

On the Use of Piecewise Linear Models in Nonlinear Programming

Richard H. Byrd * Jorge Nocedal † Richard A. Waltz ‡ Yuchen Wu†

April 29, 2010

Abstract

This paper presents an active-set algorithm for large-scale optimization that occupies the middle ground between sequential quadratic programming (SQP) and sequential linear-quadratic programming (SL-QP) methods. It consists of two phases. The algorithm first minimizes a piecewise linear approximation of the Lagrangian, subject to a linearization of the constraints, to determine a working set. Then, an equality constrained subproblem based on this working set and using second derivative information is solved in order to promote fast convergence. A study of the local and global convergence properties of the algorithm highlights the importance of the placement of the interpolation points that determine the piecewise linear model of the Lagrangian.

1 Introduction

Much research has been devoted in the last 30 years to the development of active-set methods for nonlinear programming. Yet, none of the methods proposed so far is entirely satisfactory for solving very large problems. Sequential quadratic programming (SQP) methods, which are the methods of choice for small and medium size applications, are limited in the large scale case by the significant cost of solving an inequality constrained quadratic program at every iteration. In contrast, interior point methods typically solve only one linear system per iteration and scale up well with the dimension of the problem.

To try to overcome this drawback of SQP methods, Fletcher and Sainz de la Maza [14] proposed in 1987 a two-phase algorithm that uses linear programming techniques to

*Department of Computer Science, University of Colorado, Boulder. This author was supported by National Science Foundation grant CMMI 0728190 and Department of Energy grant DE-SC0001774.

†Department of Electrical Engineering and Computer Science, Northwestern University. These authors were supported by Department of Energy grant DE-FG02-87ER25047-A004 and National Science Foundation grant DMS-0810213.

‡Department of Industrial and Systems Engineering, University of Southern California. This author was supported by National Science Foundation grant CMMI 0728036.

estimate the optimal active set and employs an equality constrained phase to yield rapid convergence. This so-called sequential linear-quadratic programming (SL-QP) method has recently received much attention [2, 3, 8, 10, 15, 21], and numerical experience suggests that it holds promise for large-scale applications.

There are, however, important limitations in using only first-order information to estimate the active set. First, the linear programming subproblem used in the active-set prediction phase of the SL-QP method must include a trust region constraint to ensure boundedness, but the choice of the trust region radius Δ is delicate. If Δ is not of the right magnitude, the active-set identification can be significantly delayed because the trust region constraint is typically active at the solution of the linear program. The delicate role played by the trust region constraint is particularly noticeable in large or degenerate problems. A second drawback of the SL-QP approach is that it is difficult to employ warm-starts when solving the linear programming subproblems (due to the presence of the trust region constraint), and as a result, the computational cost of the active-set prediction phase is not as low as one would expect from a first-order method [2, 23].

Motivated by the promise of SL-QP methods and their limitations, we propose a new method that can be considered to lie in between SQP and SL-QP methods. By using a piecewise linear (as opposed to linear) model in the active-set prediction phase, the new method becomes insensitive to the choice of the trust region constraint — in fact, one can dispense with this constraint altogether. The curvature information contained in the piecewise linear approximation also improves the active-set identification properties of the algorithm, compared with the SL-QP approach, and leads to overall savings in iterations. On the other hand, the new algorithm is computationally less demanding than SQP methods because, instead of solving a quadratic programming subproblem with inequality constraints at every iteration, it solves a linear program.

To define a piecewise linear model, we start by constructing a *separable* quadratic model of the Lagrangian function. For each variable, we then define a set of interpolation points and compute a piecewise-linear interpolant of the quadratic model. The active set prediction of the new method minimizes this separable piecewise-linear model subject to a linearization of the constraints (this problem can be formulated as a linear program). Like the SL-QP method of Fletcher and Sainz de la Maza, an equality constrained quadratic programming (EQP) phase then uses this active-set estimate to compute a step that yields fast local convergence.

The crucial components of the new algorithm are the definition of the separable quadratic model and the placement of the interpolation points that define the piecewise linear model. We present guidelines for the design of these components and show that the new algorithm enjoys global convergence guarantees and fast asymptotic convergence.

In addition to the recent contributions to the development of SL-QP methods already mentioned, there have been a number of recent proposals on how to improve SQP methods by using a simpler quadratic program to predict the active set. Gould and Robinson [17, 18] propose a trust region SQP method that solves two inequality constrained quadratic programs at every iteration. The first quadratic program is convex and provides a predictor step that ensures global convergence; the second quadratic program may be nonconvex

and is capable of incorporating the exact Hessian of the Lagrangian. Friedlander, Gould, Leyffer and Munson [15] describe a two-phase SQP method in which the quadratic term that defines the active-set prediction phase is a multiple of the identity matrix, and this multiple is adjusted to control the size of the step. The two SQP approaches just mentioned are designed for the large scale case. Morales, Nocedal and Wu [19] and Gould and Robinson [17, 18] propose an approach we call SQP+ that is aimed at improving the performance of classical SQP methods that use quasi-Newton approximations. This method accelerates the SQP iteration by adding an EQP phase that employs exact second-order information, and is designed for those problems that are amenable for solution by contemporary SQP methods.

All these new SQP methods have been tested only on prototype implementations and their relative efficiency with respect to SL-QP methods and established SQP methods such as [13, 16] is yet to be explored in production-quality software.

The paper is organized in 7 sections. In Section 2 we describe the proposed method and in Section 3 we consider in detail the construction of the piecewise linear model. The global and local convergence properties of the new algorithm are studied in Sections 4 and 5. Section 6 discusses the use of a quasi-Newton approximation, and Section 7 gives some concluding remarks.

Notation. Throughout the paper $\|\cdot\|$ denotes the 2-norm, unless indicated otherwise.

2 Outline of the Algorithm

The nonlinear programming problem under consideration is stated as

$$\begin{aligned} \min_x \quad & f(x) \\ \text{s.t.} \quad & h(x) = 0, \\ & g(x) \geq 0, \end{aligned} \tag{2.1}$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$, $h : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and $g : \mathbb{R}^n \rightarrow \mathbb{R}^t$ are smooth functions. We write the Lagrangian of this problem as

$$L(x, \lambda, \mu) = f(x) - \lambda^T h(x) - \mu^T g(x), \tag{2.2}$$

where λ and μ are vectors of multipliers. The KKT conditions for the nonlinear program (2.1) are given by

$$\nabla f(x) - \nabla h(x)\lambda - \nabla g(x)\mu = 0 \tag{2.3a}$$

$$h(x) = 0 \tag{2.3b}$$

$$g(x) \geq 0 \tag{2.3c}$$

$$\mu \geq 0 \tag{2.3d}$$

$$\mu^T g(x) = 0, \tag{2.3e}$$

where $\nabla h^T, \nabla g^T$ are the Jacobian matrices of h and g , respectively.

The algorithm proposed in this paper is a two phase method that consists of an active-set prediction phase and an equality constrained phase. At the current iterate x_k , the prediction phase generates a step d_k by solving the problem

$$\min_d m_k(d) \triangleq \nabla f(x_k)^T d + \Gamma_k(d), \quad (2.4a)$$

$$\text{s.t. } h(x_k) + \nabla h(x_k)^T d = 0, \quad (2.4b)$$

$$g(x_k) + \nabla g(x_k)^T d \geq 0, \quad (2.4c)$$

where $\Gamma_k(d)$ is a *convex piecewise linear* function that contains useful curvature information about the nonlinear program (2.1) and yet is simple enough to make the solution of (2.4) tractable in the large-scale case. In this paper, we study two choices for Γ_k , one in which it is a diagonal approximation of the Hessian of the Lagrangian, and one in which Γ_k approximates a quasi-Newton matrix. Due to the piecewise linear nature of Γ_k , problem (2.4) can be formulated as a linear program. We refer to (2.4) as the piecewise linear approximation (PLA) subproblem and denote its solution by d_k .

The KKT conditions of (2.4) at x_k state that there is a vector v_k and Lagrange multipliers λ_{k+1}, μ_{k+1} such that

$$\nabla f(x_k) - \nabla h(x_k)\lambda_{k+1} - \nabla g(x_k)\mu_{k+1} + v_k = 0, \quad (2.5a)$$

$$h(x_k) + \nabla h(x_k)^T d_k = 0, \quad (2.5b)$$

$$g(x_k) + \nabla g(x_k)^T d_k \geq 0, \quad (2.5c)$$

$$\mu_{k+1} \geq 0, \quad (2.5d)$$

$$\mu_{k+1}^T (g(x_k) + \nabla g(x_k)^T d_k) = 0, \quad (2.5e)$$

$$v_k \in \partial\Gamma_k(d_k), \quad (2.5f)$$

where $\partial\Gamma_k(d_k)$ denotes the subdifferential set of $\Gamma_k(d)$ at d_k .

Having computed the PLA step d_k , we perform the second phase of the algorithm by first defining the working set at iteration k as

$$\mathcal{W}_k = \{i \in \mathcal{E}\} \cup \{i \in \mathcal{I} \mid g^i(x_k) + \nabla g^i(x_k)^T d_k = 0\}, \quad (2.6)$$

where \mathcal{E} denotes the set of indices for the equality constraints, \mathcal{I} is the set of indices for the inequality constraints, and $h^i(x_k)$ and $g^i(x_k)$ denote the i -th components of the vectors $h(x_k)$ and $g(x_k)$. The equality constrained quadratic program (EQP) is then given by

$$\min_d (\nabla f(x_k) + W_k d_k)^T d + \frac{1}{2} d^T W_k d \quad (2.7a)$$

$$\text{s.t. } \nabla h^i(x_k)^T d = 0, \quad i \in \mathcal{E}, \quad (2.7b)$$

$$\nabla g^i(x_k)^T d = 0, \quad i \in \mathcal{I} \cap \mathcal{W}_k, \quad (2.7c)$$

where

$$W_k \triangleq \nabla_{xx}^2 L(x_k, \lambda_k, \mu_k) \quad (2.8)$$

is the Hessian of the Lagrangian (2.2) at iteration k . We denote a solution to (2.7) by d_k^E and the corresponding Lagrange multipliers by $\lambda_{k+1}^E, \mu_{k+1}^E$. (Practical procedures for computing an approximate solution of (2.7) are discussed, for example, in [2, 14, 18].) The use of second derivative information in the Hessian W_k endows the algorithm with a fast asymptotic rate of convergence.

The EQP step d_k^E may be shortened as βd_k^E , where $\beta \in [0, 1]$, to ensure that it satisfies the linearized inequalities that are not in the working set, i.e. so that

$$g^i(x_k) + \nabla g^i(x_k)^T(d_k + \beta d_k^E) \geq 0, \quad i \notin \mathcal{W}_k. \quad (2.9)$$

The algorithm promotes global convergence by imposing decrease in the ℓ_1 merit function

$$\phi_\pi(x) = f(x) + \pi \|h(x)\|_1 + \pi \|g(x)^-\|_1, \quad (2.10)$$

where $g(x)^- \triangleq \max\{0, -g(x)\}$ and $\pi > 0$ is a penalty parameter. One can consider trust region or line search implementations of our approach, and in this paper we focus only on the latter. After the steps d_k and d_k^E have been computed, the algorithm performs a backtracking line search along the piecewise linear segment that starts at x_k , goes through $x_k + d_k$ and ends at $x_k + d_k + \beta d_k^E$.

Specifically, the line search first attempts to find a steplength $\hat{\alpha}_k \in (0, 1]$ that satisfies the sufficient decrease condition

$$\phi_{\pi_k}(x_k + d_k + \hat{\alpha}_k \beta d_k^E) \leq \phi_{\pi_k}(x_k) - \sigma qred(d_k), \quad \sigma \in (0, 1), \quad (2.11)$$

where

$$qred(d) = q_\pi(0) - q_\pi(d) \quad (2.12)$$

is the change in the following convex model of the merit function ϕ_π :

$$\begin{aligned} q_\pi(d) &= f(x_k) + \nabla f(x_k)^T d + \Gamma_k(d) + \pi \|h(x_k) + \nabla h(x_k)^T d\|_1 \\ &\quad + \pi \|[g(x_k) + \nabla g(x_k)^T d]^-\|_1. \end{aligned} \quad (2.13)$$

If this line search is successful by making a small number of attempts (say $\hat{\alpha}_k = 1, \frac{1}{2}$ and $\frac{1}{4}$), we define the total step as

$$p_k = d_k + \hat{\alpha}_k \beta d_k^E. \quad (2.14)$$

Otherwise, we choose a constant $\tau \in (0, 1)$ and let $\bar{\alpha}_k \in (0, 1]$ be the first member of the sequence $\{1, \tau, \tau^2, \dots\}$ such that

$$\phi_{\pi_k}(x_k + \bar{\alpha}_k d_k) \leq \phi_{\pi_k}(x_k) - \sigma \bar{\alpha}_k qred(d_k), \quad (2.15)$$

where σ is the same constant as in (2.11); we then set

$$p_k = \bar{\alpha}_k d_k. \quad (2.16)$$

Regardless of whether p_k is defined by (2.14) or (2.16), the new primal iterate is given by

$$x_{k+1} = x_k + p_k. \quad (2.17)$$

The new Lagrange multipliers are defined at the EQP point, as follows. First, we set the multipliers corresponding to the inactive linearized constraints to zero (i.e. the constraints with indices $i \notin W_k$). The rest of the multipliers are set to the EQP multipliers $\lambda_{k+1}^E, \mu_{k+1}^E$ — except that if any multipliers in μ_{k+1}^E are negative, they are set to zero.

We outline the new algorithm using the pseudo-code in Algorithm 2.1.

Algorithm 2.1: PLA Algorithm

Initial data: x_0

For $k = 0, 1, 2, \dots$ until the KKT conditions for the nonlinear program (2.1) are satisfied, perform the following steps.

1. Construct a convex piecewise linear function $\Gamma_k(d)$; see Section 3.
2. Solve the PLA subproblem (2.4) to obtain d_k .
3. Determine the working set \mathcal{W}_k as given in (2.6).
4. Compute the EQP step d_k^E and multipliers $\lambda_{k+1}^E, \mu_{k+1}^E$ by finding a primal-dual solution of problem (2.7).
5. Compute the largest number $\beta \in [0, 1]$ such that $d_k + \beta d_k^E$ satisfies (2.9).
6. Compute the penalty parameter π_k ; see (2.20), (2.21) below.
7. Compute the steplength $\hat{\alpha}_k$ or $\bar{\alpha}_k$, define p_k by (2.14) or (2.16), and set $x_{k+1} = x_k + p_k$.
8. Set

$$[\mu_{k+1}^E]^{\mathcal{W}_k} \leftarrow \max(0, [\mu_{k+1}^E]^{\mathcal{W}_k}), \quad \text{and} \quad [\mu_{k+1}^E]^{\mathcal{W}_k^c} \leftarrow 0,$$

where $[\mu]^{\mathcal{W}_k}$ denotes the subvector of μ corresponding to the components in the working set \mathcal{W}_k , i.e., $[\mu]^{\mathcal{W}_k} = [\mu^i]_{i \in \mathcal{W}_k}$, and \mathcal{W}_k^c denotes the complement of \mathcal{W}_k .

9. Compute $W_{k+1} = \nabla_{xx}^2 L(x_{k+1}, \lambda_{k+1}^E, \mu_{k+1}^E)$.

In practice, various enhancements and modifications are needed to ensure the applicability of Algorithm 2.1 in a wide range of applications. For example, the working set (2.6) must be defined so that the gradients of the constraints in \mathcal{W}_k are linearly independent; this may preclude us from including all active constraints in the working set. Also, the constraints in subproblem (2.4) could be incompatible in which case they would have to be relaxed [6, 9, 16, 18, 22] to ensure that the PLA step is well defined. Since our main interest in this paper is the design of the piecewise linear model and the development of the theoretical foundations for the method, we will assume that these complications do not arise. An extension of Algorithm 2.1 that is capable of handling the difficulties just mentioned, as well as its software implementation, are currently under development [7].

Let us describe the procedure for updating the penalty parameter π_k in Step 6 of Algorithm 2.1. As is now common [5, 20], we require that at every iteration π_k be large enough such that

$$qred(d_k) \geq \rho\pi_k (\|h(x_k)\|_1 + \|g(x_k)^-\|_1), \quad (2.18)$$

for some prescribed constant $\rho \in (0, 1)$. For a step that satisfies the linearized constraints (2.4b)-(2.4c), we have from (2.12) and (2.13) that

$$qred(d_k) = -\nabla f(x_k)^T d_k - \Gamma_k(d_k) + \pi_k (\|h(x_k)\|_1 + \|g(x_k)^-\|_1). \quad (2.19)$$

Therefore, condition (2.18) is equivalent to the requirement

$$\pi_k \geq \frac{\nabla f(x_k)^T d_k + \Gamma_k(d_k)}{(1 - \rho) (\|h(x_k)\|_1 + \|g(x_k)^-\|_1)} \triangleq \pi_k^{\text{trial}}. \quad (2.20)$$

We can enforce this condition by updating the penalty parameter at every iteration k by means of the following rule:

$$\pi_k = \begin{cases} \pi_{k-1} & \text{if } \pi_{k-1} \geq \pi_k^{\text{trial}}, \\ \pi_k^{\text{trial}} + \pi_b & \text{otherwise,} \end{cases} \quad (2.21)$$

where $\pi_b > 0$ is a given constant.

This update strategy ensures that the PLA step d_k is a descent direction for the merit function ϕ_{π_k} . Specifically, the directional derivative of ϕ_{π_k} at a point x_k along the direction d_k , denoted as $D\phi_{\pi_k}(x_k; d_k)$, satisfies

$$D\phi_{\pi_k}(x_k; d_k) \leq \nabla f(x_k)^T d_k - \pi_k \|h(x_k)\|_1 - \pi_k \|g(x_k)^-\|_1.$$

Comparing the right hand side of this expression with (2.19), we obtain

$$D\phi_{\pi_k}(x_k; d_k) \leq -qred(d_k) - \Gamma_k(d_k). \quad (2.22)$$

By noting that $qred(d_k) > 0$ when x_k is not a KKT point of the nonlinear program (2.1), and that $\Gamma_k(d_k) \geq 0$ (by construction), we conclude that d_k is a descent direction for $\phi_{\pi_k}(x)$ at x_k . Therefore, the line search terminates in a finite number of steps and the algorithm is well defined.

3 The Piecewise Linear Model $\Gamma(d)$

Recall that the objective of the PLA subproblem (2.4) is given by

$$m_k(d) = \nabla f(x_k)^T d + \Gamma_k(d), \quad (3.1)$$

where Γ_k is a nonnegative convex and piecewise linear function. The specific form of Γ_k plays a crucial role in the algorithm. Suppose, for example, that we define

$$\Gamma_k(d) = \begin{cases} 0 & \|d\|_\infty \leq \Delta_k, \\ \infty & \text{otherwise,} \end{cases} \quad (3.2)$$

where Δ_k is a trust region radius that is updated at each iteration and $\|\cdot\|$ is some norm. For this choice of Γ_k , the algorithm reduces to an SL-QP method since the active set prediction problem (2.4) can be formulated as

$$\min_d \quad \nabla f(x_k)^T d \tag{3.3a}$$

$$\text{s.t.} \quad h(x_k) + \nabla h(x_k)^T d = 0, \tag{3.3b}$$

$$g(x_k) + \nabla g(x_k)^T d \geq 0, \tag{3.3c}$$

$$\|d\|_\infty \leq \Delta_k. \tag{3.3d}$$

This simple form of Γ_k does not, however, contain curvature information about the problem and leads to the difficulties in choosing Δ_k discussed in Section 1.

Therefore, in this paper we let Γ_k be a piecewise linear approximation of a quadratic model:

$$\Gamma_k(d) \approx \frac{1}{2} d^T B_k d.$$

Ideally, B_k would be defined as the Hessian of the Lagrangian (2.8), but this Hessian may not be positive definite, which could cause the PLA subproblem to be nonconvex — a situation we wish to avoid. Therefore, we choose B_k to be a symmetric positive definite matrix, and for our approach to be practical, B_k must have a simple structure. In the next subsection, we consider the case when B_k is a diagonal matrix, and in Section 6 we discuss the option of defining it through quasi-Newton updating.

3.1 Diagonal Hessian Approximation B_k

Constructing a piecewise linear approximation of a high dimensional quadratic function $d^T B_k d$ is not simple. In particular, the accuracy of the approximation degrades with dimensionality. Fortunately, in the case where $d^T B_k d$ is *separable*, we can write $d^T B_k d$ as the sum of n one-dimensional quadratic functions, each of which can easily be approximated by a univariate piecewise linear function.

One way to achieve separability is to define B_k to be a diagonal matrix, i.e.,

$$B_k = \text{diag}(b_k^i), \tag{3.4}$$

so that

$$\frac{1}{2} d^T B_k d = \sum_{i=1}^n \frac{1}{2} b_k^i (d^i)^2, \tag{3.5}$$

where d^i denotes the i -th component of the vector d . We approximate each one-dimensional function $\frac{1}{2} b_k^i (d^i)^2$ by a piecewise linear function $\Gamma_k^i(d^i)$ composed of $r_k^i + 1$ segments, where r_k^i is a small integer (say, less than 10). Thus,

$$\Gamma_k^i(d^i) \approx \frac{1}{2} b_k^i (d^i)^2.$$

The linear functions forming $\Gamma_k^i(d^i)$ are denoted by $\ell_{k,j}^i(d^i)$, and we define

$$\Gamma_k^i(d^i) = \max_{j \in J_k^i} \{\ell_{k,j}^i(d^i)\}, \quad \text{with} \quad J_k^i \triangleq \{0, \dots, r_k^i\}. \tag{3.6}$$

The component functions $\ell_{k,j}^i$ are chosen by Hermite interpolation at a set of designated nodes, which we denote as

$$t_{k,j}^i, \quad j \in J_k^i. \quad (3.7)$$

Specifically, the j -th linear function $\ell_{k,j}^i$ interpolates $\frac{1}{2}b_k^i(d^i)^2$ and its derivative at the node $t_{k,j}^i$, which implies that $\ell_{k,j}^i$ is given by

$$\ell_{k,j}^i(d^i) = -\frac{1}{2}b_k^i(t_{k,j}^i)^2 + (b_k^i t_{k,j}^i)d^i. \quad (3.8)$$

To ensure that $\Gamma_k(d)$ is nonnegative, we choose one interpolation point at the origin, i.e.,

$$t_{k,0}^i = 0 \quad \text{for all } i \in \{1, 2, \dots, n\}. \quad (3.9)$$

This condition, and the fact that we are employing Hermite interpolation, means that each function Γ_k^i is non-negative and has the form given in Figure 1.

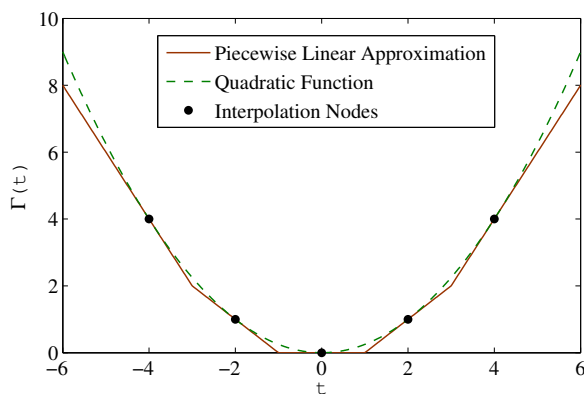


Figure 1: The univariate function Γ_k^i

The multivariate function Γ_k in (2.4a) is then defined as

$$\Gamma_k(d) \triangleq \sum_{i=1}^n \Gamma_k^i(d^i), \quad (3.10)$$

and constitutes an outer approximation of the quadratic form, i.e., $\Gamma_k(d) \leq \frac{1}{2}d^T B_k d$. Note that, by construction, $\Gamma_k(d) \geq 0$ for all d .

One more requirement we impose on Γ_k is that it must guarantee that the objective of the PLA subproblem (2.4) is bounded. (We will assume henceforth that this subproblem is feasible.) We can ensure that the PLA steps are bounded by constructing the model m_k so that

$$\lim_{\alpha \rightarrow \infty} m_k(\alpha d) = +\infty \quad \text{for all } d \neq 0. \quad (3.11)$$

We can achieve this behaviour by an appropriate placement of the interpolation points.

To see this, note from (3.1) and (3.10) that

$$m_k(d) = \sum_{i=1}^n m_k^i(d^i), \quad \text{where} \quad m_k^i(d^i) = \frac{\partial f(x_k)}{\partial x^i} d^i + \Gamma_k^i(d^i). \quad (3.12)$$

Now, given a coordinate i and any $j \in J_k^i$, we have from (3.6) and (3.8) that

$$m_k^i(d^i) \geq -\frac{1}{2}b_k^i(t_{k,j}^i)^2 + \left[b_k^i t_{k,j}^i + \frac{\partial f(x_k)}{\partial x^i} \right] d^i. \quad (3.13)$$

In other words, each interpolation point generates a linear function (given by the right hand side of (3.13)) that bounds $m_k^i(d^i)$ from below. If we make sure that one of these lower bounding functions has a positive slope and another one has a negative slope, then m_k^i will have property (3.11). Since the term inside the square brackets in (3.13) gives the slope of the lower bounding function, we can achieve these goals by selecting one interpolation point, say $t_{k,u}^i$, so that

$$b_k^i t_{k,u}^i + \frac{\partial f(x_k)}{\partial x^i} > 0, \quad (3.14)$$

and choosing another interpolation point, say $t_{k,l}^i$, so that

$$b_k^i t_{k,l}^i + \frac{\partial f(x_k)}{\partial x^i} < 0. \quad (3.15)$$

Clearly, since for each coordinate direction $i \in \{1, 2, \dots, n\}$ the function m_k^i constructed in this manner satisfies $\lim_{|d^i| \rightarrow \infty} m_k^i(d^i) = +\infty$, it follows from the convexity of (3.12) that m_k satisfies condition (3.11). Thus, by choosing a pair of interpolation points satisfying (3.14)–(3.15) along each coordinate i , the PLA subproblem (2.4) will have a bounded solution.

We summarize the observations made in this subsection as follows.

Procedure 3.1: Construction of Piecewise Linear Function Γ_k

Initial data: (i) k , the iteration number; (ii) r_k^i , $i = 1, \dots, n$, the number of interpolation points for each coordinate direction i .

1. Define a positive definite diagonal matrix $B_k = \text{diag}(b_k^i)$.
2. For each $i = 1, \dots, n$, define the interpolation points $t_{k,j}^i$, $j \in J_k^i = \{0, \dots, r_k^i\}$ so that one of these points is given by (3.9), and so that the PLA problem (2.4) is bounded.
3. For each $i = 1, \dots, n$, define the component functions $\ell_{k,j}^i$, $j \in J_k^i$, by (3.8).
4. Define Γ_k as in (3.6).

In Sections 4 and 5, we impose a few more requirements on the placement of the interpolation points (in addition to (3.14), (3.15)) to ensure that the algorithm has the desired global and local convergence properties. Practical choices for the diagonal matrix B_k are discussed in [7].

Given the quadratic form $\frac{1}{2}d^T B_k d$ and the set of nodes (3.7), the Hermite interpolation process given in Procedure 3.1 uniquely determines the function Γ_k . There are, however, various ways of expressing this function algebraically. In (3.6) we defined each Γ_k^i to be the maximum of the component functions $\ell_{k,j}^i$, and in this case the PLA subproblem can be formulated as a linear program that includes additional constraints to account for the max functions; see e.g. [1]. An alternative representation of the piecewise linear model is obtained by first computing (for each i) the r_k^i breakpoints of Γ_k^i , and defining intervals $S_{k,j}^i$ whose endpoints are given by these breakpoints, see e.g. [12]. To evaluate Γ_k^i at d^i , we identify the interval to which d^i belongs and evaluate the corresponding linear function, i.e., $d^i \in S_{k,j}^i \Rightarrow \Gamma_k^i(d^i) = \ell_{k,j}^i(d^i)$. In this representation, there is no need to introduce additional constraints, but we must increase the number of variables. The computational advantages of these two alternative representations of Γ_k are discussed in [7].

4 Global Convergence Analysis

In this section we analyze the global convergence properties of Algorithm 2.1. We make the following assumptions about the problem and the algorithm. We recall that the PLA multipliers λ_k, μ_k are defined in (2.5).

Assumptions I

- a) The sequence $\{x_k\}$ generated by Algorithm 2.1 is contained in a convex set Ω where the functions f , h and g and their first and second derivatives are bounded.
- b) The PLA subproblem (2.4) is feasible for all k .
- c) The sequence of PLA multipliers is bounded, i.e., there exists a constant γ_m such that $\|(\lambda_k, \mu_k)\|_\infty \leq \gamma_m$ for all k .
- d) There exist constants $\beta_l > 0$ and $\beta_u > 0$ such that, for all $d \in \mathbb{R}^n$ and all k ,

$$\beta_l \|d\|^2 \leq d^T B_k d \leq \beta_u \|d\|^2. \quad (4.1)$$

In Procedure 3.1, we outlined a strategy for constructing Γ_k based on Hermite interpolation of a separable model with diagonal Hessian B_k (and showed how to ensure boundedness of the PLA subproblem by an appropriate placement of the interpolation points). Other choices for B_k and other forms of interpolation could, however, be employed and we would like for our analysis to be applicable to a wide class of piecewise linear models. In particular, in Section 6 we extend our analysis to the case when B_k is defined through BFGS updating, and is not a diagonal matrix. Therefore, we present the global convergence analysis

in a general setting that includes as a special case the procedure discussed in the previous section.

The first step in our analysis is to identify the essential properties that the model Γ_k must possess for Algorithm 2.1 to enjoy global convergence guarantees. In essence, we need to ensure that m_k grows at least linearly for large d and that Γ_k does not grow faster than a quadratic function. We state these properties as follows.

Growth Conditions

The following 3 conditions are satisfied for all k .

P1) $\Gamma_k(d)$ is convex, non-negative, and satisfies

$$\Gamma_k(0) = 0. \tag{4.2}$$

P2) There exists a constant $\beta_u > 0$, such that

$$\Gamma_k(d) \leq \beta_u \|d\|^2, \quad \text{for all } d. \tag{4.3}$$

P3) There exist constants $\Delta_G > 0$, $\varsigma > 0$ such that

$$m_k(d) \geq \varsigma \|d\|, \quad \text{for all } \|d\| \geq \Delta_G. \tag{4.4}$$

The algorithm does not have direct control over the properties of Γ_k , but given that the matrices B_k are uniformly positive definite and bounded, the algorithm can enforce the growth conditions by an appropriate placement of the interpolation points. Specifically, for the Hermite interpolation procedure described in the previous section we can impose the following requirements, which are a reinforcement of conditions (3.14)-(3.15) introduced in Section 3.1 to ensure boundedness of the PLA step.

Placement Conditions A

Given a constant $\kappa_0 > 0$, for each iteration k and for each coordinate direction i , at least two interpolation points $t_{k,l}^i, t_{k,u}^i$ satisfy

$$t_{k,l}^i \leq -\frac{\partial f(x_k)}{\partial x^i} / b_k^i - \kappa_0, \quad \text{and} \quad t_{k,u}^i \geq -\frac{\partial f(x_k)}{\partial x^i} / b_k^i + \kappa_0. \tag{4.5}$$

Furthermore, the sequences $\{t_{k,l}^i\}$ and $\{t_{k,u}^i\}$ are uniformly bounded.

It follows from Assumptions I-a,d that (4.5) can be satisfied by a bounded sequence of interpolation points. Note that Placement Conditions A impose minimal requirements and allow much freedom in the placement of the interpolation points. The following result shows that if these conditions are satisfied, the piecewise linear model generated by the Hermite interpolation procedure of the previous section has the desired behavior.

Theorem 4.1 *Suppose that Assumptions I hold, that Γ_k is defined by Procedure 3.1, and that the interpolation points satisfy Placement Conditions A. Then, the function $\Gamma_k(d)$ and the model $m_k(d)$ satisfy the Growth Conditions.*

Proof. Since $\Gamma_k(d)$ is constructed by Hermite interpolation of the convex quadratic function $\frac{1}{2}d^T B_k d$ and one interpolation point is placed at the origin (see Step 2 of Procedure 3.1), we immediately obtain the Growth Condition P1.

Because Hermite interpolation builds an underestimate of the function $\frac{1}{2}d^T B_k d$, we have

$$\Gamma_k(d) \leq \frac{1}{2}d^T B_k d \leq \beta_u \|d\|^2,$$

where the second inequality follows from (4.1). Hence, condition P2 is satisfied.

To establish P3, note that since the sequences $\{t_{k,u}^i\}$, $\{t_{k,l}^i\}$ are bounded, it follows from (3.6) and (3.8) that

$$\Gamma_k^i(d^i) + \frac{\partial f(x_k)}{\partial x^i} d^i \geq -\frac{1}{2}b_k^i (t_{k,u}^i)^2 + (b_k^i t_{k,u}^i + \frac{\partial f(x_k)}{\partial x^i}) d^i \quad (4.6)$$

$$\geq \kappa_2 + (b_k^i t_{k,u}^i + \frac{\partial f(x_k)}{\partial x^i}) d^i \quad (4.7)$$

and

$$\Gamma_k^i(d^i) + \frac{\partial f(x_k)}{\partial x^i} d^i \geq -\frac{1}{2}b_k^i (t_{k,l}^i)^2 + (b_k^i t_{k,l}^i + \frac{\partial f(x_k)}{\partial x^i}) d^i$$

$$\geq \kappa_2 + (b_k^i t_{k,l}^i + \frac{\partial f(x_k)}{\partial x^i}) d^i, \quad (4.8)$$

where $\kappa_2 = -\frac{1}{2}\kappa_1^2\beta_u$ and κ_1 is a constant such that $-\kappa_1 \leq t_{k,l}^i \leq t_{k,u}^i \leq \kappa_1$, for all k and all i . Now, if $d^i \geq 0$, we have from the second inequality in (4.1), (4.5) and (4.7) that

$$\Gamma_k^i(d^i) + \frac{\partial f(x_k)}{\partial x^i} d^i \geq \kappa_2 + \kappa_0 \beta_l d^i. \quad (4.9)$$

Similarly, if $d^i < 0$ we have that

$$\Gamma_k^i(d^i) + \frac{\partial f(x_k)}{\partial x^i} d^i \geq \kappa_2 - \kappa_0 \beta_l d^i. \quad (4.10)$$

Together, (4.9) and (4.10) imply

$$\Gamma_k^i(d^i) + \frac{\partial f(x_k)}{\partial x^i} d^i \geq \kappa_2 + \kappa_0 \beta_l |d^i|,$$

for any d^i . Recalling (3.1) and (3.10), and summing over all coordinates i , yields

$$m_k(d) = \nabla f(x_k)^T d + \Gamma_k(d) \geq n\kappa_2 + \kappa_0 \beta_l \|d\|_1 \geq n\kappa_2 + \kappa_0 \beta_l \|d\|, \quad (4.11)$$

which implies (4.4) with $\varsigma = \kappa_0 \beta_l / 2$, and $\Delta_G = -2n\kappa_2 / \kappa_0 \beta_l$. \square

Having shown that there is a practical way of ensuring the Growth Conditions, we assume for the rest of this section that Γ_k is constructed at every iteration to satisfy these conditions.

For reference below, we note that since Γ_k is convex with $\Gamma_k(0) = 0$, and if v denotes an element of its subdifferential at d , we have $\Gamma_k(0) \geq \Gamma_k(d) + v^T(0 - d)$, and hence

$$\Gamma_k(d) \leq v^T d. \quad (4.12)$$

We first establish that the PLA steps d_k are bounded.

Lemma 4.2 *If Assumptions I and the Growth Conditions hold, the PLA subproblem is well defined and there exists a constant $L_0 > 0$ such that*

$$\|d_k\| \leq L_0, \quad \text{for all } k.$$

Proof. We first note that since the step d_k solves the PLA problem (2.4), it satisfies the KKT conditions (2.5), and from (2.5a), (2.5b) and (2.5e) we have that

$$\nabla f(x_k)^T d_k + v_k^T d_k = -\lambda_{k+1}^T h(x_k) - \mu_{k+1}^T g(x_k). \quad (4.13)$$

Recalling (4.12), we have

$$\nabla f(x_k)^T d_k + \Gamma_k(d_k) \leq -\lambda_{k+1}^T h(x_k) - \mu_{k+1}^T g(x_k).$$

Due to the boundedness of $h(x_k)$ and $g(x_k)$ in Ω , and the boundedness of the multipliers, there exists a constant κ_4 such that

$$m_k(d_k) = \Gamma_k(d_k) + \nabla f(x_k)^T d_k \leq \kappa_4. \quad (4.14)$$

By combining this bound with (4.4), we obtain

$$\|d_k\| \leq \min\{\kappa_4/\varsigma, \Delta_G\} \triangleq L_0. \quad (4.15)$$

□

We can now show that the penalty parameter is bounded.

Lemma 4.3 *If Assumptions I hold, there exists an integer \bar{k} and a positive scalar $\bar{\pi}$ such that for all $k \geq \bar{k}$, the sequence $\{\pi_k\}$ generated by Algorithm 2.1 satisfies $\pi_k = \bar{\pi}$.*

Proof. From (2.5c) and (2.5e) we have both

$$\lambda_{k+1}^T \nabla h(x_k)^T d_k = -\lambda_{k+1}^T h(x_k) \quad \text{and} \quad \mu_{k+1}^T \nabla g(x_k)^T d_k = -\mu_{k+1}^T g(x_k). \quad (4.16)$$

We also have that

$$-\lambda_{k+1}^T h(x_k) \leq \|\lambda_{k+1}\|_\infty \|h(x_k)\|_1,$$

and since $\mu_{k+1} \geq 0$,

$$-\mu_{k+1}^T g(x_k) \leq \mu_{k+1}^T g(x_k)^- \leq \|\mu_{k+1}\|_\infty \|g(x_k)^-\|_1,$$

where $g(x_k)^- = \max\{0, -g(x_k)\}$. Substituting these relations in (2.5a) we obtain

$$\nabla f(x_k)^T d_k + v_k^T d_k \leq \|(\lambda_{k+1}, \mu_{k+1})\|_\infty (\|h(x_k)\|_1 + \|g(x_k)^-\|_1),$$

or equivalently,

$$\frac{\nabla f(x_k)^T d_k + v_k^T d_k}{\|h(x_k)\|_1 + \|g(x_k)^-\|_1} \leq \|(\lambda_{k+1}, \mu_{k+1})\|_\infty \leq \gamma_m, \quad (4.17)$$

where γ_m is defined in Assumptions I-c. By (2.5f), the vector v_k is an element of the subdifferential set of $\Gamma_k(d_k)$, and therefore by combining (4.12) and (4.17), and recalling (2.20), we obtain

$$\pi_k^{\text{trial}} = \frac{\nabla f(x_k)^T d_k + \Gamma_k(d_k)}{(1 - \rho)(\|h(x_k)\|_1 + \|[g(x_k)]^-\|_1)} \leq \frac{\gamma_m}{1 - \rho}. \quad (4.18)$$

Consequently, by (2.21), we have that for all k

$$\pi_k \leq \max\{\pi_0, \frac{\gamma_m}{1 - \rho} + \pi_b\},$$

which shows that π_k is bounded above. Since the sequence $\{\pi_k\}$ is nondecreasing and bounded, and when it increases it does so by at least π_b , we conclude that π_k must be constant after a finite number of iterations, and that this constant value satisfies $\bar{\pi} \leq \max\{\pi_0, \frac{\gamma_m}{1 - \rho} + \pi_b\}$. \square

The following technical result will allow us to prove that the algorithm satisfies the KKT conditions of (2.1) in the limit.

Lemma 4.4 *If Assumptions I and the Growth Conditions hold, Algorithm 2.1 yields the limit*

$$\lim_{k \rightarrow \infty} \text{qred}(d_k) = 0. \quad (4.19)$$

Proof. Let us assume that the iteration indices k are large enough such that $\pi_k = \bar{\pi}$, where $\bar{\pi}$ is given in Lemma 4.3.

By (2.18), we have $\text{qred}(d_k) \geq 0$ for all k and therefore by (2.11) and (2.15) the sequence $\{\phi_{\bar{\pi}}(x_k)\}$ is monotonically decreasing. Since $\phi_{\bar{\pi}}(x_k)$ is bounded from below on Ω by Assumptions I, the sequence $\{\phi_{\bar{\pi}}(x_k)\}$ must converge, i.e.,

$$\lim_{k \rightarrow \infty} (\phi_{\bar{\pi}}(x_{k+1}) - \phi_{\bar{\pi}}(x_k)) = 0.$$

Given that (2.11) or (2.15) hold at each iteration, this limit implies

$$\lim_{k \rightarrow \infty} \alpha_k \text{qred}(d_k) = 0, \quad (4.20)$$

where $\alpha_k = \hat{\alpha}_k$ if the line search enforced condition (2.11), and $\alpha_k = \bar{\alpha}_k$ if it enforced condition (2.15). Thus, if we can show that $\bar{\alpha}_k$ is bounded away from zero for all k , then (4.20) implies the limit (4.19).

Recalling the definition (2.13), Assumptions I, the condition $\Gamma_k(d) \geq 0$, and Lemma 4.2 we have, for any α ,

$$\begin{aligned} \phi_{\bar{\pi}}(x_k + \alpha d_k) - q_{\bar{\pi}}(\alpha d_k) &\leq f(x_k + \alpha d_k) - (f(x_k) + \alpha \nabla f(x_k)^T d_k) \\ &\quad + \bar{\pi}(\|h(x_k + \alpha d_k)\|_1 - \|h(x_k) + \alpha \nabla h(x_k) d_k\|_1) \\ &\quad + \bar{\pi}(\|[g(x_k + \alpha d_k)]^-\|_1 - \|[g(x_k) + \alpha \nabla g(x_k) d_k]^-\|_1) \quad (4.21) \\ &\leq L_1 \alpha^2 \|d_k\|_2^2 \\ &\leq L_0^2 L_1 \alpha^2, \end{aligned}$$

for some constant $L_1 > 0$. In addition, using the convexity of $q_{\bar{\pi}}(\cdot)$ and recalling (2.12), we have

$$q_{\bar{\pi}}(\alpha d_k) - q_{\bar{\pi}}(0) \leq \alpha(q_{\bar{\pi}}(d_k) - q_{\bar{\pi}}(0)) = -\alpha qred(d_k). \quad (4.22)$$

By adding (4.21) and (4.22), and noting that $\phi_{\bar{\pi}}(x_k) = q_{\bar{\pi}}(0)$, we get

$$\phi_{\bar{\pi}}(x_k + \alpha d_k) - \phi_{\bar{\pi}}(x_k) \leq L_0^2 L_1 \alpha^2 - \alpha qred(d_k). \quad (4.23)$$

Now, suppose that $\bar{\alpha}_k$ in (2.15) is less than 1. Since the line search algorithm chooses $\bar{\alpha}_k$ as the first element in the sequence $\{1, \tau, \tau^2, \dots\}$ that satisfies (2.15), this implies that (2.15) was violated for the value $\bar{\alpha}_k/\tau$:

$$\phi_{\bar{\pi}}(x_k + (\bar{\alpha}_k/\tau)d_k) - \phi_{\bar{\pi}}(x_k) > -\sigma(\bar{\alpha}_k/\tau)qred(d_k).$$

Combining this inequality with (4.23) (with $\alpha = \bar{\alpha}_k/\tau$) we have

$$L_0^2 L_1 (\bar{\alpha}_k/\tau)^2 - (\bar{\alpha}_k/\tau)qred(d_k) > -\sigma(\bar{\alpha}_k/\tau)qred(d_k).$$

This implies

$$\bar{\alpha}_k > \frac{\tau(1-\sigma)}{L_0^2 L_1} qred(d_k). \quad (4.24)$$

Therefore, since in (4.20) either $\alpha_k = \bar{\alpha}_k$ or α_k satisfies the bound (4.24), we have

$$0 = \lim_{k \rightarrow \infty} \alpha_k qred(d_k) \geq \lim_{k \rightarrow \infty} \min \left\{ 1, \frac{\tau(1-\sigma)}{L_0^2 L_1} qred(d_k) \right\} qred(d_k) \geq 0$$

from which we immediately get (4.19). \square

Corollary 4.5 *The sequence $\{x_k\}$ generated by Algorithm 2.1 is asymptotically feasible, i.e.,*

$$\lim_{k \rightarrow \infty} (\|h(x_k)\|_1 + \|g(x_k)^-\|_1) = 0.$$

Proof. The result follows immediately from (2.18) and the previous lemma. \square

Lemma 4.6 *The sequence $\{x_k\}$ generated by Algorithm 2.1 satisfies*

$$\lim_{k \rightarrow \infty} g(x_k)^T \mu_{k+1} = 0 \quad (4.25)$$

and

$$\lim_{k \rightarrow \infty} \Gamma_k(d_k) - v_k^T d_k = 0. \quad (4.26)$$

Proof. By (2.19) and Lemmas 4.3–4.5, we have

$$\lim_{k \rightarrow \infty} \nabla f(x_k)^T d_k + \Gamma_k(d_k) = 0. \quad (4.27)$$

Equivalently, for any $\epsilon > 0$ there exists $k_0 > 0$ such that for $k > k_0$, we have

$$-\epsilon \leq \nabla f(x_k)^T d_k + \Gamma_k(d_k) \leq \epsilon.$$

Since $v_k \in \partial\Gamma(d_k)$, we further have by (4.12) that

$$-\epsilon \leq \nabla f(x_k)^T d_k + v_k^T d_k.$$

Combining this bound with (4.13), we obtain

$$-\epsilon \leq -h(x_k)^T \lambda_{k+1} - g(x_k)^T \mu_{k+1} \leq (\|h(x_k)\|_1 + \|g(x_k)^-\|_1) \gamma_m,$$

where γ_m is given in Assumptions I-c. By Corollary 4.5, there exists k_1 such that if $k \geq k_1$, we have $\|h(x_k)\|_1 + \|g(x_k)^-\|_1 \leq \epsilon/\gamma_m$. Therefore, for $k \geq \max\{k_0, k_1\}$, we have

$$-\epsilon \leq -[h(x_k)^T \lambda_{k+1} + g(x_k)^T \mu_{k+1}] \leq \epsilon,$$

and hence we have established that

$$\lim_{k \rightarrow \infty} h(x_k)^T \lambda_{k+1} + g(x_k)^T \mu_{k+1} = 0. \quad (4.28)$$

According to Corollary 4.5, we have $\lim_{k \rightarrow \infty} \|h(x_k)\|_1 = 0$, and by the boundedness of λ_{k+1} , it follows that $\lim_{k \rightarrow \infty} h(x_k)^T \lambda_{k+1} = 0$. Substituting this into (4.28), we get (4.25). Finally, by combining (4.13), Corollary 4.5 and (4.25) we obtain

$$\lim_{k \rightarrow \infty} \nabla f(x_k)^T d_k + v_k^T d_k = 0,$$

and subtracting this limit from (4.27) yields

$$\lim_{k \rightarrow \infty} \Gamma_k(d_k) - v_k^T d_k = 0.$$

□

We now show that the subgradients of Γ_k at d_k converge to zero.

Lemma 4.7 *The sequence $\{v_k\}$ generated by Algorithm 2.1 satisfies*

$$\lim_{k \rightarrow \infty} v_k = 0. \quad (4.29)$$

Proof. Lemma 4.6 implies that for any $\epsilon > 0$, there exists an integer $k(\epsilon)$ such that for all $k > k(\epsilon)$,

$$0 \leq v_k^T d_k - \Gamma_k(d_k) \leq \epsilon. \quad (4.30)$$

Assuming without loss of generality that $\epsilon < 1$, then by definition of v_k , the convexity of Γ_k , and (4.30), we have

$$\Gamma_k(\sqrt{\epsilon} \frac{v_k}{\|v_k\|}) \geq \Gamma(d_k) + v_k^T(\sqrt{\epsilon} \frac{v_k}{\|v_k\|} - d_k) \geq \sqrt{\epsilon} \|v_k\| - \epsilon. \quad (4.31)$$

The growth condition (4.3) implies that

$$\Gamma_k(\sqrt{\epsilon} \frac{v_k}{\|v_k\|}) \leq \beta_u \epsilon,$$

which together with (4.31) yields

$$\|v_k\| \leq (\beta_u + 1)\sqrt{\epsilon}.$$

Hence we obtain the limit (4.29). \square

We can now establish the main result of this section, namely that the primal-dual iterates defined by x_k and the PLA multipliers $(\lambda_{k+1}, \mu_{k+1})$ satisfy the KKT conditions of the nonlinear program, in the limit.

Theorem 4.8 *Suppose that Assumptions I hold and that Γ_k and m_k satisfy the Growth Conditions on page 12. Then, any limit point of the sequence $(x_k, \lambda_{k+1}, \mu_{k+1})$ generated by Algorithm 2.1 satisfies the KKT conditions (2.3).*

Proof. We have shown in Lemma 4.5 that (2.3b) and (2.3c) hold in the limit, while the nonnegativity of μ_{k+1} is guaranteed by (2.5d). Lemma 4.6 shows that the sequence $\{(x_k, \lambda_{k+1}, \mu_{k+1})\}$ satisfies (2.3e) in the limit. As to (2.3a), it holds in the limit because of (2.5a) and (4.29). \square

This global convergence result applies to the case when Γ is constructed by the Hermite interpolation process described in Procedure 3.1 and when the interpolation points satisfy Placement Conditions A. This is because Lemma 4.1 shows that the Growth Conditions are satisfied in this case. Thus, we have identified a concrete implementation of the PLA algorithm for which global convergence can be guaranteed. Clearly, many other strategies are permitted by our analysis.

5 Local Convergence Analysis

In this section, we give conditions under which Algorithm 2.1 identifies the optimal active set as the iterates approach a solution x_* that satisfies standard regularity conditions.

We recall that the working set \mathcal{W}_k at iteration k is given by (2.6), and we denote the active set at x_* as

$$\mathcal{W}_* = \{i \in \mathcal{E}\} \cup \{i \in \mathcal{I} \mid g^i(x_*) = 0\}. \quad (5.1)$$

Given a working set \mathcal{W} , we employ a superscript $^{\mathcal{W}}$ to denote a subvector with elements in \mathcal{W} . For instance,

$$g^{\mathcal{W}}(x_k) = [g^i(x_k)]_{i \in \mathcal{W}}. \quad (5.2)$$

For convenience, we also define

$$c^{\mathcal{W}}(x_k) \triangleq \begin{bmatrix} h(x_k) \\ g^{\mathcal{W}}(x_k) \end{bmatrix} \quad \text{and} \quad A^{\mathcal{W}}(x_k) \triangleq \begin{bmatrix} \nabla h(x_k)^T \\ \nabla g^{\mathcal{W}}(x_k)^T \end{bmatrix}, \quad (5.3)$$

and let $[\nu]^{\mathcal{W}}$ denote the multiplier set corresponding to a given working set \mathcal{W} ; specifically

$$[\nu]^{\mathcal{W}} \triangleq \begin{bmatrix} \lambda \\ [\mu]^{\mathcal{W}} \end{bmatrix} \quad \text{and} \quad \nu_* \triangleq \begin{bmatrix} \lambda_* \\ [\mu_*]^{\mathcal{W}_*} \end{bmatrix}. \quad (5.4)$$

The local convergence analysis of this section is self-contained and independent of the analysis in Section 4. This is so that we can identify essential properties of the model Γ that yield the desired active-set identification results. In Section 7 we outline an implementation of the algorithm that enjoys both the global and local convergence properties.

We make the following assumptions throughout this section.

Assumptions II

- a) x_* is a KKT point of the nonlinear program (2.1), and the functions f, h and g are continuously differentiable in a neighborhood of x_* .
- b) The PLA subproblem (2.4) is feasible for all k .
- c) The linear independence constraint qualification and strict complementarity hold at x_* . Thus, $A^{\mathcal{W}_*}(x_*)$ has full rank and there is a vector ν_* such that

$$\nabla f(x_*) - A^{\mathcal{W}_*}(x_*)^T \nu_* = 0, \quad [\mu_*]^{\mathcal{W}_*} > 0, \quad \text{and} \quad g^i(x_*) > 0 \text{ for } i \notin \mathcal{W}_*. \quad (5.5)$$

- d) There exist constants $\beta_l > 0$ and $\beta_u > 0$ such that for all k ,

$$\beta_l \|d\|^2 \leq d^T B_k d \leq \beta_u \|d\|^2 \quad \text{for all } d \in \mathbb{R}^n. \quad (5.6)$$

These assumptions are fairly standard in active-set identification studies; see e.g. [11].

In order to show that the PLA step d_k identifies the optimal active set \mathcal{W}_* for x_k close to x_* , the piecewise linear function $\Gamma_k(d)$ must have an appropriate shape. In general terms, it should be sufficiently flat near the origin (so that the gradient of the model $m_k(d)$ at $d = 0$ is close to $\nabla f(x_k)$), and it should grow at least linearly for large d (so that the step is restricted). The desired properties of Γ_k are stated precisely as conditions i) and ii) in the following lemma. Later on, in Lemma 5.3, we show that, by appropriately placing the interpolation points, we can guarantee that Γ_k satisfies these properties.

Lemma 5.1 *Suppose that Assumptions II hold at x_* . There exist constants $\gamma_s > 0, \gamma_1 > 0$ such that, if Γ_k satisfies the following two properties for all x_k close to x_* :*

- i) For all $\|d\| \geq \gamma_1$

$$\Gamma_k(d) \geq \gamma_s \|x_k - x_*\| \|d\|; \quad (5.7)$$

ii) For all d ,

$$\Gamma_k(d) \leq \beta_u \|d\|^2, \quad (5.8)$$

then the PLA step d_k satisfies

$$\|d_k\| < \gamma_1 \quad \text{and} \quad \mathcal{W}_k \subseteq \mathcal{W}_*. \quad (5.9)$$

Proof. Let us consider the second condition in (5.9). To show that for x_k close to x_* there cannot exist an index $i \in \mathcal{W}_k$ such that $i \notin \mathcal{W}_*$, we need to consider only the inequality constraints since, by definition (2.6), all equality constraints are contained in each working set \mathcal{W}_k as well as in \mathcal{W}_* .

For any index $i \notin \mathcal{W}_*$, we have that $g^i(x_*) > 0$, and thus there is a positive constant ϵ_g such that

$$g^i(x) \geq \epsilon_g > 0, \quad (5.10)$$

for all x_k close to x_* . Consider now any vector d that satisfies

$$g^i(x) + \nabla g^i(x)^T d \leq 0 \quad \text{for} \quad i \notin \mathcal{W}_*. \quad (5.11)$$

Then

$$\|\nabla g^i(x)\| \|d\| \geq -\nabla g^i(x)^T d \geq g^i(x) \geq \epsilon_g, \quad (5.12)$$

which gives the lower bound

$$\|d\| \geq \epsilon_g / \|\nabla g^i(x)\| \geq \epsilon_g / \gamma_g \triangleq \gamma_1, \quad (5.13)$$

where $\gamma_g > 0$ is a constant such that $\|\nabla g(x)\| \leq \gamma_g$ for all $x \in \Omega$.

Since $A^{\mathcal{V}_*}(x_*)$ has full rank by Assumptions II-(a), we have that for x_k close to x_* , the matrix $A^{\mathcal{V}_*}(x_k)$ has full rank and thus the system

$$c^{\mathcal{V}_*}(x_k) + A^{\mathcal{V}_*}(x_k)d = 0 \quad (5.14)$$

has a solution (that is not necessarily unique). Let \hat{d}_k denote the minimum norm solution of this system. Thus

$$c^{\mathcal{V}_*}(x_k) + A^{\mathcal{V}_*}(x_k)\hat{d}_k = 0, \quad (5.15)$$

and since by definition $c^{\mathcal{V}_*}(x_*) = 0$ we have, for x_k near x_* ,

$$\hat{d}_k = O(\|x_k - x_*\|). \quad (5.16)$$

Therefore, if x_k is sufficiently close to x_* the vector \hat{d}_k cannot satisfy the lower bound (5.13). But since we have shown that all steps d satisfying (5.11) also satisfy (5.13), we deduce that \hat{d}_k cannot satisfy (5.11), and thus

$$g^i(x_k) + \nabla g^i(x_k)^T \hat{d}_k > 0, \quad \text{for} \quad i \notin \mathcal{W}_*. \quad (5.17)$$

This condition and (5.15) show that \hat{d}_k is a feasible step of the PLA problem (2.4). This allows us to obtain an upper bound for the PLA objective (2.4a) by noting from (5.16) and (5.8) that there is a constant γ_2 such that

$$\nabla f(x_k)^T \hat{d}_k + \Gamma_k(\hat{d}_k) \leq \gamma_2 \|x_k - x_*\|, \quad (5.18)$$

for all x_k in some neighborhood of x_* .

Consider now the solution d_k of the PLA subproblem (2.4). Clearly,

$$\nabla f(x_k)^T d_k = [\nabla f(x_k) - A^{\mathcal{W}_*}(x_k)^T \nu_*]^T d_k + \nu_*^T A^{\mathcal{W}_*}(x_k) d_k, \quad (5.19)$$

and by the first condition in (5.5), there is a constant γ_3 such that for all x_k near x_*

$$\|[\nabla f(x_k) - A^{\mathcal{W}_*}(x_k)^T \nu_*]^T d_k\| \leq \gamma_3 \|x_k - x_*\| \|d_k\|. \quad (5.20)$$

As to the last term in (5.19), since d_k is feasible for the PLA problem (2.4), it satisfies

$$\nabla h(x_k)^T d_k + h(x_k) = 0, \quad \text{and} \quad \nabla g^{\mathcal{W}_*}(x_k)^T d_k + g^{\mathcal{W}_*}(x_k) \geq 0, \quad (5.21)$$

and since $h(x_*) = g^{\mathcal{W}_*}(x_*) = 0$ and $[\mu_*]^{\mathcal{W}_*} \geq 0$, by continuity of h and g we have that

$$\lambda_*^T \nabla h(x_k)^T d_k \geq -\gamma_4 \|x_k - x_*\|$$

and

$$([\mu_*]^{\mathcal{W}_*})^T \nabla g^{\mathcal{W}_*}(x_k)^T d_k \geq -([\mu_*]^{\mathcal{W}_*})^T g^{\mathcal{W}_*}(x_k) \geq -\gamma_5 \|x_k - x_*\|,$$

for some constants $\gamma_4, \gamma_5 > 0$. Combining these two inequalities and recalling (5.3)-(5.4), we have

$$\nu_*^T A^{\mathcal{W}_*}(x_k) d_k \geq -(\gamma_4 + \gamma_5) \|x_k - x_*\|. \quad (5.22)$$

Substituting this bound and (5.20) in (5.19), we have that

$$\nabla f(x_k)^T d_k \geq -\gamma_3 \|x_k - x_*\| \|d_k\| - \gamma_6 \|x_k - x_*\|, \quad \text{where} \quad \gamma_6 = \gamma_4 + \gamma_5. \quad (5.23)$$

Let us define

$$\gamma_s \triangleq \frac{\gamma_2 + \gamma_6}{\gamma_1} + \gamma_3 + 1, \quad (5.24)$$

and assume that (5.7) holds in a neighborhood of x_* , where γ_1 is defined in (5.13). To prove (5.9), suppose by way of contradiction that the PLA step satisfies $\|d_k\| \geq \gamma_1$. Then, combining (5.7) and (5.23), we obtain

$$\nabla f(x_k)^T d_k + \Gamma_k(d_k) \geq (\gamma_s - \gamma_3) \|x_k - x_*\| \|d_k\| - \gamma_6 \|x_k - x_*\|. \quad (5.25)$$

We have seen that the minimum norm step \hat{d}_k satisfies (5.18), and since the PLA step d_k cannot yield a greater PLA objective value than \hat{d}_k , we have from (5.18), (5.24) and (5.25) that

$$\gamma_2 \|x_k - x_*\| \geq \left(\frac{\gamma_2 + \gamma_6}{\gamma_1} + 1\right) \gamma_1 \|x_k - x_*\| - \gamma_6 \|x_k - x_*\| \quad (5.26)$$

or, equivalently,

$$\gamma_1 \leq 0.$$

This contradicts the definition of γ_1 in (5.13). Therefore, we must have that $\|d_k\| < \gamma_1$, and consequently (5.11) does not hold for any $i \notin \mathcal{W}_*$. We have thus shown that, for any $i \notin \mathcal{W}_*$, we must have that $i \notin \mathcal{W}_k$, and this concludes the proof. \square

We can now establish the active set identification result for the new algorithm.

Theorem 5.2 *Suppose that Assumptions II hold at x_* . There exist constants $\gamma_s > 0, \gamma_1 > 0$ such that if condition (5.7) and (5.8) in Lemma 5.1 are satisfied, then*

$$\mathcal{W}_k = \mathcal{W}_* \quad (5.27)$$

for all large x_k close to x_* .

Proof. By Lemma 5.1, we know that $\mathcal{W}_k \subseteq \mathcal{W}_*$ for x_k close to x_* . Suppose by way of contradiction that for some sequence $\{x_k\} \rightarrow x_*$, we have that $\mathcal{W}_k \not\supseteq \mathcal{W}_*$, i.e., there exists an index $i \in \mathcal{W}_* \cap \mathcal{I}$ such that $i \notin \mathcal{W}_k$.

Let us define the index set $\mathcal{W}_*^- := \mathcal{W}_* \setminus \{i\}$, so that $A^{\mathcal{W}_*^-}(x)$ denotes the matrix obtained by removing row $\nabla g^i(x)^T$ from $A^{\mathcal{W}_*}(x)$. Since we assume that $i \notin \mathcal{W}_k$, we have that for x_k near x_* the PLA step d_k satisfies

$$\nabla f(x_k) - A^{\mathcal{W}_*^-}(x_k)^T \nu^{\mathcal{W}_*^-} + v_k = 0 \quad (5.28a)$$

$$c^{\mathcal{W}_*^-}(x_k) + A^{\mathcal{W}_*^-}(x_k)d_k = 0 \quad (5.28b)$$

$$\nu^{\mathcal{W}_*^-} \geq 0, \quad (5.28c)$$

where $v_k \in \partial \Gamma_k(d_k)$ and $\nu^{\mathcal{W}_*^-}$ is the vector obtained by removing μ^i from $\nu^{\mathcal{W}_*}$. Equivalently, d_k solves the following relaxation of the PLA subproblem at x_k .

$$\min_d \nabla f(x_k)^T d + \Gamma_k(d) \quad (5.29a)$$

$$\text{s.t. } g^{\mathcal{W}_*^-}(x_k) + \nabla g^{\mathcal{W}_*^-}(x_k)^T d \geq 0 \quad (5.29b)$$

$$h(x_k) + \nabla h(x_k)^T d = 0. \quad (5.29c)$$

Now, since $A^{\mathcal{W}_*}(x_*)$ has full rank and $[\mu_*]^{\mathcal{W}_*} > 0$, the equations in (5.5) can no longer be satisfied if we replace $A^{\mathcal{W}_*}(x_*)$ with $A^{\mathcal{W}_*^-}(x_*)$. In other words, the dual system

$$\nabla f(x_*) - \nabla g^{\mathcal{W}_*^-}[\mu]^{\mathcal{W}_*^-} - \nabla h(x_*)\lambda = 0, \quad [\mu]^{\mathcal{W}_*^-} \geq 0$$

is infeasible, and by Farkas's lemma we can find a direction d_*^F and a scalar $\tau > 0$ such that

$$\|d_*^F\| = 1, \quad \nabla f(x_*)^T d_*^F < -2\tau, \quad \nabla g^{\mathcal{W}_*^-}(x_*)^T d_*^F \geq 0, \quad \nabla h(x_*)^T d_*^F = 0. \quad (5.30)$$

Since $A^{\mathcal{W}_*^-}(x_*)$ has full rank and $A(x)$ is continuous, in a neighborhood of x_* we can define the direction

$$d^F(x) = d_*^F + A^{\mathcal{W}_*^-}(x)^T [A^{\mathcal{W}_*^-}(x)A^{\mathcal{W}_*^-}(x)^T]^{-1} (A^{\mathcal{W}_*^-}(x_*) - A^{\mathcal{W}_*^-}(x))d_*^F, \quad (5.31)$$

which satisfies

$$A^{\mathcal{W}_*^-}(x)d^F(x) = A^{\mathcal{W}_*^-}(x_*)d_*^F. \quad (5.32)$$

It follows by (5.30)–(5.32) and continuity of $\nabla f(x)$, $d^F(x)$ and $A(x)$ that, for x_k close to x_* , the vector $d^F(x_k)$ satisfies

$$\nabla f(x_k)^T d^F(x_k) < -\tau, \quad \nabla g^{\mathcal{W}_*^-}(x_k)^T d^F(x_k) \geq 0, \quad \nabla h(x_k)^T d^F(x_k) = 0 \quad (5.33)$$

and

$$\|d^F(x_k)\| = 1 + O(\|x_k - x_*\|). \quad (5.34)$$

By Assumptions II, there exists a direction $d^C(x) = O(\|x_k - x_*\|)$ such that

$$g^{\mathcal{V}_*^-}(x_k) + \nabla g^{\mathcal{V}_*^-}(x_k)^T d^C(x_k) = 0, \quad h(x_k) + \nabla h(x_k)^T d^C(x_k) = 0. \quad (5.35)$$

We can combine these two directions using scalar $\alpha > 0$ to define

$$\tilde{d}_k(\alpha) = d^C(x_k) + \alpha d^F(x_k).$$

Clearly, $\|\tilde{d}_k(\alpha)\| = \alpha + O(\|x_k - x_*\|)$, and

$$g^{\mathcal{V}_*^-}(x_k) + \nabla g^{\mathcal{V}_*^-}(x_k)^T \tilde{d}_k(\alpha) \geq 0, \quad h(x_k) + \nabla h(x_k)^T \tilde{d}_k(\alpha) = 0, \quad (5.36)$$

showing that $\tilde{d}_k(\alpha)$ is feasible for problem (5.29), for any α .

From (5.33), we have

$$\nabla f(x_k)^T \tilde{d}_k(\alpha) < -\alpha\tau + O(\|x_k - x_*\|). \quad (5.37)$$

In particular, if we choose $\alpha = \|x_k - x_*\|^\delta$ for $\delta \in (0, 1)$, we have

$$\nabla f(x_k)^T \tilde{d}_k(\alpha) < -\frac{1}{2}\tau\|x_k - x_*\|^\delta. \quad (5.38)$$

Since $\|\tilde{d}_k(\alpha)\| = \alpha + O(\|x_k - x_*\|) = O(\|x_k - x_*\|^\delta)$, we also have from (5.8)

$$\Gamma_k(\tilde{d}_k(\alpha)) \leq \beta_u \|\tilde{d}_k(\alpha)\|^2 = O(\|x_k - x_*\|^{2\delta}). \quad (5.39)$$

Now, since d_k is the optimal solution of the PLA subproblem (2.4), it must yield an objective value that is not larger than that given by $\tilde{d}_k(\alpha)$, i.e.,

$$\nabla f(x_k)^T d_k + \Gamma_k(d_k) \leq \nabla f(x_k)^T \tilde{d}_k(\alpha) + \Gamma_k(\tilde{d}_k(\alpha)) \leq -\frac{1}{4}\tau\|x_k - x_*\|^\delta, \quad (5.40)$$

which together with the condition $\Gamma_k(d_k) \geq 0$, implies

$$\nabla f(x_k)^T d_k \leq -\frac{1}{4}\tau\|x_k - x_*\|^\delta. \quad (5.41)$$

We now show that this leads to a contradiction. By (5.5) and continuity of ∇f and A , we have that for x_k near x_*

$$\nabla f(x_k) - \nabla g^{\mathcal{V}_*}(x_k)[\mu_*]^{\mathcal{V}_*} - \nabla h(x_k)\lambda_* = O(\|x_k - x_*\|), \quad (5.42)$$

and therefore, there exists a constant $\gamma_9 > 0$ such that

$$d_k^T \nabla f(x_k) - d_k^T \nabla g^{\mathcal{V}_*}(x_k)[\mu_*]^{\mathcal{V}_*} - d_k^T \nabla h(x_k)\lambda_* \geq -\gamma_9\|x_k - x_*\|\|d_k\| \geq -\gamma_1\gamma_9\|x_k - x_*\|, \quad (5.43)$$

where the last inequality follows from (5.9). Since the PLA step satisfies $\nabla g^{\mathcal{V}_*}(x_k)d_k + g^{\mathcal{V}_*}(x_k)^T \geq 0$, and $[\mu_*]^{\mathcal{V}_*} \geq 0$, $g^{\mathcal{V}_*}(x_*) = 0$, by continuity of $g(x)$, we have that

$$d_k^T \nabla g^{\mathcal{V}_*}(x_k)[\mu_*]^{\mathcal{V}_*} \geq -g^{\mathcal{V}_*}(x_k)^T [\mu_*]^{\mathcal{V}_*} \geq -\gamma_{10}\|x_k - x_*\| \quad (5.44)$$

for some constant $\gamma_{10} > 0$. Similarly, since $\nabla h(x_k)^T d_k + h(x_k) = 0$ and $h(x_*) = 0$, we have

$$d_k^T \nabla h(x_k) \lambda_* = \lambda_*^T h(x_k) \geq -\gamma_{11} \|x_k - x_*\| \quad (5.45)$$

for some constant $\gamma_{11} > 0$. Combining (5.43)–(5.45) yields

$$\nabla f(x_k)^T d_k \geq -(\gamma_1 \gamma_9 + \gamma_{10} + \gamma_{11}) \|x_k - x_*\|, \quad (5.46)$$

which contradicts (5.41) for x_k sufficiently small to x_* . Therefore, there cannot exist an index i such that $i \in \mathcal{W}_*$ and $i \notin \mathcal{W}_k$. \square

Theorem 5.2 gives conditions on the shape of the piecewise linear function Γ that ensure correct active set identification. However, in computation Γ is determined by placement of the interpolation points, as described in Section 3. The following result specifies two intervals that must contain interpolation points for each coordinate direction, in order for Γ_k to have the desired properties.

Theorem 5.3 *Suppose that Assumptions II hold at x_* and that the PLA model Γ_k is constructed as in Section 3. Then, there exist constants $\gamma_s > 0, \gamma_1 > 0$ such that, if both intervals*

$$\left[\frac{2\sqrt{n}\gamma_s}{\beta_l} \|x_k - x_*\|, \frac{\gamma_1}{\sqrt{n}} \right] \quad \text{and} \quad \left[-\frac{\gamma_1}{\sqrt{n}}, -\frac{2\sqrt{n}\gamma_s}{\beta_l} \|x_k - x_*\| \right] \quad (5.47)$$

contain interpolation points for all coordinate directions i and for all x_k close to x_ , condition (5.7) holds and*

$$\mathcal{W}_k = \mathcal{W}_*. \quad (5.48)$$

Proof. Since the conditions of Lemma 5.1 are satisfied, let γ_s and γ_1 be the values guaranteed by that lemma. Suppose the intervals (5.47) contain interpolation points $t_+^i > t_-^i$ respectively, and suppose $\|d\| \geq \gamma_1$. Then $\|d\|_\infty \geq \frac{\gamma_1}{\sqrt{n}}$, and equivalently, for some i , $\|d\|_\infty = |d^i| \geq \frac{\gamma_1}{\sqrt{n}}$. Suppose first that $d^i \geq \frac{\gamma_1}{\sqrt{n}}$. Then $d^i \geq t_+^i$, and by (3.6) and (3.8)

$$\begin{aligned} \Gamma_k^i(d^i) &\geq -\frac{1}{2} b_k^i (t_+^i)^2 + b_k^i t_+^i d^i \\ &\geq \frac{1}{2} \beta_l t_+^i d^i \\ &\geq \gamma_s \sqrt{n} \|x_k - x_*\| \|d\|_\infty \end{aligned} \quad (5.49)$$

$$\geq \gamma_s \|x_k - x_*\| \|d\|, \quad (5.50)$$

since $t_+^i \geq 2\sqrt{n}\gamma_s \|x_k - x_*\|/\beta_l$, and $\sqrt{n}\|d\|_\infty \geq \|d\|$. A similar argument implies (5.50) when $d^i \leq -\frac{\gamma_1}{\sqrt{n}}$. Since $\Gamma_k(d) \geq \Gamma_k^i(d^i)$, we have attained (5.7).

Because $\Gamma_k(d)$ is an under estimate of $\frac{1}{2} d^T B_k d$, by noticing Assumption II-(d) we obtain (5.8). The conclusion (5.48) follows from Theorem 5.2. \square

It is interesting to note that the intervals (5.47) are similar to the interval specified by Oberlin and Wright [21] for a trust region radius in an SL-QP method. Although the values γ_s and γ_1 are not generally available, one can ensure the conditions of Theorem 5.48 are satisfied in the limit, for example by specifying that one interpolation point $t^i = O(\|x_k - x_*\|^{1/2})$. Several quantities computed by the algorithm, such as the norm of the KKT error, are of order $O(\|x_k - x_*\|^{1/2})$.

6 A Variant Based on BFGS Updating

In the previous sections, we assumed that the Hessian approximation B_k is diagonal. One way to incorporate more useful curvature information about the problem is to define B_k through limited memory quasi-Newton updating. Due to the special structure of this matrix, we can perform a change of variables (at reasonable cost) that allows us to represent the model in separable form, and then apply the interpolation techniques discussed in Section 3 to construct a piecewise linear model. We refer to the method based on this approach as the BFGS-PLA algorithm.

At an iterate x_k , we update a BFGS approximation B_k to the Hessian of the Lagrangian $\nabla_{xx}^2 L(x_k, \lambda_k, \mu_k)$. The matrix B_k is defined in terms of $l/2$ correction pairs $\{s_i, y_i\}$, $i = 1, \dots, l/2$, where each s_i stores the change in the variables x at a previous iteration and y_i stores the corresponding change in the gradient of the Lagrangian; see e.g., [20, section 18.3]. Here l is a small even number (say $l = 10$). This limited memory BFGS approximation can be written in a so-called compact form [4],

$$B_k = \theta_k I + U_k R_k U_k^T, \quad (6.1)$$

where $\theta_k > 0$ is a scaling parameter, I is the identity matrix, U_k is an $n \times l$ dense matrix whose columns are formed by the vectors s_i and y_i , and R_k is a symmetric $l \times l$ matrix.

For the sake of numerical stability, we orthogonalize the columns of U_k to obtain an $n \times l_k$ matrix V_k such that

$$V_k^T V_k = I, \quad \text{and} \quad \text{span}(V_k) = \text{span}(U_k), \quad (6.2)$$

where $l_k \leq l$ is the rank of U_k . If we use the Gram-Schmidt process or Householder transformations, the cost of computing V_k is of the order $O(nl^2)$, which is acceptable when l is small.

The quadratic model $d^T B_k d$ defined through (6.1) is not separable (B_k is a dense matrix), but it has a special structure. In order to represent it in separable form, we first form the $l_k \times l_k$ matrix

$$S_k = V_k^T B_k V_k = \theta_k I + V_k^T U_k R_k U_k^T V_k \quad (6.3)$$

that can be diagonalized at negligible cost. Thus, we compute

$$S_k = Q_k \Sigma_k Q_k^T, \quad (6.4)$$

where Q_k is an $l_k \times l_k$ orthogonal matrix and Σ_k is a diagonal matrix.

Let us define the variables $q \in \mathbb{R}^{n+l_k}$ through the invertible linear transformation

$$\begin{bmatrix} d \\ 0 \end{bmatrix} = \begin{bmatrix} V_k Q_k & I \\ 0 & V_k^T \end{bmatrix} q \quad \text{and} \quad q = \begin{bmatrix} q_1 \\ q_2 \end{bmatrix}, \quad (6.5)$$

where $q_1 \in \mathbb{R}^{l_k}$ and $q_2 \in \mathbb{R}^n$. Then from (6.1)–(6.5), we have

$$\begin{aligned}
 \frac{1}{2}d^T B_k d &= \frac{1}{2} (V_k Q_k q_1 + q_2)^T B_k (V_k Q_k q_1 + q_2) \\
 &= \frac{1}{2} q_1^T Q_k^T S_k Q_k q_1 + q_1^T Q_k^T V_k^T B_k q_2 + \frac{1}{2} q_2^T B_k q_2 \\
 &= \frac{1}{2} q_1^T \Sigma_k q_1 + \frac{1}{2} \theta_k q_2^T q_2 \\
 &= \frac{1}{2} q^T \begin{bmatrix} \Sigma_k & 0 \\ 0 & \theta_k I \end{bmatrix} q \\
 &\triangleq \frac{1}{2} q^T C_k q.
 \end{aligned} \tag{6.6}$$

where the second equality follows from (6.5), since $V_k^T q_2 = 0$ implies $U_k^T q_2 = 0$, and consequently we have both $B_k q_2 = \theta_k q_2$ and $V_k^T B_k q_2 = 0$.

Since $\frac{1}{2}q^T C_k q$ is separable function in \mathbb{R}^{n+l_k} , we can compute a piecewise linear approximation $\Gamma_k(q)$ to it, as discussed in Section 3. For each $i \in \{1, \dots, n + l_k\}$ we define a univariate piecewise linear function Γ_k^i such that

$$\Gamma_k^i(q^i) \approx \frac{1}{2} c_k^i (q^i)^2, \tag{6.7}$$

where c_k^i denotes the i -th diagonal entry of C_k . As in Section 3, the univariate function $\Gamma_k^i(q^i)$ is composed of $(r_k^i + 1)$ linear segments denoted by $\ell_{k,j}^i(q^i)$, for $j = 0, \dots, r_k^i$. We choose $\ell_{k,j}^i(q^i)$ to be the Hermite interpolant of $\frac{1}{2}c_k^i (q^i)^2$ at a designated node $t_{k,j}^i$. This implies that

$$\ell_{k,j}^i(q^i) = -\frac{1}{2}c_k^i (t_{k,j}^i)^2 + (c_k^i t_{k,j}^i) q^i. \tag{6.8}$$

We require that one interpolation point be given by $t_{k,0}^i = 0$, as in (3.9). We now define

$$\Gamma_k^i(q^i) = \max_j \{\ell_{k,j}^i(q^i)\}, \quad j \in J_k^i = \{0, \dots, r_k^i\}. \tag{6.9}$$

This ensures $\Gamma_k^i(q^i) \geq 0$, for all q^i , and $\Gamma_k^i(0) = 0$. The multivariate function $\Gamma_k(q)$ that approximates $\frac{1}{2}q^T C_k q$ is then given as

$$\Gamma_k(q) \triangleq \sum_{i=1}^{n+l_k} \Gamma_k^i(q^i). \tag{6.10}$$

If we define the linear transformation $T_k : \mathbb{R}^{n+l_k} \rightarrow \mathbb{R}^{n+l_k}$ as follows:

$$\begin{bmatrix} d \\ 0 \end{bmatrix} = T_k q, \quad \text{with} \quad T_k = \begin{bmatrix} V_k Q_k & I \\ 0 & V_k^T \end{bmatrix}, \tag{6.11}$$

then the PLA subproblem (2.4) can now be stated as

$$\min_{q,d} \nabla f(x_k)^T d + \Gamma_k(q) \tag{6.12a}$$

$$\text{s.t. } T_k q = \begin{bmatrix} d \\ 0 \end{bmatrix} \tag{6.12b}$$

$$h(x_k) + \nabla h(x_k)^T d = 0 \tag{6.12c}$$

$$g(x_k) + \nabla g(x_k)^T d \geq 0. \tag{6.12d}$$

The first l_k columns and the last l_k rows of the constraint matrix corresponding to (6.12b) are not sparse. Thus, the memory length $l/2$ must be chosen small enough so that the cost of working with these constraints is acceptable.

The BFGS-PLA algorithm is the variant of Algorithm I that uses the limited memory BFGS approximation (6.1) and defines the piecewise linear subproblem by (6.12). A practical implementation of this algorithm is outside the scope of this paper because it must address a variety of delicate issues, including a reformulation of the linear program (6.12) that keeps the computation cost to a minimum, a procedure for safeguarding the BFGS update, and the use of a trust region to stabilize the step computation. These issues, as well as computational experiments with the BFGS-PLA algorithm are reported in a companion paper [7].

Our interest here is to show that the limited memory BFGS-PLA algorithm enjoys the global and local convergence guarantees described in Sections 4 and 5.

6.1 Convergence Analysis

Let us begin by studying the global convergence properties of the BFGS-PLA algorithm. We have seen in Section 4 that Algorithm I is globally convergent if the model m_k and the function Γ_k satisfy Growth Conditions of page 12. Note that in that analysis we did not assume that the Hessian approximation B_k is diagonal, but only that it is uniformly positive definite and bounded; see (4.1). Such generality will be very useful in the analysis of this section.

We begin by expressing the PLA subproblem (6.12) in the space of the original variables d . For this purpose, we first note that, by construction, the matrix T_k defined in (6.11) is invertible for all k . Thus, we can define the function

$$\hat{q}_k(d) = T_k^{-1} \begin{bmatrix} d \\ 0 \end{bmatrix}, \tag{6.13}$$

and the model

$$\hat{\Gamma}_k(d) = \Gamma_k(\hat{q}_k(d)) = \sum_{i=1}^{n+l_k} \Gamma_k^i(\hat{q}_k^i(d)). \tag{6.14}$$

The PLA subproblem (6.12) can then be expressed as

$$\min_d \hat{n}_k(d) \triangleq \nabla f(x_k)^T d + \hat{\Gamma}_k(d) \tag{6.15a}$$

$$\text{s.t. } h(x_k) + \nabla h(x_k)^T d = 0 \tag{6.15b}$$

$$g(x_k) + \nabla g(x_k)^T d \geq 0. \tag{6.15c}$$

If we can show that, under Assumptions I, the Growth Conditions of page 12 hold for \hat{m}_k and $\hat{\Gamma}_k$, then we will be able to conclude from Theorem 4.8 that all limit points generated by the BFGS-PLA algorithm are KKT points of the nonlinear program. By (6.8)-(6.10) and (6.14), the properties of $\hat{\Gamma}_k$ are influenced by the placement of the interpolation points $t_{k,j}^i$ in \mathbb{R}^{n+l_k} . We now present conditions on the positioning of these points that will allow us to prove that the model (6.15a) satisfies the Growth Conditions.

Placement Conditions B

Define the vector $\varphi_k \in \mathbb{R}^{n+l_k}$ by,

$$\varphi_k = \hat{q}_k \left(-B_k^{-1} \nabla f(x_k) \right), \quad (6.16)$$

where \hat{q}_k is given in (6.13). Given a constant $\chi_0 > 0$, for each iteration k and for each coordinate direction i in \mathbb{R}^{n+l_k} , at least two interpolation points $t_{k,l}^i, t_{k,u}^i$ satisfy

$$t_{k,l}^i \leq \varphi_k^i - \chi_0, \quad \text{and} \quad t_{k,u}^i \geq \varphi_k^i + \chi_0. \quad (6.17)$$

Furthermore, the sequences $\{t_{k,l}^i\}$ and $\{t_{k,u}^i\}$ are uniformly bounded .

Let us now suppose that the limited memory BFGS matrix B_k defined in (6.1) satisfies Assumption I-d on page 11, i.e., that it is uniformly positive definite and bounded. (In [7] we discuss how to safeguard the limited memory BFGS update so as to guarantee this condition in practice.) Since B_k is given by (6.1) and the rank of U_k is assume to be less than n , Assumption I-d implies $\theta_k \in [\beta_l, \beta_u]$. Similarly, by (6.3) the eigenvalues of S_k lie in the interval that contains the eigenvalues of B_k . Therefore, we have from (6.4) that the diagonal elements of Σ_k also lie in the interval $[\beta_l, \beta_u]$. Therefore, it follows from the definition of C_k in (6.6) that for all k and all $i \in \{0, \dots, n + l_k\}$,

$$\beta_l \leq c_k^i \leq \beta_u. \quad (6.18)$$

We also have from (6.2) and the orthogonality of Q_k that, if d and q are related through (6.5), then

$$\begin{aligned} \|d\|^2 &= \|V_k Q_k q_1 + q_2\|^2 \\ &= q_1^T Q_k^T V_k^T V_k Q_k q_1 + 2q_2^T V_k Q_k q_1 + q_2^T q_2 \\ &= \|q\|^2. \end{aligned}$$

Thus, for any d we have

$$\|\hat{q}_k(d)\| = \|d\|. \quad (6.19)$$

The following result shows that, if the BFGS-PLA algorithm satisfies Assumptions I, and if the interpolation points $\{t_{k,j}^i\}$ comply with Placement Conditions B, then the Growth Conditions are satisfied.

Theorem 6.1 *Suppose that Assumptions I on page 11 hold, and that the Placement Conditions B are satisfied at each iteration. Then, the function $\hat{\Gamma}_k(d)$ and the model $\hat{m}_k(d)$ satisfy the Growth Conditions on page 12.*

Proof. We first verify that the Placement Conditions B are satisfiable by showing that $\{\varphi_k\}$ is bounded. From (6.16), (6.19) and Assumption I-d we have

$$\|\varphi_k\| = \|\hat{q}_k \left(-B_k^{-1} \nabla f(x_k) \right)\| \leq \frac{1}{\beta_l} \|\nabla f(x_k)\|. \quad (6.20)$$

Since $\|\nabla f(x_k)\|$ is bounded above by Assumption I-a, there exists a constant $\bar{\varphi} > 0$ such that $\|\varphi_k\| \leq \bar{\varphi}$, for all k .

Let us show that the first Growth Condition on page 12, namely (P1), is satisfied. Since each one-dimensional function $\Gamma_k^i(q^i)$ is constructed by Hermite interpolation of the convex function $\frac{1}{2}c_k^i(q^i)^2$, with one node at the origin, we immediately have that $\Gamma_k^i(q^i)$ is convex, nonnegative, and $\Gamma_k^i(0) = 0$. Thus, by (6.10) we have that $\Gamma_k(d) \geq 0$ and $\Gamma_k(0) = 0$. From (6.14), the convexity of $\Gamma_k^i(q^i)$, and the fact that $\hat{q}_k(d)$ is a linear function of d , we conclude that $\hat{\Gamma}_k(d)$ is convex, nonnegative and vanishes at zero. Therefore, the first Growth Condition (P1) is satisfied.

To show that the second Growth Condition holds, we first note from (6.6), (6.11) and (6.13), that for any vector $d \in \mathbb{R}^n$,

$$\frac{1}{2}\hat{q}_k(d)^T C_k \hat{q}_k(d) = \frac{1}{2}d^T B_k d. \quad (6.21)$$

Since, for each i , the function $\Gamma_k^i(q^i)$ is an underestimate of $\frac{1}{2}c_k^i(q^i)^2$, i.e., $\Gamma_k^i(q^i) \leq \frac{1}{2}c_k^i(q^i)^2$ for all q^i , we have

$$\hat{\Gamma}_k(d) = \Gamma_k(\hat{q}_k(d)) \leq \frac{1}{2}\hat{q}_k(d)^T C_k \hat{q}_k(d) = \frac{1}{2}d^T B_k d \leq \beta_u \|d\|^2 \quad (6.22)$$

where we have used (6.10), (6.14) and Assumption I-d. Hence the second Growth Condition (4.3) is satisfied.

Lastly, to establish the third Growth Condition, we obtain from (6.16) and (6.21) that

$$\varphi_k^T C_k \varphi_k = \nabla f(x_k)^T B_k^{-1} \nabla f(x_k). \quad (6.23)$$

Therefore,

$$\begin{aligned} & \frac{1}{2}(\hat{q}_k(d) + \varphi_k)^T C_k (\hat{q}_k(d) + \varphi_k) + \nabla f(x_k)^T (d - B_k^{-1} \nabla f(x_k)) \\ &= \frac{1}{2}d^T B_k d - \frac{1}{2}\nabla f(x_k)^T B_k^{-1} \nabla f(x_k) + \sum_{i=1}^{n+l_k} c_k^i \varphi_k^i \hat{q}_k^i(d) + \nabla f(x_k)^T d. \end{aligned} \quad (6.24)$$

On the other hand, by linearity of \hat{q}_k and (6.16), we have $(\hat{q}_k(d) + \varphi_k) = \hat{q}_k(d - B_k^{-1} \nabla f(x_k))$. Thus we also have,

$$\begin{aligned} & \frac{1}{2}(\hat{q}_k(d) + \varphi_k)^T C_k (\hat{q}_k(d) + \varphi_k) + \nabla f(x_k)^T (d - B_k^{-1} \nabla f(x_k)) \\ &= \frac{1}{2}d^T B_k d - \frac{1}{2}\nabla f(x_k)^T B_k^{-1} \nabla f(x_k). \end{aligned} \quad (6.25)$$

Comparing (6.24) and (6.25), we obtain that for all $d \in \mathbb{R}^n$,

$$\sum_{i=1}^{n+l_k} c_k^i \varphi_k^i \hat{q}_k^i(d) + \nabla f(x_k)^T d = 0. \quad (6.26)$$

We now repeat the reasoning that led to (4.7)-(4.11). Since the sequences $\{t_{k,u}^i\}$, $\{t_{k,l}^i\}$ are bounded, it follows from (6.8), (6.9) and (6.18) that

$$\Gamma_k^i(q^i) - c_k^i \varphi_k^i q^i \geq \chi_2 + (t_{k,u}^i - \varphi_k^i) c_k^i q^i \quad (6.27)$$

and

$$\Gamma_k^i(q^i) - c_k^i \varphi_k^i q^i \geq \chi_2 + (t_{k,l}^i - \varphi_k^i) c_k^i q^i, \quad (6.28)$$

where $\chi_2 = -\frac{1}{2}\chi_1^2\beta_u$ and χ_1 is a constant such that $-\chi_1 \leq t_{k,l}^i \leq t_{k,u}^i \leq \chi_1$. Now, suppose that $q_i \geq 0$. From (6.17) we have that $(t_{k,u}^i - \varphi_k^i) \geq \chi_0$, and hence it follows from (6.27) that

$$\Gamma_k^i(q^i) - c_k^i \varphi_k^i q^i \geq \chi_2 + \chi_0 c_k^i q^i. \quad (6.29)$$

Similarly, if $q_i < 0$ and since $(t_{k,l}^i - \varphi_k^i) \leq -\chi_0$, we have from (6.28) that

$$\Gamma_k^i(q^i) - c_k^i \varphi_k^i q^i \geq \chi_2 - \chi_0 c_k^i q^i. \quad (6.30)$$

Together, (6.29) and (6.30) imply that for any q_i ,

$$\Gamma_k^i(q^i) - c_k^i \varphi_k^i q^i \geq \chi_2 + \chi_0 c_k^i |q^i|.$$

In particular, for all $d \in \mathbb{R}^n$, we have

$$\Gamma_k^i(\hat{q}_k^i(d)) - c_k^i \varphi_k^i \hat{q}_k^i(d) \geq \chi_2 + \chi_0 c_k^i |\hat{q}_k^i(d)|. \quad (6.31)$$

Combining (6.14), (6.15a), (6.26) and (6.31) yields

$$\begin{aligned} \hat{m}_k(d) &= \nabla f(x_k)^T d + \sum_{i=1}^{n+l_k} \Gamma_k^i(\hat{q}_k^i(d)) \\ &= - \sum_{i=1}^{n+l_k} c_k^i \varphi_k^i \hat{q}_k^i(d) + \sum_{i=1}^{n+l_k} \Gamma_k^i(\hat{q}_k^i(d)) \\ &\geq n\chi_2 + \chi_0 \|C_k \hat{q}_k(d)\|_1. \end{aligned}$$

Now, by (6.18) and (6.19)

$$\|C_k \hat{q}_k(d)\|_1 \geq \beta_l \|\hat{q}_k(d)\|_1 \geq \beta_l \|\hat{q}_k(d)\| = \beta_l \|d\|,$$

and thus

$$\hat{m}_k(d) \geq n\chi_2 + \chi_0 \beta_l \|d\|.$$

This implies that condition (4.4) holds with $\varsigma = \chi_0 \beta_l / 2$, and $\Delta_G = 2n|\chi_2| / \chi_0 \beta_l$. \square

We can now establish a global convergence result.

Theorem 6.2 *Suppose that Assumptions I on page 11 hold, and that the Placement Conditions B are satisfied at each iteration. Then, any limit point of the sequence $(x_k, \lambda_{k+1}, \mu_{k+1})$ generated by the BFGS-PLA version of Algorithm 2.1 satisfies the KKT conditions (2.3) of the nonlinear program.*

Proof. In Theorem 6.1 we have shown that Growth Condition is satisfied. The result follows from Theorem 4.8. \square

Let us now consider the local active-set identification properties of the algorithm. By placing the interpolation points in a similar manner as in Section 5, albeit in \mathbb{R}^{n+l_k} , we can prove the analogue of Theorem 5.3.

Theorem 6.3 *Suppose that Assumptions II on page 19 hold at x_* . There exist constants $\gamma_s > 0, \gamma_1 > 0$ such that for $\|x_k - x_*\|$ sufficiently small, if both intervals*

$$\left[\frac{2\sqrt{n+l_k}\gamma_s}{\beta_l} \|x_k - x_*\|, \frac{\gamma_1}{\sqrt{n+l_k}} \right] \quad \text{and} \quad \left[-\frac{\gamma_1}{\sqrt{n+l_k}}, -\frac{2\sqrt{n+l_k}\gamma_s}{\beta_l} \|x_k - x_*\| \right] \quad (6.32)$$

contain interpolation points for each coordinate direction i in \mathbb{R}^{n+l_k} and each iteration k , then the conditions (5.7) and (5.8) in Lemma 5.1 are satisfied, and thus

$$\mathcal{W}_k = \mathcal{W}_*. \quad (6.33)$$

Proof. Since the assumptions of Lemma 5.1 are satisfied, we let γ_s and γ_1 be the constants that Lemma 5.1 requires for condition (5.7).

Suppose the intervals (6.32) contain the interpolation points t_+^i and t_-^i , respectively. If $q^i \geq \frac{\gamma_1}{\sqrt{n+l_k}}$, then $q^i \geq t_+^i$, and we obtain from (6.8) and (6.9)

$$\begin{aligned} \Gamma_k^i(q^i) &\geq -\frac{1}{2}c_k^i(t_+^i)^2 + c_k^i t_+^i q^i \\ &\geq \frac{1}{2}c_k^i t_+^i q^i \\ &\geq \frac{\gamma_s \sqrt{n+l_k}}{\beta_l} \|x_k - x_*\| c_k^i |q^i| \\ &\geq \gamma_s \sqrt{n+l_k} \|x_k - x_*\| |q^i|, \end{aligned} \quad (6.34)$$

where the last step follows from (6.18). If $q^i \leq -\frac{\gamma_1}{\sqrt{n+l_k}}$, a similar argument also implies (6.34). Now, by (6.19) we have that for any d ,

$$\|\hat{q}_k(d)\|_\infty \geq \|\hat{q}_k(d)\| / \sqrt{n+l_k} = \|d\| / \sqrt{n+l_k}.$$

In particular, if $\|d\| \geq \gamma_1$, then $\|\hat{q}_k(d)\|_\infty \geq \frac{\gamma_1}{\sqrt{n+l_k}}$ and thus (6.34) holds for some i so that

$$\hat{\Gamma}_k(d) = \sum_{i=1}^{n+l_k} \Gamma_k^i(\hat{q}_k^i(d)) \geq \gamma_s \sqrt{n+l_k} \|x_k - x_*\| \|\hat{q}_k(d)\|_\infty \geq \gamma_s \|x_k - x_*\| \|d\|. \quad (6.35)$$

Hence, the condition (5.7) holds. Because $\hat{\Gamma}_k(d)$ is an under estimate of $\frac{1}{2}d^T B_k d$, we obtain (5.8) by noticing Assumption II-d. The conclusion (5.48) follows from Theorem 5.2. \square

7 Final Remarks

In this paper, we proposed a new algorithm for nonlinear programming that uses piecewise linear models in the step computation. One of the crucial ingredients for this approach is the placement of the interpolation points employed in the construction of the piecewise linear models. We presented two sets of placement conditions. For the case when the Hessian approximation B_k is diagonal, we have shown in Section 4 that Placement Conditions A ensure global convergence, while conditions (5.47) guarantee correct active set identification.

It is desirable that a practical algorithm satisfies both conditions. This can be achieved by many strategies, including the following procedure. Let us define

$$\omega_k = \min\{1, \|\nabla f(x_k) - \nabla h(x_k)\lambda_{k+1} - \nabla g(x_k)\mu_{k+1}\| + \|h(x_k)\| + \|g(x_k)^-\| + \|\mu_{k+1}^T g(x_k)\|\}^{1/2}$$

where $(\lambda_{k+1}, \mu_{k+1})$ are the PLA multipliers. At any iteration k , generate interpolation points $t_{k,j}^i$ by any procedure, with the stipulation that these points are uniformly bounded and that for each coordinate $i \in \{1, \dots, n\}$ one of these points is at zero. (In general, the interpolation points should be placed in a region that is expected to contain the next step.) Then check whether for each i there exist interpolation points in the two intervals

$$[c_1\omega_k, c_2\omega_k] \quad \text{and} \quad [-c_1\omega_k, -c_2\omega_k],$$

where $0 < c_1 < c_2$ (say, $c_1 = 0.1, c_2 = 10$). If not, add an interpolation point in each of these intervals, as necessary. Next, verify if conditions (4.5) are satisfied. If not, add an interpolation point in each of the intervals

$$\left[\frac{\partial f(x_k)}{\partial x^i} \frac{1}{b_k^i} - c_3\kappa_0, \frac{\partial f(x_k)}{\partial x^i} \frac{1}{b_k^i} - \kappa_0 \right], \quad \text{and} \quad \left[\frac{\partial f(x_k)}{\partial x^i} \frac{1}{b_k^i} + \kappa_0, \frac{\partial f(x_k)}{\partial x^i} \frac{1}{b_k^i} + c_3\kappa_0 \right]$$

as needed, where $c_3 > 1$ is some constant.

A similar strategy can be used when B_k is defined by the limited memory BFGS update of Section 6.

In conclusion, it is not difficult to design strategies that will satisfy the conditions of our analysis. In practice, additional safeguards are desirable that ensure efficiency and robustness over a wide range of problems, This and other implementation issues are studied in the companion paper [7].

References

- [1] R.H. Bartels, A.R. Conn, and J.W. Sinclair. Minimization techniques for piecewise differentiable functions – the L1 solution to an overdetermined linear system. *SIAM Journal on Numerical Analysis*, 15:224–241, 1978.
- [2] R. H. Byrd, N. I. M. Gould, J. Nocedal, and R. A. Waltz. An algorithm for nonlinear optimization using linear programming and equality constrained subproblems. *Mathematical Programming, Series B*, 100(1):27–48, 2004.

- [3] R. H. Byrd, N. I. M. Gould, J. Nocedal, and R. A. Waltz. On the convergence of successive linear-quadratic programming algorithms. *SIAM Journal on Optimization*, 16(2):471–489, 2006.
- [4] R. H. Byrd, J. Nocedal, and R. Schnabel. Representations of quasi-newton matrices and their use in limited memory methods. *Mathematical Programming*, 63(4):129–156, 1994.
- [5] R. H. Byrd, J. Nocedal, and R. A. Waltz. Steering exact penalty methods. *Optimization Methods and Software*, 23(2), 2008.
- [6] R.H. Byrd, J. Nocedal, and G. López-Calva. A line search exact penalty method using steering rules. Technical report, Optimization Center, Northwestern University, 2009.
- [7] R.H. Byrd, J. Nocedal, R.A. Waltz, and Y. Wu. On the implementation of a method for nonlinear programming based on piecewise linear models. Technical report, Optimization Center, Northwestern University, 2010. forthcoming.
- [8] R.H. Byrd and R.A. Waltz. An active-set algorithm for nonlinear programming using parametric linear programming. *Optimization Methods and Software*, pages 1029–4937, 2009.
- [9] M. R. Celis, J. E. Dennis, and R. A. Tapia. A trust region strategy for nonlinear equality constrained optimization. In P. T. Boggs, R. H. Byrd, and R. B. Schnabel, editors, *Numerical Optimization 1984*, Philadelphia, USA, 1985. SIAM.
- [10] C. M. Chin and R. Fletcher. On the global convergence of an SLP-filter algorithm that takes EQP steps. *Mathematical Programming, Series A*, 96(1):161–177, 2003.
- [11] A. R. Conn, N. I. M. Gould, and Ph. Toint. *Trust-region methods*. MPS-SIAM Series on Optimization. SIAM publications, Philadelphia, Pennsylvania, USA, 2000.
- [12] G.B. Dantzig. Recent advances in linear programming. *Management Science*, 2:131–144, 1956.
- [13] R. Fletcher and S. Leyffer. Nonlinear programming without a penalty function. *Mathematical Programming*, 91:239–269, 2002.
- [14] R. Fletcher and E. Sainz de la Maza. Nonlinear programming and nonsmooth optimization by successive linear programming. *Mathematical Programming*, 43(3):235–256, 1989.
- [15] M.P. Friedlander, N.I.M. Gould, S. Leyffer, and T.S. Munson. A filter active-set trust-region method. Technical Report Preprint ANL/MCS-P1456-097, Argonne National Laboratory, 2007.
- [16] P. E. Gill, W. Murray, and M. A. Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Journal on Optimization*, 12:979–1006, 2002.

- [17] N. I. M. Gould and D. P. Robinson. A second derivative SQP method: Global convergence. Technical Report NA-08/18, Oxford University Computing Laboratory, November 2008. To appear in *SIAM J. Optimization*.
- [18] N. I. M. Gould and D. P. Robinson. A second derivative SQP method: Local convergence. Technical Report NA-08/21, Oxford University Computing Laboratory, December 2008. To appear in *SIAM J. Optimization*.
- [19] J.L. Morales, J. Nocedal, and Y. Wu. A sequential quadratic programming algorithm with an additional equality constrained phase. Technical Report OTC-05, Northwestern University, 2008.
- [20] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer Series in Operations Research. Springer, second edition, 2006.
- [21] C. Oberlin and S. J. Wright. Active constraint identification in nonlinear programming. *SIAM Journal on Optimization*, 17:577–605, 2006.
- [22] M. J. D. Powell and Y. Yuan. A trust region algorithm for equality constrained optimization. *Mathematical Programming*, 49(2):189–213, 1990.
- [23] R. A. Waltz. *Algorithms for large-scale nonlinear optimization*. PhD thesis, Department of Electrical and Computer Engineering, Northwestern University, Evanston, Illinois, USA, 2002.