

IBM Research Report

Adaptive Barrier Strategies for Nonlinear Interior Methods

Jorge Nocedal*, Andreas Wächter, Richard A. Waltz*

IBM Research Division

Thomas J. Watson Research Center

P.O. Box 218

Yorktown Heights, NY 10598

*Department of Electrical and Computer Engineering

Northwestern University

Evanston, IL 60208-3118



Research Division

Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

Adaptive Barrier Strategies for Nonlinear Interior Methods

Jorge Nocedal*

Andreas Wächter[†]

Richard A. Waltz*

February 25, 2005

Abstract

This paper considers strategies for selecting the barrier parameter at every iteration of an interior-point method for nonlinear programming. Numerical experiments suggest that adaptive choices, such as Mehrotra's probing procedure, outperform static strategies that hold the barrier parameter fixed until a barrier optimality test is satisfied. A new adaptive strategy is proposed based on the minimization of a *quality function*. The paper also proposes a globalization framework that ensures the convergence of adaptive interior methods. The barrier update strategies proposed in this paper are applicable to a wide class of interior methods and are tested in the two distinct algorithmic frameworks provided by the IPOPT and KNITRO software packages.

1 Introduction

In this paper we describe interior methods for nonlinear programming that update the barrier parameter adaptively, as the iteration progresses. The goal is to design algorithms that are scale invariant and efficient in practice, and that enjoy global convergence guarantees. The adaptive strategies studied in this paper allow the barrier parameter to increase or decrease at every iteration and provide an alternative to the so-called Fiacco-McCormick approach that fixes the barrier parameter until an approximate solution of the barrier problem is computed. Our motivation for this work stems from our belief that robust interior methods for nonlinear programming must be able to react swiftly to changes of scale in the problem and to correct overly aggressive decreases in the barrier parameter.

Adaptive barrier update strategies are well established in interior methods for linear and convex quadratic programming. The most popular approach of this type is Mehrotra's

*Department of Electrical and Computer Engineering, Northwestern University, Evanston, IL, 60208-3118, USA. These authors were supported by National Science Foundation grants CCR-0219438 and DMI-0422132, and Department of Energy grant DE-FG02-87ER25047-A004.

[†]Department of Mathematical Sciences, IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA.

predictor-corrector method [17]. It computes, at every iteration, a probing (affine scaling) step that determines a target value of the barrier parameter, and then takes a primal-dual step using this target value. A corrector step is added to better follow the trajectory to the solution. Mehrotra’s method has proved to be very effective for linear and convex quadratic programming, but it cannot be shown to be globally convergent. Moreover, its reliability is heavily dependent upon an appropriate choice of the starting point.

When solving nonlinear nonconvex programming problems, much caution must be exercised to prevent the iteration from failing. Non-minimizing stationary points can attract the iteration, and aggressive decreases in the barrier parameter can lead to failure. Our numerical experience shows that the direct extension of Mehrotra’s predictor-corrector method to nonlinear programming does not result in a robust method. As we discuss below, the main source of instability is the corrector step. Other adaptive barrier update strategies have been proposed specifically for nonlinear programming [7, 10, 19, 20, 21].

In this paper we propose new strategies for updating the barrier parameter that are effective in practice and that are supported by a global convergence analysis. To show the generality of our techniques, we implement them in the two different algorithmic contexts provided by the IPOPT [23] and KNITRO [24] software packages.

The global convergence properties of interior methods for nonlinear programming have recently received much attention [2, 7, 12, 13, 15, 19, 20, 22, 28]. Some of these studies focus on the effects of merit functions or filters, and on regularization techniques. With the exception of [7, 19, 20], however, these papers do not consider the numerical or theoretical properties of adaptive barrier update techniques.

Notation. For any vector z , we denote by Z the diagonal matrix whose diagonal entries are given by z . We let e denote the vector of ones, of appropriate dimension, that is, $e = (1, 1, \dots, 1)^T$.

2 Primal-Dual Nonlinear Interior Methods

The problem under consideration will be written as

$$\min_x f(x) \tag{2.1a}$$

$$\text{s.t. } c(x) = 0 \tag{2.1b}$$

$$x \geq 0, \tag{2.1c}$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and $c : \mathbb{R}^n \rightarrow \mathbb{R}^m$ are twice continuously differentiable functions. For conciseness we will refer to interior-point methods for nonlinear programming as “nonlinear interior methods.” A variety of these methods have been proposed in the last 10 years; they differ mainly in some aspects of the step computation and in the globalization scheme. Most of the nonlinear interior methods are related to the simple primal-dual iteration described next, and therefore, our discussion of barrier parameter choices will be phrased in the context of this iteration.

We associate with the nonlinear program (2.1) the barrier problem

$$\min_x \quad \varphi_\mu(x) \equiv f(x) - \mu \sum_{i=1}^n \ln x_i \quad (2.2a)$$

$$\text{s.t.} \quad c(x) = 0, \quad (2.2b)$$

where $\mu > 0$ is the barrier parameter. As is well known, the KKT conditions of the barrier problem (2.2) can be written as

$$\nabla f(x) - A(x)^T y - z = 0 \quad (2.3a)$$

$$Xz - \mu e = 0 \quad (2.3b)$$

$$c(x) = 0, \quad (2.3c)$$

together with

$$x \geq 0, \quad z \geq 0. \quad (2.4)$$

Here $A(x)$ denotes the Jacobian matrix of the constraint functions $c(x)$.

Applying Newton's method to (2.3), in the variables (x, y, z) , gives the *primal-dual* system

$$\begin{bmatrix} \nabla_{xx}^2 \mathcal{L} & -A(x)^T & -I \\ Z & 0 & X \\ A(x) & 0 & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta z \end{bmatrix} = - \begin{bmatrix} \nabla f(x) - A(x)^T y - z \\ Xz - \mu e \\ c(x) \end{bmatrix}, \quad (2.5)$$

where \mathcal{L} denotes the Lagrangian of the nonlinear program, that is,

$$\mathcal{L}(x, y, z) = f(x) - y^T c(x) - z^T x. \quad (2.6)$$

After the step $\Delta = (\Delta x, \Delta y, \Delta z)$ has been determined, we compute primal and dual steplengths, α_p and α_d , and define the new iterate (x^+, y^+, z^+) as

$$x^+ = x + \alpha_p \Delta x, \quad y^+ = y + \alpha_d \Delta y, \quad z^+ = z + \alpha_d \Delta z. \quad (2.7)$$

The steplengths are computed in two stages. First we compute

$$\alpha_x^{\max} = \max\{\alpha \in (0, 1] : x + \alpha \Delta x \geq (1 - \tau)x\} \quad (2.8a)$$

$$\alpha_z^{\max} = \max\{\alpha \in (0, 1] : z + \alpha \Delta z \geq (1 - \tau)z\}, \quad (2.8b)$$

with $\tau \in (0, 1)$ (e.g. $\tau = 0.995$). Next, we perform a backtracking line search that computes the final steplengths

$$\alpha_p \in (0, \alpha_x^{\max}], \quad \alpha_d \in (0, \alpha_z^{\max}], \quad (2.9)$$

providing sufficient decrease of a merit function or ensuring acceptability by a filter.

The other major ingredient in this simple primal-dual iteration is the procedure for choosing the barrier parameter μ . Two types of barrier update strategies have been studied in the literature: adaptive and static. Adaptive strategies [7, 10, 20, 21] allow changes in the barrier parameter at every iteration, and often achieve scale invariance, but as already mentioned they do not enjoy global convergence properties. (The analyses presented in

[7, 20] provide certain convergence results to stationary points, but these methods do not promote convergence toward minimizers.)

The most important static strategy is the so-called Fiacco-McCormick approach that fixes the barrier parameter until an approximate solution of the barrier problem is computed. It has been employed in various nonlinear interior algorithms [1, 3, 9, 11, 23, 25, 27] and has been implemented, for example, in the IPOPT and KNITRO software packages. The Fiacco-McCormick strategy provides a framework for establishing global convergence [2, 22], but suffers from important limitations. It can be very sensitive to the choice of the initial point, the initial value of the barrier parameter and the scaling of the problem, and it is often unable to recover when the iterates approach the boundary of the feasible region prematurely. Our numerical experience with IPOPT and KNITRO suggests that more dynamic update strategies are needed to improve the robustness and speed of nonlinear interior methods.

3 Choosing the Barrier Parameter

In this section we discuss two adaptive barrier strategies proposed in the literature and compare them numerically with the static Fiacco-McCormick approach. These numerical results will motivate the discussion of the following sections.

Given an iterate (x, y, z) , consider an interior method that computes primal-dual search directions by (2.5). The most common approach for choosing the barrier parameter μ that appears on the right hand side of (2.5) is to make it proportional to the current complementarity value, that is,

$$\mu = \sigma \frac{x^T z}{n}, \quad (3.1)$$

where $\sigma > 0$ is a *centering parameter* and n denotes the number of variables. Mehrotra's predictor-corrector (MPC) method [17] for linear programming determines the value of σ using a preliminary step computation (an affine scaling step). We now describe a direct extension of Mehrotra's strategy to the nonlinear programming case.

First, we calculate an affine scaling step

$$(\Delta x^{\text{aff}}, \Delta y^{\text{aff}}, \Delta z^{\text{aff}}) \quad (3.2)$$

by setting $\mu = 0$ in (2.5), that is,

$$\begin{bmatrix} \nabla_{xx}^2 \mathcal{L} & -A(x)^T & -I \\ Z & 0 & X \\ A(x) & 0 & 0 \end{bmatrix} \begin{bmatrix} \Delta x^{\text{aff}} \\ \Delta y^{\text{aff}} \\ \Delta z^{\text{aff}} \end{bmatrix} = - \begin{bmatrix} \nabla f(x) - A(x)y - z \\ Xz \\ c(x) \end{bmatrix}. \quad (3.3)$$

We then compute α_x^{aff} and α_z^{aff} to be the longest steplengths that can be taken along the direction (3.2) before violating the non-negativity conditions $(x, z) \geq 0$, with an upper bound of 1. Explicit formulae for these values are given by (2.8) with $\tau = 1$.

Next, we define μ^{aff} to be the value of complementarity that would be obtained by a full step to the boundary, that is,

$$\mu^{\text{aff}} = (x + \alpha_x^{\text{aff}} \Delta x^{\text{aff}})^T (z + \alpha_z^{\text{aff}} \Delta z^{\text{aff}}) / n, \quad (3.4)$$

and set the centering parameter to be

$$\sigma = \left(\frac{\mu^{\text{aff}}}{x^T z/n} \right)^3. \quad (3.5)$$

This heuristic choice of σ is based on experimentation with linear programming problems, and has proved to be effective for convex quadratic programming as well. Note that when good progress is made along the affine scaling direction, we have $\mu^{\text{aff}} \ll x^T z/n$, so the σ obtained from this formula is small—and conversely.

Mehrotra’s algorithm also computes a *corrector* step, but in this section we take the view that the corrector is not part of the selection of the barrier parameter, and is simply a mechanism for improving the quality of the step. In Section 7 we study the complete MPC algorithm including the corrector step.

Other adaptive procedures of the form (3.1) have been proposed specifically for nonlinear interior methods [7, 10, 20, 21]. The strategy employed in the LOQO software package [21] is particularly noteworthy because of its success in practice. It defines σ as

$$\sigma = 0.1 \min \left(0.05 \frac{1-\xi}{\xi}, 2 \right)^3, \quad \text{where } \xi = \frac{\min_i \{x_i z_i\}}{x^T z/n}. \quad (3.6)$$

Note that ξ measures the deviation of the smallest complementarity product $x_i z_i$ from the average. When $\xi = 1$ (all individual products are equal to their average) we have that $\sigma = 0$ and the algorithm takes an aggressive step. Note that the rule (3.6) always chooses $\sigma \leq 0.8$, so that even though the value of μ may increase from one iteration to the next, it will never be chosen to be larger than the current complementarity value $x^T z/n$.

Let us compare the numerical performance of these two adaptive strategies with the static Fiacco-McCormick approach. To better measure the impact of the choice of barrier parameter, that is, to try to distinguish it from other algorithmic features, we use both the IPOPT and KNITRO software packages in our tests. These codes implement significantly different variations of the simple primal-dual iteration (2.5). IPOPT was implemented for these tests without a line search; only the fraction to the boundary rule (2.8) was imposed on the primal and dual steplengths. For KNITRO we used the Interior/Direct option (we will refer to this version as KNITRO-DIRECT henceforth), which invokes a line search approach that is occasionally safeguarded by a trust region iteration (for example, when negative curvature is encountered) [25].

The barrier parameter strategies tested in our experiments are as follows:

- *LOQO rule.* The barrier parameter is chosen by (3.1) and (3.6).
- *Mehrotra probing.* At every iteration, the barrier parameter μ is given by (3.1) and (3.5). Since this requires the computation of the affine scaling step (3.2), this strategy is more expensive than the LOQO rule. For KNITRO-DIRECT, in the iterations in which the safeguarding trust region algorithm is invoked the barrier parameter is computed by the LOQO rule instead of Mehrotra probing. This is done because Mehrotra probing is expensive to implement in the trust region algorithm, which uses a conjugate gradient iteration.




Figure 1: Results for four barrier parameter updating strategies.

- *MPC*. The complete Mehrotra predictor-corrector algorithm as described in Section 7. As in the Mehrotra probing rule, when KNITRO-DIRECT falls back on the safeguarded trust region algorithm, the barrier parameter is computed using the LOQO rule for efficiency, and no corrector step is used.
- *Monotone*. (Also known as the Fiacco-McCormick approach.) The barrier parameter is fixed, and a series of primal-dual steps is computed, until the optimality conditions for the barrier problem are satisfied to some accuracy. At this point the barrier parameter is decreased. IPOPT and KNITRO implement somewhat different variations of this monotone approach; see [23, 25] for details about the initial value of μ , the rule for decreasing μ , and the form of the barrier stop tests.

For the numerical comparison, we select all the nonlinear programming problems in the CUTEr test set that contain at least one general inequality or bound constraint¹. This gives a total of 599 problems. Figure 1(a) reports the total number of iterations for IPOPT, and Figure 1(b) shows the total number of function evaluations for KNITRO, both comparing the performance of the four barrier strategies. Measuring iterations and function evaluations give similar profiles and we provide both for greater variety. All the plots in the paper use the logarithmic performance profiles proposed by Dolan and Moré [6]. To account for the fact that different local solutions might be approached, problems with significantly different final objective function values for successful runs were excluded.

Note that the adaptive strategies are more robust than the monotone approach, and that Mehrotra probing appears to be the most successful in terms of iterations or function

¹excluding those problems that seem infeasible, unbounded, or are given with initial points at which the model functions cannot be evaluated; see [23].

evaluations. The complete MPC algorithm is very fast on some problems, but is not sufficiently robust. This is clear in Figure 1(a) which reports the performance of the pure MPC strategy implemented in IPOPT; it is less clear in Figure 1(b), but as mentioned above, in the KNITRO implementation the MPC approach is replaced by the LOQO rule (without corrector step) when the safeguarded trust region step is invoked. The reason for the lack of robustness of the MPC strategy will be discussed in Section 7. Motivated by these results, we give further attention to adaptive choices of the barrier parameter.

4 Quality Functions

We now consider an approach that selects μ by approximately minimizing a *quality function*. We let $\mu = \sigma \frac{x^T z}{n}$, where the centering parameter $\sigma \geq 0$ is to be varied, and define $\Delta(\sigma)$ to be the solution of the primal-dual equations (2.5) as a function of σ . We also let $\alpha_x^{\max}(\sigma), \alpha_z^{\max}(\sigma)$ denote the steplengths satisfying the fraction to the boundary rule (2.8) for $\Delta = \Delta(\sigma)$, and let

$$\begin{aligned} x(\sigma) &= x + \alpha_x^{\max}(\sigma)\Delta x(\sigma), \\ y(\sigma) &= y + \alpha_z^{\max}(\sigma)\Delta y(\sigma), \quad z(\sigma) = z + \alpha_z^{\max}(\sigma)\Delta z(\sigma). \end{aligned}$$

Our approach is to choose the value of σ that gives the best overall step, as measured by the KKT conditions for the nonlinear program (2.1). For example, we could try to choose σ so as to minimize the nonlinear quality function

$$\begin{aligned} q_N(\sigma) &= \|\nabla f(x(\sigma)) - A(x(\sigma))^T y(\sigma) - z(\sigma)\|^2 + \|c(x(\sigma))\|^2 \\ &\quad + \|Z(\sigma)X(\sigma)e\|^2. \end{aligned} \tag{4.1}$$

The evaluation of q_N is, however, very expensive since it requires the evaluation of problem functions and derivatives for every value of σ . A much more affordable alternative is to use the *linear quality function*

$$\begin{aligned} q_L(\sigma) &= (1 - \alpha_z^{\max}(\sigma))^2 \|\nabla f(x) - A(x)^T y - z\|^2 + (1 - \alpha_x^{\max}(\sigma))^2 \|c(x)\|^2 \\ &\quad + \|(X + \alpha_x^{\max}(\sigma)\Delta X(\sigma))(Z + \alpha_z^{\max}(\sigma)\Delta Z(\sigma))e\|^2, \end{aligned} \tag{4.2}$$

where $\Delta X(\sigma)$ is the diagonal matrix with $\Delta x(\sigma)$ on the diagonal, and similarly for $\Delta Z(\sigma)$. The function q_L has been obtained from (4.1) by assuming that primal and dual feasibility (i.e. the first two terms in (4.1)) are linear functions, as is the case in linear programming. Note that $q_L(\sigma)$ is not a convex function of σ because the steplengths $\alpha_x^{\max}, \alpha_z^{\max}$ depend on σ in a complicated manner. The dominant cost in the evaluation of q_L lies in the computation of the maximal steplengths $\alpha_x^{\max}, \alpha_z^{\max}$ and the last term in (4.2), which requires a few vector operations. We define the quality function using squared norms to severely penalize any large components in the KKT error. In Section 6 we discuss the choice of norms and other details of implementation. We have also experimented with a *quadratic quality function* that is based on the KKT conditions for a quadratic program, and that like (4.2), does not require additional function or gradient evaluations. However, it did not perform as well as the linear quality function, for reasons that require further investigation.

In Section 6 we describe a procedure for approximately minimizing the scalar function $q_L(\sigma)$. Before presenting our numerical results with the quality function, we consider the issue of how to guarantee the global convergence of nonlinear interior methods that choose the barrier parameter adaptively.

5 A Globalization Method

The adaptive strategies described in Section 3 can be seen from the numerical results in that section to be quite robust. (We show in the next section this is also the case with the quality function approach.) Yet, since the barrier parameter is allowed to change at every iteration in these algorithms, there is no mechanism that enforces global convergence of the iterates. In contrast, the monotone barrier strategy employed in the Fiacco-McCormick approach allows us to establish global convergence results by combining two mechanisms. First, the algorithms used to minimize a given barrier problem (2.2) use a line search or trust region to enforce a decrease in the merit function (as in KNITRO) or to guarantee acceptability by a filter (as in IPOPT). This ensures that an optimality test for the barrier function is eventually satisfied to some tolerance ϵ . Second, by repeating this minimization process for decreasing values of μ and ϵ that converge to zero, one can establish global convergence results [2, 8] to stationary points of the nonlinear programming problem (2.1).

We now propose a globalization scheme that monitors the performance of the iterations in reference to an optimality measure for the nonlinear program (2.1). As long as the adaptive primal-dual steps make sufficient progress towards the solution, the algorithm is free to choose a new value for the barrier parameter at every iteration. We call this the *free mode*. However, if the iteration fails to maintain progress, then the algorithm reverts to a *monotone mode*, in which a Fiacco-McCormick strategy is applied. Here, the value of the barrier parameter remains fixed, and a robust globalization technique (e.g., based on a merit function or a filter) is employed to ensure progress for the corresponding barrier problem. Once the barrier problem is approximately minimized, the barrier parameter is decreased. The monotone mode continues until an iterate is generated that makes sufficient progress for the original problem, at which point the free mode resumes.

The goal of our globalization scheme is to interfere with adaptive steps as little as possible because, as already noted, they result in fairly reliable iterations. As a measurement of the progress in the optimization of the nonlinear program (2.1), we monitor the KKT error of the original nonlinear program,

$$\Phi(x, y, z) = \|\nabla f(x) - A(x)^T y - z\|^2 + \|c(x)\|^2 + \|ZXe\|^2. \quad (5.1)$$

We require that this measure be reduced by a factor of $\kappa \in (0, 1)$ over at most a fixed number of l^{\max} iterations, as long as the algorithm is in the free mode. Note that a convergent sequence of points (x, y, z) gives $\Phi(x, y, z) \rightarrow 0$ only if the limit point satisfies the first order optimality conditions for the nonlinear program (2.1).

We now formally state the proposed globalization procedure.

Globalization Framework

Given (x_0, y_0, z_0) with $(x_0, z_0) > 0$, a constant $\kappa \in (0, 1)$ and an integer $l^{\max} \geq 0$.

Repeat

Choose a target value of the barrier parameter μ_k , based on any rule.

Compute the primal dual search direction from (2.5).

Determine step sizes $\alpha_p \in (0, \alpha_x^{\max}]$ and $\alpha_d \in (0, \alpha_z^{\max}]$.

Compute the new trial iterate $(\tilde{x}_{k+1}, \tilde{y}_{k+1}, \tilde{z}_{k+1})$ from (2.7).

Compute the KKT error $\tilde{\Phi}_{k+1} \equiv \Phi(\tilde{x}_{k+1}, \tilde{y}_{k+1}, \tilde{z}_{k+1})$.

Set $M_k = \max\{\Phi_{k-l}, \Phi_{k-l+1}, \dots, \Phi_k\}$ with $l = \min\{k, l^{\max}\}$.

If $\tilde{\Phi}_{k+1} \leq \kappa M_k$

Accept $(\tilde{x}_{k+1}, \tilde{y}_{k+1}, \tilde{z}_{k+1})$ as the new iterate, and set $\Phi_{k+1} \leftarrow \tilde{\Phi}_{k+1}$.

Set $k \leftarrow k + 1$ and return to the beginning of the loop.

else

Start Monotone Mode:

Starting from $(\tilde{x}_{k+1}, \tilde{y}_{k+1}, \tilde{z}_{k+1})$, and for an initial value $\bar{\mu}$ of the barrier parameter, solve a sequence of barrier problems with a monotonically decreasing sequence of barrier parameters less than $\bar{\mu}$, to obtain a new iterate $(x_{k+1}, y_{k+1}, z_{k+1})$ such that

$$\Phi_{k+1} \equiv \Phi(x_{k+1}, y_{k+1}, z_{k+1}) \leq \kappa M_k.$$

Set $k \leftarrow k + 1$ and resume the free mode at the beginning of the loop.

end if

End (repeat).

Since this framework ensures that the optimality measure Φ is reduced by a factor of $\kappa < 1$ in at most every l^{\max} iterations, it is clear that $\Phi_k \rightarrow 0$. Consequently, all limit points of the sequence of iterates satisfy the first order optimality conditions of the nonlinear program (2.1).

Note that even in the free mode we might want to choose steplengths α_p, α_d that are shorter than the maximal step sizes $\alpha_x^{\max}, \alpha_z^{\max}$. In our implementations, we perform a line search to enforce progress in a merit function or a filter, both of which are defined with respect to the barrier problem (2.2) corresponding to the current value μ_k . This is possible, because the adaptive primal-dual search direction is compatible with this barrier problem. The penalty parameter of the merit function, or the history in the filter, are reset at every free iteration because the barrier problem itself changes at every free iteration.

In the monotone mode, it is not required to solve each barrier problem to the specified tolerance before checking whether the method can revert to the free mode. Instead, we compute the optimality error $\Phi(x, y, z)$ for all intermediate iterates in the monotone mode, and return to the free mode, as soon as $\Phi(x, y, z) \leq \kappa M_k$.

Finally, we point out that criteria other than an optimality measure Φ can be incorporated in the above globalization framework. For example, the decision of when to switch to the monotone mode could be based on a two-dimensional filter involving the original objective function $f(x)$ and the constraint violation $\|c(x)\|$. In this setting, the algorithm would stay in the free mode as long as the trial iterate reduces either the value of the objective

function or the constraint violation sufficiently compared to all previous (free) iterates. Note that, by using this strategy, the algorithm might never return from the monotone mode. We implemented this option in IPOPT and observed that it was similar in performance to the optimality measure Φ .

6 Numerical Results

In order to determine an appropriate value of the barrier parameter (3.1) using the quality function approach, we need to (approximately) minimize $q_L(\sigma)$. Due to the complicated dependence of the steplengths $\alpha_x^{\max}(\sigma), \alpha_z^{\max}(\sigma)$ on the parameter σ , it does not seem possible to obtain an analytic solution to this one-dimensional minimization problem within a reasonable amount of computation time. Therefore, we implemented search strategies using only function evaluations of $q_L(\sigma)$. The search is relatively inexpensive since an evaluation of q_L requires only a few vector-vector operations. An important consideration is that σ be allowed to take on values greater than one so that the algorithm can recover from overly aggressive reductions of the barrier parameter.

The search heuristic implemented in IPOPT uses a golden bisection procedure (see, e.g. [16]), ignoring the fact that q_L is nonconvex in general. First, golden bisection is performed in the interval $[\sigma^{\min}, 1]$, where σ^{\min} is chosen to correspond to a minimal permissible value of the barrier parameter μ^{\min} (in our implementations, $\mu^{\min} = 10^{-9}$). If the minimizer within this search interval appears to be at $\sigma = 1$, a second bisection search is performed in the interval $[1, 100]$. Each bisection procedure terminates if 10 evaluations of the quality functions are performed, or if the search interval $[a, b]$ becomes smaller than $b \times 10^{-3}$.

The search implemented in KNITRO proceeds in two phases. First a rough global search is performed by sampling values of σ distributed in the range $[\sigma^{\min}, 1000]$ (the number of sample points depends on the value of the average complementarity at the current iterate, but a typical number of sample points is between 5 and 12). This search returns σ_1 , the trial point giving the lowest value of q_L . If $\sigma_1 \leq 0.99$ or $\sigma_1 \geq 100$, then we take σ_1 as the approximate minimizer of q_L . Otherwise, we perform a more refined search in the range $[0.5, 1]$ by testing the values $\sigma = 1 - \frac{1}{2^l}$, $l = 1..5$, and denote the best value from this search as σ_2 . We then take the better of σ_1 and σ_2 as the approximate minimizer. The idea behind the more refined search is the following. If the rough search results in a minimizer which is greater than 1, then before taking this value we want to check more thoroughly whether we can find a better value of $\sigma < 1$ which aims to decrease the current complementarity.

We now discuss the choice of norms in the quality function (4.2) and in the optimality measure (5.1). (For consistency, we use the same norms and scaling factors in (4.2) and (5.1).) In IPOPT we use the 2-norm, and each of the three terms is divided by the number of elements in the vectors whose norms are being computed. In KNITRO, the first two terms in (4.2) and (5.1) use the infinity-norm, and the complementarity term uses the 1-norm divided by n . In addition, in an attempt to make the function Φ less scale dependent, KNITRO scales the second (primal infeasibility) term by the maximum constraint violation at the initial point (if this value is greater than one), and scales the first and third terms (dual infeasibility and complementarity) by $\max(1, \|\nabla f(x_k)\|_\infty)$. These are the scaling

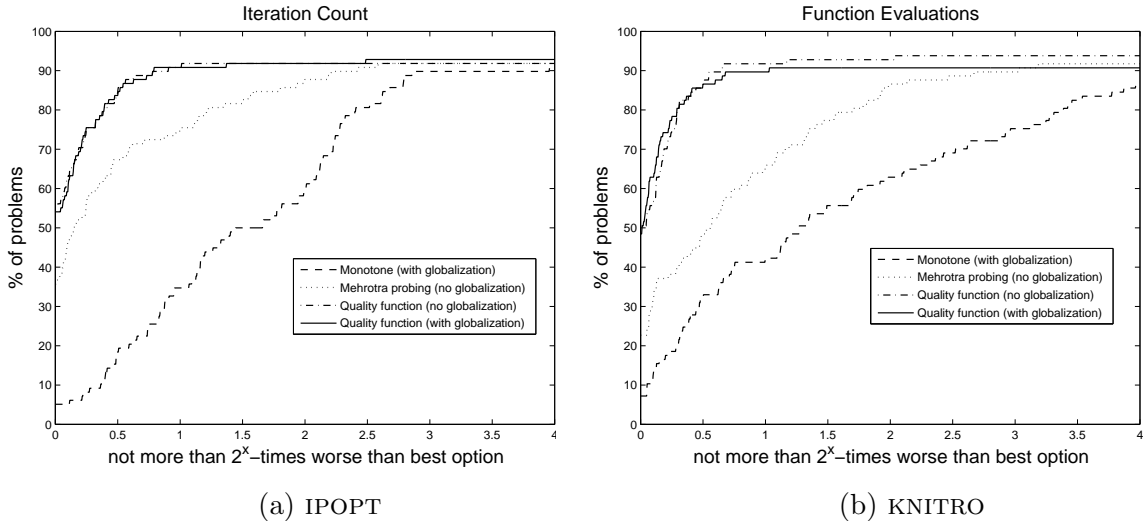


Figure 2: Results for the NETLIB test set.

factors used in the KNITRO termination test [25].

The tests involving IPOPT were run on a Dual-Pentium III, 1GHz machine running Linux. The KNITRO tests were run on a machine with an AMD Athlon XP 3200+ 2.2GHz processor running Linux. For both codes, the maximum number of iterations was set to 3000 and the time limit was set to 1800 CPU seconds. No scaling of the test problems was performed by the codes since one of our goals is to develop barrier parameter update strategies that are insensitive to problem scaling. The tests were run using the latest development versions of IPOPT and KNITRO as of February 2005.

The first results are for the linear programming problems in the NETLIB collection, as specified in the CUTER test set [14]. No preprocessing was performed, and no initial point strategy was employed (i.e., the default starting point $x_0 = (0, \dots, 0)$ was used). Figure 2 compares the performance of the quality function approach (both with and without the globalization framework of Section 5) with two of the strategies described in Section 3, namely the monotone method and the (unglobalized) Mehrotra probing heuristic. Even though our focus is on nonlinear optimization, linear programming problems are of interest since they allow us to assess the effectiveness of the quality function in a context in which they exactly predict the KKT error. It is apparent from Figure 2 that the quality function approach is very effective on the NETLIB test set, and that the globalization framework produces only a slight decrease in efficiency. We note, however, that the quality function approach requires extra work, and hence may not be the fastest in terms of CPU time.

The performance of the four barrier update strategies on nonlinear programming problems with at least one inequality or bound constraint from the CUTER collection is reported in Figure 3. The quality function approach again performs significantly better than the monotone method, but this time its advantage over Mehrotra probing is less pronounced.

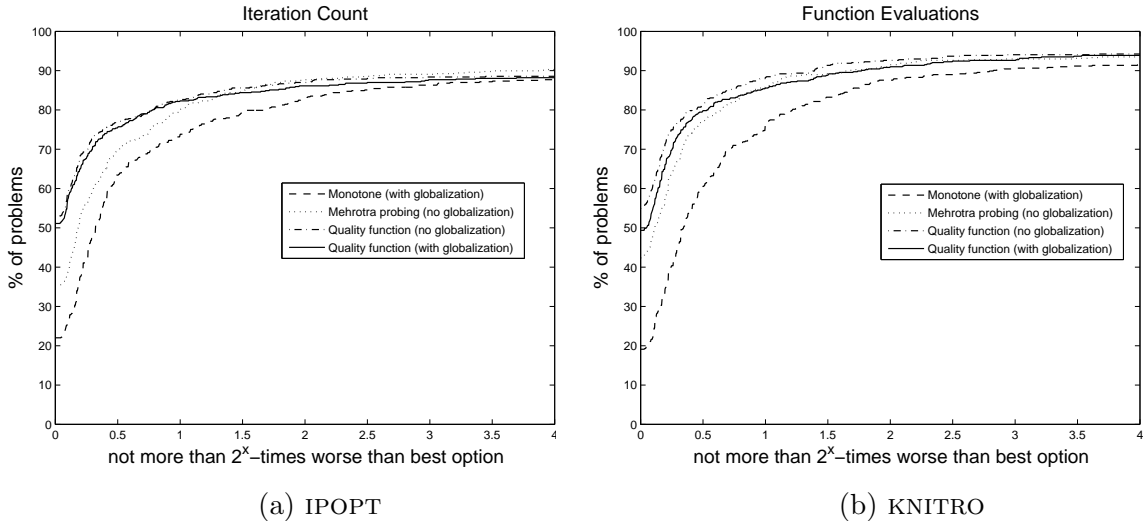


Figure 3: Results for CUTEr test set.

These results are of practical importance since the monotone method is currently the default strategy in both IPOPT and KNITRO; our tests suggest that significant improvements can be expected by using adaptive strategies. We believe that the quality function approach shows potential for future advances.

7 Corrector Steps

The numerical results of Section 3 indicate that, when solving nonlinear problems, including the corrector step in Mehrotra’s method (the MPC method) is not beneficial. This is in stark contrast with the experience in linear programming and convex quadratic programming, where the corrector step is known to accelerate the interior-point iteration without degrading its robustness. In this section we study the effect of the corrector step and find that it can also be harmful in the linear programming and quadratic programming cases *if* an initial point strategy is not used. These observations are relevant because in nonlinear programming it is much more difficult to find a good starting point.

Let us begin by considering the linear programming case. There are several ways of viewing the MPC method in this context. One is to consider the step computation as taking place in three stages (see, e.g., [26]). First, the algorithm computes the affine scaling step (3.2) and uses it to determine the target value of the barrier parameter $\mu = \sigma \frac{x^T z}{n}$, where σ is given by (3.5). Next, the algorithm computes a primal-dual step, say Δ^{pd} , from (2.5) using that value of μ . Finally, a corrector step Δ^{corr} is computed by solving (2.5) with the right hand side given by

$$(0, \Delta X^{\text{aff}} \Delta Z^{\text{aff}} e, 0)^T, \quad (7.1)$$

where ΔX^{aff} is the diagonal matrix with diagonal entries given by Δx^{aff} , and similarly for ΔZ^{aff} . The complete MPC step is the sum of the primal-dual and corrector steps. We can compute it by adding the right hand sides and solving the following system:

$$\begin{bmatrix} \nabla_{xx}^2 \mathcal{L} & -A^T(x) & -I \\ Z & 0 & X \\ A(x) & 0 & 0 \end{bmatrix} \begin{bmatrix} \Delta x^{\text{mpc}} \\ \Delta y^{\text{mpc}} \\ \Delta z^{\text{mpc}} \end{bmatrix} = - \begin{bmatrix} \nabla f(x) - A^T(x)y - z \\ Xz - \mu e + \Delta X^{\text{aff}} \Delta Z^{\text{aff}} e \\ c(x) \end{bmatrix}. \quad (7.2)$$

The new iterate (x^+, y^+, z^+) of the MPC method is given by (2.7)-(2.8) with $\Delta = (\Delta x^{\text{mpc}}, \Delta y^{\text{mpc}}, \Delta z^{\text{mpc}})$.

Alternative views of the MPC method are possible by the linearity of the step computation: we can group the right hand side in (7.2) in different ways and thereby interpret the step as the sum of different components. Yet all these views point out to an inconsistency in the MPC approach, as we now discuss.

First of all, let us justify the definition of the corrector step. In the linear programming case, primal and dual feasibility are linear functions and hence vanish at the full affine scaling step given by

$$(x, y, z) + (\Delta x^{\text{aff}}, \Delta y^{\text{aff}}, \Delta z^{\text{aff}}). \quad (7.3)$$

The complementarity term takes on the following value at the full affine scaling step:

$$(X + \Delta X^{\text{aff}})(Z + \Delta Z^{\text{aff}}) = \Delta X^{\text{aff}} \Delta Z^{\text{aff}}.$$

It follows from this equation that the value of the right hand side vector in (2.5) at the full affine scaling step (7.3) is given by (7.1). Thus the corrector step can be viewed as a modified Newton step taken from the point (7.3) and using the primal-dual matrix evaluated at the current iterate (x, y, z) .

The inconsistency in the MPC approach arises because the corrector step, which is designed to improve the full affine scaling step, is applied at the primal-dual point; see Figure 4. In some circumstances, this mismatch can cause poor steps. In particular, we have observed that if the affine scaling step is very long, in the sense that the steplengths (2.8) are very small, and if the corrector step is even larger, then the addition of the corrector step to the primal-dual step (2.7) can significantly increase the complementarity value $x^T z$. This behaviour can be sustained and lead to very slow convergence or failure, as shown in Table 1. The results in this table were obtained using PCx [5], an interior-point code for linear programming that implements the MPC method, applied to problem FORPLAN from the NETLIB collection. We disabled the initial point strategy of PCx and set the initial point to $x = e, z = e$. Note from Table 1 that the affine scaling and corrector steps appear to grow without bound, and examination of the results shows that the dual variables diverge.

To provide further support to the claim that the corrector step can be harmful we ran the complete set of test problems (94 in all) in the NETLIB collection. Using the default settings, which includes a strategy for computing a good starting point, PCx solved 90 problems, and terminated very close to the solution in the remaining 4 cases. Next we disabled the initial point strategy and set the initial point to $x = e, z = e$. PCx was now




Figure 4: An unfavorable corrector step.

Iter	Primal Obj	Dual Obj	PriInf	DualInf	α_x^{\max}	α_z^{\max}	$\log(\frac{x^T z}{n})$	$\ \Delta^{\text{aff}}\ $	$\ \Delta^{\text{mpc}}\ $
0	9.0515e+01	-4.8813e+06	1.0e-00	1.9e+00	0.0e+00	0.0e+00	0.00	0.0e+00	0.0e+00
1	9.0216e+01	-1.3664e+08	1.0e-00	1.9e+00	8.6e-13	5.0e-12	0.08	5.6e+06	1.2e+13
2	9.0403e+01	-3.3916e+08	1.0e-00	1.9e+00	7.3e-13	4.8e-13	0.16	9.1e+07	1.9e+14
3	9.0769e+01	-1.1343e+10	1.0e-00	1.9e+00	4.0e-12	1.2e-11	1.18	2.2e+08	3.9e+14
4	9.0860e+01	-1.8010e+11	1.0e-00	1.9e+00	1.5e-12	5.0e-12	2.35	8.0e+09	1.4e+16
5	9.1312e+01	-2.9307e+12	1.0e-00	1.9e+00	4.3e-12	5.1e-12	3.56	1.3e+11	2.2e+17
6	9.1710e+01	-8.2787e+13	1.0e-00	1.9e+00	6.0e-12	9.1e-12	5.01	2.1e+12	3.6e+18
7	9.2036e+01	-1.5505e+15	1.0e-00	1.9e+00	7.5e-12	6.0e-12	6.28	5.9e+13	1.0e+20
8	9.2282e+01	-6.8149e+16	1.0e-00	1.9e+00	7.0e-12	1.4e-11	7.93	1.1e+15	1.9e+21
9	9.2279e+01	-4.4155e+18	1.0e-00	1.9e+00	9.2e-12	2.1e-11	9.74	4.8e+16	8.3e+22
10	9.2244e+01	-2.8697e+20	1.0e-00	1.9e+00	6.8e-12	2.1e-11	11.55	3.1e+18	5.4e+24
11	9.2381e+01	-3.1118e+22	1.0e-00	1.9e+00	1.1e-11	3.6e-11	13.58	2.0e+20	3.5e+26
12	9.2462e+01	-7.0471e+24	1.0e-00	2.2e+01	6.2e-12	7.6e-11	15.94	2.2e+22	3.8e+28
13	9.2523e+01	-9.9820e+26	1.0e-00	2.8e+03	1.4e-11	4.7e-11	18.09	5.0e+24	8.6e+30
14	9.2605e+01	-1.1959e+30	1.0e-00	2.2e+01	2.1e-11	4.0e-10	21.17	7.1e+26	1.2e+33

Table 1: Output for NETLIB problem FORPLAN for default PCx with bad starting point




Figure 5: Results on the NETLIB test set for three corrector step strategies implemented in PCx. The initial point was set to $x = e, z = e$.

able to solve only 28 problems (and in only 3 additional cases terminated very close to the solution).

We repeated the experiment, using the initial point $x = e, z = e$, but this time removing the corrector step; this corresponds to the algorithm called *Mehrotra probing* in Section 3. We also tested a variant that we call *conditional MPC* in which the corrector step is employed in the MPC method only if it does not result in an increase of complementarity by a factor larger than 2. The results, in terms of iterations, are reported in Figure 5. Note the dramatic increase in robustness of both strategies, compared with the MPC algorithm. The conditional MPC strategy is motivated by the observation that harmful effects of the corrector steps manifest themselves in a significant increase in complementarity.

Finally we compare the monotone and quality function approaches described in Section 3 with the conditional MPC approach on the nonlinear programming problems used in that section. The conditional MPC method is now implemented so as to reject corrector steps that increase complementarity (this more conservative approach appears to be more suitable in the nonlinear case). Furthermore, if the conditional MPC step does not pass the merit function or filter acceptance test for the current barrier problem, the corrector step is also rejected, and the backtracking line search for the regular primal-dual step is executed. The results, given in Figure 6, indicate that this conditional MPC method requires fewer iterations and function evaluations, and is not less robust, than the other strategies.




Figure 6: Results for safeguarded corrector steps.

8 Final Remarks

A multi-objective quality function was used by Meszaros [18] to estimate the optimal steplengths for interior methods for quadratic programming. His solution technique is based on exploiting the properties of the efficient frontier of the multi-objective optimization problem. The numerical results show a consistent (but modest) improvement over using equal steplengths. Curtis et al. [4] study a quality function that is much closer to the one discussed in Section 4. They formulate the problem of finding optimal steplengths for quadratic programming using the quality function (4.1) specialized to the quadratic programming case. In this paper we have gone beyond the approaches in [4, 18] in that the quality function is used to determine the barrier parameter, and also indirectly, the steplengths.

Acknowledgment. We would like to thank Richard Byrd for his many valuable comments and suggestions during the course of this work.

References

- [1] J. Betts, S. K. Eldersveld, P. D. Frank, and J. G. Lewis. An interior-point nonlinear programming algorithm for large scale optimization. Technical report MCT TECH-003, Mathematics and Computing Technology, The Boeing Company, P.O. Box 3707, Seattle WA 98124-2207, 2000.

- [2] R. H. Byrd, J.-Ch. Gilbert, and J. Nocedal. A trust region method based on interior point techniques for nonlinear programming. *Mathematical Programming*, 89(1):149–185, 2000.
- [3] R. H. Byrd, M. E. Hribar, and J. Nocedal. An interior point algorithm for large scale nonlinear programming. *SIAM Journal on Optimization*, 9(4):877–900, 1999.
- [4] Curtis, F, J. Nocedal, and R. A. Waltz. Steplength selection in interior methods. Technical report, Northwestern University, Evanston, IL, USA, 2005.
- [5] Czyzyk, J., S. Mehrotra, and S. J. Wright. PCx User Guide. Technical report, Argonne National Laboratory, Argonne, IL, USA, 1996.
- [6] E. D. Dolan and J. J. Moré. Benchmarking optimization software with performance profiles. *Mathematical Programming, Series A*, 91:201–213, 2002.
- [7] A. S. El-Bakry, R. A. Tapia, T. Tsuchiya, and Y. Zhang. On the formulation and theory of the Newton interior-point method for nonlinear programming. *Journal of Optimization Theory and Applications*, 89(3):507–541, June 1996.
- [8] A. V. Fiacco and G. P. McCormick. *Nonlinear Programming: Sequential Unconstrained Minimization Techniques*. J. Wiley and Sons, Chichester, England, 1968. Reprinted as *Classics in Applied Mathematics 4*, SIAM, Philadelphia, USA, 1990.
- [9] A. Forsgren and P. E. Gill. Primal-dual interior methods for nonconvex nonlinear programming. *SIAM Journal on Optimization*, 8(4):1132–1152, 1998.
- [10] D. M. Gay, M. L. Overton, and M. H. Wright. A primal-dual interior method for nonconvex nonlinear programming. In Y. Yuan, editor, *Advances in Nonlinear Programming (Beijing, 1996)*, pages 31–56, Dordrecht, The Netherlands, 1998. Kluwer Academic Publishers.
- [11] E. M. Gertz and P. E. Gill. A primal-dual trust region algorithm for nonlinear programming. Numerical Analysis Report NA 02-1, University of California, San Diego, 2002.
- [12] D. Goldfarb and L. Chen. Interior-point ℓ_2 penalty methods for nonlinear programming with strong global convergence properties. Technical report, IEOR Dept, Columbia University, New York, NY 10027, 2004.
- [13] N. I. M. Gould, D. Orban, and Ph. Toint. An interior-point ℓ_1 -penalty method for nonlinear optimization. Technical Report RAL-TR-2003-022, Rutherford Appleton Laboratory Chilton, Oxfordshire, UK, 2003.
- [14] N. I. M. Gould, D. Orban, and Ph. L. Toint. CUTEr and sifdec: A Constrained and Unconstrained Testing Environment, revisited. *ACM Trans. Math. Softw.*, 29(4):373–394, 2003.

- [15] I. Griva, D.F. Shanno, and R.J. Vanderbei. Convergence Analysis of a Primal-Dual Method for Nonlinear Programming. Report, Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ, USA, 2004.
- [16] D. G. Luenberger. *Linear and Nonlinear Programming*. Addison-Wesley Publishing Company, Reading, Massachusetts, USA, second edition, 1984.
- [17] S. Mehrotra. On the implementation of a primal-dual interior point method. *SIAM Journal on Optimization*, 2(4):575–601, 1992.
- [18] C. Mészáros. Steplengths in infeasible primal-dual interior point algorithms of convex quadratic programming. Technical Report DOC 97/7, Imperial College, London, UK., 1997.
- [19] A. L. Tits, A. Wächter, S. Bakhtiari, T. J. Urban, and C. T. Lawrence. A primal-dual interior-point method for nonlinear programming with strong global and local convergence properties. *SIAM Journal on Optimization*, 14(1):173–199, 2003.
- [20] M. Ulbrich, S. Ulbrich, and L. Vicente. A globally convergent primal-dual interior point filter method for nonconvex nonlinear programming. Technical Report TR00-11, Department of Mathematics, University of Coimbra, Coimbra, Portugal, 2000.
- [21] R. J. Vanderbei and D. F. Shanno. An interior point algorithm for nonconvex nonlinear programming. *Computational Optimization and Applications*, 13:231–252, 1999.
- [22] A. Wächter and L. T. Biegler. Line search filter methods for nonlinear programming: Motivation and global convergence. Technical Report RC 23036, IBM T. J. Watson Research Center, Yorktown Heights, NY, USA, 2003.
- [23] A. Wächter and L. T. Biegler. On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. Technical Report RC 23149, IBM T. J. Watson Research Center, Yorktown Heights, NY, USA, March 2004.
- [24] R. A. Waltz. KNITRO 4.0 User’s Manual. Technical report, Ziena Optimization, Inc., Evanston, IL, USA, October 2004.
- [25] R. A. Waltz, J. L. Morales, J. Nocedal, and D. Orban. An interior algorithm for nonlinear optimization that combines line search and trust region steps. Technical Report 2003-6, Optimization Technology Center, Northwestern University, Evanston, IL, USA, June 2003. To appear in *Mathematical Programming A*.
- [26] S. Wright. *Primal-dual interior-point methods*. Society for Industrial and Applied Mathematics, Philadelphia, 1997.
- [27] H. Yamashita. A globally convergent primal-dual interior-point method for constrained optimization. Technical report, Mathematical System Institute, Inc., Tokyo, Japan, May 1992. Revised March 1994.

- [28] H. Yamashita, H. Yabe, and T. Tanabe. A globally and superlinearly convergent primal-dual interior point trust region method for large scale constrained optimization. Technical report, Mathematical System Institute, Inc., 2004. To appear in *Mathematical Programming A*.