

Active-set prediction for interior point methods using controlled perturbations

Coralia Cartis* and Yiming Yan†

April 15, 2014; Revised March 25, 2015

Abstract

We propose the use of controlled perturbations to address the challenging question of optimal active-set prediction for interior point methods. Namely, in the context of linear programming, we consider perturbing the inequality constraints/bounds so as to enlarge the feasible set. We show that if the perturbations are chosen appropriately, the solution of the original problem lies on or close to the central path of the perturbed problem. We also find that a primal-dual path-following algorithm applied to the perturbed problem is able to accurately predict the optimal active set of the original problem when the duality gap for the perturbed problem is not too small; furthermore, depending on problem conditioning, this prediction can happen sooner than predicting the active set for the perturbed problem or when the original one is solved. Encouraging preliminary numerical experience is reported when comparing activity prediction for the perturbed and unperturbed problem formulations.

Keywords: active-set prediction, interior point methods, linear programming.

1 Introduction

Optimal active-set prediction — namely, identifying the active inequality constraints at the solution of a constrained optimization problem — plays an important role in the optimization process by removing the difficult combinatorial aspect of the problem and reducing it to an equality-constrained one that is in general easier to solve. Active-set prediction is also crucial for efficient warmstarting and re-optimization capabilities of algorithms when a suite of closely related problems needs to be solved. Despite being state-of-the-art tools for solving large-scale Linear Programming (LP) problems [40], Interior Point Methods (IPMs) are well-known to encounter difficulties with active-set prediction due essentially to their construction. They generate iterates that progress towards the solution set through the (relative) interior of the feasible set, and thus avoid visiting possibly-many feasible vertices. This however, may also prevent IPMs from getting accurate information about the optimal active set early enough during their running. When this information is more readily predictable/available towards the end of a run, as the iterates approach the solution set, the algorithm has to solve increasingly ill-conditioned and hence difficult, subproblems. Finding ways to improve (even just partial) active set prediction for IPMs could thus be beneficial as it would allow earlier termination of an otherwise ill-conditioned and computationally expensive process by say, projecting onto the solution set (as in finite termination [42]), help with reducing the problem size or with obtaining a vertex solution at the cost of just a few additional (and less expensive) simplex method iterations.

Various ways have been devised for IPMs to predict the optimal active set during their run, with the simplest being *cut-off* [17, 23, 25] — which splits the variables into active or inactive based on whether they are less than a user-defined small value — and the most well-known being *indicators* [9] which form functions of iterates and identify the optimal active-set based on whether the values of these functions are less than a threshold. Mehrotra [26] suggests determining the active set by a simple comparison of the relative increments of primal and dual iterates, and Mehrotra and Ye [28] propose a strategy to identify the active set by comparing the primal variables with the dual slacks; see [39] for a review of active-set prediction techniques for IPMs for LP and also [32] for a more recent survey.

*Mathematical Institute, University of Oxford, Andrew Wiles Building, Radcliffe Observatory Quarter, Woodstock Road, Oxford, OX2 6GG, United Kingdom. Email: cartis@maths.ox.ac.uk.

†School of Mathematics, University of Edinburgh, James Clerk Maxwell Building, The King's Buildings, Mayfield Road, Edinburgh, EH9 3JZ, United Kingdom. Email: yiming.yan@ed.ac.uk. This work was supported by the Principal's Career Development Scholarship from the University of Edinburgh.

Here we propose the use of *controlled perturbations* [5] for active-set prediction for IPMs.¹ Namely, we perturb the inequality constraints of the LP problem (by a small amount) so as to enlarge the feasible set of the problem, then solve the resulting perturbed problem(s) using a path-following IPM while predicting on the way the active set of the original LP problem. As Figure 1.1 illustrates, provided the perturbations are chosen judiciously, the central path of the perturbed problem may pass close to the optimal solution of the original LP problem when the barrier parameter for the perturbed problem is ‘not too small’. Thus we expect that while still ‘far’ from optimality for the perturbed problem, some IPM iterates for the perturbed problem would nonetheless be close to optimality for the original LP problem (such as the third and fourth iterate in Figure 1.1) and would provide a good prediction of the original optimal active set. As it may happen that the chosen perturbations are ‘too large’ or not sufficiently effective for active-set prediction, we allow them to shrink after each IPM iteration so that the resulting perturbed feasible set is smaller but still contains the feasible set of the original LP.

Since we employ perturbed problems, albeit artificially, our proposal may be remindful of warmstarting techniques for IPMs and the related active-set prediction techniques that have been developed in that context; see for example, the surveys [11, 37]. Thus we briefly review relevant contributions here. One of the main warmstarting strategies focuses on the ‘iterates’, namely it manipulates the (IPM-computed) near optimal or optimal iterates of the initial problem to obtain a primal-dual feasible and well-centred point for the perturbed problems, see for example, [18, 20, 44, 19, 37]. Another category of approaches works on the ‘problem formulation’, namely modify the problem formulation by relaxing the nonnegativity constraints in the form of shifted logarithmic barrier variables, which has some similarity to our approach. Earlier works in this framework include Freund [14, 15, 16], Mitchell [29] and Polyak [34] with promising theoretical properties. More relevant and closer in spirit to our approach here is [3], where Benson and Shanno propose a primal-dual penalty strategy relaxing the nonnegativity constraints for both primal and dual decision variables and then penalising the relaxation variables in the objective; encouraging numerical results are also reported. Engau, Anjos and Vannelli [10, 11] apply a simplified primal-dual slack approach: instead of shifting the bounds and penalising the relaxation variables, slack variables for nonnegative constraints are introduced and penalised in the objective. One of the main differences between the above techniques and our approach is that we consider perturbations as parameters, not variables that are updated in the run of the IPM; furthermore, our focus is different as we specifically aim to predict the active set of the original LP problem by using ‘fake’ perturbations.

Another set of techniques — regularization for IPMs [35, 1, 6] — is also only loosely connected to our approach. In order to improve the conditioning of the coefficient matrix arising in calculating Newton directions in IPM iterations, regularization terms (of proximal type, weighted, and quadratic in the variables) are added to the (primal and dual) objective function. These terms result in a diagonal perturbation of the linear KKT system of interest, improving stability of factorization procedures. Note that the effect of our perturbations on the Newton system is not the same in that no similar diagonal perturbation is obtained. This is due to our formulations having no quadratic terms in the variables in the primal-dual objective, only a quadratic term in the perturbations; and to our approach perturbing the inequality constraints of the problem and allowing negative components of the primal and dual slack variables. However, the two techniques have similar aims in that they attempt to deal with the increasing ill-conditioning that affects IPMs by improved early active-set prediction (hence earlier termination and better conditioning) for our approach and by directly improving the conditioning of the linear algebra through regularization.

To validate the use of controlled perturbations for active-set prediction for IPMs, after introducing them and the associated primal-dual perturbed LP problems (Section 2), we show that, for properly chosen perturbations, the solution of the original LP problem lies on or close to the central path of the perturbed problems. Furthermore, in nondegenerate cases, the optimal active set of the perturbed problems remains the same as for the original problem (Section 3). We also prove that under certain conditions that do not necessarily require problem nondegeneracy, our predicted active sets provide a lower and upper bound for the optimal active set of the original problem, and exactly predict this set under a certain nondegeneracy assumption (but without requiring that the perturbed active set coincides with the original one). We also

¹Note that [5] proposed the use of such perturbations for creating a sequence of LPs with strict interior, converging to the original LP in the limit, so as to find the affine dimension of the feasible set of the original LP and well-centred points in Phase I of IPMs; a different focus and approach than here. A relaxation technique for Mathematical Programs with Equilibrium Constraints (MPECs) was also proposed independently in [8] that relaxes the bound constraints similarly to [5], but that also relaxes the complementarity constraint, in order to create a sequence of nonlinear programming relaxations with strict interior, which thus satisfy a constraint qualification, and also converge to the original MPEC in the limit.




Figure 1.1: Enlarge the feasible set and predict the original active set

find conditions on problem conditioning that ensure that our prediction of the optimal active set of the original LP can happen sooner than the prediction of the optimal active set of the perturbed problems (so that our approach may not need to solve the perturbed problems to high accuracy) (Section 5.1). Similarly, we characterise the situations when our approach allows an earlier prediction of the original active set as compared to the case when we solve and predict the original LP directly (Section 5.2). In our preliminary numerical experiments (Section 6), we carry out two type of tests, one comparing the accuracy of the predicted active sets and the other one exploring the case of crossover to simplex method. For verifying the accuracy of our active-set predictions, we apply an infeasible primal-dual path-following IPM to perturbed and original randomly-generated LP problems, terminate the algorithm at various iterations and compare the accuracy of predictions using certain correction comparison ratios. We observe that when using perturbations, the precision of our predictions is generally higher — namely, more than 4 times higher at certain iterations — than that when we do not use perturbations. When crossing over to simplex method, we test the efficiency of our active-set predictions by comparing the number of simplex iterations needed to solve the original problem to optimality, after some initial IPM iterations. We conduct this test on both randomly generated problems and a subset of Netlib problems. We find that when using perturbations for the IPM iterations, we can save (on average) over 30% simplex iterations compared to the case of not using any perturbations before cross-over to simplex.

2 Controlled perturbations for linear programming

Consider the following pair of primal-dual linear programming (LP) problems,

$$\begin{array}{ll}
 \text{(Primal)} & \text{(Dual)} \\
 \min_{x \in \mathbb{R}^n} & c^T x & \max_{(y,s) \in \mathbb{R}^m \times \mathbb{R}^n} & b^T y \\
 \text{s.t.} & Ax = b, & \text{s.t.} & A^T y + s = c, \\
 & x \geq 0, & & s \geq 0,
 \end{array} \tag{PD}$$

where $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$ and $c \in \mathbb{R}^n$ with $m \leq n$ are problem data, and $(x, y, s) \in \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^n$.

We enlarge the feasible set of this (PD) problem by using *controlled perturbations* [5], namely we relax the nonnegativity constraints in (PD) and consider the pair of perturbed problems²,

$$\begin{array}{ll}
 \text{(Primal)} & \text{(Dual)} \\
 \min_{x \in \mathbb{R}^n} & (c + \lambda)^T (x + \lambda) & \max_{(y,s) \in \mathbb{R}^m \times \mathbb{R}^n} & (b + A\lambda)^T y \\
 \text{s.t.} & Ax = b, & \text{s.t.} & A^T y + s = c, \\
 & x \geq -\lambda, & & s \geq -\lambda,
 \end{array} \tag{PD}_\lambda$$

for some vector of perturbations $\lambda \geq 0$. It can be checked [5] that the two problems in $(PD)_\lambda$ are dual to each other. Note that if $\lambda \equiv 0$, $(PD)_\lambda$ coincides with (PD). We denote the set of *strictly feasible points*

²Different perturbations for x and s could be used, but for simplicity, we use the same vector of perturbations for both.

of (PD_λ) ,

$$\mathcal{F}_\lambda^0 = \{(x, y, s) \mid Ax = b, A^T y + s = c, x + \lambda > 0, s + \lambda > 0\}. \quad (2.1)$$

Writing down the first order optimality conditions (KKT conditions) for (PD_λ) , according for example to [31, Theorem 12.1], we find that $(x_\lambda^*, y_\lambda^*, s_\lambda^*)$ is a *(primal-dual) solution* for (PD_λ) if and only if it satisfies the following system,

$$\begin{aligned} Ax &= b, \\ A^T y + s &= c, \\ (X + \Lambda)(S + \Lambda)e &= 0, \\ (x + \lambda, s + \lambda) &\geq 0, \end{aligned} \quad (2.2)$$

where $\Lambda = \text{diag}(\lambda)$, $X = \text{diag}(x)$, $S = \text{diag}(s)$ and $e = [1 \ \dots \ 1]^T$. Again if $\lambda \equiv 0$ in (2.2), we recover the optimality conditions for (PD) .

Equivalent formulation of (PD_λ) . Letting $p = x + \lambda$ and $q = s + \lambda$, we can write (PD_λ) in the equivalent form,

$$\begin{array}{ll} \text{(Primal)} & \text{(Dual)} \\ \min_p & c_\lambda^T p \\ \text{s.t.} & Ap = b_\lambda, \\ & p \geq 0, \end{array} \quad \begin{array}{ll} \max_{(y,q)} & b_\lambda^T y \\ \text{s.t.} & A^T y + q = c_\lambda, \\ & q \geq 0, \end{array} \quad (2.3)$$

where $c_\lambda = c + \lambda$, $b_\lambda = b + A\lambda$ and $\lambda \geq 0$. KKT conditions ensure that $(p_\lambda^*, y_\lambda^*, q_\lambda^*)$ is the (primal-dual) solution of (2.3) if and only if it satisfies

$$\begin{aligned} Ap &= b_\lambda, \\ A^T y + q &= c_\lambda, \\ PQe &= 0, \\ (p, q) &\geq 0, \end{aligned} \quad (2.4)$$

where $P = \text{diag}(p)$ and $Q = \text{diag}(q)$. It is easy to show that $(x_\lambda^*, y_\lambda^*, s_\lambda^*)$ is a (PD_λ) solution if and only if $(p_\lambda^*, y_\lambda^*, q_\lambda^*)$, where $p_\lambda^* = x_\lambda^* + \lambda$ and $q_\lambda^* = s_\lambda^* + \lambda$, is a solution of (2.3). Thus we can construct an optimal solution for (PD_λ) from an optimal solution of (2.3) and vice versa.

The central path of (PD_λ) . Following [40, Chapter 2], we derive the central path equations for (PD_λ) to be

$$\begin{aligned} Ax &= b, \\ A^T y + s &= c, \\ (X + \Lambda)(S + \Lambda)e &= \mu e, \\ (x + \lambda, s + \lambda) &> 0, \end{aligned} \quad (2.5)$$

where $\mu > 0$ is the barrier parameter for the perturbed problem (PD_λ) . The central path of (PD_λ) is well defined under mild assumptions, including

$$\textbf{Assumption:} \quad A \text{ has full row rank } m. \quad (2.6)$$

Lemma 2.1 ([5, Lemma 5.1]) *Let (2.6) hold and $\lambda \geq 0$. Then the central path of the perturbed problem (PD_λ) is well defined, namely, the system (2.5) has a unique solution for each $\mu > 0$, provided \mathcal{F}_λ^0 in (2.1) is nonempty. In particular, if $\lambda > 0$, \mathcal{F}_λ^0 is nonempty whenever (PD) has a nonempty primal-dual feasible set.*

Note that if $\lambda > 0$, the condition required for the existence of the perturbed central path is weaker than that for the central path of (PD) . The latter requires (PD) to have a nonempty *strictly* feasible set, namely, for there to be (PD) feasible points that strictly satisfy all problem inequality constraints.

3 Perturbed problems and their properties

3.1 Perfect and relaxed perturbations

Geometrically, the original optimal solution (x^*, y^*, s^*) of (PD) may lie on or near the central path of the perturbed problem (PD_λ) for carefully chosen perturbations; see Figures 3.1 and 3.2. Algebraically, this happens if (x^*, y^*, s^*) satisfies the third relation in (2.5) exactly or approximately. We make these considerations precise in the next two theorems.




Figure 3.1: Perfect perturbations.




Figure 3.2: Relaxed perturbations.

Theorem 3.1 (Existence of ‘perfect’ perturbations) Assume (2.6) holds and (x^*, y^*, s^*) is a solution of (PD). Let $\hat{\mu} > 0$. Then there exists a vector of perturbations

$$\hat{\lambda} = \hat{\lambda}(x^*, s^*, \hat{\mu}) > 0,$$

such that the perturbed central path (2.5) with $\lambda = \hat{\lambda}$ passes through (x^*, y^*, s^*) exactly when $\mu = \hat{\mu}$.

Proof. Since (x^*, y^*, s^*) is an optimal solution of (PD), it is also primal-dual feasible, and so $(x^*, y^*, s^*) \in \mathcal{F}_\lambda^0$ for any $\lambda > 0$. Thus, according to Lemma 2.1, the perturbed central path is well defined. Furthermore, if there exists a $\hat{\lambda} > 0$ such that

$$(X^* + \hat{\Lambda})(S^* + \hat{\Lambda})e = \hat{\mu}e, \quad (3.1)$$

then (x^*, y^*, s^*) is the unique solution of the perturbed central path equations (2.5) with $\lambda = \hat{\lambda}$ and $\mu = \hat{\mu}$, which implies the central path of perturbed problems passes through (x^*, y^*, s^*) . It remains to solve (3.1) for $\hat{\lambda} = [\hat{\lambda}_1 \dots \hat{\lambda}_n]^T$. Since $x_i^* s_i^* = 0$, $i = 1, \dots, n$, we have that (3.1) is equivalent to

$$\hat{\lambda}_i^2 + (x_i^* + s_i^*)\hat{\lambda}_i - \hat{\mu} = 0, \quad i = 1, \dots, n,$$

whose positive root for each i gives the corresponding component of the required $\hat{\lambda}$. \square

It is a stringent and impractical requirement to force the optimal solution of the original problem to be exactly on the central path of the perturbed problems. Thus we relax this requirement to allow for the original solution to belong to a small neighbourhood of this path.

Theorem 3.2 (Existence of relaxed perturbations) Assume (2.6) holds and (x^*, y^*, s^*) is a (PD) solution and let $\hat{\mu} > 0$ and $\xi \in (0, 1)$. Then there exist vectors $\hat{\lambda}_L = \hat{\lambda}_L(x^*, s^*, \hat{\mu}, \xi) > 0$ and $\hat{\lambda}_U = \hat{\lambda}_U(x^*, s^*, \hat{\mu}, \xi) > 0$ such that for $\hat{\lambda}_L \leq \lambda \leq \hat{\lambda}_U$, (x^*, y^*, s^*) is strictly feasible for (PD $_\lambda$) and satisfies

$$\xi \hat{\mu} e \leq (X^* + \Lambda)(S^* + \Lambda)e \leq \frac{1}{\xi} \hat{\mu} e. \quad (3.2)$$

Proof. Clearly, (x^*, y^*, s^*) satisfies (2.1) and so $(x^*, y^*, s^*) \in \mathcal{F}_\lambda^0$ for any $\lambda = [\lambda_1 \dots \lambda_n]^T > 0$. The inequalities (3.2) are equivalent to

$$\begin{cases} \lambda_i^2 + (x_i^* + s_i^*)\lambda_i - \xi \hat{\mu} \geq 0 \\ \lambda_i^2 + (x_i^* + s_i^*)\lambda_i - \frac{1}{\xi} \hat{\mu} \leq 0, \end{cases} \quad (3.3)$$

for all $i \in \{1, \dots, n\}$ and $\xi \in (0, 1)$. Solving (3.3) for λ_i , we obtain

$$\begin{cases} \lambda_i \geq \frac{-(x_i^* + s_i^*) + \sqrt{(x_i^* + s_i^*)^2 + 4\xi\hat{\mu}}}{2} = \frac{2\xi\hat{\mu}}{x_i^* + s_i^* + \sqrt{(x_i^* + s_i^*)^2 + 4\xi\hat{\mu}}} = (\hat{\lambda}_L)_i, \\ 0 < \lambda_i \leq \frac{-(x_i^* + s_i^*) + \sqrt{(x_i^* + s_i^*)^2 + \frac{4\hat{\mu}}{\xi}}}{2} = \frac{\frac{2\hat{\mu}}{\xi}}{x_i^* + s_i^* + \sqrt{(x_i^* + s_i^*)^2 + \frac{4\hat{\mu}}{\xi}}} = (\hat{\lambda}_U)_i, \end{cases} \quad (3.4)$$

for all $i \in \{1, \dots, n\}$. For any $\xi \in (0, 1)$, it is easy to see that (3.4) yields a well-defined interval for λ_i , $i \in \{1, \dots, n\}$. \square

From the above theorem, we see that by choosing the perturbations judiciously, we can bring any solution of the original problem into a ‘neighbourhood’ of the perturbed central path.

3.2 Preserving the optimal active set

Since we are interested in predicting the optimal active set of the original problem, this section addresses the relation between the active set of the perturbed problem and that of the original LP. We find that for sufficiently small perturbations, these two active sets remain the same provided the original problem is nondegenerate.

Theorem 3.3 *Assume (2.6) holds and the original pair of (PD) problems has a unique and nondegenerate primal solution x^* .³ Then there exists a positive scalar $\hat{\lambda} = \hat{\lambda}(A, b, c, x^*)$ such that the pair of perturbed problems (PD_λ) with $0 \leq \|\lambda\| < \hat{\lambda}$ has a strictly complementary solution $(x_\lambda^*, y_\lambda^*, s_\lambda^*)$ with the same active and inactive sets as x^* , where $\|\cdot\|$ denotes the Euclidean norm.*

Proof. Since (PD) has a unique and nondegenerate primal solution, it must have a unique primal-dual nondegenerate solution (x^*, y^*, s^*) [36, Theorem 4.5 (b)], which must be strictly complementary and so $x^* + s^* > 0$. Thus, letting

$$\mathcal{A} = \{i \in \{1, \dots, n\} \mid x_i^* = 0\} \quad \text{and} \quad \mathcal{I} = \{i \in \{1, \dots, n\} \mid s_i^* = 0\}, \quad (3.5)$$

the KKT conditions for (PD) at (x^*, y^*, s^*) —namely, (2.2) with $\lambda = 0$ —become

$$x_{\mathcal{A}}^* = 0, \quad x_{\mathcal{I}}^* > 0 \quad \text{and} \quad s_{\mathcal{I}}^* = 0, \quad s_{\mathcal{A}}^* > 0, \quad (3.6a)$$

$$A_{\mathcal{I}} x_{\mathcal{I}}^* = b, \quad A_{\mathcal{I}}^T y^* = c_{\mathcal{I}}, \quad A_{\mathcal{A}}^T y^* + s_{\mathcal{A}}^* = c_{\mathcal{A}}, \quad (3.6b)$$

where $A = [A_{\mathcal{I}} \ A_{\mathcal{A}}]$, $(x^*)^T = [(x_{\mathcal{A}}^*)^T \ (x_{\mathcal{I}}^*)^T]$ and $(s^*)^T = [(s_{\mathcal{A}}^*)^T \ (s_{\mathcal{I}}^*)^T]$. As the (PD) solution is also nondegenerate, we must have $|\mathcal{I}| = m$ and $\text{rank}(A_{\mathcal{I}}) = m$, namely, $A_{\mathcal{I}}$ is nonsingular. We work with the equivalent form (2.3) of problems (PD_λ) , and construct a solution $(\hat{p}, \hat{y}, \hat{q})$ of (2.3) such that $\hat{p} + \hat{q} > 0$, $\hat{p}_{\mathcal{A}} = 0$ and $\hat{q}_{\mathcal{I}} = 0$, namely,

$$\hat{p}_{\mathcal{A}} = 0, \quad \hat{p}_{\mathcal{I}} = x_{\mathcal{I}}^* + \lambda_{\mathcal{I}} + A_{\mathcal{I}}^{-1} A_{\mathcal{A}} \lambda_{\mathcal{A}}, \quad (3.7a)$$

$$\hat{y} = y^* + (A_{\mathcal{I}}^T)^{-1} \lambda_{\mathcal{I}}, \quad \hat{q}_{\mathcal{I}} = 0, \quad \hat{q}_{\mathcal{A}} = s_{\mathcal{A}}^* + \lambda_{\mathcal{A}} - (A_{\mathcal{I}}^{-1} A_{\mathcal{A}})^T \lambda_{\mathcal{I}}. \quad (3.7b)$$

Using (3.6), it is straightforward to show that $(\hat{p}, \hat{y}, \hat{q})$ in (3.7) satisfies all linear and nonlinear equality constraints in the KKT conditions (2.4). It remains to prove that $\hat{p}_{\mathcal{I}} > 0$ and $\hat{q}_{\mathcal{A}} > 0$. Let σ_{\max} be the largest singular value of $A_{\mathcal{I}}^{-1} A_{\mathcal{A}}$, and define a positive scalar $\hat{\lambda}$ as

$$\hat{\lambda} = \frac{\min \{[x_{\mathcal{I}}^* \ s_{\mathcal{A}}^*]\}}{\sigma_{\max}},$$

where $\min \{[x_{\mathcal{I}}^* \ s_{\mathcal{A}}^*]\}$ is a scalar that denotes the smallest element of $x_{\mathcal{I}}^*$ and $s_{\mathcal{A}}^*$. From $\lambda \geq 0$ and from norm properties, we have that

$$\hat{p}_{\mathcal{I}} \geq x_{\mathcal{I}}^* - \|A_{\mathcal{I}}^{-1} A_{\mathcal{A}} \lambda_{\mathcal{A}}\| e_{\mathcal{I}} \geq x_{\mathcal{I}}^* - \|A_{\mathcal{I}}^{-1} A_{\mathcal{A}}\| \cdot \|\lambda_{\mathcal{A}}\| e_{\mathcal{I}} \geq x_{\mathcal{I}}^* - \|A_{\mathcal{I}}^{-1} A_{\mathcal{A}}\| \cdot \|\lambda\| e_{\mathcal{I}}$$

³An equivalent non-degeneracy assumption that would be sufficient here is to require that all (PD) solutions are primal-dual nondegenerate [21, Section 5].

and

$$\hat{q}_{\mathcal{A}} \geq s_{\mathcal{A}}^* - \|(A_{\mathcal{I}}^{-1} A_{\mathcal{A}})^T \lambda_{\mathcal{I}}\| e_{\mathcal{A}} \geq s_{\mathcal{A}}^* - \|(A_{\mathcal{I}}^{-1} A_{\mathcal{A}})^T\| \cdot \|\lambda_{\mathcal{I}}\| e_{\mathcal{A}} \geq s_{\mathcal{A}}^* - \|(A_{\mathcal{I}}^{-1} A_{\mathcal{A}})^T\| \cdot \|\lambda\| e_{\mathcal{A}}.$$

Using matrix norm properties, we obtain that $\|A_{\mathcal{I}}^{-1} A_{\mathcal{A}}\| = \|(A_{\mathcal{I}}^{-1} A_{\mathcal{A}})^T\| = \sigma_{\max}$. This and $0 < \|\lambda\| < \hat{\lambda}$ now imply

$$\hat{p}_{\mathcal{I}} > x_{\mathcal{I}}^* - \sigma_{\max} \hat{\lambda} e_{\mathcal{I}} \geq x_{\mathcal{I}}^* - \min\{[x_{\mathcal{I}}^* \ s_{\mathcal{A}}^*]\} e_{\mathcal{I}} \geq 0,$$

and

$$\hat{q}_{\mathcal{A}} > s_{\mathcal{A}}^* - \sigma_{\max} \hat{\lambda} e_{\mathcal{A}} \geq s_{\mathcal{A}}^* - \min\{[x_{\mathcal{I}}^* \ s_{\mathcal{A}}^*]\} e_{\mathcal{A}} \geq 0,$$

where we also use the definition of $\hat{\lambda}$. \square

Remark on the assumptions and proof of Theorem 3.3. We have assumed in this theorem that (PD) is primal-dual nondegenerate and has a unique solution, which guarantees $A_{\mathcal{I}}$ is nonsingular. Considering the general case when (x^*, y^*, s^*) is a possibly non-unique strictly complementary solution, to construct the desired solution $(\hat{p}, \hat{y}, \hat{q})$ of (2.3) with the same active set and strictly complementary partition, one needs to satisfy exactly primal-dual feasibility requirements such as

$$A_{\mathcal{I}} \hat{p}_{\mathcal{I}} = b + A\lambda = b + A_{\mathcal{A}} \lambda_{\mathcal{A}} + A_{\mathcal{I}} \lambda_{\mathcal{I}}. \quad (3.8)$$

Clearly, one can only guarantee (3.8) to be consistent for $\lambda > 0$ if $A_{\mathcal{A}} \lambda_{\mathcal{A}}$ belongs to the range space of $A_{\mathcal{I}}$. Alternatively, one could consider satisfying (3.8) only approximately and look for a solution \hat{p} of the form

$$\hat{p}_{\mathcal{A}} = 0 \quad \text{and} \quad \hat{p}_{\mathcal{I}} = x_{\mathcal{I}}^* + \lambda_{\mathcal{I}} + \hat{u}, \quad (3.9)$$

where \hat{u} is the least-squares/minimal norm solution of $A_{\mathcal{I}} \hat{u} = A_{\mathcal{A}} \lambda_{\mathcal{A}}$. For instance in the case when $|\mathcal{I}| \leq m$, we have $\|A_{\mathcal{I}} \hat{u} - A_{\mathcal{A}} \lambda_{\mathcal{A}}\| \leq \|A_{\mathcal{A}} \lambda_{\mathcal{A}}\|$. The right-hand side of the latter inequality goes to zero as $\lambda \rightarrow 0$ and so primal feasibility can be approximately achieved. It can also be shown that $\hat{p}_{\mathcal{I}}$ in (3.9) stays positive. \square

Note that the nondegeneracy assumption in Theorem 3.3 is not required in the results of the next section or in our implementations and numerical experiments. Thus this theorem and its assumptions do not restrict our algorithmic or even main theoretical approach of predicting the optimal active set of the (PD) problem by solving a perturbed (PD $_{\lambda}$) problem.

4 Using perturbations to predict the original optimal active set

Recalling our main aim, we now present results for predicting the optimal active set of (PD). The idea is to solve the perturbed problem instead of the original one using IPMs, but attempt to predict the active set for the original problem during the run of the algorithm. Without assuming that the original and perturbed problems have the same optimal active set, we prove that under certain conditions and given proper perturbations, when the duality gap of (PD $_{\lambda}$) is sufficiently small, the predicted (strictly) active set for (PD) coincides with the actual optimal (strictly) active set of (PD) (Theorems 4.5, 4.6).

4.1 Some useful results

We first derive a bound on the distance between the original optimal solution set and strictly feasible points of the perturbed problems.

Lemma 4.1 (An error bound for (PD)) *Let $(x, y, s) \in \mathcal{F}_{\lambda}^0$, where \mathcal{F}_{λ}^0 is defined in (2.1), and $\lambda \geq 0$. Then there exists a (PD) solution (x^*, y^*, s^*) such that*

$$\|x - x^*\| \leq \tau_p (r(x, s) + w(x, s)) \quad \text{and} \quad \|s - s^*\| \leq \tau_d (r(x, s) + w(x, s)), \quad (4.1)$$

where $\tau_p > 0$ and $\tau_d > 0$ are problem-dependent constants independent of (x, y, s) and (x^*, y^*, s^*) , and

$$r(x, s) = \|\min\{x, s\}\| \quad \text{and} \quad w(x, s) = \|(-x, -s, x^T s)_+\|, \quad (4.2)$$

and where $\min\{x, s\} = (\min(x_i, s_i))_{i=1, \dots, n}$ and $(x)_+ = (\max(x_i, 0))_{i=1, \dots, n}$.

See Appendix A for a proof of this lemma.

Lemma 4.2 [40, Lemma 5.13] *For any $(x, y, s) \in \mathcal{F}_\lambda^0$, where \mathcal{F}_λ^0 is defined in (2.1), we have*

$$0 < x_i + \lambda_i \leq \frac{\mu_\lambda}{C_1} \quad (i \in \mathcal{A}_\lambda) \quad \text{and} \quad 0 < s_i + \lambda_i \leq \frac{\mu_\lambda}{C_1} \quad (i \in \mathcal{I}_\lambda), \quad (4.3)$$

where

$$\mu_\lambda = \frac{(x + \lambda)^T (s + \lambda)}{n} \quad (4.4)$$

and

$$C_1 = \frac{\epsilon(A, b_\lambda, c_\lambda)}{n} \quad (4.5)$$

with

$$\epsilon(A, b_\lambda, c_\lambda) = \min \left(\min_{i \in \mathcal{I}_\lambda} \sup_{x_\lambda^* \in \Omega_\lambda^P} \{ (x_\lambda^*)_i + \lambda_i \}, \min_{i \in \mathcal{A}_\lambda} \sup_{(y_\lambda^*, s_\lambda^*) \in \Omega_\lambda^D} \{ (s_\lambda^*)_i + \lambda_i \} \right) > 0,^4 \quad (4.6)$$

and Ω_λ^P and Ω_λ^D are the primal and dual solution sets of (PD_λ) respectively, and where $(\mathcal{A}_\lambda, \mathcal{I}_\lambda)$ is the strictly complementary active and inactive partition of the solution set of (PD_λ) .

Proof. Apply [40, Lemma 5.13] to (2.3) and recall $x = p - \lambda$ and $s = q - \lambda$. Note that Lemma 5.13 is a more complex result that also assumes loose proximity to the problem central path, but only strict feasibility is required to prove the required inequalities in (4.3). \square

Lemma 4.3 *Let $(x, y, s) \in \mathcal{F}_\lambda^0$, where \mathcal{F}_λ^0 is defined in (2.1) for some $\lambda \geq 0$. Then there exists a (PD) solution (x^*, y^*, s^*) and problem-dependent constants τ_p and τ_d that are independent of (x, y, s) and (x^*, y^*, s^*) , such that*

$$\|x - x^*\| < \tau_p (C_2 \mu_\lambda + 4\|\lambda\| \max(\|\lambda\|, 1)) \quad \text{and} \quad \|s - s^*\| < \tau_d (C_2 \mu_\lambda + 4\|\lambda\| \max(\|\lambda\|, 1)), \quad (4.7)$$

where

$$C_2 = \frac{n\sqrt{n}}{\epsilon(A, b_\lambda, c_\lambda)} + n, \quad (4.8)$$

$\epsilon(A, b_\lambda, c_\lambda)$ is defined in (4.6) and μ_λ in (4.4).

Proof. Since $x + \lambda > 0$ and $s + \lambda > 0$, we have $-x < \lambda$ and $-s < \lambda$, which implies

$$0 \leq (-x)_+ < \lambda \quad \text{and} \quad 0 \leq (-s)_+ < \lambda. \quad (4.9)$$

Using (4.4), $\lambda \geq 0$ and $(x + \lambda, s + \lambda) \geq 0$, we have

$$x^T s = n\mu_\lambda + \lambda^T \lambda - \lambda^T (x + \lambda) - \lambda^T (s + \lambda) \leq n\mu_\lambda + \|\lambda\|^2. \quad (4.10)$$

From (4.2), (4.9) and (4.10), we obtain

$$w(x, s) \leq \|(-x)_+\| + \|(-s)_+\| + (x^T s)_+ \leq n\mu_\lambda + 2\|\lambda\| + \|\lambda\|^2. \quad (4.11)$$

It remains to find an upper bound for $r(x, s)$ in (4.2). If $i \in \mathcal{A}_\lambda$, from (4.3) we have $\min(x_i + \lambda_i, s_i + \lambda_i) \leq x_i + \lambda_i \leq \frac{\mu_\lambda}{C_1}$. Similarly, we also have $\min(x_i + \lambda_i, s_i + \lambda_i) \leq \frac{\mu_\lambda}{C_1}$ for $i \in \mathcal{I}_\lambda$. Thus $0 < \min\{x + \lambda, s + \lambda\} \leq \frac{\mu_\lambda}{C_1} e$, and so from (4.2),

$$r(x, s) = \|\min\{x + \lambda, s + \lambda\} - \lambda\| \leq \|\min\{x + \lambda, s + \lambda\}\| + \|\lambda\| \leq \frac{\mu_\lambda}{C_1} \sqrt{n} + \|\lambda\|. \quad (4.12)$$

The bound (4.7) now follows from (4.1), (4.11) and (4.12). \square

⁴When the feasible set of (PD) is nonempty, that of (PD_λ) is also nonempty, and so $\epsilon(A, b_\lambda, c_\lambda) > 0$.

4.2 Predicting the original optimal active set using perturbations

Assume (x^*, y^*, s^*) is a (PD) solution. We denote by $\mathcal{A}(x^*)$ the optimal active set at x^* and by $\mathcal{A}_+(s^*)$, the ‘strictly’ active set at s^* , namely,

$$\mathcal{A}(x^*) = \{i \in \{1, \dots, n\} \mid x_i^* = 0\} \quad \text{and} \quad \mathcal{A}_+(s^*) = \{i \in \{1, \dots, n\} \mid s_i^* > 0\}. \quad (4.13)$$

Let

$$\bar{\mathcal{A}}(x) = \{i \in \{1, \dots, n\} \mid x_i < C\} \quad \text{and} \quad \bar{\mathcal{A}}_+(s) = \{i \in \{1, \dots, n\} \mid s_i \geq C\}, \quad (4.14)$$

where $C > 0$ is some constant threshold. $\bar{\mathcal{A}}(x)$ is considered as the predicted active set and $\bar{\mathcal{A}}_+(s)$, the predicted ‘strictly’ active set at a primal-dual pair (x, y, s) for (PD $_\lambda$).

Theorem 4.4 *Let $C > 0$ and fix the perturbation λ such that*

$$0 < \|\lambda\| < \min\left(1, \frac{C}{8 \max(\tau_p, \tau_d)}\right), \quad (4.15)$$

where τ_p and τ_d are the problem-dependent constants in (4.7). Let $(x, y, s) \in \mathcal{F}_\lambda^0$ with μ_λ sufficiently small, namely,

$$\mu_\lambda < \frac{C}{2C_2 \max(\tau_p, \tau_d)}, \quad (4.16)$$

where \mathcal{F}_λ^0 is defined in (2.1), μ_λ in (4.4) and $C_2 > 0$ in (4.8) is a problem-dependent constant when λ is fixed. Then there exists a (PD) solution (x^*, y^*, s^*) such that

$$\bar{\mathcal{A}}_+(s) \subseteq \mathcal{A}_+(s^*) \subseteq \mathcal{A}(x^*) \subseteq \bar{\mathcal{A}}(x).$$

Proof. From $\|\lambda\| < 1$ and (4.7), we have $\|x - x^*\| \leq \tau_p (C_2 \mu_\lambda + 4\|\lambda\|)$ and $\|s - s^*\| \leq \tau_d (C_2 \mu_\lambda + 4\|\lambda\|)$, which imply

$$x_i^* - \tau_p (C_2 \mu_\lambda + 4\|\lambda\|) \leq x_i \leq x_i^* + \tau_p (C_2 \mu_\lambda + 4\|\lambda\|) \quad (4.17)$$

and

$$s_i^* - \tau_d (C_2 \mu_\lambda + 4\|\lambda\|) \leq s_i \leq s_i^* + \tau_d (C_2 \mu_\lambda + 4\|\lambda\|), \quad (4.18)$$

for all $i \in \{1, \dots, n\}$. If $i \in \mathcal{A}(x^*)$, from (4.15), (4.16) and (4.17), we have $x_i < C$, namely $i \in \bar{\mathcal{A}}(x)$. So $\mathcal{A}(x^*) \subseteq \bar{\mathcal{A}}(x)$. If $i \notin \mathcal{A}_+(s^*)$, $s_i^* = 0$. Then from (4.15), (4.16) and (4.18), we have $s_i < C$, namely, $i \notin \bar{\mathcal{A}}_+(s)$. Thus $\bar{\mathcal{A}}_+(s) \subseteq \mathcal{A}_+(s^*)$. From $x_i^* s_i^* = 0$ for all $i \in \{1, \dots, n\}$, we have $\mathcal{A}_+(s^*) \subseteq \mathcal{A}(x^*)$. \square

Theorem 4.4 shows that $\bar{\mathcal{A}}(x)$ and $\bar{\mathcal{A}}_+(s)$ serve as a pair of approximations that bound $\mathcal{A}(x^*)$. Next we go a step further and show that $\bar{\mathcal{A}}(x)$ is equivalent to $\mathcal{A}(x^*)$ under certain conditions.

Theorem 4.5 *Let*

$$\psi_p = \inf_{x^* \in \Omega^P} \min_{i \notin \mathcal{A}(x^*)} (x_i^*) \quad (4.19)$$

where Ω^P is the solution set of the primal problem in (PD) and $\mathcal{A}(x^*)$ is defined in (4.13). Assume $\psi_p > 0$. Fix λ and C such that

$$0 < \|\lambda\| < \min\left(1, \frac{\psi_p}{16 \max(\tau_p, \tau_d)}\right) \quad \text{and} \quad C = \frac{\psi_p}{2}, \quad (4.20)$$

where τ_p and τ_d are the problem-dependent constants defined in (4.7). Let $(x, y, s) \in \mathcal{F}_\lambda^0$ with μ_λ sufficiently small, namely,

$$\mu_\lambda < \frac{\psi_p}{4C_2 \max(\tau_p, \tau_d)}, \quad (4.21)$$

where \mathcal{F}_λ^0 is defined in (2.1), μ_λ in (4.4) and $C_2 > 0$ in (4.8). Then there exists a (PD) solution (x^*, y^*, s^*) such that

$$\bar{\mathcal{A}}(x) = \mathcal{A}(x^*),$$

where $\bar{\mathcal{A}}(x)$ is defined in (4.14).

Proof. From Theorem 4.4 we have $\mathcal{A}(x^*) \subseteq \bar{\mathcal{A}}(x)$. It remains to prove $\bar{\mathcal{A}}(x) \subseteq \mathcal{A}(x^*)$. If $i \notin \mathcal{A}(x^*)$, from the left inequality in (4.17), (4.20) and (4.21), we have

$$x_i > x_i^* - \frac{\psi_p}{2} \cdot \frac{\tau_p}{\max(\tau_p, \tau_d)} \geq \inf_{x^* \in \Omega^D} \min_{i \notin \mathcal{A}(x^*)} (x_i^*) - \frac{\psi_p}{2} = \psi_p - \frac{\psi_p}{2} = C.$$

Thus $i \notin \bar{\mathcal{A}}(x)$, which implies $\bar{\mathcal{A}}(x) \subseteq \mathcal{A}(x^*)$. \square

Next, we show that $\bar{\mathcal{A}}_+(s)$, the predicted strictly active set at a strictly feasible point (x, y, s) of (PD_λ) , is the same as $\mathcal{A}_+(s^*)$ at some (PD) solution (x^*, y^*, s^*) .

Theorem 4.6 *Let*

$$\psi_d = \inf_{(y^*, s^*) \in \Omega^D} \min_{i \in \mathcal{A}_+(s^*)} (s_i^*)$$

where Ω^D is the solution set of the dual problem in (PD) and $\mathcal{A}_+(s^*)$ is defined in (4.13). Assume $\psi_d > 0$. Fix λ and C such that

$$0 < \|\lambda\| < \min \left(1, \frac{\psi_d}{16 \max(\tau_p, \tau_d)} \right) \quad \text{and} \quad C = \frac{\psi_d}{2}, \quad (4.22)$$

where τ_p and τ_d are the problem-dependent constants in (4.7). Let $(x, y, s) \in \mathcal{F}_\lambda^0$ with μ_λ sufficiently small, namely

$$\mu_\lambda < \frac{\psi_d}{4C_2 \max(\tau_p, \tau_d)}, \quad (4.23)$$

where \mathcal{F}_λ^0 is defined in (2.1), μ_λ in (4.4) and $C_2 > 0$ in (4.8). Then there exists a (PD) solution (x^*, y^*, s^*) such that

$$\bar{\mathcal{A}}_+(s) = \mathcal{A}_+(s^*),$$

where $\bar{\mathcal{A}}_+(s)$ is defined in (4.14).

Proof. From Theorem 4.4, we have $\bar{\mathcal{A}}_+(s) \subseteq \mathcal{A}_+(s^*)$. If $i \in \mathcal{A}_+(s^*)$, $s_i^* > 0$. This, (4.18), (4.22) and (4.23) give us

$$s_i > s_i^* - \frac{\psi_d}{2} \cdot \frac{\tau_d}{\max(\tau_p, \tau_d)} \geq \inf_{(y^*, s^*) \in \Omega^D} \min_{i \in \mathcal{A}_+(s^*)} (s_i^*) - \frac{\psi_d}{2} = \psi_d - \frac{\psi_d}{2} = C,$$

namely $\mathcal{A}_+(s^*) \subseteq \bar{\mathcal{A}}_+(s)$. \square

Remarks on Theorems 4.4–4.6.

- We require μ_λ , the mean value of the complementary products, to be sufficiently small in Theorems 4.4–4.6. This choice is possible since we have $\mu_\lambda = 0$ at any optimal solution of (PD_λ) and μ_λ can be decreased to zero (such as in an IPM framework).
- In Theorems 4.5 and 4.6, we do not require that the optimal active set of (PD_λ) is the same as that of (PD) in order to be able to predict the original optimal active set of (PD) .
- ψ_p in (4.19) is positive if the primal problem in (PD) has a unique (degenerate or nondegenerate) solution, but we expect that it may often be zero in the case of multiple solutions. (Clearly, in our implementations, we do not choose the cut-off value based on the theoretical quantity ψ_p .) Similarly to ψ_p , if the dual problem in (PD) has a unique (degenerate or nondegenerate) solution, we have $\psi_d > 0$.
- Fix λ sufficiently small and let (x^k, y^k, s^k) be iterates of a primal-dual path-following IPM applied to (PD_λ) . Then assuming these iterates belong to some good neighbourhood of the central path of (PD_λ) and that the barrier parameter is decreased appropriately, we have $\mu_\lambda^k \rightarrow 0$ as $k \rightarrow \infty$ [40, Theorem 5.11]. So, by applying Theorem 4.5, for each k sufficiently large, there exists a (PD) solution (x^*, y^*, s^*) such that $\bar{\mathcal{A}}(x^k) = \mathcal{A}(x^*)$ (see also Lemma 5.1 below).

5 Comparing perturbed and unperturbed active-set predictions

5.1 Comparing with active-set prediction for (PD_λ)

Consider the ‘large’ neighbourhood of the perturbed central path

$$\mathcal{N}_{-\infty}(\gamma, \lambda) = \{ (x, y, s) \in \mathcal{F}_\lambda^0 \mid (x_i + \lambda_i)(s_i + \lambda_i) \geq \gamma \mu_\lambda, i = 1, \dots, n \}, \quad (5.1)$$

where \mathcal{F}_λ^0 is defined in (2.1) and μ_λ is defined in (4.4); see [40, (1.16)] for the definition (5.1) in the case of $\lambda \equiv 0$.

Next we rephrase Lemma 5.13 in [40] as an active-set prediction result for (PD_λ) .

Lemma 5.1 *Let (x, y, s) in $\mathcal{N}_{-\infty}(\gamma, \lambda)$ and μ_λ defined in (4.4). Assume C in (4.14) is set to $C = \frac{\epsilon(A, b_\lambda, c_\lambda)\gamma}{n}$, where $\epsilon(A, b_\lambda, c_\lambda)$ is defined in (4.6). Then when $\mu_\lambda < \bar{\mu}_\lambda^{\max}$, where*

$$\bar{\mu}_\lambda^{\max} = \frac{\epsilon^2(A, b_\lambda, c_\lambda)\gamma}{n^2}, \quad (5.2)$$

for any strictly complementary solution $(x_\lambda^*, y_\lambda^*, s_\lambda^*)$ of (PD_λ) we have

$$\bar{\mathcal{A}}(x + \lambda) = \mathcal{A}(x_\lambda^* + \lambda),$$

where $\bar{\mathcal{A}}(x + \lambda)$ is defined in (4.14) with x replaced by $x + \lambda$ and $\mathcal{A}(x_\lambda^* + \lambda)$ is defined in (4.13) with x^* replaced by $x_\lambda^* + \lambda$.

Proof. We work with the equivalent form (2.3) of (PD_λ) . Given (5.2), apply [40, Lemma 5.13] to (2.3), recalling that $x = p - \lambda$ and $s = q - \lambda$, and then we have

$$\begin{aligned} i \in \mathcal{A}_\lambda : \quad & 0 < x_i + \lambda_i \leq \frac{\mu_\lambda}{C_1} < C_1 \gamma \leq s_i + \lambda_i, \\ i \in \mathcal{I}_\lambda : \quad & 0 < s_i + \lambda_i \leq \frac{\mu_\lambda}{C_1} < C_1 \gamma \leq x_i + \lambda_i, \end{aligned} \quad (5.3)$$

where $(\mathcal{A}_\lambda, \mathcal{I}_\lambda)$ is the strictly complementary active and inactive partition of the solution set of (2.3). For any strictly complementary solution $(x_\lambda^*, y_\lambda^*, s_\lambda^*)$ of (PD_λ) , $(x_\lambda^* + \lambda, y_\lambda^*, s_\lambda^* + \lambda)$ is a strictly complementary solution of (2.3). This and the definition of $\mathcal{A}(x_\lambda^* + \lambda)$ give us that $\mathcal{A}(x_\lambda^* + \lambda) = \mathcal{A}_\lambda$. From (5.3) and the definition of $\bar{\mathcal{A}}(x + \lambda)$, we also have $\bar{\mathcal{A}}(x + \lambda) = \mathcal{A}_\lambda$. \square

Substituting (4.8) into (4.21), we obtain the following threshold value

$$\mu_\lambda^{\max} := \frac{\psi_p \epsilon(A, b_\lambda, c_\lambda)}{4n \max(\tau_p, \tau_d) (\sqrt{n} + \epsilon(A, b_\lambda, c_\lambda))}, \quad (5.4)$$

where $\epsilon(A, b_\lambda, c_\lambda)$ is defined in (4.6), ψ_p in (4.19), and τ_p and τ_d are the positive constants in the bounds (4.7). Theorem 4.5 provides that when $\psi_p > 0$ and λ is sufficiently small and fixed, if $\mu_\lambda < \mu_\lambda^{\max}$, we can predict the optimal active set of (PD) . Lemma 5.1 shows that when $\mu_\lambda < \bar{\mu}_\lambda^{\max}$, where $\bar{\mu}_\lambda^{\max}$ is defined in (5.2), we can provide the strictly complementary partition of the solution set of (PD_λ) from any primal-dual pair in the neighbourhood $\mathcal{N}_{-\infty}(\gamma, \lambda)$ of the perturbed central path. To verify if our approach can predict the optimal active set of (PD) before the strictly complementary partition of (PD_λ) , we determine conditions under which $\mu_\lambda^{\max} > \bar{\mu}_\lambda^{\max}$.

Theorem 5.2 *In the conditions of Theorem 4.5, let*

$$\rho = \frac{\psi_p}{\max(\tau_p, \tau_d)}. \quad (5.5)$$

If

$$\epsilon(A, b_\lambda, c_\lambda) \leq \mathcal{O}(\sqrt{n\rho} \min(\sqrt{\rho}, 1)), \quad (5.6)$$

then

$$\mu_\lambda^{\max} > \bar{\mu}_\lambda^{\max},$$

where $\epsilon(A, b_\lambda, c_\lambda)$ is defined in (4.6), μ_λ^{\max} in (5.4) and $\bar{\mu}_\lambda^{\max}$ in (5.2).

Proof. Note that $\mu_\lambda^{\max} > \bar{\mu}_\lambda^{\max}$ is equivalent to $\epsilon^2(A, b_\lambda, c_\lambda) + \sqrt{n}\epsilon(A, b_\lambda, c_\lambda) - \frac{\rho}{4\gamma} < 0$, which is satisfied if

$$0 < \epsilon(A, b_\lambda, c_\lambda) \leq \frac{\sqrt{n}}{2\sqrt{\gamma}} \cdot \frac{\rho}{\sqrt{\gamma} + \rho + \sqrt{\gamma}}. \quad (5.7)$$

Since $\gamma \in (0, 1)$ and $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for any a and b nonnegative scalars, we have

$$\frac{\rho}{\sqrt{\gamma} + \rho + \sqrt{\gamma}} \geq \frac{\rho}{\sqrt{\rho} + 2\sqrt{\gamma}} \geq \frac{1}{3\sqrt{\gamma}} \frac{\rho}{\max(\sqrt{\rho}, 1)} \geq \frac{\sqrt{\rho}}{3\sqrt{\gamma}} \min(\sqrt{\rho}, 1).$$

The result follows from (5.7) and the above inequalities. \square

Theorem 5.2 implies that when solving the perturbed problems (PD_λ) , if $\epsilon(A, b_\lambda, c_\lambda)$ is sufficiently small, we can predict the optimal active set of (PD) before μ_λ gets so small that we can even obtain the strictly complementary partition of (PD_λ) . To see an example when (5.6) is satisfied, see our remarks after Theorem 5.5.

Remark. In Theorem 5.2, we do not require the optimal active set of (PD_λ) to be the same as the optimal active set of (PD) . In fact, we will show that, in the numerical tests for the randomly generated problems (degenerate or nondegenerate), the optimal active sets of most perturbed problems are different from those of the original problems, but we can still predict sooner/better for (PD) . In particular, the numerical experiments show that we are not solving (PD_λ) to high accuracy and there are iterations where we can predict the active set for (PD) but we are not close to the solution set of (PD_λ) or able to predict the active set of (PD_λ) ; see page 17. \square

5.2 Comparing with active-set prediction for (PD)

Similarly to Lemma 5.1, when we solve the original (PD) problems we can predict the optimal (PD) active set when the (PD) duality gap is smaller than some threshold. In this section, we intend to compare this threshold with the threshold value of μ_λ when we are able to predict the optimal active set of (PD) by solving (PD_λ) and show that the latter could be greater than the former under certain conditions (Theorem 5.5).

Lemma 5.13 in [40] yields an active-set prediction result for (PD) . In fact this result can be obtained by setting $\lambda = 0$ in Lemma 5.1, but for clarity, we restate it here.

Lemma 5.3 [40, Lemma 5.13] *Let (x, y, s) in $\mathcal{N}_{-\infty}(\gamma)$, where $\mathcal{N}_{-\infty}(\gamma)$ is the neighbourhood $\mathcal{N}_{-\infty}(\gamma, \lambda)$ in (5.1) with $\lambda = 0$, and let μ as in (4.4) with $\lambda = 0$. Let the cut-off value C in (4.14) be set to $C = \frac{\epsilon(A, b, c)\gamma}{n}$, where*

$$\epsilon(A, b, c) = \min \left(\min_{i \in \mathcal{I}} \sup_{x^* \in \Omega^P} x_i^*, \min_{i \in \mathcal{A}} \sup_{(y^*, s^*) \in \Omega^D} s_i^* \right) > 0,^5 \quad (5.8)$$

Ω^P and Ω^D are the primal and dual solution sets of (PD) respectively, and $(\mathcal{A}, \mathcal{I})$ is the strictly complementary active and inactive partition of the solution set of (PD) . When $\mu < \mu^{\max}$, where

$$\mu^{\max} = \frac{\epsilon^2(A, b, c)}{n^2} \gamma, \quad (5.9)$$

then for any strictly complementary solution (x^*, y^*, s^*) of (PD) we have

$$\bar{\mathcal{A}}(x) = \mathcal{A}(x^*),$$

where $\bar{\mathcal{A}}(x)$ is defined in (4.14) and $\mathcal{A}(x^*)$ is defined in (4.13).

Before we deduce a relationship between μ_λ^{\max} in (5.4) and μ^{\max} in (5.9), we first relate two other important quantities, $\epsilon(A, b_\lambda, c_\lambda)$ and $\epsilon(A, b, c)$.

⁵ $\epsilon(A, b, c) > 0$ provided that the feasible set of (PD) is nonempty.

Lemma 5.4 Assume (2.6) holds and (PD) has a unique and nondegenerate solution (x^*, y^*, s^*) . Then there exists a sufficiently small $\bar{\lambda}(A, b, c, x^*, s^*) > 0$ such that

$$\epsilon(A, b_\lambda, c_\lambda) > \epsilon(A, b, c) \quad (5.10)$$

for all λ such that $0 \leq \lambda = \alpha \bar{\lambda} < \bar{\lambda}$, where $\alpha \in (0, 1)$, and where $\epsilon(A, b_\lambda, c_\lambda)$ is defined in (4.6) and $\epsilon(A, b, c)$ in (5.8).

The proof of this lemma is given in Appendix B.

Theorem 5.5 In the conditions of Theorem 4.5, assume (2.6) holds and (PD) has a unique and nondegenerate solution (x^*, y^*, s^*) . Provided

$$\epsilon(A, b, c) \leq \mathcal{O}(\sqrt{n\rho} \min(\sqrt{\rho}, 1)), \quad (5.11)$$

where ρ is defined in (5.5), there exists a sufficiently small $\bar{\lambda}(A, b, c, x^*, s^*) > 0$ such that

$$\mu_\lambda^{\max} > \mu^{\max},$$

for all $0 < \lambda = \alpha \bar{\lambda} < \bar{\lambda}$, where $\alpha \in (0, 1)$ and where μ_λ^{\max} is defined in (5.4) and μ^{\max} in (5.9).

Proof. Applying Theorem 5.2 with $\lambda = 0$ and so replacing $\epsilon(A, b_\lambda, c_\lambda)$ with $\epsilon(A, b, c)$, we deduce

$$\mu^{\max} < \frac{\rho \epsilon(A, b, c)}{4n(\sqrt{n} + \epsilon(A, b, c))}.$$

From Lemma 5.4, we have $\epsilon(A, b_\lambda, c_\lambda) > \epsilon(A, b, c)$. This and the definition of μ_λ^{\max} in (5.4) give

$$\frac{\rho \epsilon(A, b, c)}{4n(\sqrt{n} + \epsilon(A, b, c))} < \mu_\lambda^{\max}.$$

□

Theorem 5.5 implies that if $\epsilon(A, b, c)$ is sufficiently small, we may find the optimal active set of (PD) ‘sooner’ if we solve (PD) $_\lambda$ using a primal-dual path-following IPM than if we solve (PD).

Remark. When (PD) has a unique solution (x^*, y^*, s^*) , we have

$$\epsilon(A, b, c) = \min \left(\min_{i \in \mathcal{I}} x_i^*, \min_{i \in \mathcal{A}} s_i^* \right) \leq \min_{i \in \mathcal{I}} x_i^* = \psi_p.$$

Note that according to [24] $\tau_p, \tau_d = \mathcal{O}(1)$ numerically. Thus provided $\psi_p > 1$ or n is sufficiently large, (5.11) is satisfied. We illustrate this in an example next.

A simple example of predicting the optimal (PD) active set using perturbations. To illustrate our results in this section, consider the following simple example

$$\min x_1 + 2x_2 \quad \text{subject to} \quad x_1 + x_2 = 1, \quad x_1 \geq 0, \quad x_2 \geq 0, \quad (5.12)$$

with the optimal solution $x^* = (1, 0)$ and $y^* = 1, s^* = (0, 1)$. Thus (5.12) has a unique and primal-dual nondegenerate solution with optimal active set $\mathcal{A}(x^*) = \{2\}$, and so $\psi_p = \epsilon(A, b, c) = 1$. Let the vector of perturbations be $\lambda = \alpha(1, 5)$ where $\alpha = 10^{-2}$. The perturbed problems (PD) $_\lambda$ also have a unique solution $x_\lambda^* = (1+5\alpha, -5\alpha)$, $y_\lambda^* = 1+\alpha$ and $s_\lambda^* = (-\alpha, 1-\alpha)$. So $\epsilon(A, b_\lambda, c_\lambda) = \min(1+6\alpha, 1+4\alpha) = 1+4\alpha = 1.04$.

First we verify the conditions in Theorem 4.5, which are needed in both Theorems 5.2 and 5.5. Since it is not clear how to deduce the value of τ_p and τ_d , we estimate them numerically⁶ and it turns out that $\tau_p \approx \tau_d \approx 0.8$. We set the cut-off constant C that separates the active and inactive constraints to be $C = \frac{\psi_p}{2} = 0.5$ and verify that $\|\lambda\| = \sqrt{26}\alpha < \frac{\psi_p}{16 \max(\tau_p, \tau_d)} < 1$. Thus the conditions in (4.20) are satisfied. Based on Theorem 4.5, we can predict the original optimal active set when μ_λ is less than $\mu_\lambda^{\max} \approx 0.0662$.

⁶We estimate τ_p and τ_d from their definition in (4.1), namely, we solve the following optimisation problem in MATLAB, $\max \|x - x^*\|/(r(x, s) + w(x, s))$ subject to $(x, y, s) \in \mathcal{F}_\lambda^0$, where $r(x, s)$ and $w(x, s)$ are defined in (4.2) and \mathcal{F}_λ^0 in (2.1); similarly for τ_d .

Next we verify Theorems 5.2 and 5.5. From (5.2), we get $\rho \approx 1.25$, and so $\sqrt{n\rho} \min(\sqrt{\rho}, 1) \approx 1.58$. Thus $0 < \epsilon(A, b, c) < \epsilon(A, b_\lambda, c_\lambda) < \sqrt{n\rho} \min(\sqrt{\rho}, 1)$, which implies that conditions (5.6) and (5.11) are satisfied. For the constant γ , it is common to choose a small value to have a large neighbourhood of the central path; set $\gamma = 0.01$. Then from (5.2) and (5.9), we have $\bar{\mu}_\lambda^{\max} \approx 0.0027 < \mu_\lambda^{\max}$ and $\mu^{\max} = 0.0025 < \mu_\lambda^{\max}$. This implies that when we use perturbations, we can predict the original optimal active set sooner than the perturbed active set or the original active set without perturbations. Furthermore, the threshold values (constant C) needed to separate the active constraints from the inactive ones for predicting the perturbed active set and the original active set without perturbations are 0.0052 and 0.005 respectively, both of which are much smaller than the cut-off $C = \frac{\psi_p}{2} = 0.5$ for predicting the original optimal active set using perturbations.

6 Numerical results

6.1 The perturbed algorithm and its implementation

All numerical experiments in this section employ an infeasible primal-dual path-following interior point method structure [40, Chapter 6] whether applied to (PD $_\lambda$) or (PD). The perturbed algorithm is summarised in Algorithm 6.1.

Algorithm 6.1: Perturbed Algorithm Framework.

Given perturbations $(\lambda^0, \phi^0) > 0$ and a starting point (x^0, y^0, s^0) with $(x^0 + \lambda^0, s^0 + \phi^0) > 0$;
for $k = 0, 1, 2, \dots$

solve the perturbed system (2.5) using Newton's method, namely

$$\begin{bmatrix} A & 0 & 0 \\ 0 & A^T & I \\ S^k + \Phi^k & 0 & X^k + \Lambda^k \end{bmatrix} \begin{bmatrix} \Delta x^k \\ \Delta y^k \\ \Delta s^k \end{bmatrix} = - \begin{bmatrix} Ax^k - b \\ A^T y^k + s^k - c \\ (X^k + \Lambda^k)(S^k + \Phi^k)e - \sigma^k \mu_\lambda^k e \end{bmatrix}, \quad (6.1)$$

 where $\sigma^k \in [0, 1]$ and

$$\mu_\lambda^k = \frac{(x^k + \lambda^k)^T (s^k + \phi^k)}{n}; \quad (6.2)$$

set $x^{k+1} = x^k + \alpha_p^k \Delta x^k$ and $(y^{k+1}, s^{k+1}) = (y^k, s^k) + \alpha_d^k (\Delta y^k, \Delta s^k)$, where

(α_p^k, α_d^k) is chosen such that $(x^{k+1} + \lambda^k, s^{k+1} + \lambda^k) > 0$;

predict the optimal active set of (PD) and denote by \mathcal{A}^k ;

terminate if some termination criterion is satisfied;

calculate $(\lambda^{k+1}, \phi^{k+1})$ possibly by shrinking (λ^k, ϕ^k) so that

$(x^{k+1} + \lambda^{k+1}, s^{k+1} + \phi^{k+1}) > 0$;

end (for).

Algorithm without perturbations. For comparison purposes, we refer to the algorithm with no perturbations (Algorithm 6.1 with $\lambda = \phi = 0$) as **Algorithm 6.2**. We denote the duality gap for Algorithm 6.2 as μ^k , which is equivalent to μ_λ^k in (6.2) with $\lambda^k = \phi^k = 0$.

Starting point. We use the starting point proposed by Mehrotra for (PD) [27, Section 7] as starting point for both Algorithms 6.1 and 6.2.⁷

Solving the Newton system (6.1). We follow [40, Chapter 11] and solve the augmented system form of (6.1). Also we set $\sigma^k = \min(0.1, 100\mu_\lambda^k)$.

⁷We have also tested the case when Algorithm 6.1 is initialised from Mehrotra's starting point for (PD $_\lambda$). This change did not affect our results in any significant way, suggesting some level of robustness.

Choice of perturbations. In our theory, we used the same vector of perturbations for both primal and dual variables. For better numerical efficiency, we have different perturbations λ and ϕ for primal and dual variables respectively. We set the initial perturbations to be $\lambda^0 = \phi^0 = 10^{-2}e$, where e is a vector of ones.⁸

Choice of stepsize. We choose a fixed, close to 1, fraction of the stepsize to the nearest constraints' boundary in the primal and dual spaces, respectively.

Shrinking the perturbations. One possible reason for getting a poor prediction of the active set is that the current perturbations are too large. So after we get the new iterate $(x^{k+1}, y^{k+1}, s^{k+1})$, we shrink the perturbations accordingly. Assume $t^{k+1} = \min(x^{k+1})$ and $v^{k+1} = \min(s^{k+1})$. We update the perturbations as follows,

$$\lambda^{k+1} = \begin{cases} \eta\lambda^k, & \text{if } t^{k+1} > 0 \\ (1 - \zeta)\lambda^k + \zeta(-t^{k+1})e, & \text{if } t^{k+1} \leq 0 \end{cases},$$

and

$$\phi^{k+1} = \begin{cases} \eta\phi^k, & \text{if } v^{k+1} > 0 \\ (1 - \zeta)\phi^k + \zeta(-v^{k+1})e, & \text{if } v^{k+1} \leq 0 \end{cases},$$

where $\eta \in (0, 1]$ and $\zeta \in (0, 1)$. It follows that $x^{k+1} + \lambda^{k+1} > 0$ and $s^{k+1} + \phi^{k+1} > 0$. We observed in our numerical experiments that when solving nondegenerate problems, it is better to shrink faster, roughly keeping the perturbations to be $\mathcal{O}(\mu_\lambda)$. When solving degenerate problems however, it is better to shrink slower, at a rate of $\mathcal{O}(\sqrt{\mu_\lambda})$. It is difficult and often impossible to distinguish a priori between degenerate and nondegenerate cases. After several numerical trials, we chose to set $\eta = 1$ and $\zeta = 0.5$.

Active-set prediction. In our theory, we considered that all variables less than a threshold are active at the solution. In practice, we apply a more complex strategy, inspired by [9, Step 3 in Procedure 8.1]. We partition the index set $\{1, 2, \dots, n\}$ into three sets, \mathcal{A}^k as the predicted active set, \mathcal{I}^k as the predicted inactive set and $\mathcal{Z}^k = \{1, 2, \dots, n\} \setminus (\mathcal{A}^k \cup \mathcal{I}^k)$ which includes all undetermined indices, and during the running of the algorithm, we move indices between these sets according to the following criteria,

$$x_i^k < C \quad \text{and} \quad s_i^k > C, \tag{6.3}$$

where C is a constant user-defined threshold. Theorem 4.4 guarantees the above criteria (6.3) are promising, as we are predicting the original optimal active set by estimating the intersection of $\bar{\mathcal{A}}_+(s_i^k)$ and $\bar{\mathcal{A}}(x_i^k)$. Initialise $\mathcal{A}^0 = \mathcal{I}^0 = \emptyset$ and $\mathcal{Z}^0 = \{1, 2, \dots, n\}$. An index is moving from \mathcal{Z}^k to \mathcal{A}^k if (6.3) is satisfied for two consecutive iterations, otherwise from \mathcal{Z}^k to \mathcal{I}^k . We move an index from \mathcal{A}^k to \mathcal{Z}^k if (6.3) is not satisfied at the current iteration. An index is moving from \mathcal{I}^k to \mathcal{Z}^k if (6.3) is satisfied at the current iteration. In our implementation, we choose $C = 10^{-5}$. Procedure C.1 in Appendix C contains a pseudocode of our active-set prediction technique. Our strategy enables us to make use of both primal and dual information which may be beneficial given Theorems 4.4–4.6.

Termination. Termination criteria will be defined for each set of tests.

6.2 Numerical results

6.2.1 Test problems

Randomly generated test problems (TS1). We first randomly generate the number of constraints $m \in (10, 200)$, the number of variables $n \in (20, 500)$ and density of nonzero entries in A within $(0.4, 0.8)$, where $m < n$, $2m < n < 7m$. Then randomly generate a matrix $A \in \mathbb{R}^{m \times n}$ of given density and a point $(x, y, s) \in \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^n$ with $x \geq 0$, $s \geq 0$ and density about 0.5. Finally we generate b and c by letting $b = Ax$ and $c = A^\top y + s$. Thus (x, y, s) serves as a feasible point. Problems generated this way are generally well-conditioned and primal nondegenerate⁹. This test set is inspired by the random problem

⁸We have done experiments to explore the sensitivity of our algorithm to the value of the initial perturbations. For example, choosing $\lambda^0 = \phi^0 = 10^{-1}e$ yields a high false-prediction ratio (proportion of mistakes). Perturbations of order 10^{-2} and 10^{-3} yield quickly a good approximation of the original (PD) active set. For $\lambda^0 = \phi^0 = 10^{-4}e$, the perturbed algorithm starts to behave similarly to the unperturbed one simply because the perturbations are too small.

⁹We checked the degeneracy of 100 test problems generated this way, by looking at the vertex solutions obtained by the MATLAB simplex solver. Majority of these problems are dual degenerate.

generation approach in [13, Section 8.3.4]. Whenever we use this test set, (the same) 100 problems are generated.

Randomly generated primal-dual degenerate test problems (TS2). Instead of generating a feasible point as for TS1, we generate (x, y, s) with $x \geq 0$, $s \geq 0$, $x_i s_i = 0$ for all $i \in \{1, \dots, n\}$ so that the number of nonzeros of x is strictly less than m and that of s is strictly less than $n - m$. Then get A , b , c as for TS1. Thus (x, y, s) serves as a primal-dual degenerate solution. 100 problems are also generated for this test set.

Netlib problems (TS3). Most Netlib¹⁰ test problems are not in the standard form. We reformulate them into the standard form by introducing slacks. Since our implementation is basic, in MATLAB, and mainly for illustration, we choose a subset of problems in Netlib with the number of primal variables less than 5000 (including the slack variables). See Table D.1 for the list of the 37 Netlib problems selected.

6.2.2 On the accuracy of active-set predictions using prediction ratios

Assume \mathcal{A}^k is the predicted active set at iteration k and \mathcal{A} is the actual optimal active set. To compare the accuracy of the predictions, we introduce the following three prediction ratios.

- False-prediction ratio = $\frac{|\mathcal{A}^k \setminus (\mathcal{A}^k \cap \mathcal{A})|}{|\mathcal{A}^k \cup \mathcal{A}|}$.
- Missed-prediction ratio = $\frac{|\mathcal{A} \setminus (\mathcal{A}^k \cap \mathcal{A})|}{|\mathcal{A}^k \cup \mathcal{A}|}$.
- Correction ratio = $\frac{|\mathcal{A}^k \cap \mathcal{A}|}{|\mathcal{A}^k \cup \mathcal{A}|}$.

False-prediction ratio measures the degree of incorrectly identified active constraints, missed-prediction ratio measures the degree of incorrectly rejected active constraints and correction ratio shows the accuracy of the prediction. All three ratios range from 0 to 1. If the predicted set is the same as the actual optimal active set, correction ratio is 1. The main task for this test is to compare the three measures for Algorithms 6.1 and 6.2.

When an LP problem has multiple solutions, the active set of a vertex solution is different from that of the strictly complementary solutions (about 17% difference on average for TS1 and 21% for TS2). To understand which active set do the (perturbed) Algorithm 6.1 and the (unperturbed) 6.2 predict, we terminate both algorithms at the same iteration and compare the predicted active sets with the actual optimal active sets obtained from an interior point solver and a simplex solver¹¹.

Prediction ratios for test sets TS1 and TS2. In Figures 6.1 and 6.2, we present the results for TS1 (left) and TS2 (right). The x-axis shows the number of interior point iterations at which we terminate the algorithms. In each figure, the first three plots (from left to right, top to bottom) show the average value of the three measures mentioned above for the test problems in question. The last plot at the bottom right corner presents the corresponding \log_{10} scaled relative KKT residuals. We measure the relative residual by

$$\text{relRes}^k = \frac{\|(Ax^k - b, A^T y^k + s^k - c, (X^k + \Lambda^k)(S^k + \Phi^k)e - \mu_{\lambda}^k e)\|}{1 + \max(\|b\|, \|c\|)}. \quad (6.4)$$

There are four lines in each plot, representing the prediction ratios by comparing the active set from Algorithm 6.1 with that from MATLAB's simplex solver (solid red line with circle) and from MATLAB's IPM (solid black line with square sign), and Algorithm 6.2 with simplex (dashed green line with diamond sign) and with IPM (dashed blue line with star) respectively.

- Figures 6.1 and 6.2 show that the average correction ratios for Algorithm 6.1 are at least as good and generally better than those for Algorithm 6.2. Thus it seems that using perturbations can only improve the active-set prediction capabilities of IPMs.
- Algorithm 6.2 is in fact an interior point solver applied to (PD) which approaches a strictly complementary (PD) solution. This is confirmed by having better correction ratio when comparing Algorithm 6.2 with the IPM than when comparing it with the simplex.

¹⁰Netlib is a collection of standard LP test problems. The original test problems can be found at <http://www.netlib.org/lp/data/>. We use a MATLAB version of the same test set, which is obtained from http://www.math.ntu.edu.tw/~wwang/cola_lab/test_problems/netlib_lp/.

¹¹We obtain the 'actual optimal active set' by solving the problem using MATLAB's solver LINPROG with the 'algorithm' option set to interior point or simplex and considering all variables less than 10^{-5} as active.

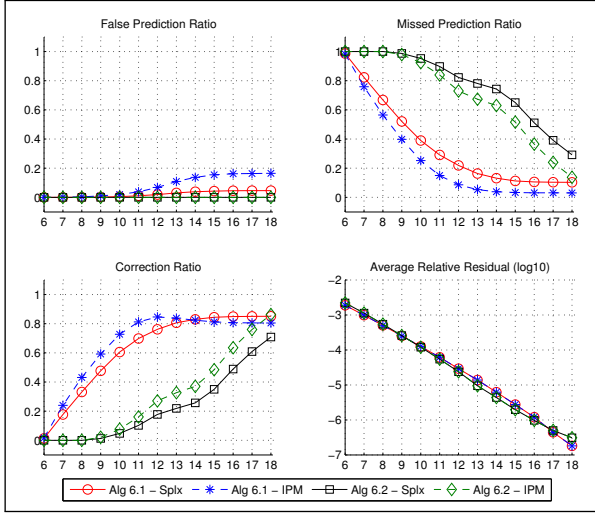


Figure 6.1: Prediction ratios for randomly generated problems

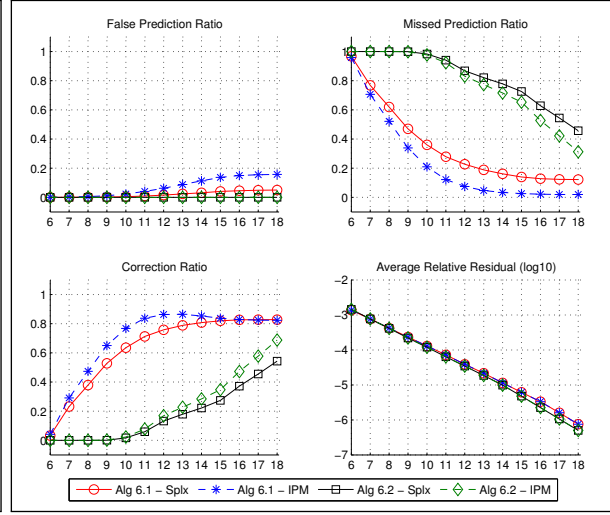


Figure 6.2: Prediction ratios for randomly generated primal-dual degenerate

- Due to the fact that the active set from the IPM (the strictly complementary partition) contains less elements than that from the simplex (vertex solution), the correction ratio of Algorithm 6.1 compared with the IPM is higher than that compared with the simplex at the early stage. However the false-prediction ratios of the former climb up to about 0.16 at the end for both test cases. Thus the corresponding correction ratios go down. The false-prediction ratios of comparing Algorithm 6.1 with simplex are much less, about 0.05 for both cases. The behaviours of the false-prediction ratios seem to imply that Algorithm 6.1 predicts the active set of a vertex solution (that may not be the same vertex as obtained by the simplex solver).
- After 18 iterations, the correction ratios do not reach 1. This is due to ill-conditioning which prevents us from solving any further. For this 18th iteration, the perturbations are not zero, they are about $\mathcal{O}(10^{-2})$ for problems in TS1 and $\mathcal{O}(10^{-3})$ for the degenerate problems in TS2, and on average the relative residual (6.4) is lower than 10^{-6} .

Can Algorithm 6.1 predict the optimal active set of (PD) sooner than it obtains the strictly complementary partition of (PD_λ) ? In Figures 6.3 and 6.4, besides comparing the predicted active set of (PD) with the actual active set of (PD) , we also compared the predicted active set of (PD_λ) ¹² with the actual active set of (PD_λ) obtained from a simplex solver (solid purple line with downward-pointing triangle) and an IPM solver (dashed brown line with upward-pointing triangle), respectively; see Footnote 11 on the choice of solvers. We again use the test sets TS1 and TS2.

We can see that on average Algorithm 6.1 can predict a better active set for (PD) than when applying Algorithm 6.2 to predict the active set of (PD_λ) . Furthermore, for test case TS1, before iteration 12, Algorithm 6.2 cannot predict much concerning the active set of (PD_λ) while Algorithm 6.1 already has an increasingly accurate prediction for the active set of (PD) (approximately 80% of the active set of (PD) at iterations 12). We can draw similar conclusions for TS2.

On the difference between the optimal active set of (PD_λ) and that of (PD) . Note that, to yield good performance, we do not need to force the active set of (PD_λ) (as defined in Footnote 12) to be the same as the (original) active set of (PD) . In fact, for most test problems in both TS1 and TS2, this does not hold. When perturbations are not so small, namely $\mathcal{O}(10^{-2})$ or $\mathcal{O}(10^{-3})$, which is the case even in the last IPM iterations in Figures 6.3 and 6.4, the perturbed optimal active set is different from the original optimal active set for 98% of the test problems in TS1 and all test problems in TS2. Furthermore, for problems in TS1, the average difference between the strictly complementary partition of (PD_λ) and

¹² Here, for each of the test problems, we set λ in (PD_λ) to be the value of the perturbations when terminating Algorithm 6.1 at the 18th iteration. We then apply Algorithm 6.2 to the equivalent form (2.3) of (PD_λ) , which means we solve the perturbed problem using an IPM method and predict the active set of the perturbed problem on the way.

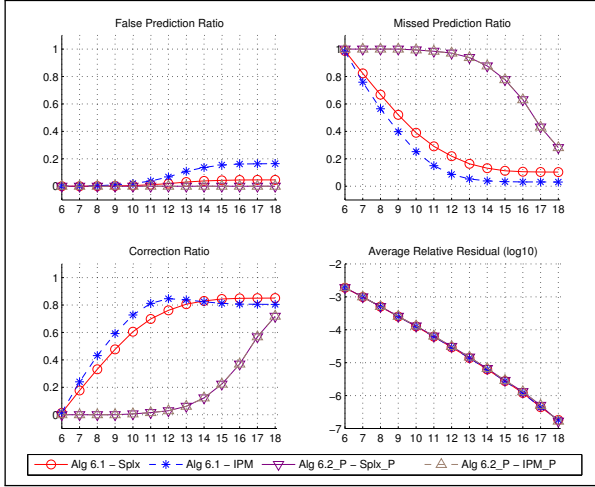


Figure 6.3: Comparing perturbed active-set predictions for (TS1)

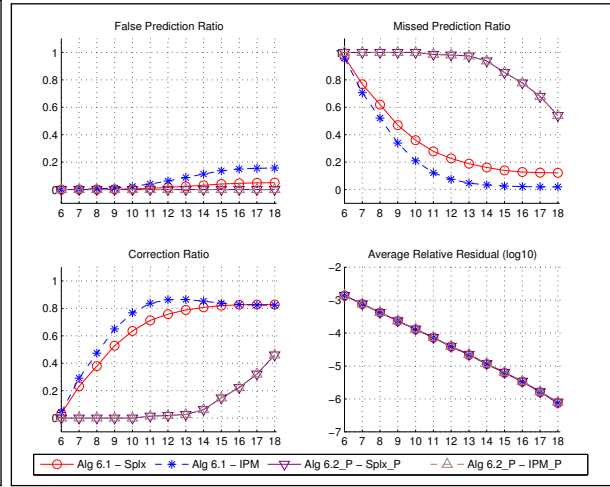


Figure 6.4: Comparing perturbed active-set predictions for (TS2)

that of (PD) is as high as 33% and the difference between the active set at a vertex solution of (PD) $_{\lambda}$ and that of (PD) is about 15% on average; for TS2, the average difference between the strictly complementary partitions of (PD) $_{\lambda}$ and (PD) is about 29% and the difference between active sets at vertex solutions is 17% on average. Another interesting observation is that, for both TS1 and TS2, over 90% of the perturbed problems have a unique and nondegenerate solution, regardless of the uniqueness or degeneracy of the original test problems. This is the reason why the predictions of the perturbed active set when comparing with simplex and IPM are identical in Figures 6.3 and 6.4.

Prediction ratios for test set TS3 (Netlib test problems). Figure 6.5 gives the prediction ratios for the Netlib test problems in TS3. In contrast to the randomly generated problems in TS1 and TS2, the number of iterations required by Netlib test problems to reach sufficient accuracy to allow meaningful predictions varies significantly from problem to problem (and so we cannot in general compare the prediction ratios at some fixed, predefined iterations). Thus to test the prediction ratios on TS3 problems, we follow a slightly different procedure¹³. For each TS3 test problem, we first solve it to optimality using Algorithm 6.2, requiring the relative residual in (6.4) (with $\lambda^k = \phi^k = 0$) to be less than 10^{-8} and we record the total number of iterations needed to reach this accuracy, say M . Then we calculate the prediction ratios on this same test problem for Algorithms 6.1 and 6.2 (with MATLAB's simplex and IPM output) over the last 10 iterations preceding (and including) the M^{th} iteration. We then average the prediction ratios for each algorithm on all TS3 test problems at each of the $M - i$ iterations for $i \in \{0, \dots, 9\}$. Again, in Figure 6.5, there are four lines in each plot, representing the prediction ratios by comparing the active set from Algorithm 6.1 with that from MATLAB's simplex solver (solid red line with circle) and from MATLAB's IPM (solid black line with square sign), and Algorithm 6.2 with simplex (dashed green line with diamond sign) and with IPM (dashed blue line with star), respectively. The bottom right figure plots the corresponding average relative residual (6.4) on a \log_{10} scale.

- Figure 6.5 shows that using perturbations can only improve the active-set prediction capabilities of IPMs on the tested Netlib problems, especially in the earlier stages of the runs, when the relative residuals are not too small. For example, the average correction ratios when using Algorithm 6.1 are about three times better than those of using Algorithm 6.2 at iteration $M - 5$, when the average relative residual is just slightly less than 10^{-3} .
- The correction ratios for both algorithms are slightly worse when compared with the vertex solution from MATLAB's simplex solver than with the strictly complementary solution from MATLAB's IPM solver; thus it is unclear in this case whether Algorithm 6.1 gets us closer to a vertex or an interior solution of the original problem (the cross-over to simplex results for TS3 in the next section seem to indicate the former is still the case).

¹³A similar methodology was also used in [9] to compare the performance of different indicators.

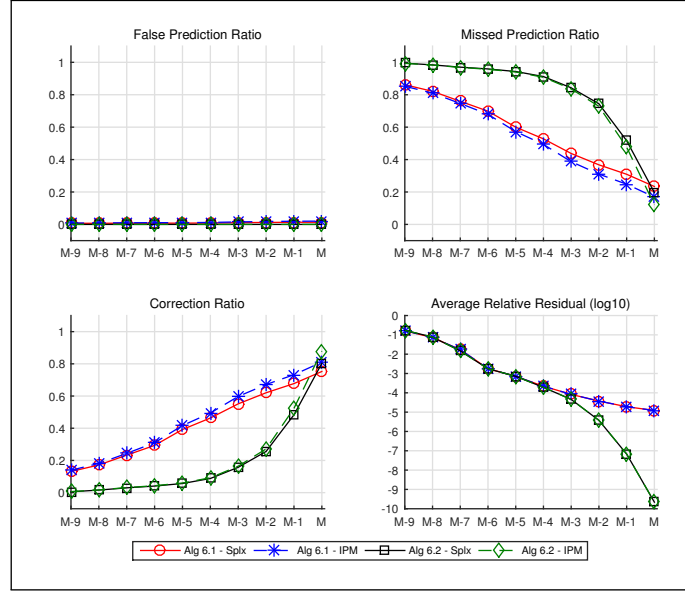


Figure 6.5: Comparing prediction ratios for the Netlib problems in TS3. M denotes the (variable) total number of iterations required to solve each test problem to a given accuracy (thus M is generally different for each test problem).

- The average relative residual in the bottom right plot is still quite large over the last few iterations for Algorithm 6.1 indicating that we have not solved the perturbed problems to high accuracy while still being able to predict well the optimal active set of the original problem, as desired.

6.2.3 Crossover to simplex

In this section, we test the efficiency of our active-set predictions using perturbations when crossing over to a simplex method after some IPM iterations. We choose LP_SOLVE [4] as our simplex solver¹⁴.

Initial basis for the simplex method. Assume we terminate the perturbed algorithm Algorithm 6.1 at the k^{th} iteration, with the predicted active set \mathcal{A}^k . To generate an initial basis \mathcal{B} from \mathcal{A}^k , we first obtain all independent columns in $A_{\mathcal{T}^k}$. If this submatrix is not of rank m , we choose a column from $A_{\mathcal{A}^k}$ and append it to the submatrix provided it is independent of existing columns in the submatrix. The order in which columns are added back in is decided by dual information, namely we keep trying a series of columns $\{A_{i_t}\}$, where $i_t \in \mathcal{A}^k$ and $s_{i_1}^k \leq s_{i_2}^k \leq \dots \leq s_{i_{|\mathcal{A}^k|}}^k$, until a full rank square matrix is obtained. Since A is full row rank¹⁵, this procedure is finite. A similar approach has been used in [38, Section 7] to form a basis of A .

To conduct the tests, we first choose a threshold μ_λ^{cap} , run Algorithm 6.1, terminate the algorithm when $\mu_\lambda^k < \mu_\lambda^{\text{cap}}$, record the number of interior point iterations, say K , generate an initial basis \mathcal{B} by the above procedure and finally start the simplex solver LP_SOLVE from the initial basis \mathcal{B} . For comparison purposes we perform exactly K iterations of Algorithm 6.2, and generate a new basis for (PD) by the same procedure, without constraining the value of μ^k . All tests in this part are run with $\mu_\lambda^{\text{cap}} = 10^{-3}$.

We compare the number of simplex iterations used to get an optimal solution after crossover from Algorithms 6.1 and 6.2, visualising the results via a relative performance profile [30]. Namely, we consider the following relative iteration count,

$$\text{rl}_i = -\log_2 \frac{\text{Iter}_i^p}{\text{Iter}_i^0}, \quad (6.5)$$

¹⁴Although there are many different simplex implementations, many of which are probably more efficient and powerful than LP_SOLVE, we chose LP_SOLVE because its MATLAB interface has the functionality that allows us to set the initial basis. To the best of our knowledge, this is the only such open source simplex solver.

¹⁵A matrix A can always be reduced to a full row rank matrix [40, Page 31-32]. In our tests, we apply the preprocessing code from LIPSOL [45] to ensure this.

where i stands for the i^{th} problem, the numerator stands for the number of simplex iterations performed after Algorithm 6.1 and the denominator measures the same but after Algorithm 6.2. If, for problem i , Algorithm 6.1 uses fewer simplex iterations, we get a positive valued bar with height $= rl_i$. If Algorithm 6.2 wins, we obtain a negative valued bar with height defined as $-rl_i$. The value of the bar will be 0 if these two yield the same simplex iterations or LP_SOLVE fails for both algorithms. If LP_SOLVE fails to solve problem i for Algorithm 6.1, we have a negative valued bar with height of $\max_i (|rl_i|)$, otherwise a positive valued bar with the same height. It is clear that the winner outperforms the loser by $2^{|rl_i|}$ times and one algorithm outperforms the other by having more bars (or larger area of bars) in its direction.

Crossover to simplex for randomly generated test problems (TS1 and TS2). In Figures 6.6 and 6.7, we show the profiles for TS1 (left) and TS2 (right), with bars sorted from largest to smallest in height. We can see that, counting the number of simplex iterations after each algorithm, the performance of Algorithm 6.1 dominates that of Algorithm 6.2 in both cases.




Figure 6.6: Simplex iteration count for randomly generated problems




Figure 6.7: Simplex iteration count for randomly generated primal-dual degenerate problems

In Table 6.1, we show the average number of simplex iterations, the average IPM iterations and the average μ_λ^k and μ^k when we terminate Algorithms 6.1 and 6.2 for both test sets (TS1 and TS2). On average, using perturbations saves about 34% simplex iterations for the test case TS1 and about 37% for TS2. Due to our experimental setup, the number of IPM iterations are the same for Algorithms 6.1 and 6.2, and the average final μ_λ^k and μ^k before crossover are of order 10^{-4} .¹⁶

Table 6.1: Crossover to simplex when $\mu_\lambda^k < 10^{-3}$ for random problems.

	Primal nondegenerate (TS1)		PD degenerate (TS2)	
	Algorithm 6.1	Algorithm 6.2	Algorithm 6.1	Algorithm 6.2
Avg simplex iterations	287	436	292	464
Avg IPM iterations	10	10	10	10
Avg μ_λ^k and μ^k when crossover	7.33×10^{-4}	6.80×10^{-4}	7.53×10^{-4}	7.14×10^{-4}

We also tracked the difference between the initial bases generated from Algorithms 6.1 and 6.2. We use relative difference¹⁷ to measure the degree of difference between two bases. On average the relative difference is over 60%, and over 90% of the test problems have greater than 50% relative difference. Thus our preliminary numerical experiments illustrate that using perturbations is likely to improve the efficiency when crossing over to simplex.

¹⁶The definition of μ_λ^k and μ^k in Algorithms 6.1 and 6.2, respectively, as well as the choice of (x^0, s^0) to be identical for (PD $_\lambda$) and (PD), imply that $\mu_\lambda^0 > \mu^0$, with the difference being essentially dictated by the level of perturbations (λ^0, ϕ^0) . Thus we are not making it any easier for Algorithm 6.1 compared to Algorithm 6.2 in the choice of starting point.

¹⁷The number of elements in either basis generated from Algorithms 6.1 or 6.2 but not both divided by the cardinality of the union of two bases.

Netlib test problems (TS3). The good prediction performance of the perturbed algorithm is not only obtained for randomly generated problems, but also for the subset of Netlib problems (TS3). Here we add an additional termination criterion, namely we terminate both algorithms when μ_λ^k and μ^k are less than 10^{-3} or when the relative residual (6.4) is less than 10^{-6} , whichever occurs first¹⁸.

Figure 6.8 presents the relative performance profile generated the same way as for the random tests (see (6.5) and accompanying explanation). From this figure, we can see that for over half of the test problems, Algorithm 6.1 outperforms Algorithm 6.2 by over 1.5 times. Algorithm 6.1 ‘loses’ for only 7 problems.




Figure 6.8: Crossover to simplex for 37 Netlib problems

We also summarise the results in Table 6.2. On average, we save about 38% simplex iterations by applying perturbations. The average numbers in the table exclude the data for SHIP08S, since LP_SOLVE fails to solve it when we do not apply perturbations.

Table 6.2: Crossover to simplex when $\mu_\lambda^k < 10^{-3}$ for 37 Netlib problems (TS3).

	Algorithm 6.1	Algorithm 6.2
Avg simplex iterations	358	612
Avg IPM iterations	22	22

We do not give the average value of μ_λ^k in the table as it is more involved than for random problems. In particular, for the problems with very large component in b (problems marked by * in Table D.1), the value of μ_λ^k is greater than 10^{-3} for both Algorithms 6.1 and 6.2. There are 8 additional problems, including 25FV47, BNL1, BRANDY, KB2, SCFXM2, SCRS8, SCATP1 and STAIR, for which the value of μ_λ^k is less than 10^{-3} only when we apply perturbations. This seems to imply that using perturbations can somehow accelerate the interior point method procedure or yield better conditioning. Except for these particular problems, the average value of μ_λ^k is of order 10^{-4} . For detailed data, see Table D.1.

As for randomly generated problems, we also tracked and compared the differences between initial bases obtained from Algorithms 6.1 and 6.2. We use the same relative difference measure (see Footnote 17). The average difference is about 40%, but there are 9 problems¹⁹ with relative difference less than 10%. Algorithm 6.1 is no better than Algorithm 6.2 for these problems. Generally, for small problems with small relative differences between bases, the simplex iterations are quite similar; for large problems, even small relative difference can yield quite different simplex iterations (such as for SEBA and STOCFOR2). The dis-

¹⁸This is because some problems have very large components in the right hand side b with $\max(b) > 10^3$. For these problems, even when $\mu_\lambda^k > 10^{-3}$, the relative residual may already be less than 10^{-6} and this causes numerical problems when trying to decrease μ_λ^k further. There are five problems of this kind, AGG3, FORPLAN, GROW7, ISRAEL and SHARE1B, and we have marked those problems by * in Table D.1.

¹⁹AFIRO, AGG3, GROW7, ISREAL, SC50B, SCFXM2, SCFXM3, SEBA and STOCFOR2.

appointing small relative difference of initial bases may be the result of inappropriate initial perturbations, improper shrinking speed of perturbations or ill-conditioning.

7 Conclusions and future work

We have proposed the use of controlled perturbations for improving active set prediction capabilities of IPMs for LP. The perturbations are chosen so as to slightly enlarge the feasible set in the hope that the central path of the perturbed problems passes through or close to the original solution set when the perturbed barrier parameter is not too small. Our approach solves a (sequence of) perturbed problems using a standard primal-dual path-following method and predicts using cut-off, the optimal active set of the original problem on the way. We have provided theoretical and preliminary numerical evidence that this approach to active-set prediction for IPMs looks promising in that the perturbed problems are not being solved to high accuracy before the original optimal active set can be accurately predicted and that the perturbations help with the accuracy and speed of the activity prediction for the original solution set.

There are several issues remaining for full validation of the proposed approach; such as the choice of the initial perturbations which we currently set to a fixed small value that we then adjust, but that may be more suitably set to some problem-dependent value. At present, we have used cut-off to predict the original optimal active set when solving the perturbed problems (PD_λ) ; we plan to explore other suitable techniques for the prediction such as the *identification function* proposed originally for nonlinear programming [12]. Note that indicators [9] are not suitable for our purposes as they can only predict the perturbed optimal active-set when calculated in the context of an IPM applied to (PD_λ) . Finally, a large-scale implementation and testing of the perturbed approach and prediction is needed to complete our numerical experiments for LP. From a theoretical point of view, it would be important to show polynomial complexity of a safeguarded (say, long-step [40]) variant of Algorithm 6.1; this seems achievable since one can think of Algorithm 6.1 as a standard primal-dual path-following IPM applied to solving (to some accuracy) a sequence of LP problems (PD_λ) , and so each solve of a (PD_λ) could be shown to have polynomial complexity.

There are several interesting/important areas that may benefit from the application of the controlled perturbations approach for active-set prediction for IPMs. For example, it may prove useful for improving *warmstarting capabilities* [11, 37] of IPMs. Furthermore, it could be applied to the *Homogeneous Self-Dual (HSD) embedding model* [43, 41] for an LP, which is a very useful re-formulation that allows assessing whether the given LP problem has a solution, as well as finding this solution, depending on whether some auxiliary variables are active or inactive at the solution of the HSD problem. Thus early and accurate activity prediction for HSD models could save significant computational effort especially if the original LP problem is infeasible/unbounded. Note that the HSD model needs to solve a monotone LCP problem and so one would need to first extend our controlled perturbations approach to such problems, which seems plausible. We have not explored the use of our active-set prediction strategy in conjunction with *iterative linear solvers* for IPMs [22]; predicting the active constraints sooner may help alleviate the powerful effect that the ill-conditioning (of the IPM subproblem) has on the performance of iterative solvers for this subproblem close to the optimal solution of the LP. Extending our activity prediction proposal to the *convex quadratic programming* case is another potential future research direction.

Acknowledgement. We are grateful to Nick Gould for useful discussions and insights. We also thank three anonymous referees for instructive comments that have improved the quality of the paper.

References

- [1] A. Altman and J. Gondzio. Regularized symmetric indefinite systems in interior point methods for linear and quadratic optimization. *Optimization Methods and Software*, 11(1-4):275–302, 1999.
- [2] M. S. Bazaraa, J. J. Jarvis, and H. D. Sherali. Special simplex implementations and optimality conditions. In *Linear Programming and Network Flows*, pages 201–257. John Wiley & Sons, 2009.
- [3] H. Benson and D. Shanno. An exact primal—dual penalty method approach to warmstarting interior-point methods for linear programming. *Computational Optimization and Applications*, 38(3):371–399, 2007.
- [4] M. Berkelaar, K. Eikland, and P. Notebaert. *lp_solve : Open source (Mixed-Integer) Linear Programming system*.

- [5] C. Cartis and N. I. M. Gould. Finding a point in the relative interior of a polyhedron. Technical Report RAL 2006-016, Rutherford Appleton Laboratory, 2006.
- [6] J. Castro and J. Cuesta. Existence, uniqueness, and convergence of the regularized primal-dual central path. *Operations Research Letters*, 38(5):366–371, 2010.
- [7] R. W. Cottle, J.-S. Pang, and R. E. Stone. *The Linear Complementarity Problem*. SIAM, 2009.
- [8] V. Demiguel, M. P. Friedlander, F. J. Nogales, and S. Scholtes. A two-sided relaxation scheme for mathematical programs with equilibrium constraints. *SIAM Journal on Optimization*, 16(2):587–609, 2005.
- [9] A. S. El-Bakry, R. Tapia, and Y. Zhang. A study of indicators for identifying zero variables in interior point methods. *SIAM Review*, 36(1):45–72, 1994.
- [10] A. Engau, M. F. Anjos, and A. Vannelli. A primal-dual slack approach to warmstarting interior-point methods for linear programming. In *Operations Research and Cyber-Infrastructure*, volume 47, pages 195–217. Springer, 2009.
- [11] A. Engau, M. F. Anjos, and A. Vannelli. On interior-point warmstarts for linear and combinatorial optimization. *SIAM Journal on Optimization*, 20(4):1828–1861, 2010.
- [12] F. Facchinei, A. Fischer, and C. Kanzow. On the accurate identification of active constraints. *SIAM Journal on Optimization*, 9(2):14–32, 1998.
- [13] M. Ferris, O. Mangasarian, and S. J. Wright. *Linear Programming with Matlab*. SIAM, 2007.
- [14] R. M. Freund. A potential-function reduction algorithm for solving a linear program directly from an infeasible “warm start”. *Mathematical Programming*, 52(1-3):441–466, 1991.
- [15] R. M. Freund. Theoretical efficiency of a shifted-barrier-function algorithm for linear programming. *Linear Algebra and its Applications*, 152:19 – 41, 1991.
- [16] R. M. Freund. An infeasible-start algorithm for linear programming whose complexity depends on the distance from the starting point to the optimal solution. *Annals of Operations Research*, 62(1):29–57, 1996.
- [17] P. E. Gill, W. Murray, M. A. Saunders, J. Tomlin, and M. H. Wright. On projected Newton barrier methods for linear programming and an equivalence to Karmarkar’s projective method. *Mathematical Programming*, 36(2):183–209, 1986.
- [18] J. Gondzio. Warm start of the primal-dual method applied in the cutting-plane scheme. *Mathematical Programming*, 83(1-3):125–143, 1998.
- [19] J. Gondzio and A. Grothey. A new unblocking technique to warmstart interior point methods based on sensitivity analysis. *SIAM Journal on Optimization*, 19(3):1184–1210, 2008.
- [20] J. Gondzio and J.-P. Vial. Warm start and ϵ -subgradients in a cutting plane scheme for block-angular linear programs. *Computational Optimization and Applications*, 14(1):17–36, 1999.
- [21] O. Güler, D. den Hertog, C. Roos, T. Terlaky, and T. Tsuchiya. Degeneracy in interior point methods for linear programming: a survey. *Annals of Operations Research*, 46-47:107–138, 1993.
- [22] G. Jacek. Interior point methods 25 years later. *European Journal of Operational Research*, 218:587–601, 2012.
- [23] N. Karmarkar and K. Ramakrishnan. Computational results of an interior point algorithm for large scale linear programming. *Mathematical Programming*, 52(1-3):555–586, 1991.
- [24] O. Mangasarian and J. Ren. New improved error bounds for the linear complementarity problem. *Mathematical Programming*, 66(2):241–255, 1994.

- [25] K. A. McShane, C. L. Monma, and D. Shanno. An implementation of a primal-dual interior point method for linear programming. *ORSA Journal on Computing*, 1(2):70–83, 1989.
- [26] S. Mehrotra. Finite termination and superlinear convergence in primal-dual methods. Technical Report 91-13, Northwestern University, 1991.
- [27] S. Mehrotra. On the implementation of a primal-dual interior point method. *SIAM Journal on Optimization*, 2(4):575–601, 1992.
- [28] S. Mehrotra and Y. Ye. Finding an interior point in the optimal face of linear programs. *Mathematical Programming*, 62(1-3):497–515, 1993.
- [29] J. E. Mitchell. An interior point column generation method for linear programming using shifted barriers. *SIAM Journal on Optimization*, 4(2):423–440, 1994.
- [30] J. L. Morales. A numerical study of limited memory BFGS methods. *Applied Mathematics Letters*, 15(4):481–487, 2002.
- [31] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, 2006.
- [32] C. Oberlin and S. J. Wright. Active set identification in nonlinear programming. *SIAM Journal on Optimization*, 17(2):577–605, 2006.
- [33] J.-S. Pang. Error bounds in mathematical programming. *Mathematical Programming*, 79(1):299–332, 1997.
- [34] R. Polyak. Modified barrier functions (theory and methods). *Mathematical Programming*, 54(1-3):177–222, 1992.
- [35] M. Saunders and J. Tomlin. Solving regularized linear programs using barrier methods and KKT systems. Technical Report SOL 96-4, Department of Operations Research, Stanford University, 1996.
- [36] G. Sierksma. *Linear and Integer Programming: Theory and Practice*. CRC Press, second edition, 2001.
- [37] A. Skajaa, E. Andersen, and Y. Ye. Warmstarting the homogeneous and self-dual interior point method for linear and conic quadratic problems. *Mathematical Programming Computation*, 5:1–25, 2013.
- [38] K. Tone. An active-set strategy in an interior point method for linear programming. *Mathematical Programming*, 59(1-3):345–360, 1993.
- [39] P. J. Williams. *Effective finite termination procedures in interior-point methods for linear programming*. PhD thesis, Department of Computational and Applied Mathematics, Rice University, 1998.
- [40] S. J. Wright. *Primal-Dual Interior-Point Methods*. SIAM, 1997.
- [41] X. Xu, P.-F. Hung, and Y. Ye. A simplified homogeneous and self-dual linear programming algorithm and its implementation. *Annals of Operations Research*, 62(1):151–171, 1996.
- [42] Y. Ye. On the finite convergence of interior-point algorithms for linear programming. *Mathematical Programming*, 57(1):325–335, 1992.
- [43] Y. Ye, M. J. Todd, and S. Mizuno. An $\mathcal{O}(n!)$ -iteration homogeneous and self-dual linear programming algorithm. *Mathematics of Operations Research*, 19(1):53–67, 1994.
- [44] E. Yildirim and S. J. Wright. Warm-start strategies in interior-point methods for linear programming. *SIAM Journal on Optimization*, 12(3):782–810, 2002.
- [45] Y. Zhang. Solving large-scale linear programs by interior-point methods under the Matlab Environment. *Optimization Methods and Software*, 10(1):1–31, 1998.

A Proof of Lemma 4.1

An error bound for an optimization problem bounds the distance from a given point to the solution set in terms of a residual function [33]. In this section, we first formulate an LP problem as a monotone Linear Complementarity Problem (LCP) and then apply a global error bound for the monotone LCP to the reformulated LP problem in order to derive an error bound for the LP, and so prove Lemma 4.1.

By setting $s = c - A^\top y$ and $y = y^+ - y^-$, where $y^+ = \max(y, 0)$ and $y^- = -\min(y, 0)$, the first order optimality conditions (2.2) with $\lambda = 0$ for (PD) can be reformulated as

$$\begin{aligned} Ax - b &\geq 0, \quad -Ax + b \geq 0, \\ c - A^\top y^+ + A^\top y^- &\geq 0, \\ x^\top (c - A^\top y^+ + A^\top y^-) &= 0, \\ x &\geq 0, \quad y^+ \geq 0, \quad y^- \geq 0. \end{aligned} \tag{A.1}$$

Let

$$M = \begin{bmatrix} 0 & -A^\top & A^\top \\ A & 0 & 0 \\ -A & 0 & 0 \end{bmatrix}, \quad z = \begin{bmatrix} x \\ y^+ \\ y^- \end{bmatrix} \quad \text{and} \quad q = \begin{bmatrix} c \\ -b \\ b \end{bmatrix}, \tag{A.2}$$

where A , b and c are (PD) problem data and $(x, y, s) \in \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^n$. Then finding a solution of (A.1) is equivalent to solving the following problem,

$$Mz + q \geq 0, \quad z \geq 0, \quad z^\top (Mz + q) = 0, \tag{A.3}$$

where M , q and z are defined in (A.2), and z is considered to be the vector of variables.

Lemma A.1 (PD) is equivalent to the LCP in (A.3) with M and q defined in (A.2), namely,

1. If (x, y^+, y^-) is a solution of the LCP (A.3), then (x, y, s) is a (PD) solution, where $y = y^+ - y^-$ and $s = c - A^\top y$.
2. If (x, y, s) is a (PD) solution, then (x, y^+, y^-) is a solution of the LCP (A.3).

Next we show that our LP problem can be viewed as a monotone LCP [7].

Lemma A.2 The matrix M , defined in (A.2), is positive semidefinite, and so (A.3) is a monotone LCP.

Proof. For all $v = (v_1, v_2, v_3)$, where $v_1 \in \mathbb{R}^n$, $v_2 \in \mathbb{R}^m$ and $v_3 \in \mathbb{R}^m$, $v^\top Mv = v_2^\top Av_1 - v_3^\top Av_1 - v_1^\top A^\top v_2 + v_1^\top A^\top v_3 = 0$, since $v_2^\top Av_1 = (v_2^\top Av_1)^\top = v_1^\top A^\top v_2$ and $v_3^\top Av_1 = (v_3^\top Av_1)^\top = v_1^\top A^\top v_3$. Thus M is positive semidefinite. \square

A global error bound for a monotone LCP [24] is given next.

Lemma A.3 (Mangasarian and Ren [24, Corollary 2.2]) Let z be any point away from the solution set of the monotone LCP (A.3) and z^* be the closest solution of (A.3) to z under the Euclidean norm $\|\cdot\|$. Then $r(z) + w(z)$ is a global error bound for (A.3), namely,

$$\|z - z^*\| \leq \tau(r(z) + w(z)),$$

where τ is some problem-dependent constant, independent of z and z^* , and

$$r(z) = \|z - (z - Mz - q)_+\| \quad \text{and} \quad w(z) = \|(-Mz - q, -z, z^\top (Mz + q))_+\|. \tag{A.4}$$

Lemma A.4 Given the monotone LCP (A.3) with M and q defined in (A.2), let (x, y^+, y^-) be any point away from the solution set of this problem and $(x^*, (y^*)^+, (y^*)^-)$ be the closest solution of this LCP to (x, y^+, y^-) under the Euclidean norm $\|\cdot\|$. Then we have

$$\|(x, y^+, y^-) - (x^*, (y^*)^+, (y^*)^-)\| \leq \tau(r(x, y^+, y^-) + w(x, y^+, y^-)),$$

where τ is some problem-dependent constant, independent of (x, y^+, y^-) and $(x^*, (y^*)^+, (y^*)^-)$,

$$r(x, y^+, y^-) = \|(\min\{x, c - A^\top y\}, \min\{y^+, Ax - b\}, \min\{y^-, b - Ax\})\|, \tag{A.5}$$

and

$$w(x, y^+, y^-) = \|(-(c - A^\top y), b - Ax, Ax - b, -x, -y^+, -y^-, c^\top x - b^\top y)_+\|, \tag{A.6}$$

and where $y = y^+ - y^-$.

Proof. Substituting (A.2) into (A.4) and noting that $u - (u - v)_+ = \min\{u, v\}$ for any u, v vectors, we have

$$r(x, y^+, y^-) = \left\| \begin{pmatrix} x - (x - (c - A^T(y^+ - y^-)))_+, y^+ - (y^+ - (Ax - b))_+, y^- - (y^- - (b - Ax))_+ \end{pmatrix} \right\|,$$

and

$$w(x, y^+, y^-) = \left\| \begin{pmatrix} -(c - A^T(y^+ - y^-)), b - Ax, Ax - b, -x, -y^+, -y^-, c^T x - b^T(y^+ - y^-) \end{pmatrix}_+ \right\|.$$

Recalling $y = y^+ - y^-$, (A.5) and (A.6) follow directly from the above equations. \square

Theorem A.5 (Error bound for LP) *Let $(x, y, s) \in \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^n$ where $s = c - A^T y$. Then there exist a (PD) solution (x^*, y^*, s^*) ²⁰ and problem-dependent constants τ_p and τ_d , independent of (x, y, s) and (x^*, y^*, s^*) , such that*

$$\|x - x^*\| \leq \tau_p (r(x, y) + w(x, y)) \quad \text{and} \quad \|s - s^*\| \leq \tau_d (r(x, y) + w(x, y)),$$

where

$$r(x, y) = \left\| \begin{pmatrix} \min\{x, s\}, \min\{y^+, Ax - b\}, \min\{y^-, -Ax + b\} \end{pmatrix} \right\|, \quad (\text{A.7})$$

and

$$w(x, y) = \left\| \begin{pmatrix} -s, b - Ax, Ax - b, -x, c^T x - b^T y \end{pmatrix}_+ \right\|, \quad (\text{A.8})$$

and where $y^+ = \max\{y, 0\}$ and $y^- = -\min\{y, 0\}$.

Proof. Consider the monotone LCP (A.3) with M and q defined in (A.2) and $z = (x, y^+, y^-)$. Let $z^* = (x^*, (y^*)^+, (y^*)^-)$ be the closest solution to z in the solution set of this LCP. From Lemma A.1, (x^*, y^*, s^*) with $y^* = (y^*)^+ - (y^*)^-$ and $s^* = c - A^T y^*$ is a (PD) solution. From $(y^+, y^-) \geq 0$, $s = c - A^T y$ and Lemma A.4, we have

$$\|(x, y^+, y^-) - (x^*, (y^*)^+, (y^*)^-)\| \leq \tau(r(x, s) + w(x, s)),$$

where $r(x, s)$ and $w(x, s)$ are defined in (A.7) and (A.8), respectively. This and norm properties give

$$\max(\|x - x^*\|, \|y^+ - (y^*)^+\|, \|y^- - (y^*)^-\|) \leq \tau(r(x, s) + w(x, s)),$$

and so letting $\tau_p = \tau$, we deduce $\|x - x^*\| \leq \tau_p(r(x, s) + w(x, s))$. Since $s^* = c - A^T y^*$, we also have

$$\|s - s^*\| \leq \|A^T\| \|y - y^*\| \leq \|A^T\| (\|y^+ - (y^*)^+\| + \|y^- - (y^*)^-\|) \leq \tau_d(r(x, s) + w(x, s)),$$

where $\tau_d = 2\tau\|A\|$. \square

Proof of Lemma 4.1. Since $(x, y, s) \in \mathcal{F}_\lambda^0$, (2.1) gives $Ax = b$ and $A^T y + s = c$. Then the result follows directly from Theorem A.5. \square

B Proof of Lemma 5.4

Theorem 3.3 shows that we are able to preserve the optimal strict complementarity partition after perturbing the problems if the original (PD) has a unique and nondegenerate solution. Actually, we can take a step further and show that then (PD_λ) will also have a unique and nondegenerate solution.

Theorem B.1 *Assume (2.6) holds and the (PD) problems have a unique and nondegenerate solution (x^*, y^*, s^*) . Let \mathcal{A} and \mathcal{I} denote the corresponding optimal active and inactive sets. Then there exists $\hat{\lambda} = \hat{\lambda}(A, b, c, x^*, s^*) > 0$ such that the perturbed problems (PD_λ) with $0 \leq \|\lambda\| < \hat{\lambda}$ have a unique and nondegenerate solution and the optimal active set is the same as that of the original (PD) problems.*

²⁰Note that we may lose the property that this is the closest solution to the given point.

Proof. We consider the equivalent perturbed problem (2.3). From Theorem 3.3, we know there exists a $\hat{\lambda}(A, b, c, x^*, s^*) > 0$ such that (2.3) with $0 \leq \|\lambda\| < \hat{\lambda}$ has a strictly complementary solution $(\hat{p}, \hat{y}, \hat{q})$ with the same optimal active and inactive sets \mathcal{A} and \mathcal{I} , namely we have

$$\hat{p}_{\mathcal{I}} > 0, \quad \hat{p}_{\mathcal{A}} = 0, \quad \hat{q}_{\mathcal{A}} > 0, \quad \text{and} \quad \hat{q}_{\mathcal{I}} = 0,$$

and also

$$A_{\mathcal{I}} \hat{p}_{\mathcal{I}} = b_{\lambda}. \quad (\text{B.1})$$

Next we are about to show that $(\hat{p}, \hat{y}, \hat{q})$ is the unique solution of (2.3). Assume there exists another solution $\bar{p} \neq \hat{p}$. Then $(\bar{p}, \hat{y}, \hat{q})$ satisfies the optimality conditions (2.4). From the complementarity equations (the third term) in (2.4) and $\hat{q}_{\mathcal{A}} > 0$ we have $\bar{p}_{\mathcal{A}} = 0 = \hat{p}_{\mathcal{A}}$. Then we have $A_{\mathcal{I}} \bar{p}_{\mathcal{I}} = b_{\lambda}$. It follows from this and (B.1) that $A_{\mathcal{I}}(\bar{p}_{\mathcal{I}} - \hat{p}_{\mathcal{I}}) = 0$. As the (PD) solution is unique and nondegenerate, we must have $|\mathcal{I}| = m$ and $\text{rank}(A_{\mathcal{I}}) = m$, namely, $A_{\mathcal{I}}$ is nonsingular, which implies $\bar{p}_{\mathcal{I}} = \hat{p}_{\mathcal{I}}$. Then (2.3) has a unique and nondegenerate primal solution, which also implies unique and nondegenerate dual solution. \square

To prove Lemma 5.4, we also need the following series of useful lemmas.

Lemma B.2 (Farkas' Lemma [2, Lemma 5.1]) *One and only one of the following two systems has a solution:*

$$\begin{aligned} \text{System 1: } & Tw \geq 0 \quad \text{and} \quad b^T w < 0, \\ \text{System 2: } & T^T y = b \quad \text{and} \quad y \geq 0, \end{aligned}$$

where $T \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, $w \in \mathbb{R}^n$ and $y \in \mathbb{R}^m$.

Lemma B.3 *Given $i \in \{1, \dots, n\}$, the following system*

$$\begin{cases} y + Ax \geq 0 \\ x - A^T y \geq 0 \quad \text{and} \quad x_i - A_i^T y > 0 \\ (x, y) \geq 0 \end{cases}$$

always has a solution, where $A \in \mathbb{R}^{m \times n}$, $x \in \mathbb{R}^n$, $y \in \mathbb{R}^m$ and A_i is the i^{th} column of A .

Proof. Without losing generality, we can choose $i = 1$. Partition x and A as $x = [x_1 \quad \bar{x}^T]^T$ and $A = [A_1 \quad \bar{A}]$, where $\bar{x} = [x_2 \quad \dots \quad x_n]^T$ and $\bar{A} = [A_2 \quad \dots \quad A_n]$.

We need to prove the following system has a solution

$$\begin{cases} y + A_1 x_1 + \bar{A} \bar{x} \geq 0 \\ -\bar{A}^T y + \bar{x} \geq 0 \\ y \geq 0 \\ x_1 \geq 0 \\ \bar{x} \geq 0 \\ A_1^T y - x_1 < 0 \end{cases}. \quad (\text{B.2})$$

From Lemma B.2, we know (B.2) has a solution if and only if

$$\begin{cases} \begin{bmatrix} I_m & -\bar{A} & I_m & 0 & 0 \\ A_1^T & 0 & 0 & 1 & 0 \\ \bar{A}^T & I_{n-1} & 0 & 0 & I_{n-1} \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \\ u_5 \end{bmatrix} = \begin{bmatrix} A_1 \\ -1 \\ 0 \end{bmatrix}, \\ (u_1, u_2, u_3, u_4, u_5) \geq 0 \end{cases}, \quad (\text{B.3})$$

has no solution, where $u_1 \in \mathbb{R}^m$, $u_2 \in \mathbb{R}^{n-1}$, $u_3 \in \mathbb{R}^m$, $u_4 \in \mathbb{R}$ and $u_5 \in \mathbb{R}^{n-1}$.

Assume (B.3) has a solution $(u_1, u_2, u_3, u_4, u_5) \geq 0$. Then we get

$$u_1 - \bar{A}u_2 - A_1 = -u_3 \leq 0, \quad (\text{B.4a})$$

$$A_1^T u_1 = -1 - u_4 < 0, \quad (\text{B.4b})$$

$$\bar{A}^T u_1 = -u_2 - u_5 \leq 0. \quad (\text{B.4c})$$

Multiplying both sides of (B.4a) by $u_1^T \geq 0$, we have

$$u_1^T u_1 - (\bar{A}^T u_1)^T u_2 - A_1^T u_1 = -u_1^T u_3.$$

From (B.4b), (B.4c) and the nonnegativity of the variables, we know

$$u_1^T u_1 - (\bar{A}^T u_1)^T u_2 - \bar{A}_1^T u_1 > 0 \quad \text{but} \quad -u_1^T u_3 \leq 0.$$

Thus (B.3) has no solution, which implies (B.2) has a solution. \square

Lemma B.4 *The system*

$$\begin{cases} y + Ax \geq 0 \\ x - A^T y > 0 \\ (x, y) \geq 0 \end{cases} \quad (\text{B.5})$$

always has a solution, where $A \in \mathbb{R}^{m \times n}$, $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^m$.

Proof. From Lemma B.3, we know for any $i \in \{1, \dots, n\}$, there exists $(x^i, y^i) \geq 0$ where $x^i \in \mathbb{R}^n$ and $y^i \in \mathbb{R}^m$, such that

$$\begin{cases} y^i + Ax^i \geq 0, \\ x^i - A^T y^i \geq 0 \quad \text{and} \quad x_i^i - A_i^T y^i > 0. \end{cases} \quad (\text{B.6})$$

Set $\hat{x} = \sum_{i=1}^n x^i \geq 0$ and $\hat{y} = \sum_{i=1}^n y^i \geq 0$. Then from (B.6), we have

$$\hat{y} + A\hat{x} = \sum_{i=1}^n y^i + A\left(\sum_{i=1}^n x^i\right) = \sum_{i=1}^n (y^i + Ax^i) \geq 0,$$

and

$$\hat{x} - A^T \hat{y} = \sum_{i=1}^n (x^i - A^T y^i) = \begin{bmatrix} (x_1^1 - A_1^T y^1) + (x_1^2 - A_1^T y^2) + \dots + (x_1^n - A_1^T y^n) \\ \vdots \\ (x_n^1 - A_n^T y^1) + (x_n^2 - A_n^T y^2) + \dots + (x_n^n - A_n^T y^n) \end{bmatrix} > 0.$$

\square

Lemma B.5 *The system*

$$\begin{cases} y + Ax > 0 \\ x - A^T y \geq 0 \\ (x, y) \geq 0 \end{cases} \quad (\text{B.7})$$

always has a solution, where $A \in \mathbb{R}^{m \times n}$, $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^m$.

Proof. Replace A in Lemma B.4 by $-A^T$. \square

Lemma B.6 *The system*

$$\begin{cases} y + Ax > 0 \\ x - A^T y > 0 \\ (x, y) \geq 0. \end{cases}$$

always has a solution, where $A \in \mathbb{R}^{m \times n}$, $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^m$.

Proof. From Lemmas B.4 and B.5, we know there exist (\hat{x}, \hat{y}) and (\tilde{x}, \tilde{y}) such that (B.5) and (B.7) hold respectively. Set $\bar{x} = \hat{x} + \tilde{x}$ and $\bar{y} = \hat{y} + \tilde{y}$ and deduce

$$\bar{y} + A\bar{x} = \underbrace{(\hat{y} + A\hat{x})}_{\geq 0} + \underbrace{(\tilde{y} + A\tilde{x})}_{> 0} > 0 \quad \text{and} \quad \bar{x} - A^T\bar{y} = \underbrace{(\hat{x} - A^T\hat{y})}_{> 0} + \underbrace{(\tilde{x} - A^T\tilde{y})}_{\geq 0} > 0.$$

□

Proof of Lemma 5.4. Assume \mathcal{A} and \mathcal{I} are the optimal active and inactive sets at the unique solution (x^*, y^*, s^*) of (PD). Then from (5.8), we have

$$\epsilon(A, b, c) = \min \left(\min_{i \in \mathcal{I}}(x_i^*), \min_{i \in \mathcal{A}}(s_i^*) \right). \quad (\text{B.8})$$

From Theorem B.1, we know there exists a $\hat{\lambda} = \hat{\lambda}(A, b, c, x^*, s^*)$ such that (PD_λ) with $0 \leq \|\lambda\| < \hat{\lambda}$ has a unique and nondegenerate solution and \mathcal{A} and \mathcal{I} are the optimal active and inactive sets. Since $(\hat{p}, \hat{y}, \hat{q})$ defined in (3.7) is a solution of (2.3), $(x_\lambda^*, y_\lambda^*, s_\lambda^*) = (\hat{p} - \lambda, \hat{y}, \hat{q} - \lambda)$ is a solution of (PD_λ) and also unique, with \mathcal{A} and \mathcal{I} being the optimal active and inactive sets. This and (4.6) give

$$\epsilon(A, b_\lambda, c_\lambda) = \min \left(\min_{i \in \mathcal{I}}(\hat{x}_i + \lambda_i), \min_{i \in \mathcal{A}}(\hat{s}_i + \lambda_i) \right). \quad (\text{B.9})$$

From (3.7), recalling that $\hat{p} = \hat{x} + \lambda$ and $\hat{q} = \hat{s} + \lambda$, we have

$$(\hat{x}_\mathcal{I} + \lambda_\mathcal{I}) - x_\mathcal{I}^* = \lambda_\mathcal{I} + A_\mathcal{I}^{-1} A_\mathcal{A} \lambda_\mathcal{A} \quad \text{and} \quad (\hat{s}_\mathcal{A} + \lambda_\mathcal{A}) - s_\mathcal{A}^* = \lambda_\mathcal{A} - (A_\mathcal{I}^{-1} A_\mathcal{A})^T \lambda_\mathcal{I}.$$

This, (B.8) and (B.9) give us that $\epsilon(A, b_\lambda, c_\lambda) > \epsilon(A, b, c)$, provided

$$\begin{cases} \lambda_\mathcal{I} + A_\mathcal{I}^{-1} A_\mathcal{A} \lambda_\mathcal{A} > 0 \\ \lambda_\mathcal{A} - (A_\mathcal{I}^{-1} A_\mathcal{A})^T \lambda_\mathcal{I} > 0 \\ \lambda_\mathcal{I}, \lambda_\mathcal{A} \geq 0 \end{cases}. \quad (\text{B.10})$$

It remains to find a solution of (B.10) whose norm is less than $\hat{\lambda}$. From Lemma B.6, we know (B.10) always has a solution, say $\bar{\lambda}$. Since (B.10) is homogeneous, $\frac{\hat{\lambda}}{2\|\bar{\lambda}\|} \bar{\lambda}$ is also a solution, and $\|\frac{\hat{\lambda}}{2\|\bar{\lambda}\|} \bar{\lambda}\| < \hat{\lambda}$. Without losing generality, we denote this solution as $\bar{\lambda}$. Furthermore, (B.10) holds for all λ with $0 < \lambda = \alpha \bar{\lambda} < \bar{\lambda}$ where $\alpha \in (0, 1)$. □

C An Active-set Prediction Procedure

Note that in Procedure C.1, \mathcal{A}^k is the predicted active set, \mathcal{I}^k , the predicted inactive set and $\mathcal{Z}^k = \{1, 2, \dots, n\} \setminus (\mathcal{A}^k \cup \mathcal{I}^k)$, the set of all undetermined indices at the k^{th} iteration.

D Results for crossover to simplex on selected Netlib problems

From the left to the right, we give the name of the test problems, number of equality constraints, number of variables, the value of duality gap μ_λ^K when we terminate the (perturbed) Algorithm 6.1, the value of duality gap μ^K when we terminate the (unperturbed) Algorithm 6.2, number of IPM iterations, the relative difference (see Footnote 17 on Page 20) between two bases generated from Algorithms 6.1 and 6.2, simplex iterations for Algorithm 6.1 and the simplex iterations for Algorithm 6.2. Since the algorithm without perturbations is terminated at the same IPM iteration as Algorithm 6.1, we show only the number of IPM iterations for the latter. Problems on which Algorithm 6.1 loses are marked in bold font. ‘—’ means the simplex solver fails for a particular test problem.

Procedure C.1 An Active-set Prediction Procedure

Initialise: $A^0 = \mathcal{I}^0 = \emptyset$ and $\mathcal{Z}^0 = \{1, 2, \dots, n\}$.
At k^{th} **iteration**, $k > 1$,
for $i = 1, \dots, n$ **do**
 if $i \in \mathcal{Z}^k$ **then**
 if the threshold test (6.3) is satisfied for iterations $k - 1$ and k **then**
 $\mathcal{A}^k = \mathcal{A}^k \cup \{i\}$ and $\mathcal{Z}^k = \mathcal{Z}^k \setminus \{i\}$;
 else
 $\mathcal{I}^k = \mathcal{I}^k \cup \{i\}$ and $\mathcal{Z}^k = \mathcal{Z}^k \setminus \{i\}$.
 end if
 if $i \in \mathcal{A}^k$ and the threshold test is not satisfied **then**
 $\mathcal{A}^k = \mathcal{A}^k \setminus \{i\}$ and $\mathcal{Z}^k = \mathcal{Z}^k \cup \{i\}$;
 end if
 if $i \in \mathcal{I}^k$ and the threshold test is satisfied **then**
 $\mathcal{I}^k = \mathcal{I}^k \setminus \{i\}$ and $\mathcal{Z}^k = \mathcal{Z}^k \cup \{i\}$.
 end if
end if
end for

Table D.1: Crossover to simplex test on a selection of Netlib problems.

Probs	m	n	μ_{λ}^K	μ^K	IPM Iter	Basis Diff	splxIter Per	splxIter Unp
25FV47	798	1854	9.38e-04	1.34e-03	35	0.15	4193	6951
ADLITTLE	55	137	3.79e-04	2.23e-04	16	0.45	18	119
AFIRO	27	51	3.68e-04	6.84e-06	11	0.07	9	9
AGG3*	516	758	9.05e-02	6.39e-02	25	0.07	112	123
BLEND	74	114	6.55e-04	7.21e-04	10	0.37	35	59
BNL1	632	1576	5.41e-04	1.96e-02	28	0.31	1583	1632
BRANDY	149	259	4.83e-04	1.09e-03	18	0.38	76	278
CZPROB	737	3141	4.00e-04	1.67e-04	56	0.77	106	1822
E226	220	469	6.13e-04	6.98e-04	18	0.54	428	319
FIT1D	1050	2075	4.81e-04	2.00e-04	22	0.39	53	787
FIT1P	1026	2076	5.55e-04	4.04e-04	20	0.36	259	760
FORPLAN*	157	485	4.67e-03	1.33e-02	29	0.45	119	341
GROW7*	420	581	4.56e-02	5.56e-02	15	0.06	226	190
ISRAEL*	174	316	5.39e-02	1.87e-02	32	0.01	164	143
KB2	52	77	3.83e-04	1.15e-02	21	0.24	44	27
SC50A	49	77	1.64e-04	6.42e-05	10	0.12	22	27
SC50B	48	76	5.37e-04	1.59e-04	8	0.00	37	37
SCAGR7	129	185	2.11e-04	2.80e-04	18	0.43	21	65
SCFXM1	322	592	6.19e-04	4.35e-04	24	0.38	188	413
SCFXM2	644	1184	5.48e-04	1.04e-03	27	0.02	690	672
SCFXM3	966	1776	8.73e-04	8.77e-04	28	0.01	1062	1074
SCRS8	485	1270	7.65e-04	1.42e-03	29	0.39	320	315
SCSD1	77	760	5.54e-04	5.54e-04	7	0.95	125	214
SCSD6	147	1350	5.86e-04	5.91e-04	8	0.94	346	411
SCSD8	397	2750	4.88e-04	5.10e-04	11	0.90	366	965
SCTAP1	300	660	5.31e-04	3.64e-03	19	0.31	114	179
SCTAP2	1090	2500	6.81e-04	2.30e-07	21	0.43	145	344
SCTAP3	1480	3340	7.81e-04	1.11e-07	22	0.48	54	451
SEBA	1029	1550	5.92e-04	3.09e-04	23	0.04	43	70
SHARE1B*	112	248	2.87e-03	8.18e-02	27	0.24	176	204
SHARE2B	96	162	3.19e-04	3.90e-04	14	0.41	57	126
SHIP04L	356	2162	6.55e-04	2.92e-04	27	0.61	13	215
SHIP08L	688	4339	6.55e-04	5.41e-04	29	0.81	441	1056
SHIP08S	416	2171	6.34e-04	3.56e-04	26	0.76	70	—
SHIP12S	466	2293	2.12e-04	2.74e-05	33	0.71	18	541
STAIR	362	544	6.56e-04	1.11e-02	16	0.29	292	294
STOCFOR2	2157	3045	5.63e-04	4.74e-05	39	0.08	1213	796