Warm Start of the Primal-Dual Method Applied in the Cutting-Plane Scheme*†

Jacek Gondzio‡
Logilab, HEC Geneva, Section of Management Studies, University of Geneva,
102 Bd Carl Vogt, CH-1211 Geneva 4, Switzerland
e-mail: gondzio@divsun.unige.ch
URL: http://ecolu-info.unige.ch/~logilab/gondzio/

Logilab Technical Report 96.3
May 19, 1996, revised April 14, 1997

Abstract

A practical warm-start procedure is described for the infeasible primal-dual interior-point method employed to solve the restricted master problem within the cutting-plane method.

In contrast to the theoretical developments in this field, the approach presented in this paper does not make the unrealistic assumption that the new cuts are shallow. Moreover, it treats systematically the case when a large number of cuts are added at one time.

The technique proposed in this paper has been implemented in the context of HOPDM, the state of the art, yet public domain, interior-point code. Numerical results confirm a high degree of efficiency of this approach: regardless of the number of cuts added at one time (can be thousands in the largest examples) and regardless of the depth of the new cuts, reoptimizations are usually done with a few additional iterations.

Key words. Warm start, primal-dual algorithm, cutting-plane methods.

*Supported by the Fonds National de la Recherche Scientifique Suisse, grant #12-42503.94.

†to appear in *Mathematical Programming.*

‡On leave from the Systems Research Institute, Polish Academy of Sciences, Newelska 6, 01-447 Warsaw, Poland.

1 Introduction

We are concerned in this paper with using the infeasible primal-dual interior-point method to solve the restricted master problem in the cutting-plane scheme. We do not make any assumptions on the solution technique for the restricted master problem; we may seek its optimal solution (cf. Kelley [16]) or only an approximate solution with a predetermined distance to optimality (cf. Gondzio and Sarkissian [12]). We assume, however, that either the size of the master problem is very large (motivating the use of interior-point methods as solution techniques) or that the cutting-plane algorithm is able to exploit a "central" solution. Goffin, Haurie and Vial's [8, 7] analytic center variant is a scheme of the latter type.

We are concerned with the following restricted master problem

$$\begin{array}{ll} \text{maximize} & b^T y \\ \text{subject to} & A^T y \leq c, \end{array} \tag{1}$$

where $b, y \in \mathcal{R}^m, c \in \mathcal{R}^n$ and each column of the matrix $A \in \mathcal{R}^{m \times n}$ corresponds to a cut. We assume that this problem has been solved and its solution $\hat{y}$ has been used to generate a set of k new cuts

$$\bar{A}^T y \leq \bar{c},$$

where $\bar{c} \in \mathcal{R}^k$ and $\bar{A} \in \mathcal{R}^{m \times k}$. Some (possibly all) of these new constraints are deeply violated by the old solution $\hat{y}$, that is, $\bar{c} - \bar{A}^T \hat{y} \ll 0$. The new restricted master problem has the form

$$\begin{array}{ll} \text{maximize} & b^T y \\ \text{subject to} & A^T y \leq c, \\ & \bar{A}^T y \leq \bar{c}, \end{array} \tag{2}$$

and we would like to solve it rapidly by exploiting the solution of (1).

The general issue of using the solution of one linear program as an advanced starting point for another closely related linear program arises in many applications. For years the problem has been considered difficult for interior-point methods. We suggest that this view arose from too-literal interpretation of the hot-starting technique. Reoptimizations were understood literally as starting from the *optimal* solution of one problem to find an optimal solution of the next problem in the sequence.

The optimal solution of the preceding problem is usually at or near the boundary of the feasible set, making it a particularly difficult starting point for an interior-point algorithm. Research in this field concentrated on practical techniques attempting to move the point away from the boundary. Hipolito [14] used directions which were especially designed for this purpose. Lustig, Marsten and Shanno [17] added a considerable perturbation to the preceding solution to enable the primal-dual algorithm to avoid this undesirable behavior. To the best of our knowledge, these methods have not been applied in practice.

The warm start in the analytic center cutting-plane method [7] is understood in a different way since the restricted master problems (1) and (2) are not solved to optimality. Instead, we seek analytic centers of the localization sets (see [11] and the references therein). This approach strictly follows the theory [20]. Moreover, it works quite well in practice, except when a very

large number of cuts are added at one time, or when the localization sets are very small and the analytic centers stay too close to the boundary of the feasible region.

We conclude that the use of the preceding optimal solution as a starting point cannot be expected to work well in practice. We therefore define an alternative starting point that takes account of the shape of the feasible polytope and is close to optimality, yet is sufficiently far from the boundary of the feasible region.

We call this point a *close-to-optimality approximate analytic center* and show in subsequent discussions that it can be computed efficiently. The theory of interior-point methods indicates that relatively fast convergence of primal-dual methods can be expected when the iterates stay close to the central path [23]. Since we usually require 6 to 8 digits of accuracy, for an "optimal" solution, a reasonable definition of "close-to-optimal" would be 3 to 4 digits of accuracy. From such a point, we would expect to reach optimality in just a few iterations. The term "approximate analytic center" indicates that the point in question is nearly feasible and has well-balanced complementarity products (see Section 2.1). Points of this type are referred to extensively in the theory of infeasible primal-dual methods (see [23], Chapter 6).

Whenever we anticipate the need of performing a warm start to solve a closely related problem, we shall split the solution strategy into two stages. In the first stage, we iterate until an approximate analytic center corresponding to a nearly optimal point (3-digits accuracy) is found. After storing this point for the future warm start, we continue with the second stage, in which we proceed to the optimal solution of the original problem with the required accuracy. The first phase usually requires just one or two additional centering steps, so the overall cost of this strategy increases only marginally over the solution of a single isolated problem.

We want to draw the reader's attention to the fact that the saved point should be a good approximation of the analytic center. This is a necessary condition for the success of the warm start procedure. When the problem is only a minor perturbation of its near relative, this condition is also a sufficient one to ensure efficiency. When the perturbation is significant, however, we need a more sophisticated technique.

The latter situation occurs frequently in solving the restricted master problem in the the cutting-plane method. New cuts appended to the problem are almost always deeply violated by the proposed initial solution (which, as described above, is the approximate analytic center of the preceding restricted master problem). In this situation, we apply a procedure that recovers dual feasibility in these new constraints in just one step before the actual optimization starts. The direction for recovering dual feasibility is derived from the target following algorithm, the theory of which has been studied by Jansen, Roos, Terlaky and Vial [15]. We compute the target in such a way that the full step removes infeasibility from all new cuts. Unfortunately, there is no guarantee that a full step toward this target will be feasible. Often, feasibility is recovered only asymptotically, as is generally the case with infeasible primal-dual interior-point methods.

The paper is organized as follows. In Section 2 we recall some well known facts about the primal-dual path-following algorithm. As it is the most efficient interior-point method (IPM), it is very well studied in the literature. We shall thus concentrate only on the notion of the neighborhood of the central path and the notion of a target in the space of complementarity products, both used extensively in the following section. In Section 3 we give the fundamentals of our warm-start procedure and in Section 4 we address some issues important for its implementation. In Section 5 we give the computational results that demonstrate the performance of the proposed approach, and in Section 6 we give our conclusions.

Notation. We shall deal with two closely related problems in this paper: the old restricted master problem and the new one that differs from the precedent with a set of new columns (cuts). The variables corresponding to the old problem are denoted by x, s. The variables corresponding to new columns are denoted by *barred* vectors $\bar{x}, \bar{s}$ and the variables of the new master problem are denoted with *tilde*: $\tilde{x} = (x, \bar{x})$, $\tilde{s} = (s, \bar{s})$. The remaining vectors and matrices follow the same scheme.

2 Fundamentals of the primal-dual method

In this section we briefly recall a few basic ideas of the primal-dual algorithm that shall be extensively used in the following parts of this paper.

2.1 Central path neighborhood

The primal-dual path-following method for linear programming is considered to be the most efficient interior-point algorithm in practice. There exists a rich literature devoted to it; hence we omit its detailed discussion. The fundamental computational developments were done by Lustig, Marsten and Shanno [18]. The interested reader may also consult an up-to-date survey [1] devoted to the implementation of the primal-dual method and the book of Wright [23].

Consider a pair of linear programs: the primal

$$\begin{aligned} \text{minimize} \quad & c^T x \\ \text{subject to} \quad & Ax = b, \\ & x \geq 0, \end{aligned} \tag{3}$$

where $c, x \in \mathcal{R}^n, b \in \mathcal{R}^m$ and $A \in \mathcal{R}^{m \times n}$ and its dual

$$\begin{aligned} \text{maximize} \quad & b^T y \\ \text{subject to} \quad & A^T y + s = c, \\ & s \geq 0, \end{aligned} \tag{4}$$

where $y \in \mathcal{R}^m$ and $s \in \mathcal{R}^n$.

After replacing nonnegativity constraints in the primal problem with the logarithmic barrier function, we can write the following Lagrangean

$$L(x, s, \mu) = c^T x - \mu \sum_{j=1}^{n} \ln x_j - y^T (Ax - b), \tag{5}$$

where $\mu > 0$ is the barrier parameter. Next, we derive the first-order optimality conditions for it

$$\begin{aligned} Ax &= b, \\ A^T y + s &= c, \\ XSe_n &= \mu e_n, \end{aligned} \tag{6}$$

where X and S are diagonal matrices with the elements x_j and s_j, respectively and e_n is the n-vector of all ones. If the feasible sets of the primal and dual problems (3) and (4) have nonempty

interiors ($\mathcal{F}^0 = \{(x,y,s)|Ax = b,\ A^Ty + s = c,\ (x,s) > 0\} \neq \emptyset$), then for any $\mu > 0$, there exists a uniquely defined point $(x(\mu), y(\mu), s(\mu))$ $x(\mu) > 0, s(\mu) > 0$ that satisfies (6). Such a point is called a μ-center or an analytic center.

The infeasible primal-dual algorithm follows the path of centers, i.e., in subsequent iterations it makes (one) damped step in the Newton direction towards the solution of (6) and reduces the barrier parameter. The iterates of this algorithm stay in a large neighborhood of the central path (cf. [23], page 109). In our implementation of the primal-dual method, this neighborhood is defined by the following inequalities

$$\begin{array}{rcll} \|Ax - b\| & \leq & \epsilon_p(\mu)(||\xi_b^0|| + 1), & \\ \|A^Ty + s - c\| & \leq & \epsilon_d(\mu)(||\xi_c^0|| + 1), & \\ \beta\mu \leq & x_j s_j & \leq \frac{1}{\beta}\mu, & i = 1, 2, \ldots, n. \end{array} \tag{7}$$

The two vectors $\xi_b^0 = Ax^0 - b$ and $\xi_c^0 = A^Ty^0 + s^0 - c$ are the violations of the primal and dual constraints at the initial point (x^0, y^0, s^0), respectively. The relative feasibility tolerances $\epsilon_p(\mu)$ and $\epsilon_d(\mu)$ decrease with μ and reach zero at $\mu = 0$. The parameter $\beta \in (0,1)$ controls the discrepancy between the largest and the smallest complementarity products.

Let us observe that an exact μ-center (6) satisfies both primal and dual feasibility constraints and the parameter μ controls its distance to optimality. Really, the duality gap at this point is

$$c^Tx - b^Ty = c^Tx - x^TA^Ty = x^T(c - A^Ty) = x^Ts = n\mu. \tag{8}$$

In the usual applications, the algorithm terminates when the duality gap drops below a predetermined relative optimality tolerance ϵ, i.e., when

$$|c^Tx - b^Ty| \leq \epsilon(|c^Tx| + 1). \tag{9}$$

However, instead of using (9) as the stopping criterion, we can use (7). Here the algorithm stops at an approximate analytic center corresponding to the predetermined barrier parameter μ (thereby controlling the required distance to optimality). Once we obtain a rough estimate of the optimal objective value $f_0 = c^Tx$, e.g. when (9) is already satisfied for $\epsilon_0 = 10^{-1}$, we limit the decrease of the barrier to

$$\mu = \epsilon\frac{|f_o|}{n}. \tag{10}$$

From (7), (8) and (10) we see that an approximate analytic center corresponding to such a μ is an approximate optimal solution with the relative tolerance ϵ.

2.2 Targets in the space of complementarity products

When efficiency of the primal-dual method is the main concern, we relax the requirement of approaching the analytic center which corresponds to the current barrier parameter. In practice primal-dual methods make only one Newton step towards the μ-center and, regardless of how successful this step was, update the barrier. As a result, the barrier parameter is reduced which will push the algorithm closer to optimality. Consequently, the iterates of the algorithm are not exact analytic centers. They are rather rough approximations of such points. Both the theory [15] and the computational practice [9] take full advantage of this observation and allow the

choice of much more general *targets* rather than only the analytic centers. The requirement $Xs = \mu e_n$ in (6) is thus replaced with a more flexible condition

$$Xs = t, \tag{11}$$

where $t \in \mathcal{R}^n$ is some target in the space of complementarity products.

A wise choice of the targets, hopefully one that can be more easily reached than the perfectly centered (but often unreachable) analytic centers, led to a very efficient implementation of the primal-dual algorithm [9]. For algorithms which intend to reach a good approximation of the analytic center with a predetermined proximity ϵ to the optimal solution (cf. equation (10)), the technique of multiple centrality correctors, which is described in the above paper, has proven to be particularly useful.

3 Warm starts in the cutting-plane method

In this section we shall present our warm-start procedure. We begin by describing an analytic-center formulation. Next, we discuss the specific properties of the reoptimization problem and describe in detail the new approach proposed to solve it.

3.1 Problem formulation

Suppose that at a given iteration of the cutting-plane method, the relaxed master problem (1) (now formulated with the explicit dual slacks)

$$\begin{array}{ll} \text{maximize} & b^T y \\ \text{subject to} & A^T y + s = c, \\ & s \geq 0, \end{array} \tag{12}$$

has been solved. Regardless of the type of "solution" that may be looked for in the method (optimal in Kelley's approach [16], analytic center of the localization set in Goffin, Haurie and Vial's approach [7], or the approximate μ-center in a simple scheme proposed in [12]) we shall identify two steps in the solution process. In the first one, we shall look for the approximate μ-center corresponding to the moderate accuracy ϵ, say, $\epsilon = 10^{-2}$ or $\epsilon = 10^{-3}$ (cf. equation (10)). Although in practice, we always deal with an approximation (7), in order to simplify the presentation, we assume that an exact μ-center has been found, i.e., the point (x_0, y_0, s_0) with $x_0 > 0$ and $s_0 > 0$ is known such that

$$\begin{aligned} Ax_0 &= b, \\ A^T y_0 + s_0 &= c, \\ X_0 s_0 &= \mu e_n. \end{aligned} \tag{13}$$

Having stored this point, we continue the optimization until the required "solution" of the relaxed master problem is obtained. The resulting proposition is then supplied to the Oracle that generates a set of new cuts

$$\bar{A}^T y \leq \bar{c},$$

which completes the iteration of the cutting-plane method. Computational experience shows that in nontrivial applications of the cutting-plane method, these new cuts are seriously violated by the earlier saved point (x_0, y_0, s_0), i.e.,

$$z = \bar{c} - \bar{A}^T y_0 < 0.$$

The following iteration requires a solution of the new relaxed master problem

$$\begin{array}{rl} \text{maximize} & b^T y \\ \text{subject to} & A^T y + s = c, \\ & \bar{A}^T y + \bar{s} = \bar{c}, \\ & s \geq 0,\ \bar{s} \geq 0. \end{array} \tag{14}$$

As before, we split the solution process into two steps. In the first one, we look for an approximation (7) of the new $\bar{\mu}$-center:

$$\begin{array}{rcl} Ax + \bar{A}\bar{x} & = & b, \\ A^T y + s & = & c, \\ \bar{A}^T y + \bar{s} & = & \bar{c}, \\ Xs & = & \bar{\mu} e_n, \\ \bar{X}\bar{s} & = & \bar{\mu} e_k. \end{array} \tag{15}$$

The *warm-start* problem consists of the exploitation of the solution (x_0, y_0, s_0) to efficiently find an approximate solution of (15).

3.2 Initial point

Let us observe that y_0 from (13) can seriously violate the new cuts. This is usually the case in practice as the new cuts are supposed to shrink the feasible region and, in particular, to cut off the earlier solution. Moreover, we cannot predict how large the violation will be, and any assumption that it will be small is unjustified.

Consequently, we have to accept that the new problem is a significant perturbation of the earlier one. We shall use the following starting point $(\tilde{x}, \tilde{y}, \tilde{s}) = (x_0, \bar{x}, y_0, s_0, \bar{s})$, where $\tilde{x}, \tilde{s} \in \mathcal{R}^{n+k}$ and, for consistency of notation, $\tilde{A} = [A,\ \bar{A}]$. Such a point violates both the primal constraint

$$\begin{array}{rcl} \xi_b & = & b - \tilde{A}\tilde{x} \\ & = & b - Ax_0 - \bar{A}\bar{x} = -\bar{A}\bar{x}, \end{array} \tag{16}$$

and the dual one

$$\begin{array}{rcl} \xi_{\tilde{c}} & = & \tilde{c} - \tilde{A}^T y - \tilde{s} \\ & = & \left[\begin{array}{c} c - A^T y - s \\ \bar{c} - \bar{A}^T y - \bar{s} \end{array} \right] = \left[\begin{array}{c} 0 \\ z - \bar{s} \end{array} \right] \leq \left[\begin{array}{c} 0 \\ z \end{array} \right], \end{array} \tag{17}$$

but the choice of

$$\bar{x} = \mu \bar{S}^{-1} e_k, \tag{18}$$

yields (for an arbitrary $\bar{s}$) that all complementarity products are equal to the former barrier

$$Xs = X_0 s_0 = \mu e_n \quad \text{and} \quad \bar{X}\bar{s} = \mu e_k. \tag{19}$$

The choice of $\bar{x}$ and $\bar{s}$ has important consequences for computational practice. Below we describe how we choose these variables in our implementation. First, we determine the depth of the new cut. We use a relative measure based on the characteristics of the saved approximate analytic center (x_0, y_0, s_0) and the corresponding barrier parameter μ. We consider that a new cut j is *deep* if

$$\bar{a}_j^T y_0 - \bar{c}_j = z_j < -\mu^{1/2}, \tag{20}$$

and that it is *shallow* otherwise. This is in contrast to most of the theoretical developments in the field that treat any cut for which $z_j < 0$ as a deep one. Our definition of deep cut corresponds to the computational experience that shows clearly that a small infeasibility of a new cut ($-\mu^{1/2} \leq z_j < 0$) can be absorbed easily (usually in merely one step) by the infeasible primal-dual algorithm.

Once the depths of the new cuts have been determined, we now choose $\bar{x}$ and $\bar{s}$. The dual slack variables are set to

$$\bar{s}_j = \max(\mu^{1/2}, |z_j|). \tag{21}$$

Hence, for a deep cut, $\bar{s}_j = -z_j > \mu^{1/2}$ while for a shallow cut, $\bar{s}_j = z_j$ (unless $|z_j|$ is too small, i.e., $|z_j| < \mu^{1/2}$, in which case $\bar{s}_j$ is set to the threshold value of $\mu^{1/2}$).

Finally, the primal variables $\bar{x}_j$ are computed from (18). Note that the primal variables defined in this way cannot exceed $\mu^{1/2}$. In the following section we shall see the advantages of keeping the primal variables $\bar{x}$ small.

Let us observe that the violations of the primal and the dual constraints are different in nature (cf. (16) and (17)). We can keep ξ_b small by the choice of small $\bar{x}$ but we have little influence on the violation of the dual constraint determined by the depth of cuts returned by the Oracle. Consequently, we find it necessary to treat the two problems of restoring primal and dual feasibility independently.

In the following two sections we discuss the approach that we have applied. We aim at restoring feasibility in one step but we are aware that it does not always have to be possible, especially when the violation of dual constraints, z is large. Hence the first step is supposed to reduce both infeasibilities by a considerable factor. If the new iterate after this step still violates the primal or dual constraints, then the remaining infeasibility will be removed together with the progress towards optimization (as is the usual case in the infeasible primal-dual algorithm).

3.3 Restoring primal feasibility

Let us observe that primal feasibility would be ensured if we could choose $\bar{x} = 0$. However, due to the use of the logarithmic barrier, we need a strictly positive solution $\bar{x} > 0$. Even a very small $\bar{x}$ would satisfy our needs. Following [7], let us generalize the direction of Mitchell [21] to handle multiple cuts

$$\Delta x = -\Theta A^T (A\Theta A^T)^{-1} \bar{A}\bar{x} \in \mathcal{R}^n, \tag{22}$$

where $\Theta = XS^{-1}$ is a diagonal primal-dual scaling matrix.

It is easy to verify that the point $(x + \alpha\Delta x, \alpha\bar{x}) \in \mathcal{R}^{n+k}$ satisfies the primal feasibility constraints. As the old solution $x = x_0$ was strictly positive, there exists a strictly positive α such that $x + \alpha\Delta x > 0$. We may now do a ratio test to determine the maximum feasible stepsize in the primal space

$$\alpha_{max} = \min_{j:\Delta x_j < 0} \{-x_j / \Delta x_j\},$$

and choose any value $\alpha \in (0, \alpha_{max})$. It is useful, for example, to choose an α that minimizes the primal potential of the point $(x + \alpha\Delta x, \alpha\bar{x})$. Such a choice was applied with great success in the implementation [11, 20] of the analytic center cutting-plane method [7]. As this method applies a primal potential reduction interior-point algorithm to find an approximate analytic center, it does not really have to bother about restoring dual feasibility.

In the method presented in this paper, we do not use any potential function. We let the infeasible primal-dual method take classical steps in Δx with the stepsize $0.9995 \times \alpha_{max}$.

3.4 Restoring dual feasibility

Restoring dual feasibility is a significantly more difficult problem. In particular, it is not realistic to believe that we can do it and at the same time keep the centrality condition satisfied with the previous barrier parameter. Would this be the case, we would have a new approximate analytic center with a very small duality gap

$$x^T s + \bar{x}^T \bar{s} = \mu(n + k). \tag{23}$$

Thus we are ready to accept an increase of the complementarity products (the loss in the "optimality") so as to give more chance to retrieve dual feasibility. An argument for this is straightforward: the old dual solution y did not take into account the new cuts. It satisfied the old dual constraints but violated the new ones. We would like to remove this inconsistency as much as possible, i.e., to adjust y so as it is feasible for all constraints.

We accept the violation of the centrality condition as a price for that. Consequently, we shall use a target in the space of complementarity products such that a step towards this target restores feasibility in the new dual constraints. Such a target is not known in advance, but we can determine it with an acceptable amount of computational effort.

Assume we choose the following target

$$\tilde{t} = \begin{bmatrix} Xs \\ \bar{t} \end{bmatrix} \in \mathcal{R}^{n+k}, \tag{24}$$

i.e., we want to keep the old complementarity products unchanged, and we allow the increase of complementarity products corresponding to the new variables $\bar{x}$ and $\bar{s}$. We shall determine the Newton direction that takes the current iterate to the target. As we shall only use the dual part of this direction, we ignore the violation of the primal constraint. The Newton equation system is then written as

$$\begin{bmatrix} \tilde{A} & & \\ & \tilde{A}^T & I \\ \tilde{S} & & \tilde{X} \end{bmatrix} \begin{bmatrix} \Delta\tilde{x} \\ \Delta y \\ \Delta\tilde{s} \end{bmatrix} = \begin{bmatrix} 0 \\ \xi_{\tilde{c}} \\ \tilde{t} - \tilde{X}\tilde{s}. \end{bmatrix}. \tag{25}$$

After the elimination of

$$\Delta\tilde{s} = \tilde{X}^{-1}\tilde{t} - \tilde{s} - \tilde{X}^{-1}\tilde{S}\Delta\tilde{x}$$

and the application of (17) and (24) we obtain

$$\begin{bmatrix} -\Theta^{-1} & & A^T \\ & -\bar{\Theta}^{-1} & \bar{A}^T \\ A & \bar{A} & \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta\bar{x} \\ \Delta y \end{bmatrix} = \begin{bmatrix} 0 \\ z - \bar{X}^{-1}\bar{t} \\ 0 \end{bmatrix}, \tag{26}$$

where $\tilde{\Theta} = [\Theta,\ \bar{\Theta}] = [XS^{-1},\ \bar{X}\bar{S}^{-1}]$ is a diagonal primal-dual scaling matrix. We reduce this system further: after an elimination of

$$\Delta x = \Theta A^T \Delta y \quad \text{and} \quad \Delta\bar{x} = \bar{\Theta}\bar{A}^T\Delta y - \Theta(z - \bar{X}^{-1}\bar{t})$$

we obtain the normal-equations system

$$(A\Theta A^T + \bar{A}\bar{\Theta}\bar{A}^T)\Delta y = \bar{A}\bar{\Theta}(z - \bar{X}^{-1}\bar{t}). \tag{27}$$

In the following part of the analysis we shall look for such a target $\bar{t}$ so that a full step generates feasibility in the dual space. If we could find such a direction, we would remove dual infeasibility in one step. Unfortunately, in the general case, there is no guarantee of finding such a direction. However, we can relax our requirement, i.e., replace it with the following one: find a target $\bar{t}$ such that at least in the new components of the dual slack variables $\bar{s}$, a full step is feasible. We then require that

$$\bar{s} + \Delta\bar{s} = 0.$$

Adding this condition to (25) leads to the following linear system

$$\begin{bmatrix} A & \bar{A} & & & & \\ & & A^T & I & & \\ & & \bar{A}^T & & I & \\ S & & & X & & \\ & \bar{S} & & & \bar{X} & -I \\ & & & & I & \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta\bar{x} \\ \Delta y \\ \Delta s \\ \Delta\bar{s} \\ \bar{t} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ z - \bar{s} \\ 0 \\ \bar{X}\bar{s} \\ -\bar{s} \end{bmatrix}. \tag{28}$$

It is a matter of linear algebra to solve (28) for $\bar{t}$. We do these computations below. The uninterested reader may skip them and pass immediately to the solution (31).

From (25) and (17) we get

$$\bar{A}^T\Delta y + \Delta\bar{s} = \xi_{\bar{c}} = z - \bar{s},$$

which gives

$$\bar{A}^T\Delta y = z.$$

Applying (27) we obtain the equation

$$\bar{A}^T(A\Theta A^T + \bar{A}\bar{\Theta}\bar{A}^T)^{-1}\bar{A}\bar{\Theta}(z - \bar{X}^{-1}\bar{t}) = z$$

that determines the target $\bar{t}$. Under the realistic conditions that the matrix $A\Theta A^T$ is nonsingular (this is the normal-equations matrix corresponding to the approximate analytic center found in the earlier reduced master problem) and that matrix $\bar{A}$ has rank k, this system can be solved for $\bar{t}$.

Let us transform it to

$$\begin{aligned}\bar{\Theta}^{1/2}\bar{A}^T(A\Theta A^T + \bar{A}\bar{\Theta}\bar{A}^T)^{-1}\bar{A}\bar{\Theta}^{1/2}\bar{\Theta}^{1/2}\bar{X}^{-1}\bar{t} = \\ -(I_k - \bar{\Theta}^{1/2}\bar{A}^T(A\Theta A^T + \bar{A}\bar{\Theta}\bar{A}^T)^{-1}\bar{A}\bar{\Theta}^{1/2})\bar{\Theta}^{1/2}z \qquad (29)\end{aligned}$$

and denote

$$LL^T = A\Theta A^T \quad \text{and} \quad V = L^{-1}\bar{A}\bar{\Theta}^{1/2}.$$

Applying the Sherman-Morrison-Woodbury formula

$$(I_k + V^TV)^{-1} = I_k - V^T(I_m + VV^T)^{-1}V,$$

from (29) we obtain

$$\begin{aligned}\bar{\Theta}^{1/2}\bar{X}^{-1}\bar{t} &= -(V^T(I_m + VV^T)^{-1}V)^{-1}(I_k - V^T(I_m + VV^T)^{-1}V)\bar{\Theta}^{1/2}z \\ &= -(V^T(I_m + VV^T)^{-1}V)^{-1}(I_k + V^TV)^{-1}\bar{\Theta}^{1/2}z \\ &= -((I_k + V^TV)(V^T(I_m + VV^T)^{-1}V))^{-1}\bar{\Theta}^{1/2}z \\ &= -(V^T(I_m + VV^T)(I_m + VV^T)^{-1}V)^{-1}\bar{\Theta}^{1/2}z \\ &= -(V^TV)^{-1}\bar{\Theta}^{1/2}z. \qquad (30)\end{aligned}$$

As

$$V^TV = \bar{\Theta}^{1/2}\bar{A}^T(A\Theta A^T)^{-1}\bar{A}\bar{\Theta}^{1/2},$$

we finally get

$$\begin{aligned}\bar{t} &= -\bar{X}\bar{\Theta}^{-1/2}(V^TV)^{-1}\bar{\Theta}^{1/2}z \\ &= -\bar{X}\bar{\Theta}^{-1}(\bar{A}^T(A\Theta A^T)^{-1}\bar{A})^{-1}z \\ &= -\bar{S}(\bar{A}^T(A\Theta A^T)^{-1}\bar{A})^{-1}z. \qquad (31)\end{aligned}$$

Summing up, for any depth of the new cuts z and for an arbitrary value of the dual slacks $\bar{s}$, we can determine a target $\bar{t}$ such that a full step in the direction $(\Delta y, \Delta\bar{s})$ in the dual space restores the feasibility of the new cuts.

The computation of such a target requires an inversion of the matrix

$$B = \bar{A}^T(A\Theta A^T)^{-1}\bar{A},$$

which may become a nontrivial computational task, especially if the number of cuts added at one time is large. In the following section we will show, however, how to compute an approximate solution of an equation with B in an inexpensive way, regardless of the size of k.

Given the definition of the target (31) we compute from (27) the direction in the dual space. Next, after simple (though long) calculations, we obtain

$$\Delta y = (A\Theta A^T)^{-1}\bar{A}(\bar{A}^T(A\Theta A^T)^{-1}\bar{A})^{-1}z. \qquad (32)$$

We easily verify that a full step in this direction absorbs total infeasibility in the new cuts ($\bar{A}^T \Delta y = z$). Moreover, Δy does not depend on the choice of $\bar{x}$ and $\bar{s}$.

Unfortunately, we have no guarantee of being able to make a full step in the dual space since the step in

$$\begin{aligned} \Delta s &= -A^T \Delta y \\ &= -A^T (A\Theta A^T)^{-1} \bar{A} (\bar{A}^T (A\Theta A^T)^{-1} \bar{A})^{-1} z \end{aligned}$$

may be blocked for variables $j \leq n$ referring to the old components of s.

Let us observe that the choice of the approximate analytic center of the previous relaxed master problem (x_0, y_0, s_0) as the initial point for reoptimization was the best of what we could have done. When this point was chosen, we did not yet know the new cuts and their depth z, hence we could not predict direction Δs. The choice of the analytic center was the safest then: if an exact analytic center exists, it is the center of an ellipsoid inscribed in the dual feasible region. Hence, there is a guarantee to be able to make a small step from this point in an arbitrary direction, see, e.g. Vial [22].

Unfortunately, the aforementioned ellipsoid may be very small in practice. This is the reason why we have decided to store an approximate analytic center corresponding to only 2 or 3 digits of accuracy. In this way we keep the ellipsoid reasonably large and increase the chance of absorbing deeper cuts although, clearly, we have no guarantee to be able to do it for arbitrarily large z.

4 Implementation issues

The approach presented in this paper has been implemented and applied in the primal-dual cutting-plane method proposed by Gondzio and Sarkissian [12]. The implementation was based on the public domain LP code HOPDM (cf. [9, 10]). The sources of the HOPDM code can be retrieved from the Web site `http://ecolu-info.unige.ch/~logilab/software/`. In this section we discuss several issues that seem crucial for the success of our warm-start routine.

4.1 Efficient centering with the primal-dual algorithm

One of the important features of our warm-start procedure is the use of a good approximation of the analytic center of one problem as the initial point in the solution of the following problem. It is then essential to be able to find such a point.

The theory of the path-following algorithms [13, 22, 23] requires that the iterates stay in a neighborhood of the central path. The general purpose primal-dual implementations [9, 18, 19] allow this neighborhood to be quite large. Keeping all complementarity products $x_j s_j, j = 1, 2, ..., n$ close to the current μ is a desirable but not a necessary condition to optimize fast. The computational experience of Gondzio [9] confirmed that additional effort to reduce the discrepancy between the largest and the smallest complementarity products results in a significant reduction of the number of iterations to reach optimality. The modification of the infeasible primal-dual algorithm presented in [9] perfectly suits the search of the approximate analytic center defined by conditions (7). We recall some of the basic ideas.

The primal-dual algorithms make steps in the Newton direction for the first-order optimality condition, i.e., they solve the system

$$\begin{bmatrix} A & & \\ & A^T & I \\ S & & X \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta s \end{bmatrix} = \begin{bmatrix} b - Ax \\ c - A^T y - s \\ \mu e_n - Xs \end{bmatrix}. \tag{33}$$

Such a step is supposed to reduce both primal and dual infeasibilities and place the next iterate close to the analytic center. In fact, after a full step in the Newton direction is made, the complementarity becomes

$$(X + \Delta X)(s + \Delta s) = Xs + (X\Delta s + S\Delta x) + \Delta X \Delta s = \mu e_n + \Delta X \Delta s.$$

Hence, as $\Delta X \Delta s$ is large, the new iterate may seriously violate the centrality condition. The predictor-corrector technique [19, 18] tries to compensate for the error caused by the second order term.

The centrality correctors of Gondzio [9] are less aggressive. Their effort concentrates on increasing the feasible stepsizes in the primal and the dual spaces from α_P and α_D to $\tilde{\alpha}_P = \min(\alpha_P + \delta_\alpha, 1)$ and $\tilde{\alpha}_D = \min(\alpha_D + \delta_\alpha, 1)$, respectively, and taking all complementarity pairs $x_j s_j$ to the vicinity of the central path

$$\beta_{min}\mu \leq x_j s_j \leq \beta_{max}\mu, \qquad j = 1, 2, ..., n, \tag{34}$$

where $\delta_\alpha = 0.1$, $\beta_{min} = 0.1$ and $\beta_{max} = 10$ (cf. [9] for details). Such correctors allow the algorithm to reach a well centered solution with the small cost of one or two additional pure re-centering iterations at the end of optimization. An important practical question to be addressed is the choice of stopping criteria that measures the quality of the approximation of the analytic center. We started the experiments with a relatively weak condition (7) in which we put $\beta = 0.1$ like in (34). Its advantage is that an approximate analytic center is usually reached after only one additional recentering iteration. However, the quality of this point was sometimes insufficient to make the following warm start efficient. Thus we increased the requirements imposed on the point saved for future warm start. In our implementation this approximate analytic center satisfies conditions (7) with $\beta = 0.5$ and $\epsilon_p = \epsilon_d = 10^{-3} \times \epsilon$, where ϵ is its relative optimality tolerance.

Another important conclusion that follows from computational practice is that the approximate analytic center is reached faster if the primal-dual method does not apply Mehrotra's correctors but if it works solely with the centrality correctors [9].

4.2 Computation of the target

In Section 3.2 we derived formula (31) for a target that should be used to compute a direction that restores dual feasibility of the new cuts. The computation of this target requires the solution of the following system of equations

$$\bar{A}^T (A\Theta A^T)^{-1} \bar{A}\, q = z. \tag{35}$$

Recall that matrix $A\Theta A^T \in \mathcal{R}^{m \times m}$ is the normal-equations matrix corresponding to the approximate analytic center (x_0, y_0, s_0) used in our warm-start procedure, and $\bar{A} \in \mathcal{R}^{m \times k}$ is the matrix built of new cuts. Matrix $A\Theta A^T$ is nonsingular; hence, its Cholesky factorization exists

$$LL^T = A\Theta A^T.$$

(We could make an additional assumption that this factorization was stored so that it could be used in the warm-start procedure.)

Two different cases can appear in computational practice. If the number of cuts is small (note that there exist applications of the cutting-plane method in which only one cut is added at every outer iteration), then the cost of forming matrix $\bar{A}^T(A\Theta A^T)^{-1}\bar{A}$ remains acceptable. We then write

$$\bar{A}^T(A\Theta A^T)^{-1}\bar{A} = F^T F,$$

where the computation of the matrix

$$F = L^{-1}\bar{A} \in \mathcal{R}^{m\times k}$$

requires k backsolves with the Cholesky factor (one backsolve per column of $\bar{A}$). Once the matrix $F^T F$ is computed, its Cholesky or QR factorization is determined and is applied to solve equation (35).

Unfortunately, this approach becomes prohibitively expensive in the second case, when k gets larger. Below we propose a technique for computing an approximate solution to (35). Its computational cost is comparable with that of one IPM iteration. Instead of the explicit formulation of matrix $F^T F$, we solve the following system

$$\begin{bmatrix} -\Theta^{-1} & & A^T \\ & -D^{-1} & \bar{A}^T \\ A & \bar{A} & \end{bmatrix} \begin{bmatrix} r \\ \tilde{q} \\ p \end{bmatrix} = \begin{bmatrix} 0 \\ -z \\ 0 \end{bmatrix}, \tag{36}$$

for a special diagonal matrix D. After the elimination, $r = \Theta A^T p$, this system reduces to

$$\begin{bmatrix} -D^{-1} & \bar{A}^T \\ \bar{A} & A\Theta A^T \end{bmatrix} \begin{bmatrix} \tilde{q} \\ p \end{bmatrix} = \begin{bmatrix} -z \\ 0 \end{bmatrix}. \tag{37}$$

We may continue the pivoting with the diagonal term $-D^{-1}$, i.e., eliminate

$$\tilde{q} = D\bar{A}^T p + Dz, \tag{38}$$

and obtain

$$(A\Theta A^T + \bar{A}D\bar{A}^T)p = -\bar{A}Dz, \tag{39}$$

that is the well known normal-equations system corresponding to the new restricted master problem and a special diagonal scaling matrix $\tilde{\Theta} = [\Theta,\ D] \in \mathcal{R}^{(n+k)\times(n+k)}$. This is the way in which system (37) is solved in our implementation.

An alternative way to solve it is to eliminate

$$p = -(A\Theta A^T)^{-1}\bar{A}\tilde{q},$$

and write

$$(\bar{A}^T(A\Theta A^T)^{-1}\bar{A} + D^{-1})\ \tilde{q} = z. \tag{40}$$

It is easy to observe that by the choice of sufficiently large D, we can obtain a good approximation $\tilde{q}$ of q that solves (35).

In our implementation we set

$$D = dI_k, \quad \text{where} \quad d = 100 \times \max_{j=1,2,\ldots,n} \{\Theta_j\}, \tag{41}$$

and solve (37) for p using the normal-equations system (39) and compute $\tilde{q}$ from (38). Any IPM code that employs the augmented system solver [2] to compute Newton direction would of course solve (36) directly.

5 Numerical results

The warm-start approach proposed in this paper has been tested in the context of the primal-dual cutting-plane method of Gondzio and Sarkissian [12]. The reader interested in more detail in this particular cutting-plane method is referred to the aforementioned paper.

Table 1. Problems statistics.

Problem	Network characteristics			Subprobs	m
	Nodes	Arcs	Commodities		
NDO22	14	22	23	45	67
NDO148	61	148	122	270	418
Random12	300	600	1000	1600	2200
Random16	300	1000	1000	2000	3000
Random20	400	1000	5000	6000	7000
Random31	700	2000	2000	4000	6000
Random32	1000	3000	2000	5000	8000
Random33	1000	4000	2000	6000	10000
Random41	1000	5000	2000	7000	12000
Random42	1000	5000	3000	8000	13000
Random51	1000	2000	6000	8000	10000
Random52	1000	2000	8000	10000	12000
Random53	1000	2000	10000	12000	14000
Random54	1000	3000	6000	9000	12000
Random55	1000	3000	8000	11000	14000

We applied the method to solve large-scale nonlinear multicommodity network flow problems formulated as in [6]. An important feature of the problems is that they require the addition of a large number of deep cuts. As a result, they form really difficult test examples for the approach proposed in this paper.

Our test examples are characterized in Table 1. There are 2 well known problems NDO22 and NDO148 from [5] and 13 randomly generated problems. A description of their (public domain) generator can be found in [6]. Table 1 gives network statistics: the number of nodes, arcs and commodities, respectively, and reports the number of subproblems to facilitate the determination of the dimension, m, of the restricted master problem. In the formulation of nonlinear multicommodity network flow problems with disaggregated cuts [6], the number of coupling constraints equals the number of arcs while the number of subproblems equals to the sum of numbers of arcs and commodities. Thus the dimension of the dual space

$$m = 2 \times n_{arcs} + n_{commodities}. \tag{42}$$

Table 2. Overall efficiency of the primal-dual cutting-plane method.

Problem	Cuts			Outer	Inner Iterations		
	Total	Deep	Final		AC	Opt	Total
NDO22	410	181	359	17	63	36	99
NDO148	3870	1813	1940	16	57	37	94
Random12	16211	5612	10941	18	97	62	159
Random16	23974	6485	15276	21	117	74	191
Random20	54391	12849	28420	19	129	107	236
Random31	51137	13214	28543	20	128	85	213
Random32	68081	21093	40797	21	125	71	196
Random33	91783	28371	60356	25	153	89	242
Random41	111362	33402	70089	27	171	94	265
Random42	99677	21399	62017	21	139	75	214
Random51	75571	23367	40930	18	135	95	230
Random52	98445	36101	51006	19	146	112	258
Random53	108631	38457	58301	18	146	124	270
Random54	101311	31612	59051	19	148	101	249
Random55	128620	43529	66149	22	166	140	306

In Table 2 we report the overall solution statistics for all problems solved. The second, third and fourth columns give the number of cuts added throughout the solution process: their total number, the number of deep cuts and the final number of cuts corresponding to the last restricted master problem solved, respectively. Recall that by deep cut we mean a constraint that is strongly violated by the precedent approximate analytic center, i.e., a constraint for which

$$\bar{a}_j^T y_0 - \bar{c}_j = z_j < -\mu^{1/2},$$

where $\bar{a}_j$ is a given column of matrix $\bar{A}$, $j \in \{1, 2, ..., k\}$ and μ is a barrier parameter corresponding to the precedent point. Let us observe that the number of cuts present in the last restricted master problem is always smaller than the total number of cuts introduced during the optimization process. All cuts mentioned in the second column of the table are appended to the restricted master problem but some of them are eliminated in later iterations because they are inactive (see [12] for details).

The remaining columns of Table 2 report the number of iterations of the cutting-plane method (Outer) and the number of interior-point iterations (Inner). In the latter we have distinguished the iterations needed to reach the approximate analytic center (to be saved for the future warm start), AC, the iterations to reach the desired accuracy of solution to the restricted master problem, Opt, and their sum, respectively.

From the results collected in Table 2 one can see that we really deal with nontrivial warm-start examples. The sizes of the restricted master problems always reach tens of thousand columns with, on the average, thousands of new cuts to be accommodated at every reoptimization.

The results collected in Table 2 confirm the overall good performance of the primal-dual analytic center cutting-plane method, but they do not give much insight into the behavior of the warm-start procedure proposed in this paper. Such insight is given by the results reported in Table 3. Its columns contain the name of the problem, the average number of interior-

point iterations to reach an approximate analytic center (AC), the average number of additional iterations necessary to obtain the required solution in the restricted master problem (we shall simply call them "optimization" iterations — Opt). The following columns of Table 3 report performance details for the last reoptimizations done during the solution of each problem. They give initial size of the last restricted master problem, m and n, the number of new cuts added to it, the number of deep cuts, and the iteration numbers to reach a 4-digit approximation of the analytic center and the final, 7-digit exact solution, respectively.

Table 3. Efficiency of warm-start procedure.

Problem	Inner		Last outer iteration					
	AC	Opt	m	n	Cuts	Deep	AC	Opt
NDO22	4	2	67	381	22	22	3	4
NDO148	4	2	418	2088	148	148	4	5
Random12	5	4	2200	11652	489	481	4	5
Random16	6	3	3000	16599	677	668	5	4
Random20	7	5	7000	29712	708	701	3	5
Random31	6	5	6000	31077	1466	1357	3	4
Random32	6	3	8000	44640	2157	1993	3	3
Random33	6	4	10000	65830	2526	2417	3	4
Random41	6	4	12000	77559	2530	2489	4	5
Random42	7	3	13000	69311	2706	2345	4	3
Random51	8	5	10000	43389	1540	1412	4	6
Random52	8	6	12000	53472	1534	1472	3	5
Random53	8	7	14000	60803	1498	1493	4	5
Random54	8	5	12000	62975	2076	2027	5	5
Random55	8	6	14000	70108	2041	1997	4	5

Let us analyze these results in detail. Certainly, the average number of iterations necessary to compute the approximate analytic center in a modified restricted master problem is quite low, which is a remarkable result. It confirms that the use of *close-to-optimality* approximate analytic center as an initial solution for reoptimization has great advantages.

The average number of iterations needed to reoptimize the restricted master problems (AC + Opt) varies between 6 and 15; hence, it is significantly smaller than the expected number of iterations if one would like to solve these problems from cold start. (Note that the number of columns in subsequent restricted master problems grows up to almost 100,000 so the primal-dual method is supposed to make at least 40-60 iterations to solve them from scratch.)

Our final remark concerns the very good behavior of the approach proposed in this paper in the last iterations of the cutting-plane method. The warm start procedure we propose reveals exceptional robustness near the end of optimization.

6 Conclusions

We have presented a practical approach for performing reoptimizations with the infeasible primal-dual interior-point algorithm employed in the cutting-plane method. The method addresses the needs which arise in computational practice since it systematically treats the case when a large number of deep cuts are added to the problem at one time.

Computational results show that it works well including situations when the number of cuts

added at each time reaches into the thousands.

The following factors seem crucial for efficient warm starting with the primal-dual algorithm:

1. The starting point should not be the optimal solution. Instead, it should be an approximate analytic center sufficiently close to optimality (so as to be able to reach solution fast) but also sufficiently far from optimality so as to be able to accommodate significant perturbations in the subsequent problem.

2. The last steps of the algorithm applied to find an approximate analytic center should not employ Mehrotra's predictor-corrector technique [19], but they should work solely with the centrality correctors of Gondzio [9].

3. Restoring primal and dual feasibility should be treated separately. If feasibility is to be achieved in a single step, one has to accept a serious violation of the centrality condition and use the target (31).

Acknowledgement
The use of an approximate analytic center as a promising initial point to reoptimize grew out of the author's discussion with Professor Jean-Philippe Vial, who should share credit for the idea. The author is grateful to Doctor Joe Czyzyk for careful reading of this paper and to the anonymous referees and the associate editor for their constructive comments and suggestions.

References

[1] E. Andersen, J. Gondzio, C. Meszaros and X. Xu, "Implementation of Interior Point Methods for Large Scale Linear Programming", in T. Terlaky, ed., *Interior Point Methods in Mathematical Programming* (Kluwer Academic Publisher, 1996) 189–252.

[2] M. Arioli, I. S. Duff, and P. P. M. de Rijk. "On the Augmented System Approach to Sparse Least–Squares Problems", *Numerische Mathematik* 55 (1989) 667–684.

[3] J.F. Benders, "Partitioning Procedures for Solving Mixed-variables Programming Problems", *Numerische Mathematik* 4 (1962) 238–252.

[4] G.B. Dantzig and P. Wolfe, "The Decomposition Algorithm for Linear Programming", *Econometrica* 29 (4) (1961) 767–778.

[5] J.-L. Goffin, "The Ellipsoid Method and Its Predecessors", in: L. Contesse, R. Correa and A. Weintraub, eds., *Recent Advances in System Modelling and Optimization*, Lecture Notes in Control and Information Sciences 87, (Springer-Verlag, Berlin, 1987) 127–141.

[6] J.-L. Goffin, J. Gondzio, R. Sarkissian and J.-P. Vial, "Solving Nonlinear Multicommodity Flow Problems by the Analytic Center Cutting Plane Method", *Mathematical Programming* 79 (1997) 131–154.

[7] J.-L. Goffin, A. Haurie and J.-P. Vial, "Decomposition and nondifferentiable optimization with the projective algorithm", *Management Science* 38 (2) (1992) 284–302.

[8] J.-L. Goffin and J.-P. Vial, "Cutting Planes and Column Generation Techniques with the Projective Algorithm", *Journal of Optimization Theory and Applications* 65 (1990) 409–429.

[9] J. Gondzio, "Multiple Centrality Corrections in a Primal-Dual Method for Linear Programming", *Computational Optimization and Applications* 6 (1996) 137–156.

[10] 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) 221–225.

[11] J. Gondzio, O. du Merle, R. Sarkissian and J.-P. Vial, "ACCPM - A Library for Convex Optimization Based on an Analytic Center Cutting Plane Method", *European Journal of Operational Research* 94 (1996) 206-211.

[12] J. Gondzio and R. Sarkissian, "Column Generation with the Primal-Dual Method", Logilab Technical Report 96.4, Department of Management Studies, University of Geneva, Switzerland, June 1996.

[13] C. C. Gonzaga, "Path–Following Methods for Linear Programming", *SIAM Review* 34 (1992) 167–224.

[14] A.L. Hipolito, "A Weighted Least Squares Study of Robustness in Interior Point Linear Programming", *Computational Optimization and Applications* 2 (1993) 29–46.

[15] B. Jansen, C. Roos, T. Terlaky, and J.-P. Vial. "Primal–Dual Target Following Algorithms for Linear Programming", *Annals of Operation Research* 62 (1996) 197–231.

[16] J.E. Kelley, "The Cutting Plane Method for Solving Convex Programs", *Journal of the SIAM* 8 (1960) 703–712.

[17] 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) 123–135.

[18] I.J. Lustig, R. E. Marsten and D. F. Shanno. "Interior Point Methods for Linear Programming: Computational State of the Art", *ORSA Journal on Computing* 6 (1994) 1–14.

[19] S. Mehrotra, "On the Implementation of a Primal–Dual Interior Point Method", *SIAM Journal on Optimization* 2 (4) (1992) 575–601.

[20] O. du Merle, J.-L. Goffin and J.-P. Vial, "On the Comparative Behavior of Kelley's Cutting Plane Method and the Analytic Center Cutting Plane Method", Technical Report 1996.4, Logilab, Department of Management Studies, University of Geneva, Switzerland, March 1996.

[21] J.E. Mitchell, "Karmarkar's Algorithm and Combinatorial Optimization Problems", PhD. Thesis, Cornell University, 1988.

[22] J.-P. Vial, "A Generic Path-following Algorithm with a Sliding Constraint and its Application to Linear Programming and the Computation of Analytic Centers", Technical Report 1996.8, Logilab, Department of Management Studies, University of Geneva, Switzerland, February 1996.

[23] S.J. Wright, *Primal-Dual Interior-Point Methods* (SIAM, Philadelphia, 1997).