UNIVERSITY OF CALIFORNIA, SAN DIEGO

Second-Derivative Sequential Quadratic Programming Methods for Nonlinear Optimization

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Mathematics with specialization in Computational Science

by

Vyacheslav Kungurtsev

Committee in charge:

Professor Philip E. Gill, Chair
Professor Randolph E. Bank
Professor Michel Holst
Professor Tara Javidi
Professor Michael Norman

2013

The dissertation of Vyacheslav Kungurtsev is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California, San Diego

2013

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

VITA

2003	Research Assistant at SUNY-Stony Brook
2006	Research Assistant at UCLA
2007	B.S. in Mathematics, Duke University
2007	Research Intern at Microsoft Research Asia
2007–2013	Graduate Teaching Assistant and Graduate Student Researcher, University of California, San Diego
2011	Research Intern at Argonne National Lab
2013	Ph.D. in Mathematics, University of California, San Diego

ABSTRACT OF THE DISSERTATION

Second-Derivative Sequential Quadratic Programming Methods for Nonlinear Optimization

by

Vyacheslav Kungurtsev

Doctor of Philosophy in Mathematics with specialization in Computational Science

University of California, San Diego, 2013

Professor Philip E. Gill, Chair

Sequential Quadratic Programming (SQP) methods are a popular and successful class of methods for minimizing a generally nonlinear function subject to nonlinear constraints. Under a standard set of assumptions, conventional SQP methods exhibit a fast rate of local convergence. However, in practice, a conventional SQP method involves solving an indefinite quadratic program (QP), which is NP hard in general. As a result, approximations to the second-derivatives are often used, which can slow the rate of local convergence and reduce the chance that the algorithm will converge to a local minimizer instead of a saddle point. In addition, the standard assumptions required for convergence often do not hold in practice. For such problems, regularized SQP methods, which also require second-derivatives, have been shown to have good local convergence properties; however, there are few regularized SQP methods that exhibit convergence to a minimizer from an arbitrary

initial starting point. This thesis considers the formulation, analysis and implementation of SQP methods with the following properties. (i) The solution of an indefinite QP is not required. (ii) Regularization is performed in such a way that global convergence can be established under standard assumptions. (iii) Implementations of the method work well on degenerate problems.

Chapter 1

Introduction

1.1 Introduction

This thesis concerns the formulation and analysis of methods for the solution of a smooth nonlinear optimization problem with equality and inequality constraints. Let $f\colon \mathcal{D} \subseteq \mathbb{R}^n \mapsto \mathbb{R}^1$ be a differentiable scalar-valued function of the components of an n-vector x. Similarly, let $c\colon \mathcal{D} \subseteq \mathbb{R}^n \mapsto \mathbb{R}^1$ be a differentiable vector-valued function of x. Consider the optimization problem

$$\begin{aligned} \underset{x\in\mathbb{R}^n}{\text{minimize}} \quad & f(x) \\ \text{subject to} \quad & c_i(x) = 0, \quad i \in \mathcal{E} \\ & c_i(x) \geq 0, \quad i \in \mathcal{I}, \end{aligned} \tag{1.1}$$

where $\mathcal{E}$ and $\mathcal{I}$ are non-intersecting index sets such that $\mathcal{E} \cup \mathcal{I} = \{1, 2, \ldots, m\}$. Any point x satisfying the constraints of (1.1) is called a *feasible point*, and the set of all such points is the *feasible region*. The *Lagrangian function* associated with problem (1.1) is given by

$$L(x, y) = f(x) - c(x)^T y = f(x) - \sum_{i=1}^{m} y_i c_i(x) \tag{1.2}$$

with Hessian $H(x, y) = \nabla^2_{xx} L(x, y)$. The components of the vector y are known as the *dual variables*.

Most algorithms seek to find a local minimizer of (1.1) by finding a point that satisfies the *first-order Karush-Kuhn-Tucker (KKT) conditions*. These conditions state that under certain regularity assumptions on the constraints, at a local solution x^* of (1.1), there must exist at least one m-vector y such that

$$\begin{aligned} \nabla f(x^*) - \sum_{i=1}^{m} y_i \nabla c_i(x^*) &= 0, && \\ c_i(x^*) &= 0, && i \in \mathcal{E}, \\ c_i(x^*) &\geq 0, && i \in \mathcal{I}, \\ y_i &\geq 0, && i \in \mathcal{I}, \\ y_i c_i(x^*) &= 0, && i \in \mathcal{I}. \end{aligned} \tag{1.3}$$

The components of any y satisfying these conditions are known as *Lagrange multipliers.* In general, there may be a set $\mathcal{M}_y(x^*)$ of Lagrange multipliers that satisfy the KKT conditions (1.3). A point x^* satisfying the KKT conditions (1.3) for some $y \in \mathcal{M}_y(x^*)$ is known as a *first-order KKT point.*

The first-order KKT conditions are satisfied at a local minimizer only if certain regularity conditions hold. These regularity assumptions are known as *constraint qualifications.* There are a number of alternative constraint qualifications, but almost all of them involve the properties of a subset of the constraints that includes all the equality constraints together with the inequality constraints that are satisfied exactly at x^*. The set of indices of the inequality constraints that are satisfied exactly at a point is known as the *active set.* The active set at the point x is denoted by $\mathcal{A}(x)$, i.e.,

$$\mathcal{A}(x) = \{i \mid c_i(x) = 0, \;\; i \in \mathcal{I}\}.$$

The *linear independence constraint qualification* requires that the gradients of the equality constraints and active constraints are linearly independent.

Second-order conditions for optimality also involve the properties of the active set. For any KKT point x and multiplier $y \in \mathcal{M}_y(x)$, we define the index sets

$$\mathcal{A}_0(x,y) = \{i \in \mathcal{A}(x) \mid y_i = 0\} \;\text{ and }\; \mathcal{A}_+(x,y) = \{i \in \mathcal{A}(x) \mid y_i > 0\}.$$

As $y \geq 0$, it follows that $\mathcal{A}(x) = \mathcal{A}_0(x.y) \cup \mathcal{A}_+(x,y)$ and $\mathcal{A}_0(x,y) \cap \mathcal{A}_+(x,y) = \emptyset$. Given a point (x,y) satisfying the first-order conditions, the sets $\mathcal{A}_+(x,y)$ and $\mathcal{A}_0(x,y)$ define a cone $\widetilde{\mathcal{F}}(x)$ that consists of directions emanating from x such that

$$\widetilde{\mathcal{F}}(x) = \{d \mid \nabla c_i(x)^T d = 0, i \in \mathcal{E},\; \nabla c_i(x)^T d = 0, i \in \mathcal{A}_+(x,y),\; \nabla c_i(x)^T d \geq 0, i \in \mathcal{A}_0(x,y)\}.$$

The *second-order necessary conditions* hold if there is a (x,y) satisfying the first-order conditions (1.3) such that $d^T \nabla^2_{xx} L(x,y) d \geq 0$ for all $d \in \widetilde{\mathcal{F}}(x)$. The *second-order sufficiency conditions* hold if there exists a strictly positive σ such that $d^T H(x,y) d \geq \sigma ||d||^2$ for all

$d \in \widetilde{\mathcal{F}}(x)$, where $H(x,y)$ denotes the Hessian of the Lagrangian function $L(x,y)$, i.e., $H(x,y) = \nabla^2_{xx}L(x,y)$. Some subtle variations of these conditions are discussed in Chapters 2 and 5.

Optimization problems of the form (1.1) may be solved using a *sequential quadratic programming (SQP) method*. Conventional SQP methods solve a sequence of quadratic programming (QP) subproblems defined in terms of a quadratic model of the objective function and a linearization of the constraints. Given estimates (x_k, y_k) of primal and dual variables satisfying (1.3), the QP subproblem for a conventional SQP is given by

$$\begin{aligned} \underset{x\in\mathbb{R}^n}{\text{minimize}} \quad & f(x_k) + \nabla f(x_k)^T(x - x_k) + \tfrac{1}{2}(x - x_k)^T H(x_k, y_k)(x - x_k) \\ \text{subject to} \quad & c_i(x_k) + \nabla c_i(x)^T(x - x_k) = 0, \quad i \in \mathcal{E}, \\ & c_i(x_k) + \nabla c_i(x)^T(x - x_k) \ge 0, \quad i \in \mathcal{I}. \end{aligned}$$

If the change in variables is written as $p = x - x_k$, then an equivalent QP is given by

$$\begin{aligned} \underset{p\in\mathbb{R}^n}{\text{minimize}} \quad & \nabla f(x_k)^T p + \tfrac{1}{2} p^T H(x_k, y_k) p \\ \text{subject to} \quad & c_i(x_k) + \nabla c_i(x)^T p = 0, \quad i \in \mathcal{E}, \\ & c_i(x_k) + \nabla c_i(x)^T p \ge 0, \quad i \in \mathcal{I}. \end{aligned} \tag{1.4}$$

(Note that the constant term in the objective may be omitted without changing the solution of the QP.) In the original SQP method proposed by Wilson [85], the next iterate is $x_{k+1} = x_k + p_k$, where p_k is a solution of (1.4). The new estimate y_{k+1} of the Lagrange multipliers are the Lagrange multipliers of the QP subproblem.

SQP methods use some method to ensure global convergence. This usually takes the form of a line-search, trust-region or filter to either reject a subproblem step or limit it by defining $x_{k+1} = x_k + \alpha p_k$, for some $\alpha > 0$. For example, a line-search method defines the new iterate as $x_{k+1} = x_k + \alpha_k p_k$, where α_k is a nonnegative scalar step length. In this case, α_k is chosen to reduce the value of a *merit function*, whose value provides a measure of the distance of the point $x_k + \alpha_k p_k$ to a solution x^* of the nonlinear problem.

If the Hessian $H(x_k, y_k)$ is indefinite, then the QP subproblem is non-convex. Non-convex quadratic programming is NP-hard and there may be no positive α that gives a reduction in a merit function. As a result, many practical SQP methods use a positive-definite *approximation* to the Hessian. However, the use of an approximate Hessian may have a detrimental effect on the convergence properties of an algorithm (see Chapter 3). In this thesis we consider methods that use exact second derivatives in the QP subproblem but do not require the solution of an indefinite quadratic program.

Proofs of convergence generally rely on a set of assumptions called constraint qualifications (CQs) that describe the geometry of the feasible region. Most SQP convergence proofs have relied on the *linear independence constraint qualification (LICQ)*, a relatively strong assumption. There has been some theoretical work in the development of new CQs, convergence theory and SQP methods that converge under weaker assumptions. Much of that work, however, is limited to local convergence, and relies on finding a solution of the QP subproblem, even though the subproblem may be indefinite. In this thesis, we formulate and analyze effective computational methods that are able to use exact Hessian matrices while avoiding the difficulties that are usually associated with them.

1.2 Previous Work

There are a few recent examples in the literature of second-derivative SQP methods. In most of these methods, each iteration has two parts. The first part involves the solution of a convex QP subproblem based an approximate Lagrangian Hessian. The optimal active set of this QP is then used to define the constraints of an equality-constrained QP that is defined in terms of the exact Hessian. For example, Byrd et al. [15] use a convex piecewise linear problem to estimate the optimal active set, followed by an EQP step wherein the constraints corresponding to the estimated active set are constrained to be equalities. They prove that limit points of the sequence of iterates satisfy the first-order optimality conditions, and also that the method eventually estimates the active set at the solution. (No results concerning the rate-of-convergence are given.)

Chin and Fletcher [17] estimate the active set by first solving a linear program, and then define a EQP based on the final linear programming active set. Their algorithm uses a filter to ensure global convergence. They do not assume any constraint qualifications, however, their (strictly global) convergence results are weak. In particular, the sequence of the algorithm either ends at a point not satisfying the filter and failing to generate a solvable LP, a Fritz John point, or generates an infinite sequence with all cluster points feasible and one satisfying the Fritz John conditions.

Gould and Robinson ([49], [48], [47]) analyze an SQP method with several different components, some involving convex subproblems and some using the exact Hessian. They begin with a "predictor" step wherein a standard convex QP is solved. In [48] they propose a method that uses some combination of 1) a "Cauchy step" that minimizes the merit function with the exact Hessian using the predictor step, 2) an inequality-constrained SQP step using the exact Hessian, 3) an equality-constrained SQP step using the active set estimated by the predictor step, and 4) an implicitly constrained SQP step that includes the constraints as a penalty term in the objective. They use a trust-region approach and heuristics to take the additional steps only if they demonstrate descent for the merit function. They prove global convergence for the overall framework, never requiring a full solution to a non-convex problem. Even though they provide a complete picture of how to incorporate second-derivatives in a method that solves two or more subproblems, it remains unclear as to the benefit of solving each particular additional subproblem.

In [47], Gould and Robinson analyze the local convergence of this framework. In this case they use a predictor step that is the solution of a convex QP and an accelerator step that consists of either an equality-constrained problem with the active set at the predictor solution set as equality constraints, or an inequality constrained problem with a descent constraint, both using the exact Hessian. The accelerator step here is meant to speed convergence. Using strong assumptions (LICQ and second-order sufficiency at the limit point) they are able to show a quadratic rate of convergence of the iterates. This points to the feasibility of such a scheme in terms of using second derivatives as an additional element to a conventional quasi-Newton SQP method to ensure rapid local convergence.

The methods discussed in this thesis differ from Gould and Robinson in 1) discussing global convergence to second-order optimal points 2) only solving a modified convex problem as necessary, and 3) ensuring local convergence under weaker assumptions.

Even though there are a number of papers on local convergence theory for SQP methods under weak regularity assumptions, there are few fully-developed methods in this area. Qi and Wei [78] demonstrate global and superlinear local convergence of the Panier-Tits method under weak constraint qualification assumptions. Although this thesis will reference their results in the subsequent convergence theory of the algorithm SQP2d, the Panier-Tits method [75] on which their results are based relies on feasibility of the iterates and has not performed particularly well in practice.

Regularized SQP methods are a class of modified SQP methods for which superlinear convergence has been established for degenerate problems. Wright [87] discusses the theoretical aspects of several regularized SQP methods, including the *stabilized SQP method*. This method was formulated by Wright [86] as one of a class of inexact SQP methods. The analysis is purely local, and requires obtaining the closest solution to the current iterate of the possibly indefinite SQP subproblem at each iteration. In [89], Wright presents a more practical algorithm, wherein active-constraints are identified and an equality-constrained problem is solved with the active constraints enforced as equalities. However, the analysis is still purely local, and one would still expect that, far away from the solution, numerous incorrect guesses as to the active set would require the solution of a large quantity of EQPs whose results would then be subsequently thrown out. The numerical results of this class of methods are mixed (see, e.g., Mostafa et al. [68]).

This thesis will demonstrate how to take advantage of the local convergence properties of these methods in practical implementations involving a complete algorithm with global and local convergence properties. In addition, it should be noted that much of these strong local convergence results rely on the assumption of second-order sufficiency holding at the limit point. By proving global convergence to a point satisfying second-order *necessary* conditions, this assumption can be considerably weakened to just requiring that the reduced Hessian at the solution is nonsingular, a condition generally expected to hold for almost all NLPs.

1.3 Contributions of this Thesis

In this thesis we formulate and analyze sequential quadratic programming algorithms that use second-derivatives for the definition of the quadratic programming subproblem. Two methods are considered for *convexifying* the QP subproblem. Convexification obviates the need to solve an indefinite quadratic programming subproblem. Three main SQP algorithms are presented for general (i.e., nonconvex) nonlinear optimization. Algorithm SQP2d is a conventional primal second-derivative line-search method that uses a primal-dual merit function to force global convergence. Algorithm pdSQP is a regularized method based on minimizing a bound-constrained primal-dual augmented Lagrangian function. Algorithm pdSQP2 is an extension of Algorithm pdSQP that provides convergence to points that satisfy the second-order necessary conditions for optimality. Results are presented for both methods that establish the global convergence and local superlinear convergence degenerate and nondegenerate optimization problems. Preliminary numerical results are presented using an implementation in Matlab.

1.4 Organization of the Thesis

This thesis is organized as follows. The next chapter gives a background of nonlinear programming optimality conditions and constraint qualifications. Chapter 3 discusses conventional SQP methods. Chapter 4 discusses convergence theory and stability results and provides some classical examples from the literature with regards to applications of these results for proofs of convergence of algorithms. Chapter 5 discusses regularized SQP methods, a class of theoretical SQP methods that have been proven to exhibit superlinear convergence under weak CQ assumptions. Chapter 6 discusses convexification, the primary methodology by which second derivatives are used in the algorithms discussed in this thesis. Chapter 7 provides an application of convexification and degenerate convergence theory to SQP2d, an inexact SQP method with an augmented Lagrangian merit function. Chapter 8 provides an application of convexification to a primal-dual augmented Lagrangian implementation of SQP (pdSQP) , wherein the subproblems are equivalent to the familiar regularized stabilized SQP method. Chapter 9 discusses directions of negative curvature and convergence to second-order optimal points for pdSQP.

Chapter 10 includes numerical results for the algorithm pdSQP on a set of example problems. The final chapter concludes and discusses the main results of the thesis.

1.5 Notation

1.5.1 Vectors

The i-th component of a vector labeled with a subscript will be denoted by $[\,\cdot\,]_i$, e.g., $[v_N]_i$ is the i-th component of the vector v_N. Similarly, a subvector of components with indices in the index set $\mathcal{S}$ is denoted by $(\,\cdot\,)_{\mathcal{S}}$, e.g., $(v_N)_{\mathcal{S}}$ is the vector with components $[v_N]_i$ for $i \in \mathcal{S}$. The vector e will be used to denote the vector of all ones with length determined by the context. The vector with components $\max\{-x_i, 0\}$ (i.e., the magnitude of the negative part of x) is denoted by $[\,x\,]_-$. Similarly, the magnitude of the positive part of a vector x is denoted by $[\,x\,]_+$.

Unless explicitly indicated otherwise, $\|\cdot\|$ denotes the vector two-norm or its induced matrix norm. Given vectors a and b with the same dimension, the vector with i-th component $a_i b_i$ is denoted by $a \cdot b$. The component-wise maximum of two vectors x and y is denoted by $\max(x, y)$, i.e., $[\max(x, y)]_i = x_i$ if $x_i \geq y_i$, and $[\max(x, y)]_i = y_i$ otherwise.

A set of vectors $\{v_i\}$ is said to be *positively linearly dependent* if there exists a set of scalars $\{\alpha_i\}$ such that $\alpha_i \geq 0$ for all i, there is at least one i for which $\alpha_i > 0$, and $\sum_i \alpha_i v_i = 0$.

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

Given a point x and a set $\mathcal{S}$, the distance of x to $\mathcal{S}$ is the distance of x to the closest point in $\mathcal{S}$, i.e., $\operatorname{dist}(x, \mathcal{S}) = \inf_{z\in\mathcal{S}} \|x - z\|$.

1.5.2 Matrices

The symbol I is used to denote an identity matrix with dimension determined by the context. The j-th column of I is denoted by e_j. Unless explicitly indicated otherwise, $\|\cdot\|$ denotes the vector two-norm or its induced matrix norm. The inertia of a real symmetric matrix A, denoted by $\operatorname{In}(A)$, is the integer triple (a_+, a_-, a_0) giving the number of positive, negative and zero eigenvalues of A. Given vectors a and b with the same dimension, the vector with i-th component a_ib_i is denoted by $a \cdot b$. The columns of the matrix $N(A)$ form a basis for $\operatorname{null}(A)$, the null-space of A. The column space of A is denoted by $\operatorname{range}(A)$.

The ith eigenvalue of a symmetric matrix A will be denoted by $\lambda_i(A)$, and the eigenvalues will be ordered in decreasing order $\lambda_1(A) \geq \lambda_2(A) \geq \cdots \geq \lambda_n(A)$. $\operatorname{In}(A)$ denotes the inertia of A, which is the integer triple (a_+, a_-, a_0) indicating the number of positive, negative, and zero eigenvalues of A. Finally a sequence of matrices $\{A_k\}$ is said to be bounded if the sequence of norms $\{||A_k||\}$ is bounded.

1.6 Some Useful Results

The first result appears prominently in this thesis.

Lemma 1.6.1 (Debreu's Lemma [21])**.** *Let H and A be matrices of order $n \times n$ and $m \times n$, respectively, with H symmetric. If $v^THv > 0$ for all v such that $Av = 0$, then there exists a finite ρ_c such that $H + \rho A^TA$ is positive-definite for all $\rho > \rho_c$.*

The next result is an important characterization of a class of matrices associated with second-order conditions for optimality.

Theorem 1.6.1. *Let H and A be matrices of order $n \times n$ and $m \times n$, respectively, with H symmetric. Assume that A has rank m. The matrix $\begin{pmatrix} H & A^T \\ A & 0 \end{pmatrix}$ has inertia $(n, m, 0)$ if and only if Z^THZ is positive-definite, where Z is a matrix whose columns form a basis for the null-space of A.*

Chapter 2

Optimality Conditions

This chapter provides a brief overview of optimality conditions and constraint qualification for a nonlinear optimization problem written in the form:

$$\begin{array}{lll} \underset{x}{\text{minimize}} & f(x) & \\ \text{subject to} & c_i(x) = 0, & i \in \mathcal{E} \\ & c_j(x) \geq 0, & j \in \mathcal{I}. \end{array} \tag{2.1}$$

Almost all algorithms attempt to find a point that satisfies a set of nonlinear equations, constraint qualifications provide a link between the local optimality of that point and these nonlinear equations.

2.1 First-Order Necessary Conditions

The *Fritz John conditions* hold at a point x if there exists an m-vector y and a scalar y_0 such that

$$\begin{aligned} y_0 \nabla f(x) - \sum_{i=1}^{m} y_i \nabla c_i(x) = 0, & \\ c_j(x) = 0, & \quad \text{for every } j \in \mathcal{E} \cup \mathcal{A}(x), \\ c_j(x) \geq 0, \;\; y_j \geq 0, \;\; y_j c_j(x) = 0, & \quad \text{for every } j \in \mathcal{I}, \end{aligned}$$

where $\mathcal{A}(x) \subset \mathcal{I}$ is the index set of the active inequality constraints at x. At a local minimizer of (2.1), the Fritz John conditions always hold (see John [58], and Mangasarian and Fromowitz [63]). Notice that y_0 may be zero. If $y_0 = 0$, then the Fritz John conditions are satisfied by a stationary point of the sum of the constraints. This implies that there are many problems for which there are points that are not local minimizers but satisfy the Fritz John condition.

The *Karush-Kuhn-Tucker (KKT) conditions* hold at a point x if there exists an m-vector y such that

$$\begin{aligned} \nabla f(x) - \sum_{i=1}^{m} y_i \nabla c_i(x) = 0, & \\ c_j(x) = 0, \quad & \text{for every } \ j \in \mathcal{E} \cup \mathcal{A}(x), \\ c_j(x) \geq 0, \ \ y_j \geq 0, \ \ y_j c_j(x) = 0, \quad & \text{for every } \ j \in \mathcal{I} \end{aligned} \tag{2.2}$$

(see Karush [59], and Kuhn and Tucker [61]). A point x^* satisfying the KKT conditions (2.2) is known as a *first-order KKT point*. In general, there is an infinite set of Lagrange multipliers associated with a first-order KKT point.

Definition 2.1.1. *The set of Lagrange multipliers associated with a first-order KKT point x (i.e., an x that satisfies* (2.2)*) is denoted by $\mathcal{M}_y(x)$. A pair of vectors (x, y) with $y \in \mathcal{M}_y(x)$ is called a* primal-dual first-order KKT pair*.*

The KKT conditions are just the Fritz John conditions with the value $y_0 \equiv 1$. The Fritz John conditions always hold at a local minimizer, but the first-order KKT conditions will satisfied at a local minimizer only if certain regularity conditions hold. These regularity assumptions are known as *constraint qualifications*.

Since y_0 must be nonzero at a KKT point, most algorithms seek to generate a sequence of points that converge to a point x^* that satisfies the first-order KKT conditions.

2.2 Overview of First-Order Constraint Qualifications

Constraint qualifications (CQs) are regularity assumptions about the geometry of the feasible region of a constrained problem that guarantee that the Karush-Kuhn-Tucker (KKT) conditions hold at a local minimizer. These conditions are necessary assumptions in the convergence proofs of most algorithms for constrained nonlinear programming.

Global convergence results define the conditions under which a sequence of iterates converge to a point that either satisfies the Karush-Kuhn-Tucker (KKT) conditions or fails to satisfy a constraint qualification. *Local* convergence results state that, once a point in the sequence of iterates is sufficiently close to its limit, the iterates converge at a superlinear rate or that limit fails to satisfy the constraint qualification. The weaker the conditions associated with a constraint qualification, the stronger the result of the convergence proof.

Constraint qualifications specify geometric details of the local feasible region and the constraint gradients. The condition introduced by Guignard [50] is the most general (i.e., the weakest) constraint qualification. However, it is too theoretical for use in the convergence proofs of algorithms.

The constraint qualification used in many proofs of convergence is the linear independence constraint qualification (LICQ), which is a relatively strong condition that fails to hold for a large class of problems. The LICQ has the advantage that it can be verified using a finite linear algebraic procedure. Moreover, it is possible to prove strong stability results when the LICQ holds (see Chapter 4).

The Mangasarian-Fromovitz constraint qualification (MFCQ), a weaker condition than the LICQ that involves the existence of an interior feasible direction. The satisfaction of the MFCQ is sufficient to ensure local and global and local convergence for a large class of algorithms.

The constant-rank constraint qualification (CRCQ) is weaker than the LICQ, though it is neither stronger nor weaker than the MFCQ. In this chapter it is shown that different kinds of geometric irregularities present problems for the CRCQ and the MFCQ.

The constant positive linear dependence (CPLD) condition is a constraint qualification that is weaker than both the MFCQ and the CRCQ. It was first introduced in Qi and Wei [78] and used in the context of the convergence theory of an SQP method. In that paper Qi and Wei conjecture that CPLD is a general constraint qualification and proved the strongest result for convergence of SQP methods among the literature at the time. Andreani, Martínez, and Schuverdt [5] proved the conjecture of CPLD being a constraint qualification by showing that CPLD implies quasinormality, an established constraint qualification. In subsequent papers, they proved global convergence of a standard augmented Lagrangian method, assuming the CPLD condition.

Weaker constraint qualifications that define “relaxed” variations of the CPLD and CRCQ have been developed by Andreani et al. [3]. In a separate paper, Andreani et al. [4] introduce the weakest constraint qualifications with associated convergence results among all of the ones discussed.

Research into second-order constraint qualifications, for which at a local minimizer second-order (curvature) relations can be expected to hold, has been more limited. Far fewer algorithms attempt to generate a sequence of points that converge to a second-order optimal point. This is because algorithms rarely use explicit second derivatives, and so accurate curvature information is not available, and with global convergence safeguards, convergence to saddle points and local maxima is rare. However, for more sophisticated methods utilizing negative curvature and exact (as opposed to approximate) Lagrangian Hessians, second-order constraint qualification theory is still useful for their convergence theory.

2.3 First-Order Constraint Qualifications

In general, constraint qualifications involve structural relationships concerning the constraint gradients and locally feasible points in a neighborhood of the point in consideration, x^*. Designating the feasible region of the NLP as $\mathbb{F}$, consider the definition of the tangent cone at at a feasible point $x \in \mathbb{F}$.

Definition 2.3.1. *A nonzero vector p is a* tangent *at $x \in \mathbb{F}$ if there exists a feasible sequence $\{x_k\}_{k\geq 0}$ such that $x_k \neq x$ for all k and*

$$\lim_{k\to\infty} (x_k - x)/\|x_k - x\| = p/\|p\|.$$

The set of all tangent vectors at x is denoted by $\mathcal{T}_+(x)$. The set $\mathcal{T}(x) \triangleq \mathcal{T}_+(x) \cup \{0\}$ is called the tangent cone *at x.*

Definition 2.3.2. *The* polar *of the tangent cone $\mathcal{T}(x)$ is the set*

$$\mathcal{T}(x)^\circ = \{w \mid w^T p \leq 0, \ \ \textit{for all} \ \ p \in \mathcal{T}(x)\}.$$

The most succinct necessary condition for optimality is

$$-\nabla f(x) \in \mathcal{T}(x)^\circ.$$

This condition implies that every feasible direction at x has a positive directional derivative with respect to f.

Definition 2.3.3. *Let $\mathcal{F}(x)$ denote the set of directions*

$$\mathcal{F}(x) = \{d \mid \nabla c_i(x)^T d = 0, \ i \in \mathcal{E}, \ \ \textit{and} \ \ \nabla c_i(x)^T d \geq 0, \ i \in \mathcal{A}(x)\}.$$

The normal cone *of $\mathcal{F}(x)$ is the set*

$$\mathcal{N}(x) = \left\{ z \mid z = -\sum_{i\in\mathcal{I}} \mu_i \nabla c_i(x) + \sum_{i\in\mathcal{E}} \lambda_i \nabla c_i(x) \right\} \tag{2.3}$$

defined for all λ such that $\lambda_i = 0$, $i \notin \mathcal{E}$, and all μ such that $\mu_i \geq 0$, $i \in \mathcal{A}(x)$ and $\mu_i = 0$, $i \notin \mathcal{A}(x)$. The normal cone is the polar of $\mathcal{F}(x)$, i.e., $\mathcal{N}(x) = \mathcal{F}(x)^\circ$.

Note the resemblance to the KKT conditions. In particular, $-\nabla f(x) \in \mathcal{N}(x)$ is precisely the first-order condition that the gradient of the Lagrangian is zero. From this it can be implied that $\mathcal{N}(x) = \mathcal{T}(x)^\circ$ serves as a suitable constraint qualification.

Definition 2.3.4. *The* Guignard constraint qualification *holds at x if $\mathcal{N}(x) = \mathcal{T}(x)^\circ$.*

The Guignard constraint qualification is the most general constraint qualification possible. Gould and Tolle [44] show that this condition is equivalent to the KKT conditions being necessary conditions for a local minimum for any objective function.

Definition 2.3.5. *The* Adabie constraint qualification *holds at x if $\mathcal{F}(x) = \mathcal{T}(x)$.*

It may appear that the Guignard and Abadie constraint qualifications are equivalent. However, $\mathcal{T}(x)$ is an arbitrary (possibly disconnected) set, whereas $\mathcal{F}(x)$ is a convex cone.

As an example, consider the equality constraint $c(x) = (x_2^2 - x_1)(x_2 - x_1^2) = 0$ at $x = 0$. The gradient is given by

$$\nabla c(x) = \begin{pmatrix} -(x_2 - x_1^2) - 2x_1(x_2^2 - x_1) \\ 2x_2(x_2 - x_1^2) + (x_2^2 - x_1) \end{pmatrix}, \quad \text{with} \quad \nabla c(0) = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

This implies that every $d \in \mathbb{R}^2$ satisfies $\nabla c(0)^T d = 0$ and hence $\mathcal{F}(0) = \mathbb{R}^2$. However, note that a point is feasible with respect to this constraint if and only if it lies on one of the parabolas $x_2 = x_1^2$ or $x_2 = x_1^2$. At $x = 0$, the tangent cone is the finite set

$$\mathcal{T}(0) = \left\{ \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ -1 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -1 \\ 0 \end{pmatrix} \right\} \neq \mathbb{R}^2 = \mathcal{F}(0).$$

However, notice that $\mathcal{N}(0) = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$, and $\mathcal{T}^\circ(0) = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$ since the elements of $\mathcal{T}(0)$ span $\mathbb{R}^2$, so $\mathcal{N}(0) = \mathcal{T}(0)^\circ$ and the Guignard constraint qualification holds.

The Abadie and Guignard constraint qualifications are not useful for proving the convergence of algorithms because they are too abstract. However, Adabie's condition is frequently used in theoretical analyses because is does not involve the use of the polar operation on a nonconvex set.

2.3.1 List of first-order constraint qualifications

One of the most commonly used constraint qualifications is the linear independence constraint qualification.

Definition 2.3.6. *The* linear independence constraint qualification (LICQ) *holds at x if the set $\{\nabla c_i(x)\}_{i\in\mathcal{E}\cup\mathcal{A}}$ is linearly independent.*

If LICQ holds, then at a local minimizer, the set of multipliers satisfying the KKT conditions contains a single point, i.e., $\mathcal{M}_y(x^*) = \{y^*\}$. If LICQ does not hold at x, then the point x is said to be *degenerate.* In the degenerate case, the Lagrange multipliers are not unique in general.

Another strong constraint qualification used in convex programming is the Slater constraint qualification (see Slater [84]).

Definition 2.3.7. *The* Slater constraint qualification *holds if $\{c_i(x)\}_{i\in\mathcal{E}}$ are affine, $\{c_j(x)\}_{j\in\mathcal{I}}$ are concave and there is a $\bar{x}$ feasible with $c_j(\bar{x}) > 0$ for every $j \in \mathcal{I}$.*

The Slater constraint qualification is a constraint qualification used in convex optimization. It is the only condition that is defined for the whole feasible region instead of a specific point, and implies that there is no duality gap for a convex problem.

The MFCQ is a weaker condition than LICQ that is also frequently used in the literature.

Definition 2.3.8. *The* Mangasarian-Fromovitz constraint qualification (MFCQ) *holds at x if x is feasible, $\{\nabla c_{i\in\mathcal{E}}(x)\}$ is linearly independent and there exists a direction $d \in \mathbb{R}^n$ such that*

$$\nabla c_i(x)^T d = 0,\ i \in \mathcal{E}, \quad \textit{and} \quad \nabla c_i(x)^T d > 0,\ i \in \mathcal{A}(x).$$

Gauvin [34] has shown that the MFCQ is equivalent to the set $\mathcal{M}_y(x^*)$ being bounded.

The constant-rank constraint qualification is another condition that is weaker than the LICQ. It has been used in some convergence theory, and is neither weaker nor stronger than the MFCQ in terms of both the geometric conditions involved and the strength of the convergence results implied. Essentially it states that the linear dependence of subsets of the active constraint gradients is preserved locally.

Definition 2.3.9. *The* constant-rank constraint qualification (CRCQ) *holds at a feasible x if there is a neighborhood $\mathcal{B}(x)$ such that the linear dependence of $\{\nabla c_i(\bar{x})\}_{i\in\mathcal{U}}$ for any index set $\mathcal{U} \subset \mathcal{E} \cup \mathcal{A}(x)$ implies the linear dependence of $\{\nabla c_i(\bar{x})\}_{i\in\mathcal{U}}$ for* every *$\bar{x} \in \mathcal{B}(x)$.*

The relaxed constant-rank constraint qualification states that just subsets of the active constraint gradients that include all of the equality constraints, rather than all possible subsets of $\{\nabla c_i(x)\}_{i\in\mathcal{E}\cup\mathcal{A}(x)}$ need to maintain their linear dependence.

Definition 2.3.10. *The* relaxed constant-rank (RCR) *constraint qualification holds at a feasible x if there is a neighborhood $\mathcal{B}(x)$ such that for every subset $\mathcal{U} \subset \mathcal{A}(x)$, the matrix composed of constraint gradients $\{\nabla c_i(\bar{x})\}_{i\in\mathcal{E}\cup\mathcal{U}}$ has the same rank for every $\bar{x} \in \mathcal{B}(x)$.*

The constant positive linear dependence constraint qualification is similar to the CRCQ but involves the *positive* linear dependence of the active constraint gradients. This constraint qualification is weaker than the CRCQ.

Definition 2.3.11. *A feasible x satisfies the* constant positive linear dependence (CPLD) *condition if it satisfies the MFCQ, or if the positive linear dependence of $\{\nabla c_{i\in\mathcal{U}}(x)\}$ for any $\mathcal{U} \subset \mathcal{E} \cup \mathcal{A}(x)$ implies positive linear dependence of $\{\nabla c_{i\in\mathcal{U}}(\bar{x})\}$ for all $\bar{x}$ in a neighborhood of x.*

Just as CRCQ has a weaker version requiring linear dependence preservation in a smaller set, there is a similarly weaker relaxed constant positive linear dependence condition, also involving only sets that include all of the equality constraints.

Definition 2.3.12. *Consider a feasible point x. Let $\mathcal{V} \subset \mathcal{E}$ be the index set such that the gradients $\{\nabla c_i(x)\}_{i\in\mathcal{V}}$ are linearly independent and span the same subspace as $\{\nabla c_i(x)\}_{i\in\mathcal{E}}$. The point x satisfies the* relaxed constant positive linear dependence (RCPLD) *constraint qualification if there is a neighborhood $\mathcal{B}(x)$ of x such that the following conditions hold.*

- *The matrix with columns $\{\nabla c_i(\bar{x})\}_{i\in\mathcal{E}}$ has constant rank for every $\bar{x} \in \mathcal{B}(x)$.*
- *For all subsets $\mathcal{U} \subset \mathcal{A}(x)$, if the gradients $\{\nabla c_i(x)\}_{i\in\mathcal{U}\cup\mathcal{V}}$ are positively linearly dependent, then the gradients $\{\nabla c_i(x)\}_{i\in\mathcal{U}\cup\mathcal{V}}$ are positively linearly dependent for all $\bar{x} \in \mathcal{B}(x)$.*

The next two constraint qualifications use more specific index sets of the constraints. The first identifies the set of inequality constraints that behave like equalities. For example, inequality constraints $c_1(x) \geq 0$ and $c_2(x) \geq 0$ such that $c_1(x) = -c_2(x)$ may be replaced by the single equality constraint $c(x) = c_1(x) = c_2(x) = 0$. As $\nabla c_1(x) = -\nabla c_2(x)$, both constraint gradients are in the normal cone and $\mathcal{I}_-$. This behavior can be generalized to be a local property that is characterized by $i \in \mathcal{A}(x)$ and $\nabla c_i(x) \in \mathcal{N}(x)$. Let $\mathcal{I}_-$ denote the index set

$$\mathcal{I}_- = \{j \in \mathcal{A}(x) \mid \nabla c_j(x) \in \mathcal{N}(x)\}.$$

Indices in the set $\mathcal{I}_+ = \mathcal{A}(x) \setminus \mathcal{I}_-$ correspond to the well-behaved inequality constraints. Let $\mathcal{E}' \subset \mathcal{E}$ be such that $\{\nabla c_i(x)\}_{i\in\mathcal{E}'\cup\mathcal{I}_+}$ is a positively linearly independent spanning set of $\{\nabla c_i(x)\}_{i\in\mathcal{E}\cup\mathcal{A}(x)}$. It can be shown that the MFCQ implies that $\mathcal{I}_-$ is empty.

An an example, consider the constraints $c_1(x) = x \geq 0$ and $c_2(x) = -x \geq 0$, for which $\nabla c_1 = 1$ and $\nabla c_2 = -1$. Clearly these are both in the normal cone, as $\nabla c_1(x) = -\nabla c_2(x)$. For this set of constraints, $\mathcal{I}_- = \{1, 2\}$ and $\mathcal{I}_+$ is empty.

These sets attempt to separate the two geometric patterns of constraint qualification —a subspace of equality constraint gradients with a locally stable dimension and a tangent cone of inequalities with a locally constant topology. The following two constraint qualifications make this formal by requiring on constant rank of the *de facto* equality constraints and the local maintenance of the positive linear basis for all of the constraint gradients.

Since the constraints in $\mathcal{I}_-$ behave like equalities, the next condition is the weakest generalization of the constant rank conditions.

Definition 2.3.13. *The* constant rank of the subspace component (CRSC) *holds at x if there is a neighborhood $\mathcal{B}(x)$ of x in which the rank of $\{\nabla c_i(\bar{x}) \mid i \in \mathcal{E} \cup \mathcal{I}_-\}$ is constant for all $\bar{x} \in \mathcal{B}(x)$.*

The following condition is the weakest generalization of the constant positive linear dependence conditions.

Definition 2.3.14. *The* constant positive generator (CPG) *constraint qualification holds at x if there is a neighborhood $\mathcal{B}(x)$ of x wherein $\{\nabla c_i(\bar{x}) \mid i \in \mathcal{E}' \cup \mathcal{I}_+\}$ positively spans $\{\nabla c_i(\bar{x}) \mid i \in \mathcal{E} \cup \mathcal{A}(x)\}$ for all $\bar{x} \in \mathcal{B}(x)$.*

There are two final constraint qualifications mentioned in the literature. These are more for theoretical interest and have not been used in the convergence theory of algorithms.

Definition 2.3.15. *The* quasinormality constraint qualification *holds at x if either x satisfies the MFCQ or there are* no *nonzero vectors λ and μ such that*

- $\mu_i \geq 0$, *with* $\mu_i = 0$ *for* $i \notin \mathcal{A}(x)$;
- $\sum_{i \in \mathcal{E}} \lambda_i \nabla c_i(x) + \sum_{i \in \mathcal{A}(x)} \mu_i \nabla c_i(x) = 0$; *and*
- *there is no sequence $x_k \to x$ such that for all i such that $\lambda_i \neq 0$ and $\mu_i \neq 0$,*

$$\lambda_i c_i(x_k) > 0 \quad \text{and} \quad \mu_i c_i(x_k) > 0,$$

for all k.

Definition 2.3.16. *The* pseudonormality constraint qualification *holds at x if either x satisfies the MFCQ or there is no $(\lambda, \mu\}$ such that*

- *$\mu_i \geq 0$, with $\mu_i = 0$ for $i \notin \mathcal{A}(x)$;*
- $\sum_{i\in\mathcal{E}} \lambda_i \nabla c_i(x) + \sum_{i\in\mathcal{I}} \mu_i \nabla c_i(x) = 0$;
- *there is no sequence $\{x_k\}$ such that $x_k \to x$ with*

$$\sum \lambda_i c_i(x_k) + \sum \mu_i c_i(x_k) > 0,$$

for all k.

Bertsekas and Ozdaglar [9] showed that these two conditions imply that there exists an exact penalty function for the original problem as well as implying the existence of strong and informative Lagrange multipliers, which can identify redundant constraints.

2.3.2 Relationships between first-order constraint qualifications

The LICQ implies the CRCQ holds. If the LICQ holds, then none of the constraint gradients are linearly dependent at x, therefore the implication that linear dependence at x implies linear dependence in a neighborhood of x is trivially true.

The CRCQ does not imply the LICQ. Consider the inequality constraints $c_1(x) = x_1 + x_2 \geq 0$ and $c_2(x) = -x_1 - x_2 \geq 0$ at $x = 0$. The gradients are $\nabla c_1(x) = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and $\nabla c_2(x) = \begin{pmatrix} -1 \\ -1 \end{pmatrix}$ which are linear dependent for all x.

The LICQ implies the MFCQ. Let J denote the matrix of active constraint gradients. If the LICQ holds, then J has full row-rank and the system $J(x)p = r$ has a solution for every vector r. The right-hand side r such that $r_i = 0$, $i \in \mathcal{E}$ and $r_i = 1$ for $i \in \mathcal{A}(x)$ gives a p such that $\nabla c_i(x)^T p = 0$ for $i \in \mathcal{E}$ and $\nabla c_i(x)^T p = 1 > 0$ for $i \in \mathcal{A}(x)$.

The Slater constraint qualification implies the MFCQ. If the equality constraints are linear, then their gradients must be linearly independent, otherwise they are redundant constraints or there is no feasible point.

If the inequality constraints are concave and there is an $\bar{x}$ for which $c_i(\bar{x}) > 0$ for all $i \in \mathcal{A}(x)$, then at a feasible point x, every point along $x + tp = x + t(\bar{x} - x)$ must be feasible since $-c_i(x + tp) \leq -c_i(x) - c_i(\bar{x}) \leq 0$. This implies that p is a strictly feasible direction.

The MFCQ does not imply the LICQ. Consider the inequality constraints $c_1(x) = x_1 \geq 0$, $c_2(x) = x_1 - x_2^2 \geq 0$ and $c_3(x) = x_1 - x_2 \geq 0$ at $x = (0, 0)$. All three constraints are active at x, which implies that LICQ does not hold since the constraint gradients are three vectors in $\mathbb{R}^2$ and must be dependent. However $p = (1, 0)$ is a strictly feasible direction at x.

The CRCQ does not imply the MFCQ. Consider the constraints $c_1(x) = x \geq 0$ and $c_2(x) = -x \geq 0$ at $x = 0$. The constraints are linearly dependent for every feasible x. However, there is no strictly feasible direction.

The MFCQ does not imply the CRCQ. Consider the constraints $c_1(x) = x_1 \geq 0$ and $c_2(x) = x_1 - x_2^2 \geq 0$ at $x = 0$. The vector $p = (1, 0)$ is a feasible direction. However, the constraint gradients $\nabla c_1(x) = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $\nabla c_1(x) = \begin{pmatrix} 1 \\ -x_2 \end{pmatrix}$ are linearly dependent only when $x_2 = 0$.

The CRCQ implies the CPLD. The rank of $\{\nabla c_{i\in\mathcal{U}}(x)\}$ is the same as the rank of $\{\{\nabla c_{i\in\mathcal{U},i\neq j}(x)\} \cup -\nabla c_{j\in\mathcal{U}}(x)\}$.

Let $\nabla c_{i\in\mathcal{U}}(x)$ be positively linearly dependent. Then all subsets of

$$\{\nabla c_1(x), \ldots, \nabla c_l(x), -\nabla c_1(x), \ldots, -\nabla c_l(x)\}$$

that include one positive or one negative gradient for each i are linearly dependent. Then for every neighborhood of x, the CRCQ implies that the same subsets among

$$\{\nabla c_1(\bar{x}), \ldots, \nabla c_l(\bar{x}), -\nabla c_1(\bar{x}), \ldots, -\nabla c_l(\bar{x})\}$$

are linearly dependent for $\bar{x}$ in a neighborhood of x. This implies that $\nabla c_1(x)$, ..., $\nabla c_l(x)$ are positively linearly dependent.

The MFCQ implies the CPLD. This follows directly from the definition of the CPLD.

The CPLD implies neither the MFCQ nor the CRCQ. Consider the inequalities $c_1(x) = x_1 \geq 0$, $c_2(x) = x_1 + x_2^2 \geq 0$, $c_3(x) = x_1 + x_2 \geq 0$, and $c_4(x) = -x_1 - x_2 \geq 0$ at $x^* = 0$. The CPLD holds because c_1, c_3, and c_4 are linear, so $\nabla c_1(x^*)$, $\nabla c_3(x^*)$, and $\nabla c_4(x^*)$ being positively linearly dependent implies that they are positively linearly dependent for all $x \in \mathbb{R}^2$.

However, since $\nabla c_3 = -\nabla c_4$, the MFCQ cannot hold.

Also $\nabla c_1(x^*)$ and $\nabla c_2(x^*)$ are linearly dependent but not $\nabla c_1(x)$ and $\nabla c_2(x)$ for any point, for instance, at which $x_1 = x_2 \neq 0$. This example is in Andreani, Martínez and Schuverdt [5].

The CRCQ implies the RCRCQ. This is trivial, since the set of subsets for which constant rank must be maintained is strictly larger for the CRCQ as compared to the RCRCQ.

The CPLD implies the RCPLD. This is trivial, since the set of subsets for which constant positive linear dependence must be maintained is strictly larger for the CPLD as compared to the RCPLD.

The RCRCQ implies the RCPLD. A proof is given by Andreani et al. [3].

The RCPLD implies the CRSC. A proof is given by Andreani et al. [4].

The CRSC implies the CPG. A proof is given by Andreani et al. [4].

The CPG implies the Abadie constraint qualification. A proof is given by Andreani et al. [4]. As many of the constraint qualifications imply CPG, this implies that CPG is the weakest of a hierarchy of constraint qualifications starting with the Abadie constraint qualification.

CPLD implies quasinormality. A proof is given by Andreani, Martínez and Schuverdt [5].

Quasinormality does not imply CPLD. Consider the equality constraints $c_1(x) = x_2 e^{x_1} = 0$ and $c_2(x) = x_2 = 0$ at $x = 0$. Then c_1 and c_2 are the same sign so take $\lambda_1 < 0$, $\lambda > 0$ to satisfy quasinormality. However, the gradients are trivially positively linearly dependent at $x^* = 0$ but not at any point in a neighborhood of x^*. This example is given in Andreani, Martínez and Schuverdt [5].

Consider the following illustrative examples to understand the nuances of the hierarchy of constraint qualifications. The first example is an illustration of the excessive strictness of LICQ. Here there are three constraints that are active at $x = (0, 0)$. Since

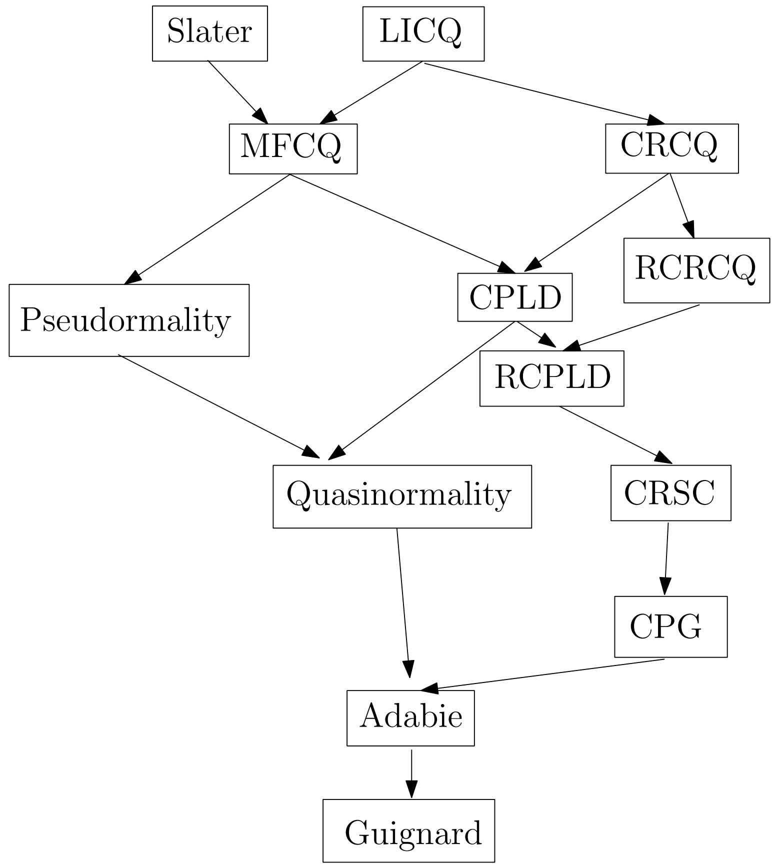

Figure 2.1: Hierarchy of constraint qualifications

these constraints live in $\mathbb{R}^2$ their gradients are trivially linearly dependent. Hence the LICQ does not hold. However, all other constraint qualifications hold.

In the second example, there are two inequalities that are in actuality an equality constraint. Notice in this case MFCQ and the stronger conditions fail, however, the constraint qualifications corresponding to constant rank in a neighborhood hold.

The next example illustrates a situation where locally the two active inequality constraints are identical while otherwise entirely different functions. In this case the constant rank conditions fail, while as there is a clear feasible direction, MFCQ holds.

In the final example, the challenges presented by the previous two example simultaneously hold, so there are two inequality constraints that act as equalities, and there are two inequality constraints with the same gradient at the point of interest. Here, MFCQ and CRCQ both fail. However, both pseudonormality and CPLD hold, since positive linear-dependence is maintained locally.

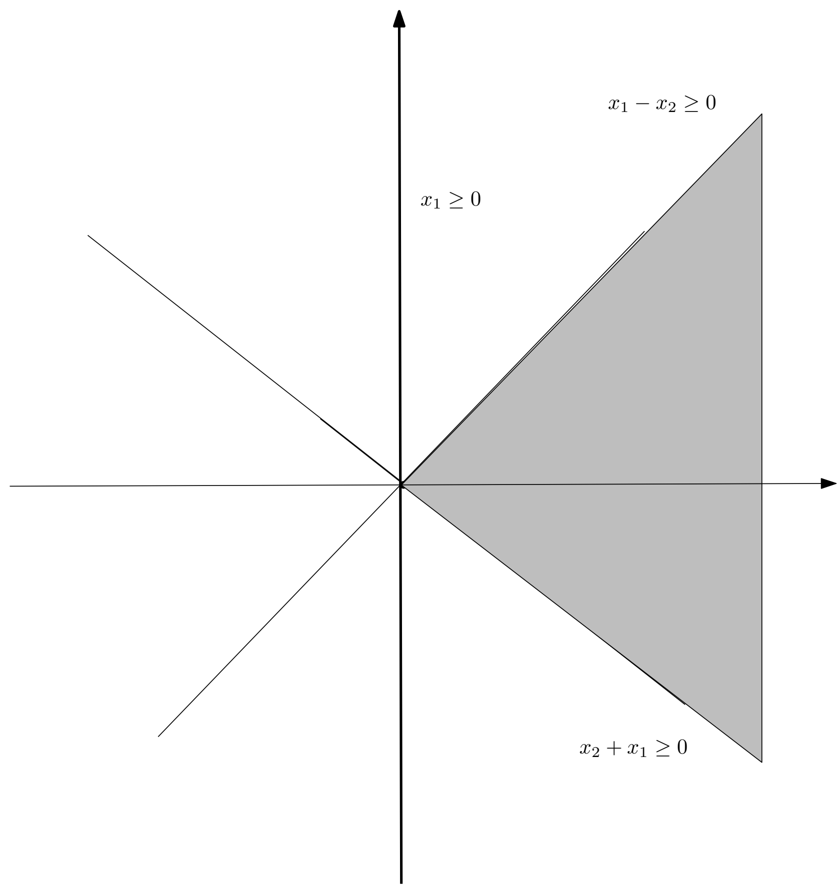

Figure 2.2: Many (three in 2D) constraints active at a point resulting in linear dependence of the constraint gradients. The constraint qualifications *Slater, MFCQ, CRCQ, RCRCQ, CPLD, RCPLD, pseudonormality, quasinormality, CRSC, CPG, Adabie, and Guignard* all hold. *LICQ* does not hold.

2.4 Second-Order Optimality Conditions

Second-order constraint qualifications define conditions on the geometry of the feasible region that certify that a local minimizer is a *second-order KKT point*. Broadly speaking, a second-order KKT point is a point at which the objective function exhibits positive curvature along certain feasible directions. Consider the problem

$$\underset{x \in \mathbb{R}^2}{\text{minimize}} \ \ \tfrac{1}{2}(x_1^2 - x_2^2) \quad \text{subject to} \ \ -1 \leq x_1 \leq 1 \ \text{ and } \ -1 \leq x_2 \leq 1. \tag{2.4}$$

Notice that at $x^* = (0, 0)$, it holds that $\nabla f(x^*) = 0$. Since both constraints are feasible and inactive, x^* satisfies the first-order KKT conditions. Yet, clearly x^* is a saddle point, and movements in the feasible directions $p = (0, 1)$ or $p = (0, -1)$ would result in a decrease of the objective function. This highlights the importance of curvature in determining optimality.

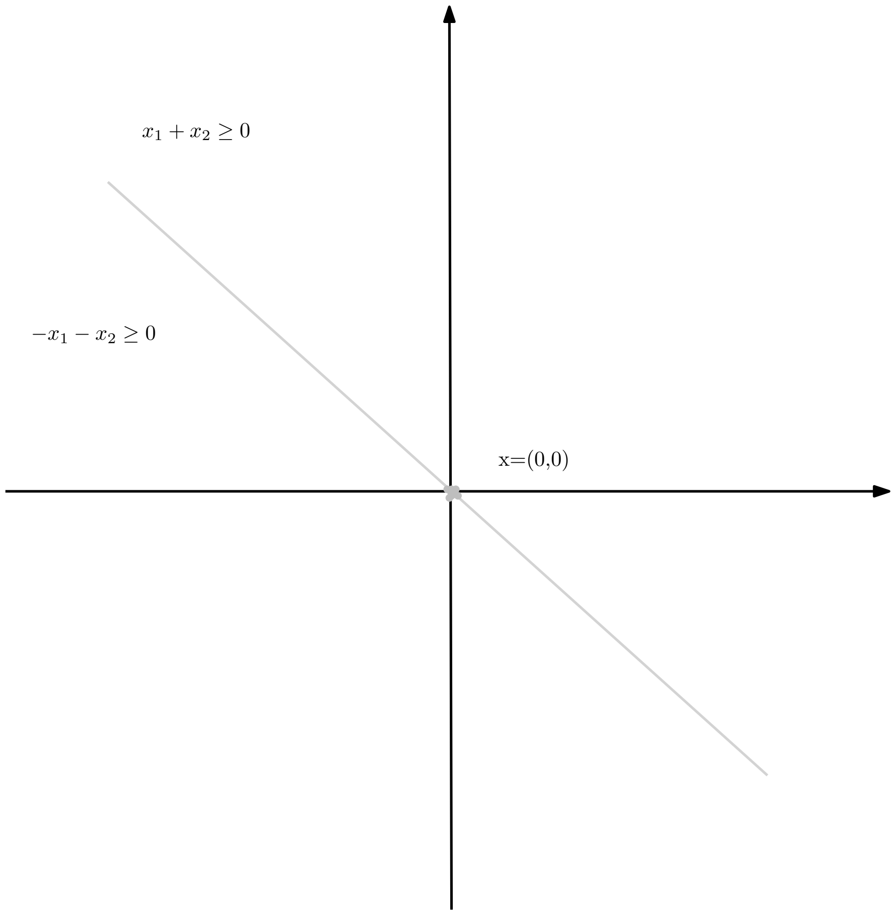

Figure 2.3: Equality constraint written as two inequalities. The feasible region is strictly along the line. The constraint qualifications *CRCQ, RCRCQ, CPLD, RCPLD, quasinormality, CRSC, CPG, Adabie, and Guignard* all hold. On the other hand, *LICQ, Slater, MFCQ, and pseudonormality* fail to hold.

In addition, consider the problem

$$\underset{x\in\mathbb{R}^2}{\text{minimize}} \;\; \tfrac{1}{2}(-x_1^2 - x_1) + x_2^2 \quad \text{subject to} \;\; -1 \leq x_1 \leq 1 \;\; \text{and} \;\; -1 \leq x_2 \leq 1, \tag{2.5}$$

which has a local (and global) minimizer at $x^* = (1, 0)$. In this case, the curvature of the objective function is negative in a direction along $e_1 = (1, 0)$ close to x^*, suggesting that a very particular set of directions have to be chosen in order to assess the relationship between optimality and curvature.

For these examples the constraints are linear, which implies that the curvature of f on the constraints with indices in $\mathcal{E} \cup \mathcal{A}(x)$ does not involve the curvature of the constraints. However, when the constraints are nonlinear, the curvature of f is characterized by the

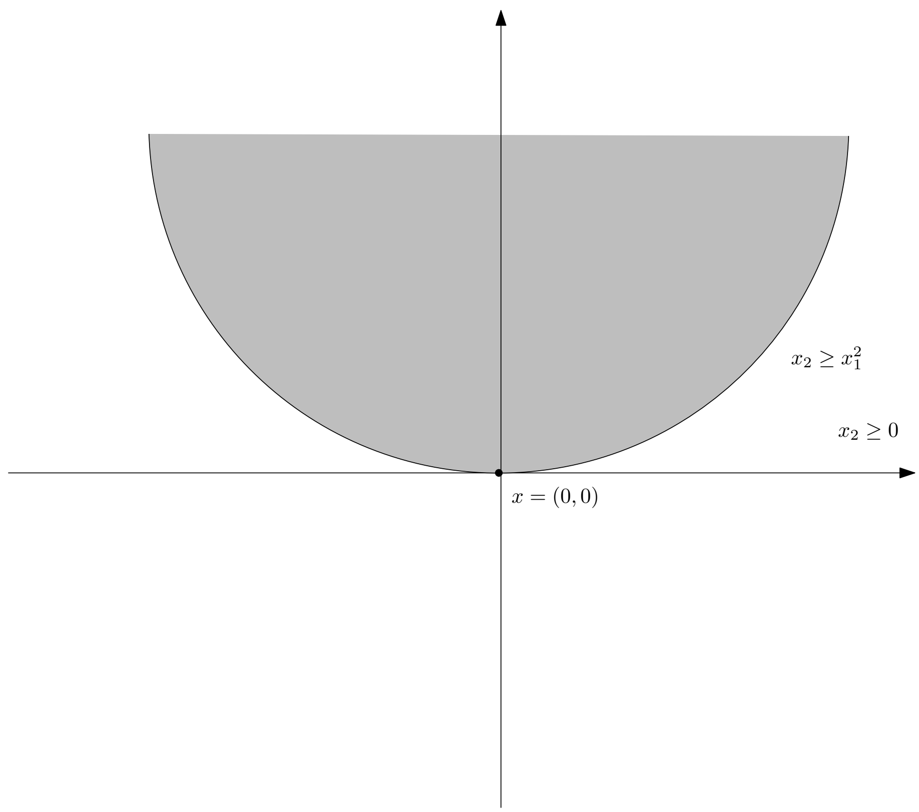

Figure 2.4: Locally linearly dependent constraints. The constraint qualifications *MFCQ, pseudonormality, CPLD, RCPLD, quasinormality, CRSC, CPG, Adabie, and Guignard* all hold. However, *LICQ, Slater, CRCQ, and RCRCQ* all do not hold.

Hessian of the Lagrangian function, i.e.,

$$H(x, y) = \nabla^2 f(x) - \sum_{i \in \mathcal{E}} y_i \nabla^2 c_i(x) - \sum_{i \in \mathcal{A}(x)} y_i \nabla^2 c_i(x),$$

which depends on the value of the multiplier vector y. In this case, the analysis is complicated by the fact that there may be many y associated with a first-order KKT point, with each y giving a Hessian with different properties.

2.4.1 Second-order necessary conditions

The set of all first-order feasible directions is

$$\mathcal{F}(x) = \{d \mid \nabla c_i(x)^T d = 0,\ i \in \mathcal{E},\ \text{ and }\ \nabla c_i(x)^T d \geq 0,\ i \in \mathcal{A}(x)\}$$

(see Definition 2.3.3, page 12). Consider any $d \in \mathcal{F}(x)$. If $\nabla f(x)^T d > 0$, then d is a direction of increase for f, and the curvature is irrelevant. The inequality $\nabla f(x^*)^T d < 0$ cannot hold

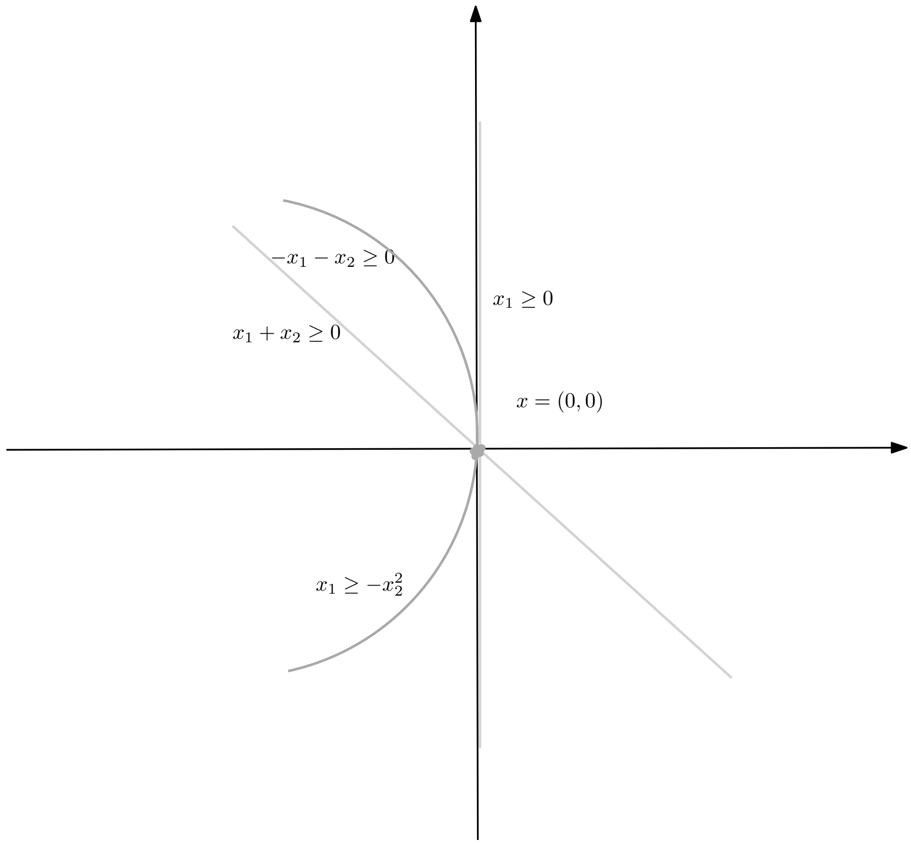

Figure 2.5: Equality constraint written as two inequalities and locally linearly dependent constraints. The feasible region is the origin. The constraint qualifications *CPLD, RCPLD, pseudonormality, quasinormality, RCPLD, CRSC, CPG, Adabie, and Guignard* hold, whereas *LICQ, Slater, MFCQ, CRCQ, and RCRCQ* all fail to hold.

because the first-order stationarity condition, since by taking inner product of d with the vanishing Lagrangian, it must hold that $y_i \nabla c_i(x^*)^T d < 0$, $i \in \mathcal{E}$ or $y_i \nabla c_i(x^*)^T d < 0$, $i \in \mathcal{I}$ for some i. It follows that the directions of interest have the property that $\nabla f(x^*)^T d = 0$, i.e., if the objective is flat along a feasible direction at a local minimizer, the curvature must be positive.

Definition 2.4.1. *Given a KKT point x, the set of directions*

$$\mathcal{C}(x) = \{\, d \mid \nabla f(x^*)^T d = 0, \;\; \nabla c_i(x^*)^T d = 0, \; i \in \mathcal{E}, \;\; \nabla c_j(x^*)^T d \geq 0, \; i \in \mathcal{A}(x) \,\} \qquad (2.6)$$

is called the critical cone *at x.*

Alternatively, in practice the appropriately named weak reduced semidefiniteness property (WSRP) is sometimes used (see Andreani, Martínez, and Schuverdt [6]) where the

Table 2.1: List of second-order necessary conditions of optimality

Condition	Cone	Curv	Second-Order Necessary Conditions
WSRP	$\widetilde{\mathcal{C}}(x^*)$	L	$\exists y \in \mathcal{M}_y(x^*)$ s.t. $d^T \nabla^2_{xx} L d \geq 0$, for every $d \in \widetilde{\mathcal{C}}(x^*)$
Fritz John	$\mathcal{C}(x^*)$	L_G	for every $d \in \mathcal{C}(x^*)$ $\exists \{y_0, y\} \in \mathcal{M}_y(x^*)$ s.t. $d^T \nabla^2_{xx} L_G d \geq 0$
SO-2	$\mathcal{C}(x^*)$	L	for every $d \in \mathcal{C}(x^*)$, $\exists y \in \mathcal{M}_y(x^*)$ s.t. $d^T \nabla^2_{xx} L d \geq 0$
SO-3	$\mathcal{C}(x^*)$	L	$\exists y \in \mathcal{M}_y(x^*)$ s.t. $d^T \nabla^2_{xx} L d \geq 0$, for every $d \in \mathcal{C}(x^*)$

cone on which positive-semidefiniteness of the Hessian is required is

$$\widetilde{\mathcal{C}}(x^*) = \{d \mid \nabla c_i(x^*)^T d = 0,\ i \in \mathcal{E},\ \ \nabla c_j(x^*)^T d = 0,\ i \in \mathcal{A}(x)\}. \tag{2.7}$$

If strict complementarity holds at x^*, then $\widetilde{\mathcal{C}}(x^*) = \mathcal{C}(x^*)$.

Consider the generalized Lagrangian

$$L_G(x, y_0, y) = y_0 f(x) - \sum_{i \in \mathcal{E}} y_i c_i(x) - \sum_{j \in \mathcal{A}(x)} y_j c_j(x).$$

Without any constraint qualifications, a local minimizer satisfies the *Fritz John second-order necessary conditions*, which state that there exist $\{y_0, y\}$ satisfying the first-order Fritz John conditions and for every $d \in \mathcal{C}(x^*)$, there are multipliers $\{y_0, y\}$ such that $d^T \nabla^2{}_{xx} L_G(x^*, y_0, y) d \geq 0$.

This condition is difficult to verify computationally. First, y_0 may be zero. Second, there may be many multiplier vectors satisfying the conditions, and an algorithm may generate multiplier estimates that do not approach those that satisfy the second-order conditions. Finally, the optimality condition require that $L_G(x, y_0, y)$ has positive curvature at least one $y \in \mathcal{M}_y$, but there may be other $y \in \mathcal{M}_y$ for which $L_G(x, y_0, y)$ has negative or zero curvature.

A stronger condition, which avoids the first issue, requires that for every $d \in \mathcal{C}(x^*)$, there is a y for which $d^T \nabla^2{}_{xx} L(x^*, y, y) d \geq 0$. This condition will be referred to as SO-2.

Finally, the condition SO-3 is said to hold if there exists a multiplier y such that for all $d \in \mathcal{C}(x^*)$, $d^T \nabla^2{}_{xx} L(x^*, y) d \geq 0$. Here there is a multiplier for which the Hessian is positive semidefinite on the entire cone. This condition is commonly associated with the existence of a single unique set of multipliers at x^*, with some exceptions (see [7]).

The second-order necessary conditions are summarized below:

2.4.2 Second-order sufficient conditions

Definition 2.4.2 (Second-Order Sufficient Condition (SOSC)). *The second-order sufficient condition is said to be satisfied at a KKT point x^* if there is a positive σ such that*

$$\textit{for every } y \in \mathcal{M}_y(x^*), \textit{ it holds that } d^T \nabla^2_{xx} L(x^*, y) d \ge \sigma \|d\|^2 \textit{ for all } d \in \mathcal{C}(x^*).$$

Definition 2.4.3 (Relaxed Second-Order Sufficient Condition (RSOSC)). *A relaxed second-order sufficient condition is said to be satisfied at a KKT point x^* if there is a positive σ such that*

$$\textit{for some } y^* \in \mathcal{M}_y(x^*), \textit{ it holds that } d^T \nabla^2_{xx} L(x^*, y^*) d \ge \sigma \|d\|^2 \textit{ for all } d \in \mathcal{C}(x^*).$$

The verification of the RSOSC at a KKT point requires finding the global minimizer of a possibly indefinite quadratic form over a cone, which is an NP-hard problem. However, stronger conditions may be formulated that can be verified using linear algebraic procedures.

For any $y \in \mathcal{M}_y(x^*)$, consider the inner product of any $d \in \mathcal{C}(x^*)$ with the gradient of the Lagrangian evaluated at (x^*, y). As x^* is a first-order KKT point, it must hold that

$$\nabla f(x^*)^T d - \sum_{i \in \mathcal{E}} y_i \nabla c_i(x^*)^T d - \sum_{i \in \mathcal{A}(x^*)} y_i \nabla c_i(x^*)^T d = 0.$$

The first two terms are zero, so that

$$\begin{aligned} 0 = \sum_{i \in \mathcal{A}(x^*)} y_i \nabla c_i(x^*)^T d &= \sum_{i \in \mathcal{A}_0(x^*, y)} y_i \nabla c_i(x^*)^T d + \sum_{i \in \mathcal{A}_+(x^*, y)} y_i \nabla c_i(x^*)^T d \\ &= \sum_{i \in \mathcal{A}_+(x^*, y)} y_i \nabla c_i(x^*)^T d, \end{aligned} \tag{2.8}$$

where $\mathcal{A}_0(x, y)$ and $\mathcal{A}_+(x, y)$ are sets containing the indices of the constraints with zero and positive multipliers, i.e.,

$$\mathcal{A}_0(x, y) = \{ i \in \mathcal{A}(x) \mid y_i = 0 \} \quad \text{and} \quad \mathcal{A}_+(x, y) = \{ i \in \mathcal{A}(x) \mid y_i > 0 \}.$$

As each term in the sum (2.8) is nonnegative, it must hold that $\nabla c_i(x^*)^T d = 0$, for every $i \in \mathcal{A}_+(x, y)$. This restriction and the overall requirement that $\nabla c_i(x^*)^T d \ge 0$, for every $i \in \mathcal{A}(x^*)$, implies that the critical cone (2.6) may be written in the alternative form:

$$\begin{aligned} \mathcal{C}(x^*) = \{ \, d \mid \; & \nabla c_i(x^*)^T d = 0, \; i \in \mathcal{E}, \\ & \nabla c_i(x^*)^T d = 0, \; i \in \mathcal{A}_+(x, y), \\ & \nabla c_i(x^*)^T d \ge 0, \; i \in \mathcal{A}_0(x, y) \; \}. \end{aligned} \tag{2.9}$$

It must be emphasized that although the critical cone $\mathcal{C}(x^*)$ can written in terms of any $y \in \mathcal{M}_y(x^*)$, the definition (2.6) implies that $\mathcal{C}(x^*)$ is independent of the choice of y. The characterization of the critical cone in the form (2.9) implies that $\mathcal{C}(x) = \mathcal{C}_+(x,y) \cap \mathcal{C}_0(x,y)$, where $\mathcal{C}_+$ and $\mathcal{C}_0$ are the index sets

$$\begin{aligned} \mathcal{C}_+(x,y) &= \{d \mid \nabla c_i(x)^T d = 0,\ i \in \mathcal{E} \cup \mathcal{A}_+(x,y)\} \\ \text{and} \quad \mathcal{C}_0(x,y) &= \{d \mid \nabla c_i(x)^T d \geq 0,\ i \in \mathcal{A}_0(x,y)\}. \end{aligned}$$

As $\mathcal{C}(x)$ is contained in both $\mathcal{C}_+(x,y)$ and $\mathcal{C}_0(x,y)$, it follows that by working with the set $\mathcal{C}_+(x,y)$ only, (which is larger and hence leads to more restrictive conditions than those imposed by $\mathcal{C}(x)$), we may derive conditions that are verifiable computationally.

Definition 2.4.4 (Strong Second-Order Sufficient Condition (SSOSC))**.** *A strong second-order sufficient condition is said to be satisfied at a first-order KKT point x^* if there is a positive σ such that:*

$$\textit{for every } y \in \mathcal{M}_y(x^*), \ \textit{ it holds that } d^T \nabla^2_{xx} L(x^*, y) d \geq \sigma \|d\|^2 \ \textit{ for all } d \in \mathcal{C}_+(x^*, y).$$

This condition was proposed by Robinson [82]. It is equivalent to the locally strong second-order sufficient condition (LSSOSC) considered by Wright [87].

Definition 2.4.5 (Relaxed Strong Second-Order Sufficient Condition (RSSOSC))**.** *The relaxed strong second-order sufficient condition is said to be satisfied at a first-order KKT point x^* if there is a positive σ such that:*

$$\textit{for some } y^* \in \mathcal{M}_y(x^*), \ \textit{ it holds that } d^T \nabla^2_{xx} L(x^*, y^*) d \geq \sigma \|d\|^2 \ \textit{ for all } d \in \mathcal{C}_+(x^*, y^*).$$

2.4.3 Second-order conditions involving strongly active constraints

Definition 2.4.6. *A constraint $c_i(x) \geq 0$ is said to be* strongly active *at a KKT point x^* if $i \in \mathcal{A}(x^*)$ and there exists at least one $y^* \in \mathcal{M}_y(x^*)$ with $y_i^* > 0$. Similarly, the ith constraint is said to be* weakly active *at x^* if $i \in \mathcal{A}(x^*)$ and $y_i^* = 0$ for* all *$y^* \in \mathcal{M}_y(x^*)$; in this case, we say that constraint i has a* null multiplier*.*

For example, the problem

$$\underset{x \in \mathbb{R}^2}{\text{minimize}} \ \ \tfrac{1}{2}(x_1^2 + x_2^2) \quad \text{subject to} \ \ x_1 + x_2 \geq 4 \ \text{ and } \ x_2 \geq 2,$$

has a unique solution $x^* = (2,2)$ with (unique) multiplier $y^* = (2,0)^T$, which implies that the second constraint has a null multiplier.

The set of strongly active constraints at x^* is denoted by $A_+(x^*)$ (without a second argument y), i.e.,

$$A_+(x^*) = \cup_{y \in \mathcal{M}_y} \mathcal{A}_+(x^*, y^*)$$

With this definition, the set

$$A_0(x^*) = \mathcal{A}(x^*) \setminus \mathcal{A}_+(x^*)$$

is the set of all weakly active constraints.

Definition 2.4.7. *The property of* strict complementarity *holds at the KKT point x^* if there is a multiplier $y^* \in \mathcal{M}_y(x^*)$ such that $y_i^* > 0$ for all $i \in \mathcal{A}(x^*)$.*

If strict complementarity holds at x^*, then $\mathcal{A}_+(x^*) = \mathcal{A}(x^*)$, in which case $\mathcal{A}_0(x^*)$ is empty.

Consider the set of critical directions defined in terms of the sets $\mathcal{A}_+$ and $\mathcal{A}_0$, i.e.,

$$\widehat{\mathcal{C}}(x) = \left\{ \, d \mid \nabla c_i(x)^T d = 0, \; i \in \mathcal{E} \cup \mathcal{A}_+(x), \;\; \nabla c_i(x)^T d \geq 0, \; i \in \mathcal{A}_0(x) \, \right\},$$

which is the intersection of the sets

$$\begin{aligned} \widehat{\mathcal{C}}_+(x) &= \{ d \mid \nabla c_i(x)^T d = 0, \; i \in \mathcal{E} \cup \mathcal{A}_+(x) \} \\ \text{and} \quad \widehat{\mathcal{C}}_0(x) &= \{ d \mid \nabla c_i(x)^T d \geq 0, \; i \in \mathcal{A}_0(x) \} \end{aligned}$$

(cf. the critical set (2.9)).

Definition 2.4.8 (Second-Order Sufficient Condition (SOSC2))**.** *The second second-order sufficient condition is said to be satisfied at a KKT point x^* if there is a positive σ such that*

$$\textit{for every } \; y \in \mathcal{M}_y(x^*), \;\; \textit{it holds that } \; d^T \nabla^2_{xx} L(x^*, y) d \geq \sigma \|d\|^2 \; \textit{ for all } \; d \in \widehat{\mathcal{C}}(x^*).$$

Note that any $d \in \widehat{\mathcal{C}}(x^*)$ must also be in $\mathcal{C}(x^*)$, but $\mathcal{C}(x^*)$ may include directions from the set

$$\{ d \mid \nabla c_j(x^*)^T d > 0, \;\; \text{for} \;\; y_j = 0, \; j \in \mathcal{A}_+(x^*), \;\; \text{and} \;\; \nabla c_i(x^*)^T d \geq 0, \;\; \text{for all} \;\; i \in \mathcal{A}(x^*) \},$$

which has no intersection with $\widehat{\mathcal{C}}(x^*)$. It follows that $\mathcal{C}(x^*)$ is a larger set than $\widehat{\mathcal{C}}(x^*)$, so that the SOSC is a stronger condition than the SOSC2.

2.4.4 Second-order constraint qualifications

Definition 2.4.9. *Let x denote a first-order KKT point, and let y be any multiplier $y \in \mathcal{M}_y(x)$. A nonzero p is a* second-order tangent *at x if there exists a feasible sequence $\{x_k\}_{k\geq 0}$ with $x_k \neq x$ such that*

$$c_i(x_k) = 0,\ i \in \mathcal{E}, \quad c_i(x_k) = 0,\ i \in \mathcal{A}_+(x,y), \quad c_i(x_k) \geq 0,\ i \in \mathcal{A}_0(x,y),$$

and $\lim_{k\to\infty} (x_k - x)/\|x_k - x\| = p/\|p\|$. *The set of all second-order tangent vectors at x is denoted by $\mathcal{T}_+^{(2)}(x)$. The set $\mathcal{T}^{(2)}(x) \triangleq \mathcal{T}_+^{(2)}(x) \cup \{0\}$ is called the* second-order tangent cone *at x.*

For a point x^* to be a local minimizer, it must hold that $d^T H(x^*, y)d \geq 0$ for some $y \in \mathcal{M}_y(x^*)$ and all $d \in \mathcal{T}^{(2)}(x^*)$.

From the discussion above, it can be seen that if it is possible to establish that a method generates a sequence of iterates that converge to a point where strong rather than weak, second-order conditions hold, then the method will converge on a wider class of problems. This class includes problems for which strict complementarity does not hold and problems having positive-curvature hold for all multipliers satisfying the first-order condition allows for robustness in primal-dual algorithms that converge to a point for which the set of optimal multipliers is nonunique. However, for stronger results, it is necessary for stronger constraint qualifications to be assumed. In the next section we define a number of different second-order constraint qualifications and discuss the second-order conditions that they imply.

There are far more first-order than second-order constraint qualifications. On the other hand, most of the theoretical results have been with regards to the stronger second-order condition, in particular showing SO-3.

Analogous to the Adabie first-order constraint qualification, the *second-order Adabie* constraint qualification holds if

$$\mathcal{C}(x) = \mathcal{T}^{(2)}(x),$$

which implies that the critical cone is the same as the second-order tangent cone.

The CRCQ was shown to be a second-order constraint qualification in Andreani, Echagiie and Schuverdt [2], implying that, at a local minimum, SO-3 holds. Andreani et al. [3] also remark that the RCRCQ has similar enough geometric properties to the CRCQ so that the results for CRCQ can also be applied to RCRCQ to show the same property.

In general, MFCQ is not a second-order constraint qualification. However, MFCQ together with some additional assumptions does become a second-order constraint qualification. Baccari and Trad [8] discuss this in detail. They give an example for which MFCQ holds but WSRP does not hold. They show that if the set of multipliers $\mathcal{M}_y(x^*)$ satisfying the first order conditions is a bounded line segment and there is at most one constraint which is weakly active for all multiplier vectors in $\mathcal{M}_y(x^*)$, SO-3 holds. Under MFCQ, $\mathcal{M}_y(x^*)$ is a bounded line segment under any one of the following conditions:

1. $n \leq 2$,

2. $|\mathcal{A}(x)| \leq 2$,

3. f_i affine for all $i \in \mathcal{E} \cup \mathcal{I}$,

4. f is convex, c_i, $i \in \mathcal{E}$ affine, c_j, $j \in \mathcal{I}$ convex, and x satisfies the Slater constraint qualification,

5. there are multipliers (y, y) such that (x^*, y) is a saddle point for the Lagrangian function, or

6. the rank of $\{\nabla c_i(x^*), \nabla c_j(x^*)\}_{i \in \mathcal{E}, j \in \mathcal{A}(x)}$ is $|\mathcal{E}| + |\mathcal{A}(x)| - 1$ and there is only one j for which $y_j = 0$ for every $y \in \mathcal{M}_y(x^*)$

In addition, Andreani, Martínez, and Schuverdt [6] introduce a condition called the WCR (weak constant-rank) which holds if $\{\nabla c_i(\bar{x}), \nabla c_j(\bar{x})\}_{i \in \mathcal{E}, j \in \mathcal{A}(x)}$ has constant rank for $\bar{x}$ in a neighborhood of x^*. This, while not on its own even a first-order constraint qualification, in conjunction with the MFCQ is a second-order constraint qualification that implies SO-3 holds at a local minimizer.

As the MFCQ is not a second-order constraint qualification, the set of constraint qualifications that hold if the MFCQ holds, in particular, the CPLD, RCPLD, quasinormality, CRSC, CPG, Adabie, and Guinard constraint qualifications, cannot be second-order constraint qualifications. This implies that the extensive research associated with constant positive linear dependence geometry for developing constraint qualifications is applicable only to first-order conditions.

Chapter 3

SQP Methods

3.1 SQP Methods for Constraints in All-Inequality Form

As discussed briefly in Chapter 1, a conventional sequential quadratic programming (SQP) method involves the solution of a sequence of quadratic programming subproblems in which a quadratic model of the objective function is minimized subject to a linearization of the constraints. Without loss of generality, the description of an SQP method may be simplified considerably by assuming that there are no equality constraints, i.e., all the constraints are inequalities. In this case, the QP subproblem (1.4) (Page 3) is given by

$$\begin{array}{ll} \underset{x\in\mathbb{R}^n}{\text{minimize}} & f(x_k) + \nabla f(x_k)^T(x-x_k) + \frac{1}{2}(x-x_k)^T H(x_k,y_k)(x-x_k) \\ \text{subject to} & c_i(x_k) + \nabla c_i(x)^T(x-x_k) \geq 0, \quad i = 1,2,\ldots,m, \end{array}$$

where (x_k, y_k) is the kth primal-dual iterate and $H(x,y)$ denotes the Hessian with respect to x of the Lagrangian function (1.2). If the change in variables is written as $p = x - x_k$, then an equivalent QP is given by

$$\begin{array}{ll} \underset{p\in\mathbb{R}^n}{\text{minimize}} & \nabla f(x_k)^T p + \frac{1}{2}p^T H(x_k,y_k)p \\ \text{subject to} & c_i(x_k) + \nabla c_i(x)^T p \geq 0, \quad i = 1,2,\ldots,m. \end{array} \tag{3.1}$$

This problem may be written in matrix form as

$$\begin{array}{ll} \underset{p\in\mathbb{R}^n}{\text{minimize}} & g(x_k)^T p + \frac{1}{2}p^T H(x_k,y_k)p \\ \text{subject to} & J(x_k)p \geq -c(x_k), \end{array} \tag{3.2}$$

where $g(x) = \nabla f(x)$, and $J(x)$ is the Jacobian matrix of the vector-valued function $c(x)$. The Lagrange multipliers for the QP subproblem (3.2) may be regarded as estimates of the

Lagrange multipliers of the nonlinear problem and are called the *QP multipliers*. If the QP multipliers are written in the form $y_k + q_k$ (so that q_k defines the *change* to the dual variables analogous to the change p_k in the primal variables), then the first-order conditions for $(x_k + p_k, y_k + q_k)$ to be a primal-dual solution of the QP (3.2) are:

$$\begin{aligned} g_k + H_k p_k &= J_k^T(y_k + q_k), & y_k + q_k &\geq 0, \\ (y_k + q_k) \cdot (J_k p_k + c_k) &= 0, & J_k p_k + c_k &\geq 0, \end{aligned} \tag{3.3}$$

where $g_k = g(x_k)$, $H_k = H(x_k, y_k)$, $c_k = c(x_k)$, and $J_k = J(x_k)$. The second-order necessary conditions require the additional condition that the QP Hessian H_k has nonnegative curvature on the surface of the active constraints at $x_k + p_k$. Let $\mathcal{A}(x)$ denote the index set of the active constraints of the QP subproblem (3.1). Similarly, let J_A denote the matrix of rows of J_k with indices in $\mathcal{A}(x)$. Then the curvature condition is equivalent the reduced Hessian $Z_A^T H_k Z_A$ being positive semidefinite, where the columns of Z_A form a basis for the null-space of J_A, i.e., $J_A Z_A = 0$.

In a line-search SQP method, the new iterate is defined as $x_{k+1} = x_k + \alpha_k p_k$, where p_k is a solution of (3.2) and α_k is a nonnegative scalar step length. The step length is chosen to reduce the value of a *merit function*, whose value provides a measure of the distance of the point $x_k + \alpha_k p_k$ to a solution x^* of the nonlinear problem. Other SQP methods use a trust-region or filter strategy to define the new point x_{k+1} based on the QP step p_k. Various trust-region and filter SQP methods are described in [10, 43, 71].

The step length α_k is included to ensure "progress" at every iteration, since the current approximation of the Lagrangian function and/or the constraints may be inaccurate when the current iterate is far from x^*. In line search methods for unconstrained and linearly constrained optimization, the value of the objective function f alone provides a "natural" measure to guide the choice of α_k. Matters are more complicated when solving a nonlinearly constrained problem. Except in a few special cases, it is impossible to generate a feasible sequence of iterates with decreasing values of the objective function.

The most common approach is to choose the step length α_k to yield a "sufficient decrease" (in the sense of Ortega and Rheinboldt [73, 74]) in a *merit function* M that measures progress toward the solution of the constrained problem. Typically, a merit function is a combination of the objective and constraint functions. An essential property of a merit function is that it should always be possible to achieve a sufficient decrease in M when the search direction is defined by the QP subproblem (3.2). A desirable feature is that the merit function should not restrict the "natural" rate of convergence of the SQP method, i.e., $\alpha_k = 1$ should be accepted at all iterations "near" the solution, in order to achieve quadratic convergence (see Section 3.1.1 below). An intuitively appealing feature is that x^* should be an unconstrained minimizer of M. A feature with great practical importance is that calculation of M should not be "too expensive" in terms of evaluations of the objective and constraint functions and/or their gradients.

A commonly used merit function is the ℓ_1 *penalty function*:

$$M(x) = P_1(x,\rho) = f(x) + \rho \sum_{i=1}^{m} \max\{-c_i(x), 0\} = f(x) + \rho \|c(x)^-\|_1, \tag{3.4}$$

where ρ is a nonnegative penalty parameter. (Han [52], first suggested use of this function as a means of "globalizing" an SQP method.) This merit function has the property that, for ρ sufficiently large, x^* is an unconstrained minimizer of $M_1(x,\rho)$. In addition, ρ can always be chosen so that the SQP search direction p_k is a descent direction for $P_1(x,\rho)$. However, requiring a decrease in M_1 at every iteration can lead to the inhibition of superlinear convergence (the "Maratos effect"; see Maratos [64]), and various strategies have been devised to overcome this drawback (see, e.g., Chamberlain et al. [16]). In practice, the choice of penalty parameter in (3.4) can have a substantial effect on efficiency.

3.1.1 Convergence rate of SQP

An important issue associated with an SQP method is the rate of convergence of the iterates $\{x_k\}$ to a local minimizer. Given an m-vector v, let v_A denote the vector of elements v_i such that $i \in \mathcal{A}(x_k + p_k)$, where $\mathcal{A}(x_k + p_k)$ denotes the active set at an optimal solution $x_k + p_k$ of the QP subproblem (3.2). Similarly, let $J_A(x_k)$ denote the matrix of rows of $J(x_k)$ with indices in $\mathcal{A}(x_k + p_k)$. With this notation, the matrix-vector quantities associated with the first-order optimality conditions (3.3) for the QP subproblem may be written in the matrix form:

$$\begin{pmatrix} H(x_k, y_k) & -J_A(x_k)^T \\ J_A(x_k) & 0 \end{pmatrix} \begin{pmatrix} p_k \\ q_A \end{pmatrix} = - \begin{pmatrix} g(x_k) - J_A(x_k)^T y_A \\ c_A(x_k) \end{pmatrix}. \tag{3.5}$$

These conditions are the Newton equations for finding a stationary point of the equality constrained problem defined in terms of the active set, i.e., (3.5) are the Newton equations for solving the equations $F(x, y_A) = 0$, where

$$F(x, y_A) = \nabla L_A(x, y_A) = \begin{pmatrix} g(x) - J_A(x)^T y_A \\ -c_A(x) \end{pmatrix}.$$

Robinson [81] has shown that if x_k is sufficiently close to a local minimizer x^* that satisfies the linear independence constraint qualification (LICQ), then the QP solution has the same active set as the solution of the nonlinear problem. The LICQ implies that the active set remains constant in a neighborhood of a solution and $J_A(x_k)$ has full row rank. If, in addition, the second-order sufficient conditions hold, then the equations (3.5) are nonsingular and the QP direction is equivalent to that defined by Newton's method associated with an equality-constraint nonlinear problem defined by the optimal active set. Under these circumstances, Newton's method is quadratically convergent. Hence, under ideal conditions, a conventional SQP method with a global optimization procedure that does not suffer from the Maratos effect is quadratically convergent.

3.1.2 Difficulties associated with conventional SQP methods

Nonconvex QP subproblems. As mentioned in the introductory section of Chapter 1, the problem of finding a local solution of a nonconvex QP is an NP-hard problem. The principal difficulty occurs at so-called *dead-points*, where the QP second-order necessary conditions for optimality hold, but the second-order sufficient conditions do not. At such points, the verification of optimality is equivalent to verifying the copositivity of a symmetric matrix (see, e.g., Forsgren, Gill and Murray [31]). Because of this, many existing SQP methods implemented as software use a positive-definite approximations to the Hessian. Unfortunately, if second derivatives are not used to define the QP direction the property of quadratic convergence is lost. Furthermore, since any negative curvature of the reduced Hessian is discarded, there is no readily available away to check whether, and ensure progress towards, a stationary point satisfying the *second-order* optimality conditions. This implies that the sequence of SQP iterates may converge to a local maximizer or saddle point.

The effects of rank-deficiency (I). For the Newton equations to be well defined, the active-constraint gradients must be linearly independent, otherwise the matrix is singular. Linear independence of the active-constraint gradients implies that the LICQ holds at the limit point of the sequence of iterates. For many practical problems, the LICQ does not hold at a local minimizer, in which case the equations (3.5) are singular with no unique solution. In this situation, one remedy is to use a *stabilized SQP method*, which is based on solving the QP subproblem:

$$\begin{array}{ll} \underset{x,y}{\text{minimize}} & g_k^T(x-x_k) + \frac{1}{2}(x-x_k)^T H(x_k,y_k)(x-x_k) + \frac{1}{2}\mu_k\|y\|^2 \\ \text{subject to} & c_k + J_k(x-x_k) + \mu_k(y-y_k) \geq 0, \end{array} \tag{3.6}$$

where μ_k is a scalar parameter such that $\mu_k \to 0$ (see, e.g., Wright [86], Hager [51], Li and Qi [62], and Fernández and Solodov [25]).

The effects of rank-deficiency (II). If the active set $\mathcal{A}$ associated with a solution of the subproblem (3.2) is known, then $x_k + p_k$ may be found by solving the Newton equations (3.5). In general, however, the optimal QP active set is not known in advance. The general class of *active-set QP methods* are based on the observation that the equations (3.5) represent the optimality conditions for an equality-constrained quadratic program (EQP) in which the constraints in the active set are fixed as equalities. This suggests an algorithm in which an estimate of the QP active set is used to define an EQP with solution satisfying a system analogous to (3.5). If a component of the Lagrange multiplier vector for the EQP is negative, then the estimate of the optimal active set is updated by excluding the constraint with the negative multiplier.

If equations of the form (3.5) are to be used to solve for estimates of p_k and q_A, then it is necessary that J_A have full rank, which is probably the greatest outstanding issue associated with methods that solve systems of the form (3.5). Two remedies are available.

- *Rank-enforcing active-set methods* maintain a set of row indices $\mathcal{W}$ associated with a matrix J_W of full rank, i.e., the rows of J_W are linearly independent. The set $\mathcal{W}$ defines a "working set" of indices that estimates the set $\mathcal{A}$ at a solution of (3.5). If $\mathcal{N}$ is a subset of $\mathcal{A}$, then the system analogous to (3.5) is given by

$$\begin{pmatrix} H(x_k, y_k) & -J_W(x_k)^T \\ J_W(x_k) & 0 \end{pmatrix} \begin{pmatrix} p_k \\ q_W \end{pmatrix} = -\begin{pmatrix} g(x_k) - J_W(x_k)^T y_W \\ c_W(x_k) \end{pmatrix}. \tag{3.7}$$

  which is nonsingular because of the linear independence of the rows of J_W. (More details of rank-enforcing active-set methods are given in Section 3.1.3 below.)

- *Regularized active-set methods* include a nonzero diagonal regularization term in the $(2,2)$ block of has (3.5). The magnitude of the regularization is generally based on heuristic arguments that give mixed results in practice. The formulation and analysis of regularized methods are discussed in Chapter 5.

3.1.3 Active-set methods for a QP subproblem in all-inequality-form

The SQP methods described in thesis utilize the properties of *primal-feasible inertia-controlling active-set methods* for quadratic programming (see Gill et al. [37], and Gill and

Wong [40]). An important feature of these methods is that once a feasible iterate is found, all subsequent iterates are feasible. The methods have two phases. In the first phase (called the *feasibility phase* or *phase one*), a feasible point is found by minimizing the sum of infeasibilities. In the second phase (the *optimality phase* or *phase two*), the quadratic objective function is minimized while feasibility is maintained. Each phase generates a sequence of inner iterates $\{p_j\}$ such that $J(x_k)p_j \geq -c(x_k)$. The new iterate p_{j+1} is defined as $p_{j+1} = p_j + \sigma_j d_j$, where the *step length* σ_j is a nonnegative scalar, and d_j is the *QP search direction*. For efficiency, it is beneficial if the computations in both phases are performed by the same underlying method. The two-phase nature of the algorithm is reflected by changing the function being minimized from a function that reflects the degree of infeasibility to the quadratic objective function. For this reason, it is helpful to consider methods for the optimality phase first.

As the equations (3.5) may be singular for a general QP, inertia-controlling methods approximate $\mathcal{A}$ by a *working set* $\mathcal{W}$ of m_W row indices associated with a *linearly independent subset* of the rows of J. Analogous to the active-constraint matrix J_A, the m_W by n working-set matrix J_W contains the gradients of the constraints in $\mathcal{W}$. Once feasible, inertia-controlling active-set methods solve the equations

$$\begin{pmatrix} H & -J_W^T \\ J_W & 0 \end{pmatrix} \begin{pmatrix} d_j \\ q_W \end{pmatrix} = -\begin{pmatrix} g - J_W^T y_W \\ c_W - Jp_j \end{pmatrix}, \tag{3.8}$$

where $g = g(x_k)$, $H = H(x_k, y_k)$, $J = J(x_k)$, and c_W are the elements of the vector $c = c(x_k)$ corresponding to indices in the working set. At each QP iteration the working set is chosen in such a way that the implicit equality constrained problem associated with the working set constraints has a well-defined minimizer.

Definition 3.1.1 (Subspace stationary point)**.** *Let $\mathcal{W}$ be a working set defined at a point p. Then p is a* subspace stationary point with respect to $\mathcal{W}$ (or, equivalently, with respect to J_W) *if $g + Hp \in \operatorname{range}(J_W^T)$, i.e., there exists a vector y such that $g + Hp = J_W^T y$. Equivalently, p is a subspace stationary point with respect to the working set $\mathcal{W}$ if the reduced gradient $Z_W^T(g + Hp)$ is zero, where the columns of Z_W form a basis for the null-space of J_W.*

At a subspace stationary point, the components of y are the Lagrange multipliers associated with a QP with equality constraints $J_W p = -c_W$. The identity $g + Hp = J_W^T y$ holds at a subspace stationary point.

To classify subspace stationary points based on curvature information, we define the terms *second-order-consistent working set* and *subspace minimizer*.

Definition 3.1.2 (Second-order-consistent working set)**.** *Let $\mathcal{W}$ be a working set associated with a point p, and let the columns of Z_W form a basis for the null-space of J_W. The working set $\mathcal{W}$ is* second-order-consistent *if the reduced Hessian $Z_W^T H Z_W$ is positive definite.*

The inertia of the reduced Hessian is related to the inertia of the $(n + m_W) \times (n + m_W)$ KKT matrix $K = \begin{pmatrix} H & J_W^T \\ J_W & 0 \end{pmatrix}$ through the identity $\mathrm{In}(K) = \mathrm{In}(Z_W^T H Z_W) + (n, m_W, 0)$ (see Gould [45]). It follows that an equivalent characterization of a second-order-consistent working set is that K has inertia $(n, m_W, 0)$. A KKT matrix K associated with a second-order-consistent working set is said to have "*correct inertia*". It is always possible to impose sufficiently many *temporary constraints* that will convert a given working set into a second-order consistent working set. For example, a temporary vertex formed by fixing variables at their current values will always provide a KKT matrix with correct inertia.

Definition 3.1.3 (Subspace minimizer)**.** *If p is a subspace stationary point with respect to a second-order-consistent basis $\mathcal{W}$, then p is known as a* subspace minimizer with respect to $\mathcal{W}$*. If every constraint in the working set is active, then p is called a* standard subspace minimizer*; otherwise p is called a* nonstandard subspace minimizer*.*

Suppose that p_j is a subspace minimizer. If y_W is nonnegative, then p_j is the solution of the QP subproblem because conditions (3.3) are satisfied and $Z_W^T H Z_W$ is positive definite. Otherwise, there is at least one strictly negative component of y_W (say, the i-th), and there exists a feasible descent direction d_j, such that $g^T d_j + d_j^T H d_j < 0$ and $J_W d_j = e_i$, where e_i is the i-th column of the identity matrix. Movement along d_j causes the i-th constraint in the working set to become strictly satisfied. The direction d_j is computed by solving equations (3.8) with the shifted right-hand side $c_W + e_i - Jp_j$. At this stage, two situations are possible. The point $p_j + \sigma_j^* d_j$ such that σ_j^* minimizes the quadratic along d_j may *violate* a constraint (or several constraints) not currently in the working set. In order to remain feasible, a nonnegative step $\bar{\sigma}_j < \sigma_j^*$ is determined such that $\bar{\sigma}_j$ is the largest step that retains feasibility. A constraint that becomes satisfied exactly at $p_j + \bar{\sigma}_j d_j$ (defined as a blocking constraint) is then "added" to the working set (i.e., J_W includes a new row). If the blocking constraint is dependent on the constraints in $\mathcal{W}_j$ then the blocking constraint is swapped with the working-set constraint corresponding to $(y_W)_i$. In either of these situations, it can be shown that $p_{j+1} = p_j + \bar{\sigma}_j d_j$ is a subspace minimizer with respect to the new working set (see Gill and Wong [40]).

If there is no blocking constraint, then the feasible point $p_{j+1} = p_j + \sigma_j^* d_j$ is a subspace minimizer with respect to J_W and the iteration may be repeated at p_{j+1} with the constraint corresponding to $(y_W)_i$ deleted from the working set. Methods of this general structure will converge to a local solution of the QP subproblem in a finite number of iterations if at every iteration the active set has full rank,

It is always possible to impose sufficiently many *temporary constraints* that will covert a given working set into a second-order consistent working set. For example, a temporary vertex formed by fixing variables at their current values will always provide a KKT matrix with correct inertia. Alternatively, if $\bar{J}$ is a full-rank matrix of active constraints such that $\bar{J}Z = 0$, and

$$Z^THZ = \begin{pmatrix} Z_1 & Z_2 \end{pmatrix}^T H \begin{pmatrix} Z_1 & Z_2 \end{pmatrix} = \begin{pmatrix} Z_1^THZ_1 & Z_1^THZ_2 \\ Z_2^THZ_1 & Z_2^THZ_2 \end{pmatrix}, \text{ with } Z_1^THZ_1 \text{ positive definite,}$$

then $Z_2^Tp = 0$ is an appropriate set of temporary constraints.

3.2 SQP Methods for Constraints in Standard Form

Every nonlinear program may be defined in the form

$$\underset{x\in\mathbb{R}^n}{\text{minimize}}\ f(x) \quad \text{subject to } c(x) = 0, \quad \ell \le x \le u,$$

where $c(x)$ is a vector of m nonlinear constraint functions $c_i(x)$, and ℓ and u are vectors of lower and upper bounds. For example, an all-inequality constraint problem may be written in the form

$$\underset{x\in\mathbb{R}^n}{\text{minimize}}\ f(x) \quad \text{subject to } c(x) - s = 0, \ \ s \ge 0, \tag{3.9}$$

where s are a set of *slack* variables. Without loss of generality, we will consider a simpler form of problem in which the upper bounds on x are omitted, i.e.,

$$\underset{x\in\mathbb{R}^n}{\text{minimize}}\ f(x) \quad \text{subject to } c(x) = 0, \quad x \ge 0. \tag{3.10}$$

This problem format is known as *standard form*. The vector-pair (x^*, y^*) is a first-order solution to this problem if it satisfies

$$c(x^*) = 0 \ \text{ and } \ \min\left(x^*, z^*\right) = 0, \tag{3.11}$$

where y^* and z^* are the Lagrange multipliers associated with the constraints $c(x) = 0$ and $x \ge 0$ respectively, with $z^* = g(x^*) - J(x^*)^Ty^*$.

Given an estimate (x_k, y_k) of a primal-dual solution of (3.10), a line-search SQP method computes a search direction p_k such that $x_k + p_k$ is the solution (when it exists) of the quadratic program

$$\begin{array}{ll} \underset{x}{\text{minimize}} & g_k^T(x - x_k) + \tfrac{1}{2}(x - x_k)^T H_k(x - x_k) \\ \text{subject to} & c_k + J_k(x - x_k) = 0, \quad x \ge 0, \end{array} \tag{3.12}$$

where c_k, g_k, J_k and H_k denote the quantities $c(x)$, $g(x)$, $J(x)$ and $H(x, y)$ evaluated at (x_k, y_k). If the Lagrange multiplier vector associated with the constraint $c_k + J_k(x - x_k) = 0$ is written in the form $y_k + q_k$, then a solution $(x_k + p_k, y_k + q_k)$ of the QP subproblem (3.12) satisfies the optimality conditions

$$c_k + J_k p_k = 0 \quad \text{and} \quad \min\left(x_k + p_k,\, g_k + H_k p_k - J_k^T(y_k + q_k)\right) = 0,$$

which are analogous to (3.11). Given any $x \ge 0$, let $\mathcal{A}$ and $\mathcal{F}$ denote the index sets

$$\mathcal{A}(x) = \{i \mid x_i = 0\} \;\text{ and }\; \mathcal{F}(x) = \{1, 2, \ldots, n\}/\mathcal{A}(x). \tag{3.13}$$

If x is feasible for the constraints $c_k + J_k(x - x_k) = 0$, then $\mathcal{A}(x)$ is the *active set* at x. With these definitions, the optimality conditions analogous to (3.5) for all-inequality form are

$$\begin{pmatrix} H_F & -J_F^T \\ J_F & 0 \end{pmatrix} \begin{pmatrix} p_F \\ q_k \end{pmatrix} = -\begin{pmatrix} [g_k - J_k^T y_k]_F \\ c_k \end{pmatrix}, \tag{3.14}$$

where p_F, H_F and J_F are the components of p_k, J_k and H_k associated with the indices in $\mathcal{F}(x_k + p_k)$. These conditions represent the Newton equations for finding a stationary point of the equality constrained problem defined in terms of the free variables. If H_k and J_F have full rank in a neighborhood of a solution, then Newton's method converges at a quadratic rate.

For problems in standard form, rank-enforcing active-set methods maintain a set of indices $\mathcal{B}$ associated with a matrix of columns J_B with rank m, i.e., the rows of J_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 of (3.12). In this case, the system analogous to (3.14) is given by

$$\begin{pmatrix} H_B & -J_B^T \\ J_B & 0 \end{pmatrix} \begin{pmatrix} p_B \\ q_k \end{pmatrix} = -\begin{pmatrix} [g_k - J_k^T y_k]_B \\ c_k \end{pmatrix}, \tag{3.15}$$

which is nonsingular because of the linear independence of the rows of J_B. More details are provided in the following section.

3.2.1 Active-set methods for QPs in Standard Form

The QP subproblem associated with the NLP (3.10) with constraints written as in standard form is given by

$$\begin{array}{ll} \underset{p\in\mathbb{R}^n}{\text{minimize}} & g(x_k)^T p + \tfrac{1}{2} p^T H(x_k, y_k) p \\ \text{subject to} & J(x_k)p = -c(x_k), \quad x_k + p \geq 0. \end{array} \tag{3.16}$$

Let J be $m \times n$. An active set method keeps a set of m_W fixed *nonbasic* variables with indices in $\mathcal{N}$ and $n_B = n - m_W$ free *basic* variables with indices in $\mathcal{B}$. Let E_N^T denote the rows of the identity corresponding to the nonbasic variables. The working-set matrix is

$$J_W = \begin{pmatrix} J \\ E_N^T \end{pmatrix}. \tag{3.17}$$

If P is a permutation such that $E_N^T P = \begin{pmatrix} 0 & I_N \end{pmatrix}$, then

$$J_W P = \begin{pmatrix} J_B & J_N \\ 0 & I_N \end{pmatrix} \quad \text{and} \quad P^T H P = \begin{pmatrix} H_B & H_D \\ H_D^T & H_N \end{pmatrix}.$$

Definition 3.2.1 (Subspace stationary point)**.** *If the following conditions hold,*

$$\begin{aligned} g + Hp &= J^T \pi + z \\ z_B &= [g + Hp]_B - J_B^T \pi = 0 \\ z_N &= [g + Hp]_N - J_N^T \pi, \end{aligned}$$

then p is a subspace stationary point*. Equivalently, p is a subspace stationary point if the reduced gradient $Z_W^T(g + Hp)$ is zero, where Z_W is a basis for the null-space of J_W.*

Definition 3.2.2 (Second-order consistent working set)**.** *If $Z_B^T H Z_B$ is positive definite, or, equivalently, the matrix*

$$K_B = \begin{pmatrix} H_B & J_B^T \\ J_B & \end{pmatrix}$$

has inertia $(n_B, m, 0)$, then the set $\mathcal{N}$ of nonbasic indices forms a second-order consistent working-set*.*

Definition 3.2.3 (Subspace minimizer)**.** *If p is a subspace stationary point, and the current working-set is second-order consistent , then p is also a* subspace minimizer*.*

Moreover, if $z_N \geq 0$, then p is a local minimizer of(3.16).

3.2.2 Temporary bounds

Let p_0 be a stationary point for (3.16) and H be indefinite. In general the point p_0 may not be at a second-order consistent working set, however, one may be formed by introducing temporary bounds on the constraints.

Decompose the null-space of J_W:

$$Z = N(J_W P) = \begin{pmatrix} N(J_B) \\ 0_{m_W} \end{pmatrix}.$$

Attempt to form a inertia-revealing factorization (such as Cholesky or block diagonal with interchanges) of $Z^T H_F Z$. At some point the factorization stops, resulting in the partially factorized matrix

$$MPZ^T H_F Z P^T M^T = \begin{pmatrix} D & 0 \\ 0 & U \end{pmatrix},$$

where, after permutation, the first m_z columns of Z satisfy $[Z\bar{P}^T M^T]_1^T H [Z\bar{P}^T M^T]_1 = D$ which is a positive diagonal matrix. Let $Z\bar{P}^T = \begin{pmatrix} Z_1 & Z_2 \end{pmatrix}$. Similarly as in the all-inequality case, we can make $Z_2^T p = 0$ an temporary constraint to construct a second-order consistent working set.

Chapter 4

Stability and Convergence

4.1 Background

For a given scalar or vector quantity ζ, consider an optimal solution $x^*(\zeta)$ of the parameterized problem

$$\begin{array}{lll} \underset{x}{\text{minimize}} & f(x,\zeta) & \\ \text{subject to} & c_i(x,\zeta) = 0, & i \in \mathcal{E}, \\ & c_i(x,\zeta) \geq 0, & i \in \mathcal{I}, \end{array} \tag{4.1}$$

associated with the problem (2.1). The study of the stability of the problem (2.1) is the study of the solutions of (4.1) as $\zeta \to \zeta^*$, where ζ^* is a value such that $x^* = x^*(\zeta^*)$. In general, let ζ be such that $\zeta \in \Pi$, with Π a Banach space and f and c are both redefined as $f : \mathbb{R}^n \times \Pi \mapsto \mathbb{R}$ and $c : \mathbb{R}^n \times \Pi \mapsto \mathbb{R}^m$. Throughout this chapter, f and c are assumed to be three-times Frechet differentiable in both x and ζ.

As an example, consider ζ, $\zeta^* \in \mathbb{R}^{n+m}$, with $\zeta^* = (x^*, y^*)$, $\zeta = (\bar{x}, \bar{y})$, and functions

$$\begin{aligned} f(x,\zeta) &= f(\bar{x}) + g(\bar{x})^T(x-\bar{x}) + \tfrac{1}{2}(x-\bar{x})^T H(\bar{x},\bar{y})(x-\bar{x}) \\ c_i(x,\zeta) &= c_i(\bar{x}) + \nabla c_i(\bar{x})^T(x-\bar{x}), \quad i = 1,2,\ldots,m. \end{aligned}$$

In this case, (4.1) is a quadratic program that has a local solution (x^*, y^*) for $\zeta = \zeta^*$.

If x^* satisfies the optimality conditions of (4.1) for some ζ^*, an important topic is the existence and characterization of solutions $x(\zeta)$ for (4.1) for ζ in a neighborhood of ζ^*, including existence, uniqueness, and bounds on $\|x(\zeta) - x^*\|$ and $\|y(\zeta) - y^*\|$.

Stability, as the study of solutions to perturbed problems is called, assists in the convergence theory of algorithms. For global convergence, the stability properties that hold under different constraint qualifications can be used to show that as the steps p_k in some sequential algorithm ($x_{k+1} = x_k + p_k$) approach zero, the iterate x_k approaches a KKT point. For local convergence, continuity assumptions together with stability are used to show convergence rates, as the subproblem begins to resemble a framework or other specific algorithm known to converge at the desired rate.

A *multifunction* is a multivalued (set-valued) function. A multifunction $r : \mathbb{R} \mapsto \mathbb{R}$ may take on one value, no values, or multiple values for any x. A multifunction $r(x)$ is *upper semicontinuous at* x_0 if for every neighborhood $\mathcal{B}_u$ of $r(x_0)$, there exists a neighborhood $\mathcal{B}_x$ such that for every $x \in \mathcal{B}_x$, it holds that $r(x) \subset \mathcal{B}_u$. If a multifunction is upper semicontinuous for every point $x \in \mathcal{B}_x$, then it is said to be upper semicontinuous on the set $\mathcal{B}_x$ (see, e.g., Bonnans and Shapiro [12]).

Let $\mathcal{U}$ be a neighborhood $\mathcal{U} \subset \mathbb{R}^n$ and consider the quantity

$$\begin{aligned}\operatorname{cdist}(\zeta,\mathcal{U}) = \sup_i \sup_{x\in\mathcal{U}} \max\Big\{&|c_i(x,\zeta) - c_i(x^*)|, \|\nabla c_i(x,\zeta) - \nabla c_i(x^*)\|,\\ &\|\nabla^2 c_i(x,\zeta) - \nabla^2 c_i(x^*)\|\Big\},\end{aligned}$$

which measures the perturbation in the constraints and their derivatives induced by the perturbation ζ.

The set $\mathcal{A}_+(y^*)$ contains the indices corresponding to the positive elements of y^* (see Section 2.4.2, page 26). Let $\mathcal{A}_+$ and $\mathcal{A}_0$ denote the sets $\mathcal{A}_+ = \cup_{y^*\in\mathcal{M}_y(x^*)} \mathcal{A}_+(y^*)$ and $\mathcal{A}_0 = \mathcal{A} \setminus \mathcal{A}_+$.

Definition 4.1.1. *The* second-order sufficiency condition (SOSC) *holds at* x^* *if for all* $y^* \in \mathcal{M}_y(x^*)$, $w^T\nabla^2_{xx}L(x^*, y^*)w \ge \sigma\|w\|^2$ *for all* w *such that* $\nabla c_i(x^*)^T w = 0$ *for* $i \in \mathcal{E} \cup \mathcal{A}_+$, *and* $\nabla c_i(x^*)^T w \ge 0$ *for* $w \in \mathcal{A}_0$.

Definition 4.1.2. *The* strong second-order sufficiency condition (SSOSC) *holds at* x^* *if for all* $y^* \in \mathcal{M}_y(x^*)$, $w^T\nabla^2_{xx}L(x^*, y^*)w \ge \sigma\|w\|^2$ *for all* w *such that* $\nabla c_i(x^*)^T w = 0$ *for* $i \in \mathcal{A}_+ \cup \mathcal{E}$.

4.2 Robinson's Theorem and the LICQ

Robinson [81] establishes a set of stability results that hold under the linear independence constraint qualification (LICQ). Nonlinear optimization convergence theorems frequently refer to these results because of their strength and reliability. As these results are the strongest stability results in the literature, and are relatively simple to apply, many proofs of global and local convergence assume that the LICQ holds at cluster points of the sequence of iterates.

Robinson analyzes the solutions of a perturbed problem near a point x^* satisfying the LICQ, strict complementarity, and the sufficient second-order optimality conditions. He proves the following main result, which implies that that for small perturbations, there is always a solution of the perturbed problem that is close to the original solution.

Theorem 4.2.1 (Robinson [81, Theorem 2.1])**.** *Let x^* satisfy the first-order necessary and the second-order sufficient conditions for optimality of* (4.1) *with $\zeta = \zeta^*$. Furthermore, let the LICQ and strict complementarity hold at x^*. Then there are open neighborhoods $\mathcal{V}(\zeta^*, \epsilon_\zeta) \subset \Pi$ and $\mathcal{U}(x^*, y^*, \epsilon_x) \subset \mathbb{R}^n \times \mathbb{R}^m$ with $\zeta^* \in \mathcal{V}$ and $(x^*, y^*) = z^* \in \mathcal{U}$ and a continuous function $G : \mathcal{V} \mapsto \mathcal{U}$ such that for each $\zeta \in \mathcal{V}$, $G(\zeta) = (x, y) = z(\zeta)$ uniquely satisfies the first-order KKT conditions of the perturbed problem* (4.1)*. Furthermore, second-order sufficiency, complementary slackness and the LICQ hold at $z(\zeta) = (x(\zeta), y(\zeta))$.*

The proof relies on the implicit function theorem and the local stability of linear independence. The implicit function theorem requires a nonsingular Jacobian at x^*, which implies that it is not possible to relax the assumption that the LICQ holds.

Robinson also derives a bound for points close to solutions of (4.1). Specifically, consider the constraints arranged sequentially with the first $|\mathcal{E}|$ constraints being equalities and the remaining constraints inequalities. Consider the vector defined as a concatenation of all of the quantities that are required to be zero at a KKT point,

$$e(x, y, \zeta) = [L_x(x, y, \zeta), c_1(x, \zeta), \ldots, c_{|\mathcal{E}|}(x, \zeta), y_{|\mathcal{E}|+1} c_{|\mathcal{E}|+1}(x, \zeta), \ldots, y_m c_m(x, \zeta)]^T,$$

measuring the distance of the current values of the problem functions from optimality. Denote the Jacobian matrix $de(z, \zeta)/dz|_{(\bar{z}, \bar{\zeta})} = e'_z(\bar{z}, \bar{\zeta})$.

The following holds:

Theorem 4.2.2 (Robinson [81, Theorem 2.2])**.** *Consider $\mathcal{U}$, $\mathcal{V}$, and G defined as in Theorem 4.2.1. For any scalar $\epsilon > 0$, there are open neighborhoods $\mathcal{U}_\epsilon \subset \mathcal{U}$ and $\mathcal{V}_\epsilon \subset \mathcal{V}$, such that for any $\tilde{\zeta} \in \mathcal{V}_\epsilon$ and $\tilde{z} \in \mathcal{U}_\epsilon$,*

$$\|\tilde{z} - G(\tilde{\zeta})\| \leq \frac{1}{1 - \epsilon} \|e'_z(z^*, \zeta^*)^{-1}\| \|e(\tilde{z}, \tilde{\zeta})\|.$$

This theorem implies that the distance of a point $\tilde{z}$ to the solution to a problem perturbed by $\tilde{\zeta}$ is bounded by the size of the violation of the first-order optimality conditions of the perturbed problem at the point $\tilde{z}$. This is later used to derive convergence rates of iterative algorithms, since, as illustrated by the introductory examples, subproblems can be defined as perturbations of the original nonlinear problem. Specifically, Robinson [81] then makes the following Lipschitz continuity assumption:

$$\|e(G(z), G(z)) - e(G(z), z))\| \le \alpha \|G(z) - z\|^{\lambda},$$

and is able to prove linear (for $\lambda = 1$) and superlinear ($\lambda > 1$) convergence for a class of Newton-like recursive algorithms.

4.2.1 Application to convergence theory

As an illustrative example of the applicability of these results, consider the sequential quadratic programming of Gill et al. [36]. Their algorithm solves nonlinear programs by solving a series of quadratic program subproblems

$$\begin{array}{ll} \underset{p}{\text{minimize}} & g(x_k)^T p + \frac{1}{2} p^T H(x_k, y_k) p \\ \text{subject to} & J(x_k)p + c(x_k) \ge 0. \end{array} \tag{4.2}$$

where x_k is the point at the current iteration. In addition, to force iterates towards a local minimizer from an arbitrary starting point, they include an augmented Lagrangian merit function,

$$\phi(x, y, s; \rho) = f(x) + y^T(c(x) - s) + \tfrac{1}{2}\rho\|c(x) - s\|^2,$$

where s are the slack variables and ρ a penalty parameter. Choosing an α that reduces the merit function sufficiently, the algorithm sets $x_{k+1} = x_k + \alpha p_k$, $y_{k+1} = y_k + \alpha q_k$ and $s_{k+1} = s_k + \alpha r_k$, where p_k is the solution, q_k are the change in the multipliers $q_k = \pi_k - y_k$ for (4.2), and r the change in the slack variables.

For global convergence, Gill et al. [36] cite Robinson [81] to show that if a step p_k is sufficiently small, then the active set at x_k is the same as the active set at x^*, the nearest KKT point. They then use the full row rank of the linearization matrices of the SQP subproblem (which holds under the LICQ) to bound $\|x^* - x_k\|$ by $\|p_k\|$. Finally, they use the properties of the merit function to show that $\|p_k\| \to 0$. Local convergence is shown by proving that eventually $\alpha = 1$ and, since the correct active set is estimated, the subproblem solves Newton's equations for the optimality conditions, as described in Chapter 3 (Page 33). Since the LICQ holds, Newton's method on the first-order KKT conditions is well-defined.

4.3 Stability under the MFCQ

4.3.1 Robinson's results for the MFCQ

Robinson also analyzed the stability of a perturbed problem under the Mangasarian-Fromovitz constraint qualification and the second-order sufficiency condition [83]. The optimality conditions are considered in the context of a generalized equation

$$0 \in F(z) + \mathcal{T}(z),$$

where $F : \mathbb{R}^l \mapsto \mathbb{R}^l$ and $\mathcal{T}$ is a multifunction $\mathcal{T} : \mathbb{R}^l \mapsto \mathbb{R}^l$. In the context of nonlinear optimization

$$0 \in \begin{pmatrix} \nabla f(x) + J(x)^T y \\ -c(x) \end{pmatrix} + \begin{pmatrix} 0 \\ \mathcal{R}(y) \end{pmatrix},$$

where $\mathcal{R}(y)$ is the cone of nonnegative vectors that are complementary to y, i.e.,

$$\mathcal{R}(y) = \begin{cases} \{b \in \mathbb{R}^m \mid b \geq 0,\, b_i y_i = 0 \text{ for all } i\} & \text{if } y \geq 0, \\ \emptyset, & \text{otherwise.} \end{cases}$$

Let x^* solve the KKT conditions for the perturbed problem (4.1) with perturbation ζ_0. The following continuity of perturbed problem solutions results hold.

Theorem 4.3.1 (Robinson [83, Theorem 3.1]). *If the MFCQ and second-order sufficiency conditions hold at x^*, then there are neighborhoods $\mathcal{U}$ and $\mathcal{V}$ of $x^* \in \mathcal{U}$ and $\zeta_0 \in \mathcal{V}$ such that $\zeta \in \mathcal{V}$ implies that the perturbed problem* (4.1) *with perturbation ζ has a solution $x \in \mathcal{U}$.*

Theorem 4.3.2 (Robinson [83, Theorem 3.2]). *Let the MFCQ and the SOSC hold at x^*. There are neighborhoods $\mathcal{U}$ and $\mathcal{V}$ of x^* and ζ_0 such that $x^* \in \mathcal{U}$, $\zeta_0 \in \mathcal{V}$ and the multifunctions $SU : \mathcal{U}_2 \times \mathcal{V}_2 \mapsto \mathbb{R}^m$ and $SP : \mathcal{V}_2 \mapsto \mathcal{U}_2$ defined as*

$$\begin{aligned} U(x,\zeta) &:= \{y \in \mathbb{R}^m \mid (x,y) \text{ satisfies the KKT conditions for} \\ &\qquad \text{the problem (4.1) with perturbation } \zeta\}, \\ SP(\zeta) &:= \{x \in \mathcal{U}_2 \mid \text{there is a multiplier } y \text{ such that } (x,y,\zeta) \\ &\qquad \text{satisfies the KKT conditions for the problem (4.1) with perturbation } \zeta\}, \end{aligned}$$

are upper semicontinuous.

Robinson establishes the following result.

Theorem 4.3.3 (Robinson [83, Theorem 3.3]). *Let $LM : \mathit{\Pi} \mapsto \mathbb{R}^n$ be defined as*

$$LM(\zeta) = \{x \mid x \text{ is a local minimizer of (4.1)}\},$$

There are neighborhoods $\mathcal{U}_3$ of x^* and $\mathcal{V}_3$ of ζ_0, for which then $SP \cap \mathcal{U}_3$, with SP as in Theorem 4.3.2, is continuous and for each $\zeta \in \mathcal{V}_3$, $LM(\zeta) \subset SP(\zeta) \cap \mathcal{U}_3$.

These results establish local existence of solutions for perturbed problems as well as provide some bound as to the distance of the perturbed problem solutions to the original solution. While weaker than Lipschitz continuity, the local boundedness associated with upper semicontinuity does provide tools for convergence theory. Robinson's results are applied by Wright [87] for his analysis of the inexact SQP subproblems discussed in Chapter 5.

4.3.2 Kojima's results for the MFCQ

Kojima [60] gives the following stability result for nonlinear programs for which the MFCQ holds at a solution x^*.

Theorem 4.3.4 (Kojima [60, Theorem 7.1])**.** *Let x^* satisfy the first-order KKT conditions and the MFCQ hold at x^*. Then the strong second order sufficiency (SSOSC) condition holds at x^** if and only if *there exists a δ^* such that for all $\delta \in (0, \delta^*]$, there is an ϵ such that when*

$$\text{cdist}(\zeta, \mathcal{B}_\delta(x^*)) \leq \epsilon,$$

$\mathcal{B}_\delta(x^*)$ contains a KKT solution $\bar{x}$ to (4.1) unique in $\mathcal{B}_\delta(x^*)$.

Corollary 4.3.1. *The MFCQ and strong second-order sufficiency conditions hold at $\bar{x}$.*

Corollary 4.3.2. *The minimizer x^* is an isolated KKT point.*

The proof relies on the fact that under the MFCQ, the set of multipliers is bounded, and in particular, it is a convex polyhedron. At an extreme point of the polyhedron, the active constraint gradients for the equalities and the inequality constraints corresponding to positive multipliers are linearly independent, and so he is able to apply the same results he had developed earlier in his paper using degree theory for the case of the LICQ.

Kojima does not provide rate of convergence results. In particular, Kojima shows continuity but not Lipschitz continuity of the solutions subject to perturbations. In the proof of Lipschitz continuity under the LICQ in Robinson [81], the Lipschitz constant is the inverse of the local linearization of the optimality conditions. As the constraint gradients may be linearly dependent, this inverse may or may not exist. Robinson [83] gives an example where Lipschitz continuity in the solutions does not hold when MFCQ holds but the LICQ does not hold.

The results of Kojima [60] and Robinson [83] imply that, while under the MFCQ, for small perturbations, solutions should be expected to exist, no precise bound can be made in proportion to the perturbation, regardless of the smoothness of the original problem functions. To contrast the two results, Robinson [83] shows local boundedness of perturbed problem solutions, while Kojima [60] shows local uniqueness, the reverse implication that local existence implies the SSOSC and the local preservation of the MFCQ and SSOSC, under the stronger assumption of the SSOSC.

4.3.3 Convergence theory for the MFCQ

Qi and Wei [78] provide convergence theory for both generic and specific SQP methods that use a constraint qualification that is weaker than the LICQ. Consider the *approximate KKT sequence*,

Definition 4.3.1. *A sequence of primal-dual iterations* $\{(x_k, y_k)\}$ *is an* approximate KKT sequence *if it holds that*

$$\begin{aligned} \nabla f(x_k) + J(x_k)^T y_k &= \epsilon_k \\ c_i(x_k) &\geq \delta_k, \quad i \in \mathcal{I} \\ [y_k]_i &\geq 0, \quad i \in \mathcal{I} \\ [y_k]_i c_i(x_k) - \delta_k &= 0, \quad i \in \mathcal{I} \\ \|c_{i\in\mathcal{E}}(x_k)\| &\leq \nu_k, \end{aligned}$$

where $\{\epsilon_k, \delta_k, \nu_k\}$ *converges to zero as* $k \to \infty$.

If an approximate KKT sequence converges to x^* and the constant positive linear dependence condition (CPLD) holds at x^* then x^* is a KKT point. This implies that for a generic SQP method with a well-defined line-search, if the sequence has a cluster point x^* and, along the sequence converging to x^*, $\liminf \|p_k\| \to 0$, then x^* is a KKT point.

They require an additional assumption, namely that the second order sufficiency condition holds at x^*, for showing results for sequences as opposed to subsequences. This is because the second-order sufficiency conditions together with the CPLD imply that KKT points are isolated. This fact is central to the proof that if $\|p_k\| \to 0$ then $x_k \to x^*$.

They then analyze the Panier-Tits [75] SQP method, wherein a generic SQP step is combined with a step of descent and a second-order correction, using a line-search to ensure global convergence. They assume that the CPLD holds at a limit point x^* and that the MFCQ holds at all other points. The MFCQ condition implies that feasible sufficient descent directions at non-KKT points will always exist. This fact is used by Qi and Wei to show global convergence (i.e., $x_k \to x^*$, where x^* is a KKT point).

Qi and Wei [78] then adjust the algorithm and decompose the constraints into a subset with linearly independent constraint gradients and the rest, and assuming, in addition, the CRCQ (implying that these sets are locally maintained), they show two-step superlinear convergence, using the results of Kojima [60] to show that the perturbed problems arising from selecting subsets of constraints are still feasible.

Qi and Wei [78] are able to use a number of weaker-than-LICQ assumptions to show that the KKT conditions hold at limit points, global convergence, and local convergence, with increasingly stronger conditions. As such, the results are some of the strongest among the convergence theory in the literature.

4.4 Convergence Under Weak Constraint Qualifications

This section describes additional global convergence results that rely on constraint qualifications that are weaker than the MFCQ.

In their paper on the constant rank subspace component (CRSC) and constant positive generator (CPG) conditions, Andreani et al.[4] are able to show that a convergent approximate KKT sequence converges to a KKT point x^* if the CPG holds at x^*. This allows them to generalize the results from Qi and Wei to assume the CPG instead of the CPLD where appropriate. Specifically, if there is a convergent subsequence $x_k \to x^*$, $\liminf d_k = 0$ and the CPG holds at x^*, then x^* is a KKT point. Likewise, if the CPG holds at x^* and the MFCQ at all other points, then the generic SQP algorithm converges ($x_k \to x^*$ with x^* a KKT point).

Andreani et al. [4] are also able to show a stronger result (assuming just that the CPG holds at a cluster point x^*) for an augmented Lagrangian method. They apply the standard Conn, Gould, and Toint [19] augmented Lagrangian for problems with general equality constraints and bound constraints on the variables. They show that by the inherent properties of the augmented Lagrangian function, the algorithm converges to a stationary point of $\|c(x)\|^2$. This implies that one of two cases occur, convergence to a feasible point or convergence to an infeasible minimum of the sum of squares of the constraints. It is shown that a sequence with a feasible limit point is an approximate KKT sequence, which implies that the feasible limit point is a KKT point.

As there are no continuity-of-solution estimates, however, there are no known local convergence results for these weaker constraint qualifications. However, in Chapter 5 it will be seen that there are stability and resultant convergence results for a class of algorithms that rely strictly on assumptions of second-order sufficiency.

4.5 Second-order Convergence Theory

Traditionally, algorithms have been constructed strictly to generate sequences of iterates converging to first-order optimal points. Most algorithms calculate directions of descent that decrease an objective or merit function until there is no feasible descent direction without being concerned with the curvature along that direction. However, it is also possible to calculate directions of negative curvature that decrease the objective or merit function until the curvature in the proper space is positive-definite. Computing a direction of negative curvature and verifying second-order optimality is far more computationally challenging than finding a descent direction and checking for first-order optimality. Ultimately, since both types of directions result in a decrease in the objective or a merit function, in practice most strictly first-order algorithms end up rarely terminating at a saddle-point or local maximum. Nevertheless, as algorithms in constrained nonlinear optimization become increasingly successful, progress in the field entails solving a wider class of problems to full optimality. Furthermore, it has been demonstrated that directions of negative curvature can considerably decrease the required number of iterations for convergence to solutions of unconstrained problems. A few algorithms demonstrating convergence to second-order local minima for constrained problems have been formulated and will be discussed here.

4.5.1 Computing directions of negative curvature

Recall from Chapter 2 (Page 25) that $\widetilde{\mathcal{C}}(x) = \{d \mid \nabla c_i(x)^T d = 0,\, i \in \mathcal{E},\, \nabla c_j(x)^T d = 0, j \in \mathcal{A}\}$. Convergence to a second-order optimal point requires computing directions of negative curvature for the reduced Hessian matrix $Z^T H Z$, where H is L_{xx}, and Z is a basis for the space corresponding to $\widetilde{\mathcal{C}}(x)$.

This is done by performing a factorization of $Z^T H Z$ which reveals its inertia, calculating a direction of negative curvature from the remaining portion of the matrix to be factorized, then performing a curvilinear line search or trust-region step using both the direction of descent and the direction of negative curvature. The two most common such factorizations are the partial Cholesky and the LBL^T symmetric indefinite factorization.

Forsgren et al. [33] describe a method in which a Cholesky factorization with pivoting is performed until there are no remaining potential positive pivots. The remaining pivot element is a direction of negative curvature, whereas the completed portions of the factorization correspond to directions of positive curvature for the particular reduced Hessian.

Forsgren [29] describes a method using a modification of the symmetric indefinite LBL^T factorization to reveal the inertia of a KKT matrix. By specific selection of pivots, the inertia of the initial phase of the factorization remains “correct” and modification in the form of adding the norm of the remaining Schur complement to the components of the Hessian matrix remaining to be factorized corrects the inertia while only changing H and not the other blocks of the KKT matrix.

These procedures are discussed in more detail in Chapter 9.

4.5.2 A second-order exact Lagrangian method

DiPillo et al. [22] present a second-order augmented Lagrangian algorithm for inequality constrained problems. They use the following augmented Lagrangian function, an alteration of the standard Hestenes-Powell-Rockafeller function to make it exact:

$$L_a(x,y,\rho) = f(x) + y^T \max(c_i(x), -\rho r(x,y)y) + \frac{\|\max(c(x), -\rho r(x,y)y)\|^2}{2\rho r(x,y)}$$
$$+ \|\nabla c(x)^T \nabla_x L(x,y) + \sum \nabla^2 c_i(x) y_i\|^2,$$

where the function $r(x,y) = (\alpha - \|c(x)_-\|_p^p)/(1+\|y\|^2)$, where $\|a\|_s$ is defined as $(\sum_i [a]_i^p)^{1/p}$ for some integer p.

Assuming the LICQ, they show that if there is a positive-definite matrix $W \in \partial_B^2 L_a(x^*, y^*; \rho)$ (notice that L_a is not twice continuously differentiable), then (x^*, y^*) satisfy the second order necessary optimality conditions. They define a the matrix Q as capturing an appropriate amount of second-order information of the problem approximating W. For each step of the inner augmented Lagrangian iteration, they perform a curvilinear line-search:

$$x + \alpha^2 d + \alpha s,$$

where $d^T \nabla L_a \le 0$ and $s^T Q s \le 0$ and $s^T \nabla L_a \le 0$. They derive that $s_k^T Q s_k \to 0$ implies that the underlying sequence of iterates converges to a second-order optimal point, and show that the properties of the line-search ensures this limit holds.

4.5.3 A second-order standard augmented Lagrangian method under weak regularity assumptions

Andreani et al. [1] manage to also prove second-order convergence for an augmented Lagrangian algorithm, but use much weaker assumptions, in particular the MFCQ and the WCR condition.

In solving the augmented Lagrangian subproblem, their inner iterations calculate a descent and a negative-curvature step. They then choose the one that provides for the largest reduction in the value of the augmented Lagrangian. The inner iterates stop when the augmented Lagrangian has both vanishing gradient and positive curvature in the reduced space. The outer iteration is the standard augmented Lagrangian multiplier, penalty parameter and tolerances updating procedure.

In their convergence proofs, they use a technical lemma from Andreani et al. [6] that states that for any element of $d \in \widetilde{\mathcal{C}}(x^*)$, there is a sequence of $d_k \in \widetilde{\mathcal{C}}(x_k)$ converging to d (where $x_k \to x^*$). The lemma allows them to prove that limit points of the sequences of iterates generated by the algorithm satisfy the second-order necessary conditions for optimality, and will be used in the second-order global convergence proofs of SQP2d (Chapter 7) and pdSQP (Chapter 9).

Note that the cone in which positive semidefinite curvature is shown to hold is $\widetilde{\mathcal{C}}(x^*)$, so a constraint qualification certifying the weak semidefinite reduced property (WSRP) is sufficient. There are currently no known factorization methods of computing directions of negative on a cone as opposed to a subspace (as in, directions satisfying $\nabla c_i(x)^T p \geq 0$ for $i \in \mathcal{A}_0$).

Chapter 5

Regularized SQP Methods

5.1 Introduction

Regularized SQP methods are a class of methods that involve slightly altered quadratic programs as subproblems. They have been shown to converge at a superlinear rate for degenerate problems. There are several regularized SQP algorithms, the most common one being the *stabilized SQP*.

Except where otherwise noted, this chapter will concern problems with inequality constraints only. The problem is written in the form

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

At any x, the gradient of f and the Jacobian of c will be denoted by $g(x)$ and $J(x)$ respectively. Similarly, the Hessian of the Lagrangian function with respect to x for given multipliers y will be denoted by $H(x, y)$.

The convergence theory for stabilized SQP requires some form of second-order sufficiency assumption. The three principal forms that appear in the literature are presented below in order of increasing strictness.

As defined in Section 2.4.2, page 26, $\mathcal{A}_+(x^*)$ contains the indices corresponding to the positive elements of y^*. Similarly, $\mathcal{A}_+$ and $\mathcal{A}_0$ are the sets $\mathcal{A}_+ = \cup_{y^* \in \mathcal{M}_y(x^*)} \mathcal{A}_+(y^*)$ and $\mathcal{A}_0 = \mathcal{A} \setminus \mathcal{A}_+$.

Definition 5.1.1. *The* second-order sufficiency condition (SOSC) *holds at x^* if for all $y^* \in \mathcal{M}_y(x^*)$, $d^T H(x^*, y^*) d \geq \sigma \|d\|^2$ for all d such that $\nabla c_i(x^*)^T d = 0$ for $i \in \mathcal{A}_+$, and $\nabla c_i(x^*)^T d \geq 0$ for $d \in \mathcal{A}_0$.*

Definition 5.1.2. *The* relaxed second-order sufficiency condition (RSOSC) *holds at* (x^*) *if for some* $y^* \in \mathcal{M}_y(x^*)$, $d^T H(x^*, y^*) d \geq \sigma \|d\|^2$ *for all* d *such that* $\nabla c_i(x^*)^T d = 0$ *for* $i \in \mathcal{A}_+$, *and* $\nabla c_i(x^*)^T d \geq 0$ *for* $d \in \mathcal{A}_0$.

Sometimes, it will also be said that the RSOSC holds at a particular (x, y), implying that the RSOSC holds and, in particular, this particular y is among the multipliers satisfying the condition.

Definition 5.1.3. *The* strong second-order sufficiency condition (SSOSC) *holds at* x^* *if for all* $y^* \in \mathcal{M}_y(x^*)$, $d^T H(x^*, y^*) d \geq \sigma \|d\|^2$ *for all* d *such that* $\nabla c_i(x^*)^T d = 0$ *for* $i \in \mathcal{A}_+ \cup \mathcal{E}$.

Definition 5.1.4. *The* relaxed strong second-order sufficiency condition (RSSOSC) *holds at* x^* *if for some* $y^* \in \mathcal{M}_y(x^*)$, $d^T H(x^*, y^*) d \geq \sigma \|d\|^2$ *for all* d *such that* $\nabla c_i(x^*)^T d = 0$ *for* $i \in \mathcal{A}_+ \cup \mathcal{E}$.

Note that these cones also correspond to cones in the definitions for the weak and strong second-order necessary conditions.

The relaxed conditions are clearly weaker, but this is only because the Hessian $H(x^*, y^*)$ is a different matrix for each $y^* \in \mathcal{M}_y(x^*)$, the space is directions are actually equivalent.

Lemma 5.1.1 (Wright [89, Lemma 2.1])**.** *The set of directions defined by:*

$$\{d \mid \nabla c_i(x^*)^T d = 0 \text{ for } i \in \mathcal{A}_+ \cup \mathcal{E}\}$$

is equivalent to the set:

$$\{d \mid \nabla c_i(x^*)^T d = 0 \text{ for } i \in \mathcal{A}_+(y^*) \cup \mathcal{E}\}$$

for all $y^* \in \mathcal{M}_y(x^*)$.

Regularized SQP methods often use a regularization parameter μ. The parameter μ is most commonly defined in terms of the optimality condition violation:

$$\eta(x, y) = \left\| \begin{pmatrix} g(x) - J(x)^T y \\ [\, c(x)\,]_- \\ c(x)^T y \end{pmatrix} \right\|. \tag{5.1}$$

Typically $\mu = \eta(x, y)^\tau$, where $0 < \tau \leq 1$. Some methods require τ to be strictly less than one, but if $\tau = 1$ is permitted, then the local convergence rate is quadratic. If $\tau < 1$, the local convergence rate is superlinear.

It can be shown that $\eta = \Theta\big(\delta(x,y)\big)$ where $\delta(x,y)$ is the distance to the nearest first-order KKT point,

$$\delta(x,y)^2 = \|x - x^*\|^2 + \inf_{y^* \in \mathcal{M}_y(x^*)} \|y - y^*\|^2.$$

In what follows, we define $\delta_1(x)$ and $\delta_2(y)$ to be $\delta_1(x) = \|x - x^*\|$ and $\delta_2(y) = \operatorname{dist}(y, \mathcal{M}_y(x^*))$.

5.2 The Inexact SQP Method iSQP

First introduced by Wright [87], there is a class of methods called inexact SQP (iSQP) that involve iterated solutions of subproblems perturbed in some way from the standard SQP subproblem.

A number of methods can be fit into the iSQP framework including the first regularized iSQP method of Fischer [26], and the stabilized SQP method described below. In addition, many commercial SQP solvers do not solve the conventional subproblem exactly, and have certain features that can be formalized as involving stabilized working-sets [87].

Throughout this section, assume that the MFCQ holds at all KKT points.

The iSQP subproblem at a given (x, y) is:

$$\begin{aligned} &\underset{p \in \mathbb{R}^n}{\text{minimize}} \;\; (g(x) + t)^T p + \tfrac{1}{2} p^T H(x,y) p \\ &\text{subject to} \;\; J(x)p + r \geq -c(x), \end{aligned} \tag{5.2}$$

where $t \in \mathbb{R}^n$ and $r \in \mathbb{R}^m$ are perturbations that may also depend on x and y..

Assume there is a local minimizer x^*. Consider an initial primal-dual estimate (x_0, y_0).

Theorem 5.2.1 (Wright [87, Lemma 5.1])**.** *If the SOSC and the MFCQ hold at (x^*, y^*), a first-order KKT point, then for all (x_0, y_0, t, r) with $\delta(x_0, y_0)$ and $\|(t, r)\|$ sufficiently small, the iSQP subproblem has a local solution $(p, y_1) = (x_1 - x_0, y_1)$ near $(0, y_0)$ that satisfies*

$$\|p\| + \delta_1(y_1) = O\big(\delta_1(x_0)\big) + O\big(\|(t,r)\|\big).$$

This result may be derived from Robinson's general stability theorem (see Robinson [83, Theorem 3.2]), which is given as Theorem 4.3.2 in Section 4.3.1, page 46. Consider the perturbation parameter as $\zeta = (x_0, y_0, t, r)$ and let the base perturbation be $\zeta_0 = (x^*, y^*, 0, 0)$, the results of the Theorem implies the existence of solutions to the perturbed problem, and the locally bounded continuity of solutions estimate bounds $\|p\|$ and $\|y_1 - y_0\|$.

Consider a general iSQP algorithm wherein at each iteration, the iSQP subproblem is solved and the step p closest to zero is chosen when there are multiple solutions. The following local convergence rate holds.

Theorem 5.2.2 (Wright [87, Theorem 5.3])**.** *For* $\delta(x, y)$ *and* $\|(t, r)\|$ *sufficiently small,*

$$\delta(x_{k+1}, y_{k+1}) = \|y_{k+1} - y_k\| \, O\big(\delta_1(x_k)\big) + O\big(\delta_1(x_k)^2\big) + O\big(\|(t, r)\|\big).$$

Fischer [26] introduces a general framework of generalized modified Newton equations and a method that computes a step from the solution of two quadratic programs. It can be shown that the method fits into the iSQP framework and superlinear convergence can be proven. Fischer's method requires the solution to two QPs, and whereas it is one of the earliest regularized SQP methods in the literature, it does not resemble the others and requires stronger assumptions for convergence.

Stabilized SQP, which also fits the context of the iSQP framework, will be discussed in detail in the next few sections. The specific results achieved by using the iSQP framework for stabilized SQP is presented as reference Wright [87] in Table 5.1.

There is one additional algorithm to discuss that fits in the iSQP framework, *stabilized working sets.*

5.2.1 Stabilized working sets

In practice, the state-of-the-art SQP optimization packages rarely converge linearly or worse to optimal points for most mildly degenerate problems. It appears as though this is a result of two standard features of these algorithms:

- Use of the working set from the solution of the previous QP subproblem as an initial working set for the current one.
- Allowing constraints not in the working set to be violated by small tolerances.

The stabilized working-set algorithm is an attempt to formalize these features [87]. A stack $\{\mathcal{B}_s, \mathcal{B}_{s-1}, \ldots, \mathcal{B}_0 = \{1, \ldots, m\}\}$ of working sets is maintained with $\mathcal{B}_l \subset \mathcal{B}_{l-1}$. At each major iteration, the subproblem

$$\begin{array}{ll} \underset{p \in \mathbb{R}^n}{\text{minimize}} & g(x_k)^T p + \frac{1}{2} p^T H(x_k, y_k) p \\ \text{subject to} & c_i(x_k) + \nabla c_i(x_k)^T p \geq 0, \quad i \in \mathcal{B}_s. \end{array} \tag{5.3}$$

is solved. If the primal solution p violates any of the constraints not in $\mathcal{B}_s$ by more than $\mu_k^{1+\tau}$, the subproblem is solved again with the constraints from the working set $\mathcal{B}_{s-1}$ and the process is repeated.

Otherwise, $\mathcal{B}_{s+1}$ is defined as the final working set for the subproblem (5.3), and added to the stack of working sets. The multipliers for the constraints not in $\mathcal{B}_{s+1}$ are set to zero.

Consider the *extended iSQP* subproblem, which is defined to be the iSQP subproblem (5.2) defined with a subset of constraints $\mathcal{B}$. The following result holds:

Lemma 5.2.1 (Wright [87, Lemma 8.2])**.** *If the SOSC and the MFCQ hold at x^*, then there exists a value $\bar{\delta}$ such that if $\delta(x,y) \le \bar{\delta}$ and $\|(t, r_{\mathcal{B}})\| \le \bar{\delta}$ and $y_i = 0$ for $i \notin \mathcal{B}$, for $\mathcal{B} = \mathcal{A}_+(y^*)$ for some y^* such that $\{\nabla c_{\mathcal{B}}(x^*))\}$ are linearly independent, then the corresponding extended iSQP subproblem has at least one solution that satisfies, for this y^*,*

$$\|p\| + \|y - y^*\| = O\big(\|x - x^*\|\big) + O\big(\|t, r_{\mathcal{B}}\|\big).$$

The next theorem concerning convergence of the stabilized working sets algorithm follows from this result.

Theorem 5.2.3 (Wright [87, Theorem 8.3])**.** *Assume the SOSC and the MFCQ hold at x^*, then there exists a $\bar{\delta}$ such that if $\delta(x,y) \le \bar{\delta}$, and there is a $\mathcal{B}$ in the stack of the stabilized working sets algorithm satisfying the assumptions of the previous lemma, then this $\mathcal{B}$ remains in the stack for all subsequent iterations and the algorithm converges superlinearly to (x^*, y^*), y^* the multiplier referenced in the previous lemma, to order $1+\tau$.*

Consider an additional procedure of dropping linearly dependent constraints to ensure linear independence of the working sets at each iteration. Upon adding this feature to SQPsws, the following result holds.

Theorem 5.2.4 (Wright [87, Corollary 8.4])**.** *Suppose that the LSSOSC, the MFCQ, and the the CRCQ are satisfied at x^*. There is a $\bar{\delta}$ such that if $\delta(x,y) \le \bar{\delta}$, the stabilized working sets algorithm converges superlinearly to a primal-dual solution.*

It should be noted that for the SQPsws method, the parameter τ must satisfy $0 < \tau < 1$, which implies that only superlinear, and not quadratic, convergence can be proven.

5.3 Stabilized SQP

The most common and well known regularized SQP method is the *stabilized SQP method*, introduced first by Wright [86]. Dropping the dependence of $g(x)$ and $J(x)$ on x and $H(x,y)$ on (x,y) in the notation, the sSQP algorithm involves the following subproblem:

$$\min_{p} \max_{y \geq 0} p^T g + \tfrac{1}{2} p^T H p + (y+q)^T (c + Jp) - \tfrac{1}{2}\mu \|q\|^2,$$

which may be rewritten as

$$\begin{aligned} \underset{p\in\mathbb{R}^n, q\in\mathbb{R}^m}{\text{minimize}} \quad & g^T p + \tfrac{1}{2} p^T H p + \tfrac{1}{2}\mu \|y+q\|^2 \\ \text{subject to} \quad & Jp + \mu q \geq -c. \end{aligned} \tag{5.4}$$

Let the initial point (x_0, y_0) be sufficiently close to a first-order KKT point (x^*, y^*). There are a number of convergence results on stabilized SQP. In the first result, quadratic convergence is shown to hold if $\mu = \eta$ for both exact and finite arithmetic under the assumptions of the MFCQ, strict complementarity, and the SOSC (assumption 5.1.1) condition holding at (x^*, y^*), a first-order KKT point (see Wright [86]). Subsequently, quadratic convergence was proven assuming just that the strong second order sufficiency condition (assumption 5.1.3) holds (see Hager [51]).

Hager [51] proves quadratic convergence under the assumption that the relaxed strong second order sufficiency condition (RSSOSC) holds (see Definition 2.4.5, page 27). Instead of requiring μ to be the measure of optimality violation η explicitly, μ is relaxed to require that $\sigma_0\|x - x^*\| \leq \mu \leq \sigma_1$, where σ_0 is sufficiently large and σ_1 depends on σ_0. With $\mu = \eta^\tau$, Hager requires τ to satisfy $0 < \tau \leq 1$. Since τ may equal one, quadratic convergence can be proven.

The strength of the stabilized SQP method relies on the properties of the perturbed KKT matrix

$$\begin{pmatrix} H(x,y) & -J(x)^T \\ J(x) & \mu I \end{pmatrix}$$

which can be nonsingular even when $J(x)$ is not full rank. The smaller the value of μ, the closer the method is to a conventional SQP method, while at the same time the closer the KKT matrix is to being ill-conditioned if the constraints are degenerate.

5.4 Indefinite sSQP Subproblems

One issue that many of the convergence results do not address is the fact that in using the exact Hessian, the sSQP subproblem,

$$\begin{aligned} \underset{p\in\mathbb{R}^n, q\in\mathbb{R}^m}{\text{minimize}} \;\; & g^Tp + \tfrac{1}{2}p^THp + \tfrac{1}{2}\mu\|y+q\|^2 \\ \text{subject to} \;\; & Jp + \mu q \geq -c, \end{aligned}$$

may be nonconvex, even close to a solution, and hence could have multiple, if not an unbounded set of solutions. This can be the case even close to a point satisfying the second-order sufficiency conditions, for which the reduced Hessian at the optimal active set may be positive-definite, but not the Hessian of the subproblem.

This implies that there there are implicit assumptions in the strong local convergence results of the sSQP literature in terms what the "sSQP method" means, such as that at each iteration, either the global minimizer or the minimizer closest to the current point or just any minimizer is taken as the iteration step. On the other hand, perhaps in penalizing large multipliers, sSQP results in the active-set estimate being maintained to be optimal once y is close to a y^* satisfying the optimality conditions.

This section will discuss an illustrative example and discuss the results in the literature related to this issue.

5.4.1 Example

Consider the NLP:

$$\begin{aligned} \underset{x\in\mathbb{R}^2}{\text{minimize}} \;\; & x_1^2 - x_2^2 \\ \text{subject to} \;\; & -2 \leq x_2 \leq 2. \end{aligned} \tag{5.5}$$

The objective gradient is $g = \begin{pmatrix} 2x_1 \\ -2x_2 \end{pmatrix}$ and the Hessian is $\begin{pmatrix} 2 & 0 \\ 0 & -2 \end{pmatrix}$ and hence is always indefinite. Consider the bound constraints to be general inequality constraints, and consider the inequality-constrained sSQP subproblem. The Jacobian matrix is $J = \begin{pmatrix} 0 & 1 \\ 0 & -1 \end{pmatrix}$.

The stabilized SQP subproblem becomes:

$$\begin{aligned} \underset{p\in\mathbb{R}^n, q\in\mathbb{R}^m}{\text{minimize}} \;\; & \begin{pmatrix} 2x_1 & -2x_2 \end{pmatrix}^T p + \tfrac{1}{2}p^T \begin{pmatrix} 2 & 0 \\ 0 & -2 \end{pmatrix} p + \tfrac{1}{2}\mu\|y+q\|^2 \\ \text{subject to} \;\; & p_2 + (x_2+2) + \mu q_1 \geq 0, \\ & -p_2 + (2-x_2) + \mu q_2 \geq 0. \end{aligned} \tag{5.6}$$

Since $x_1 = 0$ is clearly optimal, and does not appear in the constraints, reduce the dimension of the problem by setting $x_1 = 0$, and it is clear that any solution to the subproblem has $p_1 = 0$.

The optimality conditions for the sSQP under this assumption are:

$$-2x_2 - 2p_2 = y_1 + q_1 - y_2 - q_2, \tag{5.7a}$$

$$p_2 + (x_2 + 2) + \mu q_1 \geq 0, \tag{5.7b}$$

$$-p_2 - (x_2 - 2) + \mu q_2 \geq 0, \tag{5.7c}$$

$$(y_1 + q_1)(\ \ p_2 + (x_2 + 2) + \mu q_1) = 0, \tag{5.7d}$$

$$(y_2 + q_2)(-p_2 - (x_2 - 2) + \mu q_2) = 0, \tag{5.7e}$$

$$y_1 + q_1 \geq 0, \tag{5.7f}$$

$$y_2 + q_2 \geq 0. \tag{5.7g}$$

Solving the stationarity condition (5.7a) for p_2 yields

$$p_2 = -x_2 + \tfrac{1}{2}(y_2 + q_2 - y_1 - q_2).$$

As defined earlier in this chapter (Page 54), μ is defined as the optimality condition violation:

$$\eta(x, y) = \left\| \begin{pmatrix} g(x) - J(x)^T y \\ [\, c(x)\,]_- \\ c(x)^T y \end{pmatrix} \right\|.$$

For superlinear convergence, μ must be set to η^τ, with $\tau > 0$. At the same time, τ must be set to satisfy $\tau \leq 1$ for the stability results to hold (see Hager [51, Theorem 1]). For this example, set $\tau = 1$, which a typical value that guarantees quadratic convergence.

Calculate μ, assuming x stays feasible, explicitly to be

$$\mu = \sqrt{(-2x_2 + y_1 - y_2)^2 + (y_1(x_2 + 2) - y_2(2 - x_2))^2}.$$

Consider, for example, $x_2 = 2$, $y_2 = 4$. In this case, $(x, y) = (x^*, y^*)$ and all the optimality conditions for the NLP are satisfied and $\mu = 0$. In this case, the subproblem reduces to a conventional QP. This is unavoidable, since as discussed, μ must be $\mu = \eta^\tau$. As expected, it shall be seen that there are two solutions.

The problem is illustrated in Figure 5.1.

Consider the optimality conditions, using the substitution for p_2 obtained from the stationarity condition, the complementarity conditions become

$$\tfrac{1}{2}(q_1)\,(8 + q_2 - q_1) = 0,$$

$$\tfrac{1}{2}(4 + q_2)\,(-q_2 + q_1) = 0.$$

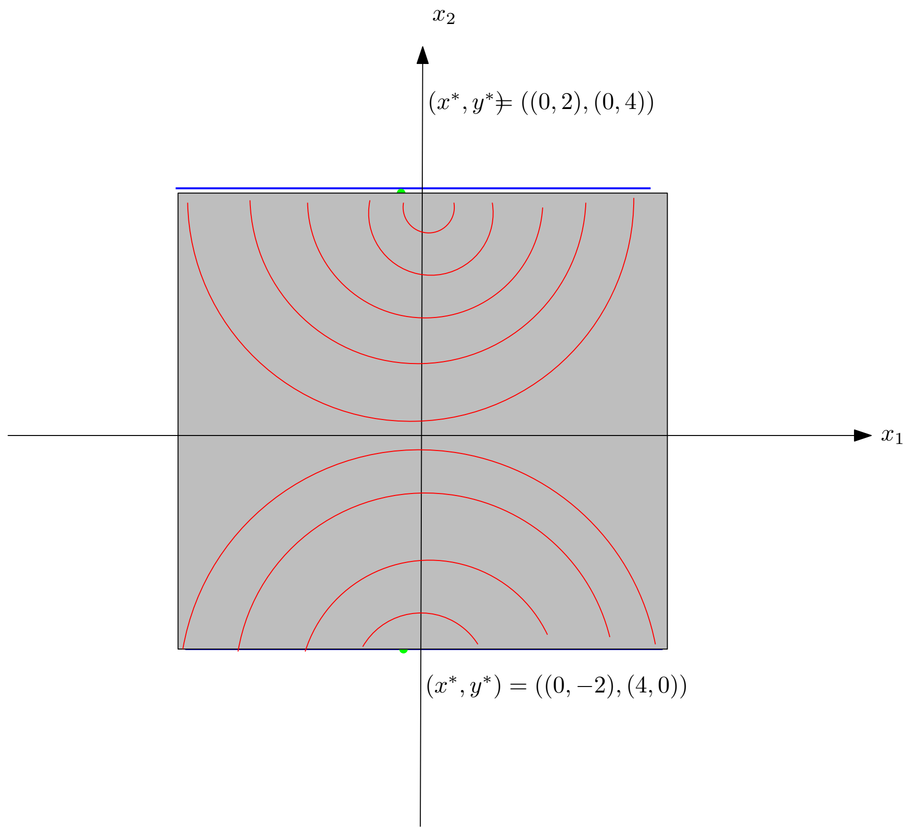

Figure 5.1: Example of nonconvex problem

Note that, in addition, each product in both equations must be nonnegative. Consider the graph of the lines in (q_2, q_1) space in Figure 5.2.

Note that a solution needs to be both feasible, and be a point at which a line A intersects a line B. It is clear that there are two solutions, $(0, 0)$, and $(-4, 4)$, corresponding to staying at the current point, and jumping to the other solution to the NLP.

It can be observed that the shape of these lines depends heavily on the value of μ. A slight perturbation of y or x changes μ considerably. For instance, if $x_2 = 1.9$, $y_1 = 0.01$, and $y_2 = 3.9$, then $\mu = 0.42$. Furthermore, there is a factor of $(1 - 2\mu)$ for the slope and intercept of the upper B line, so the shape of the line changes considerably as μ approaches one half.

If the same values for all the other quantities are maintained, but μ is set to be 0.45, then the diagram in (q_2, q_1) space changes to appearing as in Figure 5.3.

In this case, the second solution is $(-4, 40)$, a greater deviation from the original point.

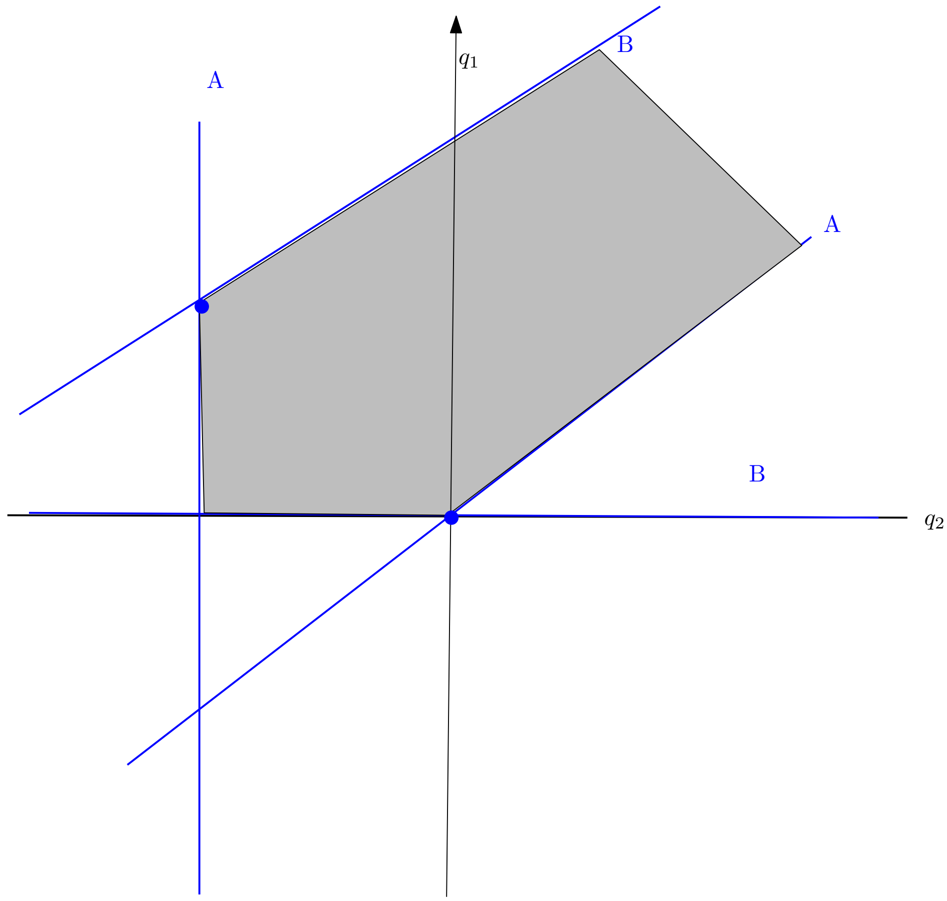

Figure 5.2: sSQP/SQP solution for example with $(x, y) = (x^*, y^*)$, $\mu = 0$.

If $\mu = 0.55$ the qualitative nature of the solutions change completely, see Figure 5.4.

Note that the second solution disappears entirely, and, for $\mu > 0.5$, there is only one unique solution to this nonconvex QP.

This suggests that practically, μ does serve as a parameter penalizing jumps between optimal active sets once y is sufficiently close to y^*. However, as $\mu \to 0$, the stabilized SQP subproblem approaches the conventional one, and is consequently more likely to contain multiple solutions.

Unfortunately, however, artificially raising μ invalidates the fast local convergence rates. Yet, as a theoretical point, it does appear that for some relatively large set of cases, it could well be that stabilized SQP subproblems, even with an indefinite Hessian, have unique solutions, and uniquely define a sequence of iterations towards an optimal point. It still remains open as to the theoretical universality of this idea, and an open problem as to

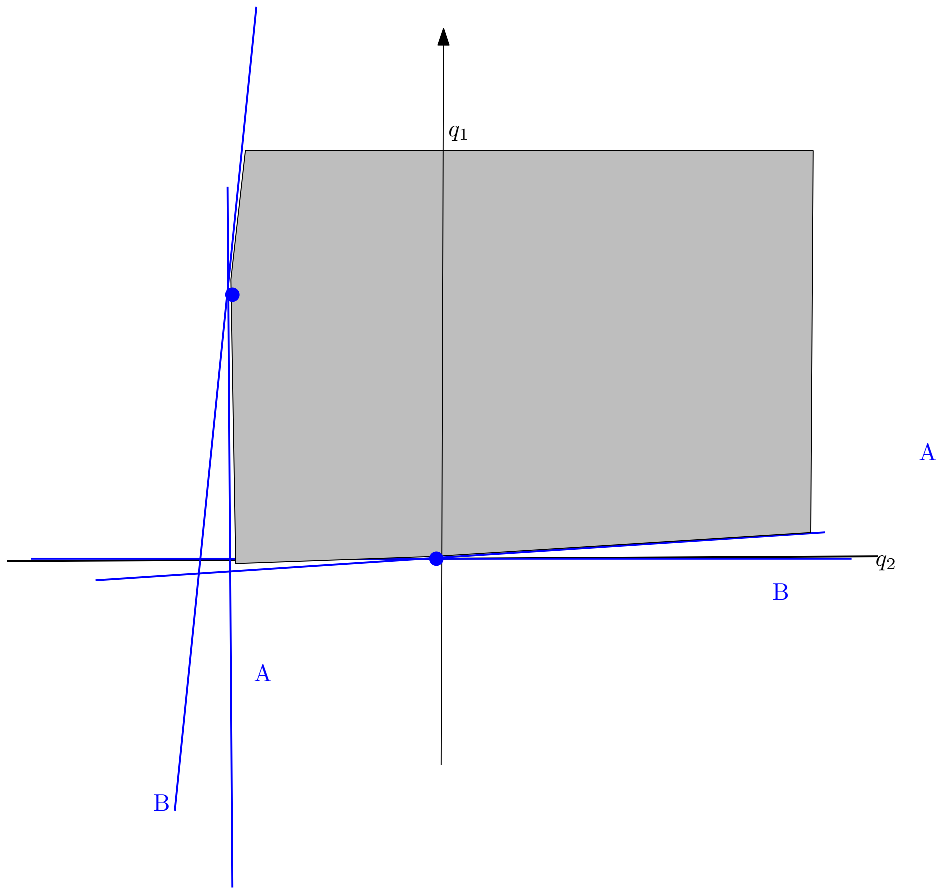

Figure 5.3: sSQP solution for example with $(x, y) = (x^*, y^*)$, $\mu = 0.45$.

how to ensure unique solutions while maintaining superlinear local convergence. However, in practice, it appears that stabilized SQP may not suffer from this drawback frequently.

5.4.2 Uniqueness results

This illustration suggests that there are implicit assumptions and possible limitations to the presentation of stabilized SQP. If there is no precisely defined "sSQP method" in a practical sense, as there could be more than one solution to each subproblem, the question arises as to what in particular the convergence results refer to.

The fine detail is that the results show that the solution satisfying the convergence estimate is *locally unique*, in that there is a neighborhood around the current point for which there is only solution to the subproblem.

In the one convergence result for sSQP, Wright [86] demonstrates that if strict

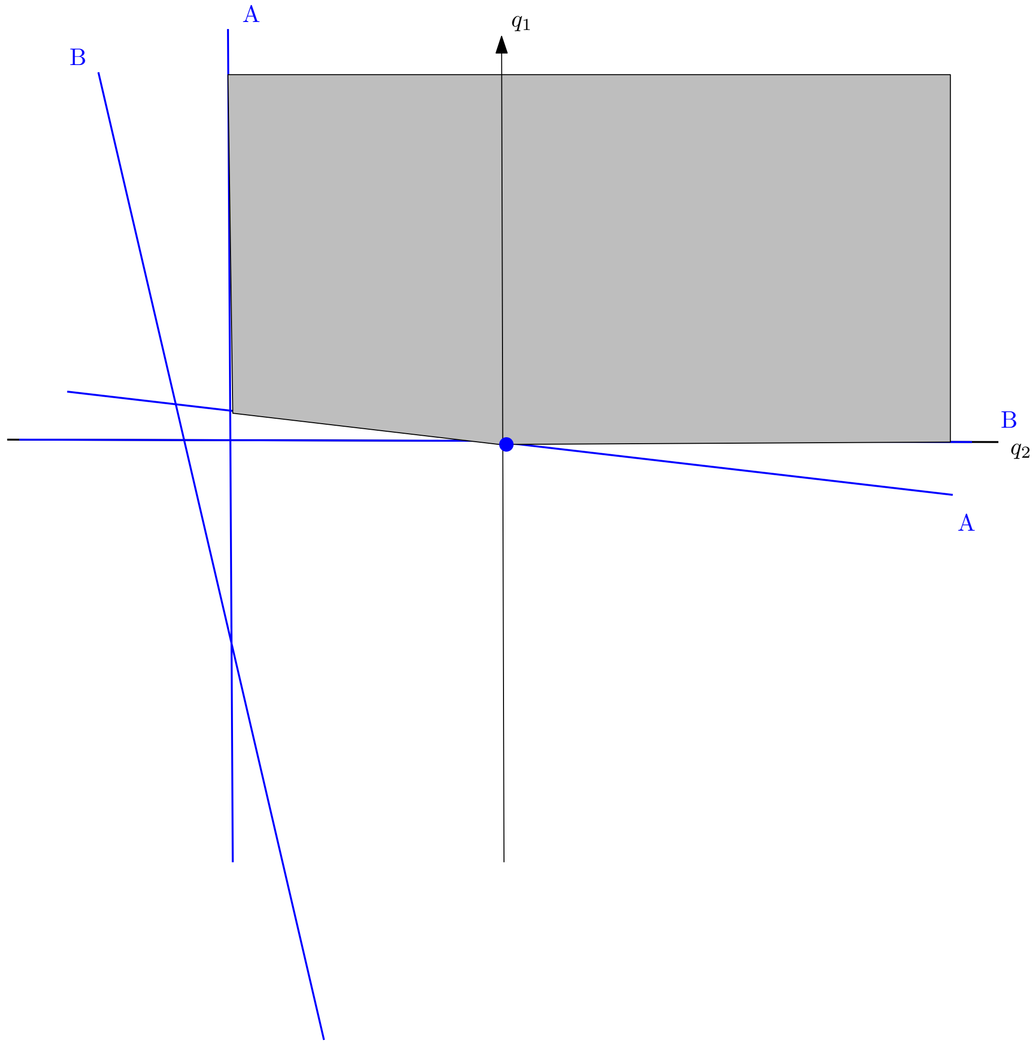

Figure 5.4: sSQP solution for example with $(x, y) = (x^*, y^*)$, $\mu = 0.55$.

complementarity holds at the x^*, then if y_0 is sufficiently strictly complementary, as in, it is close to a multiplier $y^* \in \mathcal{M}_y(x^*)$ with $\min_i y_i^* > \gamma$ for some γ, the solution to the optimality conditions with $[y_k]_{\mathcal{A}} > 0$ and $[y_k]_{\mathcal{A}^c} = 0$ satisfies

$$\|(p, y_{k+1} - y_k)\| \leq C\mu$$

and is the only solution to the stabilized SQP problem that satisfies this estimate.

On a practical level, this suggests that, if one adds a trust region constraint that limits the step in x and y, with the trust region radius being $C\mu$, then the solution to the stabilized SQP subproblem is unique. While, as an example of a set of nonconvex QPs with unique solutions, this is theoretically interesting, it is practically useless, since C depends on both the Lipschitz constants of the objective and constraint functions as well as the distance to the optimal solution, the latter of which is obviously unknown a priori.

In a different convergence result Hager [51] shows the local uniqueness of the sSQP solution. In particular, it is shown that there is a locally unique point satisfying the first-order optimality conditions for sSQP given the conditions of RSSOSC at (x^*, y^*), (x_0, y_0) being sufficiently close to (x^*, y^*) and μ bounding that distance from above. This is accomplished by considering a similar result for variational problems in Dontchev and Hager [23] and applying it for the nonlinear optimization case.

It is important to note that the required result only concerns points in a neighborhood of (x_k, y_k). It is assumed that the strictly positive multipliers *remain* strictly positive in the local linearization, which permits separating the active and inactive constraints as done in Hager [51]. This suggests that estimating the optimal active set, and maintaining that estimate, is crucial for obtaining a well-defined sequence of sSQP iterates. The only numerical implementation of sSQP [68] uses an active-set estimation procedure including several heuristics to maintain strongly active multipliers.

The implication of the local uniqueness results' dependence on constraint identification is that if the iterates converge to x^* *normal* to the active constraints, the subproblems may be nonconvex at every iteration regardless of how close the iterations get to a local minimizer satisfying any second-order sufficiency assumption.

5.5 Active Set Identification

It is a general feature of SQP methods that once the optimal active-set $\mathcal{A}(x^*)$ is identified, convergence is usually reliable and rapid. This is also evident in the first two convergence results of stabilized SQP, as the locally unique solution had the property of having the same active set as the previous iterate which, in turn, had the same active set as the closest local minimizer. This required the strong assumptions of either strict complementarity or the strong second-order sufficiency condition. This implies that weakly active constraints can make correctly identifying the active set in a well-defined algorithm more difficult.

5.5.1 Multiplier adjustment

Facchinei, Fischer and Kanzow [24] propose a method for identifying active constraints. The method does not rely on any complementarity or multiplier uniqueness. It is shown that if there is a function $\rho(x,y)$ such that $\rho(x,y) \to 0$ as $\delta(x,y) \to 0$, with $\rho(x,y)/\delta(x,y) \to \infty$ as $\delta(x,y) \to 0$, then the classification of a constraint as active if $c_i(x) \le \rho(x,y)$ is correct provided that $\delta(x,y)$ is sufficiently small.

This idea may be expanded to identify weakly and strongly active constraints. Let $\mathcal{A}(x,y)$ denote the set of estimated active constraints. Similarly, let $\widehat{\mathcal{A}} \subset \mathcal{A}(x,y)$ denote an estimate of the weakly active constraints. The ID0 Algorithm proposed by Wright [88] solves a sequence of linear programs of the form:

$$\begin{array}{ll} \underset{\widehat{y}}{\text{minimize}} & -\sum_{i\in\widehat{\mathcal{A}}} \widehat{y}_i \\ \text{subject to} & -\eta(x,y)^\tau e \le g(x) + \sum_{i\in\mathcal{A}(x,y)} \widehat{y}_i \nabla c_i(x) \le \eta(x,y)^\tau e, \quad \widehat{y}_i \ge 0 \ \text{ for } \ i \in \mathcal{A}. \end{array}$$

The estimated set of weakly active multipliers $\widehat{\mathcal{A}}$ is initialized to be $\mathcal{A}(x,y)$. Each step of the algorithm involves solving the LP above and removing all multipliers from $\widehat{\mathcal{A}}$ such that $\widehat{y}_i \ge \eta(x,y)^\tau$. The steps are repeated until no multipliers are removed from $\widehat{\mathcal{A}}$. It can be shown that Algorithm ID0 estimates the weakly and strongly active sets correctly. In the situation where it is necessary to solve many LPs, the computational cost may be reduced by solving the dual of the LP and using the final working set from one problem as the initial working set for the next.

At each iteration of the algorithm sSQPa (Wright [88]) a standard sSQP inequality constrained problem is solved. If the reduction of the optimality measures is not sufficiently large, an ID0 algorithm is used to identify the weakly and strongly active multipliers. In addition, an LP is solved to compute a multiplier such that the multiplier $\operatorname{argmin}_i \ y_i$ for $i \in \mathcal{A}_+$ is as large as possible (the so-called *interior* multiplier). The sSQP subproblem is then re-solved.

It can be shown that sSQPa converges superlinearly under the assumptions of the MFCQ and the SOSC.

The algorithm sSQPa has been implemented by Mostafa, Vicente, and Wright [68]. These numerical results, which are obtained for a small number of problems, are relatively inconsistent, with some exhibiting the desired fast local convergence properties relative to competitors, and others showing slower convergence, or no convergence at all. (As algorithm sSQPa includes no procedure to force global convergence, some of these results are to be expected.) The constraint identification procedure is relatively reliable, but does not seem to be effective for all problems, and the overall performance of sSQPa does depend on the identification procedure. While the numerical results are illustrative and show some potential for sSQP methods, they clearly indicate a need for further analysis and development.

Note that no mention is made of stability in the theoretical results, and no mention is made of addressing multiple subproblem solutions in the numerical results or Wright's presentation of sSQPa.

5.5.2 Equality-constrained subproblems

As mentioned in Chapter 1, many of the second-derivative SQP methods involve two-phase procedures of active-set estimation and equality-constrained subproblems. Noting that the sSQP subproblem matrix,

$$\begin{pmatrix} H & J^T \\ J & -\mu I \end{pmatrix},$$

is nonsingular due to the regularization parameter, once the proper active set is identified, convergence should be fast and reliable.

In one method, due to Izmailov and Solodov [54], the active set is estimated, and then the solution is sought for the equality constrained problem on the projection onto the kernel of the active constraint Jacobian. However, the method uses an expensive singular value decomposition that can introduce dense matrices in the large-scale case, and provides no safeguards in the event of inaccurate identification of the active set. However, if a point is sufficiently close to a local minimizer to the extent that the active set is correctly identified, then this algorithm converges at a quadratic rate under the assumption of SOSC only.

A more robust procedure that implements active set estimation is given by Wright [89]. This method is to be used with any reliable outer iteration procedure. Once an iterate from the outer procedure satisfies a certain tolerance for optimality, he estimates the active set and performs a sequence of equality-constrained Newton-Lagrange iterations. If, at some point, the optimality measures from these iterations are not sufficiently reduced, a constraint not in the estimated active set becomes infeasible, the multipliers become negative, or the step is too large, then the algorithm moves back to the outer iteration and decreases the required threshold to come back to the equality-constrained problem. It can be shown that once an iterate is sufficiently close to a KKT point and the correct active-set is identified, the iterations never leave the equality-constrained phase and converge quadratically to the solution. The only assumption is the second-order sufficiency condition. This relatively strong result points to the potential of integrating stabilized SQP with a robust outer convergence procedure.

5.6 Convergence and Stability of sSQP

By utilizing stability theory, it is possible to prove superlinear convergence for stabilized SQP using only second-order sufficiency assumptions. As discussed above, stabilized SQP subproblems have the effect of calming the dual iteration in the presence of nonunique dual solutions of the nonlinear problem. This statement may be made more precise in terms of specific statements about the solvability of the subproblems and relations between the distances of the solutions to the previous iterate and the nearest KKT point. These properties, in turn, are shown to be responsible for the superlinear convergence rate of stabilized SQP, independent of other stability results or constraint qualifications.

5.6.1 Fischer's iterative framework

This section discusses the result by Fischer [27], which provides a link between certain stability properties and superlinear convergence. Fischer considers the same generalized equation analyzed by Robinson [83], i.e.,

$$0 \in F(z) + \mathcal{T}(z),$$

where $F : \mathbb{R}^{l_1} \mapsto \mathbb{R}^{l_2}$ is continuous, and $\mathcal{T} : \mathbb{R}^{l_1} \rightrightarrows \mathbb{R}^{l_2}$ is a closed multifunction. Let Σ_* denote the solution set of the generalized equation, with $\Sigma_0 \subset \Sigma_*$ a nonempty closed subset of the solution set. Let $\Sigma(\zeta)$ denote the set of solutions of the perturbed equation $0 \in F(z) + \mathcal{T}(z) + \zeta$.

Consider a general iterative framework where subproblems of the form

$$0 \in Q(z, z^k) + \mathcal{T}(z),$$

are solved, where $Q(z, z^k)$ is an approximation to F. As the focus is on problems for which the solutions are not unique, it is required that the iterative method define z_{k+1} to be such that

$$\|z_{k+1} - z_k\| \le \sigma \operatorname{dist}(z_k, \Sigma_*).$$

This is an important restriction that reduces the chance of reaching a different local solution when the subproblem is not convex. Nevertheless, this condition does not imply local uniqueness of the solutions of the subproblem. The following main result holds for this procedure,

Theorem 5.6.1 (Fischer [27, Theorem 1])**.** *If the following conditions hold:*

1. *There are* ϵ_1, γ, $t > 0$ *such that, with* $Q = \Sigma_0 + \epsilon_1 B$, *it holds that*

$$\Sigma(\zeta) \cap Q \subseteq \Sigma_* + t\|\zeta\| B.$$

2. *There is a* $\epsilon_2 > 0$ *and* $c > 0$ *such that, writing* $R(w, s) = F(w) - \mathcal{A}(w, s)$,

$$\sup\{\|R(w, s)\| \mid w \in s + c \operatorname{dist}(s, \Sigma_*), B\} \leq o(\operatorname{dist}(s, \Sigma_*))\}.$$

3. *There is an* ϵ_4 *such that* $Z_c(s) \neq 0$ *for all* $s \in \Sigma_0 + \epsilon_3 B$.

then the iterative procedure of solving $0 \in Q(z, z^k) + T(z)$ *is well-defined, and converges superlinearly to some* w^*.

If the order in Condition 2 *is such that* $o(t) \leq c_0 t^\beta$, *with* $\beta = 2$, *then the convergence rate is quadratic.*

The first assumption is called *upper Lipschitz continuity* of the subproblem. The second condition quantifies the precision of the approximation of F by Q; and the third condition stipulates that the subproblems are solvable.

The three properties can be shown to hold for a class of algorithms and problems, including stabilized SQP under the assumption of the SSOSC (Assumption 5.1.3, see Fischer [27]). The discussion that follows focuses on the specific results of Fernández and Solodov [25], which are slightly stronger,

5.6.2 Convergence under the SOSC only

It can be shown that stabilized SQP satisfies the three properties required for superlinear convergence under Fischer's analysis, using only the assumption of the general second-order sufficiency condition (Assumption 5.1.1, see Fernández and Solodov [25]).

The solvability of the subproblem is demonstrated by showing that under the stabilized SOSC, it holds that the quantity

$$p^T H p + \mu \|q\|^2$$

is bounded away from zero over all (p, q) in the reduced space

$$\left\{(p, q) \mid \nabla c_i(x)^T p - \mu q_i = 0,\, i \in \mathcal{A}_+(x^*, y^*), \text{ and } \nabla c_i(x)^T p - \mu q_i \geq 0,\, i \in \mathcal{A}_0(x^*, y^*)\right\}.$$

This implies that the stabilized SQP KKT matrix

$$\begin{pmatrix} H & J^T \\ J & -\mu I \end{pmatrix},$$

is nonsingular, which is used in the existence result for the solution set.

Upper Lipschitz continuity of the subproblems can be shown using standard analysis, and superlinear local convergence follows. Furthermore, under the SSOSC, the solution of the subproblems are locally unique.

Note that, under the SOSC, although the sSQP method is well-defined and generates superlinearly convergent iterates within a certain ball around the current point, the choice of iterates is not unique. Under the SSOSC there is only one choice of iterate, and in practice, this implies that choosing the *closest* subproblem solution to the current point produces local superlinear (or quadratic) convergence, giving a precise convergence statement. However, choosing the *closest* subproblem minimizer to the current point is NP-hard.

5.6.3 Convergence to noncritical multipliers

Izmailov and Solodov [57] also analyze the properties of sSQP subproblems in more detail, analyzing the specific relationships between the notion of a *critical* multiplier, upper Lipschitz stability, the SOSC, and subproblem solvability for equality, inequality and slack-reformulated problems.

Equality constraints: Consider the following definition of a *critical multiplier.*

Definition 5.6.1. *In the case of equality constraints, a multiplier $\bar{y}$ is* critical *if there exists a nonzero $d \in \operatorname{null}(J)$ such that $H(x^*, \bar{y})d \in \operatorname{range}(J^T)^{\perp}$, and* noncritical *otherwise.*

In this case, the existence of a critical multiplier implies that the Hessian is singular in the first-order tangent feasible cone, and so the SOSC does not hold. A critical multiplier not existing in the set $\mathcal{M}_y(x^*)$ is not equivalent, but is a weaker condition that the SOSC.

It is then shown that upper Lipschitz continuity around a solution $(x^*, \bar{y})$ is equivalent to $\bar{y}$ being noncritical, and that either of these two properties imply solvability of the sSQP subproblem. These two facts imply that sSQP satisfies the assumptions of Fischer's iterative framework.

Inequality constraints: The definition of a critical multiplier may be extended to problems with inequality constraints.

Definition 5.6.2. *A multiplier is noncritical if and only if the problem*

$$\begin{array}{ll} \underset{d}{\text{minimize}} & d^T H(x^*, y^*) d \\ \text{subject to} & \nabla c_i(x^*)^T d = 0, \; i \in \mathcal{E}, \\ & \nabla c_i(x^*)^T d \geq 0, \; i \in \mathcal{A}_0, \\ & \nabla c_i(x^*)^T d = 0, \; i \in \mathcal{A}_+, \end{array}$$

has the unique solution $d = 0$.

If $H(x^*, y^*)$ is either positive-definite or negative-definite in the cone of directions d defined above, then y^* is noncritical. However, this is not a necessary condition.

Upper-Lipschitz continuity of solutions to the perturbed problem is again equivalent to the current iterate being sufficiently close to a KKT pair with a noncritical multiplier.

However, neither of these two conditions implies solvability.

When combined with the second-order *necessary* condition, upper Lipschitz continuity, or the existence of a noncritical multiplier close to the current iterate, does imply local solvability.

Slack variables: Finally, a mixed equality-inequality problem reformulated so that the inequalities are defined as equalities and bounds on slack variables is analyzed. In this case, the two properties of upper Lipschitz continuity of solutions and solvability hold if the primal-dual point is close to a solution wherein the multipliers are noncritical and the multipliers corresponding to the slacks (and so, the inequality constraints) satisfy strict complementarity.

5.6.4 Presence of critical multipliers

In general, the set of critical multipliers are of measure zero in $\mathcal{M}_y(x^*)$. However, it appears that there is evidence that Newton-like iterative methods, including stabilized SQP, have inherent tendencies of generating dual iterates that converge to critical multipliers. Izmailov and Solodov [56] point out that a relatively unlikely set of analytical conditions have to hold for the dual sequence to not converge to critical multipliers. Numerical tests in Izmailov and Solodov [55] confirm the tendency for dual iterates to converge to critical multipliers, and numerical tests in [55] and [56] confirm that such iterations do not converge superlinearly. These results show that this problem is pervasive with standard state-of-the-art optimization software, specifically SNOPT and MINOS, and occurs far less frequently, although still commonly (roughly half of the time) for sSQP.

Note that this is still consistent with the previously mentioned convergence results. Even if the set of critical multipliers is a set of measure zero, a trajectory of iterates could converge towards one normal to the set of noncritical multipliers and never enter into a domain of attraction for the noncritical multipliers.

This suggests that the research into stabilized SQP is by no means resolved, and care must be taken to prevent convergence to critical multipliers if the superlinear convergence results are to be realized.

5.7 Summary and Discussion

Table 5.1 summarizes the properties of the various regularized SQP methods that have appeared in the literature. The column headed "Num" indicates if numerical results have been given. Table 5.2 provides a legend for all of the abbreviations.

In Chapter 8, we discuss a primal-dual SQP method (pdSQP) that incorporates an sSQP algorithm with a global optimization procedure. It is shown that pdSQP drives iterates towards a KKT point and reproduces the fast local convergence properties of stabilized SQP for degenerate problems.

The second derivative SQP method SQP2d, discussed in Chapter 7 is an algorithm that is more similar to conventional SQP methods. It implements some of the ideas presented in Wright's method of stabilized working sets. The method also has fast good local convergence properties.

Table 5.1: Summary of the properties of various regularized SQP methods

Reference	Algorithm	Assumptions	Stable?	Unique?	Order r	Num
Fischer [26]	Two QPs	WCC, SOSC, MFCQ, WCR	Yes	Yes	Quad	No
Wright [86]	sSQP ineq	MFCQ, RSOSC, SC	No	Yes	Quad	No
Hager [51]	sSQP ineq	RSSOSC	Yes	Yes	Quad	No
Wright [87]	iSQP & sSQP	MFCQ and SSOSC	Yes	No	Super	No
Wright [87]	SQPsws	MFCQ, SSOSC, CRCQ	No	Yes	Super	No
Wright [88]	sSQP with ID0	MFCQ and SOSC	No	No	Super	Yes
Izmailov & Solodov [54]	AS and proj	RSOSC	Yes	Yes	Quad	No
Wright [89]	AS and eq sSQP	RSOSC on compact $\mathcal{S} \subset \mathcal{M}_y$	No	Yes	Quad	Yes
Fernández & Solodov [25]	sSQP ineq	RSOSC	Yes	SSOSC	Quad	Yes*
Izmailov & Solodov [57]	sSQP eq	Noncritical $\bar{y}$	Yes	No	Super	No
Izmailov & Solodov [57]	sSQP ineq	Noncritical y and SONC	Yes	No	Super	No
Izmailov & Solodov [57]	sSQP slacks	Noncritical y and SC	Yes	No	Super	No

* for one problem.

Table 5.2: Legend

iSQP	Inexact QP (page 55)	AS	Active-set estimation (page 65)
WCC	range($J_{\mathcal{A}_+}^T$) is independent of y	WCR	Weak Constant Rank condition (page 29)
ineq	inequality-constrained	eq	equality-constrained
proj	Projection-based algorithm (page 65)	Noncritical	Noncritical multipliers (page 70)

Chapter 6

Convexification

6.1 Introduction

This chapters considers the "convexification" of various general quadratic programs (QPs) arising in SQP methods for nonlinear optimization. It considers both the all-inequality and standard form QP subproblem.

As noted in earlier chapters, the Lagrangian Hessian is not guaranteed to be positive definite, and so a conventional SQP method could require the solution of an indefinite QP. Such problems are NP-hard, and may be directions of ascent for a merit function. As a result, many SQP methods use positive-definite approximations to the Hessian. However, by avoiding second derivatives, they do not have the potential Newton local convergence rate.

The proposed convexifications modify the Hessian matrix in the objective function $g^T x + \frac{1}{2}x^T H x$ so that it is positive definite. This procedure makes extensive use of *Debreu's Lemma*, first defined in Chapter 1, which states that if H is positive definite on the null-space of J, then there exists a finite ρ_c for which $H + \rho J^T J$ is positive definite for $\rho > \rho_c$.

6.2 Convexification in Standard Form

For a standard form quadratic program, a second-order consistent working set is one at which the basic components of $Z^T H Z$ is positive-definite, where Z is the null-space basis matrix of the equality constraint matrix J. This implies that by Debreu's Lemma, there is a ρ for which $\widetilde{H} = H + \rho J_W^T J_W = H + \rho J^T J + \rho P_N P_N^T$ is positive definite, and the quadratic program with H replaced by $\widetilde{H}$ is convex. Consider the QP,

$$\begin{array}{ll} \underset{p\in\mathbb{R}^n}{\text{minimize}} & g^Tp + \tfrac{1}{2}p^T(H + \rho J_W^T J_W)p \\ \text{subject to} & Jp = -c, \quad x_0 + p \geq 0. \end{array} \tag{6.1}$$

Consider obtaining the unique solution of (6.1) using an active-set method. Denote this solution by $\widehat{p}$. Let $\bar{\mathcal{N}}$ and $\bar{\mathcal{B}}$ denote the set of the nonbasic and basic indices at $\widehat{p}$. The stationarity condition for $\widehat{p}$ implies:

$$\begin{aligned} z &= g + (H + \rho J^TJ + \rho P_N P_N^T)\widehat{p} - J^T\bar{y} \\ &= g + H\widehat{p} - J^T(\bar{y} + \rho c) + P_N\widehat{p}_N. \end{aligned}$$

Writing this equation in terms of the basic components gives

$$z_{\bar{B}} = [g + H\widehat{p}]_{\bar{B}} - J_{\bar{B}}^T(\bar{y} + \rho c) + \rho\,[P_N\widehat{p}_N]_{\bar{B}} = 0. \tag{6.2}$$

Similarly, the nonbasic components are

$$z_{\bar{N}} = [g + H\widehat{p}]_{\bar{N}} - J_{\bar{N}}^T(\bar{y} + \rho c) - \rho\,[P_N\widehat{x}_0]_{\bar{N}} \geq 0. \tag{6.3}$$

6.2.1 Relationship between solutions

Note that, from (6.2), if $\bar{\mathcal{B}} \cap \mathcal{N}$ is empty, then this implies $\widehat{p}$ is also a stationary point of (3.16) with $y = \bar{y} + \rho c$.

Let $\widetilde{Z}$ denote a basis for the null-space for the final working set $J_{\bar{w}} = \begin{pmatrix} J \\ E_{\bar{N}}^T \end{pmatrix}$. Since $\widehat{p}$ is a subspace minimizer, $\widetilde{Z}^T H\widetilde{Z} + \rho\widetilde{Z}^T J^T J\widetilde{Z} + \rho\widetilde{Z}^T E_N E_N^T\widetilde{Z} = \widetilde{Z}^T H\widetilde{Z} + \rho\widetilde{Z}^T E_N E_N^T\widetilde{Z}$ is positive definite. This implies that for every $\widetilde{z} \in \text{null}(A_{\bar{w}})$ either $\widetilde{z}^T H\widetilde{z} > 0$ or $\widetilde{z}_N \neq 0$. The latter condition implies $\mathcal{N} \neq \bar{\mathcal{N}}$.

Finally, from (6.3), it holds that

$$[g + H\widehat{p}]_{\bar{N}} - J_{\bar{N}}^T y \geq \rho\,[P_N\widehat{x}]_{\bar{N}} \geq 0,$$

so if $\widehat{p}$ is a subspace minimizer of 3.16, the reduced costs are nonnegative and $\widehat{p}$ is also a local solution.

Convexification at a QP local minimizer

Assume the convexification is formed at a local minimizer p^* of (3.16). Let $\bar{y} = y - \rho c$. Then

$$\begin{aligned} 0 &= [g + Hp^*]_B - J_B^T y \\ &= [g + Hp^*]_B - J_B^T(\bar{y} + \rho c) + \rho(E_N p_N^*)_B, \end{aligned}$$

so p^* is also a stationary point of (6.1). Furthermore, p^* is a subspace minimizer of (6.1) since the problem is convex. Finally,

$$\begin{aligned} 0 &\le [g + Hp^*]_N - J_N^T y \\ &= [g + Hp^*]_N - J_N^T(\bar{y} + \rho c). \end{aligned}$$

Comparing this to (6.3), p^* may or may not be a local minimizer of (6.1) depending on the size of $(x_0)_N$.

6.2.2 Effect on slack-variable merit function

The directional derivative of the slack-variable augmented Lagrangian merit function can be expressed as:

$$\phi'(\alpha, \widetilde{\rho})\big|_{\alpha=0} = p^T g - p^T J^T y_0 + \widetilde{\rho} p^T J^T (c - s_0) - (c - s_0)^T q + y_0^T r - \widetilde{\rho} r^T (c - s_0), \qquad (6.4)$$

where $\widetilde{\rho}$ to denotes that a different parameter is used for the subproblem and the merit function.

Take the product p^T with the stationarity condition for the convexified subproblem:

$$p^T g + p^T H p + \rho p^T J^T J p + \rho p^T P_N P_N^T p = p^T J^T y + p^T z.$$

Input this into (6.4) to get

$$\begin{aligned} \phi'(\alpha, \widetilde{\rho})\big|_{\alpha=0} = &-(p^T H p + \rho p^T J^T J p + \rho p^T P_N P_N^T p) + p^T z \\ &+ \widetilde{\rho} p^T J^T (c - s) - (c - s)^T q + y^T r - \widetilde{\rho} r^T (c - s). \end{aligned}$$

Let $\theta(p)$ be defined such that

$$\theta(p) = p^T H p + \rho p^T J^T J p + \rho p^T P_N P_N^T p.$$

The properties of q, r and $c - s$ imply that

$$\phi'(\alpha, \widetilde{\rho})\big|_{\alpha=0} = -\theta(p) + p^T z + r^T y - 2(c - s)^T q - \widetilde{\rho}(c - s)^T (c - s).$$

Since $r^T y \le 0$, $\theta(p) > 0$ and $(c - s)^T (c - s) > 0$, for large enough $\widetilde{\rho}$ and ρ, we can make this expression negative or even bounded away from 0, below 0, regardless of the sign of $p^T z + (c - s)^T q$.

6.3 Convexification of all-inequality QPs

At a second-order consistent point, construct the convex subproblem

$$\begin{aligned} \underset{p\in\mathbb{R}^n}{\text{minimize}} \quad & g^Tp + \tfrac{1}{2}p^T(H+\rho J_W^TJ_W)x \\ \text{subject to} \quad & Jp \geq -c. \end{aligned} \tag{6.5}$$

By Debreu's Lemma, there is a ρ such that $H + \rho J_W^T J_W$ is positive definite and (6.5) is convex.

Let $\widehat{p}$ be a local minimizer of (6.5). Stationarity implies:

$$g + (H + \rho J_W^T J_W)\widehat{p} = J_A^T \bar{y}_A. \tag{6.6}$$

Optimality implies $\bar{y}_A \geq 0$.

6.3.1 Relationship between solutions

Note that from (6.6), if $J_W = J_A$ then

$$\begin{aligned} g + H\widehat{p} &= J_A^T\bar{y}_A - \rho J_A^T J_A \widehat{p} \\ &= J_W^T(\bar{y}_A + \rho c_A). \end{aligned}$$

This implies that $\widehat{p}$ is a stationary point for the original QP with $y_A = \bar{y}_A + \rho c_A$.

Since $\widehat{p}$ is a subspace minimizer, letting $\widetilde{Z} = \text{null}(J_A)$, $\widetilde{Z}^T H\widetilde{Z} + \rho\widetilde{Z}^T J_W^T J_W \widetilde{Z}$ is positive definite. This implies that either $\widetilde{Z}^T H \widetilde{Z}$ is positive definite or $J_A \neq J_W$.

Finally, the sign of y, and hence the optimality of the original QP for $\widehat{p}$ depends on the sign and magnitude of c_A.

Convexification at a QP local minimizer

Let p^* be a local minimizer for (6.5), and form the convexification at p^*. Consider the stationarity condition,

$$g + (H + \rho J_W^T J_W)p^* = J_W^T y + \rho J_W^T J_W p^* = J_W^T(y_W - \rho c_W),$$

which implies that p^* is also a stationary point for (6.5) with multiplier $\bar{y}_W = y_W - \rho c_W$.

By construction, p^* is a subspace minimizer of (6.5). However, the sign of $\bar{y}_W$ depends on ρc_W. This implies that a local minimum of the indefinite QP may not be a local minimum of (6.5).

6.4 Implementation

The implementation of convexification involves two primary considerations: 1) choosing ρ and 2) the stage in the algorithm at which to apply the change to the Hessian matrix. The challenges revolve around attempting to achieve the desired properties of a convex QP while minimizing computational cost and the perturbation to the original indefinite problem. Three strategies will be discussed: 1) perfect convexification, 2) convexification using Gershgorin circles, 3) streamlined-perfect convexification and 4) heuristic approaches.

6.4.1 Perfect convexification

Debreu's lemma states that if H is positive-definite on the null-space of J there is a $\bar{\rho}$ such that $H + \rho J^T J$ is positive-definite for $\rho > \bar{\rho}$. The original proof of Debreu's lemma [21] suggests how to obtain the exact quantity $\bar{\rho}$. From the proof, $\bar{\rho}$ can be understood as the objective value at the solution to

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \ \frac{x^T H x}{x^T J^T J x} \quad \text{subject to} \quad \|x\|_2 = 1, \quad Jx \neq 0. \tag{6.7}$$

This nonlinear program is indefinite, however, and hence potentially just as difficult to solve as the original QP.

Alternatively, note that the assumptions of the Lemma imply that under the constraint $x^T J^T J x = 0$, $x^T H x$ has a minimum reached at only $x = 0$. This implies that $\rho \geq \bar{\rho}$ is the set of Lagrange multipliers for that problem. Hence, ρ_C is the minimum ρ for which

$$\max_{\rho} \ \min_{x} \ \ x^T H x + \rho x^T J^T J x,$$

has the unique solution $x = 0$. This can be found by solving the problem

$$\max_{\rho} \ \min_{x} \ \ x^T H x + \rho x^T J^T J x - \mu \rho^2,$$

for an increasing sequence of μ-values until $x = 0$ at the solution. At the solution for ρ to the above problem,

$$\bar{\rho} = \frac{x^T J^T J x}{2\mu}$$

and so this problem becomes

$$\min_{x} x^T H x + \frac{(x^T J^T J x)^2}{2\mu}.$$

6.4.2 Standard form convexification with Gershgorin circles

Consider the Hessian decomposed as:

$$\begin{pmatrix} H_B & H_{BN}^T \\ H_{NB} & H_N \end{pmatrix}$$

where H_B is positive definite and consists of the free indices. Consider adding a diagonal,

$$\begin{pmatrix} 0 & 0 \\ 0 & D_A \end{pmatrix},$$

to H to make the entire Hessian matrix positive-definite, where D_A is diagonal, but unlikely previously, could have different values along the diagonal.

It holds that the entire matrix is positive definite if the *Schur complement* $S = H_N - H_{BN}^T H_B^{-1} H_{BN}$ is positive definite. The Gershgorin circle theorem will be used for the following analysis.

Theorem 6.4.1. Gerhgorin circle theorem*: All eigenvalues of a matrix A lie in at least one circle D_i, where D_i is defined to have center a_{ii} and radius $R_i = \sum_{j\neq i} |a_{ij}|$.*

This implies that if the entries of the Schur complement are such that $s_{ii} - \sum_{j\neq i} |s_{ij}| > 0$ for all i, then H is positive definite. This motivates a bound that would imply that all of the eigenvalues of S are positive.

Write

$$s_{ii} = [H_N]_{ii} - \sum_j [H_{BN}]_{ij} [H_F^{-1} H_{BN}]_{ji} \geq [H_N]_{ii} - \frac{\|H_{BN}\|}{\|H_B\|} \sum_j [H_{BN}]_{ij}.$$

and

$$-|s_{ij}| = -|[H_{BN}]_{ij} + \sum_k [H_{BN}]_{ik} [H_F^{-1} H_{BN}]_{kj}| \geq -|H_{BN}|_{ij} - \frac{\|H_{BN}\|}{\|H_B\|} \sum_k |[H_{BN}]_{ik}|.$$

This implies that

$$\begin{aligned} G_{\text{low}} &= s_{ii} - \textstyle\sum_{j\neq i} |s_{ij}| \\ &\geq [H_N]_{ii} - \tfrac{\|H_{BN}\|}{\|H_B\|} \textstyle\sum_j [H_{BN}]_{ij} - \sum_{j\neq i} \left(|H_{BN}|_{ij} + \tfrac{\|H_{BN}\|}{\|H_B\|} \sum_k |[H_{BN}]_{ik}| \right) \\ &\geq [H_N]_{ii} - \|H_{BN}\| (1 + (m_A + 1) \tfrac{\|H_{BN}\|}{\|H_B\|}), \end{aligned}$$

where the norm can be taken to be the one- or $\infty-$norm (which are the same since H is symmetric).

This implies that if

$$[H_N]_{ii} + D_i > ||H_{BN}||(1 + (m_A + 1)\frac{||H_{BN}||}{||H_B||})$$

then the matrix has been sufficiently convexified.

6.4.3 Streamlined-perfect convexification

In general, it is preferable to limit the degree to which the problem is altered in order for the solution of the convex QP to be as nearest as possible to a solution to the original indefinite QP. Consider a problem with two constraints $\begin{pmatrix} a_1^T \\ a_2^T \end{pmatrix} p \geq -c$. If, at a point $\bar{p}$, both constraints are active and the working set is second-order consistent, a convexification would add $\rho \begin{pmatrix} a_1^T \\ a_2^T \end{pmatrix}^T \begin{pmatrix} a_1^T \\ a_2^T \end{pmatrix}$ to H in the objective. However, an active-set method moves off one constraint at a time. This implies that if a_1 is held active, i.e., $p \in N(a_1^T)$, then the step would be identical as if the convexification only altered H by the term $\rho \left(a_2^T\right)^T \left(a_2^T\right)$. This suggests that it is possible to add one constraint at a time to H, i.e., $\widetilde{H} = \widetilde{H} + \rho_s a_s a_s^T$, s denoting the constraint that the algorithm moves off of and ρ_s chosen appropriately. Additionally, ρ_s can be chosen as in the perfect convexification for the entire J or a series of ρ_s can be chosen for each s. In the latter case, ρ_s can vary for each s, allowing a smaller overall change in the Hessian matrix of the subproblem.

6.4.4 Heuristic convexification

If every step p_k that the QP algorithm computes is a direction of positive curvature for H, then the QP algorithm executes as if the problem was convex. One approach this suggests is to convexify only if the calculated step is a direction of nonpositive curvature for H, calculating ρ to be just large enough to change the curvature along the step. Such an algorithm is summarized in Algorithm 6.4.1

One immediately noticeable drawback to this specific approach is that a potentially excessive number of steps may have to be calculated. It'd be desirable to ascertain *a priori* whether a direction moving off a constraint will be a direction of nonpositive curvature and the magnitude of the negative curvature (and hence the necessary magnitude of ρ to correct for it).

Consider a step p moving off of constraint s at a second-order consistent point. The step p may be decomposed into $p = \alpha a_s + Z p_z$, where Z is the null-space of the current working set. Consider the curvature along p,

$$p^T H p = \alpha^2 a_s^T H a_s + 2\alpha a_s^T H Z p_z + p_z^T Z^T H Z p_z.$$

$p_z^T Z^T H Z p_z > 0$ by assumption. If $a_s^T H Z \geq 0$ and $a_s^T H a_s \geq 0$ then for any value of α and p_z, the step p is a direction of positive curvature, and so a convexification is unnecessary.

Algorithm 6.4.1. [Heuristic Convexification]

Let $\widetilde{H}_0 = H$, $\rho_0 = \rho_{min}$, $\nu > 0$, $\eta > 0$, $\widetilde{J}$ as empty.

while Not converged **do**

Let $j = \text{argmin}_k\, y_k$;

if $y_j \geq 0$ **then stop**;

Calculate step p_{k+1};

Let $\rho_{k+1} = \rho_k$, $\widetilde{H}_{k+1} = \widetilde{H}_k$;

if $p_{k+1}^T \widetilde{H}_{k+1} p_{k+1} < \nu$, **then**

Let $\widetilde{J}_{k+1} = \begin{pmatrix} \widetilde{J}_k \\ a_j^T \end{pmatrix}$;

Let $\widetilde{H}_{k+1} = H + \rho_k \widetilde{J}_{k+1}^T \widetilde{J}_{k+1}$;

Recalculate step p_{k+1} using the new $\widetilde{H}_{k+1}$;

if $(p_{k+1}^T \widetilde{H}_{k+1} p_{k+1}) < \nu$ **then**

Let $\rho_{k+1} = -\frac{p_{k+1}^T H p_{k+1}}{\rho_{k+1} p_{k+1}^T J_{k+1}^T J_{k+1} p_{k+1}} + \eta$;

else

break

end

end

Remove j from the working set.

Calculate α, take the step, and add any blocking constraints;

end do

Otherwise, ρ may be calculated as some combination of $-a_s^T HZ$ and $-a_s^T Ha_s$. Since α and p_z is unknown a priori, an algorithm could initialize ρ by setting $\rho_0 = -\frac{a_s^T HZe}{||a_s||_1} - \frac{a_s^T Ha_s}{||a_s||^2}$, and then scale $\rho_{k+1} = \gamma\rho_k$, $\gamma > 1$ as in Algorithm 6.4.1 if the direction turns out to still have nonpositive curvature along $\widetilde{H}$. Additionally, an initial scaling can be introduced based on the size of $\|g\|$, an estimate on the size of the feasible region or simply the size of the current step p_k.

Chapter 7

SQP with a Smooth Augmented Lagrangian Merit Function

7.1 Introduction

This chapter discusses the convergence properties of a sequential quadratic programming algorithm in which the QP objective is defined using the exact Hessian of the Lagrangian. The algorithm uses the augmented Lagrangian merit function of Gill et al. [36], and follows Murray and Prieto [69] in decomposing the QP subproblem into several simpler parts where one requires only an *approximate* solution of the QP.

In addition, rather than assuming the strong and commonly violated assumption of the linear independence constraint qualification (LICQ) holding at the solution, as is standard in the literature, the weaker Mangasarian-Fromovitz constraint qualification (MFCQ) and weak constant rank (WCR) conditions are assumed (see Chapter 2, page 12). The global convergence theory utilizes results of Qi and Wei [79]. The local convergence theory involves the properties of stabilized working sets proposed by Wright [87, Section 8] (see Chapter 5). The local convergence results in Wright [87] require the satisfaction of the second-order sufficiency conditions at a solution. However, as our analysis implies global convergence to a point satisfying the second-order necessary conditions, it is necessary only that the reduced Hessian be nonsingular at the solution.

7.2 Description of Outer Algorithm

The following problem is to be solved:

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \;\; f(x)$$
$$\text{subject to} \;\; c(x) \geq 0.$$

The following smooth merit function is used:

$$L_A(x, y, s, \rho) = f(x) - y^T(c(x) - s) + \tfrac{1}{2}\rho(c(x) - s)^T(c(x) - s), \tag{7.1}$$

where $s \geq 0$ are slack variables, and the scalar ρ is known as the penalty parameter. In other SQP methods the choice of merit functions is largely dictated by considerations of efficiency. Here, however, it is essential that the merit function have continuous second derivatives. The use of slack variables in the merit function is the feature that makes the merit function smooth. Without slack variables it is necessary to define a merit function in terms of only the "active" constraints. Such a merit function does not have continuous second derivatives. Another virtue of this merit function is that it does not suffer from the Maratos effect. A comparative discussion of this and other merit functions may be found in Gill and Wong [41].

The choice of merit function requires a search direction defined in the x, y and s variables and a search is performed on this expanded space. The approach adopted for the slack variables is to set them at their optimal value (this is trivial since the merit function is a quadratic function of the slack variables). Given a search direction in the x and y variables, the search direction in the slack variables is then chosen to ensure that the slack variables remain at their optimal values when the constraints are linear.

An important result of this chapter is to show how a direction of negative curvature for the Hessian of the Lagrangian is transformed into a direction of negative curvature for the Hessian of the merit function in the triple space of x, y and s.

For the curvilinear search used in the algorithm, the value of the merit function as a function of the steplength α will be denoted by

$$\phi(\alpha; x, y, s, p, u, q, r, \widehat{r}, \rho) \equiv L_A(x + \alpha^2 p + \alpha u, y + \alpha^2 q, s + \alpha^2 r + \alpha \widehat{r}, \rho). \tag{7.2}$$

In what follows, only those arguments relevant to the current discussion will be included in the notation for ϕ. At the kth iteration, the value of ϕ will be denoted by ϕ_k. The first, second and third derivatives of ϕ with respect to α will be denoted by ϕ', ϕ'' and ϕ''', respectively.

The following notation will be used when appropriate:

$$g_k \equiv \nabla f(x_k), \qquad J_k \equiv c'(x_k), \qquad c_k \equiv c(x_k), \qquad H_k \equiv \nabla^2_{xx} L(x_k, y_k),$$

where L denotes the Lagrangian function.

The norm of the Lagrangian gradient is used to limit the size of some values in the algorithm. Let $\delta_k^L = \delta^L(x_k, y_k, s_k; \rho_k)$, where

$$\begin{aligned}\delta^L(x, y, s; \rho) = \|g(x) - J(x)^T(y - \rho(c(x) - s)\| + \|c(x) - s\| \\ + \sum_{j:s_j>0} |y_j - \rho(c_j(x) - s_j)|.\end{aligned} \tag{7.3}$$

The value of $\delta^L(x, y, s; \rho)$ provides a measure of the accuracy of $(x, y, s; \rho)$ as an estimate of a first-order KKT point for NP (assuming s is defined as in (7.4) below). In particular, (x_k, y_k, s_k) is a first-order KKT point at x_k if and only if $\delta_k^L = 0$.

Finally, symbols of the form β_{ab} indicate constants related to properties of the problem, or the implementation of the algorithm, where "ab" identifies the specific scalar represented.

The computation of the search direction p_k will be described in the next section. Once p_k has been computed, the slack variables are adjusted to minimize the merit function for a fixed value of x_k. This gives the optimal value s_k such that

$$s_k = \begin{cases} \max\{0, c_k\} & \text{if } \rho_{k-1} = 0, \\ \max\{0, c_k - y_k/\rho_{k-1}\} & \text{otherwise.} \end{cases} \tag{7.4}$$

The following inequality will prove useful (see Murray and Prieto [70, Equation (2.9)].

$$\|c(x)^-\| \le \|c(x) - s\|. \tag{7.5}$$

Let (p_k, π_k) be the primal-dual solution of the quadratic subproblem. If $p_k = 0$, then x_k is a second-order KKT point for problem NP and the algorithm is terminated with $y_k = \pi_k$. Otherwise, the dual search direction is given by

$$q_k = \pi_k - y_k. \tag{7.6}$$

Let $\widehat{u}_k$ be a direction of negative curvature for H_k. The definition of $\widehat{u}_k$ is discussed in a subsequent section.

The algorithm adjusts the penalty parameter to ensure that a sufficient reduction in the merit function is possible. Let

$$\omega_k \equiv \tfrac{1}{2}(\|c_k - s_k\|^2 + p_k^T \bar{B}_k p_k), \tag{7.7}$$

where $\bar{B}_k$ is a positive-definite matrix described in the next section. If

$$\phi_k''(0; \rho_{k-1}, \widehat{u}_k) \le -\omega_k + \tfrac{1}{2}\widehat{u}_k^T H_k \widehat{u}_k, \tag{7.8}$$

then $\rho_k = \rho_{k-1}$. Otherwise, ρ_k is chosen to satisfy $\rho_k \ge \widehat{\rho}_k$, where

$$\widehat{\rho}_k = \frac{\omega_k + g_k^T p_k + (2y_k - \pi_k)^T (c_k - s_k)}{\|c_k - s_k\|^2}. \tag{7.9}$$

When $\rho_k = \rho_{k-1}$ the algorithm sets $u_k = \widehat{u}_k$, otherwise $u_k = 0$ since the direction of negative curvature $\widehat{u}_k$ may no longer be valid. It will be shown in Section 7.4.2 that when $\rho_k \ne \rho_{k-1}$, the choice of ρ_k ensures that $\phi_k''(0; \rho_k, u_k) \le -\omega_k$, with $u_k = 0$. If the condition (7.8) does not hold and ρ_k is updated, then $\phi_k''(0; \rho_{k-1}, u_k) > -\omega_k$ with $u_k = 0$.

The search directions for the slack variables are

$$r_k = J_k p_k + c_k - s_k, \quad \text{and} \quad \widehat{r}_k = J_k u_k. \tag{7.10}$$

If a constraint is linear, these directions maintain the corresponding slack variable at its optimal value.

When $u_k \ne 0$ a curvilinear search is performed to obtain a step length α_k such that $x_{k+1} - x_k$ becomes parallel to a direction of negative curvature as $\alpha_k \to 0$. This type of search gives a method that makes fewer adjustments to ρ. It is possible that no simple *linear* combination of p_k and u_k gives a direction of descent because p_k is not a direction of descent for ϕ_k defined with the current value of ρ. In contrast, if u is a direction of negative curvature, the merit function can be reduced by using a curvilinear search with the current value of ρ.

The curvilinear search computes a steplength $\alpha_k > 0$ such that the new iterate

$$\begin{pmatrix} x_k \\ y_k \\ s_k \end{pmatrix} + \alpha_k \begin{pmatrix} u_k \\ 0 \\ \widehat{r}_k \end{pmatrix} + \alpha_k^2 \begin{pmatrix} p_k \\ q_k \\ r_k \end{pmatrix}$$

gives a sufficient reduction of the merit function L_A while keeping the constraint violation bounded.

The following termination criteria are used. If

$$\phi_k(1) - \phi_k(0) \le \sigma \left(\phi_k'(0) + \tfrac{1}{2}\phi_k''(0) \right), \tag{7.11}$$

set $\widehat{\alpha} = 1$. Otherwise, $\widehat{\alpha} \in (0,1)$ is determined such that

$$\phi_k(\widehat{\alpha}) - \phi_k(0) \le \sigma \left(\widehat{\alpha}\phi_k'(0) + \tfrac{1}{2}\widehat{\alpha}^2 \phi_k''(0) \right) \tag{7.12a}$$

$$\phi_k'(\widehat{\alpha}) \ge \eta_W \left(\phi_k'(0) + \widehat{\alpha}\phi_k''(0) \right), \tag{7.12b}$$

where $0 < \sigma < \frac{1}{2}$ and $\frac{1}{2} < \eta_W < 1$. Let μ_k $(0 < \mu_k < 1)$ denote a stabilization parameter defined in the next section. If the condition

$$c(x_k + \widehat{\alpha} u_k + \widehat{\alpha}^2 p_k) \geq -\max(\beta_c, \mu_k)e \tag{7.13}$$

holds, let $\alpha_k = \widehat{\alpha}$; otherwise α_k is computed by performing a backtracking linesearch from $\widehat{\alpha}$ until (7.13) and (7.12a) are simultaneously satisfied. These conditions have been shown by Olivares et al. [72] to be appropriate for combining descent and negative curvature directions when solving unconstrained problems. It will be shown in Section 7.4.2 that such step lengths exist, and that Algorithm **SQP2D** is well defined.

Given a feasible point x for linear constraints $Ax \geq b$, the maximum feasible step γ such that $A(x + \gamma p) \geq b$ is given by

$$\gamma = \min \gamma_i, \quad \text{with} \quad \gamma_i = \begin{cases} \dfrac{a_i^T x - b_i}{-a_i^T p} & \text{if } a_i^T p < 0; \\ +\infty & \text{otherwise.} \end{cases}$$

For brevity, the calculation of γ for the constraints $Ax \geq b$ will be summarized as $\gamma = \texttt{maxStep}(A, b, x, p)$.

The slacks s_{k+1} for the next iteration are recomputed from (7.4) based on the values of x_{k+1} and y_{k+1}.

Algorithm 7.2.1. [Algorithm SQP2D]
Given x_0 and y_0, choose $\rho_{-1} \geq 0$
repeat
  Compute g_k, J_k and c_k;
  if $\rho_{k-1} = 0$ **then** [Optimize the slack variables]
    Compute s_k from $s_k = \max\{0, c_k\}$;
  else
    Compute s_k as $s_k = \max\{0, c_k - y_k/\rho_{k-1}\}$;
  end if
  Compute a feasible step p_F; Compute H_k; p_S; $\widehat{p}$; $\widehat{u}$;
  $p_k = p_S + \widehat{p}$;
  $q = \pi - y_k$;
  $\omega = \frac{1}{2}(\|c_k - s_k\|^2 + p_k^T \bar{B}_k p_k$;
  if $\phi_k''(0; \rho_{k-1}, u) \leq -\omega + \frac{1}{2}\widehat{u}^T H_k \widehat{u}$ **then**

$u = \widehat{u}$; $\rho_k = \rho_{k-1}$;

else

$$u = 0; \quad \rho_k = \max\left(2\rho_{k-1}, \frac{\frac{1}{2}\omega + g_k^T p_k + (2y_k - \pi)^T(c_k - s_k)}{\|c_k - s_k\|^2}, \beta_p\right);$$

end if

$r = J_k p_k + c_k - s_k$; $\quad \widehat{r} = J_k u$;

Perform curvilinear search to compute α using the merit function;

$x_{k+1} = x_k + \alpha u + \alpha^2 p_k$; $\quad y_{k+1} = y_k + \alpha^2 q$; $\quad k \leftarrow k + 1$;

until converged;

7.3 Solving the Subproblem

7.3.1 Calculating a feasible step, warm starts and stable active sets

A rank-enforcing active-set QP method maintains a linearly independent estimate of the active set called the working set. The stabilized working-set algorithm uses the working set from the previous subproblem as an initial estimate of the working set for the next subproblems. In addition, there is a degree of infeasibility permitted for the constraints not in the working-set. The details of such a procedure are discussed in this section. The procedure described here is similar to, but not identical to, Wright's stabilized working-set framework [87], discussed in Chapter 5.

At the initial QP feasible point, the linearized constraints corresponding to the previous working set are made active. If the final working set from the previous iteration is $\mathcal{W}_{k-1}$, then $[J]_{\mathcal{W}_{k-1}}$ is checked for linearly independence. If it is linearly independent, then the initial working set $\mathcal{W}_k$ is set to be $\mathcal{W}_{k-1}$. If it is not linearly independent, then the algorithm finds a new working set $\mathcal{W}_k$ by removing linearly dependent constraints. Subsequently, a feasible point satisfying the following conditions is found:

$$\begin{aligned} \nabla c_i^T p_F + c_i &= \quad 0,\ i \in \mathcal{W}_k, \\ \nabla c_i^T p_F + c_i &\geq -\mu_k, \text{ otherwise}, \\ \|p_F\| &\leq \beta_{pf}\|\widetilde{c}^-\|, \\ g^T p_F &\leq \beta_{pf}\|\widetilde{c}^-\|, \end{aligned} \tag{7.14}$$

for some positive constant β_{p0}, where $\widetilde{c}_j$ denotes the vector of normalized constraints, $\widetilde{c}_j = c_j/(1 + \|\nabla c_j\|)$. These conditions imply

$$\|p_F\| \leq \beta_{pf}\|c^-\|, \quad g^T p_F \leq \beta_{pf}\|c^-\|. \tag{7.15}$$

The parameter μ_k is defined as

$$\mu_k = \begin{cases} \eta(x_k, y_k)^\tau & \text{if } \eta(x_k, y_k) \le \mu_{k-1}; \\ \bar{\mu}\mu_{k-1} & \text{otherwise} \end{cases}$$

where $\bar{\mu}$ and τ are preassigned constants such that $\frac{1}{2} < \bar{\mu} < 1$ and $0 < \tau \le 1$, and $\eta(x_k, y_k)$ is an estimate of the violation of the KKT conditions, i.e.,

$$\eta(x, y) = \left\| \begin{pmatrix} g(x) - J(x)^T y \\ \min(c(x), y) \end{pmatrix} \right\|.$$

This definition implies that μ_k monotonically converges to zero. Wright [87] shows that $\eta(x_k, y_k) = \Theta(\delta(x_k, y_k))$, where $\delta(x_k, y_k)$ is the distance to the nearest KKT point (x^*, y^*). Let $\widehat{\mu}$ be a vector such that $[\widehat{\mu}]_i = 0$ for $i \in \mathcal{W}$ and $[\widehat{\mu}]_i = \mu$ otherwise.

If a point p_F satisfying the conditions of (7.14) cannot be found, define p_F to satisfy $Jp_F \ge -c$, instead of the first two conditions of (7.14). In this case, define $\widehat{\mu} = 0$.

7.3.2 Definition of the QP stationary point p_S

The direction p_S is defined as a stationary point of the indefinite QP and satisfies the conditions

$$\begin{aligned} g + Hp_S &= J^T\pi_S, \\ c + Jp_S &\ge -\widehat{\mu}, \\ \pi_S^T(c + Jp_S) &= 0. \end{aligned} \tag{7.16}$$

The stationary point p_S is found using the exact Hessian. The QP iterations are initialized at p_F. It will be shown later that as $k \to \infty$, the working set corresponding to p_F and p_S are the same and the stationary point is also a local minimizer for the QP defined by the conditions of (7.16), i.e., $\pi_S \ge 0$ and the reduced Hessian is positive definite. When this is the case and p_S is also a local minimizer for the the convex QP defined in (7.17), the subproblem returns $(p, \pi) = (p_S, \pi_S)$. This ensures the fast Newton convergence rate associated with using exact second-derivatives.

7.3.3 Convex QP

Convexification

A convex QP is defined with a Hessian $\bar{B}$ that must be sufficiently positive-definite, i.e., $\bar{B}$ must satisfy the condition $p^T\bar{B}p \ge \beta_{Bp}\|p\|^2$ for all p. Debreu's lemma [21] states that if H is positive-definite on the null-space of J there is a $\bar{\rho}$ such that $H + \rho J^T J$ is positive-definite for $\rho > \bar{\rho}$. This means that a positive-definite $\bar{B}$ can be formed from H using the active constraints in the working set, with additional artificial constraints as necessary. See Chapter 6 for details.

Solution to the convex QP

The convexified subproblem is given by:

$$\begin{aligned} \underset{\widehat{p}}{\text{minimize}} \quad & g^T(p_S + \widehat{p}) + (p_S + \widehat{p})^T \bar{B}(p_S + \widehat{p}) \\ \text{subject to} \quad & J(p_S + \widehat{p}) \geq -c - \widehat{\mu}. \end{aligned} \tag{7.17}$$

with solution $(\widehat{p}, \pi)$.

Let the subscript w correspond to the entries or rows corresponding to the final working set at p_S. At the initial point for the convex QP, if there are no artificial constraints, then

$$g + \bar{B}p_S = g + Hp_S + \rho J_W^T J_W p_S = J_W^T \pi_S + \rho J_W^T J_W p_S = J_W^T(\pi_S - c_W - \widehat{\mu}).$$

This implies that p_S is also a stationary point for (7.17).

If $\pi_S - \widehat{c}_W - \widehat{\mu} \geq 0$, where $\widehat{c}_W$ is such that $[\widehat{c}_W]_i = [c_W]_i$ for $i \in \mathcal{W}$ and zero otherwise, and there are no artificial constraints, then the convex QP is not solved. In this case, if $\pi_S \geq 0$, use $(p, \pi) = (p_S, \pi_S)$ as the estimates for the line search, otherwise use $(p, \pi) = (p_S, \pi_S - c_W - \widehat{\mu})$.

Let $\mathcal{A}_k$ be the final active set and $\mathcal{W}_k$ the final working set from the subproblem.

7.3.4 Computation of the direction of negative curvature

Let Z_0 denote a basis for the null space of the active QP constraint matrix at p_S, with bounded condition number. If $Z_0^T H Z_0$ is sufficiently indefinite, in the sense that

$$\lambda_{\min}(Z_0^T H Z_0) \leq -\beta_\delta \delta^L, \tag{7.18}$$

holds for δ^L defined in (7.3), x is in a region that satisfies

$$\delta^L \|x\| \leq \beta_h, \tag{7.19}$$

and the violation of the constraints is not too large, that is, if

$$c(x) \geq -\tfrac{1}{2}\max(\beta_c, 2\mu_k)e, \tag{7.20}$$

compute $\bar{u} = Z_0\widetilde{u}$, a feasible direction of sufficient negative curvature; otherwise, let $\bar{u} = 0$. More precisely, $\widetilde{u}$ is required to satisfy

$$\begin{gathered} g^T Z_0 \widetilde{u} \le 0, \qquad \widetilde{u}^T Z_0^T H Z_0 \widetilde{u} \le \beta_{u1} \lambda_{\min}(Z_0^T H Z_0) \|\widetilde{u}\|^2, \\ \beta_{u3} |\lambda_{\min}(Z_0^T H Z_0)| \le \|\widetilde{u}\| \le \beta_{u2} |\lambda_{\min}(Z_0^T H Z_0)|, \end{gathered} \tag{7.21}$$

for positive constants β_{u1}, β_{u2} and β_{u3}, where $\lambda_{\min}(H)$ denotes the smallest eigenvalue of H. Note that the first condition is trivial to satisfy by adjusting the sign of any suitable direction of negative curvature.

The direction of negative curvature $\bar{u}$ is scaled so that the scaled value $\widehat{u}$ is a feasible step from p_S, and also satisfies $[J\widehat{u} + c]_j \ge 0$ for $c_j > 0$. Let γ' be

$$\gamma' = \min \Big\{ \min_j \Big(\frac{[c + Jp_S]_j}{-[J\bar{u}]_j} \mid [J\bar{u}]_j < 0 \Big), \min_j \Big(\frac{c_j}{-[J\bar{u}]_j} \mid [J\bar{u}]_j < 0, c_j > 0 \Big) \Big\}. \tag{7.22}$$

Let $\widehat{u}$ be defined as

$$\widehat{u} = \min(1, \gamma')\bar{u}. \tag{7.23}$$

The actual direction of negative curvature used (if it is used at all) is a scaled version of $\bar{u}$. The final scaling can be determined only after the descent step has been computed.

Details of the computation of a direction of negative curvature that satisfies the requirements above are discussed in Chapter 9, which also describes another regularized SQP algorithm incorporating directions of negative curvature.

7.3.5 Statement of the QP algorithm

Algorithm 7.3.1. [Algorithm QP2D]

Compute p_F to satisfy (7.14);

Let $\widehat{J}_0$ denote the working-set matrix at p_F and $w_0 = p_F$;

repeat

Solve $\begin{pmatrix} H & \widehat{J}_i^T \\ \widehat{J}_i & 0 \end{pmatrix} \begin{pmatrix} d_i \\ -\pi_S \end{pmatrix} = -\begin{pmatrix} g + Hw_i \\ -\widehat{\mu} \end{pmatrix}$;

$\gamma = \min\big(1, \texttt{maxStep}(\, J, -c, w_i, d_i \,)\big)$; $\quad w_{i+1} = w_i + \gamma d_i$;

Set $\widehat{J}_{i+1}$ to be the working-set Jacobian at w_{i+1};

$i \leftarrow i + 1$;

until (w_i, π_S) satisfy (7.16)

$p_S \leftarrow w_i$;

Let Z_0 be a basis for the null space at J_i;

if $\lambda_{\min}(Z_0^T H Z_0) < -\beta_\delta \delta^L$ **and** $\delta^L \|x\| \le \beta_h$ **and** $c(x) \ge -\frac{1}{2}\beta_c e$ **then**

Compute $\bar{u} = Z_0 \widetilde{u}$, where $\widetilde{u}$ is a direction of negative curvature of $Z_0^T H Z_0$;

$\gamma' \leftarrow \texttt{maxStep}(\,J, -c, p_S, \bar{u}\,)$; $\quad \widehat{u} \leftarrow \min(1, \gamma')\bar{u}$;

else

$\widehat{u} \leftarrow 0$;

end if

Construct a positive-definite approximate Hessian $\bar{B}$.

if $\pi_S - \widehat{\mu} - c \ge 0$ **then**

if $\pi_S \ge 0$ **then**

$(p_k, \pi) = (p_S, \pi_S)$;

else

$(p_k, \pi) = (p_S, \pi_S - \widehat{\mu} - c)$;

end

exit;

else

Compute $(\widehat{p}, \pi)$, the primal-dual solution of

$$\begin{array}{ll} \underset{\widehat{p}}{\text{minimize}} & g^T(p_S + \widehat{p}) + (p_S + \widehat{p})^T \bar{B}(p_S + \widehat{p}) \\ \text{subject to} & J(p_S + \widehat{p}) \ge -c - \widehat{\mu}. \end{array}$$

Set $p_k = p_S + \widehat{p}$;

end if

7.4 Global Convergence

7.4.1 Assumptions and preliminaries

Assumption 7.4.1. *The functions f and c are three-times Lipschitz continuously differentiable.*

Assumption 7.4.2. *For some constant β_{pf}, a feasible point p_F exists for every QP subproblem satisfying,*

$$\|p_F\| \le \beta_{pf} \|\tilde{c}_k^-\|, \quad g_k^T p_F \le \beta_{pf} \|\tilde{c}_k^-\|.$$

Assumption 7.4.3. *All first-order KKT points satisfy the Mangasarian-Fromovitz constraint qualification (MFCQ) and the weak constant rank (WCR) conditions.*

Assumption 7.4.4. *The multipliers π and π_S are uniformly bounded for every subproblem.*

Note that this implies that the sequence of multipliers $\{y_k\}$ are uniformly bounded.

Assumption 7.4.5. *All iterates x_k belong to a bounded, convex set.*

In this case, the second-order necessary condition to be proven to hold at the solution is the weak reduced semi-definiteness property (WSRP) which holds if the Hessian $H(x^*, y^*)$ is positive definite on the subspace $\{d \mid J_{\mathcal{A}(x^*)}(x^*)d = 0\}$ (see Chapter 2 for details).

7.4.2 Existence of the iterates

In this section it is shown that each iteration of the algorithm is well defined. In particular, it is shown that the penalty parameter and the steplength are well defined at every iterate.

Lemma 7.4.1. *Under Assumptions* 7.4.1–7.4.5*, the algorithm for computing the values of the penalty parameter ρ_k and the steplength α_k are well defined.*

Proof. In order to simplify notation, the subscript k associated with the iteration number will be omitted. Consider the definition (7.9) of the penalty parameter.

The gradient of L_A, with respect to x, y and s is given by

$$\nabla L_A(x, y, s) \equiv \begin{pmatrix} g(x) - J(x)^T y + \rho J(x)^T (c(x) - s) \\ -(c(x) - s) \\ y - \rho(c(x) - s) \end{pmatrix}. \tag{7.24}$$

From (7.10) and (7.24) it follows that

$$\begin{aligned} \phi'(0) &= \begin{pmatrix} u^T & 0 & \widehat{r}^T \end{pmatrix} \nabla L_A \\ &= g^T u - y^T J u + \rho\, u^T J^T (c - s) + y^T \widehat{r} - \rho \widehat{r}^T (c - s) \\ &= g^T u, \end{aligned} \tag{7.25}$$

where g, J, and c are evaluated at x. Note that from (7.21) it holds that $\phi'(0) \le 0$.

Consider now the Hessian of L_A with respect to x, y and s,

$$\nabla^2 L_A(x, y, s) = \begin{pmatrix} \nabla^2 L(x, y) + \rho J(x)^T J(x) & -J(x)^T & -\rho J(x)^T \\ -J(x) & 0 & I \\ -\rho J(x) & I & \rho I \end{pmatrix}. \tag{7.26}$$

From (7.24) and (7.26) it follows that

$$\begin{aligned}\phi''(0) &= 2\begin{pmatrix} p^T & q^T & r^T \end{pmatrix} \nabla L_A + \begin{pmatrix} u^T & 0 & r^T \end{pmatrix} \nabla^2 L_A \begin{pmatrix} u \\ 0 \\ \widehat{r} \end{pmatrix} \\ &= 2p^T g - 2p^T J^T y + 2\rho p^T J^T (c-s) - 2q^T (c-s) + 2r^T y - 2\rho r^T (c-s) + u^T H u \\ &\qquad + \rho u^T J^T J u - \rho u^T J^T \widehat{r}.\end{aligned}$$

Using (7.6) and (7.10), this expression becomes

$$u^T H u + 2g^T p + 2(2y - \pi)^T (c-s) - 2\rho \|c-s\|^2. \tag{7.27}$$

The penalty parameter (7.9) is well defined whenever $\|c - s\| > 0$. For the case when $c - s = 0$, from (7.15) and (7.5) p_F must simply move on to the constraints of the linearly independent working set, and the definitions in algorithm **QD2D** together with the positive definiteness of $\bar{B}$ on the relevant subspaces imply that p is obtained by taking descent steps from p_F. It follows that $g^T p + \frac{1}{2} p^T \bar{B} p \le 0$, and $g^T p \le -\frac{1}{2} p^T \bar{B} p$. It follows from the definition of p that if $c - s = 0$,

$$\phi''(0) = 2g^T p + u^T H u \le -p^T \bar{B} p + u^T H u = -2\omega + u^T H u,$$

implying that if $c - s = 0$, then (7.8) is satisfied and the penalty parameter need not be modified.

If $\|c - s\| > 0$, then from (7.27) with $u = 0$ (if $u \ne 0$, condition (7.8) is satisfied by the current value of ρ) and $\rho = \widehat{\rho}$ (see (7.9)) it holds that

$$\phi''(0) = -2\omega,$$

which implies that (7.8) is satisfied for all $\rho \ge \widehat{\rho}$.

It also must be shown that α_k introduced in the algorithm is well defined. It will be shown that a steplength $\widehat{\alpha}$ that satisfies either condition (7.11) or conditions (7.12a) always exists (see, for example, Moré and Sorensen [67]).

Assume that condition (7.11) is not satisfied. Let $\bar{\phi}$ be defined as,

$$\bar{\phi}(\alpha; \sigma) \equiv \phi(\alpha) - \phi(0) - \sigma(\alpha \phi'(0) + \tfrac{1}{2}\alpha^2 \phi''(0)),$$

where $\bar{\phi}(0;\sigma) = 0$, $\bar{\phi}'(0,\sigma) = (1-\sigma)\phi'(0) \le 0$ and $\bar{\phi}''(0;\sigma) = (1-\sigma)\phi''(0) < 0$ (all derivatives are taken with respect to the first argument).

As (7.11) is not satisfied, it must hold that $\bar{\phi}(1, \sigma) > 0$, and there must exist an $\widehat{\alpha} \in (0, 1]$ for which $\bar{\phi}'(\widehat{\alpha}; \eta_W) \geq 0$, where $\frac{1}{2} < \eta_W < 1$ is introduced in (7.12b). (If no such $\widehat{\alpha}$ exists, then $\bar{\phi}'(\alpha; \eta_W) < 0$ for all $\alpha \in (0, 1]$, and from the condition that $\eta_W > \sigma$, $\phi'(0) \leq 0$ and $\phi''(0) < 0$ it must hold that $\bar{\phi}'(\alpha; \sigma) < 0$ for all $\alpha \in (0, 1]$, and so (7.11) would necessarily hold, contradicting our assumption that there is no such $\widehat{\alpha}$.)

Suppose that $\widehat{\alpha}$ is the smallest such point, which implies that $\bar{\phi}'(\alpha; \eta_S) < 0$ for all $\alpha \in (0, \widehat{\alpha})$. Integrating this inequality between 0 and $\widehat{\alpha}$ yields

$$\phi(\widehat{\alpha}) \leq \phi(0) + \eta_W(\widehat{\alpha}\phi'(0) + \tfrac{1}{2}\widehat{\alpha}^2\phi''(0)),$$

and from the conditions $\phi'(0) \leq 0$, $\phi''(0) < 0$ and $\sigma < \eta_W$, (7.12a) must be satisfied at $\widehat{\alpha}$. In addition, the inequality $\bar{\phi}'(\widehat{\alpha}; \eta_W) \geq 0$, implies that (7.12b) is satisfied at $\widehat{\alpha}$.

It remains to be shown that (7.13) can also be satisfied. The function

$$h(\alpha) \equiv c(x + \alpha u + \alpha^2 p) + \beta_c e \tag{7.28}$$

has derivatives

$$h'(0) = Ju, \qquad h_j''(0) = u^T \nabla^2 c_j u + \nabla c_j^T p.$$

If $-\frac{1}{2}\beta_c > c_j \geq -\beta_c$, then condition (7.20) is not satisfied and u must be zero; from (7.28) it holds that $h_j(0) \geq 0$, $h_j'(0) = 0$ and $h_j''(0) = \nabla c_j^T p_L \geq -c_j \geq \frac{1}{2}\beta_c > 0$. If $c_j \geq -\frac{1}{2}\beta_c$ then $h_j(0) \geq \frac{1}{2}\beta_c > 0$ and in any case there exists a value $\widetilde{\alpha} > 0$ such that $h_j(\alpha) \geq 0$ (and $c_j(x+\alpha p) \geq -\beta_c$) for all j and all $\alpha \in [0, \widetilde{\alpha}]$, implying that for $\alpha \in [0, \min\{\widehat{\alpha}, \widetilde{\alpha}\}]$ conditions (7.13) and (7.12a) hold simultaneously. □

The strategy for the selection of the penalty parameter ρ_k is to define its value to be large enough to satisfy (7.8), while remaining small enough to be bounded by a multiple of $\widehat{\rho}_k$. The selection rule is as follows: Let

$$\rho_k = \begin{cases} \rho_{k-1} & \text{if } \phi_k''(0) \leq -\omega + \frac{1}{2}u_k^T H_k u_k, \\ \max(\widehat{\rho}_k, 2\rho_{k-1}, \beta_\rho) & \text{otherwise,} \end{cases} \tag{7.29}$$

where $\widehat{\rho}_k$ is defined as in (7.9). Then, for any iteration k_l in which the parameter needs to be increased, it holds that $\rho_{k_l} \geq 2\rho_{k_{l-1}}$. It follows from this result and (7.9) that the penalty parameter goes to infinity if and only if its value is increased in an infinite number of iterations.

When $\rho_k \neq \rho_{k-1}$ from (7.9) and (7.27) it holds that $\widehat{\rho}_k > \rho_{k-1}$, and the definition in (7.29) satisfies $\rho_k \leq 2\widehat{\rho}_k$.

The algorithm may generate an infinite sequence of iterates, or it may find a solution after a finite number of iterations. For the rest of the proofs, and to simplify the arguments, it will be assumed that there is always an infinite sequence $\{x_k\}$, and in the case of finite termination the sequence is completed by repeating an infinite number of times the last point computed. Lemma 7.4.1 implies that in both cases all the quantities associated with the algorithm are well defined.

7.4.3 Properties of the merit function

Lemma 7.4.2. *For any iteration k_l in which the value of ρ is modified,*

$$\rho_{k_l}\|c_{k_l} - s_{k_l}\| \le N \tag{7.30}$$

and

$$\rho_{k_l}(\|p_{k_l}\|^2 + \|u_{k_l}\|^3) \le N \tag{7.31}$$

Proof. It holds that

$$\widehat{\rho}\|c_k - s_k\|^2 = \omega + g_k^T p_k + (2y_k - \pi_k)^T(c_k - s_k)$$

and so

$$\widehat{\rho}\|c_k - s_k\|^2 = (\|c_k - s_k\|^2 + p_k^T\bar{B}p_k) + g_k^T p_k + (2y_k - \pi_k)^T(c_k - s_k).$$

Considering that $g^Tp = -p^T\bar{B}p + p^TJ_W^T\pi$, w corresponding to the final working set constraints, and $J_Wp = -c_W - \widehat{\mu}_W$, the expression becomes

$$\widehat{\rho}\|c_k - s_k\|^2 = \|c_k - s_k\|^2 - c_W^T\pi_W - \widehat{\mu}^T\pi_k + (2y_k - \pi_k)^T(c_k - s_k).$$

It holds that π_k and y_k are uniformly bounded. Hence $-c_W^T\pi \le \|c\|\|\pi\| \le C_\pi\|c - s\|$. Also, $\widehat{\mu} \le \eta_S$ and so $\widehat{\mu}^T\pi \le C\|c - s\|$.

Finally, the continuity of c and boundedness of x imply that $\|c - s\| \le C_{cs}$. These facts combined with the boundedness of multipliers and positive-definiteness of $\bar{B}$ give the result

$$\widehat{\rho}\|c - s\|^2 \le C_1\|c - s\| + C_2\|c - s\| + C_3\|c - s\|,$$

bounding $\widehat{\rho}\|c - s\|$ by a constant.

Since ρ is being modified, equations (7.7), (7.8) and (7.27) imply that,

$$\phi_k''(0;\rho_{k-1},u) = u_k^TH_ku_k + 2g_k^Tp_k + 2(2y_k - \pi_k)^T(c_k - s_k) > -\tfrac{1}{2}\|c_k - s_k\|^2 - \tfrac{1}{2}p_k^T\bar{B}_kp_k + \tfrac{1}{2}u_k^TH_ku_k.$$

Considering that $g^T p = -p^T \bar{B} p - c_W^T \pi \leq -p^T \bar{B} p + C_\pi \|c - s\|$, and the fact that $\|c - s\|$ is bounded, it holds that

$$C_4 \|c_k - s_k\| > \frac{3}{2} p_k^T \bar{B}_k p_k - u_k^T H_k u_k.$$

By construction, $p^T \bar{B} p \geq \beta_{Bp} \|p\|^2$ and Forsgren et al. [32, Lemma 2.4], $u_k^T H_k u_k \leq -C_u \|u_k\|^3$, the expression above becomes,

$$(\|p_k\|^2 + \|u_k\|^3) \leq C_5 \|c_k - s_k\|.$$

Applying equation (7.30) and $2\widehat{\rho} \geq \rho$, the desired result (7.31) follows. □

The subsequent results require the following assumption:

Assumption 7.4.6. *There exists a $\bar{\alpha}$ such that $\alpha_k \geq \bar{\alpha}$ for all k.*

Lemma 7.4.3. *$L_A(x_k, y_k, s_k; \rho_k)$ is bounded from below*

Proof. This follows immediately from assumptions 7.4.4, 7.4.1, and 7.4.5. □

Lemma 7.4.4. *For iterations in which ρ is not changed,*
$\phi_k(\alpha) - \phi_k(0) \leq C\alpha^2(\|p_k\|^2 + \|u_k\|^2)$.

Proof. Using the definition of u based on the penalty parameter update (7.8), the positive-definiteness of $\bar{B}$, and the properties of the direction of negative curvature, it follows that:

$$\begin{aligned}
\phi_k(\alpha) - \phi_k(0) &\leq \sigma(\alpha \phi_k'(0) + \tfrac{1}{2}\alpha^2 \phi_k''(0)) \\
&\leq \sigma(\alpha g_k^T u_k + \tfrac{1}{2}\alpha^2(-\omega + \tfrac{1}{2} u_k^T H_k u_k)) \\
&\leq -\tfrac{1}{2}\sigma\alpha^2(\omega - \tfrac{1}{2} u_k^T H_k u_k) \\
&= -\tfrac{1}{2}\sigma\alpha^2(\|c_k - s_k\|^2 + p_k^T \bar{B}_k p_k - \tfrac{1}{2} u_k^T H_k u_k) \\
&\leq \alpha^2 C(\|p_k\|^2 + \|u_k\|^2).
\end{aligned}$$

□

Theorem 7.4.1. *The algorithm generates a cluster point x^* with $x_{k_l} \to x^*$, $p_{k_l} \to 0$ and $u_{k_l} \to 0$.*

Proof. If ρ_k grows without bound, then the subsequence of iterates $\{x_{k_l}\}$ at which the penalty parameter is changed is bounded and hence has a cluster point. As in Lemma 7.4.2, it holds that $\liminf \|p_{k_l}\| \to 0$ and $\liminf \|u_{k_l}\| \to 0$. It follows that there is a subsequence that satisfies the requirements of Theorems 7.4.2 and 7.4.4.

Now assume that $\rho_k = \bar{\rho}$ for all $k \geq K$. Then, similarly, by Lemmas 7.4.6, 7.4.3, and 7.4.3, that $\liminf \|p_{k_l}\| \to 0$ and $\liminf \|u_{k_l}\| \to 0$. Since the underlying sequence x_{k_l} is Cauchy, it must converge to a cluster point. □

7.4.4 Properties of limit points

The following assumption is necessary for the results of this section:

Assumption 7.4.7. *$\{\bar{B}_k\}$ is bounded.*

An *approximate KKT sequence*, first defined in Section 4.3.3, is defined as:

Definition 7.4.1. *A primal-dual sequence $\{(x_k, y_k)\}$ is an approximate KKT sequence if,*

$$\begin{aligned} g(x_k) + J(x_k)^T y_k &= \epsilon_k \\ c_i(x_k) &\geq \delta_k \\ [y_k]_i &\geq 0, \\ [y_k]_i(c_i(x_k) - \delta_k) &= 0. \end{aligned}$$

where $\{\epsilon_k, \delta_k\}$ converges to zero as $k \to \infty$.

This first result shows that the limit point satisfies the first-order optimality conditions.

Theorem 7.4.2. *The cluster point x^* from Theorem 7.4.1 is a first-order KKT point.*

Proof. By the KKT conditions of (7.17)

$$\begin{aligned} g(x_{k_l}) + \bar{B}_{k_l} p_{k_l} - J(x_{k_l})^T \pi_{k_l} &= 0 \\ c(x_{k_l}) + J(x_{k_l}) p_{k_l} &\geq -\widehat{\mu}_{k_l} \\ \pi_{k_l} &\geq 0 \\ \pi_{k_l}^T (c(x_{k_l}) - J(x_{k_l}) p_{k_l}) &= \delta_{k_l}, \end{aligned}$$

with $\widehat{\mu}_k \to 0$ and $|\delta| \leq \pi^T \widehat{\mu} \to 0$. Theorem 7.4.1 and Assumption 7.4.7 imply that $\bar{B}_{k_l} p_{k_l} \to 0$ and by continuity of J, $J(x_{k_l}) p_{k_l} \to 0$. This implies that the sequence $\{p_{k_l}, \pi_{k_l}\}$ is an approximate KKT sequence. Since the MFCQ implies the CPLD, by Qi and Wei [78, Theorem 2.7], x^* is a first-order KKT point. □

Theorem 7.4.3. *There is a subsequence* $\{k_{l_m}\}$ *such that* $\|\pi_{k_{l_m}} - y^*\| \to 0$ *for some* $y^* \in \mathcal{M}_y(x^*)$*, where* $\mathcal{M}_y(x^*)$ *is the set of multipliers satisfying the first-order KKT conditions. Moreover, the subsequence of* $H_{k_{l_m}} \to H(x^*, y^*)$*.*

In addition, there is an integer K *such that for all* $k_l \geq K$*, the correct active set is identified, in the sense that* $\pi_j = 0$ *for* $j \notin \mathcal{A}(x^*)$*.*

Proof. The first part of the theorem is proved in Qi and Wei [78, Theorem 2.7]. Convergence of the Hessians follows from x_k lying on a compact set and the continuity of H.

Next it will be shown that for any constraint j for which $c_j(x^*) = \delta_1 > 0$, it holds that $[\pi_{k_l}]_j = 0$ for sufficiently large l. If $p_{k_l} + u_{k_l} \to 0$, then the continuity of $J(x)$ implies that $\|p_{k_l} + u_{k_l}\| \leq \delta_1/(4\delta_2)$, where $\delta_2 = \|\nabla c_j(x^*)\|$. It follows that

$$\nabla c_j(x_{k_l})^T (p_{k_l} + u_{k_l}) + c_j(x_{k_l}) \geq \tfrac{1}{2}\delta_1 > 0,$$

which implies that $[\pi_{k_l}]_j = 0$, as required. □

Theorem 7.4.3 implies that $\mathcal{A}_*^c \subset \mathcal{A}_k^c$, which implies $\mathcal{A}_k \subset \mathcal{A}_*$.

Lemma 7.4.5. $\|p_F\|_{k_l} \to 0$, $\|c^-\|_{k_l} \to 0$.

Proof. Without loss of generality, label elements of the subsequence k_l by k.

Since $J(p_k + u_k) \geq -\widehat{\mu}_k - c_k$, $p_k \to 0$ and $u_k \to 0$ imply that $c_k^- \to 0$, which also implies that $p_F \to 0$.

□

Lemma 7.4.6. *For some* K*, for* $k_l \geq K$*,* p_F *identifies the active set.*

Proof. The convergence of $c(x_{k_l}) \to c(x^*)$ implies that for constraints inactive at x^*, $c(x_{k_l})$ is eventually bounded away from zero. $p_F \to 0$ and J bounded implies $c_k + J_k p_F > -\widehat{\mu}_{k_l}$ for the inactive constraints, for large enough k_l. □

Theorem 7.4.4. x^* *satisfies the second order necessary optimality conditions.*

Proof. It has been shown that the limit point x^* is a first-order KKT point by Theorem 7.4.2.

By Theorem 7.4.3 there exists a subsequence k_{l_m} such that $H_{k_{l_m}} \to H(x^*, y^*)$.

Without loss of generality, label elements of the subsequence k_{l_m} by k.

Let $d \in T(x^*) \equiv \{d \mid J_{\mathcal{A}^*}(x^*)d = 0\}$ with $\|d\| = 1$. By Andreani et al. [6, Lemma 3.1] there exists $\{d_k\}$ such that $d_k \in T(x_k) \equiv \{d \mid J_{\mathcal{A}^*}(x_k)d = 0\}$ and $d_k \to d$. Without loss of generality, let $\|d_k\| = 1$.

It holds that

$$d_k^T H_k d_k \geq \lambda_{\min}(Z_k^T H_k Z_k),$$

and by (7.21) (which holds under the computational procedures for finding a direction of negative curvature, e.g., see Forsgren et al. [32, Theorem 3.1]). The limit $u_k^T Z_k^T H_k Z_k u_k \to 0$ implies that $\lambda_{\min}(Z_k^T H_k Z_k) \to 0$. Since H_k converges and $u_k \to 0$ by assumption, $d_k^T H_k d_k \to d^T H^* d \geq 0$. □

7.5 Local Convergence

7.5.1 Additional assumptions

Assumption 7.5.1. *The strong second-order sufficiency condition holds at all points satisfying the second-order necessary conditions.*

Assumption 7.5.2. *The constant rank constraint qualification (CRCQ) holds at all second order KKT points.*

7.5.2 Convergence

Theorem 7.5.1. *For K sufficiently large, $k \geq K$ implies $u_k = 0$, $x \to x^*$, and $y \to y^*$, and $p_k = p_{Sk}$.*

Proof. Let K be sufficiently large such that for the convergent subsequence, $\{k_l\}$, for $k_l \geq K$, Theorem 7.4.3, Lemma 7.4.6 and Assumption 7.5.2 are satisfied. In addition, invoke Qi and Wei [78, Theorem 3.2] to assert that there is only one second-order optimal point in a region around x^*.

Furthermore, from Theorem 7.4.3 it holds that $\mathcal{A}_k \subset \mathcal{A}_*$. Let k_m be the first iterate of the convergent subsequence such that $k_m \geq K$. Possibly by increasing K, the convex subproblem (7.17) can be expressed as an inexact SQP subproblem with $t = (\bar{B} - H)(p_S + \widehat{p})$ (see Wright [87] or Section 5.2 in Chapter 5). Invoke Wright [87, Lemma 4.1] to claim that there exists a $y^* \in \mathcal{M}_y(x^*)$ such that $\mathcal{A}_+(y^*) \subset \mathcal{W}_k(x_{k_m} + p_{k_m})$ (the final working-set of the QP iterations). By the CRCQ, since $\mathcal{A}_+(y^*) \subset \mathcal{W}_k(x_{k_m} + p_{k_m}) \subset \mathcal{A}^*$, and $\mathcal{A}_+(y^*)$ is a maximally linearly independent subset of $\mathcal{A}_+$, $\mathcal{W}_k(x_{k_m} + p_{k_m}) = \mathcal{A}_+(y^*)$.

At iteration $k_m + 1$, p_F estimates the active set with this $\mathcal{W}_{k_m+1}$ as the initial working set, since by the CRCQ, $\mathcal{W}_{k_m+1} = \mathcal{A}_+(y^*)$ is linearly independent. By Kojima [60, Lemma 7.4], the full subproblem solution at x_{k_m+1} satisfies $\mathcal{A}_k \subset \mathcal{A}_*$. Since only linearly dependent constraints would be added in the process of calculating a stationary point, there is a stationary point such that $g_{k_m+1} + H_{k_m+1} p_{k_m+1} = J_W^T \pi_S$ with $\pi_S \geq 0$, where w denotes the rows corresponding to the working set. This implies that a convex subproblem need not be formed and solved and that $(p_{k_m+1}, \pi_{k_m+1}) = (p_S, \pi_S)$.

By Assumption 7.4.1, the QP subproblem is a perturbation of the original NLP. Possibly by increasing K so as to make x_{k_m+1} sufficiently close to x^*, invoke Kojima [60, Lemma 7.5] to assert that the reduced Hessian of the problem is positive definite, so there are no directions of negative curvature for iterations $k \geq k_m$.

The full step taken satisfies the conditions:

$$\begin{aligned} H_{k_m+1} p_{k_m+1} + g_{k_m+1} &= J_{k_m+1}^T \pi_{S,k_m+1} \\ c_{k_m+1} + J_{k_m+1} p_{k_m+1} &\geq -\widehat{\mu}_{k_m+1} \\ (c_{k_m+1} + J_{k_m+1} p_{k_m+1} + \widehat{\mu}_{k_m+1}) \cdot \pi_{S,k_m+1} &= 0. \end{aligned}$$

which is of the form of the inexact SQP problem (4.4) in [87] with $t = \widehat{\mu}$ (alternatively, see (5.2) in Chapter 5). Possibly by making K larger so that x_{k_l} is sufficiently close to x^* and $\widehat{\mu}$ is small enough, invoke Wright [87, Theorem 5.3] (or 5.2.2 in Chapter 5) to assert that

$$\begin{aligned} \|x_{k_m+1} + p_{k_m+1} - x^*\| + \|\pi_{k_m+1} - y^*\| = \|\pi_{k_m+1} - y_{k_m+1}\| O(\|x_{k_m+1} - x^*\|) \\ + O(\|x_{k_m+1} - x^*\|) + O(\|\widehat{\mu}_{k_m+1}\|). \end{aligned}$$

Noting that $\|\widehat{\mu}\|_\infty \leq \eta(x,y) = O(\|x - x^*\|)$, and using Assumption 7.4.6, let the original K be large enough such that

$$\|x_{k_m+1} + \alpha p - x^*\| \leq \gamma \|x_{k_m+1} - x^*\|, \text{ and } \|\alpha \pi + (1-\alpha) y_{k_m+1} - y^*\| \leq \gamma \|y_{k_m+1} - y^*\|.$$

with $\gamma < 1$.

By a similar argument, the equivalent statement for an inductive step can be shown i.e. this estimate holding at $k_m + j$ implies that it also holds for $k_m + j + 1$. This implies that $x_k \to x^*$ and $y_k \to y^*$, where y^* is unique by the fact that $\mathcal{W}_{k_m} = \mathcal{W}_{k_m+1} = \mathcal{W}_{k_m+2} = \cdots = \mathcal{W}_{k_m+l} = \ldots$. □

Theorem 7.5.2. *There exists some K for which $k \geq K$ implies $\alpha_k = 1$.*

Proof. From Theorem 7.5.1, $u = 0$. Now, as in Gill et al. [36] and Powell and Yuan [77], assume

$$\begin{aligned} x_k + p_k - x^* &= o(\|x_k - x^*\|), \\ y_k + q_k - y^* &= o(\|y_k - y^*\|). \end{aligned} \tag{7.32}$$

which may be used to derive the expansions

$$f(x+p) = f(x) + \tfrac{1}{2}(g(x) + g(x^*))^T p + o(\|p\|^2),$$
$$c(x+p) = c(x) + \tfrac{1}{2}(J(x) + J(x^*))^T p + o(\|p\|^2).$$

Consider

$$\phi_k(1) = f_k + \tfrac{1}{2}(g_k + g^*)^T p_k - (y_k + q_k)^T (c(x_k + p_k) - s_k) + \tfrac{1}{2}\rho_k \|c(x_k + p_k) - s_k\|^2$$

and, using the definition of s, as in Gill et al. [36], write the following expression:

$$\phi_k(1) - \phi_k(0) = \tfrac{1}{2}(g_k + g^*)^T p_k - \tfrac{1}{2}\pi_k^T (J^* - J_k) p_k - y_k^T (c_k - s_k) - \tfrac{1}{2}\rho_k (c_k - s_k)^T (c_k - s_k) + o(\|p_k\|^2).$$

From (7.32), $g^* - J_*^T \pi = o(\|p\|)$, which enables this expression to be modified to

$$\phi_k(1) - \phi_k(0) = \tfrac{1}{2} g_k^T p_k + \tfrac{1}{2}\pi_k^T J_k p_k - y_k^T (c_k - s_k) - \tfrac{1}{2}\rho_k (c_k - s_k)^T (c_k - s_k) + o(\|p_k\|^2).$$

To set $\alpha = 1$, it must hold that

$$\phi_k(1) - \phi_k(0) \le \sigma(\phi_k'(0) + \tfrac{1}{2}\phi_k''(0)) \le -\frac{\sigma}{2}(\|c_k - s_k\|^2 + p_k^T \bar{B}_k p_k),$$

where the inequalities follow from the fact that $\phi''(0) \le -\omega$ and the definition of ω (7.7).

Let $\hat{\sigma}$ be such that $\sigma + \hat{\sigma} < \frac{1}{2}$. The last expression becomes:

$$\begin{aligned}\phi_k(1) - \phi_k(0) - \sigma(\phi_k'(0) + \phi_k''(0)) &= \frac{1}{2} g_k^T p_k + (\frac{\sigma}{2} + \hat{\sigma} - \hat{\sigma}) p_k^T \bar{B}_k p_k + \tfrac{1}{2}\pi_k^T J_k p_k - y_k^T (c_k - s_k) \\ &\quad + (\frac{\sigma}{2} - \frac{\rho_k}{2})\|c_k - s_k\|^2 + o(\|p_k\|^2) \\ &\le -\tfrac{1}{2}\pi_k^T (c_k + \widehat{\mu}_k) - \hat{\sigma} p_k^T \bar{B}_k p_k - y^T (c_k - s_k) + o(\|p_k\|^2),\end{aligned}$$

where the equations $\sigma < \frac{1}{2}$ and $p_k^T g_k + \frac{1}{2} p_k^T H_k p_k \le 0$ were used, as well as the QP complementarity conditions $\pi_k^T J_k p_k = -\pi_k^T (c_k + \widehat{\mu}_k)$.

Let k be large enough that $\mathcal{W}_k$ estimates the active set at x^*. Then $\pi_k^T c_k = \pi_k^T c_W$.

Since $c_W(x^*) = 0$, and the linearization of constraint i is always feasible, $c_W = o(\|p_k\|^2)$.

The expression $[y_k^T (c_k - s_k)]_i$ is clearly eventually zero for $i \in \mathcal{A}^c$. By the same argument as above, $[y_k^T (c_k - s_k)]_i$ is $o(\|p_k\|^2)$ for $i \in \mathcal{W}_k$, the stabilized working set. Finally, for $i \in \mathcal{A} \setminus \mathcal{W}_k$, if $[c_k]_i > 0$, then $[s]_i = [c]_i$, and $[c_k - s_k]_i = 0$. Otherwise, since the linearization is always feasible, $[c_k]_i$ must be within $o(\|p_k\|^2)$ of $\widehat{\mu}_k$, and since $[y_k]_i = o(\|p_k\|)$, the entire expression is $o(\|p_k\|^2)$.

This implies that for some K, $k \ge K$ implies

$$\phi_k(1) - \phi_k(0) - \sigma(\phi_k'(0) + \phi_k''(0)) \le 0,$$

and $\alpha = 1$ thereafter. □

Theorem 7.5.3. *There exists a K such that for $k \geq K$, (x_k, y_k) converges superlinearly to (x^*, y^*).*

Proof. Theorem 7.5.1 implies that $x_k \to x^*$ and $y_k \to y^*$. By Theorem 7.5.2, $\alpha_k = 1$ for K sufficiently large. Hence, eventually $y_k = \pi_k$, with $\pi_W^c = 0$ for the consistent working set $\mathcal{W}_k$. This implies that constraints not in $\mathcal{W}_k$ do not appear in the Lagrangian Hessian and the subproblem solutions (p_k, π_k) must satisfy:

$$\begin{aligned} H_k p_k + g_k &= J_W^T \pi_W, \\ c_W + J_W p_k &= 0, \\ \pi_W &\geq 0. \end{aligned}$$

These are the optimality conditions for the extended iSQP problem (8.3) of [87] with $(r, t) = 0$ (see Section 5.2.1 of Chapter 5). Hence we apply Wright [87, Lemma 8.2] (5.2.1 in Chapter 5) and Theorem 7.5.2 to conclude that

$$\|x_k + p_k - x^*\| + \|\pi_k - y^*\| \leq \|x_k - x^*\|^{1+\tau} + \|y_k - y^*\|^{1+\tau}.$$

□

7.6 Discussion

This chapter presented an SQP algorithm that is globally and superlinearly convergent under relatively weak assumptions. It drew from a number of results in the convergence theory literature to utilize the augmented Lagrangian merit function's strong global convergence properties in conjunction with the fast local convergence exhibited by SQP using linearly independent working sets. To fully establish the latter results, which depend on the exact Hessian being used in the subproblem, the method of convexification enabled the algorithm to a) never attempt to solve a nonconvex problem to completion, b) retain some second-derivative information even in the altered problem and c) use exact second-derivatives once the iterates are in a local neighborhood of a solution satisfying the second-order sufficiency conditions for optimality.

Regardless of the results of upcoming numerical tests of the algorithm described in this Chapter, these results should aid in the formulation of robust and efficient SQP algorithms for solving NLPs.

Chapter 8

A Primal-Dual Stabilized SQP Method

8.1 Introduction

In this chapter we focus on optimization problems with constraints written in so-called “standard form”, i.e.,

$$\begin{array}{ll} \underset{x}{\text{minimize}} & f(x) \\ \text{subject to} & c(x) = 0, \qquad x \geq 0, \end{array} \tag{8.1}$$

where $c : \mathbb{R}^n \mapsto \mathbb{R}^m$ and $f : \mathbb{R}^n \mapsto \mathbb{R}$ are twice-continuously differentiable (see Section 3.2, page 38). The purpose is to extend both the scope and the convergence theory of the stabilized primal-dual SQP method proposed by Gill and Robinson [39]. The discussion is in three parts. The first part provides a description of the overall algorithm. The second part is concerned with the convexification of nonconvex QP subproblems and includes the formulation and analysis of a new “concurrent” convexification that is performed during the solution of the QP subproblem. The final part provides an overview of the first-order local and global convergence results. Some changes are proposed that allow local superlinear convergence to be established under weaker assumptions. An extension of the algorithm that converges to points satisfying certain second-order necessary conditions for optimality is considered in Chapter 9.

8.2 Background

8.2.1 The regularized primal-dual line-search SQP algorithm

The regularized SQP line-search method is based on the primal-dual augmented Lagrangian merit function

$$M^\nu(x, y; y^E, \mu) = f(x) - c(x)^T y^E + \frac{1}{2\mu}\|c(x)\|^2 + \frac{\nu}{2\mu}\|c(x) + \mu(y - y^E)\|^2, \tag{8.2}$$

where ν is a scalar, μ is a penalty parameter, and y^E is an estimate of an optimal Lagrange multiplier vector y^*. This function, proposed by Robinson [80], and Gill and Robinson [38], may be derived by applying the primal-dual penalty function of Forsgren and Gill [30] to a problem in which the constraints are shifted by a constant vector (see Powell [76]). 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

$$\nabla M^\nu(x, y; y^E, \mu) = \begin{pmatrix} g - J^T\big((1+\nu)(y^E - \frac{1}{\mu}c) - \nu y\big) \\ \nu\big(c + \mu(y - y^E)\big) \end{pmatrix} \tag{8.3a}$$

$$= \begin{pmatrix} g - J^T\big(\pi + \nu(\pi - y)\big) \\ \nu\mu(y - \pi) \end{pmatrix}, \tag{8.3b}$$

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

$$\pi(x; y^E, \mu) = y^E - \frac{1}{\mu}c(x). \tag{8.4}$$

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

$$\nabla^2 M^\nu(x, y; y^E, \mu) = \begin{pmatrix} H\big(x, \pi + \nu(\pi - y)\big) + \frac{1}{\mu}(1+\nu)J^T J & \nu J^T \\ \nu J & \nu\mu I \end{pmatrix}. \tag{8.5}$$

The terms $M^\nu(x, y)$, $\nabla M^\nu(x, y)$, and $\nabla^2 M^\nu(x, y)$ are used to denote M^ν, ∇M^ν, and $\nabla^2 M^\nu$ evaluated with parameters y^E and μ.

8.2.2 Definition of the search direction

Consider a quadratic approximation to the primal-dual function M^ν based on an approximate Hessian $H_M^\nu \approx \nabla^2 M^\nu$ such that

$$H_M^\nu(x, y; \mu) = \begin{pmatrix} \bar{H}(x, y) + \frac{1}{\mu}(1+\nu)J(x)^T J(x) & \nu J(x)^T \\ \nu J(x) & \nu\mu I \end{pmatrix}, \tag{8.6}$$

where $\bar{H}(x,y)$ is a symmetric approximation to $H\big(x, \pi + \nu(\pi - y)\big) \approx H(x,y)$ such that $\bar{H}(x,y) + \frac{1}{\mu}J(x)^T J(x)$ is positive definite. The approximation $\pi + \nu(\pi - y) \approx y$ is valid provided $\pi \approx y$. The restriction on the inertia of $\bar{H}$ implies that $H_M^{\nu}(x, y; \mu)$ is positive definite for $\nu > 0$ and positive semidefinite for $\nu = 0$.

Using this definition of H_M^{ν} at the kth primal-dual iterate $v_k = (x_k, y_k)$, consider the convex QP subproblem

$$\underset{\Delta v = (p,q)}{\text{minimize}} \; \nabla M^{\nu}(v_k)^T \Delta v + \tfrac{1}{2}\Delta v^T H_M^{\nu}(v_k)\Delta v \quad \text{subject to} \;\; x_k + p \geq 0, \tag{8.7}$$

where $M^{\nu}(v)$ denotes the merit function evaluated at $v = (x, y)$.

The following result provides a useful equivalent definition for the search direction.

Theorem 8.2.1 (Gill and Robinson [39, Theorem 3.3]). *For any primal-dual QP solution $\Delta v_k = (p_k, q_k)$, the first-order conditions associated with the variables in the free part of $x_k + p_k$ may be written in matrix form as:*

$$\begin{pmatrix} \bar{H}_F & -J_F^T \\ J_F & \mu I \end{pmatrix} \begin{pmatrix} p_F \\ q_k \end{pmatrix} = -\begin{pmatrix} [g_k - J_k^T y_k - \bar{H} s_k]_F \\ c_k + \mu(y_k - y^E) - J_k s_k \end{pmatrix}, \tag{8.8}$$

where c_k, g_k and J_k denote the quantities $c(x)$, $g(x)$ and $J(x)$ evaluated at x_k, and s_k is a nonnegative vector such that

$$[s_k]_i = \begin{cases} [x_k]_i & \text{if } i \in \mathcal{A}(x_k + p_k); \\ 0 & \text{if } i \in \mathcal{F}(x_k + p_k). \end{cases}$$

□

(The assumption of positive-definiteness of $\bar{H}_k + \frac{1}{\mu}J_k^T J_k$ implies that the matrix associated with the equations (8.8) is nonsingular.) It follows that if $\mathcal{A}(x_k + p_k) = \mathcal{A}(x_k)$, then s_k is zero and (p_k, q_k) satisfies the perturbed Newton equations

$$\begin{pmatrix} H_F & -J_F^T \\ J_F & \mu I \end{pmatrix} \begin{pmatrix} p_F \\ q_k \end{pmatrix} = -\begin{pmatrix} [g_k - J_k^T y_k]_F \\ c_k + \mu(y_k - y^E) \end{pmatrix}.$$

Given an iterate $v_k = (x_k, y_k)$ and Lagrange multiplier estimate y_k^E, the primal-dual search direction $\Delta v_k = (p_k, q_k)$ is defined such that $v_k + \Delta v_k = (x_k + p_k, y_k + q_k)$ is a solution of the convex QP problem

$$\begin{aligned} &\underset{v=(x,y)}{\text{minimize}} \quad (v - v_k)^T \nabla M^{\nu}(v_k; y_k^E, \mu_k^R) + \tfrac{1}{2}(v - v_k)^T H_M^{\nu}(v_k; \mu_k^R)(v - v_k) \\ &\text{subject to} \;\; x \geq 0, \end{aligned} \tag{8.9}$$

where μ_k^R is a small parameter, and $H_M^\nu(v_k;\mu_k^R)$ is the matrix (8.6) evaluated at $v_k = (x_k, y_k)$. 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.

In general, augmented Lagrangian-based methods keep the penalty parameter μ as large as possible (see, e.g., [18, 35]), whereas stabilized SQP methods keep μ small (see Chapter 5). This motivates using a small μ for the quadratic subproblem, and a larger μ for the merit function.

Finally, note that if $v = v_k$ is a solution of the QP (8.9), then v_k is a first-order solution of

$$\underset{v=(x,y)}{\text{minimize}} \;\; M^\nu(v; y_k^E, \mu_k^R) \quad \text{subject to} \;\; x \ge 0. \tag{8.10}$$

The following result provides an important link between the primal-dual SQP method and regularized SQP methods,

Theorem 8.2.2 (Gill and Robinson [39, Theorem 3.1])**.** *The primal-dual vector* $v_k + \Delta v_k = (x_k + p_k, y_k + q_k)$ *is a solution of problem* (8.9) *if and only if it solves the stabilized SQP problem,*

$$\begin{aligned} \underset{x,y}{\text{minimize}} \quad & g_k^T(x - x_k) + \tfrac{1}{2}(x - x_k)^T \bar{H}(x_k, y_k)(x - x_k) + \tfrac{1}{2}\mu_k^R \|y\|^2 \\ \text{subject to} \quad & c_k + J_k(x - x_k) + \mu_k^R(y - y_k^E) = 0, \;\; x \ge 0. \end{aligned} \tag{8.11}$$

□

8.2.3 Definition of the new iterate

This algorithm uses a "flexible penalty function" as defined in Curtis and Nocedal [20]. Let $\alpha_k = 2^{-j}$, where j is the smallest nonnegative integer such that

$$M^\nu(v_k + \alpha_k \Delta v_k; y_k^E, \mu_k^F) \le M^\nu(v_k; y_k^E, \mu_k^F) + \alpha_k \eta_S \delta_k \tag{8.12}$$

for some value $\mu_k^F \in [\mu_k^R, \mu_k]$, and the scalar

$$\delta_k \triangleq \max\left(\Delta v_k^T \nabla M^\nu(v_k; y_k^E, \mu_k^R), -10^{-3}\|\Delta v_k\|^2\right) \le 0 \tag{8.13}$$

is a sufficiently negative value used in the proof of global convergence. Once an appropriate value for α_k is found, the new primal-dual iterate is given by

$$x_{k+1} = x_k + \alpha_k p_k \quad \text{and} \quad y_{k+1} = y_k + \alpha_k q_k.$$

8.2.4 Updating the multiplier estimate

Consider the functions

$$\phi^S(v) = \xi(x) + 10^{-5}\omega(v) \quad \text{and} \quad \phi^L(v) = 10^{-5}\xi(x) + \omega(v), \tag{8.14}$$

where

$$\xi(x) = \|c(x)\| \quad \text{and} \quad \omega(x,y) = \left\|\min\left(x, g(x) - J(x)^T y\right)\right\| \tag{8.15}$$

are feasibility and stationarity measures at the point (x, y), respectively. These optimality measures are based on the optimality conditions for problem (8.1) rather than for minimizing the merit function M^ν. Both measures are bounded below by zero, and are equal to zero if v is a first-order solution to problem (8.1).

The multiplier estimate y_k^E is updated when v_k satisfies either $\phi^S(v_k) \leq \frac{1}{2}\phi^S_{\max}$ or $\phi^L(v_k) \leq \frac{1}{2}\phi^L_{\max}$, where $\phi^S_{\max}$ and $\phi^L_{\max}$ are bounds that are updated throughout the solution process. To ensure global convergence, the update to y_k^E is accompanied by a decrease in either $\phi^S_{\max}$ or $\phi^L_{\max}$.

Finally, y_k^E is also updated if an approximate first-order solution of the problem

$$\underset{x,y}{\text{minimize}} \;\; M^\nu(x, y; y_k^E, \mu_k^R) \quad \text{subject to} \;\; x \geq 0 \tag{8.16}$$

has been found. The test for optimality is

$$\|\nabla_y M^\nu(v_{k+1}; y_k^E, \mu_k^R)\| \leq \tau_k \;\; \text{and} \;\; \left\|\min\left(x_{k+1}, \nabla_x M^\nu(v_{k+1}; y_k^E, \mu_k^R)\right)\right\| \leq \tau_k \tag{8.17}$$

for some small tolerance $\tau_k > 0$. Numerical experiments have shown that it is rare for an iterate to not satisfy $\phi^S(v_k) \leq \frac{1}{2}\phi^S_{\max}$ or $\phi^L(v_k) \leq \frac{1}{2}\phi^L_{\max}$ and satisfy this condition, however, the condition is still necessary for the convergence theory. It was shown in Gill and Robinson [39] that a sequence of iterates converges to either a point satisfying the KKT conditions or generates a sequence of M-iterates, converging to a stationary point of M^ν. It will be shown in this Chapter that under certain, conditions, a stationary point of M^ν is also a local minimizer.

If the condition (8.17) is satisfied, y_k^E is updated with the *safeguarded* estimate

$$y_{k+1}^E = \left(-10^6, y_{k+1}, 10^6\right).$$

8.2.5 Updating the penalty parameters

Since μ_k^R should only be decreased when "close" to optimality (ignoring locally infeasible problems), use the definition

$$\mu_{k+1}^R = \begin{cases} \min\left(\frac{1}{2}\mu_k^R, \|\eta_k\|^\gamma\right), & \text{if (8.17) is satisfied;} \\ \min\left(\ \mu_k^R, \|\eta_k\|^\gamma\right), & \text{otherwise,} \end{cases} \tag{8.18}$$

where $0 < \gamma < 1$ and, as in the regularized SQP methods, η is defined to a measure of the primal-dual optimality violation,

$$\eta_{k+1} \equiv \eta_{opt}(v_{k+1}) \triangleq \begin{pmatrix} c(x_{k+1}) \\ \min\left(x_{k+1}, g(x_{k+1}) - J(x_{k+1})^T y_{k+1}\right) \end{pmatrix}. \tag{8.19}$$

This choice of η_{k+1} ensures that, as v_k approaches a primal-dual solution, and the asymptotically superlinearly convergent region is reached, μ_k^R is eventually equal to η_k^γ, which is a typical value of μ^R used in the local convergence analysis of stabilized SQP methods. Far from the solution, μ is held to be small so as to not perturb the problem too far from a conventional SQP step, while monotonically converging to zero along iterates that improve the merit function.

The update to the penalty parameter μ_k is motivated by the goal of decreasing μ_k only when the trial step indicates that the merit function *increases*. The algorithm uses the update,

$$\mu_{k+1} = \begin{cases} \mu_k, & M^\nu(v_{k+1}; y_k^E, \mu_k) \le M^\nu(v_k; y_k^E, \mu_k) + \widehat{\alpha}_k \eta_S \delta_k \\ \max\left(\frac{1}{2}\mu_k, \mu_{k+1}^R\right), & \text{otherwise,} \end{cases} \tag{8.20}$$

where δ_k is defined in (8.13) and $\widehat{\alpha}_k = \min(\alpha_{\min}, \alpha_k)$ for some positive $\alpha_{\min}$. The use of the scalar $\alpha_{\min}$ increases the likelihood that μ_k will not be decreased.

8.2.6 Formal statement of the algorithm

In this section pdSQP is formally stated as Algorithm 8.2.1 and include some additional details. During each iteration, the trial step is computed as described in Section 8.2.2, the solution estimate is updated as in Section 8.2.3, y_k^E is updated as in Section 8.2.4, and the penalty parameters are updated as in Section 8.2.5. The value of y_k^E is crucial for both global and local convergence. To this end, 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^S or ϕ^L given by (8.14). These iterates are labeled as S- and L-iterates, respectively. It is to be expected that y_k^E

will be updated in this way most of the time. Second, if (x_{k+1}, y_{k+1}) is not acceptable to either of the merit functions ϕ^S or ϕ^L, check whether an approximate first-order solution to problem (8.16) has been computed by verifying conditions (8.17) for the current value of τ_k. If these conditions 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 as in (8.2.4). Finally, an iterate at which neither of the first two cases occur is called an F-iterate. The multiplier estimate y_k^E is not changed in an F-iterate.

Algorithm 8.2.1. Regularized primal-dual SQP algorithm **pdSQP**

Input (x_0, y_0);

Set algorithm parameters $\alpha_{\min} > 0$, $\eta_S \in (0,1)$, $\tau_{\text{stop}} > 0$, 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 $k = 0, 1, 2, \ldots$ **do**

  Define $\bar{H}_k \approx H(x_k, y_k)$ such that $\bar{H}_k + (1/\mu_k^R) J_k^T J_k$ is positive definite;

  Solve the QP (8.9) for the search direction $\Delta v_k = (p_k, q_k)$;

  Find an α_k satisfying (8.12) and (8.13);

  Update the primal-dual estimate $x_{k+1} = x_k + \alpha_k p_k, \quad y_{k+1} = y_k + \alpha_k q_k$;

  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 $\phi^S(x_{k+1}, y_{k+1}) \le \frac{1}{2}\phi^S_{\max}$ **then** [S-iterate]

    $\phi^S_{\max} = \frac{1}{2}\phi^S_{\max}$;

    $y^E_{k+1} = y_{k+1}$;

    $\tau_{k+1} = \tau_k$;

  else if $\phi^L(x_{k+1}, y_{k+1}) \le \frac{1}{2}\phi^L_{\max}$ **then** [L-iterate]

    $\phi^L_{\max} = \frac{1}{2}\phi^L_{\max}$;

    $y^E_{k+1} = y_{k+1}$;

    $\tau_{k+1} = \tau_k$;

  else if $v_{k+1} = (x_{k+1}, y_{k+1})$ satisfies (8.17) [M-iterate]

    $y^E_{k+1} = \text{middle}(-10^6, y_{k+1}, 10^6)$;

    $\tau_{k+1} = \frac{1}{2}\tau_k$;

  else [F-iterate]

    $y^E_{k+1} = y^E_k$;

    $\tau_{k+1} = \tau_k$;

  end if

Update μ_{k+1}^R and μ_{k+1} according to (8.18) and (8.20), respectively;

if $\|r_k\| \le \tau_{\text{stop}}$ **then** *exit*;

end (for)

8.3 Properties of the subproblem

8.3.1 Convexification

Introduction

The pdSQP convexification procedure in Gill and Robinson [39] proceeds as follows. If H_F is not positive-definite, then a value μ_H is found such that $H + \frac{1}{\mu_H} P_A P_A^T$ is positive definite, similarly to the procedure described for standard form problems in Chapter 6.

To review the results of Gill and Robinson [39], consider a subset $\mathcal{A}$ of the integers $\{1, 2, \ldots, n\}$ and let J_F and J_A denote the columns of J associated with the set $\mathcal{A}$ and its complement $\{1, 2, \ldots, n\}/\mathcal{A}$. (This notation indicates that $\mathcal{A}$ is often chosen as the set $\mathcal{F}(x)$ of free variables.) For given H and J, let K and K_F denote the matrices

$$K = \begin{pmatrix} H & J^T \\ J & -\mu I \end{pmatrix} \quad \text{and} \quad K_F = \begin{pmatrix} H_F & J_F^T \\ J_F & -\mu I \end{pmatrix}, \tag{8.21}$$

i.e., K_F is the matrix of $m + n_F$ rows and columns of K corresponding to the index set $\mathcal{A}$. A set $\mathcal{A}$ for which K_F has inertia $(n_F, m, 0)$ is called a *second-order consistent basis*.

Suppose that we wish to define a convex QP at a point x_0 at which a second-order consistent basis $\mathcal{A}$ is known.

Lemma 8.3.1. *If the KKT matrix* K_F (8.21) *is defined in terms of a second-order consistent basis* $\mathcal{A}$*, then the matrix*

$$B = \begin{pmatrix} H_F + \frac{1}{\mu}(1+\nu) J_F^T J_F & \nu J_F^T \\ \nu J_F & \nu \mu I \end{pmatrix}, \tag{8.22}$$

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

This implies the following primary result.

Theorem 8.3.1. *If the KKT matrix* K_F (8.21) *is defined in terms of a second-order consistent* $\mathcal{A}$*, then the matrix*

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

is positive definite. □

The remaining results of this section relate the solution of the convexified problem to the solution original problem. It is shown that if the active set does not change, then the iterations are the same, suggesting the practicality of a procedure that adds additional terms to the Hessian only if active constraints are dropped. It is also shown that the convexification procedure is unnecessary for convex nonlinear problems.

Relation between the convexified problem and the original problem

The pdSQP subproblem is of the form:

$$\begin{aligned} \underset{\Delta v}{\text{minimize}} \quad & g_M^T \Delta v + \Delta v^T H_M \Delta v, \\ \text{subject to} \quad & x + p \geq 0. \end{aligned} \tag{8.23}$$

At a stationary point for this subproblem, the following conditions hold:

$$[g + Hp - J^T y]_F = 0, \tag{8.24}$$

$$[g + Hp - J^T y]_A = z_A, \tag{8.25}$$

$$c + \mu(y - y^E) + Jp = 0. \tag{8.26}$$

Consider the convexified Hessian $\bar{H} = H + \frac{1}{\mu_H} P_A P_A^T$. The following statement about the QP iterations holds.

Proposition 8.3.1. *If there is no change in the active set, the solution to the convexified subproblem is a solution to the original indefinite subproblem.*

Proof. A step for the convexified problem satisfies

$$\begin{pmatrix} \bar{H}_F & -J_F^T \\ J_F & \mu_R I \end{pmatrix} \begin{pmatrix} p_F \\ q_j \end{pmatrix} = - \begin{pmatrix} [g + \bar{H} p_j - J^T y_j]_F \\ c + \mu_H (y_j - y^E) + J p_j \end{pmatrix}, \tag{8.27}$$

$\bar{H}_F = P_F^T (H + \frac{1}{\mu_H} P_A P_A^T) P_F = H_F + \frac{1}{\mu_H} P_F^T P_A P_A^T P_F$. The product $P_F^T P_A$ is just a matrix of inner products of the columns of P_A and P_F. Since P_A has columns e_i, $i \in \mathcal{A}$ and P_F has columns e_j, $j \in \mathcal{F}$, and $\mathcal{A} \cap \mathcal{F}$ is empty, $e_i^T e_j = 0$ for all $i \in \mathcal{A}$ and $j \in \mathcal{F}$, so $P_F^T P_A = 0$. The matrix on the left-hand side of (8.27) remains unchanged.

The matrix $[A]_F$ is equivalent to $P_F^T A$. Consider $[\bar{H} p_j]_F = P_F^T (H + \frac{1}{\mu_H} P_A P_A^T) p_j$. Similarly, $P_F^T P_A = 0$ and so $[\bar{H} p_k] = [H p_j]$ and the right hand side of (8.27) remains unchanged if $\bar{H}$ is replaced by H. □

However, convexification changes the reduced costs, and therefore may change the optimality of the active set. This is because, at a solution of the convex problem,

$$[g + Hp + \frac{1}{\mu_H} P_A P_A^T p - J^T y]_A = z_A + \frac{1}{\mu_H} P_A P_A^T p = z_A - \frac{1}{\mu_H} [x_0]_A,$$

which implies that the reduced costs may have different signs.

Suppose that $z_i < 0$ and a direction is computed that moves off constraint i. The direction for the convex problem is identical to the step for the problem with $H_1 = H + (1/\mu_H) e_i e_i^T$. Consider iterating this procedure until a solution $\bar{p}$ is found satisfying

$$[g + H\bar{p} + \frac{1}{\mu_H} P_A P_A^T \bar{p} - J^T \bar{y}]_F = 0, \tag{8.28}$$

$$[g + H\bar{p} + \frac{1}{\mu_H} P_A P_A^T \bar{p} - J^T \bar{y}]_A = \bar{z} \geq 0. \tag{8.29}$$

Note that $[g + H\bar{p} + (1/\mu_H) P_A P_A^T \bar{p} - J^T y]_A = [g + H\bar{p} - J^T y]_A$, so the reduced costs for the convex problem are also reduced costs for the original problem *if* $\bar{p}$ is also a stationary point for the original problem.

However, $[g + H\bar{p} - J^T y]_F = -[(1/\mu_H) P_A P_A^T \bar{p}]_F$ is nonzero if some of the constraints originally active at the start of the QP subproblem became inactive during the algorithm iterations. Since, in this case, the free components of the reduced costs are not zero, $\bar{p}$ is not a stationary point for the original QP.

Strictly convex problems

In this section, it is shown that it is not necessary to convexify a strictly convex problem. Consider the QP subproblem in inequality form (3.2). If this problem is convex, the objective $f(x)$ is convex, the inequality constraints $c(x) \geq 0$ are concave, and the equality constraints $c(x) = 0$ are linear. In order to simplify the discussion only inequality constraints are considered, although the theory is easily extended to problems with linear equalities. If the problem is transformed to standard form, the constraints $c(x) \geq 0$ are changed to $c(x) - s = 0$, $s \geq 0$, and H_M has the form,

$$H_M = \begin{pmatrix} H + \frac{1}{\mu}(1+\nu) J^T J & -\frac{1}{\mu}(1+\nu) J^T & J^T \\ -\frac{1}{\mu}(1+\nu) J & \frac{1}{\mu}(1+\nu) I & -I \\ J & -I & \nu\mu I \end{pmatrix}. \tag{8.30}$$

If the Hessian of the Lagrangian is positive definite for all x and y, then the following result holds.

Lemma 8.3.2. *If H is positive definite and $\nu \geq 0$, then H_M is positive definite.*

Proof. Consider the quadratic form $d^T H_M d$ for the Hessian H_M of (8.30). If H_M is positive definite, then for any $d \neq 0$, it holds that $d^T H_M d > 0$. If d is partitioned as $d = (u, v, w)$, the definition of H_M (8.30) gives

$$\begin{aligned} d^T H_M d &= u^T H u + \frac{1}{\mu}(1+\nu)(Ju)^T(Ju) - \frac{1}{\mu}(1+\nu)u^T J^T v + \nu u^T J^T w \\ &\quad - \frac{1}{\mu}(1+\nu)v^T J u + \frac{1}{\mu}(1+\nu)v^T v - \nu v^T w + \nu w^T J u - \nu w^T v + \mu\nu w^T w \\ &> (Ju - v)^T \left(\frac{1}{\mu}(1+\nu)(Ju - v) + 2\nu w \right) + \mu\nu w^T w \\ &\geq \frac{1}{\mu}(1+\nu)(Ju-v)^T(Ju-v) - 2\nu\|Ju - v\|\|w\| + \mu\nu\|w\|^2 \\ &\geq \left(\frac{1}{\sqrt{\mu}}\|Ju - v\| - \sqrt{\mu}\|w\| \right)^2 \\ &\geq 0. \end{aligned}$$

This proves the positive definiteness of H_M. □

If f is not strictly convex but simply convex, then H can only be said to be positive-semidefinite everywhere and H_M is everywhere positive-semidefinite.

Lemma 8.3.3. *The dual solution of* (8.23) *satisfies* $(y+q)_i \geq 0$.

Proof. By Theorem 8.2.2, the solution of the pdSQP subproblem is the same as the solution of the stabilized SQP subproblem. From the primal optimality conditions, it holds that,

$$g + Hp - J^T(y+q) - z = 0.$$

For the slack variable s_i, this condition is $(y+q)_i = z_i$. Since $z_i \geq 0$, so is $y_i + q_i$. □

Lemma 8.3.4. *For a convex problem, if y_0 is chosen such that $y_0 \geq 0$, then H_M is always positive definite.*

Proof. As mentioned above, for a convex problem, f is convex, c_i is linear for i such that $c_i(x) = 0$, concave for $c_i(x) \geq 0$ and convex for $c_i(x) \leq 0$.

Therefore $H = \nabla^2 f(x) - \sum y_i \nabla^2 c_i(x)$ is always positive definite if y is such that $y_i \geq 0$ for $c_i(x) \geq 0$.

By the assumption, $y_i \geq 0$ holds for the initial point. Now assume $y_i \geq 0$ at iteration k. Then, by Lemma 8.3.3, $(y+q)_i \geq 0$, so $(y + \alpha q)_i \geq 0$.

Hence $H(x_k, y_k) = \nabla^2_{xx} f(x_k) - \sum_i y_i \nabla^2_{xx} c_i(x_k)$ is always positive definite, since $f(x)$ is convex, $y_i \geq 0$ and $\{c_i(x)\}$ are concave. By the Lemma 8.3.2 of this section, H_M is always positive definite. □

8.4 Solving the QP subproblem

8.4.1 Definition of the QP step

The bound constrained quadratic program has the form

$$\underset{\Delta v=(p,q)}{\text{minimize}}\; g_M(v_k)^T\Delta v + \tfrac{1}{2}\Delta v^T H_M^\nu(v_k)\Delta v \quad \text{subject to} \quad x_k + p \geq 0. \tag{8.31}$$

In the convergence theory of Gill and Robinson [39], it is assumed that $\bar{H} + \frac{1+\nu}{\mu}J^TJ$ is positive-definite. If this is not the case, then a modified LBL^T factorization is used to obtain a diagonal matrix D such that the matrix of free rows and columns of $H + D + \frac{1+\nu}{\mu}J^TJ$ is positive definite (see Chapter 9 for details). If a nonzero matrix D is found, the variables corresponding to the nonzero elements of D are temporarily fixed at their currents values before the QP subproblem is started. These temporary constraints are removed during the iterations of the QP subproblem. The initial nonbasic set of "real" and temporary bounds defines a second-order consistent working set (see Section 3.2.2 of Chapter 3). The QP subproblem is solved using an inertia-controlling QP method, described below.

8.4.2 An inertia-controlling method for bound-constrained QP

An inertia-controlling method for the bound-constrained QP,

$$\underset{x\in\mathbb{R}^n}{\text{minimize}}\;\; \varphi(x) = g^T(x-x_0) + \tfrac{1}{2}(x-x_0)^TH(x-x_0) \quad \text{subject to} \quad x \geq 0, \tag{8.32}$$

generates a sequence of primal-dual iterates (x_j, z_j) such that $z_j = g(x_j)$ and every x_j is a subspace minimizer with respect to the current basis. The method generates a sequence of *sets of consecutive iterates* such that the first and last iterate of each set is a standard subspace minimizer. At the first point of each set, a nonbasic variable x_s with a negative dual variable is identified. In the sequence of subsequent intermediate iterates, there is at most one strictly positive nonbasic variable (the variable with index ν_s). The set of intermediate iterates ends at a point at which the dual variable for x_s has been driven to zero. At this point, the variable x_s is made basic, which implies that the last iterate is a standard subspace minimizer with respect to the new basis.

For the moment we focus on a set of consecutive iterates that starts at a *standard* subspace minimizer x_j. If $g(x_j)$ is nonnegative, then x_j is the solution of the QP and the algorithm is terminated. Otherwise, there is at least one strictly negative component of $g_N(x_j)$ (say, the s-th nonbasic, which corresponds to variable x_s), and hence there exists a direction p, such that $g^Tp + p^THp < 0$ and $p = e_s$. Movement along p_j causes the nonbasic variable x_s to become strictly satisfied. An appropriate direction is given by $p_j = x - x_j$, where x is the solution of the equality-constraint QP:

$$\underset{x\in\mathbb{R}^n}{\text{minimize}}\; g(x_j)^T(x - x_j) + \tfrac{1}{2}(x - x_j)^T H(x - x_j) \quad \text{subject to } [x]_N = e_s.$$

Given (x_j, z_j) we define the optimal (x, z) in the form $(x_j + p_j, z_j + q_j)$, with

$$\begin{aligned} [x_j + p_j]_B &\geq 0, \qquad & [x_j + p_j]_N &= e_s, \\ [z_j + q_j]_B &= 0, \qquad & [z_j + q_j]_N &\geq 0. \end{aligned}$$

As in the previous section, the equalities $[x_j + p_j]_N = e_s$ and $[z_j + q_j]_B = 0$ are written in terms of the equations

$$\begin{pmatrix} H_B & H_D & -I_B & \\ H_D^T & H_N & & -I_N \\ & & I_B & \\ & I_N & & \end{pmatrix} \begin{pmatrix} p_B \\ p_N \\ q_B \\ q_N \end{pmatrix} = - \begin{pmatrix} g_B(x_j) - [z_j]_B \\ g_N(x_j) - [z_j]_N \\ [z_j]_B \\ [x_j]_N - e_s \end{pmatrix}.$$

As $z_j = g(x_j)$, and x_j is a subspace minimizer, it must hold that $g_B(x_j) = 0$, in which case the equations simplify to give

$$\begin{pmatrix} H_B & H_D & -I_B & \\ H_D^T & H_N & & -I_N \\ & & I_B & \\ & I_N & & \end{pmatrix} \begin{pmatrix} p_B \\ p_N \\ q_B \\ q_N \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ e_s - [x_j]_N \end{pmatrix}.$$

If x_j is a standard subspace minimizer, then $[x_j]_N = 0$, and the equations can be written as

$$\begin{pmatrix} H_B & H_D & \\ H_D^T & H_N & -I_N \\ & I_N & \end{pmatrix} \begin{pmatrix} p_B \\ p_N \\ q_N \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ e_s \end{pmatrix}. \tag{8.33}$$

It follows that $p_N = e_s$, $q_N = [Hp]_N$ and p_B satisfies the equations

$$H_B p_B = -[He_s]_B = -[h_s]_B. \tag{8.34}$$

In practice, p_N is defined implicitly and only the components of p_B and q_N need be computed explicitly.

For any scalar step length α, the values of p_j and q_j specified by (8.33) give

$$g_B(x_j + \alpha p_j) = g_B(x_j) + \alpha[Hp_j]_B = g_B(x_j) = 0, \tag{8.35}$$

which implies that *every* point on the ray $x_j + \alpha p_j$ is a subspace minimizer with respect to $\mathcal{B}$. Moreover, the directional derivative and curvature along p_j are given by

$$g(x_j)^T p_j = z_N^T p_N = (z_N)_s, \quad \text{and} \quad p_j^T H p_j = p_N^T q_N = (q_N)_s. \tag{8.36}$$

Once the direction pair (p_j, q_j) is computed, a nonnegative step length α_j is computed so that $x_j + \alpha_j p_j$ is feasible and $\varphi(x_j + \alpha_j p_j) \le \varphi(x_j)$. If $p_j^T H p_j > 0$, the step that minimizes $\varphi(x_j + \alpha_j p_j)$ as a function of α is given by $\alpha_j^* = -g(x_j)^T p_j / p_j^T H p_j$. The identities above give

$$\alpha_j^* = -g(x_j)^T p_j / p_j^T H p_j = -(z_N)_s / (q_N)_s.$$

Since $(z_N)_s < 0$, if $(q_N)_s = p_j^T H p_j > 0$, the optimal step length α_j^* is positive. If $p_j^T H p_j \le 0$, then $(q_N)_s = p_j^T H p_j \le 0$ and φ has no bounded minimizer along p_j and $\alpha_j^* = +\infty$.

If $x_j + \alpha_j^* p_j$ is unbounded or infeasible, then α must be limited by $\bar{\alpha}_j$, the *maximum feasible step* from x_j along p_j. The feasible step length is defined as $\bar{\alpha}_j = \gamma_t = \min_i \{\gamma_i\}$, where

$$\gamma_i = \begin{cases} \dfrac{(x_j)_i}{-(p_j)_i}, & \text{if } (p_j)_i < 0, \\[2ex] +\infty, & \text{otherwise.} \end{cases}$$

The step length α_j is then $\min\{\alpha_j^*, \bar{\alpha}_j\}$. If $\alpha_j^* = +\infty$, the QP has no bounded solution and the algorithm terminates. (For brevity, the calculation of γ_r is summarized as $\gamma_r =$ **`minRatioTest`**(x_j, p_j).) Once a bounded α_j has been defined, the new iterate is $x_{j+1} = x_j + \alpha_j p_j$. The composition of the new working set and multipliers depends on the definition of α_j.

Case 1: $\alpha_j = \alpha_j^*$. In this case, the curvature $(q_N)_s$ must be positive, and the step length $\alpha_j = \alpha_j^* = -(z_N)_s/(q_N)_s$ minimizes $\varphi(x_j + \alpha p_j)$ with respect to α, giving the s-th element of $z_N + \alpha q_N$ as

$$(z_N + \alpha q_N)_s = (z_N)_s + \alpha_j^* (q_N)_s = 0.$$

This identity shows that the Lagrange multiplier associated with the nonbinding nonbasic variable is *zero* at $x_j + \alpha_j^* p_j$. This result, when used in conjunction with (8.35), implies that x_{j+1} is a subspace stationary point with respect to $\mathcal{B} + \{\nu_s\}$. The following argument shows that the new reduced Hessian is positive definite at x_{j+1} and is hence a subspace minimizer. The reduced Hessian with respect to the new basis is given by

$$\begin{pmatrix} H_B & [h_s]_B \\ [h_s]_B^T & h_{\nu_s,\nu_s} \end{pmatrix},$$

which is positive definite if the Schur complement $h_{\nu_s,\nu_s} - [h_s]_B^T H_B^{-1}[h_s]_B$ is positive. The definitions $q_N = [Hp_j]_N$ and $p_N = e_s$ imply that

$$(q_N)_s = e_s^T(H_D^T p_B + H_N p_N) = [h_s]_B^T p_B + h_{\nu_s,\nu_s} = h_{\nu_s,\nu_s} + [h_s]_B^T p_B.$$

The definition of p_B (8.34) gives

$$(q_N)_s = h_{\nu_s,\nu_s} - p_B^T H_B p_B = h_{\nu_s,\nu_s} - [h_s]_B^T H_B^{-1}[h_s]_B,$$

which is positive, as required.

Case 2: $\alpha_j = \bar{\alpha}_j$. In this case, α_j is the step to the bound on x_t. If the index t corresponds to the r-th basic variable, then the index β_r is moved from the basic set to the nonbasic set at $x_j + \alpha_j p_j$. The following argument shows that $x_j + \alpha_j p_j$ is a subspace minimizer with respect to the new basic set $\mathcal{B} - \{\beta_r\}$. The point $x_j + \alpha_j p_j$ is a subspace stationary point with respect to $\mathcal{B}$ from (8.35), and remains so when $g_t(x_{j+1})$ is moved to z_N. Moreover, every symmetric subset of the rows and columns of the positive-definite matrix H_B is positive definite, which implies that the matrix obtained by removing the β_r-th row and column of H_B is positive definite.

If temporary bounds are imposed at x_0, and the index ν_s corresponds to a temporary bound, then it is possible that $t = \nu_s$. In this case, the nonbasic set does not change at $x_j + \alpha_j p_j$, but the status of the nonbasic index ν_s is changed from being associated with a temporary bound to being associated with the "real" bound on x_s. As H_B remains the same at $x_j + \alpha p_j$, it follows that x_{j+1} is also a subspace minimizer with respect to $\mathcal{B}$. (A similar scheme is used to handle bound swaps for upper and lower bound constraints of the from $\ell \leq x \leq u$.)

Algorithm 8.4.1. [Inertia-Controlling Method for QP with Bounds.]

Choose x_0 such that $x_0 \geq 0$;

Choose $\mathcal{B}$ and $\mathcal{N}$ such that H_B is positive definite;

Set $x = x_0$; $\quad g = g + H(x - \bar{x})$; $\quad \nu_s = \text{argmin}_{i \in \mathcal{N}} \{g_i\}$;

while $g_s \neq 0$ **do**

  Solve $H_B p_B = -[h_s]_B$; $\quad p_N = e_s$;

  $p = P\begin{pmatrix} p_B \\ p_N \end{pmatrix}$; $\quad q_N = [Hp]_N$;

  $[\bar{\alpha}, t] = \texttt{minRatioTest}(x, p)$; [blocking variable x_t]

if $(q_N)_s > 0$ **then** $\alpha^* = -g_s/(q_N)_s$ **else** $\alpha^* = +\infty$;

$\alpha = \min\{\alpha^*, \bar{\alpha}\}$;

if $\alpha = +\infty$ **then** **stop**; [unbounded solution]

$x \leftarrow x + \alpha p$; $g_N \leftarrow g_N + \alpha q_N$;

if $\bar{\alpha} = \alpha^*$ **then**

$\mathcal{B} \leftarrow \mathcal{B} + \{\nu_s\}$; $\mathcal{N} \leftarrow \mathcal{N} - \{\nu_s\}$;

else if $\bar{\alpha} < \alpha^*$ **then**

if $t = \beta_r$ **then**

$\mathcal{B} \leftarrow \mathcal{B} - \{\beta_r\}$; $\mathcal{N} \leftarrow \mathcal{N} + \{\beta_r\}$;

end

end;

if $g_s = 0$ **then** $\nu_s = \operatorname{argmin}_{i \in \mathcal{N}} \{g_i\}$; [standard subspace minimizer]

$k \leftarrow k + 1$;

end do

8.4.3 ICQP as applied to the pdSQP subproblem

A subspace stationary point is found by solving,

$$[H_M]_B \begin{pmatrix} p_B \\ q \end{pmatrix} = -[g_M]_B.$$

At a standard subspace minimizer of the inertia-controlling QP method identifies an active constraint s with the largest nonoptimal component of z (i.e., x_s is on its bound but $z_s < 0$, or x_s is set on a temporary bound and z_s is nonzero.) This bound is set for deletion from the nonbasic set. At the next step, $(g_M)_B$ must remain zero, while the nonbasic variables remain fixed except for variable s, which moves off of its bound. These requirements may be summarized in the equation,

$$\begin{pmatrix} H_M & E_N^T \\ E_N & 0 \end{pmatrix} \begin{pmatrix} \Delta v \\ -r \end{pmatrix} = \begin{pmatrix} 0 \\ e_s \end{pmatrix},$$

where E_N is a matrix composed of the columns of the identity matrix corresponding to the nonbasic set, Δv is the QP iteration step, and r is the change in the reduced costs z. Expanding H_M, this is equal to

$$\begin{pmatrix} H + \frac{1+\nu}{\mu} J^T J & \nu J^T & E_N^T \\ \nu J & \mu\nu I & 0 \\ E_N & 0 & 0 \end{pmatrix} \begin{pmatrix} p \\ q \\ -r \end{pmatrix} = \begin{pmatrix} 0 \\ e_s \end{pmatrix}.$$

The second equation is $Jp + \mu\nu q = 0$, which implies that the first equation may be written as

$$Hp - J^T q - E_N^T r = 0. \tag{8.37}$$

These equations can be simplified because $E_N \Delta v = e_s$ has the effect of constraining all active variables except for the sth. In particular, they may be written as:

$$[H_M]_B \Delta v_B = -[H_M e_s]_B, \tag{8.38}$$

which is equivalent to a regularized Newton system (also the stabilized SQP equations), as shown in the following result.

Result 8.4.1. *Consider the application of the active-set method to the bound constrained QP* (8.31)*. Then, for every* $\nu \geq 0$*, there exists a positive* $\bar{\mu}$ *such that, for all* $0 < \mu < \bar{\mu}$*, the free components of the QP search direction* (p_j, q_j) *satisfy the nonsingular primal-dual system*

$$\begin{pmatrix} H_B & -J_B^T \\ J_B & \mu I \end{pmatrix} \begin{pmatrix} p_B \\ q_j \end{pmatrix} = -\begin{pmatrix} [h_s]_B \\ a_s \end{pmatrix}. \tag{8.39}$$

Proof. Consider the definition of the search direction when $\nu > 0$. In this case it suffices to show that the linear systems (8.38) and (8.39) are equivalent. For any positive ν, we may define the matrix

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

where the identity matrix I has dimension n_B, the column dimension of ${}_J0_B$. The matrix T_B is nonsingular with $n_B + m$ rows and columns. It follows that the equations

$$T_B [H_M]_B \Delta v_B = -T_B [H_M e_s]_B$$

have the same solution as those of (8.38). The primal-dual equations (8.39) follow by direct multiplication. The nonsingularity of the equations (8.39) follows from the nonsingularity of T_B, and the fact that H_M is nonsingular (as are all symmetric submatrices formed from its rows and columns).

The resulting equations (8.39) are independent of ν, but the simple proof above is not applicable when $\nu = 0$ because T_B is undefined in this case. For $\nu = 0$, the QP objective includes only the primal variables x, which implies that problem (8.31) may be written as

$$\underset{x \geq 0}{\text{minimize}} \;\; (g - J^T\pi)^T (x - x_0) + \tfrac{1}{2}(x - x_0)^T \Big(H + \frac{1}{\mu} J^T J\Big)(x - x_0),$$

with y arbitrary. The active-set equations analogous to (8.34) are then

$$\Big(H_B + \frac{1}{\mu} J_B^T J_B\Big) p_B = -\left[\Big(H + \frac{1}{\mu} J^T J\Big) e_s\right]_B. \tag{8.40}$$

Let the m-vector q be such that

$$q_j = -\frac{1}{\mu} J_B (p_B + e_s). \tag{8.41}$$

Equations (8.40) and (8.41) may be combined to give

$$\begin{pmatrix} H_B & -J_B^T \\ J_B & \mu I \end{pmatrix} \begin{pmatrix} p_B \\ q_j \end{pmatrix} = -\begin{pmatrix} [h_s]_B \\ a_s \end{pmatrix},$$

which are identical to the equations (8.39). □

8.4.4 Properties of the curvature

Consider the curvature $\Delta v^T H_M \Delta v$ for any arbitrary Δv. By definition, we have

$$\begin{aligned} H_M \Delta v &= \begin{pmatrix} H + \frac{1}{\mu}(1+\nu) J^T J & \nu J^T \\ \nu J & \nu\mu I \end{pmatrix} \begin{pmatrix} p \\ q \end{pmatrix} \\ &= \begin{pmatrix} Hp + \frac{1}{\mu} J^T J p + \frac{\nu}{\mu} J^T (Jp + \mu q) \\ \nu (Jp + \mu q) \end{pmatrix}. \end{aligned}$$

By assumption, both (x, y) and $(x+p, y+q)$ are subspace minimizers, which implies that

$$c + J(x + p - x_0) + \mu(y + q - y^E) = Jp + \mu q = 0.$$

Hence

$$\begin{aligned} H_M \Delta v = \begin{pmatrix} Hp + \frac{1}{\mu}J^TJp + \frac{\nu}{\mu}J^T(Jp+\mu q) \\ \nu(Jp+\mu q) \end{pmatrix} &= \begin{pmatrix} Hp + \frac{1}{\mu}J^TJp \\ 0 \end{pmatrix} \\ &= \begin{pmatrix} Hp - J^Tq \\ 0 \end{pmatrix}. \end{aligned} \tag{8.42}$$

This identity implies that the scalar $\Delta v^T H_M \Delta v$, which represents the curvature of φ_M along Δv, may be written as

$$\Delta v^T H_M \Delta v = p^T Hp - p^T J^T q = p^T Hp + \mu\|q\|_2^2. \tag{8.43}$$

It follows that if the objective $\varphi(x)$ for the QP is convex, then H is positive semidefinite and the curvature is positive for all positive μ.

If p is written in terms of its basic and nonbasic components, the identity (8.43) implies that the curvature $\Delta v^T H_M \Delta v$ can be expressed as

$$\begin{aligned} \Delta v^T H_M \Delta v &= p^T Hp + \mu\|q\|_2^2 \\ &= \begin{pmatrix} p_B^T & p_N^T \end{pmatrix} \begin{pmatrix} H_B & H_D \\ H_D^T & H_N \end{pmatrix} \begin{pmatrix} p_B \\ p_N \end{pmatrix} + \mu\|q\|_2^2 \\ &= p_B^T H_B p_B + p_B^T [h_s]_B + [h_s]_B^T p_B + h_{\nu_s,\nu_s} + \mu\|q\|_2^2. \end{aligned} \tag{8.44}$$

The vector p_B is independent of h_{ν_s,ν_s}, which implies that if the curvature is not sufficiently positive, it may be increased by adding a positive quantity θ_s to the ν_s-th diagonal of H. This procedure will result in the final iterate being the solution of a QP with Hessian $H + \Delta H$, where $\Delta H = \operatorname{diag}(\theta_1, \theta_2, \ldots, \theta_n)$ is positive semidefinite.

8.4.5 Calculating the step-size

In the inertia-controlling bound-constrained QP algorithm 8.4.1, the step is scaled to be either to a blocking constraint or an unconstrained one-dimensional minimizer along Δv. If the curvature along the step, $\Delta v^T H_M \Delta v$, is positive, this is given by the expression,

$$\alpha_* = -\big(g_M + H_M(v - v_0)\big)^T \Delta v / \Delta v^T H_M \Delta v,$$

where v is the current primal-dual iterate in the QP subproblem and v_0 is the base point of the QP.

As (x, y) is a subspace minimizer at every iteration, it must hold that $g_M + H_M(v - v_0) = z$, and $z_B = 0$. This implies that the numerator of α_* can be expressed as:

$$(g_M + H_M(v - v_0))^T \Delta v = \Delta v^T z = z_s,$$

because the only nonzero components of z are in $\mathcal{N}$ and x_s is the only nonzero component of x_N.

In addition, as $E_N p = e_s$, it must hold that

$$\Delta v^T H_M \Delta v = \Delta v^T (E_N^T r) = r^T E_N \begin{pmatrix} p \\ q \end{pmatrix} = [Hp - J^T q]_s,$$

where the last equality comes from equation (8.37). These expressions imply that the step length α^* may be written in the form:

$$\alpha^* = -\big(g_M(v)\big)_s / \Delta v^T H_M \Delta v = -z_s / r_s, \tag{8.45}$$

where r denotes the vector $Hp - J^T q$.

8.4.6 Concurrent convexification

As described in Chapter 6, an ideal convexification method makes minimal changes to H. In this section we propose a method that alters H only when negative curvature is encountered during the solution of the QP subproblem. The idea is to monitor the sign of the curvature $\Delta v^T H_M \Delta v$ at each step and add a quantity to H_{ss} that is large enough to make $\Delta v^T H_M \Delta v$ sufficiently positive. This idea can be extended to allow the curvature to be increased whenever a small positive value of the curvature would otherwise cause the QP solver to take a large step.

At any given QP iterate (x, y), the new reduced costs are a function of the step length α, i.e., $z(\alpha) = g + H(x - x_0 + \alpha p) - J^T(y + \alpha q)$. The rate of change of z as a function of α is given by the curvature $\Delta v^T H_M \Delta v$. The situation in which the curvature is positive is depicted in Figure 8.1. The intercept of $z(\alpha)$ is the starting value $z = z(0)$ and the slope is the curvature. In this case, the minimizer α^* of φ as a function of α satisfies $z(\alpha^*) = 0$. In the case of negative curvature, $z(\alpha)$ is unbounded below as $\alpha \to \infty$, and there is no finite α^*, as shown in Figure 8.2.

Let σ denote the scalar $\sigma = (x - x_0)_s$. Adding a positive scalar θ to the s-th diagonal of H has the effect of adding $\sigma\theta$ to the multiplier z_s, and θ to the curvature r_s.

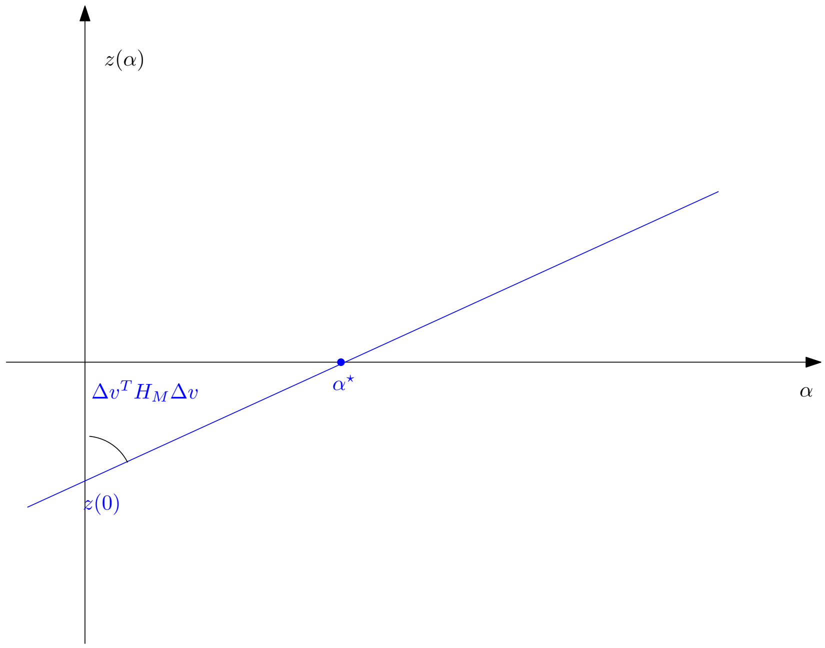

Figure 8.1: If the curvature along Δv is positive, then there is a finite α^* at which $z(\alpha) = 0$.

These modifications redefine the expression (8.45) for the step length to the minimizer along Δv as

$$\alpha(\theta) = -(z_s + \sigma\theta)/(r_s + \theta). \tag{8.46}$$

The derivative of α with respect to θ is given by

$$\alpha'(\theta) = -\frac{1}{(r_s + \theta)^2}\big(\sigma r_s + |z_s|\big).$$

There are two cases to consider for the choice of θ.

Case 1: $\sigma \le 0$**.** In this case, $\alpha(\theta)$ has a pole at $\theta = -r_s$, and decreases monotonically for $\theta > -r_s$. The inequality $\sigma \le 0$ and the assumption that $r_s \le 0$ imply that $\alpha(\theta) \to |\sigma|$ as $\theta \to +\infty$. It follows that θ can always be chosen sufficiently large that $\alpha(\theta)$ is smaller than any given value larger than $|\sigma|$. If $x(\theta)$ is the new iterate $x(\theta) = x + \alpha(\theta)p$, then

$$(x(\theta) - x_0)_s = (x + \alpha(\theta)p - x_0)_s = (x - x_0)_s + \alpha(\theta)p_s = \sigma + \alpha(\theta), \tag{8.47}$$

which implies that $(x - x_0)_s \to 0$ as $\theta \to \infty$.

Case 2: $\sigma > 0$**.** If $\sigma > 0$, then the choice of θ is complicated by the fact that if $\sigma \ge |z_s|/|r_s|$, then $\alpha(\theta)$ is not a decreasing function of θ. Moreover, even if α is decreasing, the amount

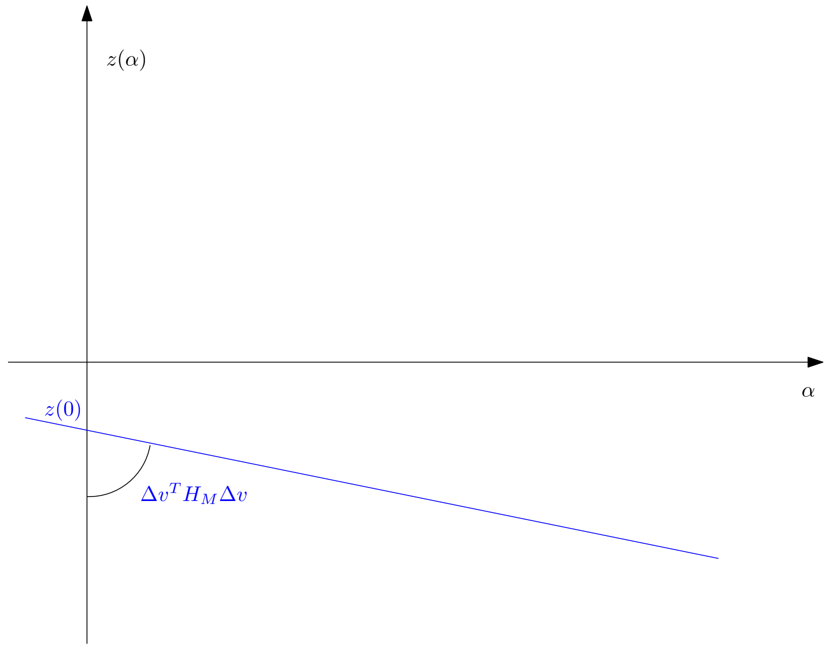

Figure 8.2: When the curvature along Δv is negative, $z(\alpha)$ is unbounded below.

that θ can be increased to reduce $\alpha(\theta)$ is limited by the fact that the modified multiplier $z_s + \sigma\theta$ is a monotonically increasing function of θ that will be positive for any $\theta > |z_s|/\sigma$. This implies that θ can be chosen sufficiently large that the multiplier z_s is positive. In this case, the variable x_s must be "deselected" as the nonbinding nonbasic variable. Gill and Wong [40, 42] show that the curvature r_s is nondecreasing during a sequence of nonstandard iterations with the same nonbinding index ν_s. This result is crucial because it means that a diagonal will be modified only at the *start* of a sequence of non-standard iterations. Therefore, if a nonbinding nonbasic variable x_s is deselected because its multiplier changes sign after the modification, then x_s must be at its bound and may be returned to being a regular nonbasic variable at its current value.

The convexification procedure attempts to modify the Hessian so that the step length taken is a particular desired value, α_c, defined heuristically as a moving average of the norms of the previous steps αp. In particular, the change to the diagonal of H required to make α_c the step length to solve $z(\alpha) = 0$ is the solution to: $\theta_{ss}[(x - x_0 + \alpha_c p)]_s = [-g - H(x - x_0 + \alpha_c p) + J^T(y + \alpha_c q)]_s$. Ideally this modification will modify the curvature to be positive so that $z(\alpha)$ resembles Figure 8.1, with $\alpha^* = \alpha_c$.

There are several limitations associated with this modification to H. First, it may be that $[x - x_0 + \alpha_c p]_s$ is zero, which makes the solution undefined. In this case, α_c is simply doubled.

Another complication is that when H is modified, then $z(0) = g + (H + \Delta H)(x - x_0) - J^T y$ is modified as well. In particular, the sth component can change sign. In particular, as the modification to the Hessian increases, $z(0)$ may change sign before the point at which $z(\alpha)$ crosses the α axis is α_c. In this case, the solution to the modification of H results in the picture depicted in Figure 8.3. In this case, the optimal modification does not result in positive curvature. In this case, the variable s set at its bound becomes optimal and *no* step should be taken along Δv, and a new constraint is found for moving off of (which may be the same s but along the opposite direction if it is a temporary bound).

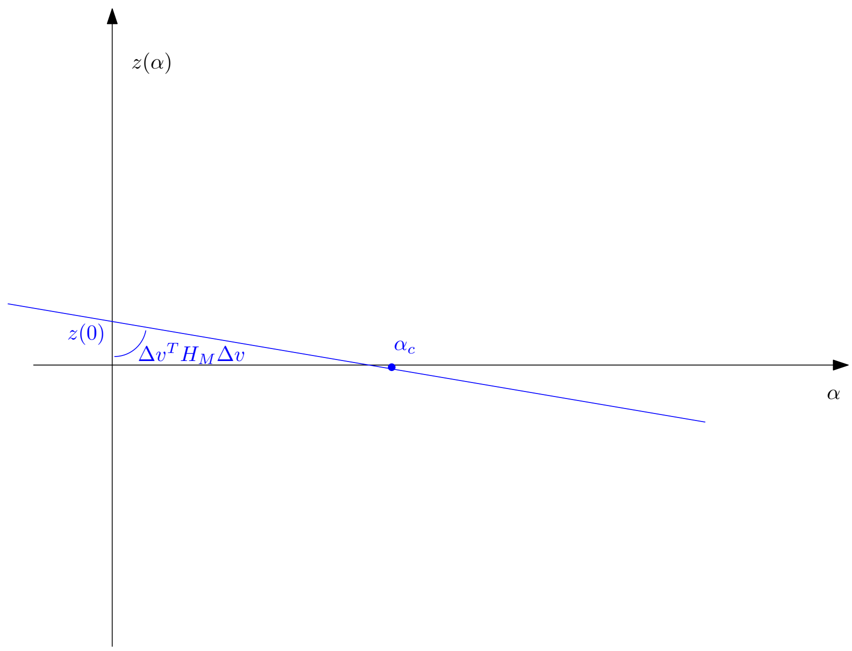

Figure 8.3: The sign of $z(0)$ changes as a result of a positive modification to H at a value less than the curvature.

Finally, it can so happen that the calculated modification in H is *negative*. This solution is depicted in Figure 8.4. The scenario which causes this to be the solution for the modification of H is depicted in Figure 8.5. Here, as the diagonal ΔH_{ss} increases, the value of $z(0)$ decreases faster than the curvature increases, and no positive value solves for $\alpha^* = \alpha_c$.

The convexification procedure monitors the sign of $\Delta H_{ss} = [-g - H(x - x_0 + \alpha_c p) + J^T(y + \alpha_c q)]_s / [(x - x_0 + \alpha_c p)]_s$, and if it is negative, adds a value so as to make the curvature

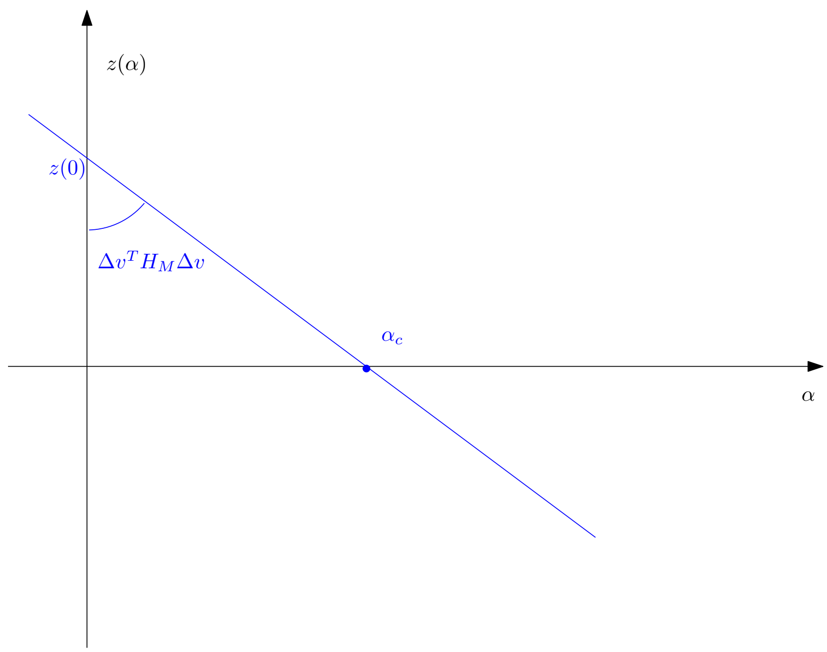

Figure 8.4: The solution for the desired modification in H is *negative*.

positive with the opposite sign of $\Delta v^T H_M \Delta v$. In this case, α^* may be considerably larger than α_c. If there are no blocking constraints, then the step length may be very large. This case is less than ideal, but should generally only occur if the QP is fundamentally unbounded.

8.4.7 An alternative concurrent convexification method

Instead of bounding the inner iteration step-size, there could be a bound on the total size of the step along variable s $|[x - x_0]_s|$. Note that this is similar to, but distinct from a trust-region strategy, which, if an inf-norm trust-region, would already be reflected in the boundary values, in which case convexification would be unnecessary as we "trust" stepping to the boundary along a direction of negative curvature. It leaves open the question as to how to define this bound. However, a method implementing this procedure $\|x - x_0\|_\infty$ is described below.

The definition of θ depends on certain “target values” τ_m and $z_{\min}$ that define the maximum change in x and the smallest positive multiplier. If $\sigma > 0$ then θ is chosen so that the multiplier at x changes from a negative value to a nonnegative value than is no smaller

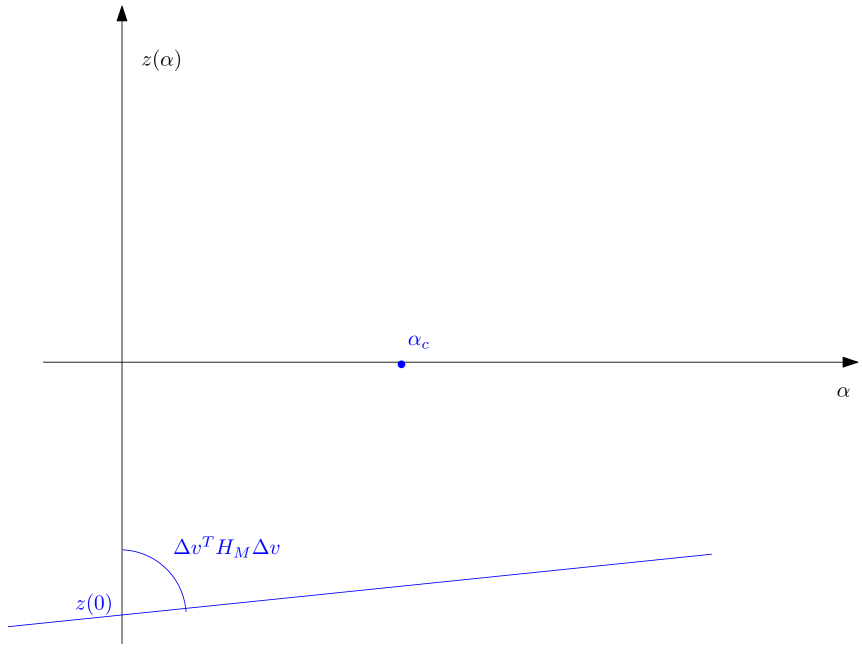

Figure 8.5: As the diagonal of H increases, the value of $z(0)$ decreases more than the value of $\Delta v^T H_M \Delta v$ increases.

than $z_{\min}$, i.e.,

$$z_s + \sigma\theta \geq z_{\min} \geq 0.$$

This inequality is satisfied for every θ such that $\theta \geq \theta_F$, where

$$\theta_F = (z_{\min} + |z_s|)/\sigma.$$

is the value that "flips the sign" of the multiplier.

If $\sigma \leq 0$, the definition of θ is based on choosing θ subject to a limit on the change in the nonbinding nonbasic variable, The expressions (8.46) and (8.47) imply that $(x(\theta) - x_0)_s > 0$ for all $\theta > -r_s$ and $(x(\theta) - x_0)_s \to 0$ as $\theta \to \infty$. Accordingly, we require that θ be chosen large enough that

$$(x(\theta) - x_0)_s \leq \tau_m, \tag{8.48}$$

where τ_m is a positive constant. This condition forces the distance of $x_s(\theta)$ to $[x_0]_s$ to be of the same order as the distance of x_s to $[x_0]_s$. An advantage of this choice of α is that the modification is mainly determined by the behavior of φ along the component of the direction along the "offending" variable x_s.

The restriction (8.48) implies that the value of θ must be larger than the critical value θ_T such that

$$(x(\theta_T) - x_0)_s = (x + \alpha(\theta_T)p - x_0)_s = \sigma + \alpha(\theta_T) = \tau_m.$$

Combining this expression with the definition of α given by (8.46) yields

$$\alpha(\theta) \le \alpha_{\max} = \tau_m + |\sigma| \quad \text{for all} \quad \theta \ge \theta_T = \big(|z_s| + \alpha_{\max}|r_s|\big)/\tau_m.$$

The value of θ must be chosen such that $\theta > \theta_L$, where $\theta_L = |r_s|$ is the smallest perturbation that makes the curvature nonnegative. Observe that

$$\theta_T = |r_s| + \big(|z_s| + |\sigma|\big)/\tau_m > \theta_L,$$

as required. In summary, the definition of θ is given by

$$\theta = \begin{cases} \theta_T & \text{if } \theta_T \le \theta_F, \\ \max\{\theta_L,\, \theta_F\}, & \text{otherwise.} \end{cases}$$

One clear advantage of this method is that the scenario depicted in figure 8.5 is not possible, and a modification that bounds the total deviation of variable s from the value at x_0 can always be found. The drawback is more subtle. This method requires the value τ_m, a de facto trust-region on the variable. Without a proper trust-region framework, the drawbacks will be inherent in whatever arbitrary heuristics are put in place. For instance, by making $\tau_m = \|x - x_0\|_\infty$, the algorithm potentially penalizes large step lengths early among the iterations, when they are expected to be *larger* since the larger magnitude $[z]_s$ are moved off of first, and these correspond to the directional derivative of the objective. Of course the previous convexification procedure suffers from an arbitrary heuristic as well, of using a moving average of inner iteration steps, but with a large collection of such steps accumulated, it can be expected that this is a relatively safe and self-correcting procedure.

8.5 Global Convergence

This section contains several results concerning the global convergence of the primal-dual iterates of pdSQP. For the primal iterates, it is shown that if the constant positive generator constraint qualification holds, then a sequence of S- and L-iterates converges to a first-order KKT point. (The constant positive generator (CPG) constraint qualification is the weakest constraint qualification that ensures that approximate KKT sequences converge to KKT points; see [4] and Chapter 2, Page 16). It is also shown that for equality-constrained problems, an infinite sequence of M-iterates converges to either a KKT point or a point failing to satisfy the quasinormality constraint qualification.

8.5.1 Convergence of S- and L- iterates

The main result of this section is that a subsequence of S- and L-iterates form an *approximate KKT sequence.* An approximate KKT sequence $\{x_k\}$, defined in Chapter 4 (Page 48) is a sequence $\{x_k\}$ such that the following conditions hold

$$g(x_k) - \sum y_j \nabla c_j(x_k) - z_k = \epsilon_k, \tag{8.49}$$

$$x_k \geq -\delta_k, \tag{8.50}$$

$$z \geq 0, \tag{8.51}$$

$$z^T(x_k - \delta_k) = 0, \tag{8.52}$$

$$\|c_j(x_k)\| \leq \nu_k, \tag{8.53}$$

with $\{\epsilon_k, \delta_k, \nu_k\} \to 0$.

Andreani et al. [4, Theorem 5.7] show that if: (i) $x_k \to x^*$; (ii) x_k is an approximate KKT sequence; and (iii) CPG holds at x^*, then x^* is a first-order KKT point.

Result 8.5.1. *An infinite subsequence $\{v_k\}$ of S- and L-iterates is an approximate KKT sequence.*

Proof. Let $\gamma_k = \max(\phi_S^{max}, \phi_L^{max})$. As $\gamma_k \to 0$, the definition of ϕ_S and ϕ_L implies that the constraint inequality (8.53) holds. In addition, the stationarity condition implies $\|\min(x_k, g(x_k) - J(x_k)^T y_k)\| \leq \gamma_k$.

For components i satisfying $\min([x_k]_i, [g(x_k) - J(x_k)^T y]_i) = [x_k]_i$, we define $[z_k]_i = ([g(x_k) - J(x_k)^T y_k]_i)_+$ and $[\delta_k]_i = [x_k]_i$. For the components i such that $\min((x_k)_i, (g(x_k) - J(x_k)^T y_k)_i) = (g(x_k) - J(x_k)^T y_k)_i$, we define $[z_k]_i = 0$ and $[\delta_k]_i = 0$. The resulting vector z_k satisfies (8.51), and x_k and δ_k satisfy (8.50) by construction.

For the indices i such that $\min(x_i, [g(x_k) - J(x_k)^T y_k]_i) = [x_k]_i$, it must hold that $[x_k]_i = [\delta_k]_i \leq \gamma_k$, which implies that (8.52) holds for the ith component. The condition $[g_k - J_k^T y_k]_i - [z_k]_i = 0$ implies that the ith component of (8.49) holds with $[\epsilon_k]_i = 0$.

In the second case, $[z_k]_i = 0$, so equation (8.52) holds and $[\epsilon_k]_i = [g_k - J_k^T y_k]_i \leq \gamma_k$, and (8.49) must hold also.

This proves that v_k is an approximate KKT sequence with $0 \leq \{\epsilon_k, \delta_k, \nu_k\} \leq \gamma_k$ and $\gamma_k \to 0$. □

8.5.2 Convergence of the M-iterates for equality constraints

Assuming very weak regularity conditions, a global convergence result can be derived for the limit sequences of M-iterates when there are only equality constraints (i.e., there are no bounds on the variables).

Let $\{v_k\}$ be a sequence of M-iterates, for which $\lim_{k\to\infty} \mu_k^R = 0$ and $\lim_{k\to\infty} \tau_k = 0$. Note that, by assumption, $\{x_k\}$ lies in a compact set, and hence $\{x_k\}$ has a limit point x^*. Without loss of generality, let k denote the indices of the corresponding subsequence.

It is assumed that the *quasinormality* constraint qualification holds at x^*. The quasinormality constraint qualification, first defined in Chapter 2 (Page 16), implies the following condition:

Definition 8.5.1. *The* quasinormality *condition holds at x^* if there is no $\{\lambda_i\}$ such that*

1. $\sum \lambda_i \nabla c_i(x^*) = 0$,
2. *the λ_i are not all zero,*
3. *For every neighborhood $\mathcal{N}$ of x^* there is an $x \in \mathcal{N}$ such that $\lambda_i c(x)_i > 0$ for all i with $\lambda_i \neq 0$.*

□

The gradient of M^ν is

$$\nabla M^\nu = \begin{pmatrix} g_k - J_k^T\left((1+\nu)\left(y_k^E - \dfrac{1}{\mu^R}c_k\right) - \nu y_k\right) \\ \nu(c_k + \mu_R(y_k - y_k^E)) \end{pmatrix} \tag{8.54}$$

As the definition of an M-iterate implies that $\|\nabla M(v_k)\| \leq \tau_k$, we define $\epsilon_k = \nabla M(v_k)$, where $\|\epsilon_k\| \leq \tau_k$. Let $\epsilon_k^x = \nabla_x M(v_k)$ and $\epsilon_k^y = \nabla_y M(v_k)$.

Proposed change to Algorithm 8.2.1: For the test that defines an M-iterate, the term $\nabla M^\nu(v_{k+1}; y_k^E, \mu_k^R)$ is replaced by $\nabla M^\nu(v_{k+1}; y_{k+1}^E, \mu_k^R)$, where y_{k+1}^E is defined as the new multiplier estimate y_{k+1}. This modification is necessary for the first result below to hold. It does not change any of the other convergence results.

Theorem 8.5.1. *If the sequence $\{y_k\}$ is bounded, then every cluster point of the sequence of M-iterates is a first-order KKT point.*

Proof. As k is an M-iterate and $\{y_k\}$ is bounded, it follows that $y_k = y_k^E$ in the definition of ∇M. This implies that $c_k \to 0$ because the second component of the gradient of M converges to zero,

The stationarity condition implies that

$$\lim_{k\to\infty} \Big(g_k - J_k^T y_k + \frac{1+\nu}{\mu_k} J_k^T c_k \Big) = 0.$$

From the definition of the second component of ∇M, it holds that that $c_k = \epsilon_k^y/\nu$, where $\epsilon_k^y \to 0$. Furthermore, since η_k, the optimality residual, is bounded from below by assumption, it holds that, for some index K, the update for μ is $\mu_{k+1} = \frac{1}{2}\mu_k$ for all $k \geq K$, where μ_k is updated if and only if τ_k is updated. This implies that $\epsilon_k^y/\mu_k \leq \tau_k/\mu_k = \tau_K/\mu_K$ is the member of a bounded sequence and must have a cluster point ϵ^*.

If c_k is written as ϵ_k^y/ν, in the stationarity condition, it holds that

$$\lim_{k\to\infty} \Big(g_k - J_k^T y_k + \frac{1+\nu}{\nu\mu_k} J_k^T \epsilon_k^y \Big) = 0,$$

which implies that, at any cluster point (x^*, y^*),

$$g_* - J_*^T\Big(y^* + \frac{1+\nu}{\nu}\epsilon^*\Big) = 0,$$

as required. ◻

Theorem 8.5.2. *If the sequence $\{y_k\}$ is not bounded, then the sequence of M-iterates has at least one cluster point that is either a first-order KKT point, or a point failing to satisfy the quasinormality constraint qualification.*

Proof. The second component of ∇M^ν may be rearranged to give y_k in the form

$$y_k = \frac{1}{\mu_k^R}\Big(\frac{1}{\nu}\epsilon_k^y - c(x_k)\Big) + y_k^E.$$

Similarly, the first component of ∇M^ν may be written as

$$\begin{aligned} g_k - J_k^T\left((1+\nu)\Big(y_k^E - \frac{1}{\mu_k^R}c_k\Big) - \frac{\nu}{\mu_k^R}\Big(\frac{1}{\nu}\epsilon_k^y - c_k\Big) - \nu y_k^E\right) &= g_k - J_k^T\Big(- \frac{1}{\mu_k^R}c_k + y_k^E - \frac{1}{\mu_k^R}\epsilon_k^y\Big) \\ &= \epsilon_k^x. \end{aligned}$$

This implies that

$$\lim_{k\to\infty} \Big(\mu_k g_k - \mu_k J_k^T y_k^E + J_k^T c_k + J_k^T \epsilon_k^y\Big) = 0.$$

As the sequences $\{J_k^T y_k^E\}$ and $\{g_k\}$ are bounded, with $\mu_k \to 0$ and $\epsilon_k^y \to 0$, the sequence $J_k^T c_k$ must converge to zero. Also, as the sequence $\{x_k\}$ is bounded, if the sequence of iterates has a cluster point x^*, it must hold that either $c(x^*) = 0$ or $c(x^*)$ is a nontrivial member of the null space of $J(x^*)$. The latter case implies that the quasinormality condition fails at x^* with $\lambda = c(x^*)$.

If $c(x^*) = 0$, then by the same argument as in Theorem 8.5.1, the sequence $\{\epsilon_k^y/\mu_k\}$ is bounded and hence has a cluster point. Since x_k and y_k^E are also bounded, by the second component of ∇M, there is a subsequence $\{k_l\}$ such that $y_{k_l} \to y^*$. The stationarity KKT condition is satisfied by the multiplier $\bar{y} = y_*^E - \epsilon^*$, where ϵ^* is a cluster point of ϵ_k^y/μ_k. $\square$

8.5.3 Convergence of the M-iterates for equalities and bounds

In the case of the general problem (8.1) with equality and bound constraints, only, weak statements about infinite sequences of M-iterates can be made. Let x_k be a sequence of M-iterates with cluster point x^*. Consider the expression,

$$\begin{aligned} 0 = \lim_{k\to\infty} \|\min(x_k, \nabla_x M_k)\| &= \lim_{k\to\infty} \|\left[(x_k - \nabla_x M_k) - x_k\right]_+\| \\ &\geq \lim_{k\to\infty} \|\left[(x_k - (\mu_k g_k - \mu_k J_k^T y_k^E + J_k^T c_k + J_k^T \epsilon_k^y)) - x_k\right]_+\| \\ &= \|\left[(x^* - J(x^*)^T c(x^*)) - x^*\right]_+\|. \end{aligned}$$

This implies that x^* is a stationary point of the problem

$$\underset{x}{\text{minimize}} \;\; \|c(x)\|^2 \quad \text{subject to} \;\; x \geq 0.$$

This is a standard result for augmented Lagrangian methods (see, e.g., Conn et al. [19]).

8.6 Local Convergence

As noted in Section 3 of Gill and Robinson [39], the subproblem in the algorithm is of the same form as the stabilized SQP subproblem. However, several other properties have to hold for the two subproblems to be fully identical. Since stabilized SQP is superlinearly convergent for degenerate problems, the two subproblems being identical would imply that pdSQP is superlinearly convergent as well.

The following assumption is made throughout this section:

Assumption 8.6.1. *There exists a subsequence of the S- and L- iterates that converges to a first-order KKT point x^*.*

In the global convergence results of Gill and Robinson [39, Theorem 4.2], the sequence of iterates generates either a subsequence of S- and L-iterates or converges to a stationary point of the primal-dual augmented Lagrangian. In view of the previous section's discussion of global convergence results, this assumption concerns the case where the algorithm does not converge to a non-optimal stationary point of $\|c\|^2$.

The following second-order sufficiency condition was defined in Chapter 2, Page 26.

Definition 8.6.1. *The* second order sufficiency condition (SOSC) *holds at a first-order KKT point* (x^*, y^*, z^*) *if* $p^T H(x^*, y^*)p > \sigma\|p\|^2$ *for all* p *such that* $J(x^*)p = 0$, $p_i = 0$ *for* $z_i^* > 0$, *and* $p_i \geq 0$ *for* $z_i^* = 0$.

The results of this section require that the SOSC and the MFCQ hold.

Assumption 8.6.2. *The second-order sufficiency conditions and the MFCQ hold at all first-order KKT limit points* (x^*, y^*, z^*) *of the iterates.*

To begin with, it will be shown that eventually, an unconstrained step length is taken, i.e., $\alpha_k = 1$ for $k \geq K$ for some K. In particular, the primal-dual augmented Lagrangian merit function does not suffer from the Maratos effect.

8.6.1 The unconstrained step

Assume that for all sufficiently large $k \in \mathcal{S}$, and x_k is sufficiently close to x_* to imply that,

$$x_k + p_k - x^* = o(\|x_k - x^*\|), \tag{8.55}$$

$$y_k + q_k - y^* = o(\|y_k - y^*\|). \tag{8.56}$$

For the rest of this section, suppress the index k. Equations (8.55) and (8.56) allow $f(x+p)$ and $c(x+p)$ to be expressed as,

$$f(x+p) = f(x) + \frac{1}{2}(g(x) + g(x^*))^T p + o(\|p\|^2),$$

$$c(x+p) = c(x) + \frac{1}{2}(J(x) - J(x^*))^T p + o(\|p\|^2).$$

Alternatively, $f(x+p)$ and $c(x+p)$ can also be written as

$$f(x+p) = f(x) + g(x)^T p + o(\|p\|^2).$$

$$c(x+p) = c(x) + J(x)p + o(\|p\|^2)$$

Consider the expression for $M(v + \Delta v)$,

$$\begin{aligned} M(v+\Delta v) = f &+ \tfrac{1}{2}(g+g^*)^T p - (c^T + \tfrac{1}{2}(J+J^*)p)^T y^E + \frac{1}{2\mu}(c+Jp)^T(c+Jp) \\ &+ \frac{\nu}{2\mu}(c + Jp + \mu(y+q-y^E))^T(c+Jp+\mu(y+q-y^E)) + o(\|p\|^2). \end{aligned}$$

Using $g^* - J^{*T}y = o(\|p\|^2)$ and subtracting $M(v)$, this expression becomes

$$\begin{aligned} M(v+\Delta v) - M(v) = \tfrac{1}{2}g^Tp - \tfrac{1}{2}p^TJ^Ty^E + \frac{1}{\mu}p^TJ^Tc \\ + \frac{\nu}{2\mu}(2c^TJp + 2\mu c^Tq + 2\mu p^TJ^T(y+q-y^E) \\ + 2\mu^2q^T(y+q-y^E)) + o(\|p\|^2). \end{aligned}$$

On the other hand,

$$\nabla M^T\Delta v = g^Tp - ((1+\nu)(y^E - \frac{1}{\mu}c) - \nu y)^TJp + \nu c^Tq + \nu\mu(y-y^E)^Tq.$$

This implies that if ΔM is defined as $\Delta M = M(v+\Delta v) - M(v)$, then

$$\begin{aligned} \Delta M - \eta_S\nabla M^T\Delta v = (\tfrac{1}{2} - \eta_S)\nabla M^T\Delta v - \tfrac{1}{2}\left(g^Tp - ((1+\nu)(y^E - \tfrac{1}{\mu}c) - \nu y)^TJp\right) \\ - \tfrac{1}{2}\left(\nu c^Tq + \nu\mu(y-y^E)^Tq\right) + \tfrac{1}{2}g^Tp - \tfrac{1}{2}p^TJ^Ty^E + \frac{1}{\mu}p^TJ^Tc \\ + \frac{\nu}{2\mu}(2c^TJp + 2\mu c^Tq + 2\mu p^TJ^T(y+q-y^E) \\ + 2\mu^2q^T(y+q-y^E)) + o(\|p\|^2). \end{aligned} \tag{8.57}$$

This expansion is used in the proof of the following result.

Theorem 8.6.1. *For some K, it holds that $\alpha_k = 1$ for all $k \geq K$.*

Proof. Consider k to be an S-, L- or M-iterate. This implies that $y^E = y_k$. The expression (8.57) becomes,

$$\begin{aligned} &(\tfrac{1}{2} - \eta_S)\nabla M^T\Delta v - \frac{1}{2}\left(g^Tp - (y - (1+\nu)\frac{1}{\mu}c)^TJp + \nu c^Tq\right) + \tfrac{1}{2}g^Tp - \tfrac{1}{2}p^TJ^Ty + \frac{1}{\mu}p^TJ^Tc \\ &\qquad + \frac{\nu}{2\mu}(2c^TJp + 2\mu c^Tq + 2\mu p^TJ^Tq + 2\mu^2q^Tq) + o(\|p\|^2) \\ &= (\tfrac{1}{2} - \eta_S)\nabla M^T\Delta v - -\frac{1+\nu}{2\mu}c^TJp - \tfrac{1}{2}\nu c^Tq + \frac{1}{\mu}p^TJ^Tc \\ &\qquad + \frac{\nu}{2\mu}(2c^TJp + 2\mu c^Tq + 2\mu p^TJ^Tq + 2\mu^2q^Tq) + o(\|p\|^2). \end{aligned}$$

From the optimality conditions of the sSQP subproblem, $Jp = -\mu q - c$. Substituting this into the above expression, it changes to

$$
\begin{aligned}
&(\frac{1}{2}-\eta_S)\nabla M^T\Delta v+\tfrac{1}{2}(1+\nu)c^Tq+\tfrac{1}{2}(1+\nu)\frac{c^Tc}{\mu}-\tfrac{1}{2}\nu c^Tq-c^Tq-\frac{1}{\mu}c^Tc-\nu c^Tq \\
&\qquad -\frac{\nu}{\mu}c^Tc+\nu c^Tq-\nu\mu q^Tq-\nu c^Tq+\nu\mu q^Tq+o(\|p\|^2) \\
&=(\tfrac{1}{2}-\eta_S)\nabla M^T\Delta v+c^Tq(\tfrac{1}{2}(1+\nu)-\tfrac{1}{2}\nu-1-\nu+\nu-\nu) \\
&\qquad +c^Tc\Big(\tfrac{1}{2}(1+\nu)-\frac{1}{\mu}-\frac{\nu}{\mu}\Big)+q^Tq(-\nu\mu+\nu\mu)+o(\|p\|^2) \\
&=(\tfrac{1}{2}-\eta_S)\nabla M^T\Delta v+\frac{1}{\mu}c^T\Big(-\tfrac{1}{2}\mu q+(1+\nu)(\tfrac{1}{2}\mu-1)c\Big)+o(\|p\|^2) \\
&=(\tfrac{1}{2}-\eta_S)\nabla M^T\Delta v-\tfrac{1}{2}c^Tq+(\tfrac{1}{2}(1+\nu)-\frac{1}{\mu})c^Tc+o(\|p\|^2).
\end{aligned}
$$

As μ becomes small, the multiple of c^Tc term is eventually negative. Since $q=o(\delta(x,y))$ by (8.56) and $c=O(\delta(x,y))$, the q^Tc term is $o(\|p\|^2)$. On the other hand $\nabla M^T\Delta v=-O(\|\Delta v\|^2)$. This implies that the entire expression eventually becomes negative. □

8.6.2 Equivalence of subproblems

To show that the superlinear convergence results of stabilized SQP are applicable for pdSQP, it must be shown that, asymptotically, the solution to the convex problem is a solution to the stabilized SQP problem. The two specific differences between the two subproblems arise in the definition of the parameter μ^R and the modified Hessian. It was shown in Section 8.3.1 that a stationary point of the convex problem is a stationary point of the unconvexified subproblem, so whether or not a local minimizer of one is also a local minimizer for the other depends on whether the reduced costs change in sign. It will be seen that if strict complementarity holds for the reduced costs, then asymptotically they do not change sign and the two subproblems produce identical local minimizers. In a subsequent section, a modification is proposed which weakens the assumption of strict complementarity.

The following result, due to Hager [51], is a generalization of Debreu's Lemma that is useful for the local convergence theory below.

Theorem 8.6.2 (Hager [51, Lemma 3])**.** *Let Q_* be a symmetric matrix. Suppose that $w^TQ_*w\ge\alpha\|w\|^2$ whenever $B_*w=0$. Then, given any $\delta>0$, there exists a $\sigma>0$ and neighborhoods $\mathcal{B}$ or B_* and $\mathcal{Q}$ of Q_* such that*

$$
v^T\Big(Q+\frac{1}{\rho}B^TB\Big)v\ge(\alpha-\delta)\|v\|^2,
$$

for all $v\in\mathbb{R}^n$, $0<\rho\le\sigma$, $B\in\mathcal{B}$, and $Q\in\mathcal{Q}$. □

Local convergence results for the original pdSQP

It will be seen that the following assumption of strict complementarity is both sufficient (along with the MFCQ and SOSC) and necessary to establish the full asymptotic equivalence of the pdSQP and sSQP subproblems.

Assumption 8.6.3. *There exists a z^* such that $z^* > 0$.* $\square$

Consider the convexification procedure that defines the positive-definite matrix $\bar{H} = H + \frac{1}{\mu_H} P_{\scriptscriptstyle A} P_{\scriptscriptstyle A}^T$ based on the index set $\mathcal{A}(x)$ such that if $i \in \mathcal{A}(x)$, then $x_i \le \mu$ (see Chapter 6, Page refc:convex).

Theorem 8.6.3. *There is an iteration index K for which $k \ge K$ implies that the pdSQP subproblem solution satisfies the optimality conditions of the stabilized SQP subproblem. In particular, $\mu_k^{\scriptscriptstyle R} = \eta_k$ and the local minimizer for the pdSQP subproblem is a local minimizer for the stabilized SQP subproblem.*

Proof. By construction $\mu_k \le \eta_k$, and since $\eta_k \to 0$, by Wright [89, Theorem 3.3], eventually $\mathcal{A}(x_k) = \mathcal{A}^*$, where $\mathcal{A}^*$ is the active set at x^*. Construct the convexification of H_k using the elements of $\mathcal{A}(x_k)$. By the SOSC and Hager [51, Lemma 3], there exists a $\mu_{\scriptscriptstyle H}$ such that, with $\bar{H}_k$ defined to be $\bar{H}_k = H_k + \frac{1}{\mu_H} P_{\scriptscriptstyle A} P_{\scriptscriptstyle A}^T$, the matrix $\bar{H} + \frac{1}{\mu_k^R} J_k^T J_k$ is positive-definite.

Since $\mu_k \le \eta_k$, the optimality residual in the sSQP literature, it holds that $\mu_k \le \eta_k = O(\|x_k - x^*\|)$. By Wright [87, Lemma 4.1], since stabilized SQP is an inexact SQP subproblem, there is a solution $(x_{k+1}, y_{k+1}, z_{k+1})$ to the non-convexified sSQP subproblem (closest to (x_k, y_k, z_k) among all the solutions) such that $\mathcal{A}_+(z^*)$ is a subset of the active set of $(x_{k+1}, y_{k+1}, z_{k+1})$ for some z^*. By strict complementarity, z^* is bounded away from zero. This implies that, if k is large such that z_k is sufficiently close to z^*, the optimal reduced cost z_{k+1} satisfies $z_{k+1} > 0$.

For the remainder of the proof, the iteration subscript k is omitted. The optimality conditions for the nonconvexified sSQP subproblem are:

$$\begin{aligned} g + Hp &= J^T(y+q) + z, \\ z^T(x+p) &= 0, \\ c + Jp + \mu q &= 0, \\ z &\ge 0. \end{aligned}$$

Consider $\mathcal{A}$ to be the active set at the start of the current iteration, e.g. the one to be used for the projection matrix in the convexification. It holds that,

$$g + Hp + \frac{1}{\mu_H} P_A P_A^T p = J^T y + z + \frac{1}{\mu_H} P_A P_A^T p.$$

Since $\frac{1}{\mu_H}$ is bounded from above, and p approaches zero, eventually $z + \frac{1}{\mu_H} P_A P_A^T p \geq 0$.

Defining $\bar{z} = z + \frac{1}{\mu_H} P_A P_A^T p$, $(p, q, \bar{z})$ is a solution to the convexified sSQP. Since a solution of a convex QP is unique, this is the solution that the QP solver will generate. For large enough k in the subsequence of S- and L- iterates, since $\alpha = 1$, the subsequence satisfies the conditions of Wright [87, Theorem 5.3], and

$$\|(x_{k+1} - x^*, y_{k+1} - y^*)\| \leq \|q\| O(\|x_k - x^*\|) + O(\|x_k - x^*\|^2) + \mu O(\|q\|).$$

Since $\eta = O(\|x - x^*\|)$, $\mu \leq \eta$ and, for the subsequence of S- and L-iterates, $q \to 0$, $x_k \to x^*$, it holds that eventually, for some K, $\eta_{k+1} \leq \frac{1}{2}\eta_k$ for $k \geq K$ and all such iterates are S- and L-iterates and $y^E = y$ (see Section 8.2.4, Page 108). Likewise, eventually $\eta_{k+1} \leq (\frac{1}{2} - \delta)\eta_k$ for a small $\delta > 0$ and by the the definition of μ (see (8.18), Page 109), eventually $\mu_k = \eta_k$, and the step is equivalent to an sSQP step.

Since the SOSC and the MFCQ are assumed to hold, superlinear convergence follows from the superlinear convergence of the sSQP subproblem. □

Example: The example in this section illustrates the necessity of sufficient strict complementarity for the equivalence of the sSQP and the convexified sSQP subproblem solutions in the formulation of sSQP in Gill and Robinson [39], as shown in Theorem 8.6.3. Consider the problem:

$$\begin{aligned} \underset{x}{\text{minimize}} \;\; & x_2 - x_1^2, \\ \text{subject to} \;\; & x_2 - 2x_1^2 \geq 0, \\ & x_2 + 2x_2^1 \geq 0. \end{aligned} \tag{8.58}$$

At the minimizer $x = (0, 0)$, it must hold that $y_1 + y_2 = 1$. SOSC holds for $y_2 > \frac{1}{2} + y_1$.

Consider the subproblem with the initial point $x = (0, 1)$, $s = (1, 1)$, $\mu = 1$. First, it is assumed that $y = (0, 1)$. If the problem is reformulated with slack variables, the Hessian is given by

$$H = \begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

The optimality conditions for the indefinite QP subproblem are:

$$\begin{aligned} 1 &= q_1 + 1 + q_2, \\ q_1 &= z_1, \\ 1 + q_2 &= z_2, \\ p_2 - p_{s1} + q_1 &= 0, \\ p_2 - p_{s2} + q_2 &= 0. \end{aligned}$$

It is clear that $p = (0, -1, -1, -1)$ and $q = (0, 0)$ solves the problem with $z = (0, 1)$.

For the convexified subproblem, with the convexification parameter as $\mu_H = .1$, it holds that

$$\begin{aligned} 1 &= q_1 + 1 + q_2, \\ q_1 + .1p_{s1} &= z_1, \\ 1 + q_2 + .1p_{s2} &= z_2, \\ p_2 - p_{s1} + q_1 &= 0, \\ p_2 - p_{s2} + q_2 &= 0. \end{aligned}$$

at which, if $p = (0, -1, -1, -1)$ and $q = (0, 0)$, it must hold that $z = (-.1, .9)$, which is not optimal. However, if $y = (.15, .85)$, then the optimality conditions become:

$$\begin{aligned} 1 &= q_1 + 1 + q_2, \\ .15 + q_1 &= z_1, \\ .85 + q_2 &= z_2, \\ p_2 - p_{s1} + q_1 &= 0, \\ p_2 - p_{s2} + q_2 &= 0, \end{aligned}$$

which hold for $p = (0, -1, -1, -1)$ and $q = (0, 0)$ with $z = (.15, .85)$.

For the convexified subproblem:

$$\begin{aligned} 1 &= q_1 + 1 + q_2, \\ .15 + q_1 + .1p_{s1} &= z_1, \\ .85 + q_2 + .1p_{s2} &= z_2, \\ p_2 - p_{s1} + q_1 &= 0, \\ p_2 - p_{s2} + q_2 &= 0. \end{aligned}$$

In this case, $p = (0, -1, 1, 1)$ and $q = (0, 0)$ is a primal-dual solution with $z = (.05, .75)$.

8.6.3 Obtaining superlinear convergence under weaker assumptions

As outlined in the last section, the assumption of strict complementarity for the reduced costs is essential for a proof of superlinear convergence of Algorithm 8.2.1. As will be observed in the numerical results, strict complementarity is a strong assumption, i.e., it fails for a large quantity of problems. From the example it can be seen that the problem is particularly acute if there is a step onto the active constraints after predicting that they should be active. This section shows that by identifying the variable indices that fail to satisfy strict complementarity and selectively convexifying, it is possible to prove superlinear convergence assuming the MFCQ and the SOSC only.

The following procedure is defined as algorithm **pdSQPid0**:

1. If $x_i \leq \mu_R$ then put $i \in \mathcal{A}(x)$. Set all x_i on their bounds for $i \in \mathcal{A}(x)$.
2. Apply Algorithm **ID0** to z to identify the weakly and strongly active bounds (see Wright [88])
3. Solve for $\widehat{z}$, the interior multiplier estimate.
4. Project x onto the bounds in $\mathcal{A}$.
5. Check if, after convexifying on only the strongly active constraints, $\bar{H}$ has the correct inertia. If it does, proceed to the next step. If it does not, discard the interior multiplier estimate and solve the QP subproblem as originally defined.
6. Solve the QP subproblem. As x changes with each step, change the value of $\widehat{z}$ by the appropriate amount ($H\Delta x$). If $\widehat{z}_j < 0$ for some $j \in \mathcal{A}_+$, then, and only then, step off of the constraint j and add the appropriate $(1/\mu_H)P_jP_j^T$ to H.

Equivalence to a stabilized SQP method

For this section, let ϵ_z be defined as

$$\epsilon_z = \max_{z \in \mathcal{M}_y(x^*)} \ \min_{i \in \mathcal{A}_+} z_i^*.$$

Theorem 8.6.4. *If $\delta(x^k, y_k)$ is sufficiently small, the procedure described above is equivalent to a sequence of iterations generated by solving the following stabilized SQP subproblem:*

$$\begin{array}{ll} \underset{x,y}{\text{minimize}} & g_k^T(x - x_k) + \frac{1}{2}(x - x_k)^T H(x_k, y_k)(x - x_k) + \frac{1}{2}\mu_k^R \|y\|^2 \\ \text{subject to} & c_k + J_k(x - x_k) + \mu_k^R(y - y_k^E) = 0, \quad x \geq 0, \quad [x]_{\mathcal{A}_+} = 0. \end{array} \tag{8.59}$$

Proof. Wright [88, Theorem 4] has shown that for δ sufficiently small, the procedure **ID0** identifies the strongly and weakly active constraints correctly. Moreover, Wright [88, Theorem 5] shows that the interior estimate for the reduced costs satisfies $[\widehat{z}]_{i\in\mathcal{A}_+} \geq \epsilon_z$.

By SOSC and Hager [51, Lemma 3] (see Theorem 8.6.2 above), $H + \frac{1}{\mu^R} J^T J + \frac{1}{\mu_H} P_A P_A^T$ is positive definite for some μ_H on the cone $\{d \mid [d]_{A_0} \geq 0\}$.

It will be shown that this implies that the subproblem is convex. As the subproblem with the exact Hessian is equivalent to stabilized SQP subject to bounds, the subproblem satisfies the inexact SQP framework. By Wright [87, Lemma 5.1] the solution satisfies the estimate:

$$\|p\| + \delta(y_k) + \delta(z_{k+1}) = O(\delta(x_k)) + O(\|(t, r)\|),$$

with $t = 0$ and $r = -\mu^R q$. This implies that it is possible to take $\delta(x_k, y_k)$ to be sufficiently small as to make p sufficiently small such that, since the sequence $\{H_k\}$ is bounded, Hp must satisfy $\|Hp\|_\infty \leq \frac{1}{2}\epsilon_z$.

This implies that the strongly active components of $\widehat{z}$ never become negative. Since $[x]_{i\in\mathcal{A}_+}$ does not change for the pdSQPid0 subproblem iteration and $[x]_{i\in\mathcal{A}_0}$ is initialized at the bounds, the step p would satisfy $[p]_{i\in\mathcal{A}_0} \geq 0$ and no step in the minor QP iterations would be a step of negative curvature. Therefore H will not be modified and the steps are identical to those for which H is not convexified. This implies that the subproblem has the same solution as the sSQP subject to the additional constraint of $[x]_{i\in\mathcal{A}_+} = 0$.

□

The sSQP subproblem

Next we focus on the sSQP subproblem for which the variables associated with the strongly active bounds are set to equality, i.e.,

$$\begin{array}{ll} \underset{x,y}{\text{minimize}} & g_k^T(x - x_k) + \frac{1}{2}(x - x_k)^T \bar{H}(x_k, y_k)(x - x_k) + \frac{1}{2}\mu_k^R \|y\|^2 \\ \text{subject to} & c_k + J_k(x - x_k) + \mu_k^R(y - y_k^E) = 0, \;\; x \geq 0, \;\; [x]_{\mathcal{A}_+} = 0. \end{array}$$

It will be shown that a sequence $\{x_k\}$ generated using this subproblem is superlinearly convergent to an optimal point. The proof involves showing that the subproblem satisfies the three properties described in Chapter 5 (see Page 68) for Fischer's generalized framework. The proof is similar to that of Fernández and Solodov [25].

The proof of the accuracy of the approximation defined by subproblem solution is identical because the only additional constraints are the fixed variables $[x]_{\mathcal{A}_+} = 0$, which are the values at the solution of the original problem. Similarly, the upper Lipschitz continuities of the solutions follows as well, since the addition of constraints cannot add solutions or make them unbounded. Local solvability of the subproblem, however, must be explicitly proven since the number of constraints has increased, potentially removing the existence of local solutions.

In the notation of Fischer's framework (see Page 68 in Chapter 5) that the generalized equation seeks to find a solution to

$$0 \in G(w) - \mathcal{T}(w),$$

where, in this case,

$$G(w) = \begin{pmatrix} g(x) - J(x)^T y \\ c(x) \end{pmatrix} \text{ and } \mathcal{T}(x) = \begin{pmatrix} \mathcal{N}(x) \\ 0 \end{pmatrix},$$

where

$$\mathcal{N}(x) = \begin{cases} \{b \in \mathbb{R}^n_+ \mid b \geq 0,\, b^T x = 0\} & \text{if } x \geq 0, \\ \emptyset, & \text{otherwise.} \end{cases}$$

The following will closely follow the proofs of the same assertions for inequality-constrained problems in Fernández and Solodov [25].

Proposition 8.6.1. *Let the SOSC hold at (x^*, y^*, z^*). There is a neighborhood $\mathcal{B}$ of (x^*, y^*, z^*) such that for all $(x, y, z) \in \mathcal{B}$, it holds that*

$$u^T H(x,y) u + \mu(x,y)\|v\|^2 \geq \gamma_1(\|u\|^2 + \mu(x,y)\|v\|^2),$$

for all (u, v) satisfying,

$$\begin{aligned} \nabla c_i(x)^T u &= \mu v_i, \\ u_i &= 0, \quad i \in \mathcal{A}_+(x^*, y^*, z^*), \\ u_i &\geq 0, \quad i \in \mathcal{A}_0(x^*, y^*, z^*). \end{aligned}$$

Proof. Assume the contrary, that there exist (x_k, y_k) and (u_k, v_k) such that

$$u_k^T H u_k + \mu_k \|v_k\|^2 < \frac{1}{k}(\|u_k\|^2 + \mu_k \|v_k\|^2).$$

Let $\xi_k = \|(u_k, \sqrt{\mu_k} v_k)\|$. Assume that

$$\frac{1}{\xi_k} \begin{pmatrix} u_k \\ \sqrt{\mu_k} v_k \end{pmatrix} \to \begin{pmatrix} \bar{u} \\ \bar{v} \end{pmatrix} \neq 0.$$

Since u_k satisfies $\nabla c_i(x)^T u_k = \mu_k v_i$ and $\mu_k \to 0$, $\nabla c_i(x)^T \bar{u} = 0$. Likewise, $[\bar{u}]_i = 0$ for $i \in \mathcal{A}_+$ and $[\bar{u}]_i \geq 0$ for $i \in \mathcal{A}_0$.

But, dividing the original assumption by ξ_k^2 and taking limits, it holds that,

$$\bar{u}^T H(x^*, y^*)\bar{u} + \|\bar{v}\|^2 = 0,$$

which contradicts the second-order sufficiency condition. □

Corollary 8.6.1. *There exists a neighborhood $\mathcal{B}$ of (x^*, y^*, z^*) such that for $(x, y, z) \in \mathcal{B}$, with $[x]_A = 0$, the matrix*

$$\begin{pmatrix} H(x,y) & -J(x)^T & -P_A \\ J(x) & \mu I & 0 \\ P_A^T & 0 & 0 \end{pmatrix},$$

is nonsingular.

Proof. Let (u, v, w) be in the kernel of this matrix. So,

$$\begin{aligned} 0 &= Hu - J^T v - P_A w, \\ 0 &= Ju + \mu v = 0, \\ [u]_A &= 0. \end{aligned}$$

By the second equation and $[u]_A = 0$, (u, v) are in the appropriate cone in Proposition 8.6.1.

Take the inner product of the first equation with u. Since $[u]_A = 0$, this results in:

$$0 = u^T H u - u^T J^T v.$$

By the second equation, $Ju = -\mu v$, so the previous equation becomes:

$$0 = u^T H u + \mu v^T v,$$

which, by Proposition 8.6.1, implies that $u = 0$ and $v = 0$, which implies, by the first equation row, $w = 0$.

□

Theorem 8.6.5 (Fernández and Solodov [25, Theorem 2])**.** *Let K be a closed convex cone. Suppose that $d = 0$ is the unique solution to the generalized complementarity problem $K \ni d \perp Md \in K^*$, and that M is copositive on K.*

Then for all q, the generalized complementarity problem of finding d such that $K \ni d \perp Md + q \in K^*$ has a nonempty compact solution set.

Proposition 8.6.2. *There is a neighborhood $\mathcal{B}$ of (x^*, y^*, z^*) such that for $(\bar{x}, \bar{y}, \bar{z})$, the mixed complementarity problem of finding (x, y, z) to satisfy*

$$\begin{aligned} 0 &= g + H(x - \bar{x}) + J(\bar{x})^T y - z, \\ 0 &= c(\bar{x}) + J(\bar{x})(x - \bar{x}) + \mu(y - \bar{y}) = 0, \\ & 0 \le z \perp x \ge 0, \\ & 0 = [x]_{\mathcal{A}_+}, \\ & 0 \le [x]_{A_0}. \end{aligned}$$

has a nonempty compact solution set.

Proof. Let

$$M = \begin{pmatrix} H & 0 & P_A \\ 0 & \mu I & 0 \\ P_A^T & 0 & 0 \end{pmatrix}, \ q = \begin{pmatrix} g - H\bar{x} \\ 0 \\ 0 \end{pmatrix},$$

and

$$b = c + J\bar{x} + \mu\bar{y}, \ a = \begin{pmatrix} J & -\mu I \end{pmatrix},$$

$$b_2 = 0, \ a_2 = \begin{pmatrix} I & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

The solution to the mixed complementarity problem above can be written as: Find $\bar{\xi} \in Q$ such that $\left\langle M\bar{\xi} + q, \xi - \bar{\xi} \right\rangle \ge 0$ for all $\xi \in Q$, where Q is:

$$Q = \left\{ \xi \mid A\xi + b = 0, \ A_{2,\mathcal{A}_+}\xi + b_{2,\mathcal{A}_+} = 0, \ A_{2,\mathcal{A}_0}\xi + b_{2,\mathcal{A}_0} \ge 0 \right\}.$$

Let $(\tilde{u}, \tilde{v}, \tilde{w}$ solve,

$$\begin{pmatrix} H & -J^T & P_A \\ J & \mu I & 0 \\ P_A^T & 0 & 0 \end{pmatrix} \begin{pmatrix} u \\ v \\ w \end{pmatrix} = \begin{pmatrix} -g - J^T y \\ c \\ 0 \end{pmatrix},$$

which has a unique solution due to Corollary 8.6.1.

It can be seen that if $\tilde{\xi} = (x + \tilde{u}, y + \tilde{v}, z + \tilde{w})$ then $A\tilde{\xi} = -b$ and $A_2\tilde{\xi} = -b_2$.

Thus $Q = \tilde{\xi} + K$, the critical cone, and so the solution of the original problem becomes $K \ni d \perp Md + M\tilde{\xi} + q \in K^*$.

Since M is copositive by Proposition 8.6.1, by Theorem 8.6.5 there is a nonempty compact solution set.

□

Proposition 8.6.3. *There is a neighborhood $\mathcal{B}$ of (x^*, y^*, z^*) and a constant $\gamma_3 > 0$ such that for all $(\bar{x}, \bar{y}, \bar{z}) \in \mathcal{B}$ with $[\bar{x}]_{\mathcal{A}_+} = 0$, it holds that any solution to the complementarity problem above satisfies*

$$\|(x - \bar{x}, y - \bar{y})\| \leq \gamma_3 \mu.$$

Proof. Suppose, to the contrary, that there is a sequence (x_k, y_k) such that $(\bar{x}_k, \bar{y}_k, \bar{z}_k) \to (x^*, y^*, z^*)$ and $\xi_k = \|(x - \bar{x}, y - \bar{y})\| > k\mu_k$, where (x, y, z) solves

$$\begin{aligned} 0 &= g + H(x - \bar{x}) + J(\bar{x})^T y - z, \\ 0 &= c(\bar{x}) + J(\bar{x})(x - \bar{x}) + \mu(y - \bar{y}), \\ 0 &\leq z \perp x \geq 0, \\ 0 &= [x]_{\mathcal{A}_+}, \\ 0 &\leq [x]_{A_0}. \end{aligned}$$

Notice that

$$\frac{\mu_k}{\xi_k} \leq \frac{1}{k} \to 0.$$

It holds that, by Lipschitz continuity of the constraint and objective functions,

$$\|c(\bar{x}_k)\| = \|c(\bar{x}_k) - c(x^*)\| \leq c_1 \mu,$$

and

$$\|J(\bar{x}_k) - J(x^*)\| \leq c_2 \mu.$$

Finally,

$$\|g - J^T y - z\| = \|g - J^T y - z - g(x^*) + J^T \hat{y} + \hat{z}\| \leq c(\|\bar{x} - x^*\| + \|y - \hat{y}\| + \|z - \hat{z}\| \leq c\mu,$$

where $\hat{y}$ and $\hat{z}$ are the projections of (y, z) onto $\bar{\mathcal{M}}_y(x^*)$, where the set $\bar{\mathcal{M}}_y$ is denoted to include the reduced costs.

Let (u, v, w) be such that,

$$\frac{1}{\xi} \begin{pmatrix} x - \bar{x} \\ y - \bar{y} \\ z - \bar{z} \end{pmatrix} \to \begin{pmatrix} u \\ v \\ w \end{pmatrix} \neq 0.$$

From the stationarity conditions of the subproblem, the constraint and objective function bounds, and the original assumption, it holds that,

$$Hu - J^T v - w = 0.$$

From the equality constraints, it holds that,

$$Ju + \mu v = 0.$$

Together, this implies that,

$$u^T H u + \frac{1}{\mu} v^T v + w^T u = 0.$$

Since $\bar{x}_A = 0$ and $z \perp x = 0$, it holds that

$$u^T H u + \frac{1}{\mu} v^T v = 0,$$

which, since $[u]_{\mathcal{A}_+} = 0$, $u_{A_0} \geq 0$, and $Ju = -\mu v$, contradicts Proposition 8.6.1. □

Theorem 8.6.6. *There is a neighborhood $\mathcal{B}$ of (x^*, y^*, z^*) such that for $(\bar{x}, \bar{y}, \bar{z})$ in $\mathcal{B}$ with $[\bar{x}]_{\mathcal{A}_+} = 0$, there is a solution to*

$$\begin{aligned} 0 &= g + H(x - \bar{x}) + J(\bar{x})^T y - z, \\ 0 &= c(\bar{x}) + J(\bar{x})(x - \bar{x}) + \mu(y - \bar{y}) = 0, \\ 0 &\leq z \perp x \geq 0, \end{aligned}$$

satisfying

$$\left\| \begin{pmatrix} x - \bar{x} \\ y - \bar{y} \\ z - \bar{z} \end{pmatrix} \right\| \leq \gamma \mu.$$

Proof. In view of Proposition 8.6.3, this holds if the solution to this problem satisfies $[x]_{\mathcal{A}_+} = 0$. However, from the proof of Theorem 8.6.4 $\bar{z}$ is such that $[z]_{\mathcal{A}_+} \geq \frac{1}{2}\epsilon_z$, so if μ is sufficiently small, this holds by the complementarity condition $z \perp x$.

□

Convergence and discussion

The results in the previous section show local existence of solutions to the sSQP subproblem with the additional constraint $[x]_{A_+} = 0$, which, together with upper-Lipschitz continuity and precise approximation, imply the final theorem:

Theorem 8.6.7. *Assume $\delta(x, y, z) < \bar{\delta}$ is sufficiently close to a first-order KKT point $(x^*, \mathcal{M}_y(x^*))$ satisfying the second-order sufficiency conditions and MFCQ. The sequence of iterations computed by pdSQPid0 is superlinearly convergent to a point (x^*, y^*, z^*).*

Thus, by convexifying only on the strongly active variable indices, the weakly active indices no longer present a problem of altering the reduced costs in a convexified solution and potentially making the solutions of stabilized SQP and pdSQP distinct.

An interesting result arising from the nature of the subproblem is:

Corollary 8.6.2. *The sequence of multipliers $\{y_k\}$ converges to a unique least-length multiplier solution to $g(x) - J(x)^T y = 0$.*

8.6.4 Active-set stabilization

Note that by the superlinear convergence estimate, it holds that,

$$\delta(x_{k+1}) = O(\delta(x)^{1+\gamma}),$$

Assuming that K is sufficiently large that μ_k^R estimates the active-set, $\mu_{k+1} = \delta(x)_{k+1}^{1+\tau} = \delta(x)_k^{\tau+\tau^2}$. Since $\delta(x)_k^{\tau+\tau^2} > \delta(x)^{1+\tau}$ and zero is the optimal value for the variables whose indices are in $\mathcal{A}$, if at the start of the QP subproblem, x_0 is set to the estimated active bounds, $[p_k]_A \leq \delta(x)^{1+\tau} < \mu_{k+1}$, which implies that it is active.

Chapter 9

Second-Order Primal-Dual SQP

9.1 Introduction

This chapter is concerned with the computation and use of a direction of negative curvature in the regularized sequential quadratic programming primal-dual augmented Lagrangian method (pdSQP) of Gill and Robinson [38, 39] for the purpose of ensuring convergence towards second-order optimal points. Section 9.2 discusses how to compute a direction of negative curvature using appropriate matrix factorizations. Section 9.3 discusses the specific relevant changes to the algorithm. Section 9.4 discusses the changes in the convergence results established by Gill and Robinson [39], showing that the desired convergence results continue to hold. In Section 9.5, global convergence to points satisfying the second-order necessary optimality conditions is established.

9.2 Direction of Negative Curvature

9.2.1 The active-set estimate

An index set $\mathcal{W}_k$ is maintained that consists of the variable indices that estimate which components of x on their bounds. This set determines the the space in which to calculate the directions of negative curvature. The tolerance for an index to be in $\mathcal{W}_k$ must converge to zero. A test such as $i \in \mathcal{W}_k$ if $[x_k]_i \leq \min\{\mu_k, \epsilon_a\}$, would be appropriate for the purpose of forming a $\mathcal{W}_k$ for convexification, initializing the QP, and obtaining a direction of negative curvature. Otherwise, it would be necessary to use three different factorizations.

9.2.2 Calculating the direction

In pdSQP, the QP must use a Lagrangian Hessian $\widetilde{H}$ such that $\widetilde{H}+\frac{1}{\mu}J^TJ$ is positive definite (see Chapter 8, Page 105). The process for forming the requisite $\widetilde{H}$, as well as calculating a direction of negative curvature begins with the inertia-controlling factorization of the KKT matrix (see Forsgren [29]). Consider the KKT matrix,

$$\begin{pmatrix} H_F & J_F^T \\ J_F & -\mu I_{|F|} \end{pmatrix}, \tag{9.1}$$

with F the set of estimated free variables (those not in $\mathcal{W}_k$), and $I_{|F|}$ the identity matrix with $|F|$ rows and columns.

The algorithm begins an LBL^T factorization of the KKT matrix, where L is lower triangular and B is a symmetric diagonal with 1×1 and 2×2 diagonal blocks. Standard pivoting strategies are described in the literature (see Bunch and Parlett [14], Fletcher [28], and Bunch and Kaufman [13]). Let the lower-right block be defined as $D = -\mu I_{|F|}$.

At step k of the factorization, let the partially factorized matrix have the following structure:

$$\begin{pmatrix} L_1 & 0 \\ L_2 & I \end{pmatrix} \begin{pmatrix} B & 0 \\ 0 & A \end{pmatrix} \begin{pmatrix} L_1^T & L_2^T \\ 0 & I \end{pmatrix},$$

with L_1 being lower triangular, I the identity of appropriate size, and A the matrix remaining to be factorized. Let A be partitioned as $A = \begin{pmatrix} a & b^T \\ b & C \end{pmatrix}$. If the top left element is chosen as a 1×1 pivot, at the next step,

$$\begin{pmatrix} L_1 & 0 & 0 \\ L_3 & 1 & 0 \\ L_4 & a^{-1}b & I \end{pmatrix} \begin{pmatrix} B & 0 & 0 \\ 0 & a & 0 \\ 0 & 0 & C - ba^{-1}b^T \end{pmatrix} \begin{pmatrix} L_1^T & L_3^T & L_4^T \\ 0 & 1 & a^{-1}b^T \\ 0 & 0 & I \end{pmatrix}.$$

Let $S = C - ba^{-1}b$ denote the *Schur complement* of the factorization. The matrix S is factorized at the next step.

For inertia control, this factorization has two stages. In the first stage, we restrict the factorization to allow only for pivots of type H^+, D^- or HD. This means that an element (i,j) of H is selected such that $H_{ij} > 0$, a diagonal element of D is selected, or (i_1, i_2, j_1, j_2) is selected such that (i_1, j_1) is an element of H, (i_2, j_2) is an element of D and $S_k[(i_1, i_2), (j_1, j_2)]$ has mixed eigenvalues. This procedure is continued until there are no such remaining pivots.

The KKT matrix can be partitioned as

$$\begin{pmatrix} H_{11} & H_{12} & J_1^T \\ H_{21} & H_{22} & J_2^T \\ J_1 & J_2 & -\mu I \end{pmatrix},$$

where, all of the pivots have come from the rows and columns of H_{11}, J_1, and $-\mu I$. At the end of the first stage, the factorization can be written as:

$$\begin{pmatrix} L_1 & 0 \\ L_2 & I \end{pmatrix} \begin{pmatrix} B & 0 \\ 0 & H_{22} - K_{21}K_{11}^{-1}K_{12} \end{pmatrix} \begin{pmatrix} L_1^T & L_2^2 \\ 0 & I \end{pmatrix}. \tag{9.2}$$

Let $S = H_{22} - K_{21}K_{11}^{-1}K_{12}$. Forsgren [29, Proposition 3] shows that if δI is added to H_{22} such that $\delta > \|S\|$ then K_F has the correct inertia. In practice this δ is excessively large for the purpose of constructing the appropriate matrix with the required eigenvalues, but this result does indicate that such a constant exists.

Instead of proceeding to the second phase of this factorization, the procedure of Forsgren et al. [32, Lemma 2.4] is applied to S to compute $\widehat{u}$, a direction of negative curvature for S. The procedure to calculate this $\widehat{u}$ is as follows:

Let $\rho = \max_{i,j} |[S]_{ij}|$ with $|[S]_{qr}| = \rho$. Define $\widehat{u}$ as:

$$\widehat{u} = \sqrt{\rho}h, \tag{9.3}$$

where

$$h = \begin{cases} e_q & \text{if } q = r, \\ \frac{1}{\sqrt{2}}(e_q - \operatorname{sgn}([S]_{qr})e_r) & \text{otherwise.} \end{cases}$$

This $\widehat{u}$ satisfies $\widehat{u}^T S\widehat{u} \le \gamma\lambda_{\min}(S)\|\widehat{u}\|^2$, with γ independent of S.

The following bounds are important for the subsequent second-order convergence theory.

Lemma 9.2.1. *Let $\widehat{u}$ be defined as in* (9.3)*, S be the Schur complement of the partially factorized matrix* (9.2)*, J_F and H_F defined as in* (9.1)*, and Z a matrix consisting of columns for the basis of the null-space of J_F, then*

$$\frac{\widehat{u}^T S\widehat{u}}{\gamma\|\widehat{u}\|^2} \le \lambda_{\min}(S) \le \lambda_{\min}(H_F + \frac{1}{\mu}J_F^T J_F) \le \lambda_{\min}(Z^T H_F Z).$$

Proof. Forsgren et al. [32, Lemma 2.4] directly implies that $\widehat{u}^T S\widehat{u}/\|\widehat{u}\|^2 \le \gamma\lambda_{\min}(S)$, where $\gamma > 0$.

The proof that $\lambda_{\min}(S) \le \lambda_{\min}(H_F + \frac{1}{\mu} J_F^T J_F)$ is given in the proof of Theorem 4.5 in Forsgren and Gill [30]. For the final inequality, let $w = Zv$, with $Z^T H_F Z v = \lambda_{\min}(Z^T H_F Z) v$ and $\|v\| = 1$. Then

$$\lambda_{\min}(H_F + \frac{1}{\mu} J_F^T J_F) \le \frac{w^T (H_F + \frac{1}{\mu} J_F^T J_F) w}{w^T w} = w^T H_F w = v^T Z^T H_F Z v = \lambda_{\min}(Z^T H_F Z).$$

□

9.3 Implementing Directions of Negative Curvature

9.3.1 Step of negative curvature

Several changes must be made to the algorithm of Gill and Robinson [39]. In order to minimize the number of factorizations, the computation of the direction of negative curvature should be followed by a test of second-order optimality. In addition, it is necessary that the direction of negative curvature is bounded, and a feasible direction with respect to both the linearized equalities and the bound constraints. Finally, the line search must be extended to allow for this additional step of negative curvature.

In the description below, the subscript k denoting the step number in the sequence of iterations is suppressed.

The following procedure satisfies these requirements.

1. The first step computes the direction of negative curvature for the free KKT-matrix as described in Section 9.2, denoted as $\widehat{u}_S$, then defines $\widehat{u}_F$ to be $[\widehat{u}_F]_{\mathcal{S}} = \widehat{u}_S$ with $\mathcal{S}$ corresponding to indices corresponding to the remaining unfactorized entries of H_F, and $[\widehat{u}_F]_{\mathcal{S}^c} = 0$. Then the step defines $\widehat{u}$ to be $[\widehat{u}]_F = \widehat{u}_F$ and $[\widehat{u}]_A = 0$. If no such direction of negative curvature exists, then $\widehat{u}$ is set to zero.

2. The second step uses $\widehat{u}$ in a test of second-order optimality. This is described in Section 9.3.2.

3. The corresponding change in the multipliers corresponding to the definition for $\widehat{u}$ is defined as $\widehat{w} = -\frac{1}{\mu} J\widehat{u}$. This ensures that the linearized equality constraints are satisfied, i.e.,
$$0 = Jp + c + \mu q = J(p + \widehat{u}) + c + \mu(q - \frac{1}{\mu} J\widehat{u}).$$
The final resulting (u, w) is shown below in Section 9.3.3 to be a direction of negative curvature for $\nabla^2 M^\nu$.

4. Since both $(\widehat{u}, \widehat{w})$ and $-(\widehat{u}, \widehat{w})$ are directions of negative curvature, the sign is chosen so that the step is a descent direction for ∇M, i.e., $\nabla M^T \begin{pmatrix} \widehat{u} \\ \widehat{w} \end{pmatrix} \leq 0$.

5. Compute $\Delta v = (p, q)$, the solution of the convex QP.

6. The direction of negative curvature is scaled so that it is both bounded by $\max(u_{\max}, 2\|p\|)$ and also, in conjunction with the definition of the QP step, satisfies the bound constraints $x \geq 0$.

   Specifically, u and w are set as $u = \beta\widehat{u}$ and $w = \beta\widehat{w}$, where

$$\beta = \left\{ \max \hat{\beta} \mid x + p + \hat{\beta}\widehat{u} \geq 0, \ \|\hat{\beta}\widehat{u}\| \leq \max(u_{\max}, 2\|p\|) \right\}.$$

   Note that this implies that if $[x + p]_i = 0$ and $[u]_i < 0$, then u is set to zero.

9.3.2 Optimality measures

In Gill and Robinson [39], an iterate is an S-iterate if $\phi_S(v) \leq \frac{1}{2}\phi_S^{max}$ and an L-iterate if $\phi_L(v) \leq \frac{1}{2}\phi_L^{max}$, where

$$\phi_S(v) = \xi(x) + 10^{-5}\omega(v) \quad \text{and} \quad \phi_L(v) = 10^{-5}\xi(x) + \omega(v),$$

with

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

Otherwise, an iterate is an M-iterate if

$$\|\nabla_y M^\nu(v_{k+1}; y_k^E, \mu_k^R)\| \leq \tau_k \quad \text{and} \quad \|\min(x_{k+1}, \nabla_x M^\nu(v_{k+1}; y_k^E, \mu_k^R))\| \leq \tau_k.$$

If none of these conditions hold, then v_k is an F-iterate.

In order to force convergence to a second-order optimal point, it is necessary to change the function $\omega(x, y)$ that appears in ϕ_S and ϕ_L, as well as the test for an iteration being an M-iterate.

Ideally, the minimum eigenvalue of H on the null-space of J_F should be computed, as well as the minimum eigenvalue of ${\nabla^2}_{xx} M^\nu$. However, this would require extensive computation. Instead, these quantities are estimated based on the value of the negative curvature. It holds that

$$\frac{\widehat{u}^T(H + \frac{1}{\mu}J^TJ)\widehat{u}}{\gamma\|\widehat{u}\|^2} \leq \lambda_{\min}\Big(H + \frac{1}{\mu}J^TJ\Big),$$

where the suffix F is omitted for clarify. As γ is bounded from below and above, if $\widehat{u}^T(H + \frac{1}{\mu}J^TJ)\widehat{u}/\|\widehat{u}\|^2 \to 0$, the estimate for $\widehat{u}$ implies $\lim \lambda_{\min}(H + \frac{1}{\mu}J^TJ) \geq 0$. Hence, the test for the optimality of an M-iterate is:

$$\begin{aligned}
\|\nabla_y M(v_{k+1}; y_k^E, \mu_k^R)\| &\leq \tau_k, \\
\|\min(x_{k+1}, \nabla_x M^\nu(v_{k+1}; y_k^E, \mu_k^R))\| &\leq \tau_k, \\
\text{and} \qquad \widehat{u}_{k+1}^T\Big(H + \frac{1}{\mu}J^TJ\Big)\widehat{u}_{k+1} &\geq -\|\widehat{u}_{k+1}\|^2\tau_k.
\end{aligned}$$

Similarly, for the filter functions,

$$\phi_S(v) = \eta(x) + 10^{-5}\omega(v) \ \text{ and } \ \phi_L(v) = 10^{-5}\eta(x) + \omega(v)$$

the optimality test functions are $\eta(x) = \|c(x)\|$, and

$$\omega(x, y) = \min\left(\|\min(x, g(x) - J(x)^Ty)\|, -\frac{\widehat{u}_{k+1}^T(H + \frac{1}{\mu}J^TJ)\widehat{u}_{k+1}}{\|\widehat{u}_{k+1}\|^2}\right).$$

9.3.3 The merit function

The line search must also be modified to include the direction of negative curvature. First, it will be shown that the full primal-dual step is a step of negative curvature for the merit function Hessian.

Lemma 9.3.1. *The vector (u, w) defined as in* Section 9.3.1 *is a direction of negative curvature for $\nabla^2 M^\nu$ for all $\nu \geq 0$.*

Proof. From the definition of $\nabla^2 M^\nu$ it holds that

$$\begin{aligned}
\begin{pmatrix} u \\ w \end{pmatrix}^T \nabla^2 M^\nu \begin{pmatrix} u \\ w \end{pmatrix} &= \begin{pmatrix} u \\ w \end{pmatrix}^T \begin{pmatrix} H + \frac{1}{\mu}(1+\nu)J^TJ & \nu J^T \\ \nu J & \nu\mu I \end{pmatrix} \begin{pmatrix} u \\ w \end{pmatrix} \\
&= u^THu + \frac{1}{\mu}(1+\nu)u^TJ^TJu + 2\nu u^TJ^Tw + \nu\mu\|w\|^2.
\end{aligned}$$

From the definition above, $u = \beta\widehat{u}$ and $\widehat{u}^T(H + \frac{1}{\mu}J^TJ)\widehat{u} \leq \gamma\lambda_{\min}(H + \frac{1}{\mu}J^TJ)\|\widehat{u}\|^2$, so multiplying both sides by β^2, the expression becomes $u^T(H + \frac{1}{\mu}J^TJ)u \leq \gamma\lambda_{\min}(H + \frac{1}{\mu}J^TJ)\|u\|^2$. If $\bar{\gamma} = \gamma\lambda_{\min}(H + \frac{1}{\mu}J^TJ)$ and $w = -\frac{1}{\mu}Ju$, then

$$\begin{aligned}
u^THu + \frac{1}{\mu}(1+\nu)u^TJ^TJu + 2\nu u^TJ^Tw + \nu\mu\|w\|^2 &\leq -\bar{\gamma}\|u\|^2 - \frac{2\nu}{\mu}u^TJ^TJu + \frac{\nu}{\mu}\|Ju\|^2 \\
&= -\bar{\gamma}\|u\|^2 - \frac{\nu}{\mu}\|Ju\|^2 \\
&\leq -\bar{\gamma}\|u\|^2 - \nu\mu\|w\|^2.
\end{aligned}$$

which is negative for all $\nu \geq 0$. ∎

In order to simplify the notation, we suppress the dependence of the merit function on ν and write $M(v; y^E, \mu) = M^\nu(v; y^E, \mu)$.

Two approaches may be considered for the definition of the line search. The first extends the method for unconstrained functions proposed by McCormick [65]. Let s denote the primal-dual negative-curvature step $s = \begin{pmatrix} u \\ w \end{pmatrix}$. Let R_k denote the curvature $s_k^T \nabla^2 M(v_k; y_k^E, \mu_k^R) s_k$, which is non positive by definition. McCormick defines $\alpha_k = 2^{-j}$ such that

$$M(v_k + \alpha_k s_k + \alpha_k^2 \Delta v_k; y_k^E, \mu_k^F) \leq M + \alpha_k^2 \eta_S N_k + \alpha_k \eta_S R_k. \tag{9.4}$$

If we denote $\bar{\alpha} \triangleq \min(\alpha_{\min}, \alpha_k)$ and $\widehat{\mu} \triangleq \max\left(\frac{1}{2}\mu_k, \mu_{k+1}^R\right)$, a suitable update for the penalty parameter is

$$\mu_{k+1} = \begin{cases} \mu_k, & M(v_{k+1}; y_k^E, \mu_k) \leq M(v_k; y_k^E, \mu_k) + \bar{\alpha}\eta_S R_k + \bar{\alpha}^2 \eta_S N_k \\ \widehat{\mu}, & \text{otherwise.} \end{cases} \tag{9.5}$$

This curvilinear line-search is expected to be effective in reaching a second-order minimizer of the primal-dual augmented Lagrangian merit function. However, there are more robust methods available. Moguerza and Prieto [66] formulate an interior-point augmented Lagrangian method that incorporates directions of negative curvature. Olivares et al. [72] give a set of tests to determine whether or not a direction of negative curvature is appropriate for a given iteration. The method presented here is based on a combination these two methods.

Consider the quadratic model of M at x_k,

$$M_2(d) = \nabla M^T d + \frac{1}{2} d^T \nabla_{xx}^2 M d.$$

Let η_1 and η_2 be constants such that $0 < \eta_2 < 2 < \eta_1$, and $\eta_3 > 0$.

If it holds that,

$$\eta_2 M_2(s) \geq \frac{\nabla M(v_k)^T \Delta v}{\|\Delta v\|_2} \geq \eta_1 M_2(s),$$

then both directions are suitable for decreasing the merit function. On the other hand, if it holds that,

$$\frac{\nabla M(v_k)^T \Delta v}{\|\Delta v\|_2} < \eta_1 M_2(s),$$

then the direction of negative curvature provides an insufficient reduction of the merit function and the standard line-search using only Δv will be used. Finally, if

$$\eta_2 M_2(s) < \frac{\nabla M(v_k)^T \Delta v}{\|\Delta v\|_2} \text{ and } \nabla M(v_k) \geq \eta_3 s^T H_M(v_k) s,$$

then it is clear that the direction of negative curvature provides for greater reduction in the merit function than the QP solution, and the one-direction line-search will be used with s.

In this case, the step length $\alpha_k = 1$ is used if it holds that

$$\phi(1) \leq \phi(0) + \tfrac{1}{2}\gamma_1 \phi''(0),$$

where $\phi(\alpha) = M(v_k + \alpha^2 \Delta v + \alpha s)$. Otherwise, α is reduced by a backtracking procedure until a value α_k is found that satisfies

$$\begin{aligned} \phi(\alpha_k) &\leq \phi(0) + \gamma_1 \phi''(0), \\ \phi'(\alpha_k) &\geq \gamma_2(\phi'(0) + \alpha_k \phi''(0)), \\ \|c(x(\alpha))\| &\leq \beta_c, \end{aligned}$$

where $0 < \gamma_1 < \frac{1}{2}$, $\frac{1}{2} < \gamma_2 < 1$, and $\beta_c = \|c(x_0)\|$.

9.4 Consistency with Established Convergence Theory

The first-order convergence analysis of Gill and Robinson [39], requires three assumptions.

Assumption 9.4.1. *Each $\bar{H}(x_k, y_k)$ is chosen so that the sequence $\{\bar{H}(x_k, y_k)\}_{k\geq 0}$ is bounded, with $\{\bar{H}(x_k, y_k) + (1/\mu_k^R)J(x_k)^T J(x_k)\}_{k\geq 0}$ uniformly positive definite.*

Assumption 9.4.2. *The functions f and c are twice continuously differentiable.*

Assumption 9.4.3. *The sequence $\{x_k\}_{k\geq 0}$ is contained in a compact set.*

As ∇M^ν does not involve any term involving the objective or constraint Hessians, much of the first-order convergence theory holds. The use of a direction of negative curvature implies that Theorem 4.1 of Gill and Robinson [39] must be restated as follows.

Theorem 9.4.1. *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:*

1. $\{\|\Delta v_k\| + \|u_k\|\}_{k\geq \widehat{k}}$ *is bounded away from zero*

2. *There exists a positive ϵ such that for all $k \geq \widehat{k}$, it holds that*

$$\nabla M^\nu(v_k; y_k^E, \mu_k^R)^T \Delta v_k \leq -\epsilon \quad \text{or} \quad s_k^T \nabla^2 M^\nu(v_k; y_k^E, \mu_k^R) s_k \leq -\epsilon.$$

Proof. If all iterates $k \geq \widehat{k}$ are F-iterates, then,

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

Proof of the first result: Assume the contrary, i.e., there exists a subsequence $\mathcal{S}_1 \subset \left\{ k \mid k \geq \widehat{k} \right\}$ such that $\lim_{k \in \mathcal{S}_1} \Delta v_k = 0$ and $\lim_{k \in \mathcal{S}_1} u_k = 0$. The solution Δv_k to the QP subproblem satisfies

$$\begin{pmatrix} z_k \\ 0 \end{pmatrix} = H_M^\nu(v_k; \mu^R)\Delta v_k + \nabla M^\nu(v_k; y^E, \mu^R) \quad \text{and} \quad 0 = \min(x_k + p_k, z_k).$$

As H_M^ν is uniformly bounded, eventually for some $k \in \mathcal{S}_1$ sufficiently large, Δv_k satisfies the first-order conditions of an M-iterate, i.e.,

$$\|\nabla_y M(v_{k+1}; y_k^E, \mu_k^R)\| \leq \tau_k \quad \text{and} \quad \|\min(x_{k+1}, \nabla_x M^\nu(v_{k+1}; y_k^E, \mu_k^R))\| \leq \tau_k.$$

In the construction of u_k, $\|u\|$ is the largest possible value, subject to an upper bound, that is feasible. This implies that if $\lim u_k \to 0$, then eventually, u is constrained by feasibility, or set to zero.

Consider the first case, i.e., the limiting upper bound constraint on u_k must be $x_k + p_k + u_k \geq 0$. Since $u_k \to 0$ and $p_k \to 0$, eventually, if i is a blocking bound for u_k, $x_i \leq \min(\mu, \epsilon_x)$ and $i \in \mathcal{W}_k$, which implies that $[u_k]_i \equiv 0$. Hence, by construction and the fact that the set of possible indices is finite, u_k is eventually identically zero. This implies that the second-order conditions of an M-iterate are also satisfied trivially, i.e.,

$$\widehat{u}_{k+1}^T \Big(H + \frac{1}{\mu} J^T J \Big) \widehat{u}_{k+1} \geq \tau_k \|\widehat{u}_{k+1}\|^2,$$

and μ_k^R is decreased. This contradicts the assumption that μ_k^R is held fixed at $\mu_k^R \equiv \mu^R$ for all $k \geq \widehat{k}$.

Proof of Part 2. Assume that the result does not hold, i.e., there exists a subsequence $\mathcal{S}_2$ of $\{k : k \geq \widehat{k}\}$ such that

$$\lim_{k \in \mathcal{S}_2} \nabla M^\nu(v_k; y^E, \mu^R)^T \Delta v_k = 0 \tag{9.6}$$

and

$$\lim_{k \in \mathcal{S}_2} s_k^T \nabla^2 M^\nu(v_k; y_k^E, \mu_k^R) s_k = 0.$$

Consider the matrix

$$L_k = \begin{pmatrix} I & 0 \\ \frac{1}{\mu^R} J_k & I \end{pmatrix}.$$

Since the $\Delta v = 0$ is feasible and Δv_k a solution for the convex problem, it follows that

$$\begin{aligned} -\nabla M^\nu(v_k; y^E, \mu^R)^T \Delta v_k &\geq \tfrac{1}{2} \Delta v_k^T H_M^\nu(v_k; \mu^R) \Delta v_k \\ &= \tfrac{1}{2} \Delta v_k^T L_k^{-T} L_k^T H_M^\nu(v_k; \mu^R) L_k L_k^{-1} \Delta v_k \\ &= \begin{pmatrix} p_k \\ q_k + \frac{1}{\mu^R} J_k p_k \end{pmatrix}^T \begin{pmatrix} \bar{H}_k + \frac{1}{\mu^R} J_k^T J_k & 0 \\ 0 & \nu \mu^R \end{pmatrix} \begin{pmatrix} p_k \\ q_k + \frac{1}{\mu^R} J_k p_k \end{pmatrix} \end{aligned}$$

As H_M^ν is bounded, it must hold that

$$\Delta v_k^T L_k^{-T} L_k^T H_M^\nu(v_k; \mu^R) L_k L_k^{-1} \Delta v_k \geq \bar{\lambda}_{\min} \|p_k\|^2 + \nu \mu^R \|q_k + (1/\mu^R) J_k p_k\|^2,$$

for some $\bar{\lambda}_{\min} > 0$. Combining this bound with (9.6) it follows that

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

in which case $\lim_{k \in \mathcal{S}_2} q_k = 0$. Hence $\Delta v_{k \in \mathcal{S}_2} \to 0$.

As $\lim_{k \in \mathcal{S}_2} s_k^T \nabla^2 M^\nu(x_k, y_k; y^E, \mu) s_k = 0$, either there exists a $\widehat{k}_2$, such that for all $k \geq \widehat{k}_2$, $\gamma s_k^T \nabla^2 M^\nu(x_k, y_k; y^E, \mu) s_k / \|s_k\|^2 > -\tau$ or $u_k \to 0$, where $\gamma > 0$ is the scalar defined in Lemma 9.2.1. The first case, by the same argument as for Part 1, together with $\Delta v_k \to 0$, implies that eventually k is an M-iterate. The latter, together with $\lim \Delta v_k = 0$, contradicts the statement of Part 1 of the theorem. This implies that Part 3 must hold. □

The proofs of the first result of Theorem 4.1, and the result of Theorem 4.2 of Gill and Robinson [39] hold for the modified algorithm.

9.5 Global Convergence to Second-order Optimal Points

9.5.1 Filter convergence

Definition 9.5.1. *The Weak Constant Rank (WCR) condition holds at x if there is a neighborhood $\mathcal{M}(x)$ for which the rank of $\begin{pmatrix} J(z) \\ E_A^T \end{pmatrix}$ is constant for all $z \in \mathcal{M}(x)$, where E_A is the columns of the identity corresponding to the indices of x active at x (as in $i \in \mathcal{A}$ if $x_i = 0$).*

Theorem 9.5.1. *Assume there is a subsequence v_k of S- and L-iterates converging to v^*, with $v^* = (x^*, y^*)$ satisfying the first-order KKT conditions. Furthermore, assume that MFCQ and WCR hold at v^*. Then v^* satisfies the necessary second-order necessary optimality conditions.*

Proof. Let $d \in \widetilde{C}(x^*) \equiv \left\{ d \mid J(x^*)d = 0 \text{ and } E_{w^*}^T d = 0 \right\}$ with $\|d\| = 1$. By Lemma 3.1 of Andreani et al. ([6]) there exists $\{d_k\}$ such that $d_k \in \widetilde{C}(x_k)$ and $d_k \to d$, where

$$\widetilde{C}(x_k) = \left\{ d \mid J(x_k)d = 0 \quad \text{and} \quad E_{w^*}^T d = 0 \right\}.$$

Without loss of generality, we may let $\|d_k\| = 1$. Since $x_k \to x^*$, eventually $\mathcal{W}_k = \mathcal{A}^*$, where $\mathcal{A}^*$ is the active set at x^*. Then, from the definitions of an S- and L-iterate, and Lemma 9.2.1, it follows that $d_k^T(\nabla^2 f(x_k) + \sum y_k \nabla^2 c(x_k))d_k > \lambda_{\min}(Z_k H_k Z_k) > -\xi_k$, where $0 < \xi_k \to 0$. Taking limits, it follows that $d^T(\nabla^2 f(x_k) + \sum y^* \nabla^2 c(x^*))d \geq 0$. $\square$

Chapter 10

Numerical Results

10.1 Standard Test Problems

The algorithms were tested on optimization problems from the CUTEr test set. These problems include cases from from industrial applications, standard academic problems, and problems designed to exploit common weaknesses of optimization algorithms (see Gould et al. [11, 46] for a more detailed description).

For many of the problems in the CUTEr test set, the size of the problem can be specified. In these cases, the number of variables and constraints were chosen to be the largest permissible values less than 500. All the problems selected have at least one constraint (not including any simple upper or lower bounds on variables). The total of 540 problems were selected.

For each algorithm tested, the total number of outer iterations was limited to 1000. If an algorithm did not converge in 1000 iterations, it was considered to have not converged. The threshold for the optimality measures was set at 10^{-8}. If the infeasibility and stationarity measures are both below this value, the algorithm was terminated and the run was considered to be a success.

10.2 First-Order pdSQP Results

10.2.1 Global convergence results

In total, pdSQP converged to a point satisfying the optimality conditions for 407 (75%) problems, compared to 452 (84%) for SNOPT, and 319 (59%) for MATLAB's `fmincon`

Table 10.1: Reliability of SNOPT and pdSQP

Algorithm	Total	Rank-deficient	SC fails	SONC	SOSC
SNOPT	452	70	198	431	337
pdSQP	407	55	127	402	366

SQP algorithm. These results are encouraging, as SNOPT and fmincon are sophisticated packages that have been developed over a number of years, whereas pdSQP is a prototype MATLAB implementation.

Degenerate problems are of interest, as part of the intention of pdSQP is to use the results of sSQP for degenerate problems. There is no *a priori* expectation of SNOPT or pdSQP performing better in terms of *global* convergence, however, it is important to investigate the practical global convergence of the primal-dual merit function compared to the standard augmented Lagrangian merit function in SNOPT. Although a test on a large number of known degenerate problems cannot be performed, as a solution needs to be known in advance, the number of convergent problems for SNOPT and pdSQP can be compared, with a higher proportion of degenerate problems suggesting the global procedure performs comparatively well or poorly on those sets of problems.

In addition, as pdSQP uses second derivatives, the number of problems satisfying second-order necessary and sufficiency conditions are included. This also provides motivation for the potential improvement in including directions of negative curvature.

There are several noticeable patterns. First, there is a distinct difference in the quantities of convergent problems failing to satisfy strict-complementarity. This is intuitively plausible, as the convexification procedure results in a solution for the original indefinite problem only if strict-complementarity holds. Furthermore, the local convergence results of sSQP rely on second-order sufficiency holding at the solution, and while there are no global convergence results involving sSQP methods, aside from pdSQP, as the convergence theory shows, local convergence results can imply global convergence if a cluster point at a local minimizer exists, so pdSQP has *stronger* theoretical global convergence results for problems satisfying the SOSC, which is corroborated by the data shown. In addition, the larger gap in the quantity of problems satisfying the SONC and SOSC for SNOPT compared to pdSQP corroborates the experiments conducted by Izmailov and Solodov [56] that indicate that sSQP is less likely to exhibit dual convergence to critical multipliers.

10.3 Results for Selected Problems

A set of 116 problems were selected for a more detailed analysis of the performance of different variations of pdSQP, in particular the Hock-Schittkowski "HS" problems from the CUTEr test set (see also, Hock and Schittkowski [53]).

The next table summarizes the convergence results for the four solvers: pdSQP, pdSQPcc, pdSQPid0 and pdSQPccnc. The solver pdSQP is the primal-dual SQP algorithm implemented with preconvexification. pdSQPcc is pdSQP with concurrent and post convexification. pdSQPid0 is the variation of pdSQP that uses the ID0 procedure to identify strongly active variables. pdSQPccnc is pdSQP with concurrent convexification and a direction of negative curvature.

Below are the results for the 71 equality-constrained problems.

Algorithm	Solved
`pdSQP`	56
`pdSQPcc`	48
`pdSQPid0`	48
`pdSQPccnc`	40

The details of the results for each solver are listed in the tables below. As there are no inequality constraints, the results for pdSQPid0 are not shown because they are identical to those of pdSQPcc. Overall, it was found that the solver pdSQPcc required post convexification in at least one subproblem for 18 (25%) of the problems.

For the problems that could not be solved, an alphabetic code is used to indicate the reason for the failure.

- m – the maximum number of iterations was exceeded
- b – the maximum number of backtracks for the line-search was reached
- i – the second order modification of the free Hessian matrix failed
- c – the preconvexification procedure failed
- q – the QP solver failed to produce a solution
- d – the QP solver failed to produce a descent direction for the merit function

Table 10.2: Results for pdSQP on equality constrained problems

Name	m	n	Result	nFun	nQPs	f-value
BT1	1	2	S	14	9	-1.00e+00
BT2	1	3	S	11	10	3.26e-02
BT3	3	5	S	4	3	4.09e+00
BT4	2	3	S	17	8	-4.55e+01
BT5	2	3	S	10	7	9.62e+02
BT6	2	5	S	21	17	4.88e+00
BT7	3	5	S	42	30	3.06e+02
BT8	2	5	S	11	10	1.00e+00
BT9	2	4	S	9	8	-1.00e+00
BT10	2	2	S	7	6	-1.00e+00
BT11	3	5	S	8	7	8.25e-01
BT12	3	5	S	8	5	6.19e+00
BYRDSPHR	2	3	S	52	30	-4.68e+00
COOLHANS	9	9	S	37	24	0.00e+00
DIXCHLNG	5	10	F^m	606	601	4.27e+03
EIGENA2	55	110	S	5	3	6.04e-24
EIGENACO	55	110	S	5	3	1.77e-20
EIGENB2	55	110	F^c	27	7	1.80e+01
EIGENBCO	55	110	F^m	1224	601	8.99e+00
EIGENC2	231	462	F^c	17	5	3.63e+02
EIGENCCO	231	462	F^c	116	53	1.84e+01
ELEC	200	600	F^i	6	2	2.84e+04
GRIDNETE	36	60	S	5	4	3.96e+01
GRIDNETH	36	60	S	5	4	3.96e+01
HS6	1	2	S	38	17	9.98e-31
HS7	1	2	S	38	18	-1.73e+00
HS8	2	2	S	6	4	-1.00e+00
HS9	1	2	S	5	4	-5.00e-01
HS26	1	3	S	21	18	1.02e-12
HS27	1	3	S	121	69	4.00e-02

Table 10.2: Results for pdSQP on equality constrained problems (continued)

Name	m	n	Result	nFun	nQPs	f-value
HS28	1	3	S	4	2	9.86e-32
HS39	2	4	S	9	8	-1.00e+00
HS40	3	4	S	5	4	-2.50e-01
HS42	2	4	S	8	6	1.39e+01
HS46	2	5	S	24	18	5.07e-12
HS47	3	5	S	23	16	1.92e-10
HS48	2	5	S	4	2	2.47e-31
HS49	2	5	S	18	14	8.43e-10
HS50	3	5	S	14	10	1.42e-22
HS51	3	5	S	4	2	1.21e-14
HS52	3	5	S	4	3	5.33e+00
HS56	4	7	S	6	5	-3.46e+00
HS61	2	3	S	19	14	-1.44e+02
HS77	2	5	S	18	16	2.42e-01
HS78	3	5	S	6	5	-2.92e+00
HS79	3	5	S	8	7	7.88e-02
HS100LNP	2	7	S	18	7	6.81e+02
HS111LNP	3	10	S	25	14	-4.78e+01
LUKVLE1	98	100	S	15	9	5.50e-16
LUKVLE3	2	100	S	18	13	2.76e+01
LUKVLE6	49	99	F^m	606	601	2.00e+06
LUKVLE7	4	100	S	23	12	-1.30e+01
LUKVLE8	98	100	F^m	602	601	4.65e+05
LUKVLE9	6	100	S	43	21	1.02e+01
LUKVLE10	98	100	S	51	31	3.48e+01
LUKVLE13	64	98	S	26	25	7.90e+02
LUKVLE14	64	98	F^m	602	601	1.06e+06
LUKVLE16	72	97	F^m	604	601	1.62e+05
LCH	1	300	F^c	59	18	-6.41e+00
LCH	1	300	F^c	59	18	-6.41e+00

Table 10.2: Results for pdSQP on equality constrained problems (continued)

Name	m	n	Result	nFun	nQPs	f-value
MARATOS	1	2	S	4	3	-1.00e+00
MWRIGHT	3	5	S	11	5	2.50e+01
ORTHRDM2	100	203	S	7	5	7.78e+00
ORTHRDS2	100	203	F^m	1206	601	8.28e+02
ORTHREGA	64	133	S	48	27	3.50e+02
ORTHREGB	6	27	S	7	5	3.02e-23
ORTHREGC	10	25	S	7	6	3.99e-01
ORTHREGD	10	23	F^b	379	377	3.25e+01
ORTHRGDM	10	23	F^b	473	399	9.65e+00
ORTHRGDS	76	155	S	239	118	2.34e+01
S316-322	1	2	S	10	9	3.34e+02

Table 10.3: Results for pdSQPcc on equality constrained problems

Name	m	n	Result	nFun	nQPs	f-value
BT1	1	2	S	14	9	-1.00e+00
BT2	1	3	S	11	10	3.26e-02
BT3	3	5	S	4	3	4.09e+00
BT4	2	3	S	19	9	-4.55e+01
BT5	2	3	S	40	18	9.62e+02
BT6	2	5	F^q	43	18	3.32e+03
BT7	3	5	S	58	38	3.06e+02
BT8	2	5	S	11	10	1.00e+00
BT9	2	4	S	9	8	-1.00e+00
BT10	2	2	S	7	6	-1.00e+00
BT11	3	5	S	8	7	8.25e-01
BT12	3	5	S	8	5	6.19e+00
BYRDSPHR	2	3	F^m	602	601	-4.97e+00
COOLHANS	9	9	S	361	320	0.00e+00

Table 10.3: Results for pdSQPcc on equality constrained problems (continued)

Name	m	n	Result	nFun	nQPs	f-value
DIXCHLNG	5	10	F^m	606	601	4.27e+03
EIGENA2	55	110	S	5	3	6.06e-24
EIGENACO	55	110	S	5	3	1.77e-20
EIGENB2	55	110	F^q	3	2	1.80e+01
EIGENBCO	55	110	F^q	3	2	9.00e+00
EIGENC2	231	462	F^q	2	2	3.63e+02
EIGENCCO	231	462	F^q	3	2	2.00e+01
ELEC	200	600	F^q	6	2	2.84e+04
GRIDNETE	36	60	S	5	4	3.96e+01
GRIDNETH	36	60	S	5	4	3.96e+01
HS6	1	2	S	38	17	0.00e+00
HS7	1	2	S	14	8	-1.73e+00
HS8	2	2	S	6	4	-1.00e+00
HS9	1	2	S	5	4	-5.00e-01
HS26	1	3	S	21	18	1.02e-12
HS27	1	3	S	133	75	4.00e-02
HS28	1	3	S	4	2	9.86e-32
HS39	2	4	S	9	8	-1.00e+00
HS40	3	4	S	5	4	-2.50e-01
HS42	2	4	S	8	6	1.39e+01
HS46	2	5	S	24	18	5.07e-12
HS47	3	5	S	23	16	1.92e-10
HS48	2	5	S	4	2	4.93e-32
HS49	2	5	S	18	14	8.43e-10
HS50	3	5	S	14	10	1.42e-22
HS51	3	5	S	4	2	1.21e-14
HS52	3	5	S	4	3	5.33e+00
HS56	4	7	S	6	5	-3.46e+00
HS61	2	3	F^m	602	601	-7.03e+00
HS77	2	5	S	40	23	2.42e-01

Table 10.3: Results for pdSQPcc on equality constrained problems (continued)

Name	m	n	Result	nFun	nQPs	f-value
HS78	3	5	S	6	5	-2.92e+00
HS79	3	5	S	8	7	7.88e-02
HS100LNP	2	7	S	18	7	6.81e+02
HS111LNP	3	10	F^q	3	2	-4.17e+01
LUKVLE1	98	100	S	15	9	5.50e-16
LUKVLE3	2	100	S	18	13	2.76e+01
LUKVLE6	49	99	F^m	606	601	2.23e+06
LUKVLE7	4	100	S	21	11	3.56e+01
LUKVLE8	98	100	F^q	45	45	4.62e+05
LUKVLE9	6	100	S	57	35	1.12e+01
LUKVLE10	98	100	F^q	88	36	4.66e+01
LUKVLE13	64	98	S	28	25	7.90e+02
LUKVLE14	64	98	F^m	602	601	1.06e+06
LUKVLE16	72	97	F^m	604	601	1.62e+05
LCH	1	300	F^q	6	2	5.42e+04
LCH	1	300	F^q	6	2	5.42e+04
MARATOS	1	2	S	4	3	-1.00e+00
MWRIGHT	3	5	S	11	5	2.50e+01
ORTHRDM2	100	203	S	7	5	7.78e+00
ORTHRDS2	100	203	F^q	9	4	1.48e+02
ORTHREGA	64	133	F^q	3	3	1.64e+02
ORTHREGB	6	27	F^q	38	10	9.22e-02
ORTHREGC	10	25	S	7	6	3.99e-01
ORTHREGD	10	23	F^q	49	12	2.18e+01
ORTHRGDM	10	23	F^b	28	4	3.23e+01
ORTHRGDS	76	155	F^b	25	2	3.55e+01
S316-322	1	2	S	9	8	3.34e+02

Table 10.4: Results for pdSQPccnc on equality constrained problems

Name	m	n	Result	nFun	nQPs	f-value
BT1	1	2	S	16	7	-1.00e+00
BT2	1	3	S	13	12	3.26e-02
BT3	3	5	S	4	3	4.09e+00
BT4	2	3	S	71	35	-4.55e+01
BT5	2	3	F^b	25	2	9.56e+02
BT6	2	5	S	12	10	2.77e-01
BT7	3	5	F^b	33	11	1.27e+00
BT8	2	5	S	15	13	1.00e+00
BT9	2	4	S	13	11	-1.00e+00
BT10	2	2	S	7	6	-1.00e+00
BT11	3	5	S	8	7	8.25e-01
BT12	3	5	S	8	5	6.19e+00
BYRDSPHR	2	3	F^b	39	2	-5.00e+00
COOLHANS	9	9	F^b	27	4	0.00e+00
DIXCHLNG	5	10	S	41	40	2.47e+03
EIGENA2	55	110	S	6	5	3.44e-25
EIGENACO	55	110	S	6	5	3.19e-18
EIGENB2	55	110	F^q	2	2	2.64e+01
EIGENBCO	55	110	F^b	24	2	1.60e+01
EIGENC2	231	462	F^q	2	2	8.27e+02
EIGENCCO	231	462	F^q	2	2	7.81e+02
ELEC	200	600	F^q	2	2	4.11e+04
GRIDNETE	36	60	S	5	4	3.96e+01
GRIDNETH	36	60	S	5	4	3.96e+01
HS6	1	2	S	77	36	0.00e+00
HS7	1	2	F^b	24	2	-2.54e+00
HS8	2	2	S	6	4	-1.00e+00
HS9	1	2	S	5	2	-5.00e-01
HS26	1	3	S	18	17	9.31e-13
HS27	1	3	S	18	11	4.00e-02

Table 10.4: Results for pdSQPccnc on equality constrained problems (continued)

Name	m	n	Result	nFun	nQPs	f-value
HS28	1	3	S	3	2	3.94e-31
HS39	2	4	S	13	11	-1.00e+00
HS40	3	4	S	5	4	-2.50e-01
HS42	2	4	S	5	4	1.39e+01
HS46	2	5	S	20	18	1.41e-11
HS47	3	5	F^m	1224	601	1.51e+00
HS48	2	5	S	3	2	2.47e-31
HS49	2	5	S	18	17	1.14e-09
HS50	3	5	S	11	10	2.02e-24
HS51	3	5	S	3	2	1.50e-14
HS52	3	5	S	4	3	5.33e+00
HS56	4	7	F^b	26	2	-3.21e+00
HS61	2	3	F^b	24	2	-9.30e-03
HS77	2	5	S	11	10	2.42e-01
HS78	3	5	S	5	4	-2.92e+00
HS79	3	5	S	44	21	7.88e-02
HS100LNP	2	7	F^m	1233	601	6.87e+02
HS111LNP	3	10	F^d	3	2	-8.50e+01
LUKVLE1	98	100	S	12	11	6.23e+00
LUKVLE3	2	100	S	11	10	2.76e+01
LUKVLE6	49	99	F^m	602	601	2.82e+07
LUKVLE7	4	100	F^q	21	7	1.20e+03
LUKVLE8	98	100	F^q	2	2	5.71e+04
LUKVLE9	6	100	F^b	32	6	3.84e+01
LUKVLE10	98	100	F^b	24	2	1.68e+02
LUKVLE13	64	98	S	18	16	7.90e+02
LUKVLE14	64	98	F^m	602	601	1.58e+06
LUKVLE16	72	97	F^m	602	601	5.90e+02
LCH	1	300	F^q	2	2	3.39e+04
LCH	1	300	F^q	2	2	3.39e+04

Table 10.4: Results for pdSQPccnc on equality constrained problems (continued)

Name	m	n	Result	nFun	nQPs	f-value
MARATOS	1	2	F^b	33	2	-1.12e+00
MWRIGHT	3	5	S	15	8	2.50e+01
ORTHRDM2	100	203	F^q	8	3	3.79e+01
ORTHRDS2	100	203	F^q	2	2	3.34e+00
ORTHREGA	64	133	F^q	4	4	3.02e+02
ORTHREGB	6	27	F^q	25	6	1.79e-01
ORTHREGC	10	25	S	7	6	3.99e-01
ORTHREGD	10	23	S	10	7	3.41e+00
ORTHRGDM	10	23	F^b	28	3	5.56e+00
ORTHRGDS	76	155	F^q	2	2	9.76e+01
S316-322	1	2	S	9	7	3.34e+02

Figure 10.1 gives performance profiles for SNOPT, pdSQP and pdSQPcc. A performance profile provide an "at a glance" comparison of a set of algorithms on a large test set. Let $r_{p,s}$ be the ratio of major iterations for solver s as compared to the best-performing solver on problem p. Let $|\mathcal{A}|$ denote the size of a set. The function,

$$\pi_s(\tau) = \frac{1}{|\mathcal{P}|}|\{p \in \mathcal{P} \mid \log_2(r_{p,s}) \leq \tau\}|,$$

expresses the proportion of problems that are solved in at worst 2^τ iterations times the number of iterations the best solver takes. The performance profile plots $\pi_s(\tau)$ as a function of τ for the different solvers. A solver starting and initially staying at a comparatively a high-value on the vertical axis indicates a fast solver, and a solver that has a comparatively high value on the vertical axis for larger values of τ represents a reliable solver. As can be seen, pdSQP outperforms SNOPT, on average, with respect to speed. However, it is less reliable in terms of being able to solve the most problems in a reasonable number of iterations. This is to be expected, as SNOPT has been tested and maintained for a long period of time.

Below is a summary of results of a set of inequality and equality constrained problems. Overall it was found that for pdSQPcc, post convexification was necessary for 48 (41%) problems.

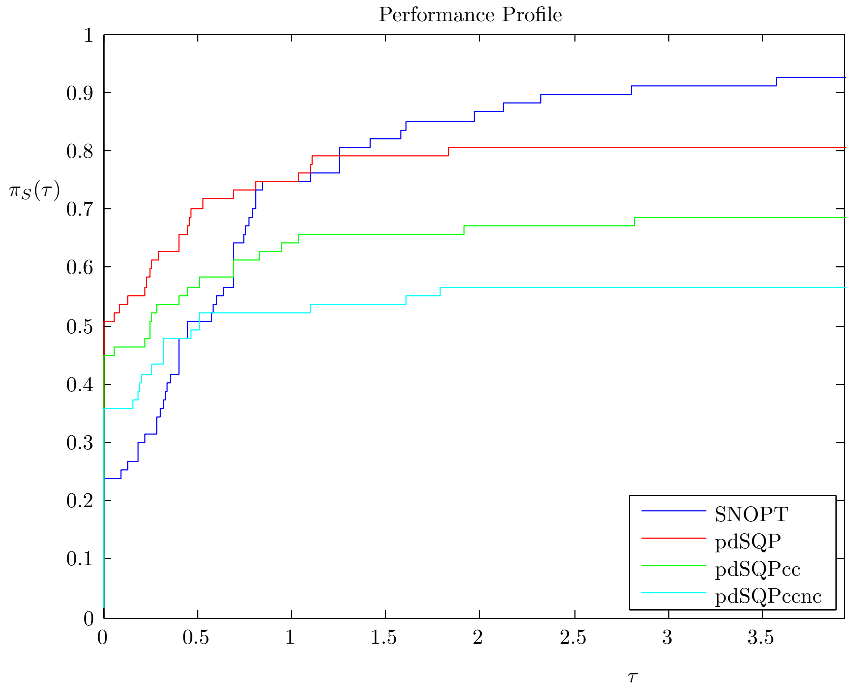

Figure 10.1: Performance profile comparing SNOPT, pdSQP and pdSQPcc

Table 10.5: Reliability of variants of pdSQP

Algorithm	Solved
pdSQP	101
pdSQPcc	84
pdSQPid0	84

Table 10.6: Results for pdSQP on inequality constrained problems

Name	m	n	Result	nFun	nQPs	f-value
HS6	1	2	S	38	17	9.98e-31
HS7	1	2	S	38	18	-1.73e+00
HS8	2	2	S	6	4	-1.00e+00
HS9	1	2	S	5	4	-5.00e-01
HS10	1	2	S	10	9	-1.00e+00
HS11	1	2	S	6	5	-8.50e+00

Table 10.6: Results for pdSQP on inequality constrained problems (continued)

Name	m	n	Result	nFun	nQPs	f-value
HS12	1	2	S	14	7	-3.00e+01
HS13	1	2	S	446	443	9.85e-01
HS14	2	2	S	6	5	1.39e+00
HS15	2	2	F^m	602	601	1.40e+01
HS16	2	2	S	33	16	2.31e+01
HS17	2	2	S	9	8	1.00e+00
HS18	2	2	S	59	30	5.00e+00
HS19	2	2	S	18	17	-6.96e+03
HS20	3	2	S	14	7	4.02e+01
HS21	1	2	S	3	2	-1.00e+02
HS21MOD	1	7	S	3	2	-9.60e+01
HS22	2	2	S	5	4	1.00e+00
HS23	5	2	S	19	11	2.00e+00
HS24	3	2	S	8	7	-1.00e+00
HS26	1	3	S	21	18	1.02e-12
HS27	1	3	S	121	69	4.00e-02
HS28	1	3	S	4	2	9.86e-32
HS29	1	3	S	16	8	-2.26e+01
HS30	1	3	S	12	11	1.00e+00
HS31	1	3	S	10	6	6.00e+00
HS32	2	3	S	11	5	1.00e+00
HS33	2	3	S	8	6	-4.00e+00
HS34	2	3	S	15	7	-8.34e-01
HS35	1	3	S	3	2	1.11e-01
HS35I	1	3	S	3	2	1.11e-01
HS35MOD	1	3	S	3	2	2.50e-01
HS36	1	3	S	4	3	-3.30e+03
HS37	2	3	S	8	6	-3.46e+03
HS39	2	4	S	9	8	-1.00e+00
HS40	3	4	S	5	4	-2.50e-01

Table 10.6: Results for pdSQP on inequality constrained problems (continued)

Name	m	n	Result	nFun	nQPs	f-value
HS41	1	4	S	7	6	1.93e+00
HS42	2	4	S	8	6	1.39e+01
HS43	3	4	S	12	6	-4.40e+01
HS44	6	4	S	6	5	-1.30e+01
HS44NEW	6	4	S	6	5	-1.50e+01
HS46	2	5	S	24	18	5.07e-12
HS47	3	5	S	23	16	1.92e-10
HS48	2	5	S	4	2	2.47e-31
HS49	2	5	S	18	14	8.43e-10
HS50	3	5	S	14	10	1.42e-22
HS51	3	5	S	4	2	1.21e-14
HS52	3	5	S	4	3	5.33e+00
HS53	3	5	S	4	3	4.09e+00
HS54	1	6	$\mathtt{F}^c$	2	2	-7.22e-34
HS55	6	6	S	4	3	6.67e+00
HS56	4	7	S	6	5	-3.46e+00
HS57	1	2	S	7	4	3.06e-02
HS59	3	2	$\mathtt{F}^m$	1206	601	2.37e+01
HS60	1	3	S	6	5	3.26e-02
HS61	2	3	S	19	14	-1.44e+02
HS62	1	3	S	67	66	-2.63e+04
HS63	2	3	$\mathtt{F}^m$	602	601	9.70e+02
HS64	1	3	$\mathtt{F}^m$	603	601	6.22e+03
HS65	1	3	S	20	10	9.54e-01
HS66	2	3	S	13	6	5.18e-01
HS67	14	3	$\mathtt{F}^b$	417	185	-9.39e+02
HS68	2	4	S	128	69	-9.20e-01
HS69	2	4	S	18	11	-9.57e+02
HS70	1	4	S	24	17	7.50e-03
HS71	2	4	S	5	4	1.70e+01

Table 10.6: Results for pdSQP on inequality constrained problems (continued)

Name	m	n	Result	nFun	nQPs	f-value
HS72	2	4	S	426	425	7.28e+02
HS73	3	4	S	5	4	2.99e+01
HS74	5	4	S	14	13	5.13e+03
HS75	5	4	S	269	268	5.17e+03
HS76	3	4	S	3	2	-4.68e+00
HS76I	3	4	S	3	2	-4.68e+00
HS77	2	5	S	18	16	2.42e-01
HS78	3	5	S	6	5	-2.92e+00
HS79	3	5	S	8	7	7.88e-02
HS80	3	5	S	5	4	5.39e-02
HS81	3	5	S	7	6	5.39e-02
HS83	3	5	S	11	10	-3.07e+04
HS84	3	5	F^{m}	604	601	-3.59e+06
HS86	10	5	S	6	4	-3.23e+01
HS88	1	2	S	224	222	1.36e+00
HS89	1	3	S	33	28	1.36e+00
HS90	1	4	S	29	27	1.36e+00
HS91	1	5	F^{m}	610	601	5.83e-01
HS92	1	6	S	27	25	1.36e+00
HS93	2	6	S	31	16	1.35e+02
HS95	4	6	S	11	6	1.56e-02
HS96	4	6	S	11	6	1.56e-02
HS97	4	6	S	5	4	4.07e+00
HS98	4	6	S	5	4	4.07e+00
HS99	2	7	F^{b}	264	247	-8.31e+08
HS100	4	7	S	18	10	6.81e+02
HS100LNP	2	7	S	18	7	6.81e+02
HS100MOD	4	7	S	44	37	6.79e+02
HS101	5	7	S	463	400	1.81e+03
HS102	5	7	S	213	184	9.12e+02

Table 10.6: Results for pdSQP on inequality constrained problems (continued)

Name	m	n	Result	nFun	nQPs	f-value
HS103	5	7	S	111	97	5.44e+02
HS104	5	8	S	28	13	3.95e+00
HS105	1	8	S	9	8	1.06e+03
HS106	6	8	F^m	1210	601	1.35e+04
HS107	6	9	S	55	54	5.06e+03
HS108	13	9	S	12	8	-5.00e-01
HS109	10	9	F^m	602	601	1.15e+03
HS111	3	10	S	25	14	-4.78e+01
HS111LNP	3	10	S	25	14	-4.78e+01
HS112	3	10	S	12	11	-4.78e+01
HS113	8	10	S	9	6	2.43e+01
HS114	11	10	S	58	46	-1.77e+03
HS116	14	13	F^c	80	63	1.66e+02
HS117	5	15	S	20	13	3.23e+01
HS118	17	15	S	4	3	6.65e+02
HS119	8	16	S	34	33	2.45e+02
HS268	5	5	S	13	5	2.91e-11

Table 10.7: Results for pdSQPcc on inequality constrained problems

Name	m	n	Result	nFun	nQPs	f-value
HS6	1	2	S	38	17	0.00e+00
HS7	1	2	S	14	8	-1.73e+00
HS8	2	2	S	6	4	-1.00e+00
HS9	1	2	S	5	4	-5.00e-01
HS10	1	2	S	10	9	-1.00e+00
HS11	1	2	S	6	5	-8.50e+00
HS12	1	2	S	41	35	-3.00e+01
HS13	1	2	F^m	603	601	9.90e-01

Table 10.7: Results for pdSQPcc on inequality constrained problems (continued)

Name	m	n	Result	nFun	nQPs	f-value
HS14	2	2	S	6	5	1.39e+00
HS15	2	2	F^m	602	601	1.57e+02
HS16	2	2	S	5	4	2.31e+01
HS17	2	2	S	13	11	1.00e+00
HS18	2	2	F^m	604	601	6.86e+00
HS19	2	2	S	17	16	-6.96e+03
HS20	3	2	S	14	7	4.02e+01
HS21	1	2	S	248	247	-1.00e+02
HS21MOD	1	7	S	248	247	-9.60e+01
HS22	2	2	S	5	4	1.00e+00
HS23	5	2	S	552	551	2.00e+00
HS24	3	2	S	63	62	-1.00e+00
HS26	1	3	S	21	18	1.02e-12
HS27	1	3	S	133	75	4.00e-02
HS28	1	3	S	4	2	9.86e-32
HS29	1	3	S	88	86	-2.26e+01
HS30	1	3	S	12	11	1.00e+00
HS31	1	3	S	10	6	6.00e+00
HS32	2	3	S	34	32	1.00e+00
HS33	2	3	S	53	52	-4.00e+00
HS34	2	3	S	15	7	-8.34e-01
HS35	1	3	S	3	2	1.11e-01
HS35I	1	3	S	3	2	1.11e-01
HS35MOD	1	3	S	3	2	2.50e-01
HS36	1	3	S	4	3	-3.30e+03
HS37	2	3	S	6	5	-3.46e+03
HS39	2	4	S	9	8	-1.00e+00
HS40	3	4	S	5	4	-2.50e-01
HS41	1	4	S	6	5	1.93e+00
HS42	2	4	S	8	6	1.39e+01

Table 10.7: Results for pdSQPcc on inequality constrained problems (continued)

Name	m	n	Result	nFun	nQPs	f-value
HS43	3	4	S	40	35	-4.40e+01
HS44	6	4	S	41	40	-1.30e+01
HS44NEW	6	4	S	10	9	-1.50e+01
HS46	2	5	S	24	18	5.07e-12
HS47	3	5	S	23	16	1.92e-10
HS48	2	5	S	4	2	4.93e-32
HS49	2	5	S	18	14	8.43e-10
HS50	3	5	S	14	10	1.42e-22
HS51	3	5	S	4	2	1.21e-14
HS52	3	5	S	4	3	5.33e+00
HS53	3	5	S	4	3	4.09e+00
HS54	1	6	S	3	2	-7.22e-34
HS55	6	6	S	4	3	6.67e+00
HS56	4	7	S	6	5	-3.46e+00
HS57	1	2	F^m	610	601	3.06e-02
HS59	3	2	F^m	602	601	2.98e+01
HS60	1	3	S	6	5	3.26e-02
HS61	2	3	F^m	602	601	-7.03e+00
HS62	1	3	S	67	66	-2.63e+04
HS63	2	3	F^m	602	601	9.76e+02
HS64	1	3	F^m	603	601	6.22e+03
HS65	1	3	S	120	110	9.54e-01
HS66	2	3	S	13	6	5.18e-01
HS67	14	3	S	408	331	-1.16e+03
HS68	2	4	S	22	16	-9.20e-01
HS69	2	4	S	18	15	-9.57e+02
HS70	1	4	S	338	334	7.50e-03
HS71	2	4	S	5	4	1.70e+01
HS72	2	4	S	426	425	7.28e+02
HS73	3	4	S	89	88	2.99e+01

Table 10.7: Results for pdSQPcc on inequality constrained problems (continued)

Name	m	n	Result	nFun	nQPs	f-value
HS74	5	4	S	15	14	5.13e+03
HS75	5	4	S	269	268	5.17e+03
HS76	3	4	S	56	55	-4.68e+00
HS76I	3	4	S	56	55	-4.68e+00
HS77	2	5	S	40	23	2.42e-01
HS78	3	5	S	6	5	-2.92e+00
HS79	3	5	S	8	7	7.88e-02
HS80	3	5	S	5	4	5.39e-02
HS81	3	5	S	7	6	5.39e-02
HS83	3	5	S	98	97	-3.07e+04
HS84	3	5	F^m	605	601	-2.63e+06
HS86	10	5	S	13	11	-3.23e+01
HS88	1	2	F^m	603	601	2.83e-04
HS89	1	3	S	35	28	1.36e+00
HS90	1	4	F^q	24	12	7.08e-01
HS91	1	5	F^q	16	8	1.87e-02
HS92	1	6	F^q	6	5	3.37e-01
HS93	2	6	S	59	29	1.35e+02
HS95	4	6	S	91	90	1.56e-02
HS96	4	6	S	91	90	1.56e-02
HS97	4	6	F^m	611	601	3.15e+00
HS98	4	6	F^m	611	601	3.15e+00
HS99	2	7	F^b	264	247	-8.31e+08
HS100	4	7	S	287	277	6.81e+02
HS100LNP	2	7	S	18	7	6.81e+02
HS100MOD	4	7	F^m	608	601	6.79e+02
HS101	5	7	F^m	1204	601	2.18e+03
HS102	5	7	F^m	714	601	1.71e+03
HS103	5	7	F^m	634	601	1.63e+03
HS104	5	8	S	33	15	3.95e+00

Table 10.7: Results for pdSQPcc on inequality constrained problems (continued)

Name	m	n	Result	nFun	nQPs	f-value
HS105	1	8	S	7	6	1.06e+03
HS106	6	8	F^b	387	342	1.44e+04
HS107	6	9	S	55	54	5.06e+03
HS108	13	9	S	18	15	-5.00e-01
HS109	10	9	F^m	602	601	6.87e+02
HS111	3	10	F^m	1215	601	-4.89e+01
HS111LNP	3	10	F^q	3	2	-4.17e+01
HS112	3	10	S	12	11	-4.78e+01
HS113	8	10	F^m	604	601	2.47e+01
HS114	11	10	F^m	801	601	-1.59e+03
HS116	14	13	F^m	1064	601	2.42e+02
HS117	5	15	F^q	2	2	1.22e+03
HS118	17	15	S	59	58	6.65e+02
HS119	8	16	S	42	41	2.45e+02
HS268	5	5	F^m	609	601	1.38e-01

Table 10.8: Results for pdSQPid0 on inequality constrained problems

Name	m	n	Result	nFun	nQPs	f-value
HS6	1	2	S	38	17	0.00e+00
HS7	1	2	S	14	8	-1.73e+00
HS8	2	2	S	6	4	-1.00e+00
HS9	1	2	S	5	4	-5.00e-01
HS10	1	2	S	10	9	-1.00e+00
HS11	1	2	S	6	5	-8.50e+00
HS12	1	2	S	41	35	-3.00e+01
HS13	1	2	S	445	443	9.85e-01
HS14	2	2	S	6	5	1.39e+00
HS15	2	2	F^m	602	601	1.58e+02

Table 10.8: Results for pdSQPid0 on inequality constrained problems (continued)

Name	m	n	Result	nFun	nQPs	f-value
HS16	2	2	S	5	4	2.31e+01
HS17	2	2	S	13	11	1.00e+00
HS18	2	2	F^m	604	601	6.86e+00
HS19	2	2	S	17	16	-6.96e+03
HS20	3	2	S	13	8	4.02e+01
HS21	1	2	F^m	602	601	-1.00e+02
HS21MOD	1	7	F^m	602	601	-9.60e+01
HS22	2	2	S	6	5	1.00e+00
HS23	5	2	S	552	551	2.00e+00
HS24	3	2	S	84	83	-1.00e+00
HS26	1	3	S	21	18	1.02e-12
HS27	1	3	S	133	75	4.00e-02
HS28	1	3	S	4	2	9.86e-32
HS29	1	3	S	88	86	-2.26e+01
HS30	1	3	S	12	11	1.00e+00
HS31	1	3	S	13	8	6.00e+00
HS32	2	3	S	34	32	1.00e+00
HS33	2	3	S	53	52	-4.00e+00
HS34	2	3	S	16	9	-8.34e-01
HS35	1	3	S	3	2	1.11e-01
HS35I	1	3	S	3	2	1.11e-01
HS35MOD	1	3	S	3	2	2.50e-01
HS36	1	3	S	4	3	-3.30e+03
HS37	2	3	S	6	5	-3.46e+03
HS39	2	4	S	9	8	-1.00e+00
HS40	3	4	S	5	4	-2.50e-01
HS41	1	4	S	6	5	1.93e+00
HS42	2	4	S	8	6	1.39e+01
HS43	3	4	S	40	35	-4.40e+01
HS44	6	4	S	48	47	-1.30e+01

Table 10.8: Results for pdSQPid0 on inequality constrained problems (continued)

Name	m	n	Result	nFun	nQPs	f-value
HS44NEW	6	4	S	29	28	-1.30e+01
HS46	2	5	S	24	18	5.07e-12
HS47	3	5	S	23	16	1.92e-10
HS48	2	5	S	4	2	4.93e-32
HS49	2	5	S	18	14	8.43e-10
HS50	3	5	S	14	10	1.42e-22
HS51	3	5	S	4	2	1.21e-14
HS52	3	5	S	4	3	5.33e+00
HS53	3	5	S	4	3	4.09e+00
HS54	1	6	S	3	2	-7.22e-34
HS55	6	6	S	4	3	6.67e+00
HS56	4	7	S	6	5	-3.46e+00
HS57	1	2	F^m	610	601	3.06e-02
HS59	3	2	F^m	602	601	2.99e+01
HS60	1	3	S	6	5	3.26e-02
HS61	2	3	F^m	602	601	-7.03e+00
HS62	1	3	S	68	67	-2.63e+04
HS63	2	3	F^m	602	601	9.76e+02
HS64	1	3	F^m	603	601	6.64e+03
HS65	1	3	S	139	132	9.54e-01
HS66	2	3	F^b	5	4	5.44e-01
HS67	14	3	S	411	337	-1.16e+03
HS68	2	4	S	22	16	-9.20e-01
HS69	2	4	S	18	15	-9.57e+02
HS70	1	4	S	338	334	7.50e-03
HS71	2	4	S	5	4	1.70e+01
HS72	2	4	S	426	425	7.28e+02
HS73	3	4	S	86	85	2.99e+01
HS74	5	4	S	15	14	5.13e+03
HS75	5	4	S	269	268	5.17e+03

Table 10.8: Results for pdSQPid0 on inequality constrained problems (continued)

Name	m	n	Result	nFun	nQPs	f-value
HS76	3	4	S	70	69	-4.68e+00
HS76I	3	4	S	70	69	-4.68e+00
HS77	2	5	S	40	23	2.42e-01
HS78	3	5	S	6	5	-2.92e+00
HS79	3	5	S	8	7	7.88e-02
HS80	3	5	S	5	4	5.39e-02
HS81	3	5	F^b	18	14	5.39e-02
HS83	3	5	S	90	89	-3.07e+04
HS84	3	5	F^m	605	601	-2.63e+06
HS86	10	5	S	13	11	-3.23e+01
HS88	1	2	F^m	603	601	2.83e-04
HS89	1	3	S	33	26	1.36e+00
HS90	1	4	F^q	38	12	4.21e+00
HS91	1	5	F^q	85	72	7.32e-01
HS92	1	6	F^q	12	6	9.95e-01
HS93	2	6	S	71	39	1.35e+02
HS95	4	6	S	91	90	1.56e-02
HS96	4	6	S	91	90	1.56e-02
HS97	4	6	S	448	441	4.07e+00
HS98	4	6	S	448	441	4.07e+00
HS99	2	7	F^m	602	601	-7.77e+08
HS100	4	7	S	287	277	6.81e+02
HS100LNP	2	7	S	18	7	6.81e+02
HS100MOD	4	7	F^m	608	601	6.79e+02
HS101	5	7	F^m	1146	601	2.15e+03
HS102	5	7	F^m	1009	601	1.99e+03
HS103	5	7	F^m	674	601	1.67e+03
HS104	5	8	S	41	23	3.95e+00
HS105	1	8	S	9	8	1.06e+03
HS106	6	8	F^b	306	298	1.47e+04

Table 10.8: Results for pdSQPid0 on inequality constrained problems (continued)

Name	m	n	Result	nFun	nQPs	f-value
HS107	6	9	S	55	54	5.06e+03
HS108	13	9	S	103	68	-5.00e-01
HS109	10	9	F^m	602	601	6.87e+02
HS111	3	10	F^m	1215	601	-4.89e+01
HS111LNP	3	10	F^q	3	2	-4.17e+01
HS112	3	10	S	12	11	-4.78e+01
HS113	8	10	F^m	604	601	2.60e+01
HS114	11	10	S	542	314	-1.77e+03
HS116	14	13	F^m	1162	601	2.50e+02
HS117	5	15	F^q	2	2	1.22e+03
HS118	17	15	S	45	44	6.65e+02
HS119	8	16	S	51	50	2.45e+02
HS268	5	5	F^b	626	595	2.72e-01

Bibliography

[1] R. Andreani, E. G. Birgin, J. M. Martínez, and M. L. Schuverdt. Second-order negative-curvature methods for box-constrained and general constrained optimization. *Comput. Optim. Appl.*, 45(2):209–236, 2010.

[2] R. Andreani, C. E. Echagüe, and M. L. Schuverdt. Constant-rank condition and second-order constraint qualification. *J. Optim. Theory Appl.*, 146(2):255–266, 2010.

[3] R. Andreani, G. Haeser, M. L. Schuverdt, and P. Silva. A relaxed constant positive linear dependence constraint qualification and applications. *Math. Program.*, pages 1–19, 2010. 10.1007/s10107-011-0456-0.

[4] R. Andreani, G. Haeser, M. L. Schuverdt, and P. J. Silva. Two new weak constraint qualifications and applications. *Optimization Online*, July 2011.

[5] R. Andreani, J. M. Martínez, and M. L. Schuverdt. On the relation between constant positive linear dependence condition and quasinormality constraint qualification. *J. Optim. Theory Appl.*, 125(2):473–485, 2005.

[6] R. Andreani, J. M. Martínez, and M. L. Schuverdt. On second-order optimality conditions for nonlinear programming. *Optimization*, 56:529–542, 2007.

[7] A. Baccari. On the classical necessary second-order optimality conditions. *J. Optim. Theory Appl.*, 123(1):213–221, 2004.

[8] A. Baccari and A. Trad. On the classical necessary second-order optimality conditions in the presence of equality and inequality constraints. *SIAM J. Optim.*, 15(2):394–408 (electronic), 2004/05.

[9] D. P. Bertsekas and A. E. Ozdaglar. Pseudonormality and a Lagrange multiplier theory for constrained optimization. *Journal of Optimization Theory and Applications*, 114:287–343, 2002. 10.1023/A:1016083601322.

[10] P. T. Boggs and J. W. Tolle. Sequential quadratic programming. In *Acta Numerica, 1995*, volume 4 of *Acta Numer.*, pages 1–51. Cambridge Univ. Press, Cambridge, 1995.

[11] I. Bongartz, A. R. Conn, N. I. M. Gould, and P. L. Toint. CUTE: Constrained and unconstrained testing environment. *ACM Trans. Math. Software*, 21(1):123–160, 1995.

[12] J. F. Bonnans and A. Shapiro. *Perturbation Analysis of Optimization Problems.* Springer Series in Operations Research. Springer, 2000.

[13] J. R. Bunch and L. Kaufman. A computational method for the indefinite quadratic programming problem. *Linear Algebra Appl.*, 34:341–370, 1980.

[14] J. R. Bunch and B. N. Parlett. Direct methods for solving symmetric indefinite systems of linear equations. *SIAM J. Numer. Anal.*, 8:639–655, 1971.

[15] R. Byrd, J. Nocedal, R. Waltz, and Y. Wu. On the use of piecewise linear models in nonlinear programming. *Math. Program.*, pages 1–36, 2010. 10.1007/s10107-011-0492-9.

[16] R. M. Chamberlain, M. J. D. Powell, C. Lemarechal, and H. C. Pedersen. The watchdog technique for forcing convergence in algorithms for constrained optimization. *Math. Programming Stud.*, (16):1–17, 1982. Algorithms for constrained minimization of smooth nonlinear functions.

[17] C. M. Chin and R. Fletcher. On the global convergence of an SLP-filter algorithm that takes EQP steps. *Math. Program.*, 96(1, Ser. A):161–177, 2003.

[18] A. R. Conn, N. I. M. Gould, and P. L. Toint. Testing a class of method for solving minimization problems with simple bounds on the variables. Preprint, 1986.

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

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

[21] G. Debreu. Definite and semidefinite quadratic forms. *Econometrica*, 20:295–300, 1952.

[22] G. Di Pillo, S. Lucidi, and L. Palagi. Convergence to second order stationary points of a primal-dual algorithm model for nonlinear programming. Technical Report DIS 10-01, Department of Computer and Systems Science, University of Rome "La Sapienza", 2001.

[23] A. L. Dontchev and W. W. Hager. Lipschitzian stability for state constrained nonlinear optimal control. *SIAM J. Control Optim.*, 36(2):698–718, 1998.

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

[25] D. Fernández and M. Solodov. Stabilized sequential quadratic programming for optimization and a stabilized Newton-type method for variational problems. *Math. Program.*, 125(1, Ser. A):47–73, 2010.

[26] A. Fischer. Modified Wilson's method for nonlinear programs with nonunique multipliers. *Math. Oper. Res.*, 24(3):699–727, 1999.

[27] A. Fischer. Local behavior of an iterative framework for generalized equations with nonisolated solutions. *Math. Program.*, 94(1, Ser. A):91–124, 2002. 10.1007/s10107-002-0364-4.

[28] R. Fletcher. Factorizing symmetric indefinite matrices. *Linear Algebra Appl.*, 14:257–272, 1976.

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

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

[31] A. Forsgren, P. E. Gill, and W. Murray. On the identification of local minimizers in inertia-controlling methods for quadratic programming, August 29–September 2 1988. Presented at the 13th International Symposium on Mathematical Programming, Tokyo, Japan.

[32] A. Forsgren, P. E. Gill, and W. Murray. Computing modified Newton directions using a partial Cholesky factorization. Report TRITA-MAT-1993-9, Division of Optimization and Systems Theory, Department of Mathematics, Royal Institute of Technology, 1993.

[33] A. Forsgren, P. E. Gill, and W. Murray. Computing modified Newton directions using a partial Cholesky factorization. *SIAM J. Sci. Comput.*, 16:139–150, 1995.

[34] J. Gauvin. A necessary and sufficient regularity condition to have bounded multipliers in nonconvex programming. *Math. Program.*, 12(1):136–138, 1977.

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

[36] P. E. Gill, W. Murray, M. A. Saunders, and M. H. Wright. Some theoretical properties of an augmented Lagrangian merit function. Report SOL 86-6R, Department of Operations Research, Stanford University, Stanford, CA, 1986.

[37] P. E. Gill, W. Murray, M. A. Saunders, and M. H. Wright. Inertia-controlling methods for general quadratic programming. *SIAM Rev.*, 33(1):1–36, 1991.

[38] P. E. Gill and D. P. Robinson. A primal-dual augmented Lagrangian. *Computational Optimization and Applications*, pages 1–25, 2010.

[39] P. E. Gill and D. P. Robinson. A globally convergent stabilized SQP method. Numerical Analysis Report 12-02, Department of Mathematics, University of California, San Diego, La Jolla, CA, 2012.

[40] P. E. Gill and E. Wong. Methods for convex and general quadratic programming. Numerical Analysis Report 10-1, Department of Mathematics, University of California, San Diego, La Jolla, CA, 2010.

[41] P. E. Gill and E. Wong. Sequential quadratic programming methods. Numerical Analysis Report 10-3, Department of Mathematics, University of California, San Diego, La Jolla, CA, 2010.

[42] P. E. Gill and E. Wong. A regularized method for convex and general quadratic programming. Numerical Analysis Report 11-2, Department of Mathematics, University of California, San Diego, La Jolla, CA, 2011.

[43] P. E. Gill and E. Wong. Sequential quadratic programming methods. In J. Lee and S. Leyffer, editors, *Mixed Integer Nonlinear Programming*, volume 154 of *The IMA Volumes in Mathematics and its Applications*, pages 147–224. Springer New York, 2012. 10.1007/978-1-4614-1927-3_6.

[44] F. J. Gould and J. W. Tolle. A necessary and sufficient qualification for constrained optimization. *SIAM J. Appl. Math.*, 20:164–172, 1971.

[45] N. I. M. Gould. On practical conditions for the existence and uniqueness of solutions to the general equality quadratic programming problem. *Math. Program.*, 32:90–99, 1985.

[46] N. I. M. Gould, D. Orban, and P. L. Toint. CUTEr and SifDec: A constrained and unconstrained testing environment, revisited. *ACM Trans. Math. Software*, 29(4):373–394, 2003.

[47] N. I. M. Gould and D. P. Robinson. A second derivative SQP method: local convergence. Numerical Analysis Report 08/21, Computational Laboratory, University of Oxford, Oxford, UK, 2008.

[48] N. I. M. Gould and D. P. Robinson. A second derivative SQP method: theoretical issues. Numerical Analysis Report 08/18, Computational Laboratory, University of Oxford, Oxford, UK, 2008.

[49] N. I. M. Gould and D. P. Robinson. A second derivative SQP method with imposed descent. Numerical Analysis Report 08/09, Computational Laboratory, University of Oxford, Oxford, UK, 2008.

[50] M. Guignard. Generalized Kuhn-Tucker conditions for mathematical programming problems in a Banach space. *SIAM J. Control*, 7:232–241, 1969.

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

[52] S. P. Han. A globally convergent method for nonlinear programming. *J. Optim. Theory Appl.*, 22:297–309, 1977.

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

[54] A. F. Izmailov and M. V. Solodov. Newton-type methods for optimization problems without constraint qualifications. *SIAM J. on Optimization*, 15(1):210–228, jan 2005.

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

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

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

[58] F. John. Extremum problems with inequalities as side conditions. In K. O. Friedrichs, O. E. Neugebaur, and J. J. Stoker, editors, *Studies and Essays, Courant Anniversary Volume*, pages 187–204, New York, 1948. Wiley (Interscience).

[59] W. Karush. Minima of functions of several variables with inequalities as side constraints. Master's thesis, Department of Mathematics, University of Chicago, 1939.

[60] M. Kojima. Strongly stable stationary solutions in nonlinear programming. In S. M. Robinson, editor, *Analysis and Computation of Fixed Points*, pages 93–138. Academic Press, London, 1980.

[61] H. W. Kuhn and A. W. Tucker. Nonlinear programming. In J. Neyman, editor, *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 481–492, Berkeley, 1951. University of California Press.

[62] 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.

[63] O. L. Mangasarian and S. Fromovitz. The Fritz John necessary optimality conditions in the presence of equality and inequality constraints. *J. Math. Anal. Appl.*, 17:37–47, 1967.

[64] N. Maratos. *Exact Penalty Function Algorithms for Finite-Dimensional and Control Optimization Problems.* PhD thesis, Department of Computing and Control, University of London, 1978.

[65] G. P. McCormick. A modification of Armijo's step-size rule for negative curvature. *Math. Program.*, 13:111–115, 1977.

[66] J. M. Moguerza and F. J. Prieto. An augmented Lagrangian interior-point method using directions of negative curvature. *Math. Program.*, 95(3, Ser. A):573–616, 2003.

[67] J. J. Moré and D. C. Sorensen. Newton's method. In G. H. Golub, editor, *Studies in Mathematics, Volume 24. MAA Studies in Numerical Analysis*, pages 29–82. Math. Assoc. America, Washington, DC, 1984.

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

[69] W. Murray and F. J. Prieto. A sequential quadratic programming algorithm using an incomplete solution of the subproblem. Report SOL 90-12, Department of Operations Research, Stanford University, Stanford, CA, 1990.

[70] W. Murray and F. J. Prieto. A sequential quadratic programming algorithm using an incomplete solution of the subproblem. *SIAM J. Optim.*, 5:590–640, 1995.

[71] J. Nocedal and S. J. Wright. *Numerical Optimization.* Springer-Verlag, New York, 1999.

[72] A. Olivares, J. M. Moguerza, and F. J. Prieto. Nonconvex optimization using negative curvature within a modified linesearch. *European J. Oper. Res.*, 189(3):706–722, 2008.

[73] J. M. Ortega and W. C. Rheinboldt. *Iterative solution of nonlinear equations in several variables.* Academic Press, New York, 1970.

[74] J. M. Ortega and W. C. Rheinboldt. *Iterative solution of nonlinear equations in several variables.* Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2000. Reprint of the 1970 original.

[75] E. R. Panier and A. L. Tits. A superlinearly convergent feasible method for the solution of inequality constrained optimization problems. *SIAM J. Control Optim.*, 25(4):934–950, 1987.

[76] M. J. D. Powell. A method for nonlinear constraints in minimization problems. In R. Fletcher, editor, *Optimization*, pages 283–298, London and New York, 1969. Academic Press.

[77] M. J. D. Powell and Y.-X. Yuan. A recursive quadratic programming algorithm that uses differentiable exact penalty functions. *Math. Program.*, 35(3):265–278, 1986.

[78] L. Qi and Z. Wei. On the constant positive linear dependence condition and its application to SQP methods. *SIAM J. Optim.*, 10(4):963–981 (electronic), 2000.

[79] L. Qi and Z. Wei. Corrigendum: "On the constant positive linear dependence condition and its application to SQP methods" [SIAM J. Optim. **10** (2000), no. 4, 963–981; MR1777075 (2001e:90100)]. *SIAM J. Optim.*, 11(4):1145–1146 (electronic), 2001.

[80] D. P. Robinson. *Primal-Dual Methods for Nonlinear Optimization.* PhD thesis, Department of Mathematics, University of California, San Diego, September 2007.

[81] S. M. Robinson. Perturbed Kuhn-Tucker points and rates of convergence for a class of nonlinear programming algorithms. *Math. Program.*, 7:1–16, 1974.

[82] S. M. Robinson. Strongly regular generalized equations. *Math. Oper. Res.*, 5:43–62, 1980.

[83] S. M. Robinson. Generalized equations and their solutions. II. Applications to nonlinear programming. *Math. Programming Stud.*, 19:200–221, 1982.

[84] M. L. Slater. A note on Motzkins's transposition theorem. *Econometrica*, 19:185–187, 1951.

[85] R. B. Wilson. *A Simplicial Method for Convex Programming.* PhD thesis, Harvard University, 1963.

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

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

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

[89] S. J. Wright. An algorithm for degenerate nonlinear programming with rapid local convergence. *SIAM J. Optim.*, 15(3):673–696, 2005.