THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

A New Unblocking Technique to Warmstart Interior Point Methods Based on Sensitivity Analysis

Citation for published version:

Gondzio, J & Grothey, A 2008, 'A New Unblocking Technique to Warmstart Interior Point Methods Based on Sensitivity Analysis', *Siam journal on optimization*, vol. 19, no. 3, pp. 1184-1210.
<https://doi.org/10.1137/060678129>

Digital Object Identifier (DOI):

[10.1137/060678129](https://doi.org/10.1137/060678129)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Siam journal on optimization

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

A NEW UNBLOCKING TECHNIQUE TO WARMSTART INTERIOR POINT METHODS BASED ON SENSITIVITY ANALYSIS*

JACEK GONDZIO[†] AND ANDREAS GROTHEY[†]

Abstract. One of the main drawbacks associated with Interior Point Methods (IPMs) is the perceived lack of an efficient warmstarting scheme which would enable the use of information from a previous solution of a similar problem. Recently there has been renewed interest in the subject. A common problem with warmstarting for IPM is that an advanced starting point which is close to the boundary of the feasible region, as is typical, might lead to blocking of the search direction. Several techniques have been proposed to address this issue. Most of these aim to lead the iterate back into the interior of the feasible region—we classify them as either “modification steps” or “unblocking steps” depending on whether the modification is taking place before solving the modified problem to prevent future problems, or during the solution if and when problems become apparent. A new “unblocking” strategy is suggested which attempts to directly address the issue of blocking by performing sensitivity analysis on the Newton step with the aim of increasing the size of the step that can be taken. This analysis is used in a new technique to warmstart interior point methods: we identify components of the starting point that are responsible for blocking and aim to improve these by using our sensitivity analysis. The relative performance of a selection of different warmstarting techniques suggested in the literature and the new proposed unblocking by sensitivity analysis is evaluated on the warmstarting test set based on a selection of NETLIB problems proposed by [Benson and Shanno, *Comput. Optim. Appl.*, 38 (2007), pp. 371–399]. Warmstarting techniques are also applied in the context of solving nonlinear programming problems as a sequence of quadratic programs solved by interior point methods. We also apply the warmstarting technique to the problem of finding the complete efficient frontier in portfolio management problems (a problem with 192 million variables—to our knowledge the largest problem to date solved by a warmstarted IPM). We find that the resulting best combined warmstarting strategy manages to save between 50 and 60% of interior point iterations, consistently outperforming similar approaches reported in current optimization literature.

Key words. interior-point methods, warm-start, quadratic programming

AMS subject classifications. 90C51, 90C20, 65K05

DOI. 10.1137/060678129

1. Introduction. Since their introduction, Interior Point Methods (IPMs) have been recognized as an invaluable tool to solve linear, quadratic, and nonlinear programming problems, in many cases outperforming traditional simplex and active set-based approaches. This is especially the case for large scale problems. One of the weaknesses of IPMs is, however, that unlike their active set-based competitors, they cannot easily exploit an advanced starting point obtained from the preceding solution process of a similar problem. Many optimization problems require the solution of a sequence of closely related problems, either as part of an algorithm (e.g., SQP, Branch & Bound) or as a direct application to a problem (e.g., finding the efficient frontier in portfolio optimization). Because of their weakness in warmstarting, IPMs have not made as big an impact in these areas.

Over the years there have been several attempts to improve the warmstarting capabilities of IPMs [5, 8, 15, 6, 1, 2, 10]. All of these, apart from [1, 2], involve

*Received by the editors December 19, 2006; accepted for publication (in revised form) April 23, 2008; published electronically November 19, 2008.

<http://www.siam.org/journals/siopt/19-3/67812.html>

[†]School of Mathematics, University of Edinburgh, Mayfield Road, Edinburgh EH9 3JZ United Kingdom (J.Gondzio@ed.ac.uk, <http://www.maths.ed.ac.uk/~gondzio/>, A.Grothey@ed.ac.uk, <http://www.maths.ed.ac.uk/~agr/>).

remembering a primal/dual iterate encountered during the solution of the original problem and using this (or some modification of it) as a starting point for the modified problem. All of these papers (apart from [2]) deal with the linear programming (LP) case, whereas we are equally interested in the quadratic programming (QP) case.

A typical way in which a ‘bad’ starting point manifests itself is *blocking*: The Newton direction from this point leads far outside the positive orthant, resulting in only a very small fraction of it to be taken. Consequently, the next iterate will be close to the previous one, and the search direction will likely block again. In our observation this blocking is usually due only to a small number of components of the Newton direction. We therefore suggest an unblocking strategy which attempts to modify these blocking components without disturbing the primal-dual direction too much. The unblocking strategy is based on performing sensitivity analysis of the primal-dual direction with respect to the components of the current primal/dual iterate.

As a separate thread to the paper, it is our feeling that a wealth of warmstarting heuristics have been proposed by various authors, each demonstrating improvements over a coldstarted IPM. However, there has been no attempt at comparing these in a unified environment, or indeed investigating how these might be combined. This paper will give an overview of some of the warmstarting techniques that have been suggested and explore what benefit can be obtained from combining them.

This will also set the scene for evaluating the new unblocking strategy derived in this paper, within a variety of different warmstarting settings.

We continue by stating the notation used in this paper. In section 3, we review traditionally used warmstart strategies. In section 4 we present the new unblocking techniques based on sensitivity analysis. Numerical comparisons as to the efficiency of the suggested techniques are reported in section 5. In section 6, we draw our conclusions.

2. Notation and background. The infeasible primal dual interior point methods applied to solve the quadratic programming problem

$$(1) \quad \begin{array}{ll} \min & c^T x + \frac{1}{2} x^T Q x \\ \text{s.t.} & Ax = b \\ & x \geq 0 \end{array}$$

can be motivated from the KKT conditions for (1)

$$\begin{array}{ll} (2a) & c + Qx - A^T y - z = 0 \\ (2b) & Ax = b \\ (2c) & XZe = \mu e \\ (2d) & x, z \geq 0, \end{array}$$

where the zero right-hand side of the complementary products has been replaced by the centrality parameter $\mu > 0$. The set of solutions to (2) for different values of μ is known as the *central path*. It is beneficial in this context to consider two neighborhoods of the central path, the N_2 neighborhood

$$N_2(\theta) := \{(x, y, z) : Ax = b, A^T y - Qx + z = c, \|XZe - \mu e\|_2 \leq \theta\}$$

and the wider N_∞ neighborhood

$$N_\infty(\gamma) := \{(x, y, z) : Ax = b, A^T y - Qx + z = c, x_i z_i \geq \gamma \mu\}.$$

Assume that at some stage during the algorithm the current iterate is (x, y, z) . Our variant of the predictor-corrector algorithm [4, 7] will calculate a predictor direction $(\Delta x_p, \Delta y_p, \Delta z_p)$ as the Newton direction for system (2) and a small μ -target ($\mu^0 \approx 0.001 \frac{x^T z}{n}$):

$$(3) \quad \begin{aligned} -Q\Delta x_p + A^T \Delta y_p + \Delta z_p &= c + Qx - A^T y - z = \xi_c \\ A\Delta x_p &= b - Ax = \xi_b \\ X\Delta z_p + Z\Delta x_p &= \mu^0 e - XZe = r_{xz}, \end{aligned}$$

which can be further condensed by using the third equation to eliminate Δz_p

$$(4a) \quad \begin{bmatrix} -Q - X^{-1}Z & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \Delta x_p \\ \Delta y_p \end{bmatrix} = \begin{bmatrix} r_x \\ r_y \end{bmatrix} := \begin{bmatrix} \xi_c - X^{-1}r_{xz} \\ \xi_b \end{bmatrix}$$

$$(4b) \quad \Delta z_p = X^{-1}r_{xz} - X^{-1}Z\Delta x_p.$$

As in Mehrotra's predictor-corrector algorithm [13], we calculate maximal primal and dual stepsizes for the predictor direction

$$\bar{\alpha}_p = \max\{\alpha > 0 : x + \alpha \Delta x_p \geq 0\}, \quad \bar{\alpha}_d = \max\{\alpha > 0 : z + \alpha \Delta z_p \geq 0\}$$

and determine a target μ -value by

$$\mu = \frac{[(x + \bar{\alpha}_p \Delta x_p)^T (z + \bar{\alpha}_d \Delta z_p)]^3}{n(x^T z)^2}.$$

With these we compute the corrector direction $(\Delta x_c, \Delta y_c, \Delta z_c)$ by

$$(5) \quad \begin{aligned} A^T \Delta y_c + \Delta z_c &= 0 \\ A\Delta x_c &= 0 \\ X\Delta z_c + Z\Delta x_c &= (\mu - \mu^0)e - \Delta X_p \Delta Z_p e, \end{aligned}$$

and finally the new primal and dual stepsizes and the new iterate (x^+, z^+) as

$$\begin{aligned} \alpha_p &= 0.995 \max\{\alpha > 0 : x + \alpha(\Delta x_p + \Delta x_c) \geq 0\} \\ \alpha_d &= 0.995 \max\{\alpha > 0 : z + \alpha(\Delta z_p + \Delta z_c) \geq 0\} \\ x^+ &= x + \alpha_p(\Delta x_p + \Delta x_c), \quad z^+ = z + \alpha_d(\Delta z_p + \Delta z_c). \end{aligned}$$

Our main interest is generating a good starting point for the QP problem (1) — the *modified problem* — from the solution of a previously solved similar QP problem

$$(6) \quad \begin{aligned} \min \quad & \tilde{c}^T x + \frac{1}{2} x^T \tilde{Q} x \\ \text{s.t.} \quad & \tilde{A}x = \tilde{b} \\ & x \geq 0, \end{aligned}$$

the *original problem*. The difference between the two problems, i.e., the change from the original problem to the second problem, is denoted by

$$(\Delta A, \Delta Q, \Delta c, \Delta b) = (A - \tilde{A}, Q - \tilde{Q}, c - \tilde{c}, b - \tilde{b}).$$

3. Warmstart heuristics. Unlike the situation in the Simplex Method, for IPMs it is not a good strategy to use the optimal solution of a previously solved problem as the new starting point for a similar problem. This is because problems are often ill-conditioned; hence the final solution of the original problem might be far away from the central path of the modified problem. Furthermore, [9] demonstrates that the predictor direction tends to be parallel to nearby constraints, resulting in difficulties to drop misidentified nonbasic variables.

Over the years numerous contributions [11, 5, 8, 15, 6] have addressed this problem, with renewed interest in the subject from [1, 2, 10] over the last year. With the exception of [1, 2] which use an L_1 -penalty reformulation of the problem that has better warmstarting capabilities, all remedies follow a common theme: They identify an advanced center [5], a point close to the central path of the original problem (usually a nonconverged iterate), and **modify** it in such a manner that the modified point is close to the central path of the new problem. Further, in the first few iterations of the reoptimization, additional techniques which address the issue of getting stuck at nearby constraints may be employed. In this paper these will be called **unblocking heuristics**. The generic IPM warmstarting algorithm is as follows:

Algorithm: Generic Interior Point Warmstart

1. Solve the original problem (6) by an Interior Point Algorithm. From it choose one of (or a selection of) the iterates $(\tilde{x}, \tilde{y}, \tilde{z}, \tilde{\mu})$ encountered during the solution process. We will assume that this iterate (or any one of these iterates) satisfies

$$\begin{aligned}\tilde{c} + \tilde{Q}\tilde{x} - \tilde{A}^T\tilde{y} - \tilde{z} &= 0 \\ \tilde{b} - \tilde{A}\tilde{x} &= 0 \\ \tilde{x}_i\tilde{z}_i &\approx \tilde{\mu} \quad \forall i = 1, \dots, n.\end{aligned}$$

2. **Modify** the chosen iterate to obtain a starting point (x, y, z, μ) for the modified problem.

3. Solve the modified problem by an Interior Point Algorithm using (x, y, z, μ) as the starting point. During the first few iterations of the IPM a special **unblocking step** might be taken.

The question arises as to what should guide the construction of modification and unblocking steps. It is well known that for a feasible method (i.e., $\xi_b = \xi_c = 0$), a well-centered point (i.e., in $N_2(\theta)$ or $N_{-\infty}(\gamma)$) and a small target decrease ($\mu \lesssim \mu^0$), and the Newton step is feasible. Analysis by [15] and [6] identifies two factors that lead to the ability of IPMs to absorb infeasibilities ξ_b, ξ_c present at the starting point. Firstly, the larger the value of μ the more infeasibility can be absorbed in one step. Secondly, the centrality of the iterate: from a well-centered point the IPM can again absorb more infeasibilities. Using these general guidelines, a number of different warmstarting techniques have been suggested. We review some of them here:

Modification Steps:

(i) *Shift small components:* [11] shift $\tilde{x}, \tilde{z}$ by $h_x = \epsilon D^{-1}e, h_z = \epsilon De$, where $D = \text{diag}\{\|a_j\|_1\}$ and a_j is the j th column of A to ensure $x_i z_i \geq \gamma \mu$ for some small $\gamma > 0$, i.e., improve centrality by aiming for a point in $N_{-\infty}(\gamma)$.

(ii) References [15, 10] suggest a *Weighted Least Squares Step* (WLS) that finds the minimum step (with respect to a weighted 2-norm) from the starting point, to

a point that is both primal and dual feasible. The WLS step does not necessarily preserve positiveness of the iterate. To overcome this, [15] suggests keeping a selection of potential warmstart iterates and retracing to one corresponding to a large μ , which will guarantee that the WLS step is feasible. Since we do not want to remember several different points from the solution of the original problem, we will take a fraction of the WLS step should the full step be infeasible. Mehrotra's starting point [13] can be seen as a (damped) WLS step from the origin.

(iii) References [15, 10] further suggest a *Newton Modification Step*, i.e., an interior point step (3) correcting only for the primal and dual infeasibilities introduced by the change of problem, with no attempt to improve centrality: (3) is solved with $r_{xz} = 0$. Again only a fraction of this step might be taken.

Unblocking Heuristics

(i) *Splitting Directions*: Reference [6] advocates computing separate search directions aimed at achieving primal feasibility, dual feasibility, and centrality separately. These are combined into the complete step by taking the maximum of each step that can be taken without violating the positivity of the iterates. A possible interpretation of this strategy is to emulate a gradual change from the original problem to the modified problem where for each change the modification step is feasible.

(ii) *Higher Order Correctors*: The $\Delta X_p \Delta Z_p$ component in (5) is a correction for the linearization error in $XZe - \mu e = 0$. A corrector of this type can be repeated several times. Reference [5] employs this idea by additionally correcting only for small complementary products to avoid introducing additional blocking. This is used in [6] as an unblocking technique with the interpretation of choosing a target complementary vector $\bar{t} \approx \mu e$ in such a way that a large step in the resulting Newton direction is feasible, aiming to absorb as much of the primal/dual infeasibility as possible in the first step.

(iii) *Change Diagonal Scaling*: Reference [9] investigates changing elements in the scaling matrix $\Theta = XZ^{-1}$ to make nearby constraints repelling rather than attracting to the Newton step. However, we are not aware of any implementation of this technique in a warmstarting context.

A number of additional interesting techniques are listed here and described below:

(i) *Dual adjustment*: Adjust advanced starting point $\tilde{z}$ to compensate for changes to c , A , and Q in the dual feasibility constraint (2a).

(ii) *Additional centering* iterations before the advanced starting point is used.

(iii) Unblocking of the step direction by *sensitivity analysis*.

We will give a brief description of the first two of these strategies. The third (unblocking by sensitivity analysis) is the subject of section 4.

Dual adjustment

Using $(\tilde{x}, \tilde{y}, \tilde{z})$ as a starting point in problem (1) will result in the initial dual infeasibility

$$\xi_c = c + Q\tilde{x} - A^T\tilde{y} - \tilde{z} = \Delta c + \Delta Q\tilde{x} - \Delta A^T\tilde{y}.$$

Setting $z = \tilde{z} + \Delta z$, where $\Delta z = \Delta c + \Delta Q\tilde{x} - \Delta A^T\tilde{y}$, would result in a point satisfying the dual feasibility constraint (2a). However, the conditions $z \geq 0$ and $x_i z_i \approx \mu$ are likely violated by this, so instead we set

$$z_i = \max\{\tilde{z}_i + \Delta z_i, \min\{\sqrt{\mu}, \tilde{z}_i/2\}\};$$

i.e., we try to absorb as much of the dual infeasibility into z as possible without decreasing z either below $\sqrt{\mu}$ or half its value.

Adjusting the saved iterate $(\tilde{x}, \tilde{y}, \tilde{z})$ in a minimal way to absorb primal/dual infeasibilities is similar in spirit to the WLS modification step. Unlike this, however, direct adjustment of z is much cheaper to compute.

Additional centering iterations

The aim of improving the centrality of the saved iterate can also be achieved by performing an additional pure centering iteration, i.e., choose $\xi_c = \xi_b = 0, \mu^0 = x^T z/n$ in (3), in the original problem before saving the iterate as a starting point for the new problem. This pure centering iteration could be performed with respect to the original or the modified problem. In the latter case, this is similar in spirit to the Newton Modification Step of [15, 10] (whereas [15, 10] use $r_{xz} = 0$, we use $r_{xz} = \mu^0 e - \tilde{X}\tilde{Z}$ with $\mu^0 = \tilde{x}^T \tilde{z}/n$). In the case of a perfectly centered saved iterate—as we hope to achieve at least approximately by the previous centering in the original problem—these two are identical). We refer to these as centering iteration at the *beginning* of solving the modified problem or at the *end* of solving the original problem.

In the next section we will derive the unblocking strategy based on sensitivity analysis.

4. Unblocking by sensitivity analysis.

4.1. Sensitivity analysis. In this section we will lay the theoretical foundations for our proposed unblocking strategy. Much of it is based on the observation that the advanced starting information (x, y, z, μ) with which to start the solution of the modified problem is to some degree arbitrary. It is therefore possible to treat it as parameters to the solution process and to explore how certain properties of the solution process change as the starting point is changed. In particular we are interested in the primal and dual stepsizes that can be taken for the Newton direction computed from this point.

At some iterate (x, y, z) of the IPM, the primal-dual direction $(\Delta x, \Delta y, \Delta z)$ is obtained as the solution to the system (3) or (4) for some target value μ^0 . If we think of (x, y, z) as the advanced starting point, the step $(\Delta x, \Delta y, \Delta z)$ can be obtained as a function of the current point (x, y, z) . The aim of this section is to derive a procedure by which the sensitivity of $\Delta x(x, y, z)$, $\Delta y(x, y, z)$, $\Delta z(x, y, z)$, that is the first derivatives of these functions can be computed.

First note that the value of y has no influence on the new step $\Delta x, \Delta z$. This is because after substituting for ξ_b, ξ_c, r_{xz} in (4a)

$$\begin{bmatrix} -Q - X^{-1}Z & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} c + Qx - A^T y - \mu X^{-1}e \\ b - Ax \end{bmatrix}$$

we can rewrite this as

$$(7) \quad \begin{bmatrix} -Q - X^{-1}Z & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ y^{(k+1)} \end{bmatrix} = \begin{bmatrix} c + Qx - \mu X^{-1}e \\ b - Ax \end{bmatrix}$$

with $\Delta y = y^{(k+1)} - y$. In effect (7) solves for the new value of $y^{(k+1)} = y^{(k)} + \Delta y$ directly, whereas all influence of y onto $\Delta x, \Delta z$ has been removed. Notice also that only the step components in x, z variables can lead to a blocking of the step; therefore we are interested only in the functional relationship and sensitivity for the functions $\Delta x = \Delta x(x, z), \Delta z = \Delta z(x, z)$. To this end we start by differentiating with respect to

x_i in (3):

$$(8a) \quad -Q \frac{d\Delta x}{dx_i} + A^T \frac{d\Delta y}{dx_i} + \frac{d\Delta z}{dx_i} = Qe_i,$$

$$(8b) \quad A \frac{d\Delta x}{dx_i} = -Ae_i,$$

$$(8c) \quad X \frac{d\Delta z}{dx_i} + Z \frac{d\Delta x}{dx_i} + \Delta Ze_i = -Ze_i.$$

Note that this result is independent of the value of μ^0 that is used as a target. Similarly, differentiating with respect to y_i yields

$$(9a) \quad -Q \frac{d\Delta x}{dy_i} + A^T \frac{d\Delta y}{dy_i} + \frac{d\Delta z}{dy_i} = -A^T e_i$$

$$(9b) \quad A \frac{d\Delta x}{dy_i} = 0$$

$$(9c) \quad X \frac{d\Delta z}{dy_i} + Z \frac{d\Delta x}{dy_i} = 0,$$

and finally differentiating with respect to z_i yields

$$(10a) \quad -Q \frac{d\Delta x}{dz_i} + A^T \frac{d\Delta y}{dz_i} + \frac{d\Delta z}{dz_i} = -e_i$$

$$(10b) \quad A \frac{d\Delta x}{dz_i} = 0$$

$$(10c) \quad X \frac{d\Delta z}{dz_i} + Z \frac{d\Delta x}{dz_i} + \Delta X e_i = -X e_i.$$

Taking all three systems together we have

$$(11) \quad \begin{bmatrix} -Q & A^T & I \\ A & 0 & 0 \\ Z & 0 & X \end{bmatrix} \begin{bmatrix} \frac{d\Delta x}{dx} & \frac{d\Delta y}{dy} & \frac{d\Delta z}{dz} \\ \frac{d\Delta x}{dy} & \frac{d\Delta y}{dy} & \frac{d\Delta z}{dz} \\ \frac{d\Delta x}{dz} & \frac{d\Delta y}{dz} & \frac{d\Delta z}{dz} \end{bmatrix} = \begin{bmatrix} Q & -A^T & -I \\ -A & 0 & 0 \\ -Z - \Delta Z & 0 & -X - \Delta X \end{bmatrix}.$$

Under the assumption that A has full row rank, the system matrix is nonsingular, therefore

$$(12a) \quad \begin{bmatrix} \frac{d\Delta x}{dx_i} \\ \frac{d\Delta y}{dx_i} \\ \frac{d\Delta z}{dx_i} \end{bmatrix} = \begin{bmatrix} -e_i \\ 0 \\ 0 \end{bmatrix} + \Delta z_i \begin{bmatrix} -Q & A^T & I \\ A & 0 & 0 \\ Z & 0 & X \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ 0 \\ -e_i \end{bmatrix}$$

$$(12b) \quad \begin{bmatrix} \frac{d\Delta x}{dy_i} \\ \frac{d\Delta y}{dy_i} \\ \frac{d\Delta z}{dy_i} \end{bmatrix} = \begin{bmatrix} 0 \\ -e_i \\ 0 \end{bmatrix}$$

$$(12c) \quad \begin{bmatrix} \frac{d\Delta x}{dz_i} \\ \frac{d\Delta y}{dz_i} \\ \frac{d\Delta z}{dz_i} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ -e_i \end{bmatrix} + \Delta x_i \begin{bmatrix} -Q & A^T & I \\ A & 0 & 0 \\ Z & 0 & X \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ 0 \\ -e_i \end{bmatrix},$$

where the system common to (12a/12c)

$$(13) \quad \begin{bmatrix} -Q & A^T & I \\ A & 0 & 0 \\ Z & 0 & X \end{bmatrix} \begin{bmatrix} \widetilde{d\Delta x} \\ \widetilde{d\Delta y} \\ \widetilde{d\Delta z} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ -e_i \end{bmatrix}$$

can be solved by using the third line to substitute for $\widetilde{d\Delta z}$ as

$$(14a) \quad \begin{bmatrix} -Q - X^{-1}Z & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \widetilde{d\Delta x} \\ \widetilde{d\Delta y} \end{bmatrix} = \begin{bmatrix} X^{-1}e_i \\ 0 \end{bmatrix}$$

$$(14b) \quad \widetilde{d\Delta z} = -X^{-1}Z\widetilde{d\Delta x} - X^{-1}e_i.$$

There are a few insights to be gained from these formulas. First, they confirm that the step $(\Delta x, \Delta z)$ does not depend on y .

Second, the *sensitivity* of the primal-dual step with respect to the current iterate (x, y, z) —unlike the step $(\Delta x, \Delta y, \Delta z)$ itself—does not depend on the target value μ^0 either. We will exploit this property when constructing a warmstart heuristic that uses the sensitivity information.

Finally we can get the complete sensitivity information with respect to (x_i, z_i) for a given component i by solving a single system of linear equations with the same augmented system matrix that has been used to obtain the step $(\Delta x, \Delta y, \Delta z)$ (and for which a factorization is available); the solution of n such systems will likewise retrieve the complete sensitivity information.

Although this system matrix is already factorized as part of the normal interior point algorithm, and backsolves are an order of magnitude cheaper than the factorization, obtaining the complete sensitivity information is prohibitively expensive. The aim of the following section is therefore to propose a warmstarting heuristic that uses the sensitivity information derived above, but requires only a few, rather than all n backsolves.

4.2. Unblocking the primal-dual direction using sensitivity information. Occasionally, despite all our attempts, a starting point might result in a Newton direction that leads to blocking: i.e., only a very small step can be taken along it. We do not want to abandon the advanced starting information at this point, but rather try to *unblock* the search direction. To this end we will make use of the sensitivity analysis presented in section 4.1. The following Lemma 1 gives conditions under which a step (d_x, d_z) can be expected to unblock based on the sensitivity analysis.

LEMMA 1. *A necessary and sufficient condition for a step (d_x, d_z) to unblock to first order to a given level ρl , i.e.,*

$$(15a) \quad x + d_x + \Delta x + \frac{d\Delta x}{dx}d_x + \frac{d\Delta x}{dz}d_z \geq \rho l,$$

$$(15b) \quad z + d_z + \Delta z + \frac{d\Delta z}{dx}d_x + \frac{d\Delta z}{dz}d_z \geq \rho l,$$

is that there exists vectors d_x, d_z, t_x, t_y, t_z of appropriate dimensions that satisfy the system of equations

$$\begin{aligned} (16a) \quad & At_x = 0 \\ (16b) \quad & -Qt_x + A^T t_y + t_z = 0 \\ (16c) \quad & Zt_x + Xt_z = -\Delta Z d_x - \Delta X d_z \\ (16d) \quad & t_x \geq -x - \Delta x + \rho l \\ (16e) \quad & t_z \geq -z - \Delta z + \rho l \end{aligned}$$

Proof. Note that the relations of (11),(12) can be more concisely written as

$$(17) \quad \begin{bmatrix} \frac{d\Delta x}{dx} + I & \frac{d\Delta x}{dy} & \frac{d\Delta z}{dz} \\ \frac{d\Delta y}{dx} & \frac{d\Delta y}{dy} + I & \frac{d\Delta y}{dz} \\ \frac{d\Delta z}{dx} & \frac{d\Delta z}{dy} & \frac{d\Delta z}{dz} + I \end{bmatrix} = - \begin{bmatrix} -Q & A^T & I \\ A & 0 & 0 \\ Z & 0 & X \end{bmatrix}^{-1} \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ \Delta Z & 0 & \Delta X \end{bmatrix}.$$

Conditions (15) are equivalent to the existence of (d_x, d_y, d_z) such that

$$(18) \quad \begin{bmatrix} \frac{d\Delta x}{dx} + I & \frac{d\Delta x}{dy} & \frac{d\Delta z}{dz} \\ \frac{d\Delta y}{dx} & \frac{d\Delta y}{dy} + I & \frac{d\Delta y}{dz} \\ \frac{d\Delta z}{dx} & \frac{d\Delta z}{dy} & \frac{d\Delta z}{dz} + I \end{bmatrix} \begin{bmatrix} d_x \\ d_y \\ d_z \end{bmatrix} \geq \begin{bmatrix} -x - \Delta x + \rho l \\ -\infty \\ -z - \Delta z + \rho l \end{bmatrix},$$

where d_y is an arbitrary vector (note that $\frac{d\Delta x}{dy} = \frac{d\Delta z}{dy} = 0$). This, on the other hand, is satisfied, if and only if there exists (t_x, t_y, t_z) such that

$$(19) \quad \begin{bmatrix} \frac{d\Delta x}{dx} + I & \frac{d\Delta x}{dy} & \frac{d\Delta z}{dz} \\ \frac{d\Delta y}{dx} & \frac{d\Delta y}{dy} + I & \frac{d\Delta y}{dz} \\ \frac{d\Delta z}{dx} & \frac{d\Delta z}{dy} & \frac{d\Delta z}{dz} + I \end{bmatrix} \begin{bmatrix} d_x \\ d_y \\ d_z \end{bmatrix} = \begin{bmatrix} t_x \\ t_y \\ t_z \end{bmatrix} \geq \begin{bmatrix} -x - \Delta x + \rho l \\ -\infty \\ -z - \Delta z + \rho l \end{bmatrix}.$$

Now using (17) to substitute for the matrix of derivatives, multiplying both sides of the equality with the augmented system matrix and multiplying out we see that (19) is equivalent to

$$\begin{bmatrix} 0 \\ 0 \\ -\Delta X d_z - \Delta Z d_x \end{bmatrix} = \begin{bmatrix} -Q & A^T & I \\ A & 0 & 0 \\ Z & 0 & X \end{bmatrix} \begin{bmatrix} t_x \\ t_y \\ t_z \end{bmatrix}, \quad \begin{bmatrix} t_x \\ t_y \\ t_z \end{bmatrix} \geq \begin{bmatrix} -x - \Delta x + \rho l \\ -\infty \\ -z - \Delta z + \rho l \end{bmatrix}$$

that is to (16). $\square$

The sensitivity analysis thus gives us conditions that an unblocking direction needs to satisfy. However it is unclear if a direction (d_x, d_z) and the corresponding (t_x, t_y, t_z) to satisfy the conditions of Lemma 1 exist. We can however prove existence of such a direction by assuming that we know the analytic center $\hat{p} = (\hat{x}, \hat{y}, \hat{z})$ of the problem (or indeed any strictly primal-dual feasible point) and denote by $\underline{\hat{p}}, \hat{\bar{p}}$ its largest and smallest component:

$$0 < \underline{\hat{p}} \leq \hat{x}_i, \hat{z}_i \leq \hat{\bar{p}}$$

LEMMA 2. For all $l : 0 < l < \min\{\underline{\hat{p}}/4, 1\}$, $\rho < 1$ and fixed μ and γ there exists a $c = c(\gamma, \mu)$ such that for all starting points (x, y, z) and corresponding blocking step

$(\Delta x, \Delta y, \Delta z)$ obtained from (3) with $\mu^0 = \mu^+$ satisfying

$$\begin{aligned} x^T z / n = \mu, \quad x_i z_i \geq \gamma \mu, \quad x_i \leq b_u, z_i \leq b_u, \quad \mu^+ \leq \frac{1}{2} \gamma \mu, \\ z + \Delta z \geq -le, \quad x + \Delta x \geq -le, \end{aligned}$$

there exists a step $(d_x, d_z) : \|d_x, d_z\|_\infty \leq c(1 + \rho)l$ that unblocks to first order to level ρL , i.e., that satisfies conditions (15).

Proof. With $\alpha = 2(1 + \rho)l/\hat{p}$ set

$$\begin{aligned} t_x &= \alpha(\hat{x} - (x + \Delta x)), \\ t_y &= \alpha(\hat{y} - (y + \Delta y)), \\ t_z &= \alpha(\hat{z} - (z + \Delta z)). \end{aligned}$$

We will show that (t_x, t_y, t_z) satisfies (16a/b/d/e), that we can construct a corresponding (d_x, d_z) satisfying (16c), and finally that $(t_x, t_y, t_z), (d_x, d_z) = \mathcal{O}((1 + \rho)l)$.

First we notice that both $\hat{p} = (\hat{x}, \hat{y}, \hat{z})$ and $(x + \Delta x, y + \Delta y, z + \Delta z)$ are primal and dual feasible (although in the latter case, of course, not positive). Hence, their difference (and therefore (t_x, t_y, t_z)) satisfies (16a/b).

To proof (16d) we need to distinguish the two cases: $x_i + \Delta x_i < \rho l$ and $x_i + \Delta x_i \geq \rho l$. In the first case $x_i + \Delta x_i < \rho l$ we have

$$\hat{x}_i - (x_i + \Delta x_i) \geq \hat{p} - \rho l \geq \frac{1}{2} \hat{p},$$

where the last inequality is due to $\rho \leq 1$ and $l \leq \hat{p}/4$. Then

$$t_{x,i} = \frac{2(1 + \rho)l}{\hat{p}}(\hat{x}_i - (x_i + \Delta x_i)) \geq \frac{2(1 + \rho)l}{\hat{p}} \frac{1}{2} \hat{p} = (1 + \rho)l \geq -x_i - \Delta x_i + \rho l.$$

In the second case $x_i + \Delta x_i \geq \rho l$, we note that $\hat{x}_i \geq \hat{p} \geq 4l \geq \rho l$. Since $2(1 + \rho) \leq 4$ we have $0 < \alpha \leq 1$, and hence

$$x_i + \Delta x_i + \alpha(\hat{x}_i - (x_i + \Delta x_i)) = (1 - \alpha)(x_i + \Delta x_i) + \alpha \hat{x}_i \geq \rho l,$$

which proves (16d). (16e) is proven in the same manner.

Next we establish a bound for $\|t_x\|, \|t_z\|$. Since $\mu^+ < \gamma\mu/2$ we have from the last equation of (3):

$$(20) \quad x_i \Delta z_i + z_i \Delta x_i = \mu^+ - x_i z_i \leq \frac{1}{2} \gamma \mu - \gamma \mu = -\frac{1}{2} \gamma \mu < 0,$$

and hence at least one of $\Delta x_i, \Delta z_i$ must be negative. Assume w.l.o.g. that $\Delta x_i < 0$. Then $x_i + \Delta x_i \geq -l$ implies

$$|\Delta x_i| = -\Delta x_i \leq l + x \leq l + b_u.$$

We can make no further assumptions on the sign of Δz_i . If $\Delta z_i < 0$, then $|\Delta z_i|$ is bounded in the same way as Δx_i . If, on the other hand, $\Delta z_i \geq 0$, then (20) together with $x_i \geq \gamma/z_i > \gamma/b_u$ implies

$$\Delta z_i < -z_i \Delta x_i / x_i < b_u(l + b_u)b_u/\gamma = b_u^2(l + b_u)/\gamma.$$

Since we can reasonably assume that $b_u^2/\gamma > 1$, we have

$$\|\Delta x\|, \|\Delta z\| \leq b_u^2(1 + b_u)/\gamma.$$

From this we get

$$\|t_x\| = \alpha \|\hat{x} - (x + \Delta x)\| \leq \frac{2(1 + \rho)l}{\hat{p}}(\hat{p} + b_u + b_u^2(1 + b_u)/\gamma) = c_1(1 + \rho)l,$$

where $c_1 = c_1(\gamma) = 2(\hat{p} + b_u + b_u^2(1 + b_u)/\gamma)/\hat{p}$. $\|t_z\| \leq c_1(1 + \rho)l$ follows in the same manner.

Finally we know from (20) that for all i at least one of $x_i \Delta z_i, z_i \Delta x_i$ must be less than $-\gamma\mu/4$. Assume w.l.o.g. $z_i \Delta x_i < -\gamma\mu/4$, and then we get

$$(21) \quad \Delta x_i < -\frac{\gamma\mu}{4z_i} < 0.$$

Therefore we can set

$$(22) \quad d_{z,i} = -\frac{z_i t_{x,i} + x_i t_{z,i}}{\Delta x_i}, \quad d_{x,i} = 0$$

(and vice versa if $x_i \Delta z_i < -\gamma\mu/4$) to construct a direction (d_x, d_z) that satisfies (16c). It remains to be shown that $(d_x, d_z) = \mathcal{O}((1 + \rho)l)$:

From (21) we know

$$|\Delta x_i| = -\Delta x_i > \frac{\gamma\mu}{4z_i} > \frac{\gamma\mu}{4b_u};$$

hence (22) gives

$$|d_{z,i}| \leq (b_u c_1(1 + \rho)l + b_u c_1(1 + \rho)l) / \frac{\gamma\mu}{4b_u} = \frac{8b_u^2 c_1}{\gamma\mu}(1 + \rho)l := c(1 + \rho)l$$

with $c = c(\gamma, \mu) = (8b_u^2 c_1)/(\gamma\mu)$. $\square$

We can now proof the main result of this section.

THEOREM 1. *There exists $L > 0$ such that for all $l : 0 < l < L$ and all starting points (x, y, z) and their corresponding blocking step $(\Delta x, \Delta y, \Delta z)$ obtained from (3) with $\mu^0 = \mu^+$ that satisfy*

$$x^T z/n = \mu, \quad x_i z_i \geq \gamma\mu, \quad x_i, z_i \leq b_u, \quad \mu^+ < \frac{1}{2}\gamma\mu, \quad x + \Delta x \geq -le, \quad z + \Delta z \geq -le,$$

there is a step (d_x, d_z) that unblocks, i.e.,

$$\begin{aligned} x + d_x + \Delta x(x + d_x, z + d_z) &\geq 0, \\ z + d_z + \Delta z(x + d_x, z + d_z) &\geq 0. \end{aligned}$$

Proof. Set $\epsilon = \frac{1}{10c}$. From the differentiability of $\Delta x(x, z), \Delta z(x, z)$ there exists a δ such that for all $(d_x, d_z) : \|(d_x, d_z)\|_\infty \leq \delta$:

$$(23) \quad \begin{aligned} \left\| \Delta x(x + d_x, z + d_z) - \Delta x - \frac{d\Delta x}{dx}d_x - \frac{d\Delta x}{dz}d_z \right\| &\leq \epsilon \|(d_x, d_z)\|, \\ \left\| \Delta z(x + d_x, z + d_z) - \Delta z - \frac{d\Delta z}{dx}d_x - \frac{d\Delta z}{dz}d_z \right\| &\leq \epsilon \|(d_x, d_z)\|. \end{aligned}$$

Now set $\rho = \frac{1}{4}$ and $L = \min\{\frac{\hat{p}}{4}, \frac{4}{5}\frac{\delta}{c}\}$. Then from Lemma 2 there exists $(d_x, d_z) : \|(d_x, d_z)\|_\infty \leq c\frac{5}{4}l \leq \delta$ such that

$$\begin{aligned} x + d_x + \Delta x + \frac{d\Delta x}{dx}d_x + \frac{d\Delta x}{dz}d_z &\geq \rho l e = \frac{1}{4}le, \\ z + d_z + \Delta z + \frac{d\Delta z}{dx}d_x + \frac{d\Delta z}{dz}d_z &\geq \rho l e = \frac{1}{4}le, \end{aligned}$$

and therefore

$$\begin{aligned} &x_i + d_{x,i} + \Delta x_i(x + d_x, z + d_z) \\ &= x_i + d_{x,i} + \Delta x_i + \frac{d\Delta x_i}{dx}d_x + \frac{d\Delta x_i}{dz}d_z \\ &\quad - \left(\Delta x_i + \frac{d\Delta x_i}{dx}d_x + \frac{d\Delta x_i}{dz}d_z - \Delta x_i(x + d_x, z + d_z) \right) \\ &\geq \frac{1}{4}l - \epsilon \|(d_x, d_z)\| = \frac{1}{4}l - \frac{1}{10c} \|(d_x, d_z)\| \\ &\geq \frac{1}{4}l - \frac{1}{10c}c\frac{5}{4}l \geq \frac{1}{8}l > 0, \end{aligned}$$

and the same for the z components. $\square$

The insight gained from this theorem is that our proposed unblocking strategy is sound in principle: If the negative components of the prospective next iterate $(x + \Delta x, z + \Delta z)$ are bounded in size by L , then there exists an unblocking perturbation (d_x, d_z) of the current iterate. The size of this perturbation is $\mathcal{O}(L)$. Unfortunately the construction of (d_x, d_z) relies on the knowledge of the analytic center $\hat{p}$ of the problem (or at least any other strictly primal/dual feasible point). Therefore the construction used in the proof cannot be implemented in practice. In the following section we will derive an implementable heuristic.

4.3. Implementation. There is a principle difficulty with finding a solution to the *unblocking equations* (16). Theorem 1 guarantees that a solution (of bounded size) exists. The system (15):

$$\begin{aligned} x + d_x + \Delta x + \frac{d\Delta x}{dx}d_x + \frac{d\Delta x}{dz}d_z &\geq \rho L, \\ z + d_z + \Delta z + \frac{d\Delta z}{dx}d_x + \frac{d\Delta z}{dz}d_z &\geq \rho L \end{aligned}$$

seems to imply that we could gather the complete sensitivity information $(\frac{d\Delta x}{dx}, \frac{d\Delta x}{dz}, \frac{d\Delta z}{dx}, \frac{d\Delta z}{dz})$, requiring n backsolves to do so, and find d_x, d_z to satisfy

$$(24) \quad \begin{bmatrix} \frac{d\Delta x}{dx} + I & \frac{d\Delta x}{dz} \\ \frac{d\Delta z}{dx} & \frac{d\Delta z}{dz} + I \end{bmatrix} \begin{bmatrix} d_x \\ d_z \end{bmatrix} \geq \begin{bmatrix} -x - \Delta x + \rho L \\ -z - \Delta z + \rho L \end{bmatrix}.$$

However, the system matrix in (24) is singular (actually of rank n) as can be seen from (17); hence it is unclear if a solution (d_x, d_z) exists at all.

In the results of Theorem 1 we get around this difficulty by assuming the knowledge of the analytic center, something that does not hold in practice. The only

solution we can suggest is to use the sensitivity information in a heuristic targeted at unblocking the search direction.

The idea is based on the observation that typically only a few components of the Newton step $(\Delta x, \Delta z)$ are blocking seriously and that these can be effectively influenced by changing the corresponding components of (x, z) only. One potential danger of aiming solely at unblocking the step direction is that we might have to accept a significant worsening of centrality or feasibility of the new iterate, which is clearly not in our interest. The proposed strategy attempts to avoid this as well by minimizing the perturbation (d_x, d_z) to the current point.

The heuristic that we are proposing is based on the assumption that a change in the i th component x_i, z_i will have a strong influence on the i th component of the step $\Delta x_i, \Delta z_i$, so changing only x_i, z_i components corresponding to blocking components of the step might be sufficient. Indeed our strategy will identify a (small) index set $\mathcal{I}$ of most blocking components, obtain the sensitivity information with respect to these components, and attempt to unblock each $(\Delta x_i, \Delta z_i)$ by changes to component i of (x, z) *only*. Since usually only Δx_i or Δz_i but not both are blocking, allowing perturbations in both x_i or z_i leaves one degree of freedom, which will be used to minimize the size of the required unblocking step.

The assumption made above can be justified as follows: according to (12), the sensitivity $d(\Delta x, \Delta z)/dx_i$ (and similarly d/dz_i) is made up of two components: the i th unit vector e_i and the solution to (13), which according to (14) is the weighted projection of the i th unit vector onto the null space of A .

Our implemented unblocking strategy is thus as follows:

Algorithm: Unblocking Strategy

- 1) Choose the size of the unblocking set $|\mathcal{I}|$, a target unblocking level $t > 1$, and bounds $0 < \underline{\gamma} < 1 < \bar{\gamma}$ on the acceptable change to a component.
 - 2) find the set $\mathcal{I}$ of most blocking components (in x or z)

for all i in 10% most blocking components do
 - 3) find sensitivity of $(\Delta x, \Delta z)$ with respect to (x_i, z_i)
 - 4) find the change $(d_{x,i}, d_{z,i})$ needed in x_i or z_i to unblock component i
 - 5) change either x_i or z_i depending on where the change would be *more effective*.

next i
 - 6) update $x = x + d_x$ and $z = z + d_z$ and recompute the affine scaling direction
-

Steps 4) and 5) of the above algorithm need further clarification: For each blocking component x_i (or z_i) we have $x_i + \alpha_x \Delta x_i < 0$ for small positive values of α_x , or $\Delta x_i/x_i \ll -1$. From the sensitivity analysis we know $\frac{d\Delta x_i}{dx_i}$, the rate of change of Δx_i when x_i changes. We are interested in the necessary change $d_{x,i}$ to x_i such that the search direction is unblocked, that is to say

$$\frac{\Delta x_i + \frac{d\Delta x_i}{dx_i} d_{x,i}}{x_i + d_{x,i}} \geq -t, \quad (t \approx 5);$$

in other words a step of $\alpha_p \geq 1/t$ ($1/t \approx 0.2$) will be possible. From this requirement

we get the provisional change

$$\widetilde{d_{x,i}} = -\frac{tx_i + \Delta x_i}{t + \frac{d\Delta x_i}{dx_i}}.$$

We need to distinguish several cases:

$$(i) \quad \frac{d\Delta x_i}{dx_i} \leq \frac{\Delta x_i}{x_i}:$$

A step in positive direction would lead to even more blocking. A negative step will unblock. However, we are not prepared to let $x_i + d_{x,i}$ approach zero; hence we choose

$$\overline{d_{x,i}} = \max\{\widetilde{d_{x,i}}, (\underline{\gamma} - 1)x_i\}.$$

$$(ii) \quad \frac{d\Delta x_i}{dx_i} > \frac{\Delta x_i}{x_i}:$$

A positive step would weaken the blocking. However, if $\frac{d\Delta x_i}{dx_i} < -t$, the target unblocking level $-t$ can never be reached (and the provisional $\widetilde{d_{x,i}}$ is negative). In this case (and also if the provisional $\widetilde{d_{x,i}}$ is very large) we choose the maximal step that we are prepared to take:

$$\overline{d_{x,i}} = \begin{cases} d_{max} & \text{if } \widetilde{d_{x,i}} < 0, \\ \min\{\widetilde{d_{x,i}}, d_{max}\} & \text{otherwise} \end{cases}$$

with $d_{max} = (\bar{\gamma} - 1)x_i$.

Alternatively we can unblock a blocking Δx_i by changing z_i . The required provisional change $\widetilde{d_{z,i}}$ can be obtained from

$$\frac{\Delta x_i + \frac{d\Delta x_i}{dz_i} d_{z,i}}{x_i} \geq -t$$

as

$$\widetilde{d_{z,i}} = -\frac{tx_i + \Delta x_i}{\frac{d\Delta x_i}{dz_i}}.$$

In this case $\widetilde{d_{z,i}}$ indicates the correct sign of the change, but for $\frac{d\Delta x_i}{dz_i}$ close to zero the provisional step might be very large. We apply the same safeguards as for the step in x to obtain

$$\overline{d_{z,i}} = \begin{cases} \max\{\widetilde{d_{z,i}}, (\underline{\gamma} - 1)z_i\} & \widetilde{d_{z,i}} < 0, \\ \min\{\widetilde{d_{z,i}}, d_{max}\} & \widetilde{d_{z,i}} \geq 0, \end{cases}$$

where $d_{max} = (\bar{\gamma} - 1)z_i$. Since our aim was to reduce the blocking level from $-\Delta x_i/x_i$ to t , we can evaluate the effectiveness of the suggested changes $\overline{d_{x,i}}, \overline{d_{z,i}}$ by

$$p_x = \frac{(\text{old blocking level}) - (\text{new blocking level})}{(\text{old blocking level}) - (\text{target blocking level})} = \frac{-\frac{\Delta x_i}{x_i} + \frac{\Delta x_i + \frac{d\Delta x_i}{dx_i} \overline{d_{x,i}}}{x_i + \overline{d_{x,i}}}}{-\frac{\Delta x_i}{x_i} + t}$$

and

$$p_z = \frac{-\frac{\Delta x_i}{x_i} + \frac{\Delta x_i + \frac{d\Delta x_i}{dz_i} \overline{d_{z,i}}}{x_i}}{-\frac{\Delta x_i}{x_i} + t}.$$

Given these quantities we use $p_x/|\overline{d_{x,i}}|, p_z/|\overline{d_{z,i}}|$ as measures of the *relative effectiveness* of changing the x_i, z_i component. Our strategy is to first change the component for which this ratio is larger, and, should the corresponding p_x, p_z be less than 1, add a proportional change in the other component, i.e., if $p_x/|\overline{d_{x,i}}| > p_z/|\overline{d_{z,i}}|$:

$$\begin{aligned} d_{x,i} &= \overline{d_{x,i}}, \\ d_{z,i} &= \min\{(1 - p_x)/p_z, 1\} \overline{d_{z,i}}. \end{aligned}$$

An analogous derivation can be performed to unblock the z -component Δz_i of the search direction.

The analysis in the previous section was aimed at unblocking the primal-dual direction corresponding to a fixed target value μ^0 . We are, however, interested in using this analysis in the context of a predictor-corrector method. This seems to complicate the situation, since the predictor-corrector direction is now the result of a two-step procedure. As pointed out earlier, however, while the primal-dual direction and subsequently the length of the step that can be taken along it does depend on the target μ^0 value, the sensitivity of this step does not depend on μ^0 . This leads us to the following strategy: We obtain the sensitivity with respect to the most blocking components after the predictor step and use these to unblock the combined predictor-corrector (and higher order corrector steps) separately following the above heuristic.

5. Numerical results. In order to evaluate the relative merit of the suggested warmstarting schemes, we have run a selection of numerical tests. In the first instance we have used a warmstarting setup based on the NETLIB LP test set as described in [1, 10] to evaluate a selection of the described heuristics.

In a second set of tests we have used the best warmstart settings from the first set and used these to warmstart the NETLIB LP test set, a selection of QP problems from [12] as well as some large scale QP problems arising from the problem of finding the efficient frontier in portfolio optimization and solving a nonlinear capacitated Multi-Commodity Network Flow problem (MCNF).

All warmstarting strategies have been implemented in our interior point solver OOPS [7]. For all tests we save the first iterate in the original problem solution process for which the relative duality gap satisfies

$$\frac{(c^T x + 0.5x^T Qx) - (b^T y - 0.5x^T Qx)}{(c^T x + 0.5x^T Qx) + 1} = \frac{x^T z}{(c^T x + 0.5x^T Qx) + 1} \leq 0.01$$

for use as a warmstarting point. We do not attempt to find an “optimal” value for $\bar{\mu}$: our motivation is primarily to evaluate unblocking techniques in order to recover from “bad” warmstarting situations; furthermore it is likely that the optimal $\bar{\mu}$ is highly problem (and perturbation) dependent. On the contrary, we assume that a 2-digit approximate optimal solution of the original problem should be a good starting point for the perturbed problem.

5.1. The NETLIB warmstarting test set. In order to compare our results more easily to other contributions, we use the NETLIB warmstarting testbed suggested by [1]. This uses the smaller problems from the NETLIB LP test set as the original problems and considers changes to the right-hand side b , the objective vector c , the system matrix A , and different perturbation sizes δ . The perturbed problem instances are randomly generated as follows:

For perturbations to b and c we first generate a uniform-[0,1] distributed random number for every vector component. Should this number be less than $\min\{0.1, 20/n\}$

TABLE 1
Higher order correctors as unblocking device.

	<i>b</i>			<i>c</i>			<i>A</i>			
	0.1	0.01	0.001	0.1	0.01	0.001	0.1	0.01	0.001	total
base	6.4	5.6	6.1	14.5	8.5	6.4	10.2	7.0	7.2	8.1
hoc	6.0	5.4	5.6	11.3	7.6	6.3	8.6	6.5	6.8	7.2

(n being the dimension of the vector), this component is marked for modification. That is, we modify on average 10% (but at most 20) of the components. For all marked components we will generate a second uniform- $[-1, 1]$ distributed random number r . The new component $\tilde{b}_i$ is generated from the old one b_i as

$$\tilde{b}_i = \begin{cases} \delta r & |b_i| \leq 10^{-6}, \\ (1 + \delta r)b_i & \text{otherwise.} \end{cases}$$

For perturbations to A we proceed in the same manner, perturbing the vector of nonzero elements in A as before. For the results presented in this paper we have solved each problem for each warmstart strategy for 10 random perturbations of each type (b , c , and A). We will use these to evaluate the merit of each of the considered modifications and unblocking heuristics. A list of the considered NETLIB problems can be obtained from Tables 6–9.

In the numerical test performed we were guided by two objectives: first to evaluate if and how the various warmstarting strategies presented in section 3 can be combined, and second to evaluate the merit of the proposed unblocking strategy. In order to save on the total amount of computation, we will use the following strategy: Every warmstarting heuristic is tested against a *base* warmstarting code and against the best combination found so far. If a heuristic is found to be advantageous, it will be added to the *best* benchmark strategy for the future tests.

5.1.1. Higher order correctors. We investigate the use of higher-order correctors as an unblocking device. The interior point code OOPS applied for these calculations uses higher-order correctors by default if the Mehrotra corrector step (5) has been successful (i.e., it leads to larger stepsizes α_P, α_D than the predictor step). When using higher order correctors as an unblocking device, we will attempt them even if the Mehrotra corrector has been rejected. Table 1 gives results with and without forcing higher order correctors (*hoc* and *base*, respectively). The numbers reported are the average number of iterations of the warmstarted problem over all problems in the test set and all 10 random perturbations. Problem instances which are infeasible or unbounded after the perturbation have been discarded. Clearly the use of higher order correctors is advantageous. We therefore recommend the use of higher order correctors in all circumstances in the context of warmstarting. All following tests are performed with the use of higher order correctors.

5.1.2. Centering steps. We explore the benefit of using centering steps as a technique to facilitate warmstarting. These are performed either at the end of the solution process for the original problem before the advanced center is returned (*end*) or at the beginning of the modified problem solution, before any reduction of the barrier μ is applied (*beg*). As pointed out earlier the latter corresponds to the Newton corrector step of [15]. We have tested several settings of *end* and *beg* corresponding to the number of steps of this type being taken. The additional centering iterations are included in the numbers reported. Results are summarized in Table 2.

TABLE 2
Additional centering iterations.

	b			c			A			total
	0.1	0.01	0.001	0.1	0.01	0.001	0.1	0.01	0.001	
base										
beg=0, end=0	6.4	5.6	6.1	14.5	8.5	6.4	10.2	7.0	7.2	8.1
beg=0, end=1	6.3	5.3	5.2	15.4	8.5	6.3	11.6	6.9	6.9	8.2
beg=1, end=0	6.1	5.4	5.9	13.9	7.9	6.3	9.7	6.7	7.1	7.8
beg=1, end=1	6.1	5.0	5.2	14.7	8.4	6.2	10.8	7.0	6.9	8.0
beg=1, end=2	6.1	5.0	5.0	14.9	8.7	6.2	11.5	7.0	6.6	8.0
best										
beg=0, end=0	6.0	5.4	5.6	11.3	7.6	6.3	8.6	6.5	6.8	7.2
beg=1, end=0	6.0	5.3	5.5	10.9	7.4	6.1	8.4	6.6	7.0	7.1
beg=0, end=1	6.0	4.9	5.1	11.9	7.6	5.9	9.2	6.4	6.5	7.2
beg=1, end=1	5.7	5.0	5.1	11.8	7.4	5.9	9.2	6.6	6.5	7.2
beg=1, end=2	5.7	4.7	5.2	11.6	7.1	5.8	9.4	6.4	6.5	7.0

TABLE 3
 z -adjustment as modification step.

		b			c			A			total
		0.1	0.01	0.001	0.1	0.01	0.001	0.1	0.01	0.001	
base	no-adj	6.4	5.6	6.1	14.5	8.5	6.4	10.2	7.0	7.2	8.1
	z -adj	6.3	5.5	5.8	12.5	7.7	6.3	9.2	7.1	7.1	7.6
	WLS-0.01	6.3	5.5	6.1	14.0	8.3	6.4	9.9	7.0	7.1	8.0
	WLS-0.1	7.0	6.6	6.9	12.7	9.1	7.4	8.1	7.1	8.5	8.1
best	no-adj	5.7	4.7	5.2	11.6	7.1	5.8	9.4	6.4	6.5	7.0
	z -adj	5.7	4.8	5.1	10.5	6.8	5.7	8.8	6.3	6.4	6.8
	WLS-0.01	5.7	4.8	5.2	11.6	7.0	5.9	9.3	6.4	6.5	7.0
	WLS-0.1	6.3	5.9	5.9	10.0	7.9	6.8	7.4	6.7	7.7	7.2

Compared with the base, strategy (1, 0) is the best, whereas compared to the best (which just includes higher-order correctors at this point), strategy (1, 2) is preferable. Due to the theoretical benefits of working with a well-centered point, we will use centering strategy (1, 2) in the *best* benchmark strategy for the following tests.

5.1.3. z -adjustment/WLS-step. We have evaluated the benefit of attempting to absorb dual infeasibilities into the z value of the warmstart vector, together with the related WLS heuristic (which attempts to find a least squares correction to the saved iterate such that the resulting point is primal/dual feasible). The results are summarized in Table 3. Surprisingly there is a clear advantage of the simple z -adjustment heuristic, whereas the (computationally more expensive and more sophisticated) WLS step (WLS-0.01) hardly improves on the base strategy. Our only explanation for this behavior is that for our fairly low saved μ -value (2-digit approximate optimal solution to the original problem) the full WLS direction is usually infeasible, so only a fractional step in it can be taken. The z -adjustment, on the other hand, has a more sophisticated fallback strategy which considers adjustment for each component separately, so it is not quite as easily affected by blocking in the modification direction. Reference [15] suggests employing the WLS step together with a backtracking strategy, which saves several iterates from the original problem for different μ and chooses one for which the WLS step does not block. We have emulated this by trying the WLS step for a larger μ (WLS-0.1). Any gain of a larger portion of the WLS step being taken, however, is offset by the starting point now being further away from optimality, resulting in an increase of the number of iterations. We have added the z -adjustment heuristic to our *best* benchmark strategy.

TABLE 4
Splitting directions.

		b			c			A			total
		0.1	0.01	0.001	0.1	0.01	0.001	0.1	0.01	0.001	
base	it=0	6.4	5.6	6.1	14.5	8.5	6.4	10.2	7.0	7.2	8.1
	it=1	6.3	5.5	6.1	14.4	8.6	6.5	10.1	6.9	7.2	8.1
	it=2	6.3	5.5	6.1	14.3	8.6	6.5	10.1	6.9	7.2	8.1
best	it=0	5.7	4.8	5.1	10.5	6.8	5.7	8.8	6.3	6.4	6.8
	it=1	5.7	4.8	5.1	10.5	6.8	5.8	8.7	6.3	6.4	6.8
	it=2	5.8	4.8	5.1	10.4	6.7	5.7	8.7	6.4	6.4	6.8

TABLE 5
Sensitivity based unblocking heuristic.

	b			c			A			
	0.1	0.01	0.001	0.1	0.01	0.001	0.1	0.01	0.001	total
base										
unblk=0	6.4	5.6	6.1	14.5	8.5	6.4	10.2	7.0	7.2	8.1
unblk=1	6.1	5.5	6.0	13.2	8.2	6.4	9.7	7.0	7.1	7.8
unblk=2	6.1	5.3	5.9	12.1	8.1	6.1	9.2	6.8	6.9	7.5
unblk=3	6.0	5.6	6.1	11.4	8.0	6.2	9.0	7.4	7.1	7.5
best: hoc, beg=1, end=2, z-adj										
unblk=0	5.7	4.8	5.1	10.5	6.8	5.7	8.8	6.3	6.4	6.8
unblk=1	5.6	4.8	5.1	9.8	6.5	5.7	8.3	6.4	6.4	6.6
unblk=2	5.7	5.1	5.5	9.4	6.8	5.9	8.2	6.4	6.1	6.7
unblk=3	5.6	5.1	5.7	9.5	6.8	5.8	8.2	6.2	6.5	6.7
beg=0, end=0, z-adj										
unblk=0	6.1	5.0	5.0	14.9	8.7	6.2	11.5	7.0	6.6	8.0
unblk=1	5.9	4.9	5.0	13.2	7.9	6.0	10.4	6.8	6.7	7.6
unblk=2	5.8	5.0	5.0	11.9	8.0	6.1	9.7	6.7	6.9	7.4
unblk=3	5.8	5.2	5.1	11.5	7.7	5.8	9.7	6.8	6.8	7.3
hoc, beg=0, end=0, z-adj										
unblk=0	5.7	4.7	5.2	11.6	7.1	5.8	9.4	6.4	6.5	7.0
unblk=1	5.5	4.8	5.3	10.7	6.9	5.6	9.1	6.5	6.4	6.8
unblk=2	5.8	4.9	5.1	9.8	7.4	5.7	8.7	6.7	5.4	6.7
unblk=3	5.6	5.0	5.5	9.4	6.7	5.7	9.0	6.3	5.5	6.6

5.1.4. Splitting directions. This analyzes the effectiveness of using the computations of separate primal, dual, and centrality correcting directions as in [6] as an unblocking strategy. The results given in Table 4 correspond to different numbers of initial iterations in the solution process of the modified problem using this technique.

As can be seen there is no demonstrable benefit from using this unblocking technique, we have therefore left it out of all subsequent tests.

5.1.5. Unblocking by sensitivity. Finally we have tested the effectiveness of our unblocking scheme based on using sensitivity information. We are considering employing this heuristic for up to the first three iterations. The parameters we have used are $|Z| \leq 0.1n$ (i.e., the worst 10% of components are unblocked), $t = 5$, $\bar{\gamma} = 10$, and $\underline{\gamma} = 0.1$. Results are summarized in Table 5. Unlike the other tests, we have not only tested the unblocking strategy against the base and the best but also against two further setups to evaluate the effectiveness of the strategy to recover from blocking in different environments.

As can be seen there is a clear benefit in employing this heuristic in all tests. The results are less pronounced when comparing against the *best* strategy, but even here there is a clear advantage of performing one iteration of the unblocking strategy.

TABLE 6
Results (best warmstart)—perturbations in b .

Problem	0.1			0.01			0.001		
	cold	warm	red	cold	warm	red	cold	warm	red
ADLITTLE	10.0	6.0	40.0	10.0	5.0	50.0	11.4	6.0	47.3
AFIRO	10.1	4.2	58.4	10.1	4.3	57.4	10.1	4.3	57.4
AGG2	16.1	4.6	71.4	16.2	4.0	75.3	16.1	4.0	75.1
AGG3	15.7	5.6	64.3	15.5	5.0	67.7	16.0	5.0	68.7
BANDM	13.8	8.2	40.5	14.0	4.1	70.7	13.5	4.0	70.3
BEACONFD	-	-	-	-	-	-	-	-	-
BLEND	9.0	4.0	55.5	9.0	4.3	52.2	9.0	4.2	53.3
BOEING1	19.3	7.2	62.6	21.5	8.3	61.3	19.1	5.1	73.2
BORE3D	-	-	-	-	-	-	-	-	-
BRANDY	-	-	-	-	-	-	-	-	-
DEGEN2	-	-	-	-	-	-	-	-	-
E226	16.0	12.8	20.0	15.8	5.0	68.3	15.0	4.8	68.0
GROW15	13.0	4.0	69.2	13.0	4.0	69.2	13.0	4.0	69.2
GROW7	12.0	4.0	66.6	12.0	4.0	66.6	12.0	4.0	66.6
ISRAEL	21.0	6.9	67.1	20.5	4.0	80.4	19.9	4.0	79.8
KB2	17.7	5.0	71.7	17.4	5.0	71.2	17.2	5.0	70.9
LOTFI	19.3	6.8	64.7	20.0	5.7	71.5	20.0	5.8	71.0
RECIPELP	14.0	7.0	50.0	14.0	7.0	50.0	14.5	10.8	25.5
SC105	12.0	5.0	58.3	12.0	5.1	57.5	12.0	5.0	58.3
SC205	12.0	5.2	56.6	12.0	5.0	58.3	12.0	5.0	58.3
SC50A	11.0	4.0	63.6	11.0	4.0	63.6	11.0	4.0	63.6
SC50B	10.0	4.2	58.0	10.0	4.0	60.0	12.1	14.2	-17.3
SCAGR25	12.0	4.8	60.0	11.9	4.1	65.5	12.7	4.0	68.5
SCAGR7	10.1	4.1	59.4	9.9	4.0	59.5	9.8	4.0	59.1
SCFXM1	14.6	5.0	65.7	15.2	5.8	61.8	14.1	4.1	70.9
SCSD1	9.9	9.5	4.0	10.3	5.9	42.7	10.2	5.1	50.0
SCTAP1	14.7	6.0	59.1	14.9	5.0	66.4	15.6	5.3	66.0
SHARE1B	21.5	5.8	73.0	20.8	5.4	74.0	21.3	5.0	76.5
SHARE2B	9.3	5.2	44.0	9.2	5.1	44.5	9.1	5.1	43.9
STOCFOR1	13.5	5.4	60.0	13.0	5.1	60.7	15.4	5.3	65.5
Average	13.8	5.8	56.3	13.8	4.9	62.6	13.9	5.3	60.0

5.2. Results for best warmstart strategy. After these tests we have combined the best setting for all of the considered warmstart heuristics and give more detailed results on the NETLIB test set as well as for a selection of large scale quadratic programming problems.

Tables 6–9 compare the best combined warmstarting strategy for all test problems with a cold start. We give in each case the average number of iterations over 10 random perturbations. Column **red** gives the average percentage iteration reduction achieved by employing the warmstart. An entry “-” denotes that all corresponding perturbations of the problem were either infeasible or unbounded. As can be seen we are able to save between 50% and 60% of iterations on all considered problems.

5.3. Comparison with LOQO results. To judge the competitiveness of our best combined warmstarting strategy, we have compared the results on the NETLIB test set with those reported by [1] which use a different warmstarting methodology. Figure 1 gives a summary of this comparison. The four lines on the left graph give the number of iterations needed for each of the 30 NETLIB problems reported in Tables 6–9 averaged over all perturbations for OOPS and LOQO [1], using a warmstart and a coldstart. As can be seen the default version of OOPS (solid line) needs fewer iterations than LOQO (dotted line). The warmstarted versions of each code

TABLE 7
Results (best warmstart)—perturbations in c .

Problem	0.1			0.01			0.001		
	cold	warm	red	cold	warm	red	cold	warm	red
ADLITTLE	10.3	7.3	29.1	10.1	5.2	48.5	10.4	5.0	51.9
AFIRO	10.3	5.3	48.5	10.3	4.8	53.3	10.7	4.8	55.1
AGG2	16.7	6.6	60.4	16.4	4.8	70.7	16.0	4.1	74.3
AGG3	16.0	6.9	56.8	16.0	5.3	66.8	15.9	4.9	69.1
BANDM	13.7	14.2	-3.6	13.9	5.2	62.5	13.6	4.0	70.5
BEACONFD	10.1	4.7	53.4	10.0	4.0	60.0	11.0	4.8	56.3
BLEND	9.4	7.3	22.3	9.0	4.6	48.8	9.0	4.3	52.2
BOEING1	19.6	24.2	-23.4	19.6	8.6	56.1	19.1	5.8	69.6
BORE3D	12.9	6.1	52.7	13.2	4.4	66.6	13.2	4.2	68.1
BRANDY	15.2	8.7	42.7	15.5	4.3	72.2	15.3	4.0	73.8
DEGEN2	9.8	4.5	54.0	10.0	4.8	52.0	10.0	5.0	50.0
E226	15.6	15.0	3.8	15.2	9.0	40.7	15.1	4.5	70.1
GROW15	22.9	13.7	40.1	22.9	9.2	59.8	17.7	11.0	37.8
GROW7	18.9	14.3	24.3	19.9	12.4	37.6	23.6	17.5	25.8
ISRAEL	20.4	7.7	62.2	21.0	4.2	80.0	21.1	4.3	79.6
KB2	17.8	6.8	61.7	17.9	5.0	72.0	18.0	5.0	72.2
LOTFI	19.0	30.7	-61.5	23.0	20.9	9.1	22.4	12.7	43.3
RECIPERP	-	-	-	-	-	-	-	-	-
SC105	11.4	15.4	-35.0	11.8	5.9	50.0	11.5	5.0	56.5
SC205	12.7	20.9	-64.5	13.1	18.2	-38.9	12.1	6.7	44.6
SC50A	11.2	6.8	39.2	11.0	4.1	62.7	11.0	4.0	63.6
SC50B	10.3	7.2	30.0	10.0	4.4	56.0	10.0	4.0	60.0
SCAGR25	12.0	4.7	60.8	12.4	4.4	64.5	13.0	4.0	69.2
SCAGR7	10.1	4.8	52.4	9.9	4.1	58.5	10.0	4.0	60.0
SCFXM1	14.4	7.4	48.6	14.0	4.0	71.4	14.0	4.0	71.4
SCSD1	9.5	5.2	45.2	9.2	5.0	45.6	9.0	5.0	44.4
SCTAP1	16.2	6.6	59.2	16.1	5.8	63.9	15.8	6.0	62.0
SHARE1B	22.6	8.9	60.6	21.9	6.0	72.6	20.9	5.5	73.6
SHARE2B	9.2	7.2	21.7	9.0	5.0	44.4	9.1	5.0	45.0
STOCFOR1	12.8	5.0	60.9	13.0	5.0	61.5	14.4	5.0	65.2
Average	14.2	9.8	31.1	14.3	6.5	54.1	14.2	5.7	59.8

(solid and dotted lines with markers, respectively) need significantly fewer iterations on average than their coldstarted siblings, with warmstarted OOPS being the most effective strategy over all. This plot indicates only the best combination of interior point code and warmstarting strategy without giving any insight into the relative effectiveness of the warmstarting approaches themselves. In order to measure the efficiency of the warmstart approaches, the second plot in Figure 1 compares the number of iterations *saved* by each warmstarting strategy as compared with its respective cold-started variant. As can be seen our suggested warmstart implemented in OOPS is able to save around 50-60% of iterations, outperforming the LOQO warmstart which averages around 30% saved iterations.

5.4. Medium scale QP problems. We realize that the NETLIB testbed proposed in [1] includes only small LP problems. While this makes it ideal for the extensive testing that we have reported in the previous section, there is some doubt over whether the achieved warmstarting performance can be maintained for quadratic and (more realistic) large scale problems. In order to counter such criticism we have conducted warmstarting tests on two selections of small to medium scale QP problems as well as two sources of large scale quadratic programming. For the small and medium scale tests we have used the quadratic programming collection of Maros

TABLE 8
Results (best warmstart)—perturbations in A.

Problem	0.1			0.01			0.001		
	cold	warm	red	cold	warm	red	cold	warm	red
ADLITTLE	10.8	9.4	12.9	10.5	5.0	52.3	10.4	5.0	51.9
AFIRO	10.1	5.0	50.4	10.0	4.1	59.0	10.0	4.0	60.0
AGG2	15.9	5.3	66.6	16.0	4.2	73.7	16.2	4.0	75.3
AGG3	15.2	6.3	58.5	15.7	5.2	66.8	16.1	5.0	68.9
BANDM	13.8	7.9	42.7	13.8	4.4	68.1	13.4	4.1	69.4
BEACONFD	10.1	4.8	52.4	10.0	4.0	60.0	10.0	4.0	60.0
BLEND	9.0	9.5	−5.5	9.2	5.3	42.3	9.0	4.4	51.1
BOEING1	19.3	5.2	73.0	19.6	5.0	74.4	19.8	5.0	74.7
BORE3D	15.0	4.0	73.3	13.9	4.0	71.2	13.6	4.0	70.5
BRANDY	14.2	14.1	0.7	17.8	15.4	13.4	28.1	18.8	33.0
DEGEN2	11.1	13.4	−20.7	29.2	30.5	−4.4	93.0	86.0	7.5
E226	15.5	10.2	34.1	15.1	4.9	67.5	15.0	4.1	72.6
GROW15	20.2	12.9	36.1	15.3	11.3	26.1	13.4	5.0	62.6
GROW7	24.0	16.1	32.9	17.1	8.8	48.5	13.5	6.4	52.5
ISRAEL	19.8	5.4	72.7	20.0	4.0	80.0	19.9	4.0	79.8
KB2	18.2	15.3	15.9	18.2	5.1	71.9	17.8	5.0	71.9
LOTFI	20.0	7.1	64.5	25.8	12.3	52.3	50.1	36.2	27.7
RECIPELP	13.9	7.1	48.9	13.9	6.6	52.5	14.0	6.0	57.1
SC105	11.8	7.1	39.8	11.5	5.0	56.5	12.0	5.0	58.3
SC205	12.6	7.7	38.8	12.0	5.0	58.3	12.0	5.0	58.3
SC50A	11.1	7.1	36.0	11.0	4.0	63.6	11.0	4.0	63.6
SC50B	10.0	5.1	49.0	10.0	4.0	60.0	10.0	4.0	60.0
SCAGR25	11.7	9.4	19.6	11.8	4.3	63.5	12.5	4.3	65.6
SCAGR7	10.1	6.5	35.6	10.0	4.0	60.0	9.7	4.0	58.7
SCFXM1	15.2	8.0	47.3	14.9	4.6	69.1	14.4	5.0	65.2
SCSD1	9.1	6.3	30.7	9.3	5.2	44.0	9.2	4.8	47.8
SCTAP1	14.2	9.5	33.0	15.6	6.2	60.2	15.1	5.2	65.5
SHARE1B	21.0	9.4	55.2	21.2	7.0	66.9	22.1	5.6	74.6
SHARE2B	9.6	9.9	−3.1	9.2	5.7	38.0	9.0	5.0	44.4
STOCFOR1	11.5	5.8	49.5	12.3	5.2	57.7	12.1	5.1	57.8
Average	14.1	8.4	38.0	14.7	6.7	55.8	17.7	8.9	58.9

and Meszaros [12]. This includes QP problems from the CUTE test set as well as quadratic modifications of the NETLIB LP test set used in the previous comparisons. We have excluded problems that either have free variables (since OOPS currently has no facility to deal with free variables effectively), or where random perturbations of the problem data yield the problem primal or dual infeasible. The same methodology in perturbing the problems as for the NETLIB LP test set has been used, apart that perturbations in the objective function will now perturb random elements of c and Q . The results are displayed in Table 10. As for the LP case we list for each problem and perturbation the average number of iterations needed by OOPS when coldstarted and when warmstarted with the best strategy found in section 5.1 over the 10 random runs and 3 perturbation sizes. We also state the percentage of iterations saved by the warmstart. A blank entry indicates that all 30 random perturbations lead to primal or dual infeasible problems. The results demonstrate a similar performance of our best combined warmstarting strategy as obtained earlier for the LP problems.

5.5. Large scale QP problems. Finally we have evaluated our warmstart strategy in the context of two sources of large scale quadratic problems. In the first

TABLE 9
Results (best warmstart)—all perturbations.

Problem	b			c			A		
	cold	warm	red	cold	warm	red	cold	warm	red
ADLITTLE	10.4	5.6	46.1	10.2	5.8	43.1	10.5	6.4	39.0
AFIRO	10.1	4.2	58.4	10.4	4.9	52.8	10.0	4.3	57.0
AGG2	16.1	4.2	73.9	16.3	5.1	68.7	16.0	4.5	71.8
AGG3	15.7	5.2	66.8	15.9	5.7	64.1	15.6	5.5	64.7
BANDM	13.7	5.4	60.5	13.7	7.8	43.0	13.6	5.4	60.2
BEACONFD	-	-	-	10.3	4.5	56.3	10.0	4.2	58.0
BLEND	9.0	4.1	54.4	9.1	5.4	40.6	9.0	6.4	28.8
BOEING1	19.9	6.8	65.8	19.4	12.8	34.0	19.5	5.0	74.3
BORE3D	-	-	-	13.1	4.9	62.5	14.1	4.0	71.6
BRANDY	-	-	-	15.3	5.6	63.3	20.0	16.1	19.5
DEGEN2	-	-	-	9.9	4.7	52.5	44.4	43.3	2.4
E226	15.6	7.5	51.9	15.3	9.5	37.9	15.2	6.4	57.8
GROW15	13.0	4.0	69.2	21.1	11.3	46.4	16.3	9.7	40.4
GROW7	12.0	4.0	66.6	20.8	14.7	29.3	18.2	10.4	42.8
ISRAEL	20.4	4.9	75.9	20.8	5.4	74.0	19.9	4.4	77.8
KB2	17.4	5.0	71.2	17.9	5.6	68.7	18.0	8.4	53.3
LOTFI	19.7	6.1	69.0	21.4	21.4	0.0	31.9	18.5	42.0
RECIPELP	14.1	8.2	41.8	-	-	-	13.9	6.5	53.2
SC105	12.0	5.0	58.3	11.5	8.7	24.3	11.7	5.7	51.2
SC205	12.0	5.0	58.3	12.6	15.2	-20.6	12.2	5.9	51.6
SC50A	11.0	4.0	63.6	11.0	4.9	55.4	11.0	5.0	54.5
SC50B	10.7	7.4	30.8	10.1	5.2	48.5	10.0	4.3	57.0
SCAGR25	12.2	4.3	64.7	12.4	4.3	65.3	12.0	6.0	50.0
SCAGR7	9.9	4.0	59.5	10.0	4.3	57.0	9.9	4.8	51.5
SCFXM1	14.6	4.9	66.4	14.1	5.1	63.8	14.8	5.8	60.8
SCSD1	10.1	6.8	32.6	9.2	5.0	45.6	9.2	5.4	41.3
SCTAP1	15.0	5.4	64.0	16.0	6.1	61.8	14.9	6.9	53.6
SHARE1B	21.2	5.4	74.5	21.8	6.8	68.8	21.4	7.3	65.8
SHARE2B	9.2	5.1	44.5	9.1	5.7	37.3	9.2	6.8	26.0
STOCFOR1	13.9	5.2	62.5	13.4	5.0	62.6	11.9	5.3	55.4
Average	13.8	5.3	59.6	14.2	7.3	48.4	15.5	8.0	50.9

instance we have solved the capacitated MCNF problem

$$\begin{aligned}
 (25) \quad & \min \sum_{(i,j) \in \mathcal{E}} \frac{x_{ij}}{K_{ij} - x_{ij}}, \\
 & \text{s.t.} \quad \sum_{k \in \mathcal{D}} x_{ij}^{(k)} \leq K_{ij}, \quad \forall (i,j) \in \mathcal{E}, \\
 & \quad Nx^{(k)} = d^{(k)}, \quad \forall k \in \mathcal{D}, \\
 & \quad x^{(k)} \geq 0, \quad \forall k \in \mathcal{D},
 \end{aligned}$$

where N is the node-arc incidence matrix of the network, $d^{(k)}$, $k \in \mathcal{D}$ are the demand points, K_{ij} is the capacity of each arc (i,j) , and x_{ij} is the flow along the arc. This is a nonlinear problem formulation. We have solved it by SQP using the interior point code OOPS as the QP solver and employing our best combined warmstart strategy between QP solutions. We have tested this on nine different MCNF models using from 4–300 nodes, up to 600 arcs, and up to 7021 commodities. The largest problem in the selection has 353,400 variables. All solutions have required more than 10 SQP iterations. Table 11 gives the average number of IPM iterations for each SQP iteration both for cold- and warmstarting the IPM.

FIG. 1. Results of LOQO and OOPS on warmstarting NETLIB problems.

As before we achieve between 50 and 60% reduction in the number of interior point iterations.

Our last test example consists of calculating the complete efficient frontier in a Markowitz Portfolio Selection problem (see [14]). A Portfolio Selection problem aims to find the optimal investment strategy in a selection of assets over time. If the value of the portfolio at the end of the time horizon is denoted by the random variable X , the Markowitz formulation of the portfolio selection problem requires one to maximize the final expected wealth $\mathbb{E}(X)$ and minimize the associated risk, measured as the variance $\text{Var}(X)$ which are combined into a single objective:

$$(26) \quad \min -\mathbb{E}(X) + \rho \text{Var}(X)$$

TABLE 10
Results QP (best warmstart)—all perturbations.

Problem	<i>b</i>			<i>c</i> and <i>Q</i>			<i>A</i>		
	cold	warm	red	cold	warm	red	cold	warm	red
AUG2DCQP	10.0	5.2	48.1	10.0	5.1	49.1	10.9	6.8	37.7
AUG2DQP	10.0	4.1	59.0	10.0	4.1	59.1	10.7	7.0	34.6
AUG3DCQP	8.0	4.0	50.0	8.0	3.8	53.0	8.0	3.7	53.9
AUG3DQP	11.0	5.1	53.4	10.2	5.2	48.9	10.9	4.8	56.3
CVXQP1_S	11.2	5.3	52.9	11.0	5.1	53.8	11.2	5.4	51.7
CVXQP2_M	15.1	4.8	68.1	15.0	5.3	64.5	15.1	5.0	67.1
CVXQP2_S	11.9	6.1	49.2	12.0	6.1	49.0	12.0	6.1	49.4
CVXQP3_M	-	-	-	-	-	-	-	-	-
CVXQP3_S	10.5	6.6	37.1	10.0	5.0	49.9	10.2	6.4	37.2
DUAL1	10.0	4.9	51.2	10.0	5.1	48.6	9.9	5.2	47.6
DUAL2	10.0	5.3	47.1	10.0	4.9	51.1	9.7	4.8	50.7
DUAL3	11.0	6.1	44.1	10.8	5.6	48.0	10.7	5.7	46.5
DUAL4	9.0	4.8	46.9	9.0	5.1	43.1	9.0	5.2	42.2
DUALC1	21.9	3.8	82.6	22.0	4.2	81.1	22.6	4.0	82.4
DUALC2	22.0	3.8	82.6	21.9	3.9	82.0	21.5	3.9	82.0
DUALC5	12.0	3.9	67.9	12.0	4.1	65.4	12.2	3.8	68.4
DUALC8	14.6	3.9	73.2	15.0	4.2	72.3	15.2	3.9	74.5
GOULDQP2	6.0	4.9	18.3	8.0	7.8	2.5	6.0	5.1	14.3
GOULDQP3	9.0	5.0	44.4	9.0	5.1	43.3	9.0	4.9	45.3
HS118	9.0	3.7	59.1	9.0	3.6	59.8	9.0	4.2	53.1
HS21	17.0	6.9	59.5	16.8	6.7	60.3	17.0	7.0	58.7
HS35MOD	9.9	5.8	41.2	9.8	6.0	39.0	9.8	5.9	40.3
HS35	7.0	4.4	37.8	7.1	4.1	42.3	7.0	3.9	44.5
HS53	6.0	5.3	11.9	6.0	5.2	13.7	6.1	4.9	18.6
HS76	7.0	4.1	40.8	7.0	4.1	41.2	7.0	4.0	43.6
HUES-MOD	14.8	5.2	64.6	15.0	4.9	67.5	17.4	12.1	30.4
LOTSCHD	7.9	5.6	29.4	8.0	5.6	29.9	7.6	5.5	28.0
MOSARQP1	7.0	3.9	44.5	7.0	4.1	41.2	7.0	4.4	37.8
MOSARQP2	8.0	4.1	49.2	8.1	4.5	44.2	8.0	3.8	52.0
QPCBOEI1	35.3	23.0	34.8	34.0	22.9	32.4	35.2	23.1	34.3
QPCBOEI2	-	-	-	20.8	7.2	65.4	28.1	12.3	56.0
STCQP1	-	-	-	15.0	5.7	61.9	-	-	-
STCQP2	-	-	-	15.0	7.1	52.7	15.0	6.9	54.1
TAME	6.0	2.2	63.9	6.0	1.9	67.6	6.0	1.9	68.7
ZECEVIC2	7.0	4.8	32.3	7.1	5.2	26.8	7.0	5.3	24.9
25FV47	38.0	7.5	80.3	38.3	7.3	80.8	38.0	8.2	78.4
ADLITTLE	10.6	5.6	47.4	10.2	5.9	42.6	10.2	5.8	43.4
AFIRO	15.1	4.7	69.2	15.0	6.9	54.0	15.0	4.2	72.0
BEACONFD	-	-	-	10.0	4.2	58.2	10.0	4.2	57.9
BORE3D	-	-	-	15.3	4.7	69.0	15.5	4.1	73.5
BRANDY	-	-	-	12.8	5.5	57.3	14.3	14.3	0.1
E226	14.6	7.3	50.11	14.0	7.9	43.7	13.8	5.0	63.5
ETAMACRO	-	-	-	31.9	10.8	66.0	39.5	18.9	52.2
FFFFFF800	63.1	8.6	86.4	61.3	7.0	88.7	56.9	6.6	88.3
GROW15	14.0	5.4	61.6	18.4	13.6	26.1	19.6	10.9	44.5
GROW22	15.9	4.8	69.7	15.6	6.3	59.5	17.0	7.4	56.6
GROW7	15.0	4.9	67.3	21.7	11.5	47.1	17.5	6.5	62.6
ISRAEL	18.9	5.1	73.2	20.1	5.5	72.4	18.4	5.2	71.6
SC205	16.8	7.1	57.6	18.9	19.3	-2.5	17.2	6.9	59.6
SCAGR25J	11.0	4.0	64.1	11.0	3.8	65.3	11.3	4.8	57.2
SCAGR25	11.1	5.0	55.3	11.2	5.2	53.8	11.6	5.3	54.5
SCAGR7	12.4	5.1	59.3	11.8	5.1	56.7	12.0	5.0	58.3
SCFXM1	21.0	7.4	64.6	20.8	7.1	66.0	21.5	7.3	66.1
SCFXM2	23.9	10.7	55.1	23.6	10.5	55.4	24.1	10.8	55.2
SCFXM3	26.6	13.5	49.2	24.5	13.2	46.2	27.8	13.6	51.1
SCORPION	-	-	-	12.1	3.1	74.3	-	-	-

TABLE 10
Continued.

Problem	b			c and Q			A		
	cold	warm	red	cold	warm	red	cold	warm	red
SCRS8	19.6	8.6	56.1	19.6	5.4	72.5	20.0	8.1	59.4
SCSD1	10.8	6.8	36.4	10.2	5.3	48.2	10.2	5.3	47.6
SCSD6	10.9	7.1	34.6	11.4	4.8	58.3	11.1	7.0	37.3
SCSD8	9.0	4.0	55.2	11.3	6.0	46.7	11.4	11.5	-0.2
SCTAP1	13.9	5.5	60.3	15.7	6.5	58.7	15.0	7.5	49.6
SCTAP2	16.0	4.4	72.2	17.9	6.0	66.1	15.1	4.7	68.9
SCTAP3	16.9	5.2	69.5	17.9	6.6	63.1	16.5	5.6	66.2
SEBA	53.3	24.3	54.3	53.7	22.8	57.6	53.5	23.7	55.6
SHARE1B	20.1	6.2	69.2	20.6	6.5	68.6	19.2	6.6	65.8
SHARE2B	24.9	15.1	39.2	24.3	14.9	39.0	25.9	16.8	35.2
SHELL	20.0	7.5	62.3	20.1	6.9	65.8	20.5	9.6	53.4
SHIP04L	-	-	-	11.9	3.7	68.7	11.6	12.0	-3.7
SHIP04S	-	-	-	12.0	4.0	66.8	11.6	11.1	4.3
SHIP08L	-	-	-	11.0	5.0	54.1	11.1	13.3	-19.7
SHIP08S	-	-	-	11.0	4.1	62.4	11.1	9.3	16.4
SHIP12L	-	-	-	16.0	5.1	67.8	14.7	11.4	22.3
SHIP12S	-	-	-	14.4	6.1	57.7	14.2	15.0	-5.1
SIERRAJG	-	-	-	37.4	5.5	85.3	-	-	-
SIERRA	-	-	-	38.3	5.1	86.7	-	-	-
STANDATA	23.0	17.8	22.7	17.4	5.2	70.0	17.9	4.9	72.8
Average	11.5	5.3	53.9	11.8	5.6	52.5	11.9	5.9	50.4

TABLE 11
Capacitated MCNF solved by warmstarted IPM-SQP.

iter	1	2	3	4	5	6	7	8	9	10
cold	12.7	11.9	13.7	15.8	16.2	15.6	14.9	14.6	14.5	15.0
warm	12.7	7.0	6.0	5.8	6.4	7.0	7.0	6.7	6.2	6.0
red	0.0	41.2	56.2	63.3	60.5	55.1	53.0	54.1	57.2	60.0

which leads to a QP problem. We use the multistage stochastic programming version of this model (described in [7]). This formulation leads to very large problem sizes.

The parameter ρ in (26) is known as the *Risk Aversion Parameter* and captures the investor's attitude to risk. A low value of ρ will lead to a riskier strategy with a higher value for the final expected wealth, but a higher risk associated with it.

Often the investor's attitude to risk is difficult to capture a priori in a single parameter. A better decision tool is the efficient frontier, a plot of $\mathbb{E}(X)$ against the corresponding $\text{Var}(X)$ values for different settings of ρ . Computing the efficient frontier requires the solution of a series of problems for different values of ρ . Apart from this all the problems in the sequence are identical, which makes them prime candidates for a warmstarting strategy (although see [3] for a different approach). Table 12 gives results for four different problem sizes with up to 192 million variables and 70 million constraints. For each problem the top line gives the number of iterations a coldstarted IPM needs to solve the problem for a given value of ρ , whereas the middle line gives the number of iterations when warmstarting each problem from the one with the next lowest setting of ρ . The last line gives the percentage saving in IPM iterations. Again we are able to save in the range of 50 and 60% of IPM iterations. As far as we are aware these are the largest problems to date for which an interior point warmstart has been employed.

TABLE 12
Computation of efficient frontier with IPM warmstarts.

variables (n) constraints (m)	$\rho =$	1e-3	5e-3	0.01	0.05	0.1	0.5	1	5	10
$n = 223.321$ $m = 76.881$	cold	14	14	14	14	14	13	17	16	17
	warm	14	5	5	5	4	5	5	8	8
	red	0.0	64.2	64.2	64.2	71.4	61.5	70.5	50.0	52.9
$n = 533.725$ $m = 198.525$	cold	14	14	14	14	14	15	18	18	17
	warm	14	5	5	5	6	5	5	9	10
	red	0.0	64.3	64.3	64.3	57.1	66.7	72.2	50.0	41.2
$n = 16.316.191$ $m = 5.982.604$	cold	24	23	24	23	25	22	24	23	24
	warm	24	8	11	13	11	13	12	12	14
	red	0.0	65.2	54.2	43.5	56.0	40.9	50.0	47.8	41.7
$n = 192.478.111$ $m = 70.575.308$	cold	52	53	45	43	44	42	44	46	46
	warm	52	13	13	15	15	16	16	23	25
	red	0.0	75.5	71.1	65.1	65.9	61.9	63.6	50.0	45.6

6. Conclusions. In this paper we have compared the effectiveness of various interior point warmstarting schemes on the NETLIB base test set suggested by [1]. We have categorized warmstarting strategies into *modification* strategies and *unblocking* strategies. *Modification* strategies are aimed at modifying an advanced iterate from a previous solution of a nearby problem before it is used to warmstart an IPM, whereas unblocking strategies aim to directly address the negative effect known as blocking which typically affects a “bad” warmstart in the first few iterations. We suggest a new unblocking strategy based on sensitivity analysis of the step direction with respect to the current point. In our numerical tests we obtain an optimal combination of modification and unblocking strategies (including the new strategy based on sensitivity analysis) and are subsequently able to save an average of 50 and 60% of interior point iterations on a range of LP and QP problems varying from the small scale NETLIB test set to problems with over 192 million variables.

Acknowledgment. We would like to thank the anonymous referees for their constructive comments that helped to improve this paper.

REFERENCES

- [1] H. Y. BENSON AND D. F. SHANNO, *An exact primal-dual penalty method approach to warm-starting interior-point methods for linear programming*, Comput. Optim. Appl., 38 (2007), pp. 371–399.
- [2] H. Y. BENSON AND D. F. SHANNO, *Interior-point methods for nonconvex nonlinear programming: Regularization and warmstarts*, Comput. Optim. Appl., 40 (2008), pp. 143–189.
- [3] J. FLIEGE, *An efficient interior-point method for convex multicriteria optimization problems*, Math. Oper. Res., 31 (2006), pp. 825–845.
- [4] J. GONDZIO, *HOPDM (version 2.12)—a fast LP solver based on a primal-dual interior point method*, European J. Oper. Res., 85 (1995), pp. 221–225.
- [5] J. GONDZIO, *Warm start of the primal-dual method applied in the cutting plane scheme*, Math. Program., 83 (1998), pp. 125–143.
- [6] J. GONDZIO AND A. GROTHEY, *Reoptimization with the primal-dual interior point method*, SIAM J. Optim., 13 (2003), pp. 842–864.
- [7] J. GONDZIO AND A. GROTHEY, *Parallel interior point solver for structured quadratic programs: Application to financial planning problems*, Ann. Oper. Res., 152 (2007), pp. 319–339.
- [8] J. GONDZIO AND J.-P. VIAL, *Warm start and ε -subgradients in cutting plane scheme for block-angular linear programs*, Comput. Optim. Appl., 14 (1999), pp. 17–36.
- [9] A. L. HIPOLITO, *A weighted least squares study of robustness in interior point linear programming*, Comput. Optim. Appl., 2 (1993), pp. 29–46.

- [10] E. JOHN AND E. A. YILDIRIM, *Implementation of warm-start strategies in interior-point methods for linear programming in fixed dimension*, Comput. Optim. Appl., 41 (2008), pp. 151–183.
- [11] I. J. LUSTIG, R. E. MARSTEN, AND D. F. SHANNO, *Computational experience with a globally convergent primal-dual predictor-corrector algorithm for linear programming*, Math. Program., 66 (1994), pp. 123–135.
- [12] I. MAROS AND C. MÉSZÁROS, *A repository of convex quadratic programming problems*, Technical report DOC 97/6, Department of Computing, Imperial College, London, U.K., 1997.
- [13] S. MEHROTRA, *On the implementation of a primal-dual interior point method*, SIAM J. Optim., 2 (1992), pp. 575–601.
- [14] M. STEINBACH, *Markowitz revisited: Mean-variance models in financial portfolio analysis*, SIAM Rev., 43 (2001), pp. 31–85.
- [15] E. A. YILDIRIM AND S. J. WRIGHT, *Warm-start strategies in interior-point methods for linear programming*, SIAM J. Optim., 12 (2002), pp. 782–810.