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

Jacek Gondzio Andreas Grothey

December 19th, 2006

MS-06-005

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

Jacek Gondzio* Andreas Grothey†

School of Mathematics
The University of Edinburgh
Mayfield Road, Edinburgh EH9 3JZ
United Kingdom.

December 19, 2006

*Email: J.Gondzio@ed.ac.uk, URL: http://www.maths.ed.ac.uk/~gondzio/
†Email: A.Grothey@ed.ac.uk, URL: http://www.maths.ed.ac.uk/~agr/

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

Abstract

One of the main drawbacks associated with Interior Point Methods (IPM) 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 [1]. 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%-60% of interior point iterations, consistently outperforming similar approaches reported in current optimisation literature.

1 Introduction

Since their introduction, Interior Point Methods (IPMs) have been recognised 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 can not 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 interior point methods [5, 8, 14, 7, 1, 2, 10]. All of these, apart from [1, 2], involve 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 LP case, whereas we are equally interested in the 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. Hence the next iteration point is close to the previous one, resulting in a repeat of the bad performance. 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 affine scaling direction too much. The unblocking strategy is based on performing sensitivity analysis of the affine scaling direction with respect to the components of the current primal/dual iterate.

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 technique are reported in Section 5. In the final Section 6, we draw our conclusions.

2 Notation & Background

The infeasible primal dual interior point methods applied to solve the quadratic programming problem

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

can be motivated from the KKT conditions for (1)

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

where the zero right hand side of the complementarity products has been replaced by the centrality parameter μ. 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 neighbourhoods of the central path, the N_2 neighbourhood

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

and the wider $N_{-\infty}$ neighbourhood

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

Assume that at some stage during the algorithm the current iterate is (x, y, z). Our variant of the predictor-corrector algorithm [4, 6] 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 ($\mu^0 \approx 0.001\frac{x^T z}{n}$):

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

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

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

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

As in Mehrotra's predictor-corrector algorithm[12] we calculate maximal primal and dual step-sizes for the predictor direction

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

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

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

and finally the new primal and dual step sizes 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

$$\begin{array}{ll} \min & \tilde{c}^T x + \frac{1}{2} x^T \tilde{Q} x \\ \text{s.t.} & \tilde{A}x = \tilde{b} \\ & x \geq 0 \end{array} \tag{6}$$

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 Interior Points Methods 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, 14, 7] 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 centre [5], a point close to the central path of the original problem (usually a non-converged 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 re-solve unblocking 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 (or a selection of) 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

$$\begin{aligned} \tilde{c} + \tilde{Q}\tilde{x} - \tilde{A}^T\tilde{y} - \tilde{z} &= 0 \\ \tilde{b} - \tilde{A}\tilde{x} &= 0 \\ \tilde{x}_i\tilde{z}_i &\approx \tilde{\mu}. \end{aligned}$$

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 ($\mu \lessapprox \mu^0$), the Newton step is feasible. Analysis by [14] and [7] 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:

- **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 $N_\infty(\gamma)$.

- [14, 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 weighted least squares step does not necessarily preserve positiveness of the iterate. To overcome this, [14] suggest 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 [12] can be seen as a (damped) weighted least squares step from the origin.
- [14, 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

- **Splitting Directions**: [7] advocate 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.
- **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. [5] employs this idea by additionally correcting only for small complementarity products to avoid introducing additional blocking. This is used in [7] as a unblocking technique with the interpretation of choosing a target complementarity vector $\bar{t} \approx \mu e$ in such a way that as much as possible of the Newton step is feasible, aiming to absorb as much as possible of the primal/dual infeasibility in the first step.
- **Change Diagonal Scaling:** [9] investigates changing elements in the scaling matrix $\Theta = XZ^{-1}$ to make nearby constraints repelling rather than attracting to the Newton step.

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

- **Dual adjustment**: Adjust advanced starting point $\tilde{z}$ to compensate for changes to c and A in the dual feasibility constraint (2a).
- **Balancing complementarity products** of the advanced starting point before use in new problem.
- **Additional centering** iterations before advanced starting point is saved.
- Unblocking of the step direction by **sensitivity analysis**.

We will give a brief description of the first three of these strategies. The fourth (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 dual feasible starting point. 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 weighted least squares modification step. Unlike this, however, direct adjustment of z is much cheaper to compute.

Balancing complementarity products
According to [14, 7] a well centered point can absorb larger primal/dual infeasibilities in a single Newton step. Also very small complementarity products $x_i z_i$ are likely to cause bad performance as observed by [11]. An appealing idea is therefore to attempt to improve the centrality of the saved iterate by small changes to x, z where possible. For this we project the complementarity products $\tilde{x}_i\tilde{z}_i$ onto the interval $[\frac{1}{\rho}\mu, \rho\mu]$, where $\rho \geq 1$ is a parameter. Since we want to change the point $(\tilde{x}, \tilde{z})$ as little as possible in order not to introduce additional infeasibilities, we do this by changing the smaller of $\tilde{x}_i, \tilde{z}_i$. If both $\tilde{x}_i, \tilde{z}_i \leq \sqrt{\mu}$ the resulting change could be very large, so instead we set

$$x_i = \tilde{x}_i + \delta, \quad z_i = \tilde{z}_i + \delta,$$

where $\delta \geq 0 : (\tilde{x}_i + \delta)(\tilde{z}_i + \delta) = \mu/\rho$. The changes introduced by this procedure are therefore bounded by $\sqrt{\mu}$.

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 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 identical to the Newton Modification Step of [14, 10]. We refer to these as centering iteration at the *beginning* of the modified problem or at the *end* of 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 solve of the modified problem is to some degree arbitrary. It is therefore

possible to treat it as parameters to the solve and to explore how certain properties of the solve change if 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 affine scaling direction $(\Delta x, \Delta y, \Delta z)$ is obtained as the solution to the system (3) or (4). If we think of (x, y, z, μ) 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, μ). The aim of this section is to derive a procedure by which the sensitivity of $\Delta x(x, y, z, \mu)$, $\Delta y(x, y, z, \mu)$, $\Delta z(x, y, z, \mu)$, that is the first derivatives of these function, 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

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

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 only interested 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):

$$-Q\frac{\mathrm{d}\Delta x}{\mathrm{d}x_i} + A^T\frac{\mathrm{d}\Delta y}{\mathrm{d}x_i} + \frac{\mathrm{d}\Delta z}{\mathrm{d}x_i} = Qe_i \tag{8a}$$

$$A\frac{\mathrm{d}\Delta x}{\mathrm{d}x_i} = -Ae_i \tag{8b}$$

$$X\frac{\mathrm{d}\Delta z}{\mathrm{d}x_i} + Z\frac{\mathrm{d}\Delta x}{\mathrm{d}x_i} + \Delta Z e_i = -Ze_i. \tag{8c}$$

Similarly differentiating with respect to y_i yields

$$-Q\frac{\mathrm{d}\Delta x}{\mathrm{d}y_i} + A^T\frac{\mathrm{d}\Delta y}{\mathrm{d}y_i} + \frac{\mathrm{d}\Delta z}{\mathrm{d}y_i} = -A^T e_i \tag{9a}$$

$$A\frac{\mathrm{d}\Delta x}{\mathrm{d}y_i} = 0 \tag{9b}$$

$$X\frac{\mathrm{d}\Delta z}{\mathrm{d}y_i} + Z\frac{\mathrm{d}\Delta x}{\mathrm{d}y_i} = 0 \tag{9c}$$

and finally differentiating with respect to z_i yields

$$-Q\frac{\mathrm{d}\Delta x}{\mathrm{d}z_i} + A^T\frac{\mathrm{d}\Delta y}{\mathrm{d}z_i} + \frac{\mathrm{d}\Delta z}{\mathrm{d}z_i} = -e_i \tag{10a}$$

$$A\frac{\mathrm{d}\Delta x}{\mathrm{d}z_i} = 0 \tag{10b}$$

$$X\frac{\mathrm{d}\Delta z}{\mathrm{d}z_i} + Z\frac{\mathrm{d}\Delta x}{\mathrm{d}z_i} + \Delta X e_i = -Xe_i. \tag{10c}$$

Taking all three systems together we have

$$\begin{bmatrix} -Q & A^T & I \\ A & 0 & 0 \\ Z & 0 & X \end{bmatrix} \begin{bmatrix} \frac{\mathrm{d}\Delta x}{\mathrm{d}x} & \frac{\mathrm{d}\Delta x}{\mathrm{d}y} & \frac{\mathrm{d}\Delta x}{\mathrm{d}z} \\ \frac{\mathrm{d}\Delta y}{\mathrm{d}x} & \frac{\mathrm{d}\Delta y}{\mathrm{d}y} & \frac{\mathrm{d}\Delta y}{\mathrm{d}z} \\ \frac{\mathrm{d}\Delta z}{\mathrm{d}x} & \frac{\mathrm{d}\Delta z}{\mathrm{d}y} & \frac{\mathrm{d}\Delta z}{\mathrm{d}z} \end{bmatrix} = \begin{bmatrix} Q & -A^T & -I \\ -A & 0 & 0 \\ -Z - \Delta Z & 0 & -X - \Delta X \end{bmatrix} \tag{11}$$

Under the assumption that A has full row rank, the system matrix is non-singular, therefore

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

$$\begin{bmatrix} \frac{\mathrm{d}\Delta x}{\mathrm{d}y_i} \\ \frac{\mathrm{d}\Delta y}{\mathrm{d}y_i} \\ \frac{\mathrm{d}\Delta z}{\mathrm{d}y_i} \end{bmatrix} = \begin{bmatrix} 0 \\ -e_i \\ 0 \end{bmatrix} \tag{12b}$$

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

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

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

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

$$\begin{bmatrix} -Q - X^{-1}Z & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \widetilde{\mathrm{d}\Delta x} \\ \widetilde{\mathrm{d}\Delta y} \end{bmatrix} = \begin{bmatrix} X^{-1}e_i \\ 0 \end{bmatrix} \tag{13a}$$

$$\widetilde{\mathrm{d}\Delta z} = -X^{-1}Z\widetilde{\mathrm{d}\Delta x} - X^{-1}e_i \tag{13b}$$

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

Secondly we can get the complete sensitivity information with respect to (x_i, z_i) for a given component i by a single backsolve with the same augmented system that has been used to obtain the step $(\Delta x, \Delta y, \Delta z)$; n backsolves 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 Affine Scaling Direction using Sensitivity Information

Occasionally, despite all our attempts, a starting point might lead to a Newton direction that blocks: 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.

Assume that at the current point (x, z) a step $(\Delta x, \Delta z)$ has been calculated by (3) that is blocking, i.e.

$$l := -\min_i \{x_i + \Delta x_i, z_i + \Delta z_i\} \gg 0. \tag{14}$$

We seek a correction (d_x, d_z) to the current point that removes the blocking, i.e.

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

Using the results from the sensitivity analysis, this suggests taking (d_x, d_z) to satisfy

$$\begin{aligned} x + d_x + \Delta x + \frac{\mathrm{d}\Delta x}{\mathrm{d}x} d_x + \frac{\mathrm{d}\Delta x}{\mathrm{d}z} d_z &\geq 0 \\ z + d_z + \Delta z + \frac{\mathrm{d}\Delta z}{\mathrm{d}x} d_x + \frac{\mathrm{d}\Delta z}{\mathrm{d}z} d_z &\geq 0 \end{aligned}$$

or after re-arranging

$$\Xi \begin{pmatrix} d_x \\ d_z \end{pmatrix} \geq \begin{pmatrix} x + \Delta x \\ z + \Delta z \end{pmatrix}^{-}, \tag{15}$$

where

$$\Xi = \begin{bmatrix} \frac{\mathrm{d}\Delta x}{\mathrm{d}x} + I & \frac{\mathrm{d}\Delta x}{\mathrm{d}z} \\ \frac{\mathrm{d}\Delta z}{\mathrm{d}x} & \frac{\mathrm{d}\Delta z}{\mathrm{d}z} + I \end{bmatrix}$$

and $(x)^- := -\min\{x, 0\}$ componentwise (i.e. the absolute value of the negative components of x). The following Lemma shows that a variation of such a step will indeed succeed in unblocking the step, providing l from (14) is not too large. To state the Lemma we need the following definitions. Let

$$\mathcal{C}_l = \{(x, z) : x + \Delta x(x, z) \geq -le, z + \Delta z(x, z) \geq -le\}$$

and observe that

$$\|(x + \Delta x, z + \Delta z)^-\|_\infty \leq l, \quad \forall (x, z) \in \mathcal{C}_l.$$

Further for $u > 0$, let Ω_u be the set of (d_x, d_z) such that

$$\Xi \begin{pmatrix} d_x \\ d_z \end{pmatrix} \geq \begin{pmatrix} x + \Delta x \\ z + \Delta z \end{pmatrix}^{-} + L\|\Xi^{-1}\|_\infty e \tag{16a}$$

$$\begin{pmatrix} d_x \\ d_z \end{pmatrix} \geq \begin{pmatrix} -x \\ -z \end{pmatrix} \tag{16b}$$

$$\|(d_x, d_z)\|_\infty \leq u. \tag{16c}$$

This allows us to state the main result of this section

Lemma.

- there exists an $L > 0$ and $u > 0$ such that for every $(x, z) \in \mathcal{C}_l, l \leq L$, any (d_x, d_z) in Ω_u will unblock the step, that is the full Newton Step from $(x + d_x, z + d_z)$ is feasible.
- $\Omega_u \neq \emptyset$. In particular

$$\begin{pmatrix} d_x \\ d_z \end{pmatrix} = \Xi^{-1} \begin{pmatrix} x + \Delta x \\ z + \Delta z \end{pmatrix}^{-} + L\|\Xi^{-1}\|_\infty \Xi^{-1} e \in \Omega_u, \tag{17}$$

Proof:
The solution to the Newton step equation $(\Delta x, \Delta z) = (\Delta x(x, z), \Delta z(x, z))$ is a function of the current point (x, z). Let $\epsilon = \kappa^{-1}(1 + \|\Xi^{-1}\|_\infty)^{-1}$ for some $\kappa > 1$.

Since $\Delta x, \Delta z$ are differentiable functions, there exists a $\delta \in (0,1)$ such that

$$\|\Delta x(x+d_x, z+d_z) - \Delta x(x,z) - \frac{\mathrm{d}\Delta x}{\mathrm{d}x}d_x - \frac{\mathrm{d}\Delta x}{\mathrm{d}z}d_z\|_\infty \leq \epsilon\|(d_x,d_z)\|_\infty \tag{18a}$$

$$\|\Delta z(x+d_x, z+d_z) - \Delta z(x,z) - \frac{\mathrm{d}\Delta z}{\mathrm{d}x}d_x - \frac{\mathrm{d}\Delta z}{\mathrm{d}z}d_z\|_\infty \leq \epsilon\|(d_x,d_z)\|_\infty \tag{18b}$$

for all $(d_x, d_z) : \|(d_x,d_z)\|_\infty \leq \delta$. Then with

$$L = \min\{\delta[\|\Xi^{-1}\|_\infty(1+\|\Xi^{-1}\|_\infty)]^{-1}, [\|\left[\begin{smallmatrix} X & 0 \\ 0 & Z \end{smallmatrix}\right]^{-1}\Xi^{-1}\|_\infty(1+\|\Xi^{-1}\|_\infty)]^{-1}\}$$

and

$$u = \min\{\delta, \kappa L(1+\|\Xi^{-1}\|_\infty)\|\Xi^{-1}\|_\infty\}$$

the assertion of the Lemma is satisfied.

To show the assertion we first show that the vector (d_x, d_z) defined in (17) satisfies (16), and then that any unblocking step in Ω_u leads to a full Newton step.

For the first part consider (d_x, d_z) from (17):

$$\begin{pmatrix} d_x \\ d_z \end{pmatrix} = \Xi^{-1}\begin{pmatrix} x+\Delta x \\ z+\Delta z \end{pmatrix}^- + L\|\Xi^{-1}\|_\infty \Xi^{-1}e. \tag{19}$$

Then

$$\Xi(d_x,d_z)^T = \begin{pmatrix} x+\Delta x \\ z+\Delta z \end{pmatrix}^- + L\|\Xi^{-1}\|_\infty e$$

which proves (16a). By premultiplying (19) with $\left[\begin{smallmatrix} X & 0 \\ 0 & Z \end{smallmatrix}\right]^{-1}$ we obtain

$$\begin{aligned} \|(X^{-1}d_x, Z^{-1}d_z)\|_\infty &\leq \|\left[\begin{smallmatrix} X & 0 \\ 0 & Z \end{smallmatrix}\right]^{-1}\Xi^{-1}\|_\infty(L + L\|\Xi^{-1}\|_\infty) \\ &= L(1+\|\Xi^{-1}\|_\infty)\|\left[\begin{smallmatrix} X & 0 \\ 0 & Z \end{smallmatrix}\right]^{-1}\Xi^{-1}\|_\infty \leq 1 \end{aligned}$$

which implies (16b), and

$$\|(d_x,d_z)\|_\infty \leq \|\Xi^{-1}\|_\infty\|Le\|_\infty + L\|\Xi^{-1}\|_\infty^2 = L(1+\|\Xi^{-1}\|_\infty)\|\Xi^{-1}\|_\infty$$

which implies (16c).

Now let (d_x, d_z) be an unblocking step satisfying (16). First note that due to (16b) $(x+d_x, z+d_z) \geq 0$, therefore the unblocking step itself can be taken. Further

$$\begin{aligned} & x + d_x + \Delta x(x+d_x, z+d_z) \\ &\geq x + d_x + \Delta x(x,z) + \frac{\mathrm{d}\Delta x}{\mathrm{d}x}d_x + \frac{\mathrm{d}\Delta x}{\mathrm{d}z}d_z - \epsilon\|(d_x,d_z)\|_\infty e \\ &= x + \Delta x(x,z) + [\frac{\mathrm{d}\Delta x}{\mathrm{d}x} + I]d_x + \frac{\mathrm{d}\Delta x}{\mathrm{d}z}d_z - \epsilon\|(d_x,d_z)\|_\infty e \\ &\geq x + \Delta x(x,z) + [x+\Delta x]^- + L\|\Xi^{-1}\|_\infty e - \epsilon\|(d_x,d_z)\|_\infty e && \text{from (16}a\text{)} \\ &\geq L\|\Xi^{-1}\|_\infty e - \epsilon\|(d_x,d_z)\|_\infty e, && \text{using } y + [y]^- \geq 0 \\ &\geq L\|\Xi^{-1}\|_\infty e - \epsilon\kappa L(1+\|\Xi^{-1}\|_\infty)\|\Xi^{-1}\|_\infty e, && \text{using (16}c\text{)} \\ &= 0, && \text{using def. of } \kappa \end{aligned}$$

and similarly for $z + d_z + \Delta z(x+d_x, z+d_z)$ that is, the full Newton Step from $(x+d_x, z+d_z)$ is feasible. □

4.3 Implementation

Clearly obtaining the complete sensitivity information and finding an unblocking direction that satisfies (16) is computationally expensive. Instead we suggest a heuristic to achieve the same effect.

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 furthermore these are most effectively influenced by changing the corresponding components of (x, z). 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 mainly influence the i-th component of the step $\Delta x_i, \Delta z_i$. 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.

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 re-compute 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 $\frac{\mathrm{d}\Delta x_i}{\mathrm{d}x_i}$, 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 affine scaling direction is unblocked, that is say

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

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

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

We need to distinguish several cases

- $\frac{\mathrm{d}\Delta x_i}{\mathrm{d}x_i} \leq \frac{\Delta x_i}{x_i}$:
  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

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

- $\frac{\mathrm{d}\Delta x_i}{\mathrm{d}x_i} > \frac{\Delta x_i}{x_i}$:
  A positive step would weaken the blocking. However if $\frac{\mathrm{d}\Delta x_i}{\mathrm{d}x_i} < -t$ 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 which 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

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

as

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

In this case $\widetilde{d_{z,i}}$ indicates the correct sign of the change, but for $\frac{\mathrm{d}\Delta x_i}{\mathrm{d}z_i}$ 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

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

and

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

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 affine scaling direction.

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 both the NETLIB test set 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[6]. For all tests we save the first iterate in the original problem solve that satisfies

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

for use as warmstarting point. Since our motivation is primarily to evaluate unblocking techniques in order to recover from 'bad' warmstarting situations, we do not attempt to find an 'optimal' value for $\bar{\mu}$. 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 and 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 δ for every vector component. Should this number be less than $\min\{0.1, 20/n\}$ (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 7-10.

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 calculation 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 `no-hoc` respectively). The numbers reported are the average number of iteration 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

	b			c			A			
	0.1	0.01	0.001	0.1	0.01	0.001	0.1	0.01	0.001	total
no-hoc	6.1	5.0	5.0	14.9	8.7	6.2	11.5	7.0	6.6	8.0
hoc	5.7	4.7	5.2	11.6	7.1	5.8	9.4	6.4	6.5	7.0

Table 1: Higher Order Correctors as Unblocking Device

recommend the use of higher order correctors in all circumstances in the context of unblocking. 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 solve, before any reduction of the barrier μ is applied (`beg`). As pointed out earlier the latter corresponds to the Newton corrector step of [14]. 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. Strategies $(0,1),(1,1),(1,2)$ are

	b			c			A			
	0.1	0.01	0.001	0.1	0.01	0.001	0.1	0.01	0.001	total
beg=0,end=0	6.1	5.4	5.3	10.2	7.4	6.4	8.0	6.3	7.1	7.0
beg=1,end=0	6.0	5.3	5.2	10.0	7.1	6.2	7.9	6.8	7.0	6.9
beg=0,end=1	6.1	4.9	5.0	10.4	6.9	5.9	8.5	6.4	6.8	6.8
beg=1,end=1	5.8	5.0	5.0	10.2	7.0	6.0	8.4	6.2	6.5	6.8
beg=1,end=2	5.7	4.8	5.1	10.5	6.8	5.7	8.8	6.3	6.4	6.8

Table 2: Additional Centering Iterations

about on par and preferable to the default of no centering iterations.

5.1.3 Splitting Directions

This analyzes the effectiveness of using the computations of separate primal, dual and centrality correcting directions as in [7] as an unblocking strategy. The results given in Table 3 correspond to different numbers of initial iterations in the solution process of the modified problem using this technique.

	b			c			A			
	0.1	0.01	0.001	0.1	0.01	0.001	0.1	0.01	0.001	total
it=0	5.7	4.7	5.2	11.6	7.1	5.8	9.4	6.4	6.5	7.0
it=1	5.7	4.8	5.2	12.0	6.9	5.8	9.4	6.5	6.5	7.1
it=2	5.8	4.8	5.2	11.9	7.2	5.7	9.8	6.5	6.5	7.2

Table 3: Splitting Directions

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

5.1.4 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 4. 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

	b			c			A			
	0.1	0.01	0.001	0.1	0.01	0.001	0.1	0.01	0.001	total
base	5.7	4.7	5.2	11.6	7.1	5.8	9.4	6.4	6.5	7.0
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
z-adj	5.7	4.8	5.1	10.5	6.8	5.7	8.8	6.3	6.4	6.8

Table 4: z-Adjustment as Modification Step

behaviour, 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 fall-back strategy which considers adjustment for each component separately, so is not quite as easily affected by blocking in the modification direction. [14] suggest employing the WLS step together with a backtracking strategy, that 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.

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 $|\mathcal{I}| \leq 0.1n$ (i.e. the worst 10% of components are unblocked), $t = 5$ and $\bar{\gamma} = 10, \underline{\gamma} = 0.1$. Results are summarized in Tables 5/6 giving results with and without the combined use of higher order correctors.

	b			c			A			
	0.1	0.01	0.001	0.1	0.01	0.001	0.1	0.01	0.001	total
no-sens	6.1	5.0	5.0	14.9	8.7	6.2	11.5	7.0	6.6	8.0
sens-1	5.9	4.9	5.0	13.2	7.9	6.0	10.4	6.8	6.7	7.6
sens-2	5.8	5.0	5.0	11.9	8.0	6.1	9.7	6.7	6.9	7.4
sens-3	5.8	5.2	5.1	11.5	7.7	5.8	9.7	6.8	6.8	7.3

Table 5: Sensitivity Based Unblocking Heuristic

	b			c			A			total
δ	0.1	0.01	0.001	0.1	0.01	0.001	0.1	0.01	0.001	
hoc.0	5.7	4.7	5.2	11.6	7.1	5.8	9.4	6.4	6.5	7.0
hoc.1	5.5	4.8	5.3	10.7	6.9	5.6	9.1	6.5	6.4	6.8
hoc.2	5.8	4.9	5.1	9.8	7.4	5.7	8.7	6.7	5.4	6.7
hoc.3	5.6	5.0	5.5	9.4	6.7	5.7	9.0	6.3	5.5	6.6

Table 6: Sensitivity Based Unblocking Heuristic (with higher-order correctors).

There is a clear benefit in employing this heuristic in the first step, and a smaller but still noticeable improvement in the next two steps.

5.2 Results for best warmstart strategy

After these tests we have combined the best setting for all 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 7/8/9/10 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

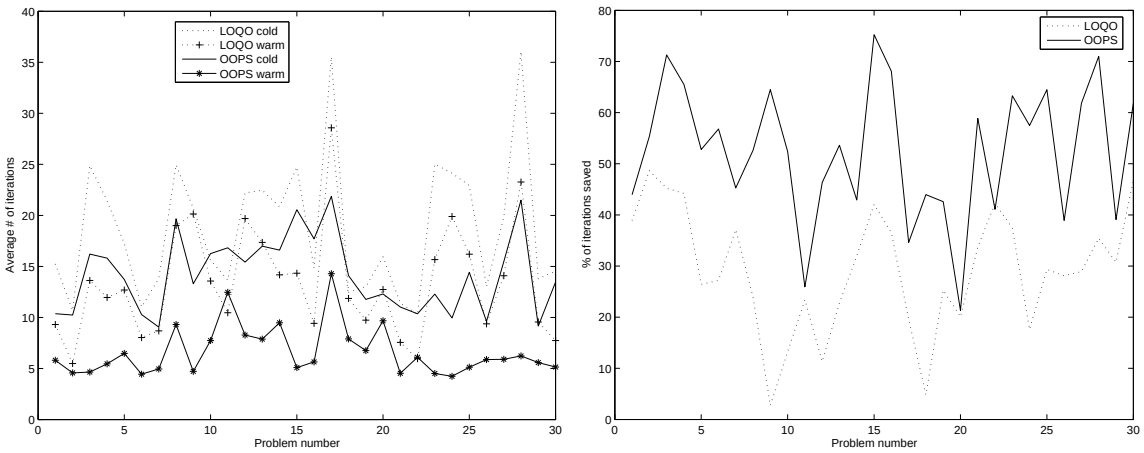

Figure 1: Results of LOQO and OOPS on warmstarting NETLIB problems

the number of iterations needed for each of the 30 NETLIB problems reported in Tables 7-10 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 (solid and dotted lines with markers respectively) need significantly fewer iterations on average than their cold started siblings, with warmstarted OOPS being the most effective strategy over all. 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%. This is especially remarkable given the fact that coldstarted OOPS already takes about 20% fewer iterations than coldstarted LOQO, which should translate into less scope for further reductions.

5.3.1 Large Scale QP Problems

We realise that the NETLIB testbed proposed in [1] only includes 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 (more realistic) large scale problems and nonlinear problems. In order to counter such criticism we have conducted warmstarting tests on two sources of large scale quadratic programming problems. In the first instance we have solved the capacitated Multi-commodity Network Flow (MCNF) problem

$$\begin{array}{lll} \min & \sum_{(i,j)\in\mathcal{E}} \frac{x_{ij}}{K_{ij}-x_{ij}} & \\ \text{s. t.} & \sum\limits_{k\in\mathcal{D}} x_{ij}^{(k)} \le K_{ij}, & \forall (i,j)\in\mathcal{E}, \\ & Nx^{(k)} = d^{(k)}, & \forall k\in\mathcal{D}, \\ & x^{(k)} \ge 0, & \forall k\in\mathcal{D}. \end{array} \tag{20}$$

where N is the node-arc incidence matrix of the network, $d^{(k)}, k \in \mathcal{D}$ are the demands, 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 solves. We have tested this on 9 different MCNF models using between 4-300 nodes, up to 600 arcs, and up to 7021 commodities. The largest problem in the selection has $353,400$ variables. All solves 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.

As before we achieve between 50%-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 [13]). 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 time horizon is denoted by the random variable X, the Markowitz formulation of the portfolio selection problem requires to maximize the final expected wealth $I\!E(X)$ and minimize the associated risk, measured as the variance $\mathrm{Var}(X)$ which are combined into a single objective:

$$\min -I\!E(X) + \rho \mathrm{Var}(X) \tag{21}$$

which leads to a QP problem. We use the multistage stochastic programming version of this model (described in [6]). This formulation leads to very large problem sizes.

The parameter ρ in (21) is known as the *Risk Aversion Parameter* and captures the investors 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 investors attitude to risk is difficult to capture *a-priori* in a single parameter. A better decision tool is the efficient frontier, a plot of $I\!E(X)$ against the corresponding $\mathrm{Var}(X)$ 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 cold started IPM needed 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%-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.

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 categorised 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%-60% of interior point iterations on a range of problems varying from the small scale NETLIB test set to problems with over 192 million variables.

References

[1] H. Y. Benson and D. F. Shanno, *An exact primal-dual penalty method approach to warmstarting interior-point methods for linear programming*, tech. rep., Drexel University, 2005. Accepted for publication in Computational Optimization and Applications.

[2] ——, *Interior-point methods for nonconvex nonlinear programming: Regularization and warmstarts*, tech. rep., Drexel University, 2005. Accepted for publication in Computational Optimization and Applications.

[3] J. Fliege, *An efficient interior-point method for convex multicriteria optimization problems*, tech. rep., University of Birmingham, 2004. Accepted for publication in Mathematics of Operations Research.

[4] J. Gondzio, *HOPDM (version 2.12) – a fast LP solver based on a primal-dual interior point method*, European Journal of Operational Research, 85 (1995), pp. 221–225.

[5] ——, *Warm start of the primal-dual method applied in the cutting plane scheme*, Mathematical Programming, 83 (1998), pp. 125–143.

[6] J. Gondzio and A. Grothey, *Parallel interior point solver for structured quadratic programs: Application to financial planning problems*, Technical Report MS-03-001, School of Mathematics, University of Edinburgh, Edinburgh EH9 3JZ, Scotland, UK, April 2003. Accepted for publication in *Annals of Operations Research*.

[7] ——, *Reoptimization with the primal-dual interior point method*, SIAM Journal on Optimization, 13 (2003), pp. 842–864.

[8] J. Gondzio and J.-P. Vial, *Warm start and ε-subgradients in cutting plane scheme for block-angular linear programs*, Computational Optimization and Applications, 14 (1999), pp. 17–36.

[9] A. L. Hipolito, *A weighted least squares study of robustness in interior point linear programming*, Computational Optimization and Applications, 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*, www.optimization-online.net, (2006). Accepted for publication in Computational Optimization and Applications.

[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*, Mathematical Programming, 66 (1994), pp. 123–135.

[12] S. Mehrotra, *On the implementation of a primal–dual interior point method*, SIAM Journal on Optimization, 2 (1992), pp. 575–601.

[13] M. Steinbach, *Markowitz revisited: Mean variance models in financial portfolio analysis*, SIAM Review, 43 (2001), pp. 31–85.

[14] E. A. Yildirim and S. J. Wright, *Warm-start strategies in interior-point methods for linear programming*, SIAM Journal on Optimization, 12 (2002), pp. 782–810.

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

Table 7: Results (Best Warmstart) - perturbations in b

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
RECIPELP	-	-	-	-	-	-	-	-	-
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

Table 8: Results (Best Warmstart) - perturbations in c

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

Table 9: Results (Best Warmstart) - perturbations in A

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

Table 10: Results (Best Warmstart) - all perturbations

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

Table 11: Capacitated MCNF solved by warmstarted IPM-SQP

constraints	variables	$\rho =$	1e-3	5e-3	0.01	0.05	0.1	0.5	1	5	10
223.321	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
533.725	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
5.982.604	16.316.191	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
70.575.308	192.478.111	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

Table 12: Computation of Efficient Frontier with IPM warmstarts