y_k+q_k) + o(\|(p_k,q_k)\|)\big]_{\mathcal{A}_0}\\
&= \big[g(\widehat{x}_k) - J(\widehat{x}_k)^T\widehat{y}_k + o(\|(p_k,q_k)\|)\big]_{\mathcal{A}_0},
\end{aligned}\tag{4.43}
$$

where $(\widehat{x}_k, \widehat{y}_k) = (x_k + p_k, y_k + q_k)$. The definition of $r(x, y)$ and part (ii) of Lemma 4.6 give

$$r(\widehat{x}_k, \widehat{y}_k) \geq \bigl|\min\bigl([\,\widehat{x}_k\,]_i, [\,g(\widehat{x}_k) - J(\widehat{x}_k)^T\widehat{y}_k\,]_i\bigr)\bigr| \\ = \bigl|\min(0, [\,g(\widehat{x}_k) - J(\widehat{x}_k)^T\widehat{y}_k\,]_i)\bigr|, \quad \text{for all } i \in \mathcal{A}_0. \tag{4.44}$$

There are two possible cases for each $i \in \mathcal{A}_0$, depending on the sign of $[\,g(\widehat{x}_k) - J(\widehat{x}_k)^T\widehat{y}_k\,]_i$. If $[\,g(\widehat{x}_k) - J(\widehat{x}_k)^T\widehat{y}_k\,]_i \geq 0$, then the property that $[\,d_k\,]_i \leq 0$ for every $i \in \mathcal{A}$ implies that $[\,g(\widehat{x}_k) - J(\widehat{x}_k)^T\widehat{y}_k\,]_i[\,d_k\,]_i \leq 0$. The expression for $[\,\nabla M_k + B_k d_k\,]_i[\,d_k\,]_i$ from (4.43), and the result that $\|(p_k, q_k)\| = O\bigl(\delta(x_k, y_k)\bigr)$ from Lemma 4.6(i) gives

$$[\,\nabla M_k + B_k d_k\,]_i[\,d_k\,]_i = \bigl[g(\widehat{x}_k) - J(\widehat{x}_k)^T\widehat{y}_k\bigr]_i[\,d_k\,]_i + o(\|(p_k, q_k)\|)[\,d_k\,]_i \\ = o\bigl(\delta(x_k, y_k)\bigr)\bigl|[\,d_k\,]_i\bigr|.$$

Alternatively, if $i \in \mathcal{A}_0$ and $[\,g(\widehat{x}_k) - J(\widehat{x}_k)^T\widehat{y}_k\,]_i < 0$, then $[\,\nabla M_k + B_k d_k\,]_i[\,d_k\,]_i$ satisfies

$$\begin{aligned} [\,\nabla M_k + B_k d_k\,]_i[\,d_k\,]_i &= [\,g(\widehat{x}_k) - J(\widehat{x}_k)^T\widehat{y}_k + o(\|(p_k, q_k)\|)\,]_i[\,d_k\,]_i \\ &\leq r(\widehat{x}_k, \widehat{y}_k)\bigl|[\,d_k\,]_i\bigr| + o\bigl(\delta(x_k, y_k)\bigr)\bigl|[\,d_k\,]_i\bigr| && \bigl((4.44) \text{ and Lemma } 4.6(\text{i})\bigr) \\ &\leq \kappa\delta(\widehat{x}_k, \widehat{y}_k)\bigl|[\,d_k\,]_i\bigr| + o\bigl(\delta(x_k, y_k)\bigr)\bigl|[\,d_k\,]_i\bigr| && (\text{Lemma } 4.1) \\ &= O\bigl(\delta(x_k, y_k)^{1+\gamma}\bigr)|[\,d_k\,]_i| + o\bigl(\delta(x_k, y_k)\bigr)\bigl|[\,d_k\,]_i\bigr| && (\text{Lemma } 4.7) \\ &= o\bigl(\delta(x_k, y_k)\bigr)\bigl|[\,d_k\,]_i\bigr|. \end{aligned}$$

A combination of the two cases provides the estimate

$$[\,\nabla M_k + B_k d_k\,]_{\mathcal{A}_0}^T[\,d_k\,]_{\mathcal{A}_0} \leq o\bigl(\delta(x_k, y_k)\bigr)\,\|[\,d_k\,]_{\mathcal{A}_0}\|. \tag{4.45}$$

It now follows from (4.31), (4.36), (4.41), (4.45), and $\lim_{k\in\mathcal{S}_*} d_k = 0$ that there exist positive constants κ_5, κ_6, and κ_7 such that

$$\nabla M_k^T d_k + \bar{\sigma} d_k^T B_k d_k \\ \leq -\kappa_5\|[\,d_k\,]_{\mathcal{F}_0}\|^2 + \kappa_6\delta(x_k, y_k)^{1-\gamma}\|[\,d_k\,]_{\mathcal{A}_+}\| - \kappa_7\|[\,d_k\,]_{\mathcal{A}_+}\| + o\bigl(\delta(x_k, y_k)\bigr)\,\|[\,d_k\,]_{\mathcal{A}_0}\|.$$

As $\lim_{k\in\mathcal{S}_*}\delta(x_k, y_k) = 0$, it must hold that $\kappa_6\delta(x_k, y_k)^{1-\gamma} \leq \frac{1}{2}\kappa_7$ for all $k \in \mathcal{S}_*$ sufficiently large, which gives

$$\nabla M_k^T d_k + \bar{\sigma} d_k^T B_k d_k \leq -\kappa_5\|[\,d_k\,]_{\mathcal{F}_0}\|^2 - \tfrac{1}{2}\kappa_7\|[\,d_k\,]_{\mathcal{A}_+}\| + o\bigl(\delta(x_k, y_k)\bigr)\,\|[\,d_k\,]_{\mathcal{A}_0}\|. \tag{4.46}$$

The next step of the proof is to show that the right-hand side of the inequality (4.46) is bounded above by a positive multiple of $-\|d_k\|^2$. Consider the sequence $v_k^p = \bigl(x^*, y_P^*(\widehat{y}_k)\bigr)$, where $y_P^*(\cdot)$ is given by (4.4) and satisfies the second-order sufficient conditions for all k. The triangle inequality and substitution of $\widehat{v}_k$ for $v_k + d_k$ yields

$$\|v_k - v_k^p\| = \|v_k + d_k - v_k^p - d_k\| = \|\widehat{v}_k - v_k^p - d_k\| \leq \|\widehat{v}_k - v_k^p\| + \|d_k\|. \tag{4.47}$$

By definition, $\|\widehat{v}_k - v_k^p\| = \delta(\widehat{x}_k, \widehat{y}_k)$, and the estimate $\delta(\widehat{x}_k, \widehat{y}_k) = o\bigl(\delta(x_k, y_k)\bigr)$ given by Lemma 4.7 implies that $\delta(\widehat{x}_k, \widehat{y}_k) \leq \frac{1}{2}\delta(x_k, y_k)$ for k sufficiently large. In addition, the definition of $\delta(x_k, y_k)$ is such that $\delta(x_k, y_k) \leq \|v_k - v_k^p\|$. If these inequalities are used to estimate $\|d_k\|$ in (4.47), then

$$-\|d_k\| \leq \|\widehat{v}_k - v_k^p\| - \|v_k - v_k^p\| \leq -\tfrac{1}{2}\delta(x_k, y_k). \tag{4.48}$$

Consider the inequality (4.46). Suppose that k is sufficiently large that $\kappa_5\|[\,d_k\,]_{\mathcal{F}_0}\| \le \frac{1}{4}\kappa_7$. Standard norm inequalities applied in conjunction with the estimates $\|d_k\| \le \|[\,d_k\,]_{\mathcal{F}_0}\| + \|[\,d_k\,]_{\mathcal{A}_+}\|$, $\|[\,d_k\,]_{\mathcal{A}_0}\| \le \|[\,d_k\,]_{\mathcal{F}_0}\|$, and $\|d_k\| \ge \frac{1}{2}\delta(x_k, y_k)$ from (4.48), give

$$\begin{aligned}
&-\kappa_5\|[\,d_k\,]_{\mathcal{F}_0}\|^2 - \tfrac{1}{2}\kappa_7\|[\,d_k\,]_{\mathcal{A}_+}\| + o\big(\delta(x_k,y_k)\big)\,\|[\,d_k\,]_{\mathcal{A}_0}\| \\
&\qquad \le -\kappa_5\|[\,d_k\,]_{\mathcal{F}_0}\|^2 - \tfrac{1}{4}\kappa_7\|[\,d_k\,]_{\mathcal{A}_+}\| - \tfrac{1}{2}\kappa_5\|[\,d_k\,]_{\mathcal{F}_0}\|\,\|[\,d_k\,]_{\mathcal{A}_+}\| + o\big(\delta(x_k,y_k)\big)\,\|[\,d_k\,]_{\mathcal{A}_0}\| \\
&\qquad \le -\tfrac{1}{2}\kappa_5\|[\,d_k\,]_{\mathcal{F}_0}\|^2 - \tfrac{1}{4}\kappa_7\|[\,d_k\,]_{\mathcal{A}_+}\| - \tfrac{1}{2}\kappa_5\|d_k\|\,\|[\,d_k\,]_{\mathcal{F}_0}\| + o\big(\delta(x_k,y_k)\big)\,\|[\,d_k\,]_{\mathcal{A}_0}\| \\
&\qquad \le -\tfrac{1}{4}\kappa_7\|[\,d_k\,]_{\mathcal{A}_+}\| - \tfrac{1}{2}\kappa_5\|[\,d_k\,]_{\mathcal{F}_0}\|^2 - \tfrac{1}{4}\kappa_5\delta(x_k,y_k)\,\|[\,d_k\,]_{\mathcal{F}_0}\| + o\big(\delta(x_k,y_k)\big)\,\|[\,d_k\,]_{\mathcal{A}_0}\| \\
&\qquad \le -\tfrac{1}{4}\kappa_7\|[\,d_k\,]_{\mathcal{A}_+}\| - \tfrac{1}{2}\kappa_5\|[\,d_k\,]_{\mathcal{F}_0}\|^2 - \tfrac{1}{4}\kappa_5\delta(x_k,y_k)\,\|[\,d_k\,]_{\mathcal{A}_0}\| + o\big(\delta(x_k,y_k)\big)\,\|[\,d_k\,]_{\mathcal{A}_0}\| \\
&\qquad \le -\tfrac{1}{4}\kappa_7\|[\,d_k\,]_{\mathcal{A}_+}\| - \tfrac{1}{2}\kappa_5\|[\,d_k\,]_{\mathcal{F}_0}\|^2 \\
&\qquad \le -\tfrac{1}{4}\kappa_7\|[\,d_k\,]_{\mathcal{A}_+}\|^2 - \tfrac{1}{2}\kappa_5\|[\,d_k\,]_{\mathcal{F}_0}\|^2.
\end{aligned}$$

These inequalities, when used with (4.46), imply that

$$\nabla M_k^T d_k + \bar{\sigma} d_k^T B_k d_k \le -\kappa_5\|[\,d_k\,]_{\mathcal{F}_0}\|^2 - \tfrac{1}{2}\kappa_7\|[\,d_k\,]_{\mathcal{A}_+}\| + o\big(\delta(x_k,y_k)\big)\,\|[\,d_k\,]_{\mathcal{A}_0}\| \le -\bar{c}\,\|d_k\|^2, \quad (4.49)$$

with $\bar{c} = \min\{\frac{1}{4}\kappa_7, \frac{1}{2}\kappa_5\}$. This establishes the first part of (4.27).

To prove the second part of (4.27), the bounds on $\nabla M_k^T d_k + \bar{\sigma} d_k^T B_k d_k$ and $d_k^T B_k d_k$ given by (4.46) and (4.36) imply that

$$\begin{aligned}
\nabla M_k^T d_k &= \nabla M_k^T d_k + \bar{\sigma} d_k^T B_k d_k - \bar{\sigma} d_k^T B_k d_k \\
&\le -\kappa_5\|[\,d_k\,]_{\mathcal{F}_0}\|^2 - \tfrac{1}{2}\kappa_7\|[\,d_k\,]_{\mathcal{A}_+}\| + o\big(\delta(x_k,y_k)\big)\,\|[\,d_k\,]_{\mathcal{A}_0}\| \\
&\qquad - \bar{\sigma}\kappa_1\|[\,d_k\,]_{\mathcal{F}_0}\|^2 + \bar{\sigma}\kappa_2\delta(x_k,y_k)^{1-\gamma}\|[\,d_k\,]_{\mathcal{A}_+}\|.
\end{aligned} \quad (4.50)$$

As $\lim_{k\in\mathcal{S}_*} d_k = 0$, there is an index k sufficiently large that $\bar{\sigma}\kappa_2\delta(x_k,y_k)^{1-\gamma} \le \frac{1}{4}\kappa_7$, and the bound (4.50) may be written in the form

$$\nabla M_k^T d_k \le -(\kappa_5 + \bar{\sigma}\kappa_1)\|[\,d_k\,]_{\mathcal{F}_0}\|^2 - \tfrac{1}{4}\kappa_7\|[\,d_k\,]_{\mathcal{A}_+}\| + o\big(\delta(x_k,y_k)\big)\,\|[\,d_k\,]_{\mathcal{A}_0}\|, \quad (4.51)$$

which is the inequality (4.46) with different positive constants. If the argument used to derive (4.49) is repeated for the inequality (4.51), it follows that there is a positive constant $\widehat{c}$ such that $\nabla M_k^T d_k \le -\widehat{c}\,\|d_k\|^2$. From Assumption 4.3, d_k is nonzero, which implies that d_k is a strict descent direction for $M(v\,; y_k^E, \mu_k^R)$ at v_k. ∎

Lemma 4.9 establishes that the last of the three conditions (2.29) needed for the acceptance of the local descent direction d_k holds for all $k \in \mathcal{S}_*$ sufficiently large.

Theorem 4.1. *For all $k \in \mathcal{S}_* \subseteq \mathcal{S}$ sufficiently large, it holds that:*

(i) *a local descent direction $d_k = (p_k, q_k)$ is computed in Step 8 of Algorithm 2;*

(ii) *$v_k + d_k$ is feasible, $[\,\nabla\mathcal{Q}_k(v_k + d_k\,; \mu_k^R)\,]_{\mathcal{A}_\epsilon} \ge -t_k e$, and $\nabla M_k^T d_k < 0$, i.e., the conditions in Step 9 of Algorithm 2 are satisfied; and*

(iii) *$\mathcal{A}_\epsilon(x_k, y_k, \mu_k^R) = \mathcal{A}(x^*) = \mathcal{A}(x_k + p_k)$.*

Proof. Part (i) follows from Lemma 4.6. Part (ii) follows from Lemmas 4.6(ii) and 4.9. It remains to prove part (iii). The equality $\mathcal{A}_\epsilon(x_k, y_k, \mu_k^R) = \mathcal{A}(x^*)$ is established in Lemma 4.2(ii). Suppose that $i \in \mathcal{A}(x^*) = \mathcal{A}_\epsilon(x_k, y_k, \mu_k^R)$. The definition of d_k in Steps 7–8 of Algorithm 2 implies that $[\,x_k + p_k\,]_i = 0$, which gives $i \in \mathcal{A}(x_k + p_k)$. For the reverse inclusion, suppose that $i \notin \mathcal{A}(x^*)$, i.e., $x_i^* > 0$. In this case, the assumption that $\lim_{k\in\mathcal{S}_*} x_k = x^*$ implies that $[\,x_k\,]_i \ge \frac{1}{2}x_i^*$ for all $k \in \mathcal{S}_*$ sufficiently large. Part (i) of Lemma 4.6 gives $\max\{\|p_k\|, \|q_k\|\} = O\big(\delta(x_k, y_k)\big)$, and the assumption $\lim_{k\in\mathcal{S}_*}(x_k, y_k) = (x^*, y^*)$ implies that $\lim_{k\in\mathcal{S}_*} \delta(x_k, y_k) = 0$. It follows that $\lim_{k\in\mathcal{S}_*} p_k = 0$, with $[\,x_k + p_k\,]_i \ge \frac{1}{2}x_i^* + [\,p_k\,]_i \ge \frac{1}{3}x_i^* > 0$ for all $k \in \mathcal{S}_*$ sufficiently large, which means that $i \notin \mathcal{A}(x_k + p_k)$. This completes the proof. ∎

The next result shows that the flexible line search returns the unit step length for all $k \in \mathcal{S}_*$ sufficiently large.

Theorem 4.2. *If Assumptions 4.1 and 4.2 hold, then $\alpha_k = 1$ for all $k \in \mathcal{S}_* \subseteq \mathcal{S}$ sufficiently large.*

Proof. Throughout the proof, the quantities $M(v\,; y_k^E, \mu_k^R)$, $\nabla M(v\,; y_k^E, \mu_k^R)$, and $\nabla^2 M(v_k\,; y_k^E, \mu_k^R)$ are denoted by $M(v)$, $\nabla M(v)$, and $\nabla^2 M_k$. Assumption 4.3 and part (vi) of Lemma 4.2 imply that the first-order line-search model is used for all $k \in \mathcal{S}_* \subseteq \mathcal{S}$ sufficiently large, i.e., the quantity ι_k is set to one in Algorithm 5. A Taylor-series expansion of $M(v_k + d_k)$ gives

$$\begin{aligned} M(v_k + d_k) &= M(v_k) + \nabla M(v_k)^T d_k + \tfrac{1}{2} d_k^T \nabla^2 M_k d_k + O\Big(\frac{1}{\mu_k^R}\|d_k\|^3\Big) \\ &= M(v_k) + \nabla M(v_k)^T d_k + \tfrac{1}{2} d_k^T \nabla^2 M_k d_k + O\big(\delta(x_k, y_k)^{1-\gamma}\|d_k\|^2\big), \end{aligned}$$

where the bound on the last term follows from the sequence of estimates

$$(1/\mu_k^R)\|d_k\| = r(x_k, y_k)^{-\gamma}\|d_k\| = O\big(\delta(x_k, y_k)^{-\gamma}\big)\|d_k\| = O\big(\delta(x_k, y_k)^{1-\gamma}\big)$$

derived in Lemmas 4.4, 4.1, and 4.6(i).

Let the scalar $\bar\sigma$ of Lemma 4.9 be defined so that $(1-\gamma_s)\bar\sigma = \frac{1}{2}$, where γ_s $(0 < \gamma_s < \frac{1}{2})$ is the parameter used for the modified Armijo condition (2.33) in the flexible line search of Algorithm 4. With this definition, $\bar\sigma$ satisfies $0 < \bar\sigma < 1$, and the application of Lemma 4.9 with $\bar\sigma = \frac{1}{2}(1-\gamma_s)^{-1}$ gives

$$\begin{aligned} &M(v_k + d_k) - M(v_k) - \gamma_s \nabla M(v_k)^T d_k \\ &= \ (1-\gamma_s)\nabla M(v_k)^T d_k + \tfrac{1}{2} d_k^T \nabla^2 M_k d_k + O\big(\delta(x_k, y_k)^{1-\gamma}\|d_k\|^2\big) \\ &\le \ [\tfrac{1}{2} - (1-\gamma_s)\bar\sigma] d_k^T B_k d_k - (1-\gamma_s)\bar c\,\|d_k\|^2 + \tfrac{1}{2}\|\nabla^2 M_k - B_k\|\,\|d_k\|^2 + O\big(\delta(x_k, y_k)^{1-\gamma}\|d_k\|^2\big) \\ &= -(1-\gamma_s)\bar c\,\|d_k\|^2 + \tfrac{1}{2}\|\nabla^2 M_k - B_k\|\,\|d_k\|^2 + O\big(\delta(x_k, y_k)^{1-\gamma}\|d_k\|^2\big), \end{aligned}$$

for all $k \in \mathcal{S}_*$ sufficiently large. The global convergence property of Assumption 4.2(i) implies that $\lim_{k\in\mathcal{S}_*} \delta(x_k, y_k) = 0$, which gives $\lim_{k\in\mathcal{S}_*} d_k = 0$ from part (i) of Lemma 4.6. In addition, Lemma 4.8 implies that $\lim_{k\in\mathcal{S}_*} \|\nabla^2 M_k - B_k\| = 0$. The combination of these results gives the estimate

$$M(v_k + d_k) - M(v_k) - \gamma_s \nabla M(v_k)^T d_k \le -(1-\gamma_s)\bar c\,\|d_k\|^2 + o(\|d_k\|^2) < 0,$$

for all $k \in \mathcal{S}_*$ sufficiently large. The fact that $\alpha_k = 1$ for all $k \in \mathcal{S}_*$ sufficiently large follows from the previous displayed inequality, $d_k \neq 0$, $s_k = 0$ (see Lemma 4.2(vi)), $\iota_k = 1$, and the structure of the line-search procedure given by Algorithm 4. ∎

The next theorem shows that, for k sufficiently large, the properties established in Lemmas 4.1–4.9 and Theorems 4.1–4.2 hold for every k, not just those in the set $\mathcal{S}_* \subseteq \mathcal{S}$.

Theorem 4.3. *For any positive ϵ sufficiently small, and any ρ such that $1 < \rho < 1+\gamma$, there exists a V-iteration index $k_V = k_V(\epsilon)$ such that the following results hold for every $k \geq k_V$:*

(i) $\|(x_k - x^*, y_k - y^*)\| \leq \epsilon$;

(ii) $\delta(x_{k+1}, y_{k+1}) \leq \delta(x_k, y_k)^\rho$;

(iii) *k is a V-iterate; and*

(iv) *the results of Lemmas 4.1–4.9 and Theorems 4.1–4.2 hold.*

Proof. Let the positive scalar ϵ be sufficiently small that the results of Lemmas 4.1–4.9 and Theorems 4.1–4.2 hold for every V-O iterate (x_k, y_k) satisfying $\|(x_k - x^*, y_k - y^*)\| \leq \epsilon$. (The proof of (iv) establishes that these results hold for *every* k sufficiently large.)

Let (x_k, y_k) be a primal-dual iterate with $k \in \mathcal{S}_*$. Theorem 4.1 implies that the unit step is accepted in the line search, in which case $(x_{k+1}, y_{k+1}) = (x_k + p_k, y_k + q_k)$. Let κ be the positive scalar defined in Lemma 4.1. Similarly, let c_1 ($c_1 > 0$) and c_2 ($c_2 \geq 1$) denote constants such that

$$\max\{\|x_{k+1} - x_k\|, \|y_{k+1} - y_k\|\} \leq c_1\delta(x_k, y_k), \quad \text{and} \quad \delta(x_{k+1}, y_{k+1}) \leq c_2\delta(x_k, y_k)^{1+\gamma}. \tag{4.52}$$

(The existence of c_1 and c_2 is implied by the results of Lemmas 4.6(i) and 4.7.)

If ρ is any scalar satisfying $1 < \rho < 1+\gamma$, let $k_V = k_V(\epsilon)$ be an index in $\mathcal{S}_* \subseteq \mathcal{S}$ that is sufficiently large that (x_{k_V}, y_{k_V}) is a V-iterate and satisfies the conditions

$$\max\{\, \|x_{k_V} - x^*\|,\ \|y_{k_V} - y^*\|,\ 2c_1\delta_V,\ 2c_1\delta_V^\rho/(1-\delta_V^\rho)\,\} \leq \tfrac{1}{4}\epsilon, \quad \text{and} \tag{4.53}$$

$$\max\Big\{\, 2\kappa^{\rho+2}\delta_V^{\rho-1}/\beta,\ c_2\delta_V^{1+\gamma-\rho},\ \delta_V^\rho \,\Big\} \leq 1, \tag{4.54}$$

where $\delta_V = \delta(x_{k_V}, y_{k_V})$, and β ($0 < \beta < 1$) is the weight used in the definitions of $\phi_V(x,y)$ and $\phi_V(x,y)$. The following argument shows that an index κ_V satisfying these conditions must exist. As $\lim_{k\in\mathcal{S}_*}(x_k, y_k) = (x^*, y^*)$, it must hold that the optimality and feasibility measures (2.12) give $\lim_{k\in\mathcal{S}_*}\phi_V(x_k, y_k) = 0$ and $\lim_{k\in\mathcal{S}_*}\phi_O(x_k, y_k) = 0$. As Assumption 4.2(i) implies that there are infinitely many V-O-iterates, and the condition $\phi_V(v_k) \leq \frac{1}{2}\phi_V^{\max}$ for a V-iteration is checked before the condition for an O-iteration, then there must be infinitely many V-iterates. In addition, as $\lim_{k\in\mathcal{S}_*}\delta(x_k, y_k) = 0$, there must be an index $k = k_V$ such that $\delta_V = \delta(x_k, y_k)$ is sufficiently small to give (4.53) and (4.54).

An inductive argument is used to prove that parts (i)–(iv) hold for all $k \geq k_V$. The base case is $k = k_V$. The definition of k_V implies that $k = k_V$ is a V-iteration index, and it follows trivially that part (iii) holds. Moreover, the assumption (4.53) and standard norm inequalities yield

$$\|(x_{k_V} - x^*, y_{k_V} - y^*)\| \leq \|x_{k_V} - x^*\| + \|y_{k_V} - y^*\| \leq \tfrac{1}{4}\epsilon + \tfrac{1}{4}\epsilon < \epsilon, \tag{4.55}$$

which establishes part (i) for $k = k_V$. It follows immediately from (4.55) and the choice of ϵ that part (iv) holds for $k = k_V$. As part (iv) holds for $k = k_V$, (4.52), and (4.54) may be combined to give

$$\delta(x_{k_V+1}, y_{k_V+1}) \le c_2 \delta_V^{1+\gamma} = c_2 \delta_V^{1+\gamma-\rho} \delta_V^{\rho} \le \delta_V^{\rho},$$

which establishes part (ii) for $k = k_V$. This completes the base case $k = k_V$.

The inductive hypothesis is that (i)–(iv) hold for every iterate k such that $k_V \le k \le k_V + j - 1$. Under this hypothesis, it must be shown (i)–(iv) hold for $k = k_V + j$. For (i), standard norm inequalities give

$$\begin{aligned}\left\| \begin{pmatrix} x_{k_V+j} - x^* \\ y_{k_V+j} - y^* \end{pmatrix} \right\| &\le \|x_{k_V+j} - x^*\| + \|y_{k_V+j} - y^*\| \\ &= \left\| \sum_{l=0}^{j-1} (x_{k_V+l+1} - x_{k_V+l}) + x_{k_V} - x^* \right\| + \left\| \sum_{l=0}^{j-1} (y_{k_V+l+1} - y_{k_V+l}) + y_{k_V} - y^* \right\| \\ &\le \sum_{l=0}^{j-1} \big(\|x_{k_V+l+1} - x_{k_V+l}\| + \|y_{k_V+l+1} - y_{k_V+l}\| \big) + \|x_{k_V} - x^*\| + \|y_{k_V} - y^*\| \\ &\le 2c_1 \sum_{l=0}^{j-1} \delta(x_{k_V+l}, y_{k_V+l}) + \tfrac{1}{2}\epsilon,\end{aligned}$$

where the first inequality of (4.52) has been used to bound each of the terms in the summation, and the term $\|x_{k_V} - x^*\| + \|y_{k_V} - y^*\|$ is estimated by (4.55). It follows from the inductive hypothesis for part (ii) and (4.53) that

$$\left\| \begin{pmatrix} x_{k_V+j} - x^* \\ y_{k_V+j} - y^* \end{pmatrix} \right\| = 2c_1 \Big[\delta_V + \sum_{i=1}^{j-1} \delta_V^{i\rho} \Big] + \tfrac{1}{2}\epsilon < 2c_1 \Big[\delta_V + \frac{\delta_V^{\rho}}{1 - \delta_V^{\rho}} \Big] + \tfrac{1}{2}\epsilon \le \epsilon,$$

which establishes that part (i) holds for $k = k_V + j$.

The next stage of the proof involves establishing that part (iii) holds for $k = k_V + j$. For all $k \ge k_V$, it holds that $\xi_k^{(1)} = 0$ and the feasibility measure ϕ_V satisfies

$$\beta r(x_k, y_k) \le \phi_V(x_k, y_k) = \eta(x_k) + \beta\omega(x_k, y_k) \le 2r(x_k, y_k) \le 2\kappa\delta(x_k, y_k),$$

where the last inequality follows from Lemma 4.1. Applying these inequalities at (x_{k_V+j}, y_{k_V+j}), together with Lemma 4.1 and the induction assumption (ii) at $(x_{k_V+j-1}, y_{k_V+j-1})$, gives

$$\begin{aligned}\phi_V(x_{k_V+j}, y_{k_V+j}) &\le 2\kappa\delta(x_{k_V+j}, y_{k_V+j}) \le 2\kappa\delta(x_{k_V+j-1}, y_{k_V+j-1})^{\rho} \\ &\le 2\kappa^{\rho+1} r(x_{k_V+j-1}, y_{k_V+j-1})^{\rho} \\ &= 2\kappa^{\rho+1} r(x_{k_V+j-1}, y_{k_V+j-1})^{\rho-1} r(x_{k_V+j-1}, y_{k_V+j-1}) \\ &\le (2\kappa^{\rho+1}/\beta) r(x_{k_V+j-1}, y_{k_V+j-1})^{\rho-1} \phi_V(x_{k_V+j-1}, y_{k_V+j-1}).\end{aligned} \tag{4.56}$$

If $\phi_{V,k}^{\max}$ denotes the value of the bound $\phi_V^{\max}$ of (2.13) at the start of iteration k, then the assumption that $(x_{k_V+j-1}, y_{k_V+j-1})$ is a V-iterate implies that $\phi_V(x_{k_V+j-1}, y_{k_V+j-1}) \le \frac{1}{2}\phi_{V,k_V+j-1}^{\max}$. This allows the bound (4.56) to be extended so that

$$\begin{aligned}\phi_V(x_{k_V+j}, y_{k_V+j}) &\le (\kappa^{\rho+1}/\beta) r(x_{k_V+j-1}, y_{k_V+j-1})^{\rho-1} \phi_{V,k_V+j-1}^{\max} \\ &\le (\kappa^{\rho+2}/\beta) \delta(x_{k_V+j-1}, y_{k_V+j-1})^{\rho-1} \phi_{V,k_V+j-1}^{\max} \\ &\le (\kappa^{\rho+2}\delta_V^{\rho-1}/\beta) \phi_{V,k_V+j-1}^{\max} \le \tfrac{1}{2}\phi_{V,k_V+j-1}^{\max}.\end{aligned}$$

The final inequality, which follows from (4.54), implies that $k_V + j$ is a V-iterate. This establishes that part (iii) holds for $k = k_V + j$, as required. Part (iv) then follows immediately from the choice of ϵ and the fact that (i) and (iii) hold at $k = k_V + j$.

It remains to show that (ii) holds for $k = k_V + j$. It follows from the bound (4.54) and definition of ρ ($\rho > 1$), that

$$c_2(\delta_V^{j\rho})^{1+\gamma-\rho} \le c_2\delta_V^{\rho(1+\gamma-\rho)} \le c_2\delta_V^{1+\gamma-\rho} \le 1. \tag{4.57}$$

This inequality, in conjunction with the induction hypotheses of parts (ii) and (iv), and Lemma 4.7, give

$$\begin{aligned}\delta(x_{k_V+j+1}, y_{k_V+j+1}) &\le c_2\delta(x_{k_V+j}, y_{k_V+j})^{1+\gamma} = c_2\delta(x_{k_V+j}, y_{k_V+j})^{1+\gamma-\rho}\delta(x_{k_V+j}, y_{k_V+j})^{\rho}\\ &\le c_2(\delta_V^{j\rho})^{1+\gamma-\rho}\delta(x_{k_V+j}, y_{k_V+j})^{\rho} \le \delta(x_{k_V+j}, y_{k_V+j})^{\rho},\end{aligned}$$

which shows that part (ii) holds for $k = k_V + j$. This completes the induction proof. ∎

It remains to establish the rate of convergence of the primal-dual iterates to (x^*, y^*). The proof is based on showing that the iterates of Algorithm 5 are equivalent to those of a stabilized SQP method for which superlinear convergence has been established.

Theorem 4.4. *If Assumptions 4.1 and 4.2 are satisfied, then* $\lim_{k\to\infty}(x_k, y_k) = (x^*, y^*)$ *and the convergence rate is superlinear.*

Proof. Since $\epsilon > 0$ was arbitrary in Theorem 4.3, it follows that $\lim_{k\to\infty}(x_k, y_k) = (x^*, y^*)$. It remains to show that the convergence rate is superlinear. Theorem 4.3(iii) shows that the iterates generated by the algorithm are all V-iterates for k sufficiently large. Moreover, Theorem 4.3(iv) implies that Lemmas 4.1–4.9 and Theorems 4.1–4.2 hold for all k sufficiently large (not just for $k \in \mathcal{S}_* \subseteq \mathcal{S}$). It follows that for all k sufficiently large: (a) $\mu_k^R = r(x_k, y_k)^\gamma$ (from Lemma 4.4); (b) $\mathcal{A}(x^*) = \mathcal{A}(x_k) = \mathcal{A}_\epsilon(x_k, y_k, \mu_k^R)$ (from Lemma 4.2(ii)); and (c) $(x_{k+1}, y_{k+1}) = (x_k + p_k, y_k + q_k)$ with every direction (p_k, q_k) a local descent direction (from Theorems 4.2 and 4.1(i)–(iii)). The combination of these results gives $[\, x_k \,]_{\mathcal{A}} = 0$ for all k sufficiently large, where the suffix "$\mathcal{A}$" denotes the components with indices in the optimal active set $\mathcal{A}(x^*)$. It follows that the sequence (x_k, y_k) is identical to the sequence generated by a conventional stabilized SQP method applied to the equality-constrained problem (4.11), i.e., the iterates correspond to performing a conventional stabilized SQP method on problem (NP) having correctly estimated the active set (the associated stabilized QP subproblem is defined in the statement of Lemma 4.5). The superlinear rate convergence of the iterates now follows, for example, from [24, Theorem 1]. ∎

5. Numerical Experiments

This section describes an implementation of algorithm `pdSQP2` and includes the results of some numerical experiments that are designed to validate the algorithm. Section 5.1 provides the details of a preliminary MATLAB implementation. Section 5.2 includes the results of testing `pdSQP2` on the Hock-Schittkowski suite of test problems, which is a commonly used set of problems for assessing the overall performance of a method. Finally, Section 5.3 focuses on the performance of `pdSQP2` on problems that exhibit various forms of degeneracy.

5.1. Implementation details

The numerical experiments were performed using a preliminary implementation of `pdSQP2` written in MATLAB [28]. The control parameter values and their initial values are specified in Table 1. If `pdSQP2` did not converge within $k_{\max} = 1000$ iterations, then it was considered to have failed. The tests used to terminate the algorithm at an approximate solution or an infeasible stationary point are given by (2.35) and (2.36), respectively.

Table 1: Control parameter and initial values required by algorithm `pdSQP2`.

Parameter	Value	Parameter	Value	Parameter	Value
ν	1.0	μ_0^R	1.0e-4	τ_{stop}	1.0e-6
ϵ_a	1.0e-6	μ_1	1.0	β	1.0e-5
γ	0.5	γ_s	1.0e-3	λ	0.2
$y_{\max}$	1.0e+6	θ	1.0e-5	τ_0	1.0
$k_{\max}$	1000	$\phi_V^{\max}, \phi_O^{\max}$	1.0e+3		

All the results included in this paper are from a variant of `pdSQP2` that does not test for a direction of negative curvature until a first-order stationary point is located. Both the global and local convergence analysis remains valid for this version. Other aspects of the implementation that require discussion are the definition of the QP Hessian matrix and the computation of the direction of negative curvature (when one exists). The positive-definite Hessian of the bound-constrained QP problem (2.19) is obtained by using a form of pre-convexification. Specifically, the positive-definite matrix $\widehat{H}$ of (2.20) has the form $\widehat{H}(x_k, y_k) = H(x_k, y_k) + E_k + D_k$ for some positive-semidefinite matrix E_k and positive-semidefinite diagonal matrix D_k, as described in [16, Section 4]. If the matrix formed from the ϵ-free rows and columns of B is positive definite (see (2.5)), then E_k is zero, in which case, the Newton equations (2.28) are not modified. The calculation of the matrix E_k is based on an LBL^T factorization of a matrix in regularized KKT form (see (2.6)). The factorization also provides the direction of negative curvature $u_k^{(1)}$ required by Algorithm 1 (see, e.g., Forsgren [11], Forsgren and Gill [12], and Kungurtsev [26, Chapter 9]). The unique minimizer of the strictly convex QP is found using a MATLAB implementation of the inertia-controlling QP solver by Gill and Wong [17]. The QP solver is called only when a local descent direction is required (see Algorithm 2).

5.2. Performance on the Hock-Schittkowski test problems

This section concerns the performance of `pdSQP2` relative to the first-derivative SQP solver `SNOPT7` [14] on a subset of the Hock-Schittkowski [20] test problems from the CUTEst [18] test collection. The problems `HS85` and `HS87` were omitted from the test set because they are nonsmooth and violate the basic assumptions required by both `pdSQP2` and `SNOPT7`.

Figure 1 depicts the performance profiles associated with the number of function evaluations needed by `SNOPT7` and `pdSQP2`. Performance profiles were proposed by Dolan and Moré [5] to provide a visual comparison of the relative performance of two or more algorithms on a set of problems. The graph associated with `pdSQP2` passes (roughly) through the point $(3, 0.97)$, which implies that on 97% of the problems, the number of function evaluations required by `pdSQP2` was less than 2^3 times the number of function evaluations required by `SNOPT7`.

It follows that the algorithm with a higher value on the vertical axis may be considered the more efficient algorithm, while the algorithm on top at the far right may be considered more reliable.

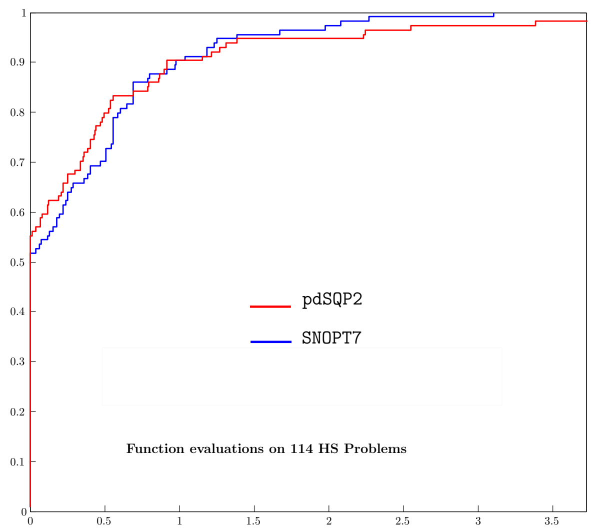

Figure 1: Performance profile for the number of function evaluations on the set of Hock-Schittkowski problems.

Figure 1 suggests that, with respect to function evaluations, `pdSQP2` is more efficient than `SNOPT7` but less reliable. This is to be expected because `SNOPT7` has been developed, tested, and maintained continuously for more than 20 years. Nonetheless, the profile provides some numerical support for the global and local convergence theory. It is likely that a more sophisticated implementation that includes a better tuning of algorithm parameters will improve the efficiency and reliability of `pdSQP2`. In particular, almost all of the Hock-Schittkowski problems requiring more function evaluations than `SNOPT7` involve a significant number of iterations in which convexification is required. Further research on efficient convexification methods is likely to have a significant impact on both robustness and efficiency.

5.3. Performance on degenerate problems

The convergence theory of Section 4 implies that `pdSQP2` is locally superlinearly convergent under relatively weak assumptions that do not include a constraint qualification. This section concerns an investigation of the observed numerical rate of convergence on a number of degenerate problems.

5.3.1. Degenerate CUTEst problems

The local rate of convergence of algorithm `pdSQP2` was investigated for a set of degenerate problems from the CUTEst [18] test set. In particular, 84 problems were identified for which the active-constraint Jacobian is numerically rank deficient at the computed solution. In addition, 56 problems have at least one negligible multiplier associated with a variable on its bound. In this case, a multiplier is considered as being negligible if it is less than τ_{stop} in absolute value. A zero multiplier associated with an active constraint implies that the property of strict complementarity does not hold. A total of 26 problems were identified that fail both the linear independence constraint qualification (LICQ) and strict complementarity.

For these degenerate problems, the order of convergence was estimated by the quantity

$$\text{EOC} = \log r(x_{k_f}, y_{k_f}) / \log r(x_{k_f-1}, y_{k_f-1}), \tag{5.1}$$

where k_f denotes the final computed iterate. The results are given in Table 2. The column with heading "Last is global" contains the statistics for problems for which the final search direction is a global descent direction. The column marked "Last is local" gives the statistics for problems for which the final direction is a local descent direction. Column headed "Last two are local" contains the statistics for problems for which the final two descent steps are local descent directions. The values in parentheses indicate the number of problems that satisfy the weak second-order sufficient optimality conditions, i.e., the Hessian of the Lagrangian is positive definite on the null space of the active constraint Jacobian matrix. In the implementation considered here, this property is considered to hold if the smallest eigenvalue of $Z^T H_{\mathcal{F}_\epsilon} Z$ is greater than $\tau_{\rm stop}$, where the columns of Z form a basis for the null space of $J_{\mathcal{F}_\epsilon}$.

Table 2: The estimated order of convergence for algorithm `pdSQP2` on the degenerate CUTEst test problems

	Last is global	Last is local	Last two are local	Total
Problems not satisfying the LICQ				
$1.25 < \text{EOC}$	20 (7)	16 (12)	33 (31)	69 (50)
$1.1 < \text{EOC} \leq 1.25$	3 (3)	1 (1)	6 (6)	10 (10)
$\text{EOC} \leq 1.1$	3 (2)	0 (0)	2 (2)	5 (4)
Problems not satisfying strict complementarity				
$1.25 < \text{EOC}$	17 (6)	4 (2)	16 (16)	37 (24)
$1.1 < \text{EOC} \leq 1.25$	4 (4)	0 (0)	3 (3)	7 (7)
$\text{EOC} \leq 1.1$	9 (7)	1 (0)	2 (1)	12 (8)
Problems not satisfying strict complementarity and the LICQ				
$1.25 < \text{EOC}$	11 (3)	4 (2)	6 (6)	21 (11)
$1.1 < \text{EOC} \leq 1.25$	2 (2)	0 (0)	2 (2)	4 (4)
$\text{EOC} \leq 1.1$	1 (1)	0 (0)	0 (0)	1 (1)

Table 2 shows that if the LICQ does not hold, but strict complementarity does, then local descent steps are computed in the final stages and they contribute to the superlinear rate of convergence. Moreover, superlinear convergence is typical even when the local descent step is not computed. This observation is consistent with [26, Chapter 8], which shows that the iterates generated by the algorithm `pdSQP` of Gill and Robinson [15] converge superlinearly when the second-order sufficient conditions for optimality hold in conjunction with the property of strict complementarity. In contrast, on those problems for which `pdSQP2` converges to a solution at which strict complementarity fails, the results indicate that linear convergence is just as likely as superlinear convergence.

5.3.2. The degenerate problems of Mostafa, Vicente, and Wright

In [29], Mostafa, Vicente and Wright analyze the performance of a stabilized SQP algorithm proposed by Wright [34] that estimates the weakly and strongly active multipliers. The authors demonstrate that the algorithm is robust in general and converges rapidly on a specified collection of 12 degenerate problems that includes some of the original Hock-Schittkowski problems; several Hock-Schittkowski problems modified to include redundant constraints; and several problems drawn from the literature (see the reference [29] for additional details). All 12 problems have either a rank-deficient Jacobian or at least one weakly active multiplier at the solution.

Algorithm pdSQP2 was tested on ten of the twelve problems that could be coded directly or obtained from other sources. Of the ten cases, pdSQP2 converges superlinearly on seven problems, converges linearly on two problems, and fails to converge on one problem. These results appear to be similar to those obtained by Mostafa, Vicente and Wright using their code sSQPa [29].

5.3.3. Degenerate MPECs

Mathematical programs with equilibium constraints (MPECs) are optimization problems that have variational inequalities as constraints. Various reformulations of MPECs as nonlinear programs (see, e.g., Baumrucker, Renfro and Biegler [1]) include complementarity constraints that do not satisfy either the LICQ or the MFCQ. This feature is generally recognized as the main source of difficulty for conventional nonlinear solvers. In the case of pdSQP2, the violation of the MFCQ implies that Theorem 3.1 cannot be used to guarantee the existence of limit points of the sequence of dual variables. As a consequence, the primal-dual iterates computed by pdSQP2 may never enter a region of superlinear convergence. Nonetheless, as MPECs constitute an important and challenging class of problems, this section includes results from pdSQP2 on a large set of MPECs.

Figure 2 gives the performance profiles that compare the number of function evaluations required by pdSQP2 and SNOPT7 on a set of 86 MPECs obtained from Sven Leyffer at the Argonne National Laboratory. Many of these problems are included in the MPECLib library [4], which is a large and varied collection of MPECs from both theoretical and practical test models. As in the previous section, the performance profiles indicate that SNOPT7 is more robust than pdSQP2 overall. Nevertheless, pdSQP2 is substantially faster than SNOPT7 in a significant number of cases.
This suggests that the stabilized SQP subproblem encourages fast local convergence on this type of degenerate problem.

As discussed above, the theoretical results of Section 4 do not guarantee that the primal-dual iterates will enter a region in which local descent steps are used. In order to study this possibility, Table 3 gives the EOC rates defined in (5.1) for all of the MPEC problems. The results indicate that, as predicted by the theory, the last search direction is a global descent direction in 23 cases. Nonetheless, 20 of these cases still converge at a superlinear rate. By comparison, of the 55 problems for which the last direction is a local descent direction, superlinear convergence occurs in 52 cases.

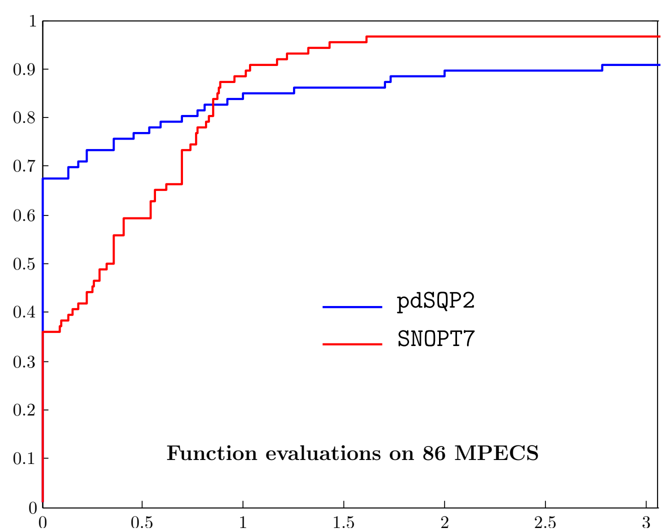

Figure 2: Performance profile for the number of function evaluations on the set of MPEC problems.

Table 3: The estimated order of convergence for pdSQP2 on the MPEC test set.

	Last is global	Last is local	Last two are local	Total
$1.25 < \text{EOC}$	18 (9)	17 (17)	31 (31)	66 (57)
$1.1 < \text{EOC} \le 1.25$	2 (2)	1 (1)	3 (3)	6 (6)
$\text{EOC} \le 1.1$	3 (2)	2 (2)	1 (1)	6 (5)

5.3.4. Degenerate problems from the DEGEN test set

In a series of numerical tests, Izmailov and Solodov [22, 23] demonstrate that Newton-like algorithms such as SQP or inexact SQP methods tend to generate dual iterates that converge to critical multipliers, when they exist. Critical multipliers are those multipliers $y \in \mathcal{Y}(x^*)$ for which the regularized KKT matrix (2.6) is singular at x^* (cf. (2.28)). This is significant because dual convergence to critical multipliers will result in a linear rate of convergence [23]. However, Izmailov [23] shows that an implementation of a conventional stabilized SQP algorithm is less likely to exhibit this behavior, although poor performance can still occur in a small number of cases. This has motivated the use of stabilized SQP subproblems as a way of accelerating local convergence in the presence of critical multipliers. However, such algorithms have had mixed results in practice (see, e.g., Izmailov [25]). The purpose of this section is to use a subset of the DEGEN test set to investigate the performance of pdSQP2 on problems with critical multipliers. The subset of problems consists of those considered by Izmailov [22], and Izmailov and Solodov [23].

The estimated order of convergence (EOC) (cf. (5.1)) for these problems are given in Table 4. The results are separated based on the following properties: (i) no critical multipliers exist; (ii) critical multipliers exist but the limit point y^* is not critical; and (iii) the limit point y^* is critical. The problem summaries indicate which optimal multipliers (if any) are critical. If the final multiplier estimate is within 10^{-3} of a critical multiplier, the multiplier is designated as critical.

Table 4: The estimated order of convergence of algorithm pdSQP2 on the DEGEN test set.

	$\text{EOC} > 1.25$	$1.25 \ge \text{EOC} > 1.1$	$\text{EOC} \le 1.1$
No critical multipliers exist	36	9	6
Critical multipliers exist but y^* is not critical	9	1	2
The limit point y^* is critical	6	29	11

6. Conclusions

This paper considers the formulation, analysis and numerical performance of a stabilized SQP method introduced by Gill, Kungurtsev and Robinson [13]. The method appears to constitute the first algorithm with provable convergence to second-order points as well as a superlinear rate of convergence. The method is formulated as a regularized SQP method with an implicit safeguarding strategy based on minimizing a bound-constrained primal-dual augmented Lagrangian. The method involves a flexible line search along a direction formed from an approximate solution of a regularized quadratic programming subproblem and, when one exists, a direction of negative curvature for the primal-dual augmented Lagrangian. With an appropriate choice of termination condition, the method terminates in a finite number of iterations under weak assumptions on the problem functions. Safeguarding becomes relevant only when the iterates are converging to an infeasible stationary point of the norm of the constraint violations. Otherwise, the method terminates with a point that either satisfies the second-order necessary conditions for optimality, or fails to satisfy a weak second-order constraint qualification. In the former case, superlinear local convergence is established by using an approximate solution of the stabilized QP subproblem that guarantees that the optimal active set, once correctly identified, remains active regardless of the presence of weakly active multipliers. It is shown that the method has superlinear local convergence under the assumption that limit points become close to a solution set containing multipliers satisfying the second-order sufficient conditions for optimality. This rate of convergence is obtained without the need to solve an indefinite QP subproblem, or impose restrictions on which local minimizer of the QP is found. For example, it is not necessary to compute the QP solution closest to the current solution estimate.

Numerical results on a variety of problems indicate that the method performs relatively well compared to a state-of-the-art SQP method. Superlinear convergence is typical, even for problems that do not satisfy standard constraint qualifications. Results are more mixed for problems that do not satisfy the property of strict complementarity.

The proposed method is based on the beneficial properties of dual regularization, which implies that it is necessary to assume a second-order sufficient condition that rules out the possibility of critical multipliers at the solution. Future research will focus on the development of *primal* regularization techniques that allow superlinear convergence when critical multipliers are present. For a local algorithm framework based on primal regularization, see Facchinei, Fischer and Herrich [6, 7].

References

[1] B. T. Baumrucker, J. G. Renfro, and L. T. Biegler. MPEC problem formulations in chemical engineering applications. *preprint, June*, 2007. 38

[2] B. Borchers and J. E. Mitchell. An improved branch and bound algorithm for mixed integer nonlinear programs. *Computers & Operations Research*, 21(4):359–367, 1994. 16

[3] F. E. Curtis and J. Nocedal. Flexible penalty functions for nonlinear constrained optimization. *IMA J. Numer. Anal.*, 28(4):749–769, 2008. 11

[4] S. Dirkse. MPECLib: A collection of Mathematical Programs with Equilibrium Constraints. http://www.gamsworld.eu/mpec/mpeclib.htm, 2004. 38

[5] E. D. Dolan and J. J. Moré. Benchmarking optimization software with COPS. Technical Memorandum ANL/MCS-TM-246, Argonne National Laboratory, Argonne, IL, 2000. 35

[6] F. Facchinei, A. Fischer, and M. Herrich. A family of Newton methods for nonsmooth constrained systems with nonisolated solutions. *Mathematical Methods of Operations Research*, 77(3):433–443, 2013. 40

[7] F. Facchinei, A. Fischer, and M. Herrich. An LP-Newton method: nonsmooth equations, KKT systems, and nonisolated solutions. *Math. Program.*, pages 1–36, 2013. 40

[8] F. Facchinei, A. Fischer, and C. Kanzow. On the accurate identification of active constraints. *SIAM Journal on Optimization*, 9(1):14–32, 1998. 19

[9] D. Fernández and M. Solodov. Stabilized sequential quadratic programming for optimization and a stabilized Newton-type method for variational problems. *Math. Program. Ser. A*, 125:47–73, 2010. 3

[10] C. A. Floudas. *Nonlinear and mixed-integer optimization: fundamentals and applications*. Oxford University Press, 1995. 16

[11] A. Forsgren. Inertia-controlling factorizations for optimization algorithms. *Appl. Num. Math.*, 43:91–107, 2002. 35

[12] A. Forsgren and P. E. Gill. Primal-dual interior methods for nonconvex nonlinear programming. *SIAM J. Optim.*, 8:1132–1152, 1998. 35

[13] P. E. Gill, V. Kungurtsev, and D. P. Robinson. A globally convergent stabilized SQP method: Convergence to second-order optimal points. Center for Computational Mathematics Report Report CCoM 13-04, University of California, San Diego, 2013. 2, 6, 11, 13, 15, 16, 19, 27, 39

[14] P. E. Gill, W. Murray, and M. A. Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Rev.*, 47:99–131, 2005. 35

[15] P. E. Gill and D. P. Robinson. A primal-dual augmented Lagrangian. *Computational Optimization and Applications*, 51:1–25, 2012. 2, 5, 21, 37

[16] P. E. Gill and D. P. Robinson. A globally convergent stabilized SQP method. *SIAM J. Optim.*, 23(4):1983–2010, 2013. 6, 8, 19, 35

[17] P. E. Gill and E. Wong. A regularized method for convex and general quadratic programming. Center for Computational Mathematics Report CCoM 13-3, Department of Mathematics, University of California, San Diego, La Jolla, CA, 2013. 35

[18] N. I. M. Gould, D. Orban, and P. L. Toint. CUTEst: a constrained and unconstrained testing environment with safe threads. Technical report, Rutherford Appleton Laboratory, Chilton, England, 2013. 35, 36

[19] W. W. Hager. Stabilized sequential quadratic programming. *Comput. Optim. Appl.*, 12(1-3):253–273, 1999. Computational optimization—a tribute to Olvi Mangasarian, Part I. 3, 19

[20] W. Hock and K. Schittkowski. *Test Examples for Nonlinear Programming Codes*. Lecture Notes in Economics and Mathematical Systems 187. Springer Verlag, Berlin, Heidelberg and New York, 1981. 35

[21] A. F. Izmailov. Solution sensitivity for Karush-Kuhn-Tucker systems with non-unique Lagrange multipliers. *Optimization: A Journal of Mathematical Programming and Operations Research*, 59(5), 2010. 22, 24

[22] A. F. Izmailov and M. V. Solodov. Examples of dual behaviour of Newton-type methods on optimization problems with degenerate constraints. *Computational Optimization and Applications*, 42:231–264, 2009. 10.1007/s10589-007-9074-4. 39

[23] A. F. Izmailov and M. V. Solodov. On attraction of linearly constrained Lagrangian methods and of stabilized and quasi-Newton SQP methods to critical multipliers. *Math. Program.*, 126(2, Ser. A):231–257, 2011. 3, 39

[24] A. F. Izmailov and M. V. Solodov. Stabilized SQP revisited. *Math. Program.*, 133:93–120, 2012. 22, 34

[25] A. F. Izmailov and E. I. Uskov. On the influence of the critical Lagrange multipliers on the convergence rate of the multiplier method. *Computational Mathematics and Mathematical Physics*, 52(11):1504–1519, 2012. 39

[26] V. Kungurtsev. *Second-Derivative Sequential Quadratic Programming Methods for Nonlinear Optimization*. PhD thesis, Department of Mathematics, University of California San Diego, La Jolla, CA, 2013. 3, 35, 37

[27] S. Leyffer. Integrating SQP and branch-and-bound for mixed integer nonlinear programming. *Comput. Optim. Appl.*, 18(3):295–309, 2001. 16

[28] MathWorks Inc., Natick, Massachusetts. *Matlab User's Guide*, 1992. 35

[29] E.-S. Mostafa, L. Vicente, and S. Wright. Numerical behavior of a stabilized SQP method for degenerate NLP problems. In C. Bliek, C. Jermann, and A. Neumaier, editors, *Global Optimization and Constraint Satisfaction*, volume 2861 of *Lecture Notes in Computer Science*, pages 123–141. Springer Berlin / Heidelberg, 2003. 37, 38

[30] C. Oberlin and S. J. Wright. Active set identification in nonlinear programming. *SIAM J. Optim.*, 17(2):577–605, 2006. 3

[31] D. P. Robinson. *Primal-Dual Methods for Nonlinear Optimization*. PhD thesis, Department of Mathematics, University of California San Diego, La Jolla, CA, 2007. 2

[32] S. J. Wright. Superlinear convergence of a stabilized SQP method to a degenerate solution. *Comput. Optim. Appl.*, 11(3):253–275, 1998. 3

[33] S. J. Wright. Modifying SQP for degenerate problems. *SIAM J. Optim.*, 13(2):470–497, 2002. 3

[34] S. J. Wright. Constraint identification and algorithm stabilization for degenerate nonlinear programs. *Math. Program.*, 95(1, Ser. B):137–160, 2003. ISMP 2000, Part 3 (Atlanta, GA). 37

[35] S. J. Wright. An algorithm for degenerate nonlinear programming with rapid local convergence. *SIAM J. Optim.*, 15(3):673–696, 2005. 3, 18, 19, 22, 43, 44, 45

A. Properties of Problem Perturbations

Several of the theorems discussed in Section 4 involve the relationship between the proximity measure $r(x,y)$, and the quantities $\eta(x,y)$ and $\bar{\eta}(x,y)$ defined by Wright [35] (and also defined below). Throughout the discussion, the scaled closed interval $[\alpha\,\alpha_\ell, \alpha\,\alpha_u]$ defined in terms of the positive scalars α_ℓ, α_u and scale factor α, will denoted by $[\alpha_\ell, \alpha_u]\cdot\alpha$.

A.1. Inequality-constraint form

The original results apply to an optimization problem with all inequality constraints. The all-inequality form of problem (NP) is

$$\begin{array}{ll} \underset{x\in\mathbb{R}^n}{\text{minimize}} & f(x) \\ \text{subject to} & c(x) \ge 0, \quad -c(x) \ge 0, \quad x \ge 0. \end{array} \tag{A.1}$$

Given multipliers y for problem (NP), the multipliers for the nonnegativity constraints $x \ge 0$ are $g(x) - J(x)^T y$ and are denoted by $z(x,y)$.

Consider the primal-dual solution set $\mathcal{V}_z(x^*)$ for problem (A.1). It follows that $\mathcal{V}_z(x^*) = \mathcal{V}(x^*) \times \mathcal{Z}(x^*)$, where

$$\mathcal{V}(x^*) = \{x^*\} \times \Lambda(x^*) \;\text{ and }\; \mathcal{Z}(x^*) = \{z : g(x^*) - J(x^*)^T y, \;\text{ for some }\; y \in \Lambda(x^*)\}$$

The distance to optimality for the problem (A.1) is

$$\begin{aligned} \operatorname{dist}\big((x,y,z),\mathcal{V}_z(x^*)\big) &= \min_{(\bar{x},\bar{y},\bar{z})\in\mathcal{V}_z(x^*)} \|(x-\bar{x}, y-\bar{y}, z-\bar{z})\| \\ &= \min_{(\bar{x},\bar{y})\in\mathcal{V}(x^*)} \|(x-\bar{x}, y-\bar{y}, z(x,y) - (g(\bar{x}) - J(\bar{x})^T\bar{y}))\|. \end{aligned}$$

Result A.1. *If* $\operatorname{dist}\big((x,y,z),\mathcal{V}_z(x^*)\big)$ *denotes the distance to optimality for the problem* (NP) *written in all-inequality form, then* $\delta(x,y) = \Theta\big(\operatorname{dist}((x,y,z),\mathcal{V}_z(x^*))\big)$.

Proof. Let $y_P^*(y)$ denote the vector that minimizes the distance from y to the compact set $\Lambda(x^*)$ (see (4.4)). Consider the quantity

$$\underline{\delta}(x,y) = \big\|\big(x - x^*,\; y - y_P^*(y),\; z(x,y) - z(x^*, y_P^*(y))\big)\big\|\,.$$

The components of the vector $(x - x^*,\; y - y_P^*(y))$ used to define $\delta(x,y)$ form the first $n+m$ components of $\underline{\delta}(x,y)$, which implies that $\delta(x,y) \le \underline{\delta}(x,y)$. For the upper bound, the Lipschitz continuity of J and g, together with the boundedness of $y_P^*(y)$ and $J(x)$ imply that

$$\begin{aligned} \|z(x,y) - z(x^*, y_P^*(y))\| &= \|g(x) - J(x)^T y - g(x^*) + J(x^*)^T y_P^*(y)\| \\ &\le L_g\|x - x^*\| + \|J(x)^T(y - y_P^*(y))\| + \|(J(x) - J(x^*))^T y_P^*(y)\| \\ &\le L_g\|x - x^*\| + C_J\|y - y_P^*(y)\| + L_2 C_y\|x - x^*\| \\ &\le C_a\delta(x,y). \end{aligned} \tag{A.2}$$

It follows that $\underline{\delta}(x,y) \le \delta(x,y) + \|z(x,y) - z(x^*, y_P^*(y))\| \le (1 + C_a)\delta(x,y)$, which implies that $\underline{\delta}(x,y) = \Theta\big(\delta(x,y)\big)$, and, equivalently, $\delta(x,y) = \Theta\big(\underline{\delta}(x,y)\big)$.

The proof is complete if it can be shown that $\delta(x,y) = \Theta\big(\,\mathrm{dist}((x,y,z),\mathcal{V}_z(x^*))\big)$. The definitions of $\mathrm{dist}((x,y,z),\mathcal{V}_z)$ and $\underline{\delta}(x,y)$ imply that $\mathrm{dist}((x,y,z),\mathcal{V}_z(x^*)) \le \underline{\delta}(x,y)$. Moreover, $\delta(x,y) = \mathrm{dist}((x,y),\mathcal{V}(x^*)) \le \mathrm{dist}((x,y,z),\mathcal{V}_z(x^*))$ because there is no third component in the definition of $\delta(x,y)$. As $\delta(x,y) = \Theta\big(\underline{\delta}(x,y)\big)$, it must hold that $\delta(x,y) = \Theta\big(\,\mathrm{dist}((x,y,z),\mathcal{V}_z(x^*))\big)$, as required. ∎

Let $\bar{\eta}(x,y,z_{\mathcal{A}})$ be the practical estimate of $\mathrm{dist}\big((x,y,z),\mathcal{V}_z(x^*)\big)$ given by

$$\eta(x,y) = \|(v_1,v_2,v_3,v_4)\|_1,$$

where $v_1 = g(x) - J(x)^T y - z(x,y)$, $v_2 = \min(x, z(x,y))$, $v_3 = \min(c(x), \max(y,0))$, and $v_4 = \min(-c(x), \max(-y,0))$. Wright [35] shows that

$$\eta(x,y) \in [1/\kappa, \kappa] \cdot \mathrm{dist}\big((x,y,z),\mathcal{V}_z(x^*)\big)$$

for all (x,y) sufficiently close to (x^*,y^*).

Result A.2. *Consider the function $\eta(x,y) = \|(v_1,v_2,v_3,v_4)\|_1$, where $v_1 = g(x) - J(x)^T y - z(x,y)$, $v_2 = \min(x, z(x,y))$, $v_3 = \min(c(x), \max(y,0))$, and $v_4 = \min(-c(x), \max(-y,0))$. The quantity $\eta(x,y)$ defines a measure of the quality of (x,y) as an approximate solution of problem* (NP) *defined in all-inequality form and satisfies $r(x,y) = \Theta\big(\eta(x,y)\big)$.*

Proof. It will be established that $\eta(x,y) = \Theta\big(r(x,y)\big)$. The vector v_1 is zero by definition. The vector v_2 is $\min(x, g(x) - J(x)^T y)$, which is the second part of $r(x,y)$.

If $c_i(x) < 0$ and $y_i \ge 0$ then $\min(c_i(x), \max(y_i,0)) = c_i(x)$ and $\min(-c_i(x), \max(-y_i,0)) = 0$. If $c_i(x) < 0$ and $y_i \le 0$ then $\min(c_i(x), \max(y_i,0)) = c_i(x)$ and $\min(-c_i(x), \max(-y_i,0)) = \min(|c_i(x)|, |y_i|)$. If $c_i(x) > 0$ and $y_i \ge 0$ then $\min(c_i(x), \max(y_i,0)) = \min(|c_i(x)|, |y_i|)$ and $\min(-c_i(x), \max(-y_i,0)) = -c_i(x)$. If $c_i(x) > 0$ and $y_i \le 0$ then $\min(c_i(x), \max(y_i,0)) = 0$ and $\min(-c_i(x), \max(-y_i,0)) = -c_i(x)$.

It follows that for every i, one or the other of the vectors v_3 or v_4 has a component equal to $|c_i(x)|$ and hence $\eta(x,y) \ge r(x,y)$. In addition, v_3 or v_4 may have a term that is $\min(|c_i(x)|, |y_i|) \le |c_i(x)|$, and so $\eta(x,y) \le 2r(x,y)$. It follows that $\eta(x,y) = \Theta\big(r(x,y)\big)$, as required. ∎

A.2. Equality-constraint form

Any solution x^* of problem (NP) is also a solution of the problem

$$\underset{x}{\text{minimize}}\; f(x) \;\; \text{subject to} \;\; c(x) = 0, \;\; \text{and} \;\; [\,x\,]_{\mathcal{A}} = E_{\mathcal{A}}^T x = 0. \tag{A.3}$$

Furthermore, any primal-dual solution (x^*,y^*) of problem (NP) must satisfy the SOSC for (A.3) because the conditions for problem (NP) imply that $p^T H(x^*,y^*)p > 0$ for all p such that $J(x^*)p = 0$ and $p_i = 0$ for every $i \in \mathcal{A}(x^*)$. The primal-dual solution set $\mathcal{U}_z(x^*)$ for problem (A.3) has the form $\mathcal{U}_z(x^*) = \mathcal{U}(x^*) \times \mathcal{Z}(x^*)$, where

$$\mathcal{U}(x^*) = \{x^*\} \times \Lambda(x^*) \;\; \text{and} \;\; \mathcal{Z}(x^*) = \{z_{\mathcal{A}} : z_{\mathcal{A}} = [\,g(x^*) - J(x^*)^T y\,]_{\mathcal{A}}, \;\; \text{for some} \;\; y \in \Lambda(x^*)\}.$$

Let y and $z_{\mathcal{A}}$ denote estimates of the multipliers for the constraints $c(x) = 0$ and $E_{\mathcal{A}}^T x = 0$. Let $\bar{\delta}(x, y, z_{\mathcal{A}})$ be the distance of $(x, y, z_{\mathcal{A}})$ to a solution of (A.3), i.e.,

$$\begin{aligned}\operatorname{dist}(x, y, z_{\mathcal{A}}, \mathcal{U}_z(x^*)) &= \min_{(\bar{x}, \bar{y}, \bar{z}_{\mathcal{A}}) \in \mathcal{U}_z(x^*)} \|(x - \bar{x}, y - \bar{y}, z_{\mathcal{A}} - \bar{z}_{\mathcal{A}})\| \\ &= \min_{(\bar{x}, \bar{y}) \in \mathcal{U}(x^*)} \|(x - \bar{x}, y - \bar{y}, [\, g(x) - J(x)^T y - (g(\bar{x}) - J(\bar{x})^T \bar{y})\,]_{\mathcal{A}})\| \\ &= \min_{\bar{y} \in \Lambda(x^*)} \|(x - x^*, y - \bar{y}, [\, g(x) - J(x)^T y - (g(x^*) - J(x^*)^T \bar{y})\,]_{\mathcal{A}})\|,\end{aligned}$$

where $\Lambda(x^*)$ is the compact subset of the set of optimal multipliers corresponding to x^* for problem (NP).

Let $\widetilde{\mu}(x, y, z_{\mathcal{A}})$ be the estimate of $\operatorname{dist}(x, y, z_{\mathcal{A}}, \mathcal{U}_z(x^*))$ given by

$$\widetilde{\mu}(x, y, z_{\mathcal{A}}) = \left\| \begin{pmatrix} g(x) - J(x)^T y - E_{\mathcal{A}} z_{\mathcal{A}} \\ c(x) \\ x_{\mathcal{A}} \end{pmatrix} \right\| = \left\| \begin{pmatrix} [\, g(x) - J(x)^T y\,]_{\mathcal{A}} - z_{\mathcal{A}} \\ [\, g(x) - J(x)^T y\,]_{\mathcal{F}} \\ c(x) \\ x_{\mathcal{A}} \end{pmatrix} \right\|_1 . \tag{A.4}$$

Wright [35] uses the notation $\bar{\eta}(x, y, z_{\mathcal{A}}) = \widetilde{\mu}(x, y, z_{\mathcal{A}})$ and shows that for all (x, y) sufficiently close to (x^*, y^*), the estimate $\widetilde{\mu}(x, y, z_{\mathcal{A}})$ satisfies

$$\widetilde{\mu}(x, y, z_{\mathcal{A}}) \in [1/\kappa, \kappa] \cdot \operatorname{dist}(x, y, z_{\mathcal{A}}, \mathcal{U}_z(x^*)), \tag{A.5}$$

where $\kappa = \kappa(\mathcal{U}_z(x^*))$ is a constant.

Result A.3. *For all (x, y) sufficiently close to (x^*, y^*), the estimate $\widetilde{\mu}(x, y, z_{\mathcal{A}}) = \|(g(x) - J(x)^T y - E_{\mathcal{A}} z_{\mathcal{A}}, c(x), x_{\mathcal{A}})\|_1$ satisfies $\widetilde{\mu}(x, y, z_{\mathcal{A}}) = O(\delta(x, y))$.*

Proof. For all (x, y) sufficiently close to (x^*, y^*), the definition of $\operatorname{dist}(x, y, z_{\mathcal{A}}, \mathcal{U}_z(x^*))$ and the Lipschitz continuity of g and J imply that

$$\begin{aligned}\operatorname{dist}(x, y, z_{\mathcal{A}}, \mathcal{U}_z(x^*)) &\le \delta(x, y) + \|[\, g(x) - J(x)^T y - (g(x^*) - J(x^*)^T y_P^*(y))\,]_{\mathcal{A}}\| \\ &\le \delta(x, y) + \|g(x) - J(x)^T y - (g(x^*) - J(x^*)^T y_P^*(y))\| \\ &\le \delta(x, y) + C_a \delta(x, y),\end{aligned}$$

for some bounded constant C_a (cf. (A.2)). The result now follows from (A.5). ∎