SIAM J. OPTIM.
Vol. 23, No. 4, pp. 1983–2010

A GLOBALLY CONVERGENT STABILIZED SQP METHOD*

PHILIP E. GILL† AND DANIEL P. ROBINSON‡

Abstract. Sequential quadratic programming (SQP) methods are a popular class of methods for nonlinearly constrained optimization. They are particularly effective for solving a sequence of related problems, such as those arising in mixed-integer nonlinear programming and the optimization of functions subject to differential equation constraints. Recently, there has been considerable interest in the formulation of *stabilized* SQP methods, which are specifically designed to handle degenerate optimization problems. Existing stabilized SQP methods are essentially local in the sense that both the formulation and analysis focus on the properties of the methods in a neighborhood of a solution. A new SQP method is proposed that has favorable global convergence properties yet, under suitable assumptions, is equivalent to a variant of the conventional stabilized SQP method in the neighborhood of a solution. The method combines a primal-dual generalized augmented Lagrangian function with a flexible line search to obtain a sequence of improving estimates of the solution. The method incorporates a convexification algorithm that allows the use of exact second derivatives to define a convex quadratic programming (QP) subproblem without requiring that the Hessian of the Lagrangian be positive definite in the neighborhood of a solution. This gives the potential for fast convergence in the neighborhood of a solution. Additional benefits of the method are that each QP subproblem is regularized and the QP subproblem always has a known feasible point. Numerical experiments are presented for a subset of the problems from the `CUTEr` test collection.

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

DOI. 10.1137/120882913

1. Introduction.

This paper is concerned with methods for the solution of optimization problems of the form

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

where $c : \mathbb{R}^n \mapsto \mathbb{R}^m$ and $f : \mathbb{R}^n \mapsto \mathbb{R}$ are twice-continuously differentiable. This problem formulation 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 $l \leq x \leq u$ (see, e.g., section 6).

Some of the most efficient algorithms for nonlinear optimization are sequential quadratic programming (SQP) methods (for a survey, see, e.g., [1, 25]). Conventional SQP methods find an approximate solution of a sequence of quadratic programming (QP) subproblems in which a quadratic model of the objective function is minimized subject to the linearized constraints. Convergence from any starting point is enforced by requiring the improvement in some merit function at each step. The merit function is usually a penalty or augmented Lagrangian function that defines some compromise

*Received by the editors June 29, 2012; accepted for publication (in revised form) July 12, 2013; published electronically October 15, 2013.

http://www.siam.org/journals/siopt/23-4/88291.html

†Department of Mathematics, University of California, San Diego, La Jolla, CA 92093-0112 (pgill@ucsd.edu). This author's research was supported in part by National Science Foundation grants DMS-0915220 and DMS-1318480 and by Department of Energy grant DE-SC0002349.

‡Department of Applied Mathematics and Statistics, Johns Hopkins University, Baltimore, MD 21218-2682 (daniel.p.robinson@jhu.edu). This author's research was supported in part by National Science Foundation grant DMS-1217153.

between reducing the objective function and satisfying the constraints. SQP methods that solve the QP subproblem using an active-set method are able to capitalize on a good initial starting point, which makes them particularly effective for solving a sequence of related problems, such as those arising in the optimization of functions subject to differential equation constraints.

SQP methods have an inner/outer iteration structure, with the work for an inner iteration being dominated by the cost of solving a system of symmetric indefinite linear equations involving a subset of the variables and constraints. Some of the most successful methods use sophisticated matrix factorization updating techniques that exploit the fact that the linear equations change by only a single row and column at each inner iteration. These updating techniques are often customized for the particular QP method being used and have the benefit of providing a uniform treatment of ill-conditioning and singularity. On the negative side, it is difficult to implement conventional SQP methods so that the second derivatives of f and c may be used efficiently and reliably. Some of these difficulties stem from the theoretical properties of the QP subproblem, which can be nonconvex when second derivatives are used. Nonconvex QP is NP-hard—even for the calculation of a local minimizer [11, 19]. The complexity of the QP subproblem has been a major impediment to the formulation of conventional second-derivative SQP methods (although methods based on indefinite QP have been proposed [15, 16]). Over the years, algorithm developers have avoided this difficulty by using a convex QP defined in terms of a positive semidefinite approximate Hessian. In some cases, this QP is used to define the search direction directly [21, 22, 33, 46, 48, 49]; in others, the QP is used to identify the constraints for an equality-constrained subproblem that uses second derivatives [30, 31, 42].

Recently, there has been considerable interest in the formulation of *stabilized* SQP methods, which are specifically designed to improve the convergence rate for degenerate problems [14, 32, 36, 44, 51, 52]. Existing stabilized SQP methods are essentially local in the sense that both the formulation and analysis focus on the properties of the methods in a neighborhood of a solution. In parallel with the development of stabilized SQP methods, *regularized* methods have been proposed that reduce the dependency on custom-built matrix factorization and updating methods for solving the QP subproblem (see, e.g., [25, 40]). Regularized methods use a so-called regularization parameter to define linear equations that are always nonsingular. This feature obviates specialized software to detect rank deficiency and allows the application of third-party linear system solvers.

A seemingly different approach from tackling problem (NP) directly is to replace the constrained problem by a sequence of bound-constrained problems in which the equality constraints are included in an augmented Lagrangian objective function [4, 6, 8, 9, 10, 34, 45]. These methods have strong global convergence properties that require relatively weak assumptions on the problem.

In this paper we formulate and analyze a new SQP method that effectively combines the use of a bound-constrained augmented Lagrangian function with the three elements of conventional, regularized, and stabilized SQP. In particular, the method has favorable global convergence properties, yet, under suitable assumptions, the method is equivalent to a conventional stabilized SQP method in the neighborhood of a solution. The method pairs the primal-dual generalized augmented Lagrangian merit function defined in [23] with a flexible line search to obtain a sequence of improving estimates of the solution. A crucial feature of the method is that the QP subproblem is based on the exact Hessian of the Lagrangian, yet has a unique bounded solution. This gives the potential for fast convergence in the neighborhood of a solution.

Additional benefits of the method include the following: (i) each QP subproblem is regularized; and (ii) the QP subproblem always has a known feasible point.

The paper is organized in seven sections. Section 1.2 is a review of some of the basic properties of SQP methods. In section 2, the steps of the proposed primal-dual SQP method are defined. Similarities with the conventional Hestenes–Powell augmented Lagrangian method are discussed and equivalence to stabilized SQP under certain assumptions is described. Global convergence results are established in section 3. In section 4, a "convexification procedure" is proposed for obtaining a QP subproblem with a bounded unique solution. In section 5, the connection between the QP subproblems associated with the bound-constrained augmented Lagrangian and stabilized SQP is extended to show that a conventional active-set method generates identical QP iterates on both problems. Finally, in section 6 some numerical experiments are presented for a simple MATLAB implementation applied to a selection of problems from the CUTEr test collection [29].

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 $\mathrm{In}(A)$, is the integer triple (a_+, a_-, a_0) giving the number of positive, negative, and zero eigenvalues of 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 long 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., $[\,v_{\scriptscriptstyle F}\,]_i$ is the ith component of the vector $v_{\scriptscriptstyle F}$. The subvector of components with indices in the index set $\mathcal{S}$ is denoted by $[\,\cdot\,]_{\mathcal{S}}$; e.g., $[\,v\,]_{\mathcal{S}}$ is the vector with components v_i for $i \in \mathcal{S}$. A local solution of an optimization problem 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 $H(x, y) = \nabla^2 f(x) - \sum_{i=1}^m y_i \nabla^2 c_i(x)$.

1.2. Background. The vector-pair (x^*, y^*) is called a first-order solution to problem (NP) if it satisfies

$$c(x^*) = 0 \ \text{ and } \ \min\left(x^*, g(x^*) - J(x^*)^T y^*\right) = 0, \tag{1.1}$$

where y^* are the Lagrange multipliers associated with the constraints $c(x) = 0$.

Given an estimate (x_k, y_k) of a primal-dual solution of (NP), a line-search SQP method computes a search direction $p_k = \widehat{x}_k - x_k$ such that $\widehat{x}_k$ is the solution (when it exists) of the quadratic program

$$\begin{array}{ll} \underset{x}{\text{minimize}} & g_k^T(x - x_k) + \frac{1}{2}(x - x_k)^T \widehat{H}_k (x - x_k) \\ \text{subject to} & c_k + J_k(x - x_k) = 0, \quad x \geq 0, \end{array} \tag{1.2}$$

where c_k, g_k, and J_k denote the quantities $c(x)$, $g(x)$, and $J(x)$ evaluated at x_k, and $\widehat{H}_k$ is some approximation of the Lagrangian Hessian $H(x_k, y_k)$. For the moment, it is assumed that the approximate Hessian $\widehat{H}_k$ is positive definite, in which case

the QP subproblem (1.2) is convex. (This assumption is not required for the method proposed in section 2.) If the Lagrange multiplier vector associated with the constraint $c_k + J_k(x - x_k) = 0$ is given by $\widehat{y}_k$, then a solution $(\widehat{x}_k, \widehat{y}_k)$ of the QP subproblem (1.2) satisfies the optimality conditions

$$c_k + J_k(\widehat{x}_k - x_k) = 0 \quad \text{and} \quad \min\left(\widehat{x}_k,\, g_k + \widehat{H}_k(\widehat{x}_k - x_k) - J_k^T\widehat{y}_k\right) = 0,$$

which are analogous to (1.1). Given any $x \geq 0$, we define the index sets

$$\mathcal{A}(x) = \{i : x_i = 0\} \quad \text{and} \quad \mathcal{F}(x) = \{1, 2, \ldots, n\} \setminus \mathcal{A}(x). \tag{1.3}$$

If x is feasible for the constraints $c_k + J_k(x - x_k) = 0$, then $\mathcal{A}(x)$ is the index set of active variables at x (the *active set* at x) and $\mathcal{F}(x)$ is the index set of free variables at x. If the set $\mathcal{A}(\widehat{x}_k)$ associated with a solution of the subproblem (1.2) is known, then $\widehat{x}_k$ may be found by solving linear equations that represent the optimality conditions for an equality-constrained QP with the inequalities $x \geq 0$ replaced by $x_i = 0$ for $i \in \mathcal{A}(\widehat{x}_k)$. In general, the optimal active set $\mathcal{A}(\widehat{x}_k)$ is not known in advance, and active-set QP methods generate a sequence of feasible estimates $(\widehat{x}^{(j)}, \widehat{y}^{(j)})$ of $(\widehat{x}_k, \widehat{y}_k)$ such that $(\widehat{x}^{(j+1)}, \widehat{y}^{(j+1)}) = (\widehat{x}^{(j)}, \widehat{y}^{(j)}) + \alpha^{(j)}(p^{(j)}, q^{(j)})$, with

$$p_{\scriptscriptstyle A}^{(j)} = 0 \ \text{ and } \ \begin{pmatrix} \widehat{H}_{\scriptscriptstyle F} & J_{\scriptscriptstyle F}^T \\ J_{\scriptscriptstyle F} & 0 \end{pmatrix} \begin{pmatrix} p_{\scriptscriptstyle F}^{(j)} \\ -q^{(j)} \end{pmatrix} = -\begin{pmatrix} [\, g_k + \widehat{H}_k(\widehat{x}^{(j)} - x_k) - J_k^T\widehat{y}^{(j)} \,]_{\scriptscriptstyle F} \\ c_k + J_k(\widehat{x}^{(j)} - x_k) \end{pmatrix}, \tag{1.4}$$

where the quantities with subscripts "F" and "A" are defined in terms of the index sets $\mathcal{F}(\widehat{x}^{(j)})$ and $\mathcal{A}(\widehat{x}^{(j)})$; i.e., $\widehat{H}_{\scriptscriptstyle F}$ is the matrix of free rows and columns of $\widehat{H}_k$, $J_{\scriptscriptstyle F}$ is the matrix of free columns of J_k, and $p_{\scriptscriptstyle F}^{(j)}$ and $p_{\scriptscriptstyle A}^{(j)}$ are the free and fixed components of $p^{(j)}$. The step length $\alpha^{(j)}$ is chosen to ensure the feasibility of every element of $\widehat{x}^{(j+1)}$.

If equations (1.4) are to be used to define $p^{(j)}$ and $q^{(j)}$, then it is necessary that $J_{\scriptscriptstyle F}$ has full rank, which is a crucial issue in the formulation of reliable methods. Two remedies are available.

- *Rank-enforcing active-set methods* maintain a set of indices $\mathcal{B}$ associated with a matrix of columns $J_{\scriptscriptstyle B}$ with rank m; i.e., the rows of $J_{\scriptscriptstyle B}$ are linearly independent. The set $\mathcal{B}$ is the complement in $\{1, 2, \ldots, n\}$ of a "working set" $\mathcal{N}$ of indices that estimates the set $\mathcal{A}$ at a solution $\widehat{x}_k$ of (1.2). If $\mathcal{N}$ is a subset of $\mathcal{A}$, then a system analogous to (1.4) may be solved with $\mathcal{F}$ replaced by $\mathcal{B}$. The system is nonsingular because of the linear independence of the rows of $J_{\scriptscriptstyle B}$.
- *Regularized active-set methods* include a nonzero diagonal regularization term in the $(2,2)$ block of (1.4). The magnitude of the regularization term is generally based on heuristic arguments.

Another important attribute associated with SQP methods is the rate of convergence to a local minimizer. If the active set at the solution $\widehat{x}_k$ of the QP subproblem (1.2) is the same as the active set at x_k, i.e., $\mathcal{A}(\widehat{x}_k) = \mathcal{A}(x_k)$ and $\mathcal{F}(\widehat{x}_k) = \mathcal{F}(x_k)$, then the QP solution $(\widehat{x}_k, \widehat{y}_k)$ is obtained in a *single* iteration. Specifically, the QP solution is given by $(\widehat{x}_k, \widehat{y}_k) = (x_k, y_k) + (p^{(0)}, q^{(0)})$, with

$$p_{\scriptscriptstyle A}^{(0)} = 0 \quad \text{and} \quad \begin{pmatrix} \widehat{H}_{\scriptscriptstyle F} & J_{\scriptscriptstyle F}^T \\ J_{\scriptscriptstyle F} & 0 \end{pmatrix} \begin{pmatrix} p_{\scriptscriptstyle F}^{(0)} \\ -q^{(0)} \end{pmatrix} = -\begin{pmatrix} [\, g_k - J_k^T y_k \,]_{\scriptscriptstyle F} \\ c_k \end{pmatrix}, \tag{1.5}$$

where $\widehat{H}_{\scriptscriptstyle F}$, $J_{\scriptscriptstyle F}$, $p_{\scriptscriptstyle F}^{(0)}$, and $p_{\scriptscriptstyle A}^{(0)}$ are as defined above, but with the index sets $\mathcal{F}(\widehat{x}_k)$ $(= \mathcal{F}(x_k))$ and $\mathcal{A}(\widehat{x}_k)$ $(= \mathcal{A}(x_k))$. The conditions (1.5) represent the Newton equations

for finding a stationary point of the equality-constrained problem defined in terms of the free variables. If $\widehat{H}_k = H(x_k, y_k)$ and J_F has full rank in a neighborhood of a solution, then Newton's method converges at a quadratic rate. However, if J_F does not have full rank, equations (1.5) are singular with no unique solution. In this case, one remedy is to use a *stabilized SQP method* in which the QP subproblem (1.2) is replaced by

$$\begin{array}{ll} \underset{x,y}{\text{minimize}} & g_k^T(x - x_k) + \frac{1}{2}(x - x_k)^T\widehat{H}_k(x - x_k) + \frac{1}{2}\mu_k\|y\|^2 \\ \text{subject to} & c_k + J_k(x - x_k) + \mu_k(y - y_k) = 0, \quad x \geq 0, \end{array} \tag{1.6}$$

where $\{\mu_k\}$ is a positive sequence such that $\mu_k \to 0$ as $x_k \to x^*$ (see, e.g., Wright [51], Hager [32], Li and Qi [39], and Fernández and Solodov [14]). The QP (1.6) is often referred to as a stabilized subproblem because of its calming effect on multiplier estimates for degenerate problems (see, e.g., [32, 51]). Under certain assumptions, stabilized SQP methods exhibit fast local convergence. However, there is no guarantee of convergence to a local solution for an arbitrary starting point. Under suitable assumptions, the method proposed in this paper is guaranteed to be globally convergent and is equivalent to stabilized SQP in the limit. Result 2.1 of section 2.1 below describes the precise relationship between the QP subproblems of the proposed method and stabilized SQP.

2. A regularized primal-dual line-search SQP algorithm. This section defines a regularized SQP line-search method based on the primal-dual augmented Lagrangian function

$$M^\nu(x, y\,; y^E, \mu) = f(x) - c(x)^T y^E + \frac{1}{2\mu}\|c(x)\|^2 + \frac{\nu}{2\mu}\|c(x) + \mu(y - y^E)\|^2, \tag{2.1}$$

where ν is a scalar, μ is the penalty parameter, and y^E is an estimate of an optimal Lagrange multiplier vector y^*. (A trust-region-based method could also be given, but we focus on line-search methods in this paper.) The function (2.1), proposed by Robinson [47] and Gill and Robinson [23], may be derived by applying the primal-dual penalty function of Forsgren and Gill [18] to a problem in which the constraints are shifted by a constant vector (see Powell [45]). With the notation $c = c(x)$, $g = g(x)$, and $J = J(x)$, the gradient of $M^\nu(x, y\,; y^E, \mu)$ may be written as

$$\nabla M^\nu(x, y\,; y^E, \mu) = \begin{pmatrix} g - J^T\big((1+\nu)(y^E - \frac{1}{\mu}c) - \nu y\big) \\ \nu\big(c + \mu(y - y^E)\big) \end{pmatrix} \tag{2.2a}$$

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

where $\pi = \pi(x\,; y^E, \mu)$ denotes the vector-valued function

$$\pi(x\,; y^E, \mu) = y^E - \frac{1}{\mu}c(x). \tag{2.3}$$

Similarly, the Hessian of $M^\nu(x, y\,; y^E, \mu)$ may be written as

$$\nabla^2 M^\nu(x, y\,; y^E, \mu) = \begin{pmatrix} H\big(x, \pi + \nu(\pi - y)\big) + \frac{1}{\mu}(1+\nu)J^TJ & \nu J^T \\ \nu J & \nu\mu I \end{pmatrix}. \tag{2.4}$$

Our approach is motivated by the following theorem, which shows that under certain assumptions, minimizers of problem (NP) are also minimizers of the bound-constrained problem

$$\underset{x,y}{\text{minimize}} \;\; M^\nu(x,y\,;y^E,\mu) \quad \text{subject to} \quad x \geq 0. \tag{2.5}$$

THEOREM 2.1 (Robinson [47, Theorem 4.6.1]). *If (x^*,y^*) satisfies second-order sufficient conditions for a solution of problem* (NP)*, then there exists a positive $\bar{\mu}$ such that, for all $0<\mu<\bar{\mu}$, $\nu>0$, and $y^E=y^*$, the point (x^*,y^*) is a minimizer of problem* (2.5).

The reader may refer to Robinson [47] and Gill and Robinson [23] for additional details. In this paper, however, Theorem 2.1 is used as motivation for the algorithm described below.

2.1. Definition of the primal-dual search direction. Given the kth iterate $v_k=(x_k,y_k)$, a Lagrange multiplier estimate y_k^E, and a positive regularization parameter μ_k^R, a symmetric matrix $\widehat{H}(x_k,y_k)\approx H(x_k,y_k)$ is defined such that $\widehat{H}(x_k,y_k)+(1/\mu_k^R)J(x_k)^TJ(x_k)$ is positive definite. One may choose $\widehat{H}(x_k,y_k)$ itself to be positive definite, but we explore a more sophisticated strategy in section 4 that allows for an *indefinite* matrix $\widehat{H}(x_k,y_k)$ that more faithfully approximates $H(x_k,y_k)$. With this assumption on the matrix $\widehat{H}$, part (i) of Lemma 2.2 given below may be applied with the quantities $H=\widehat{H}(x_k,y_k)$, $J=J(x_k)$, and $\mu=\mu_k^R$, to infer that the matrix

$$B^\nu(x_k,y_k\,;\mu_k^R)=\begin{pmatrix}\widehat{H}(x_k,y_k)+\frac{1}{\mu_k^R}(1+\nu)J(x_k)^TJ(x_k) & \nu J(x_k)^T\\ \nu J(x_k) & \nu\mu_k^R I\end{pmatrix} \tag{2.6}$$

is a positive semidefinite approximation to the Hessian of M^ν. Given an appropriate matrix $B^\nu(v_k\,;\mu_k^R)\equiv B^\nu(x_k,y_k\,;\mu_k^R)$, the primal-dual search direction is given by

$$d_k=\widehat{v}_k-v_k, \tag{2.7}$$

where $\widehat{v}_k=(\widehat{x}_k,\widehat{y}_k)$ is a solution of the convex bound-constrained QP subproblem:

$$\begin{array}{ll}\underset{v}{\text{minimize}} & \nabla M^\nu(v_k\,;y_k^E,\mu_k^R)^T(v-v_k)+\frac{1}{2}(v-v_k)^TB^\nu(v_k\,;\mu_k^R)(v-v_k)\\ \text{subject to} & v_i\geq 0,\quad i=1,2,\ldots,n.\end{array} \tag{2.8}$$

The following lemma provides the connections between the inertias of various matrices (part (i) may be used to conclude that subproblem (2.8) is convex).

LEMMA 2.2. *Let μ, ν be scalars such that $\mu>0$ and $\nu\geq 0$. Let H and J be matrices such that H is symmetric $n\times n$ and J is $m\times n$. If we define*

$$B^\nu=\begin{pmatrix}H+\frac{1}{\mu}(1+\nu)J^TJ & \nu J^T\\ \nu J & \nu\mu I_m\end{pmatrix} \quad \textit{and} \quad K=\begin{pmatrix}H & J^T\\ J & -\mu I_m\end{pmatrix},$$

then the following properties hold:

(i) *The matrix $H+\frac{1}{\mu}J^TJ$ is positive definite if and only if*

$$\mathrm{In}(B^\nu)=\begin{cases}(n+m,0,0) & \text{for } \nu>0;\\ (n,0,m) & \text{for } \nu=0.\end{cases}$$

(ii) *The matrix* $H + \frac{1}{\mu}J^TJ$ *is positive definite if and only if* $\mathrm{In}(K) = (n, m, 0)$.

Proof. It may be verified by direct multiplication that

$$L^TB^\nu L = \begin{pmatrix} H + \frac{1}{\mu}J^TJ & 0 \\ 0 & \nu\mu I_m \end{pmatrix}, \quad \text{where} \quad L = \begin{pmatrix} I_n & 0 \\ -\frac{1}{\mu}J & I_m \end{pmatrix}.$$

The matrix L is nonsingular, and Sylvester's law of inertia gives

$$\mathrm{In}(B^\nu) = \mathrm{In}(L^TB^\nu L) = \begin{cases} \mathrm{In}\left(H + \frac{1}{\mu}J^TJ\right) + (m, 0, 0) & \text{for } \nu > 0; \\ \mathrm{In}\left(H + \frac{1}{\mu}J^TJ\right) + (0, 0, m) & \text{for } \nu = 0, \end{cases}$$

which implies the result of part (i).

To prove part (ii), consider the identity

$$S^TKS = \begin{pmatrix} H + \frac{1}{\mu}J^TJ & 0 \\ 0 & -\mu I_m \end{pmatrix}, \quad \text{where} \quad S = \begin{pmatrix} I_n & 0 \\ \frac{1}{\mu}J & I_m \end{pmatrix}.$$

It now follows from the nonsingularity of S and Sylvester's law of inertia that

$$\mathrm{In}(K) = \mathrm{In}(S^TKS) = \mathrm{In}\left(H + \frac{1}{\mu}J^TJ\right) + (0, m, 0),$$

from which part (ii) follows directly. ∎

A proof similar to that used for Theorem 5.1 in section 5 may be used to show that the first-order optimality conditions for any primal-dual QP solution $\widehat{v}_k = (\widehat{x}_k, \widehat{y}_k)$ of the bound-constrained QP (2.8) may be written in matrix form

$$\begin{pmatrix} \widehat{H}_{\scriptscriptstyle F} & J_{\scriptscriptstyle F}^T \\ J_{\scriptscriptstyle F} & -\mu_k^R I \end{pmatrix} \begin{pmatrix} [\,\widehat{x}_k - x_k\,]_{\scriptscriptstyle F} \\ -(\widehat{y}_k - y_k) \end{pmatrix} = -\begin{pmatrix} [\,g_k + \widehat{H}_k s_k - J_k^T y_k\,]_{\scriptscriptstyle F} \\ c_k + J_k s_k + \mu_k^R(y_k - y_k^E) \end{pmatrix}, \tag{2.9}$$

where c_k, g_k, and J_k denote the quantities $c(x)$, $g(x)$, and $J(x)$ evaluated at x_k, and the quantities with suffix "F" are defined in terms of the index set $\mathcal{F}(\widehat{x}_k)$, i.e., $\widehat{H}_{\scriptscriptstyle F}$ is the matrix of free rows and columns of $\widehat{H}_k = \widehat{H}(x_k, y_k)$, and $J_{\scriptscriptstyle F}$ is the matrix of free columns of J_k. The vector s_k is nonpositive with components

$$[\,s_k\,]_i = \begin{cases} -[\,x_k\,]_i & \text{if } i \in \mathcal{A}(\widehat{x}_k); \\ 0 & \text{if } i \in \mathcal{F}(\widehat{x}_k). \end{cases}$$

As $\widehat{H}_k + (1/\mu_k^R)J_k^TJ_k$ is positive definite by construction, it follows immediately that the principal submatrix $\widehat{H}_{\scriptscriptstyle F} + (1/\mu_k^R)J_{\scriptscriptstyle F}^TJ_{\scriptscriptstyle F}$ is also positive definite. We may then apply part (ii) of Lemma 2.2 with values $H = \widehat{H}_{\scriptscriptstyle F}$, $J = J_{\scriptscriptstyle F}$, and $\mu = \mu_k^R$, to infer that the matrix associated with equations (2.9) is nonsingular. It follows that if $\mathcal{A}(\widehat{x}_k) = \mathcal{A}(x_k)$, then s_k is zero and $(\widehat{x}_k, \widehat{y}_k)$ satisfies the perturbed Newton equations

$$\begin{pmatrix} \widehat{H}_{\scriptscriptstyle F} & J_{\scriptscriptstyle F}^T \\ J_{\scriptscriptstyle F} & -\mu_k^R I \end{pmatrix} \begin{pmatrix} [\,\widehat{x}_k - x_k\,]_{\scriptscriptstyle F} \\ -(\widehat{y}_k - y_k) \end{pmatrix} = -\begin{pmatrix} [\,g_k - J_k^T y_k\,]_{\scriptscriptstyle F} \\ c_k + \mu_k^R(y_k - y_k^E) \end{pmatrix}. \tag{2.10}$$

A key property is that if $\mu_k^R = 0$ and $J_{\scriptscriptstyle F}$ has full rank, then this equation is identical to the equation for the conventional SQP step given by (1.5). This provides the

motivation to use a small penalty parameter μ_k^R for the step computation and a different larger penalty parameter μ_k for the merit function. In this context, μ_k^R plays the role of a *regularization* parameter rather than a *penalty* parameter, thereby providing an $O(\mu_k^R)$ estimate of the conventional SQP direction. This approach is nonstandard because a small "penalty parameter" μ_k^R is used by design, whereas conventional augmented Lagrangian-based methods attempt to keep μ as large as possible [5, 21].

The discussion above has established the relationship between the computation of the primal-dual bound-constrained step and the solution of a regularized QP. The next result formalizes the connection between the primal-dual step and the step associated with a *stabilized* SQP method.

RESULT 2.1. *Let ν and μ_k^R denote fixed scalars such that $\nu \geq 0$ and $\mu_k^R > 0$. Let $v_k = (x_k, y_k)$, $g_k = g(x_k)$, $c_k = c(x_k)$, and $J_k = J(x_k)$. Given a matrix $\widehat{H}_k = \widehat{H}(x_k, y_k)$ such that $\widehat{H}_k + (1/\mu_k^R)J_k^T J_k$ is positive definite, consider the subproblem*

$$\begin{array}{ll} \underset{x,y}{\text{minimize}} & g_k^T(x - x_k) + \frac{1}{2}(x - x_k)^T \widehat{H}_k (x - x_k) + \frac{1}{2}\mu_k^R \|y\|^2 \\ \text{subject to} & c_k + J_k(x - x_k) + \mu_k^R (y - y_k^E) = 0, \quad x \geq 0, \end{array} \tag{2.11}$$

which is the stabilized SQP subproblem (1.6) *defined with $\mu_k = \mu_k^R$ and $y_k = y_k^E$. The following results hold:*

(i) *The stabilized QP* (2.11) *has a unique bounded primal-dual solution $\widehat{v}_k = (\widehat{x}_k, \widehat{y}_k)$.*

(ii) *The unique solution $\widehat{v}_k = (\widehat{x}_k, \widehat{y}_k)$ of the stabilized QP* (2.11) *is a solution of the bound-constrained QP* (2.8) *for all $\nu \geq 0$. If $\nu > 0$, then the stabilized solution $\widehat{v}_k = (\widehat{x}_k, \widehat{y}_k)$ is the* unique *solution of* (2.8).

Proof. To simplify notation, the regularization parameter μ_k^R will be denoted by μ. For part (i), given the particular feasible point $v_0 = (x_k, \pi_k)$ with $\pi_k = y_k^E - c_k/\mu$, any feasible point $v = (x, y)$ may be written as

$$v = v_0 + Nw \;\text{ for some vector } w \in \mathbb{R}^n, \text{ where } N = \begin{pmatrix} \mu I \\ -J_k \end{pmatrix}.$$

The matrix N is $(n + m) \times n$ with rank n, and its columns form a basis for the null-space of the constraint matrix $\begin{pmatrix} J_k & \mu I \end{pmatrix}$. Applying this equivalent form of v to (2.11) gives the equivalent problem

$$\underset{w \in \mathbb{R}^n}{\text{minimize}} \;\; \frac{\mu}{2} w^T \left(\widehat{H}_k + \frac{1}{\mu} J_k^T J_k \right) w + w^T (g_k - J_k^T \pi_k) \;\; \text{subject to} \;\; \mu w \geq -x_k.$$

The matrix $\widehat{H}_k + (1/\mu)J_k^T J_k$ is positive definite by assumption, and it follows that the stabilized QP (2.11) is equivalent to a convex program with a strictly convex objective. The existence of a unique bounded solution follows directly.

For part (ii), it is sufficient to show that the optimality conditions for problems (2.11) and (2.8) are equivalent. The first-order conditions for (x, y) to be a solution of the stabilized QP (2.11) are

$$\begin{aligned} c_k + J_k(x - x_k) + \mu(y - y_k^E) = 0, \qquad & \mu y = \mu w, \\ g_k + \widehat{H}_k(x - x_k) - J_k^T w - z = 0, \qquad & z \geq 0, \\ z \cdot x = 0, \qquad & x \geq 0, \end{aligned}$$

where w and z denote the dual variables for the equality and inequality constraints of problem (2.11), respectively. Eliminating w using the equation $w = y$ gives

$$c_k + J_k(x - x_k) + \mu(y - y_k^E) = 0, \tag{2.12a}$$

$$g_k + \widehat{H}_k(x - x_k) - J_k^T y - z = 0, \qquad z \geq 0, \tag{2.12b}$$

$$z \cdot x = 0, \qquad x \geq 0. \tag{2.12c}$$

First, part (ii) is established for the case $\nu > 0$. The optimality conditions for the bound-constrained QP (2.8) are

$$\nabla M^\nu(v_k; y_k^E, \mu) + B^\nu(v_k\,; \mu)(v - v_k) = \begin{pmatrix} z \\ 0 \end{pmatrix}, \qquad z \geq 0, \tag{2.13a}$$

$$z \cdot x = 0, \qquad x \geq 0. \tag{2.13b}$$

Premultiplying the equality of (2.13a) by the nonsingular matrix T defined by

$$T = \begin{pmatrix} I_n & -\frac{1+\nu}{\nu\mu} J_k^T \\ 0 & \frac{1}{\nu} I_m \end{pmatrix},$$

and using the definitions (2.2) and (2.3) yields the equivalent conditions

$$g_k + \widehat{H}_k(x - x_k) - J_k^T y - z = 0 \quad \text{and} \quad c_k + J_k(x - x_k) + \mu(y - y_k^E) = 0,$$

which are identical to the relevant equalities in (2.12). Thus, if $\nu > 0$, the solutions of (2.11) and (2.8) are identical.

It remains to consider the case $\nu = 0$. In this situation, the objective function of the QP (2.8) includes only the primal variables x, which implies that the problem may be written as

$$\begin{aligned} &\underset{x}{\text{minimize}} \;\; (g_k - J_k^T \pi_k)^T (x - x_k) + \tfrac{1}{2}(x - x_k)^T \left(\widehat{H}_k + \frac{1}{\mu} J_k^T J_k \right)(x - x_k) \\ &\text{subject to } x \geq 0, \end{aligned} \tag{2.14}$$

with y an arbitrary vector. Although there are infinitely many solutions of (2.8) when ν is zero, the vector x associated with a particular solution (x, y) is unique because it is the solution of problem (2.14) for a positive-definite matrix $\widehat{H}_k + \frac{1}{\mu} J_k^T J_k$. The optimality conditions for (2.14) are

$$g_k - J_k^T \pi_k + \left(\widehat{H}_k + \frac{1}{\mu} J_k^T J_k \right)(x - x_k) = z, \qquad z \geq 0,$$

$$z \cdot x = 0, \qquad x \geq 0. \tag{2.15}$$

For the given y_k and optimal x, define the m-vector y such that

$$y - y_k = -\frac{1}{\mu}\big(J_k(x - x_k) + c_k + \mu(y_k - y_k^E)\big) = -\frac{1}{\mu}\big(J_k(x - x_k) + \mu(y_k - \pi_k)\big). \tag{2.16}$$

Equation (2.16) and the equality of (2.15) may be combined to give the matrix equation

$$\begin{pmatrix} g_k - J_k^T y_k + 2J_k^T(y_k - \pi_k) \\ \mu(y_k - \pi_k) \end{pmatrix} + \begin{pmatrix} \widehat{H}_k + \frac{2}{\mu} J_k^T J_k & J_k^T \\ J_k & \mu I \end{pmatrix} \begin{pmatrix} x - x_k \\ y - y_k \end{pmatrix} = \begin{pmatrix} z \\ 0 \end{pmatrix}.$$

Applying the nonsingular matrix $\left(\begin{smallmatrix} I_n & -\frac{2}{\mu}J_k^T \\ 0 & I_m \end{smallmatrix}\right)$ to both sides of this equation yields

$$\begin{pmatrix} g_k - J_k^T y_k \\ c_k + \mu(y_k - y_k^E) \end{pmatrix} + \begin{pmatrix} \widehat{H}_k & -J_k^T \\ J_k & \mu I \end{pmatrix} \begin{pmatrix} x - x_k \\ y - y_k \end{pmatrix} = \begin{pmatrix} z \\ 0 \end{pmatrix}.$$

These equations, together with the complementarity conditions $z \cdot x = 0$ of (2.15), are the optimality conditions for the stabilized QP (2.11) (cf. (2.12)). It follows that if $\nu = 0$, the unique solution of (2.11) is a solution of (2.8), which is what we wanted to show.

If $\nu > 0$, the uniqueness of the solution $v = (x, y)$ follows from part (i) of Lemma 2.2, which implies that the objective Hessian of the bound constrained QP (2.8) is positive definite, thereby ensuring a strictly convex QP. ∎

2.2. Definition of the new iterate. Once the search direction $d_k = \widehat{x}_k - x_k$ has been determined, a "flexible" backtracking line search is performed on the primal-dual augmented Lagrangian. A conventional backtracking line search defines $v_{k+1} = v_k + \alpha_k d_k$, where $\alpha_k = 2^{-j}$ and j is the smallest nonnegative integer such that

$$M^\nu(v_k + \alpha_k d_k\,; y_k^E, \mu_k) \le M^\nu(v_k\,; y_k^E, \mu_k) + \alpha_k \eta_S d_k^T \nabla M^\nu(v_k\,; y_k^E, \mu_k)$$

for a given $\eta_S \in (0, \frac{1}{2})$. However, this approach would suffer from the Maratos effect [41] simply because the penalty parameter μ_k and the regularization parameter μ_k^R used to compute the trial step have different values in general. This difficulty is avoided by using an augmented Lagrangian version of the "flexible penalty function" proposed by Curtis and Nocedal [12]. This method defines a step length of the form $\alpha_k = 2^{-j}$, where j is the smallest nonnegative integer satisfying

$$M^\nu(v_k + \alpha_k d_k\,; y_k^E, \mu_k^F) \le M^\nu(v_k\,; y_k^E, \mu_k^F) + \alpha_k \eta_S \delta_k \tag{2.17}$$

for some value $\mu_k^F \in [\mu_k^R, \mu_k]$, and δ_k such that

$$\delta_k = \max\left(d_k^T \nabla M^\nu(v_k\,; y_k^E, \mu_k^R), -\eta_D \|d_k\|^2\right) \le 0, \tag{2.18}$$

with η_D a small positive constant. The use of the second term in the definition of δ_k increases the chance that a step is accepted during the early iterations when $|d_k^T \nabla M^\nu(v_k\,; y_k^E, \mu_k^R)|$ is large. Once an appropriate value for α_k is found, the new primal-dual solution estimate is given by

$$x_{k+1} = x_k + \alpha_k(\widehat{x}_k - x_k) \quad \text{and} \quad y_{k+1} = y_k + \alpha_k(\widehat{y}_k - y_k).$$

In a practical algorithm, the step is reduced until the Armijo condition (2.17) is satisfied for one of the values $\mu_k^F = \mu_k$ or $\mu_k^F = \mu_k^R$ (where the condition for $\mu_k^F = \mu_k$ is tried first). The following simple argument shows that the acceptance criterion (2.17) is well-defined; i.e., the sequence $\{2^{-j}\}$ must terminate with an acceptable α_k. As $v = v_k$ is feasible for the strictly convex problem (2.8), the search direction $d_k = (\widehat{x}_k - x_k, \widehat{y}_k - y_k)$ is a feasible descent direction for $M^\nu(v\,; y_k^E, \mu_k^R)$ at $v_k = (x_k, y_k)$. If follows from standard theory that the weakened Armijo condition (2.17) will be satisfied for $\mu_k^F = \mu_k^R$ and all $\alpha_k > 0$ sufficiently small.

2.3. Updating the multiplier estimate. The QP equivalence established in Result 2.1, together with the definition of the stabilized SQP subproblem (1.6), implies that setting $y_k^E = y_k$ in the definition of subproblem (2.11) (or, equivalently,

in the bound-constrained QP (2.8)) makes the proposed trial step identical to that of the stabilized SQP method. This motivates an update strategy that allows the definition $y_k^E = y_k$ as often as possible. (The validity of this strategy is supported by the numerical results of section 6.) The idea is to define $y_{k+1}^E = y_{k+1}$ for the next subproblem if the line search gives an (x_{k+1}, y_{k+1}) that improves at least one of two merit functions that measure the accuracy of (x_{k+1}, y_{k+1}) as an estimate of (x^*, y^*). Let β denote a small positive parameter and consider the merit functions

$$\phi_V(x,y) = \eta(x) + \beta\omega(x,y) \quad \text{and} \quad \phi_O(x,y) = \beta\eta(x) + \omega(x,y), \tag{2.19}$$

where $\eta(x)$ and $\omega(x,y)$ are the feasibility violation and optimality measures

$$\eta(x) = \|c(x)\| \quad \text{and} \quad \omega(x,y) = \left\|\min\left(x, g(x) - J(x)^T y\right)\right\|. \tag{2.20}$$

These functions provide two alternative weighted measures of the accuracy of (x,y) as an approximate solution of problem (NP) rather than as an approximate minimizer of M^ν. Both measures are bounded below by zero, and are equal to zero if v is a first-order solution to problem (NP).

Given these definitions, the estimate y_k^E is updated when any iterate $v_k = (x_k, y_k)$ satisfies either $\phi_V(v_k) \le \frac{1}{2}\phi_V^{\max}$ or $\phi_O(v_k) \le \frac{1}{2}\phi_O^{\max}$, where $\phi_V^{\max}$ and $\phi_O^{\max}$ are bounds that are updated throughout the solution process. To ensure global convergence, an update to y_k^E forces a decrease in either $\phi_V^{\max}$ or $\phi_O^{\max}$. The idea is to choose the parameter β of (2.20) to be relatively small, say $\beta = 10^{-5}$. This allows frequent updates to y_k^E, as shown in the numerical results of section 6.

Finally, y_k^E is also updated if an approximate first-order solution to problem (2.5) has been found for the values $y^E = y_k^E$ and $\mu = \mu_k^R$. The test for optimality is

$$\|\nabla_y M^\nu(v_{k+1}; y_k^E, \mu_k^R)\| \le \tau_k \ \text{ and } \ \left\|\min\left(x_{k+1}, \nabla_x M^\nu(v_{k+1}; y_k^E, \mu_k^R)\right)\right\| \le \tau_k \tag{2.21}$$

for some small tolerance $\tau_k > 0$. This condition is rarely triggered in practice, but the test is needed to ensure global convergence (see section 6 for statistics on the frequency of this test being satisfied). Nonetheless, if condition (2.21) is satisfied, y_k^E is updated with the safeguarded estimate

$$y_{k+1}^E = \max\left(-y_{\max} e, \min(y_{k+1}, y_{\max} e)\right)$$

for some large positive scalar constant $y_{\max}$.

2.4. Updating the penalty parameters. The following definition is designed to decrease μ_k^R only in the neighborhood of an optimal point (assuming that the problem is not locally infeasible):

$$\mu_{k+1}^R = \begin{cases} \min\left(\frac{1}{2}\mu_k^R, \|r_{\text{opt}}(v_{k+1})\|^{3/2}\right) & \text{if (2.21) is satisfied;} \\ \min\left(\ \mu_k^R, \|r_{\text{opt}}(v_{k+1})\|^{3/2}\right) & \text{otherwise,} \end{cases} \tag{2.22}$$

where r_{opt} is the vector-valued function

$$r_{\text{opt}}(v) = \begin{pmatrix} c(x) \\ \min\left(\, x,\ g(x) - J(x)^T y\,\right) \end{pmatrix}. \tag{2.23}$$

The update to μ_k is motivated by a different goal. Namely, μ_k should be decreased only when the trial step indicates that the merit function defined with penalty pa-

rameter μ_k *increases*. This motivates the definition

(2.24)
$$\mu_{k+1} = \begin{cases} \mu_k & \text{if } M^\nu(v_{k+1}\,; y_k^E, \mu_k) \le M^\nu(v_k\,; y_k^E, \mu_k) + \widehat{\alpha}_k \eta_S \delta_k; \\ \max\left(\tfrac{1}{2}\mu_k, \mu_{k+1}^R\right) & \text{otherwise,} \end{cases}$$

where $\widehat{\alpha}_k = \min(\alpha_{\min}, \alpha_k)$ for some positive $\alpha_{\min}$, and δ_k is defined by (2.18). The use of the scalar $\alpha_{\min}$ increases the likelihood that μ_k will not be decreased.

2.5. Formal statement of the algorithm. This section provides a formal statement of the proposed method as Algorithm 2.1 and includes some additional details. During each iteration, the trial step is computed as described in section 2.1, the solution estimate is updated as in section 2.2, the multiplier estimate y_k^E is updated as in section 2.3, and the penalty parameters are updated as in section 2.4. It is clear from Theorem 2.1 that it is advantageous to update the value of y_k^E as often as possible in order to promote global convergence, and (from Result 2.1) to recover a stabilized SQP subproblem. There are three possibilities. First, y_k^E is set to y_{k+1} if (x_{k+1}, y_{k+1}) is acceptable to either of the merit functions ϕ_V or ϕ_O given by (2.19). These iterates are referred to as "V-iterates" and "O-iterates," respectively. The numerical results

---

Algorithm 2.1. Regularized primal-dual SQP algorithm (`pdSQP`).
Input (x_0, y_0);
Set control parameters $\alpha_{\min} > 0$, $\eta_S \in (0,1)$, $\eta_D \in (0,1)$, $\tau_{\text{stop}} > 0$, $k_{\max} > 0$,
$0 < \beta \ll 1$, $y_{\max} \gg 1$, and $\nu > 0$;
Initialize $y_0^E = y_0$, $\tau_0 > 0$, $\mu_0^R > 0$, $\mu_0 \in [\mu_0^R, \infty)$, and $k = 0$;
Compute $f(x_0)$, $c(x_0)$, $g(x_0)$, $J(x_0)$, and $H(x_0, y_0)$;
for $k = 0, 1, 2, \ldots, k_{\max}$ **do**
  Define $\widehat{H}(x_k, y_k) \approx H(x_k, y_k)$ such that $\widehat{H}(x_k, y_k) + (1/\mu_k^R) J_k^T J_k$ is positive definite;
  Solve the QP (2.11) (which is equivalent to solving (2.8)) for $(\widehat{x}_k, \widehat{y}_k)$;
  Find an α_k satisfying condition (2.17) for either $\mu_k^F = \mu_k$ or $\mu_k^F = \mu_k^R$;
  Update the primal-dual estimate $(x_{k+1}, y_{k+1}) = (x_k, y_k) + \alpha_k(\widehat{x}_k - x_k, \widehat{y}_k - y_k)$;
  Compute $f(x_{k+1})$, $c(x_{k+1})$, $g(x_{k+1})$, $J(x_{k+1})$, and $H(x_{k+1}, y_{k+1})$;
  if $\phi_V(x_{k+1}, y_{k+1}) \le \frac{1}{2}\phi_V^{\max}$ **then** [V-iterate]
    $\phi_V^{\max} = \frac{1}{2}\phi_V^{\max}$;
    $y_{k+1}^E = y_{k+1}$; $\tau_{k+1} = \tau_k$;
  else if $\phi_O(x_{k+1}, y_{k+1}) \le \frac{1}{2}\phi_O^{\max}$ **then** [O-iterate]
    $\phi_O^{\max} = \frac{1}{2}\phi_O^{\max}$;
    $y_{k+1}^E = y_{k+1}$; $\tau_{k+1} = \tau_k$;
  else if $v_{k+1} = (x_{k+1}, y_{k+1})$ satisfies (2.21) **then** [M-iterate]
    $y_{k+1}^E = \max\big(-y_{\max} e, \min(y_{k+1}, y_{\max} e)\big)$; $\tau_{k+1} = \frac{1}{2}\tau_k$;
  else [F-iterate]
    $y_{k+1}^E = y_k^E$; $\tau_{k+1} = \tau_k$;
  end if
  Update μ_{k+1}^R and μ_{k+1} according to (2.22) and (2.24), respectively;
  if $\|r_{\text{opt}}(v_{k+1})\| \le \tau_{\text{stop}}$ **then** *exit*;
end (for)

---

in section 6 show that y_k^E is updated this way most of the time. Second, if (x_{k+1}, y_{k+1}) is not acceptable to either of the merit functions ϕ_V or ϕ_O, then conditions (2.21) are used to determine if (x_{k+1}, y_{k+1}) is an approximate first-order solution of the bound-constrained problem (2.5). If conditions (2.21) are satisfied, the iterate is called an "M-iterate." In this case, the regularization parameter μ_k^R and subproblem tolerance τ_k are decreased and y_k^E is updated by $\pi(x_k\,; y_k^E, \mu_k^R)$ as in (2.3). Finally, if neither of the first two cases occurs, the multiplier estimate y_k^E is fixed at its current value and the associated iterate is designated an "F-iterate."

3. Convergence. The convergence of Algorithm 2.1 is discussed under the following assumptions.

Assumption 3.1. Each $\widehat{H}(x_k, y_k)$ is chosen so that the sequence $\{\widehat{H}(x_k, y_k)\}_{k\geq 0}$ is bounded, with $\{\widehat{H}(x_k, y_k) + (1/\mu_k^R)J(x_k)^T J(x_k)\}_{k\geq 0}$ uniformly positive definite.

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.

In the "worst" case, i.e., when all iterates are eventually M-iterates or F-iterates, Algorithm 2.1 emulates a *primal-dual* augmented Lagrangian method [7, 8, 47]. Consequently, it is possible that y_k^E and μ_k^R will remain fixed over a sequence of iterations, although this has been uncommon in our preliminary numerical results. The following result concerns this situation.

THEOREM 3.1. *Let Assumptions* 3.1–3.3 *hold. If there exists an integer* $\widehat{k}$ *such that* $\mu_k^R \equiv \mu^R > 0$ *and* k *is an F-iterate for all* $k \geq \widehat{k}$*, then the following hold for the search directions* $d_k = (\widehat{x}_k - x_k, \widehat{y}_k - y_k)$*, where* $(\widehat{x}_k, \widehat{y}_k)$ *is the solution of subproblem* (2.8)*:*

(i) $\{d_k\}_{k\geq\widehat{k}}$ *is uniformly bounded;*

(ii) $\{d_k\}_{k\geq\widehat{k}}$ *is bounded away from zero; and*

(iii) *there exists a constant* $\epsilon > 0$ *such that*

$$\nabla M^\nu(v_k\,; y_k^E, \mu_k^R)^T d_k \leq -\epsilon \;\; \textit{for all } k \geq \widehat{k}.$$

Proof. The assumptions of this theorem imply that

$$\tau_k \equiv \tau > 0, \;\; \mu_k^R = \mu^R, \;\; \text{and} \;\; y_k^E = y^E \;\; \text{for all } k \geq \widehat{k}. \tag{3.1}$$

First we prove part (i). As discussed in the proof of Result 2.1, it is known that the solution $(\widehat{x}_k, \widehat{y}_k)$ of (2.8) satisfies

$$d_k = \begin{pmatrix} \widehat{x}_k - x_k \\ \widehat{y}_k - y_k \end{pmatrix} = \begin{pmatrix} 0 \\ \pi_k - y_k \end{pmatrix} + N_k w^*, \;\; \text{where} \;\; N_k = \begin{pmatrix} \mu^R I \\ -J_k \end{pmatrix},$$

and w^* is the unique solution of

$$\underset{w\in\mathbb{R}^n}{\text{minimize}} \;\; \tfrac{1}{2}\mu^R w^T\left(\widehat{H}_k + \frac{1}{\mu^R}J_k^T J_k\right)w + w^T\big(g_k - J_k^T\pi_k\big) \;\; \text{subject to} \;\; x_k + \mu^R w \geq 0$$

for all $k \geq \widehat{k}$. It follows from Assumption 3.1 that $\{d_k\}_{k\geq\widehat{k}}$ is uniformly bounded, provided that the quantities $g_k - J_k^T\pi_k$, N_k, π_k, and y_k are all uniformly bounded for $k \geq \widehat{k}$. The boundedness of $g_k - J_k^T\pi_k$, π_k and N_k follow from Assumption 3.2, Assumption 3.3, (3.1), and (2.3). It remains to prove that the set of multipliers $\{y_k\}_{k\geq\widehat{k}}$ is bounded. As $\mu_k^R = \mu^R$ for all $k \geq \widehat{k}$, the update to μ_k given by (2.24) implies that $\mu_k \equiv \mu \geq \mu^R$ for some μ and all k sufficiently large. For all subsequent iterations the primal-dual merit function is monotonically decreasing, i.e.,

$M^\nu(x_{k+1}, y_{k+1}\,; y^E, \mu) \le M^\nu(x_k, y_k\,; y^E, \mu)$. It follows that $\{y_k\}_{k\ge\widehat{k}}$ must be bounded, since if there were a subsequence such that $\|y_k\|$ went to infinity, then for the same subsequence M^ν would also go to infinity because both $\{f_k - c_k^T y^E + \frac{1}{2\mu}\|c_k\|^2\}_{k\ge\widehat{k}}$ and $\{c_k\}_{k\ge\widehat{k}}$ are bounded from Assumptions 3.2 and 3.3. This completes the proof of part (i).

Part (ii) is established by showing that $\{\|d_k\|\}_{k\ge\widehat{k}}$ is bounded away from zero. If this were not the case, there would exist a subsequence $\mathcal{S}_1 \subseteq \{k : k \ge \widehat{k}\}$ such that $\lim_{k\in\mathcal{S}_1} d_k = 0$, where $d_k = (\widehat{x}_k - x_k, \widehat{y}_k - y_k)$ and $(\widehat{x}_k, \widehat{y}_k)$ is a solution of problem (2.8). It follows that d_k satisfies

$$\begin{pmatrix}\widehat{z}_k\\ 0\end{pmatrix} = B^\nu(v_k\,;\mu^R)d_k + \nabla M^\nu(v_k\,; y^E, \mu^R) \quad \text{and} \quad 0 = \min(\widehat{x}_k, \widehat{z}_k)$$

for all $k \in \mathcal{S}_1$. It may then be inferred from the definition of $B^\nu(v_k\,;\mu^R)$, Assumptions 3.1–3.3, and the definitions (3.1) of τ_k, μ_k^R, and y_k^E that for $k \in \mathcal{S}_1$ sufficiently large, the iterate v_k satisfies the definition (2.21) of an M-iterate, and as a consequence, μ_k^R will be decreased. This contradicts the assumption that $\mu_k^R \equiv \mu^R$ for all $k \ge \widehat{k}$. It follows that $\{\|d_k\|\}_{k\ge\widehat{k}}$ is bounded away from zero and part (ii) holds.

The proof of part (iii) is also by contradiction. Assume that there exists a subsequence $\mathcal{S}_2$ of $\{k : k \ge \widehat{k}\}$ such that

$$\lim_{k\in\mathcal{S}_2} \nabla M^\nu(v_k\,; y^E, \mu^R)^T d_k = 0, \tag{3.2}$$

where we have used (3.1) and d_k is defined as above. As the vector $v_k = (x_k, y_k)$ is feasible for the convex problem (2.8), and $(\widehat{x}_k, \widehat{y}_k)$ is the solution of problem (2.8) for $\nu > 0$ chosen in Algorithm 2.1, it must hold that

$$\begin{aligned}
-\nabla M^\nu(v_k\,; y^E, \mu^R)^T d_k &\ge \tfrac{1}{2} d_k^T B^\nu(v_k\,;\mu^R) d_k \\
&= \tfrac{1}{2} d_k^T L_k^{-T} L_k^T B^\nu(v_k\,;\mu^R) L_k L_k^{-1} d_k \\
&= \tfrac{1}{2} d_k^T L_k^{-T} \begin{pmatrix} \widehat{H}_k + \frac{1}{\mu^R} J_k^T J_k & 0 \\ 0 & \nu\mu^R \end{pmatrix} L_k^{-1} d_k,
\end{aligned}$$

where L_k denotes the nonsingular matrix

$$L_k = \begin{pmatrix} I & 0 \\ -\frac{1}{\mu^R} J_k & I \end{pmatrix}, \quad \text{with} \quad L_k^{-1} d_k = \begin{pmatrix} p_k \\ q_k + \frac{1}{\mu^R} J_k p_k \end{pmatrix},$$

with $p_k = \widehat{x}_k - x_k$ and $q_k = \widehat{y}_k - y_k$. Assumption 3.1 yields

$$\begin{aligned}
-\nabla M^\nu(v_k\,; y^E, \mu^R)^T d_k &\ge \tfrac{1}{2} p_k^T \left(\widehat{H}_k + \frac{1}{\mu^R} J_k^T J_k \right) p_k + \tfrac{1}{2}\nu\mu^R \|q_k + (1/\mu^R) J_k p_k\|^2 \\
&\ge \lambda_{\min} \|p_k\|^2 + \tfrac{1}{2}\nu\mu^R \|q_k + (1/\mu^R) J_k p_k\|^2
\end{aligned}$$

for some $\lambda_{\min} > 0$. Combining this inequality with (3.2) gives the limit

$$\lim_{k\in\mathcal{S}_2} p_k = \lim_{k\in\mathcal{S}_2} \left(q_k + \frac{1}{\mu^R} J_k p_k \right) = 0,$$

in which case $\lim_{k\in\mathcal{S}_2} q_k = 0$ follows from Assumptions 3.2 and 3.3. This contradicts the result of part (ii) and so part (iii) must hold. ∎

The following theorem states the main convergence result for Algorithm 2.1.

THEOREM 3.2. *Let Assumptions* 3.1–3.3 *hold. If v_k denotes the kth iterate generated by Algorithm* 2.1*, then either*

(i) *Algorithm* 2.1 *terminates with an approximate primal-dual first-order solution v_k satisfying $\|r_{\rm opt}(v_k)\| \le \tau_{\rm stop}$, where $r_{\rm opt}$ is defined by* (2.23)*; or*

(ii) *there exists a subsequence $\mathcal{S}$ such that $\lim_{k\in\mathcal{S}} \mu_k^R = 0$, $\{y_k^E\}_{k\in\mathcal{S}}$ is bounded, $\lim_{k\in\mathcal{S}} \tau_k = 0$, and for each $k \in \mathcal{S}$ the vector v_{k+1} is an approximate first-order solution of* (2.5) *with the choice $y^E = y_k^E$ and $\mu = \mu_k^R$ that satisfies* (2.21).

Proof. If there exists a subsequence of $\{\|r_{\rm opt}(v_k)\|\}_{k\ge0}$ that converges to zero, then clearly case (i) holds. Therefore, for the remainder of the proof, it is assumed that the sequence $\{\|r_{\rm opt}(v_k)\|\}_{k\ge0}$ is bounded away from zero.

From the definitions of a V-iterate and O-iterate, the functions ϕ_V and ϕ_O, and the update strategies for $\phi_V^{\max}$ and $\phi_O^{\max}$, we conclude that the number of V-iterates and O-iterates must be finite. We claim that there must be an infinite number of M-iterates. To prove this, assume to the contrary that the number of M-iterates is finite, so that all iterates are F-iterates for k sufficiently large. It follows from the form of the update to μ_k^R (2.22) and the assumption made in this case, that eventually μ_k^R remains constant. In this case, the update to μ_k given by (2.24) implies that eventually, μ_k also remains constant. These arguments imply the existence of an integer $\widehat{k}$ such that

$$\mu_k^R \equiv \mu^R \le \mu \equiv \mu_k, \quad y_k^E \equiv y^E, \quad \tau_k \equiv \tau > 0, \quad \text{and } k \text{ is an F-iterate for all } k \ge \widehat{k}.$$

It follows from (2.24) that

$$M^\nu(v_{k+1}\,; y^E, \mu) \le M^\nu(v_k\,; y^E, \mu) + \min(\alpha_{\min}, \alpha_k)\eta_s\delta_k \quad \text{for all } k \ge \widehat{k}, \tag{3.3}$$

where δ_k is defined by (2.18). Moreover, parts (ii) and (iii) of Theorem 3.1 ensure that $\{\delta_k\}_{k\ge\widehat{k}}$ is a negative sequence bounded away from zero. In addition, it must hold that $\{\alpha_k\}_{k\ge\widehat{k}}$ is bounded away from zero. To see this, note that parts (i) and (iii) of Theorem 3.1 and Assumption 3.2 ensure that $\{\alpha_k\}_{k\ge\widehat{k}}$ is bounded away from zero if a conventional Armijo line search is used, i.e., if $\mu_k^F = \mu^R$ and $\delta_k = d_k^T\nabla M^\nu(v_k\,; y^E, \mu^R)$ in (2.17). However, the computed value of α_k can be no smaller because the definition of δ_k is less restrictive, and the use of a flexible line search makes the acceptance of a step more likely. Combining these results with (3.3) yields

$$M^\nu(v_{k+1}\,; y^E, \mu) \le M^\nu(v_k\,; y^E, \mu) - \kappa \quad \text{for all } k \ge \widehat{k} \text{ and some } \kappa > 0,$$

so that $\lim_{k\to\infty} M^\nu(v_k\,; y^E, \mu) = -\infty$. However, Assumptions 3.2 and 3.3 ensure that this is not possible. This contradiction implies that there must exist infinitely many M-iterations, and *every* iterate is an M-iterate or F-iterate for k sufficiently large. Part (ii) now follows from (2.22) and the properties of the updates to τ_k and y_k^E used for M-iterates and F-iterates in Algorithm 2.1. ∎

The "ideal" scenario is that Algorithm 2.1 generates many V-iterates/O-iterates that converge to an approximate solution of (NP). This corresponds to the result of part (i) of Theorem 3.2. The scenario considered in part (ii) of Theorem 3.2, i.e., the generation of infinitely many M-iterates, is the "fall-back" position for Algorithm 2.1. This result would appear to be the best that can be expected without additional assumptions, such as the satisfaction of a constraint qualification at x^*. In particular,

Assumptions 3.1–3.3 do not preclude the possibility that problem (NP) is infeasible. Recent work indicates that the iterates of the stabilized SQP subproblem exhibit superlinear convergence under mild conditions (see, e.g., [13, 14, 36, 37]). In particular, neither strict complementarity nor a constraint qualification is required. However, further analysis is required to determine the conditions under which a V- or O-iterate is always generated in the limit, which is a necessary requirement for the proposed method to inherit the superlinear convergence rate associated with conventional stabilized SQP.

The following corollary illustrates how part (ii) of Theorem 3.2 relates to the original nonlinear problem (NP) with the assumption that the linear independence constraint qualification is satisfied (see, e.g., [43]).

COROLLARY 3.3. *Suppose that the linear independence constraint qualification holds at all limit points generated by Algorithm* 2.1. *Then Algorithm* 2.1 *terminates with an approximate primal-dual first-order solution v_k satisfying*

$$\|r_{\text{opt}}(v_k)\| \leq \tau_{\text{stop}}, \tag{3.4}$$

where r_{opt} is defined by (2.23).

Proof. The result is satisfied trivially if part (i) of Theorem 3.2 holds. Thus, it may be assumed that part (ii) of Theorem 3.2 holds, which ensures the existence of a subsequence $\mathcal{S}$ of M-iterates such that $\lim_{k\in\mathcal{S}} \mu_k^R = 0$, the set of multipliers estimates $\{y_k^E\}_{k\in\mathcal{S}}$ is bounded, $\lim_{k\in\mathcal{S}} \tau_k = 0$, and that for each $k \in \mathcal{S}$, the vector v_{k+1} is an approximate first-order solution of (2.5) with values $y^E = y_k^E$ and $\mu = \mu_k^R$ satisfying (2.21). Under Assumption 3.3 it may be assumed without loss of generality that $\lim_{k\in\mathcal{S}} x_k = x^*$ for some vector x^*. Consider the definition

$$y^* = \underset{y}{\operatorname{argmin}} \| [\, g(x^*) - J(x^*)^T y \,]_F \|^2,$$

where the "F" denotes the components corresponding to the free variables at x^*. It now follows from [47, Lemma 4.3.1] that (x^*, y^*) is a first-order solution of problem (NP), and that

$$\lim_{k\in\mathcal{S}} (x_k, y_k) = (x^*, y^*).$$

Moreover, it follows from Assumption 3.2 and the definition of r_{opt} that (3.4) will be satisfied for all $k \in \mathcal{S}$ sufficiently large. ∎

4. Convexification of the bound-constrained subproblem. An important aspect of the proposed method is the definition of $\widehat{H}(x_k, y_k)$, which is used to ensure that the bound-constrained QP subproblem (2.8) is convex. A conventional QP subproblem defined with the Hessian of the Lagrangian is not convex in general. To avoid solving an indefinite subproblem, most existing methods are based on solving a convex QP based on a positive-semidefinite approximation $\widehat{H}(x_k, y_k)$ of the Hessian $H(x_k, y_k)$. This convex subproblem is used to either define the search direction directly or identify the constraints for an equality-constrained QP subproblem that uses the exact Hessian (see, e.g., [3, 21, 30]).

Here we take a different approach and define a *convexified* QP subproblem in terms of the exact Hessian of the Lagrangian. The convex problem is defined in such a way that if the inner iterations do not alter the active set, then the computed direction is equivalent to a second-derivative stabilized SQP direction, provided that $y_k^E = y_k$. The method is based on defining a symmetric matrix $\widehat{H}(x_k, y_k)$ (not necessarily positive

definite) as a modification of $H(x_k, y_k)$ that gives a bounded convex primal-dual subproblem (2.8).

The remainder of this section focuses on the solution of a single QP subproblem, and the notation is simplified so that $v_k = (x_k, y_k) = (x, y)$, $J = J_k$, $H = H(x_k, y_k)$, $\widehat{H} = \widehat{H}(x_k, y_k)$, $B^\nu = B^\nu(x_k, y_k\,; \mu_k^R)$, and $\mu = \mu_k^R$. Similarly, $J_{\mathcal{F}}$ and $J_{\mathcal{A}}$ denote the columns of J associated with the index sets $\mathcal{F}(x)$ and $\mathcal{A}(x)$ of free and fixed variables at x. Throughout this section, if M is a symmetric matrix, then $M_{\mathcal{F}}$ and $M_{\mathcal{A}}$ denote the symmetric matrices with elements m_{ij} for i, j in $\mathcal{F}$ and $\mathcal{A}$, respectively. The definition of $\widehat{H}$ involves certain projections defined in terms of the gradients of the bound constraints. If $P_{\mathcal{A}}$ is the matrix whose (unit) columns are the subset of columns of the identify matrix with indices in $\mathcal{A}(x)$, then $P_{\mathcal{A}} P_{\mathcal{A}}^T x$ is the orthogonal projection of x onto the bounds in $\mathcal{A}(x)$. The complementary projection may be defined in terms of the matrix $P_{\mathcal{F}}$ with (unit) columns orthogonal to $P_{\mathcal{A}}$. With these definitions, the matrix $P = \begin{pmatrix} P_{\mathcal{F}} & P_{\mathcal{A}} \end{pmatrix}$ defines a permutation matrix such that $JP = \begin{pmatrix} J_{\mathcal{F}} & J_{\mathcal{A}} \end{pmatrix}$.

At any given x, the proposed convexification gives a matrix $\widehat{H}$ of the form

$$\widehat{H} = H + E + D, \tag{4.1}$$

where E is a symmetric positive-semidefinite matrix, and D is a positive-semidefinite diagonal. It must be emphasized that $\widehat{H}$ itself is not necessarily positive definite.

First, we consider the definition of E. Let K and $K_{\mathcal{F}}$ denote the matrices

$$K = \begin{pmatrix} H + E & J^T \\ J & -\mu I \end{pmatrix} \quad \text{and} \quad K_{\mathcal{F}} = \begin{pmatrix} H_{\mathcal{F}} + E_{\mathcal{F}} & J_{\mathcal{F}}^T \\ J_{\mathcal{F}} & -\mu I \end{pmatrix}. \tag{4.2}$$

We are particularly interested in matrices E that endow $K_{\mathcal{F}}$ with the property of *second-order consistency*.

DEFINITION 4.1. *If G is a symmetric matrix of order r and C is $m \times r$, then the matrix $\begin{pmatrix} G & C^T \\ C & -\mu I_m \end{pmatrix}$ is said to be* second-order consistent *if it has inertia $(r, m, 0)$.*

The idea is to define $E_{\mathcal{F}}$ so that the matrix $K_{\mathcal{F}}$ of (4.2) is second-order consistent. Once $E_{\mathcal{F}}$ has been defined, the full matrix E is given by

$$E = P_{\mathcal{F}} E_{\mathcal{F}} P_{\mathcal{F}}^T. \tag{4.3}$$

A suitable modification $E_{\mathcal{F}}$ may be based on some variant of the symmetric indefinite factorization of the matrix

$$\begin{pmatrix} H_{\mathcal{F}} & J_{\mathcal{F}}^T \\ J_{\mathcal{F}} & -\mu I \end{pmatrix}. \tag{4.4}$$

Appropriate methods include the following: (i) the inertia controlling LBL^T factorization (Forsgren [17], Forsgren and Gill [18]); (ii) an LBL^T factorization with pivot modification (Gould [28]); (iii) tile preordering in conjunction with pivot modification (Gill and Wong [26]); and (iv) a conventional LBL^T factorization of $H_{\mathcal{F}} + \sigma I_{\mathcal{F}}$ for some nonnegative scalar σ (Wächter and Biegler [50]). In each case, the modification E is zero if the matrix (4.4) is second-order consistent.

The following lemma establishes the main property of the matrix $H + E$.

LEMMA 4.2. *Given the matrix $K_{\mathcal{F}}$ of* (4.2)*, let $\bar{H}_{\mathcal{F}}$ denote the symmetric matrix $H_{\mathcal{F}} + E_{\mathcal{F}}$ with $E_{\mathcal{F}}$ positive semidefinite. If $K_{\mathcal{F}}$ is second-order consistent, then the matrix*

$$\begin{pmatrix} \bar{H}_{\mathcal{F}} + \frac{1}{\mu}(1+\nu) J_{\mathcal{F}}^T J_{\mathcal{F}} & \nu J_{\mathcal{F}}^T \\ \nu J_{\mathcal{F}} & \nu\mu I \end{pmatrix} \tag{4.5}$$

is positive definite for $\nu > 0$ and positive semidefinite for $\nu = 0$.

Proof. Since K_F is second-order consistent, it follows from part (ii) of Lemma 2.2 with the choice $H = \bar{H}_F$ and $J = J_F$ that $\bar{H}_F + (1/\mu)J_F^T J_F$ is positive definite. The result now follows from part (i) of Lemma 2.2. ∎

The second modification D is a positive-semidefinite diagonal matrix defined in terms of the gradients of the constraints with indices in $\mathcal{A}(x)$. In particular,

$$D = \frac{1}{\mu_A} P_A P_A^T, \tag{4.6}$$

where μ_A is a small positive scalar. The properties of the matrix $\widehat{H} = H + E + D$ are established in Theorem 4.5 below, which requires two auxiliary results.

LEMMA 4.3. *Let K_A denote the matrix*

$$K_A = \begin{pmatrix} \bar{H} + \frac{1}{\mu}(1+\nu)J^T J & \nu J^T & P_A \\ \nu J & \nu\mu I & 0 \\ P_A^T & 0 & 0 \end{pmatrix}, \quad \textit{where } \bar{H} = H + E,$$

and the n_A rows of P_A^T comprise the gradients of the bounds in $\mathcal{A}(x)$. Then

$$\mathrm{In}(K_A) = (n_A, n_A, 0) + \mathrm{In}\begin{pmatrix} \bar{H}_F + \frac{1}{\mu}(1+\nu)J_F^T J_F & \nu J_F^T \\ \nu J_F & \nu\mu I \end{pmatrix}.$$

Proof. Applying the column permutation $P = \begin{pmatrix} P_F & P_A \end{pmatrix}$ yields

$$\begin{pmatrix} P^T & 0 & 0 \\ 0 & I_m & 0 \\ 0 & 0 & I_A \end{pmatrix} K_A \begin{pmatrix} P & 0 & 0 \\ 0 & I_m & 0 \\ 0 & 0 & I_A \end{pmatrix} = \begin{pmatrix} \widetilde{H}_F & \widetilde{H}_O & \nu J_F^T & 0 \\ \widetilde{H}_O & \widetilde{H}_A & \nu J_A^T & I_A \\ \nu J_F & \nu J_A & \nu\mu I & 0 \\ 0 & I_A & 0 & 0 \end{pmatrix}, \tag{4.7}$$

where $\widetilde{H}_F$, $\widetilde{H}_A$, and $\widetilde{H}_O$ are the diagonal and off-diagonal blocks of the partition

$$\widetilde{H} \equiv P^T\Big(\bar{H} + \frac{1}{\mu}(1+\nu)J^T J\Big)P = \begin{pmatrix} \widetilde{H}_F & \widetilde{H}_O \\ \widetilde{H}_O^T & \widetilde{H}_A \end{pmatrix}.$$

Note that $\widetilde{H}_F = \bar{H}_F + \frac{1}{\mu}(1+\nu)J_F^T J_F$. The matrix of (4.7) is similar (via symmetric permutations) to

$$\begin{pmatrix} 0 & I_A & 0 & 0 \\ I_A & \widetilde{H}_A & \widetilde{H}_O^T & \nu J_A^T \\ 0 & \widetilde{H}_O & \widetilde{H}_F & \nu J_F^T \\ 0 & \nu J_A & \nu J_F & \nu\mu I \end{pmatrix} = L \begin{pmatrix} 0 & I_A & 0 & 0 \\ I_A & \widetilde{H}_A & 0 & 0 \\ 0 & 0 & \widetilde{H}_F & \nu J_F^T \\ 0 & 0 & \nu J_F & \nu\mu I \end{pmatrix} L^T,$$

where L is the nonsingular matrix

$$L = \begin{pmatrix} I_A & 0 & 0 & 0 \\ 0 & I_A & 0 & 0 \\ \widetilde{H}_O & 0 & I_F & 0 \\ \nu J_A & 0 & 0 & I \end{pmatrix}.$$

Sylvester's law of inertia gives

$$\begin{aligned} \mathrm{In}(K_A) &= \mathrm{In}\begin{pmatrix} 0 & I_A \\ I_A & \widetilde{H}_A \end{pmatrix} + \mathrm{In}\begin{pmatrix} \widetilde{H}_F & \nu J_F^T \\ \nu J_F & \nu\mu I \end{pmatrix} \\ &= (n_A, n_A, 0) + \mathrm{In}\begin{pmatrix} \bar{H}_F + \frac{1}{\mu}(1+\nu)J_F^T J_F & \nu J_F^T \\ \nu J_F & \nu\mu I \end{pmatrix}, \end{aligned}$$

as required. ∎

For a proof of the following lemma, see, e.g., Gill and Robinson [23, Theorem 3.1].

LEMMA 4.4. *Let G be a symmetric matrix of order r. Let C be an $m \times r$ matrix with rank m. If the matrix $\left(\begin{smallmatrix} G & C^T \\ C & 0 \end{smallmatrix}\right)$ has inertia $(r, m, 0)$, then $G + \frac{1}{\mu} C^T C$ is positive definite for all $\mu > 0$ sufficiently small.*

THEOREM 4.5. *If the KKT matrix $K_{\scriptscriptstyle F}$ (4.2) is second-order consistent, then the matrix*

$$B^\nu = \begin{pmatrix} \widehat{H} + \frac{1}{\mu}(1+\nu) J^T J & \nu J^T \\ \nu J & \nu\mu I \end{pmatrix}, \quad \textit{with} \quad \widehat{H} = H + E + D \ \textit{ and } \nu > 0,$$

is positive definite for all sufficiently small positive $\mu_{\scriptscriptstyle A}$, where D is defined by (4.6).

Proof. Consider the matrix

$$K_{\scriptscriptstyle A} = \begin{pmatrix} \bar{H} + \frac{1}{\mu}(1+\nu) J^T J & \nu J^T & P_{\scriptscriptstyle A} \\ \nu J & \nu\mu I & 0 \\ P_{\scriptscriptstyle A}^T & 0 & 0 \end{pmatrix}, \ \text{where } \bar{H} = H + E,$$

and $P_{\scriptscriptstyle A}$ contains the unit vectors associated with the active bounds. Then it follows from Lemmas 4.3 and 4.2 that

$$\begin{aligned} \operatorname{In}(K_{\scriptscriptstyle A}) &= (n_{\scriptscriptstyle A}, n_{\scriptscriptstyle A}, 0) + \operatorname{In} \begin{pmatrix} \bar{H}_{\scriptscriptstyle F} + \frac{1}{\mu}(1+\nu) J_{\scriptscriptstyle F}^T J_{\scriptscriptstyle F} & \nu J_{\scriptscriptstyle F}^T \\ \nu J_{\scriptscriptstyle F} & \nu\mu I \end{pmatrix} \\ &= (n_{\scriptscriptstyle A}, n_{\scriptscriptstyle A}, 0) + (n_{\scriptscriptstyle F} + m, 0, 0) \\ &= (n + m, n_{\scriptscriptstyle A}, 0). \end{aligned}$$

This identity implies that $K_{\scriptscriptstyle A}$ satisfies the conditions of Lemma 4.4 with

$$G = \begin{pmatrix} \bar{H} + \frac{1}{\mu}(1+\nu) J^T J & \nu J^T \\ \nu J & \nu\mu I \end{pmatrix} \quad \text{and} \quad C^T = \begin{pmatrix} P_{\scriptscriptstyle A} \\ 0 \end{pmatrix},$$

in which case the matrix

$$G + \frac{1}{\mu_{\scriptscriptstyle A}} C^T C = \begin{pmatrix} \bar{H} + \frac{1}{\mu}(1+\nu) J^T J & \nu J^T \\ \nu J & \nu\mu I \end{pmatrix} + \frac{1}{\mu_{\scriptscriptstyle A}} \begin{pmatrix} P_{\scriptscriptstyle A} \\ 0 \end{pmatrix} \begin{pmatrix} P_{\scriptscriptstyle A}^T & 0 \end{pmatrix} = B^\nu$$

is positive definite for $\mu_{\scriptscriptstyle A} > 0$ sufficiently small, which completes the proof. ∎

Theorem 4.5 implies that

$$B^\nu = \begin{pmatrix} \widehat{H} + \frac{1}{\mu}(1+\nu) J^T J & \nu J^T \\ \nu J & \nu\mu I \end{pmatrix}, \quad \text{with} \quad \widehat{H} = H + E + D,$$

is an appropriate positive-definite Hessian for the convexified QP. The final result of this section shows that if the final QP active set is the same as the active set at x_k, and the perturbed KKT matrix defined with the Hessian $H(x_k, y_k)$ and Jacobian $J(x_k)$ is second-order consistent (which will hold near a solution satisfying the strong second-order conditions for optimality), then the QP step is the solution of the QP subproblem defined in terms of $H(x_k, y_k)$ and $J(x_k)$.

THEOREM 4.6. *Let $\widehat{v}_k = (\widehat{x}_k, \widehat{y}_k)$ denote the unique solution of the QP subproblem* (2.8) *defined at $v_k = (x_k, y_k)$. Let $H_{\scriptscriptstyle F}$ denote the matrix of rows and columns of $H(x_k, y_k)$ corresponding to the set of free variables $\mathcal{F}(x_k)$. Similarly, let $J_{\scriptscriptstyle F}$ denote the matrix of free columns of $J(x_k)$. If the matrix*

$$\begin{pmatrix} H_{\scriptscriptstyle F} & J_{\scriptscriptstyle F}^T \\ J_{\scriptscriptstyle F} & -\mu_k^{\scriptscriptstyle R} I \end{pmatrix} \tag{4.8}$$

is second-order consistent, and $\mathcal{A}(\widehat{x}_k) = \mathcal{A}(x_k)$*, then* $(\widehat{x}_k, \widehat{y}_k)$ *satisfies the perturbed Newton equations*

$$\begin{pmatrix} H_{\scriptscriptstyle F} & J_{\scriptscriptstyle F}^T \\ J_{\scriptscriptstyle F} & -\mu_k^R I \end{pmatrix} \begin{pmatrix} [\,\widehat{x}_k - x_k\,]_{\scriptscriptstyle F} \\ -(\widehat{y}_k - y_k) \end{pmatrix} = - \begin{pmatrix} [\,g(x_k) - J(x_k)^T y_k\,]_{\scriptscriptstyle F} \\ c(x_k) + \mu_k^R (y_k - y_k^E) \end{pmatrix}.$$

Proof. As $\mathcal{A}(\widehat{x}_k) = \mathcal{A}(x_k)$, it follows from (2.10) that the QP solution $(\widehat{x}_k, \widehat{y}_k)$ satisfies the equations

$$\begin{pmatrix} \widehat{H}_{\scriptscriptstyle F} & J_{\scriptscriptstyle F}^T \\ J_{\scriptscriptstyle F} & -\mu_k^R I \end{pmatrix} \begin{pmatrix} [\,\widehat{x}_k - x_k\,]_{\scriptscriptstyle F} \\ -(\widehat{y}_k - y_k) \end{pmatrix} = - \begin{pmatrix} [\,g(x_k) - J(x_k)^T y_k\,]_{\scriptscriptstyle F} \\ c(x_k) + \mu_k^R (y_k - y_k^E) \end{pmatrix}.$$

If E_k and D_k denote the modifications (4.3) and (4.6) associated with $H(x_k, y_k)$, then

$$\widehat{H}_{\scriptscriptstyle F} = P_{\scriptscriptstyle F}^T \widehat{H}(x_k, y_k) P_{\scriptscriptstyle F} = P_{\scriptscriptstyle F}^T \big(H(x_k, y_k) + E_k + D_k\big) P_{\scriptscriptstyle F} = H_{\scriptscriptstyle F} + E_{\scriptscriptstyle F}.$$

As the perturbed Newton KKT matrix (4.8) is second-order consistent, it holds that $E_{\scriptscriptstyle F} = 0$ and $\widehat{H}_{\scriptscriptstyle F} = H_{\scriptscriptstyle F}$, as required. ∎

5. Equivalence of the iterates of an active-set method. This section discusses additional connections between the bound-constrained QP subproblem (2.8) and the stabilized SQP subproblem (2.11). Result 2.1 implies that if $\nu > 0$, then the solutions of subproblems (2.8) and (2.11) are *unique* and identical. Under the conditions of Result 2.1, $\widehat{x}_k$ is a unique solution of (2.8) when $\nu = 0$, even though the solution pair $(\widehat{x}_k, \widehat{y}_k)$ is not unique. In this case, there is a particular solution pair that is identical to the unique solution of (2.11). This analysis is extended below to establish the relationship between the iterates when an active-set method is applied to each problem.

In all that follows, the indices associated with the SQP iteration are omitted and it will be assumed that the constraints of the QP involve the constraints linearized at the point $\bar{x}$. In all cases, the suffix j will be reserved for the iteration index of the QP algorithm.

5.1. An active-set method. Both the bound-constrained QP subproblem and stabilized SQP subproblem may be considered in terms of a "conventional" active-set method on a generic convex QP with constraints written in standard form. The problem format is

$$\begin{aligned} &\underset{x}{\text{minimize}} \quad \mathcal{Q}(x) = g^T(x - \bar{x}) + \tfrac{1}{2}(x - \bar{x})^T H (x - \bar{x}) \\ &\text{subject to} \quad c + A(x - \bar{x}) = 0, \quad x \geq 0, \end{aligned} \tag{5.1}$$

where $\bar{x}$, c, A, g, and H are constant. Throughout, we assume that the constraints are feasible; i.e., there exists at least one nonnegative x such that $c + A(x - \bar{x}) = 0$.

Given a feasible x_0, active-set methods generate a feasible sequence $\{x_j\}$ such that $\mathcal{Q}(x_{j+1}) \leq \mathcal{Q}(x_j)$ with $x_{j+1} = x_j + \alpha_j p_j$. Let the index sets $\mathcal{A}$ and $\mathcal{F}$ be defined as in (1.3). At the start of the jth QP iteration, given primal-dual iterates (x_j, w_j), new estimates $(x_j + p_j, w_j + q_j)$ are defined by solving a QP formed by fixing the variables with indices in $\mathcal{A}(x_j)$ and defining p_j such that $x_j + p_j$ minimizes $\mathcal{Q}(x)$ with respect to the free variables, subject to the equality constraints. With this definition, the quantities $w_j + q_j$ are the Lagrange multipliers at the minimizer $x_j + p_j$. The

components of p_j with indices in $\mathcal{A}(x_j)$ are zero, and the free components $p_{\scriptscriptstyle F} = [\, p_j \,]_{\scriptscriptstyle F}$ are determined from the equations

$$\begin{pmatrix} H_{\scriptscriptstyle F} & -A_{\scriptscriptstyle F}^T \\ A_{\scriptscriptstyle F} & 0 \end{pmatrix} \begin{pmatrix} p_{\scriptscriptstyle F} \\ q_j \end{pmatrix} = - \begin{pmatrix} [\, g + H(x_j - \bar{x}) - A^T w_j \,]_{\scriptscriptstyle F} \\ c + A(x_j - \bar{x}) \end{pmatrix}, \tag{5.2}$$

where $[\, \cdot \,]_{\scriptscriptstyle F}$ denotes the subvector of components with indices in $\mathcal{F}(x_j)$. The choice of step length α_j is based on remaining feasible with respect to the satisfied bounds. If $x_j + p_j$ is feasible, i.e., $x_j + p_j \geq 0$, then α_j will be taken as unity. Otherwise, α is set to $\alpha_{\scriptscriptstyle M}$, the largest feasible step along p_j. Finally, the iteration index j is incremented by one and the iteration is repeated.

It must be emphasized that this active-set method is not well-defined unless equations (5.2) have a solution at every (x_j, w_j).

5.2. Solution of the bound-constrained subproblem. Next, the active-set method is applied to a bound-constrained QP of the form

$$\min_{v} \ \nabla M^T (v - \bar{v}) + \tfrac{1}{2} (v - \bar{v})^T B^\nu (v - \bar{v}) \quad \text{subject to} \quad v_i \geq 0, \ i = 1, 2, \ldots, n, \tag{5.3}$$

where v is the vector of $n + m$ primal-dual variables $v = (x, y)$, $\bar{v}$ is the constant vector $\bar{v} = (\bar{x}, \bar{y})$, and

$$\nabla M = \begin{pmatrix} g - J^T\big(\pi + \nu(\pi - \bar{y})\big) \\ \nu\big(c + \mu(\bar{y} - y^{\scriptscriptstyle E})\big) \end{pmatrix}, \quad B^\nu = \begin{pmatrix} \widehat{H} + \frac{1}{\mu}(1 + \nu) J^T J & \nu J^T \\ \nu J & \nu \mu I \end{pmatrix},$$

where $\widehat{H}$ is chosen so that $\widehat{H} + \frac{1}{\mu} J^T J$ is positive definite. (See section 4 for additional details.) It follows from Lemma 2.2 that the bound constrained problem (5.3) is a convex QP that may be solved using the conventional active-set method of section 5.1. At the jth iterate $v_j = (x_j, y_j)$, the index sets of active and free variables are given by $\widehat{\mathcal{A}}(v_j)$ and $\widehat{\mathcal{F}}(v_j)$, where

$$\widehat{\mathcal{A}}(v) = \mathcal{A}(x) = \big\{\, i : x_i = 0 \,\big\} \quad \text{and} \quad \widehat{\mathcal{F}}(v) = \{1, 2, \ldots, n + m\} \setminus \widehat{\mathcal{A}}(v)$$

(cf. (1.3)). (As the dual variables are not subject to bounds, the vector of free components of any $v = (x, y)$ has the form $v_{\widehat{\scriptscriptstyle F}} = (x_{\scriptscriptstyle F}, y)$ with $x_{\scriptscriptstyle F}$ defined in terms of $\mathcal{F}$.) Given $v_j = (x_j, y_j)$, the next QP iterate is defined as $v_{j+1} = v_j + \alpha_j d_j$, where the free components of the vector $d_j = (p_j, q_j)$ satisfy the equations

$$B^\nu_{\widehat{\scriptscriptstyle F}} d_{\widehat{\scriptscriptstyle F}} = -[\, \nabla M + B^\nu (v_j - \bar{v}) \,]_{\widehat{\scriptscriptstyle F}}, \tag{5.4}$$

with $d_{\widehat{\scriptscriptstyle F}} = (p_{\scriptscriptstyle F}, q_j)$. Equations (5.4) appear to be ill-conditioned for small μ because of the $O(1/\mu)$ term in the $(1,1)$ block of the matrix B^ν. However, this ill-conditioning is superficial. The next result shows that $d_{\scriptscriptstyle F}$ may be determined by solving an equivalent nonsingular primal-dual system with conditioning dependent on that of the original problem.

THEOREM 5.1. *Consider the application of the active-set method to the bound-constrained QP* (5.3)*. Then, for every $\nu \geq 0$, the free components of the QP search direction (p_j, q_j) satisfy the nonsingular primal-dual system*

$$\begin{pmatrix} \widehat{H}_{\scriptscriptstyle F} & -J_{\scriptscriptstyle F}^T \\ J_{\scriptscriptstyle F} & \mu I \end{pmatrix} \begin{pmatrix} p_{\scriptscriptstyle F} \\ q_j \end{pmatrix} = - \begin{pmatrix} [\, g + \widehat{H}(x_j - \bar{x}) - J^T y_j \,]_{\scriptscriptstyle F} \\ c + \mu(y_j - y^{\scriptscriptstyle E}) + J(x_j - \bar{x}) \end{pmatrix}. \tag{5.5}$$

Proof. Consider the definition of the search direction when $\nu > 0$. In this case it suffices to show that the linear systems (5.4) and (5.5) are equivalent. For any positive ν, we may define the matrix

$$T = \begin{pmatrix} I & -\frac{1+\nu}{\nu\mu} J_{\scriptscriptstyle F}^T \\ 0 & \frac{1}{\nu} I_m \end{pmatrix},$$

where the identity matrix I has dimension $n_{\scriptscriptstyle F}$, the column dimension of $J_{\scriptscriptstyle F}$. The matrix T is nonsingular with $n_{\scriptscriptstyle F} + m$ rows and columns. It follows that the equations

$$T B_{\widehat{F}}^{\nu} d_{\widehat{F}} = -T[\,\nabla M + B^{\nu}(v_j - \bar{v})\,]_{\widehat{F}}$$

have the same solution as those of (5.4). The primal-dual equations (5.5) follow by direct multiplication. The nonsingularity of equations (5.5) follows from the nonsingularity of T, and the fact that B^{ν} is positive definite (as are all symmetric submatrices formed from its rows and columns).

The resulting equations (5.5) are independent of ν, but the proof above is not applicable when $\nu = 0$ because T is undefined in this case. For $\nu = 0$, the QP objective includes only the primal variables x, which implies that problem (5.3) may be written as

$$\underset{x \geq 0}{\text{minimize}} \;\; \big(g - J^T\pi\big)^T\big(x - \bar{x}\big) + \tfrac{1}{2}\big(x - \bar{x}\big)^T\left(\widehat{H} + \frac{1}{\mu} J^T J\right)\big(x - \bar{x}\big),$$

with y arbitrary. The active-set equations analogous to (5.4) are then

$$\left(\widehat{H}_{\scriptscriptstyle F} + \frac{1}{\mu} J_{\scriptscriptstyle F}^T J_{\scriptscriptstyle F}\right) p_{\scriptscriptstyle F} = -\left[g + \left(\widehat{H} + \frac{1}{\mu} J^T J\right)\big(x_j - \bar{x}\big) - J^T\pi\right]_{\scriptscriptstyle F}. \tag{5.6}$$

For *any* choice of y_j, define the m-vector q_j such that

$$q_j = -\frac{1}{\mu}\Big(J_{\scriptscriptstyle F} p_{\scriptscriptstyle F} + \mu(y_j - \pi) + J(x_j - \bar{x})\Big), \tag{5.7}$$

where $\pi = y^{\scriptscriptstyle E} - c/\mu$ (see (2.3)). Equations (5.6) and (5.7) may be combined to give equations $K d_{\widehat{F}} = -r$, where $d_{\widehat{F}} = (p_{\scriptscriptstyle F}, q_j)$,

$$K = \begin{pmatrix} \widehat{H}_{\scriptscriptstyle F} + \frac{2}{\mu} J_{\scriptscriptstyle F}^T J_{\scriptscriptstyle F} & J_{\scriptscriptstyle F}^T \\ J_{\scriptscriptstyle F} & \mu I \end{pmatrix},$$

and the right-hand side is

$$r = \begin{pmatrix} [\,g + \widehat{H}(x_j - \bar{x})\,]_{\scriptscriptstyle F} + \frac{2}{\mu} J_{\scriptscriptstyle F}^T J(x_j - \bar{x}) - J_{\scriptscriptstyle F}^T y_j + 2 J_{\scriptscriptstyle F}^T (y_j - \pi) \\ \mu(y_j - \pi) + J(x_j - \bar{x}) \end{pmatrix}.$$

Forming the equations $TK d_{\widehat{F}} = -Tr$, where T is the nonsingular matrix

$$T = \begin{pmatrix} I & -\frac{2}{\mu} J_{\scriptscriptstyle F}^T \\ 0 & I_m \end{pmatrix},$$

gives the equivalent system

$$\begin{pmatrix} \widehat{H}_{\scriptscriptstyle F} & -J_{\scriptscriptstyle F}^T \\ J_{\scriptscriptstyle F} & \mu I \end{pmatrix} \begin{pmatrix} p_{\scriptscriptstyle F} \\ q_j \end{pmatrix} = -\begin{pmatrix} [\,g + \widehat{H}(x_j - \bar{x}) - J^T y_j\,]_{\scriptscriptstyle F} \\ c + \mu(y_j - y^{\scriptscriptstyle E}) + J(x_j - \bar{x}) \end{pmatrix},$$

which is identical to system (5.5). ∎

5.3. Solution of the stabilized SQP subproblem. Consider the application of the conventional active-set method of section 5.1 to the stabilized QP:

$$\begin{array}{ll} \underset{x,y}{\text{minimize}} & g^T(x-\bar{x}) + \frac{1}{2}(x-\bar{x})^T\widehat{H}(x-\bar{x}) + \frac{1}{2}\mu\|y\|^2 \\ \text{subject to} & c + J(x-\bar{x}) + \mu(y-y^{E}) = 0, \quad x \geq 0. \end{array} \tag{5.8}$$

In terms of the data "$(x, \bar{x}, H, g, A, c)$" for the generic QP (5.1), the variables are "x" $= (x, y)$, with "$\bar{x}$" $= (\bar{x}, \bar{y})$,

$$\text{"}H\text{"} = \begin{pmatrix} \widehat{H} & 0 \\ 0 & \mu I \end{pmatrix}, \quad \text{"}g\text{"} = \begin{pmatrix} g \\ \mu\bar{y} \end{pmatrix}, \quad \text{"}A\text{"} = \begin{pmatrix} J & \mu I \end{pmatrix}, \quad \text{and} \quad \text{"}c\text{"} = c + \mu(\bar{y} - y^{E}).$$

(The discussion of the properties of the stabilized QP relative to the generic form (5.1) is not affected by the nonnegativity constraints being applied to only a subset of the variables in (5.8).) After some simplification, the equations analogous to (5.2) may be written as

$$\begin{pmatrix} \widehat{H}_{F} & 0 & -J_{F}^T \\ 0 & \mu I & -\mu I \\ J_{F} & \mu I & 0 \end{pmatrix} \begin{pmatrix} p_{F} \\ \bar{p}_{F} \\ q_j \end{pmatrix} = - \begin{pmatrix} [\, g + \widehat{H}(x_j - \bar{x}) - J^T w_j \,]_{F} \\ \mu y_j - \mu w_j \\ c + \mu(y_j - y^{E}) + J(x_j - \bar{x}) \end{pmatrix}, \tag{5.9}$$

where p_{F} and $\bar{p}_{F}$ denote the free components of the search directions for the x and y variables, respectively. (Observe that the right-hand side of (5.9) is independent of $\bar{y}$.) The second block of equations gives $\bar{p}_{F} = q_j - y_j + w_j$, which implies that

$$y_{j+1} = y_j + \bar{p}_{F} = y_j + q_j - y_j + w_j = w_j + q_j = w_{j+1},$$

so that the primal y-variables and dual variables of the stabilized QP are identical.

Similarly, substituting for $\bar{p}_{F}$ in the third block of equations in (5.10) and using the primal-dual equivalence $w_j = y_j$ gives

$$\begin{pmatrix} \widehat{H}_{F} & -J_{F}^T \\ J_{F} & \mu I \end{pmatrix} \begin{pmatrix} p_{F} \\ q_j \end{pmatrix} = - \begin{pmatrix} [\, g + \widehat{H}(x_j - \bar{x}) - J^T y_j \,]_{F} \\ c + \mu(y_j - y^{E}) + J(x_j - \bar{x}) \end{pmatrix}, \tag{5.10}$$

which are identical to the equations associated with those for the QP subproblem (5.3).

The preceding discussion constitutes a proof of the following result.

THEOREM 5.2. *Consider the application of the active-set method to the bound-constrained QP* (5.3) *and stabilized QP* (5.8) *defined with the same quantities c, g, J, and $\widehat{H}$. Consider any x_0 and y_0 such that (x_0, y_0) is feasible for the stabilized QP* (5.8). *Then, for every $\nu \geq 0$, the active-set method generates identical primal-dual iterates $\{(x_j, y_j)\}_{j \geq 0}$.* ∎

6. Numerical results. Numerical results from a simple MATLAB implementation of `pdSQP` (Algorithm 2.1) were obtained for 158 problems from the CUTEr test collection (see Bongartz et al. [2] and Gould, Orban, and Toint [29]). The problems tested were 111 of the 126 nonlinearly constrained problems from the Hock and Schittkowski [35] test suite and the 47 CUTEr equality-constrained problems considered in [23]. A total of 15 problems were excluded from the Hock and Schittkowski set. Problems `hs1`, `hs2`, `hs3`, `hs3mod`, `hs4`, `hs5`, `hs25`, `hs38`, `hs45`, and `hs110` have no general constraints; problems `hs67`, `hs85`, and `hs87` are nonsmooth; the objective

function of `hs84` is unbounded in the feasible region; and problem `hs99exp` is poorly scaled.

Each CUTEr problem may be written in the form

$$\underset{x}{\text{minimize}} \;\; f(x) \quad \text{subject to} \quad \begin{pmatrix} x_l \\ c_l \end{pmatrix} \leq \begin{pmatrix} x \\ C(x) \end{pmatrix} \leq \begin{pmatrix} x_u \\ c_u \end{pmatrix}, \tag{6.1}$$

where $C : \mathbb{R}^n \mapsto \mathbb{R}^m$, $f : \mathbb{R}^n \mapsto \mathbb{R}$, and (x_l, c_l) and (x_u, c_u) are constant vectors of lower and upper bounds. In this format, a fixed variable or equality constraint has the same value for its upper and lower bounds. For `pdSQP`, each problem was converted to the equivalent form

$$\underset{x,s}{\text{minimize}} \;\; f(x) \quad \text{subject to} \quad c(x,s) = C(x) - s = 0, \quad \begin{pmatrix} x_l \\ c_l \end{pmatrix} \leq \begin{pmatrix} x \\ s \end{pmatrix} \leq \begin{pmatrix} x_u \\ c_u \end{pmatrix}, \tag{6.2}$$

where s is a vector of slack variables. With this formulation, the bound-constrained and stabilized QP subproblems involve simple upper and lower bounds instead of non-negativity constraints. The MATLAB implementation was initialized with parameter values given in Table 1. The primal-dual vector (x_0, y_0) was the default values supplied by CUTEr, although the code immediately projects x_0 to ensure feasibility with respect to the simple bounds on x. The sequence of iterates was terminated at a point satisfying the condition

$$\|r_{\text{opt}}(v_k)\|_\infty < \tau_{\text{stop}}, \tag{6.3}$$

where $r_{\text{opt}}(v_k)$ is the optimality measure (2.23) defined in terms of the problem format (6.2).

TABLE 1
Control parameters and initial values for Algorithm `pdSQP`.

Parameter	Value	Parameter	Value	Parameter	Value
ν	`1.0`	$\alpha_{\min}$	`1.0e-3`	μ_0^R	`1.0e-4`
$y_{\max}$	`1.0e+6`	τ_{stop}	`1.0e-6`	μ_0	`1.0`
η_S	`1.0e-2`	$k_{\max}$	`600`	τ_0	`1.0e-2`
η_D	`1.0e-3`	β	`1.0e-5`	$\phi_V^{\max}, \phi_O^{\max}$	`1.0e+3`

Algorithm `pdSQP` solves the bound-constrained QP subproblem (2.8) using a MATLAB version of the inertia-controlling QP solver of Gill and Wong [27], which solves a sequence of regularized KKT systems of the form (2.9). At the start of each subproblem, an initial QP working set is defined that contains the indices of all the fixed variables and variables within $\epsilon_A = 10^{-6}$ of their upper or lower bounds. This set is used to compute the matrix H_F and the (possibly) modified Hessian $\widehat{H} = H + E + D$ used in (2.6). In some situations, it is possible to guarantee that both E and D are zero; i.e., no convexification is needed. If the problem (6.1) is a convex program, then it can be shown that the matrix B^ν is positive definite with $\widehat{H} = H$ (see Kungurtsev [38]). Similarly, the elements of H corresponding to fixed variables need not be modified because the corresponding elements of d are always zero. This implies that $D \equiv 0$ for problems with only equality constraints (i.e., problems such that $c_l = c_u$ for the bound constraints of (6.2)). In `pdSQP`, the matrix $\widehat{H}$ is defined as described in section 4, where the matrix E_F of (4.2) is determined using the inertia-controlling LBL^T

factorization (see [17, 18] for more details). For the calculation of the modification $D = (1/\mu_A)P_A P_A^T$, we consider the properties of the matrix

$$K_{FA} = \begin{pmatrix} H_{FA} + E_{FA} & J_{FA}^T \\ J_{FA} & -\mu_A I_m \end{pmatrix},$$

where H_{FA} and E_{FA} denote the $n_{FA} \times n_{FA}$ rows and columns of H and E associated with the union of the index sets of free variables and active bound constraints (i.e., fixed variables are not included). A similar definition holds for the matrix J_{FA} formed from a subset of the columns of J. Given the value

$$\sigma = \min\big(10^{-1}, 1/\max(1, \|E\|)\big),$$

the parameter μ_A is then the first member of the sequence σ, $\sigma/10$, $\sigma/10^2$, ..., for which K_{FA} has inertia $(n_{FA}, m, 0)$. (Theorem 4.5 implies that this sequence must terminate.) This method is clearly impractical for any serious implementation. The discussion of more efficient methods that compute the modified matrix during the solution of the QP subproblem are beyond the scope of this paper (see Kungurtsev [38]). We recognize that other definitions of $\widehat{H}$ are possible, including a positive-definite quasi-Newton approximations based on the BFGS update (see, e.g., [21, 43]).

Detailed results of running the MATLAB `pdSQP` on the 158 CUTEr test problems may be found in Gill and Robinson [24]. These results include a problem-by-problem listing of the percentage number of iterations for which the Hessian modifications E and D of (4.1) were nonzero. In addition, the results give the percentages of V- or O-iterates and F-iterates required for each problem (see Algorithm 2.1), as well as the total number of M-iterates. (M-iterates generally constitute significantly less than 1% of the total iterations.)

For brevity, only a summary of the main results is presented here. The MATLAB `pdSQP` was able to satisfy the optimality measure for 150 of the 158 test problems. A run was considered to have failed if the optimality condition (6.3) could not be satisfied in $k_{\max} = 600$ iterations. Of the 8 "failures," the six problems `dixchlng`, `hs106`, `hs109`, `hs116`, `lukvle6`, and `lukvle14` terminated at infeasible local minimizers of the merit function, as measured by the optimality conditions for the infeasibility problem

$$\underset{x,s}{\text{minimize}} \ \tfrac{1}{2}\|c(x,s)\|^2 \quad \text{subject to} \quad \begin{pmatrix} x_l \\ c_l \end{pmatrix} \le \begin{pmatrix} x \\ s \end{pmatrix} \le \begin{pmatrix} x_u \\ c_u \end{pmatrix}.$$

The final point for problem `lukvle8` satisfied the complementarity measure but gave a maximum constraint violation of 1.1×10^{-4}. This problem can be solved successfully in 771 iterations, of which 22% are V- or O-iterates, 76% are F-iterates, and 8 iterations are M-iterates. Problem `mss1` can be solved in 2872 iterations, with F-iterates forming 92% of the total iterations. In this case, 99% of the iterations required some form of convexification.

Overall, the results indicate that the algorithm of `pdSQP` is robust, and that the primal-dual augmented Lagrangian provides an effective way of ensuring global convergence. For the 150 problems that were solved successfully, a grand total of 94% of the iterations computed a V- or O-iterate, and 5% of the iterations computed an F-iterate. An M-iterate was computed in only 26 of the iterations needed to solve the 150 problems.

The results also illustrate the importance of an effective convexification strategy. Overall, 77 of the 158 problems required some form of convexification. Of these 77 problems, a grand total of 44% of the iterations required the computation of a nonzero E, and 23% of the iterations required a nonzero D.

7. Summary and future work. This paper considers the formulation and analysis of an SQP method for solving general nonlinear optimization problems. An algorithm is proposed that combines the favorable properties of augmented Lagrangian methods, conventional SQP methods, and stabilized SQP methods. Numerical results given in section 6 for a simple MATLAB implementation indicate that the proposed method is robust and often exhibits fast local convergence. However, further analysis is required to determine the conditions under which the method inherits the superlinear convergence rate associated with stabilized SQP on degenerate problems.

The use of exact second derivatives presents a significant challenge to the formulation of robust and efficient SQP methods. A key contribution of this paper is a convexification procedure that provides a convex QP subproblem based on exact second derivatives. This approach provides a first step towards an effective and efficient way of incorporating exact second derivatives in the inequality-constrained QP subproblem of an SQP method.

One possible extension of the method is the use of additional regularization in the form of *explicit* bounds on the dual variables in the QP subproblem. For reasons of brevity, this refinement is not considered here. However, explicit temporary bounds on the dual variables are easily incorporated in the primal-dual QP subproblem (see, e.g., Robinson [47] and Gill and Robinson [23]). The formulation of improved update strategies for the regularization parameter μ^R is the focus of current research. It is anticipated that such strategies will allow the use of projected gradient methods for the computation of an *approximate* solution of each QP subproblem when far from a solution (see Friedlander and Leyffer [20]). Approaches such as this should allow future implementations to solve problems that are larger than those that can be solved by current SQP methods. In addition, reliable techniques that allow the rapid decrease of μ^R near a solution should give superlinear convergence under the standard assumptions.

Acknowledgments. The authors would like to thank Vyacheslav Kungurtsev for numerous discussions during the preparation of this paper. We are also grateful to the referees for constructive comments that significantly improved the presentation.

REFERENCES

[1] P. T. Boggs and J. W. Tolle, *Sequential quadratic programming*, Acta Numer., 4 (1995), pp. 1–51.

[2] I. Bongartz, A. R. Conn, N. I. M. Gould, and Ph. L. Toint, *CUTE: Constrained and unconstrained testing environment*, ACM Trans. Math. Software, 21 (1995), pp. 123–160.

[3] R. Byrd, J. Nocedal, R. Waltz, and Y. Wu, *On the use of piecewise linear models in nonlinear programming*, Math. Program., 137 (2013), pp. 289–324.

[4] A. R. Conn, N. Gould, A. Sartenaer, and Ph. L. Toint, *Convergence properties of an augmented Lagrangian algorithm for optimization with a combination of general equality and linear constraints*, SIAM J. Optim., 6 (1996), pp. 674–703.

[5] A. R. Conn, N. I. M. Gould, and Ph. L. Toint, *Testing a class of methods for solving minimization problems with simple bounds on the variables*, Math. Comp., 50 (1988), pp. 399–430.

[6] A. R. Conn, N. I. M. Gould, and Ph. L. Toint, *Global convergence of a class of trust region algorithms for optimization with simple bounds*, SIAM J. Numer. Anal., 25 (1988), pp. 433–460.

[7] A. R. Conn, N. I. M. Gould, and Ph. L. Toint, *A Comprehensive Description of LANCELOT*, Technical Report 91/10, Département de Mathématique, Facultés Universitaires de Namur, 1991.

[8] A. R. Conn, N. I. M. Gould, and Ph. L. Toint, *A globally convergent augmented Lagrangian algorithm for optimization with general constraints and simple bounds*, SIAM J. Numer. Anal., 28 (1991), pp. 545–572.

[9] A. R. Conn, N. I. M. Gould, and Ph. L. Toint, *LANCELOT: A Fortran Package for Large-Scale Nonlinear Optimization (Release A)*, Lecture Notes in Comput. Math. 17, Springer-Verlag, Berlin, Heidelberg, New York, London, Paris, Tokyo, 1992.

[10] A. R. Conn, N. I. M. Gould, and Ph. L. Toint, *Large-Scale Nonlinear Constrained Optimization*, Technical Report 92/2, Département de Mathématique, Facultés Universitaires de Namur, 1992.

[11] L. B. Contesse, *Une caractérisation complète des minima locaux en programmation quadratique*, Numer. Math., 34 (1980), pp. 315–332.

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

[13] D. Fernández and M. Solodov, *Local Convergence of Exact and Inexact Augmented Lagrangian Methods under the Second-Order Sufficiency Condition*, IMPA preprint A677, IMPA, Rio de Janeiro, Brazil, 2010.

[14] 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 (2010), pp. 47–73.

[15] R. Fletcher, *An ℓ_1 penalty method for nonlinear constraints*, in Numerical Optimization 1984, P. T. Boggs, R. H. Byrd, and R. B. Schnabel, eds., SIAM, Philadelphia, 1985, pp. 26–40.

[16] R. Fletcher and S. Leyffer, *User Manual for Filter SQP*, Technical Report NA/181, Department of Mathematics, University of Dundee, Scotland, 1998.

[17] A. Forsgren, *Inertia-controlling factorizations for optimization algorithms*, Appl. Numer. Math., 43 (2002), pp. 91–107.

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

[19] A. L. Forsgren, P. E. Gill, and W. Murray, *On the identification of local minimizers in inertia-controlling methods for quadratic programming*, SIAM J. Matrix Anal. Appl., 12 (1991), pp. 730–746.

[20] M. P. Friedlander and S. Leyffer, *Global and finite termination of a two-phase augmented Lagrangian filter method for general quadratic programs*, SIAM J. Sci. Comput., 30 (2008), pp. 1706–1729.

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

[22] P. E. Gill, W. Murray, M. A. Saunders, and M. H. Wright, *Some theoretical properties of an augmented Lagrangian merit function*, in Advances in Optimization and Parallel Computing, P. M. Pardalos, ed., North-Holland, Amsterdam, 1992, pp. 101–128.

[23] P. E. Gill and D. P. Robinson, *A primal-dual augmented Lagrangian*, Comput. Optim. Appl., 51 (2012), pp. 1–25.

[24] P. E. Gill and D. P. Robinson, *A Globally Convergent Stabilized SQP Method*, Report 13-2, Center for Computational Mathematics, University of California, San Diego, La Jolla, CA, 2013.

[25] P. E. Gill and E. Wong, *Sequential quadratic programming methods*, in Mixed Integer Nonlinear Programming, J. Lee and S. Leyffer, eds., IMA Vol. Math. Appl. 154, Springer, New York, 2012, pp. 147–224.

[26] P. E. Gill and E. Wong, *Methods for Convex and General Quadratic Programming*, Report 13-1, Center for Computational Mathematics, University of California, San Diego, La Jolla, CA, 2013.

[27] P. E. Gill and E. Wong, *A Regularized Method for Convex and General Quadratic Programming*, Report 13-3, Center for Computational Mathematics Report, Department of Mathematics, University of California, San Diego, La Jolla, CA, 2013.

[28] N. I. M. Gould, *On modified factorizations for large-scale linearly constrained optimization*, SIAM J. Optim., 9 (1999), pp. 1041–1063.

[29] N. I. M. Gould, D. Orban, and Ph. L. Toint, *CUTEr and SifDec: A constrained and unconstrained testing environment, revisited*, ACM Trans. Math. Software, 29 (2003), pp. 373–394.

[30] N. I. M. Gould and D. P. Robinson, *A second derivative SQP method: Global convergence*, SIAM J. Optim., 20 (2010), pp. 2023–2048.

[31] N. I. M. Gould and D. P. Robinson, *A second derivative SQP method: Local convergence and practical issues*, SIAM J. Optim., 20 (2010), pp. 2049–2079.
[32] W. W. Hager, *Stabilized sequential quadratic programming: Computational optimization—a tribute to Olvi Mangasarian, Part* I, Comput. Optim. Appl., 12 (1999), pp. 253–273.
[33] S. P. Han, *Superlinearly convergent variable metric algorithms for general nonlinear programming problems*, Math. Programming, 11 (1976/1977), pp. 263–282.
[34] M. R. Hestenes, *Multiplier and gradient methods*, J. Optim. Theory Appl., 4 (1969), pp. 303–320.
[35] W. Hock and K. Schittkowski, *Test Examples for Nonlinear Programming Codes*, Lecture Notes in Econom. and Math. Systems 187, Springer-Verlag, Berlin, Heidelberg, New York, 1981.
[36] 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. Ser. A, 126 (2011), pp. 231–257.
[37] A. F. Izmailov and M. V. Solodov, *Stabilized SQP revisited*, Math. Program., 133 (2012), pp. 93–120.
[38] V. Kungurtsev, *Second-Derivative Sequential Quadratic Programming Methods for Nonlinear Optimization*, Ph.D. thesis, Department of Mathematics, University of California, San Diego, La Jolla, CA, 2013.
[39] D.-H. Li and L. Qi, *A Stabilized SQP Method via Linear Equations*, Technical Report AMR00/5, School of Mathematics, University of New South Wales, Sydney, 2000.
[40] C. M. Maes, *A Regularized Active-Set Method for Sparse Convex Quadratic Programming*, Ph.D. Thesis, Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA, 2010.
[41] N. Maratos, *Exact Penalty Function Algorithms for Finite-Dimensional and Control Optimization Problems*, Ph.D. thesis, Department of Computing and Control, University of London, 1978.
[42] J. Morales, J. Nocedal, and Y. Wu, *A sequential quadratic programming algorithm with an additional equality constrained phase*, IMA J. Numer. Anal., 32 (2012), pp. 553–579.
[43] J. Nocedal and S. J. Wright, *Numerical Optimization*, Springer-Verlag, New York, 1999.
[44] C. Oberlin and S. J. Wright, *Active set identification in nonlinear programming*, SIAM J. Optim., 17 (2006), pp. 577–605.
[45] M. J. D. Powell, *A method for nonlinear constraints in minimization problems*, in Optimization, R. Fletcher, ed., Academic Press, London, New York, 1969, pp. 283–298.
[46] M. J. D. Powell, *A Fast Algorithm for Nonlinearly Constrained Optimization Calculations*, Technical Report 77/NA 2, Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge, UK, 1977.
[47] D. P. Robinson, *Primal-Dual Methods for Nonlinear Optimization*, Ph.D. thesis, Department of Mathematics, University of California, San Diego, La Jolla, CA, 2007.
[48] K. Schittkowski, *The nonlinear programming method of Wilson, Han, and Powell with an augmented Lagrangian type line search function.* I. *Convergence analysis*, Numer. Math., 38 (1981/1982), pp. 83–114.
[49] K. Schittkowski, *The nonlinear programming method of Wilson, Han, and Powell with an augmented Lagrangian type line search function.* II. *An efficient implementation with linear least squares subproblems*, Numer. Math., 38 (1981/1982), pp. 115–127.
[50] A. Wächter and L. T. Biegler, *Line search filter methods for nonlinear programming: Motivation and global convergence*, SIAM J. Optim., 16 (2005), pp. 1–31.
[51] S. J. Wright, *Superlinear convergence of a stabilized SQP method to a degenerate solution*, Comput. Optim. Appl., 11 (1998), pp. 253–275.
[52] S. J. Wright, *Modifying SQP for degenerate problems*, SIAM J. Optim., 13 (2002), pp. 470–497.