THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Implementation of an Interior Point Method with Basis Preconditioning

Citation for published version:

Schork, L & Gondzio, J 2020, 'Implementation of an Interior Point Method with Basis Preconditioning', *Mathematical Programming Computation*, vol. 12, pp. 603–635. <https://doi.org/10.1007/s12532-020-00181-8>

Digital Object Identifier (DOI):

[10.1007/s12532-020-00181-8](https://doi.org/10.1007/s12532-020-00181-8)

Link:

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

Document Version:

Peer reviewed version

Published In:

Mathematical Programming Computation

General rights

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

Take down policy

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

Implementation of an Interior Point Method with Basis Preconditioning*

Lukas Schork[†]

Jacek Gondzio[‡]

School of Mathematics, University of Edinburgh, Edinburgh EH9 3FD, Scotland, UK

Technical Report ERGO 18-014, 20 September 2018

Abstract

The implementation of a linear programming interior point solver is described that is based on iterative linear algebra. The linear systems are preconditioned by a basis matrix, which is updated from one interior point iteration to the next to bound the entries in a certain tableau matrix. The update scheme is based on simplex-type pivot operations and is implemented using linear algebra techniques from the revised simplex method. An initial basis is constructed by a crash procedure after a few interior point iterations. The basis at the end of the interior point solve provides the starting basis for a crossover method which recovers a basic solution to the linear program. Results of a computational study on a diverse set of medium to large-scale problems are discussed.

1 Introduction

The purpose of this paper is to describe a new linear programming (LP) interior point solver called IPX that is based on iterative linear algebra. IPX uses an approach originally implemented in [2, 16] that employs a basis matrix as preconditioner for the linear systems. In the previous works the basis was chosen as the first m linearly independent columns of an $m \times n$ matrix A , after ordering the columns by scaling factors from the current interior point iterate. Although this yields a perfect preconditioner close to the solution of a nondegenerate LP model, the success of the method was limited in practice.

The core of IPX is a new method for basis selection. A starting basis is constructed after a few interior point iterations and subsequently updated in an attempt to maintain “maximum volume”. The maximum volume criterion is known from rank revealing matrix factorizations (see, for example, [17]) and bounds the condition number of the preconditioned matrices in terms of the dimension of the LP model [3, 20]. Because finding a maximum volume basis is computationally impractical for large-scale problems, the implementation is based on a heuristic that efficiently finds a basis of comparable quality. The required linear algebra operations are those from the revised simplex method, for which established and highly optimized computational techniques exist.

Section 2 sketches the interior point algorithm and discusses the effect of inexact step directions, which arise from iterative linear algebra. Section 3 presents the basis preconditioner and the update scheme for maintaining a basis matrix. Section 4 is concerned with

*Supported by Google Research Award “Fast interior point method for linear programming problems”.

[†]L.Schork@ed.ac.uk

[‡]J.Gondzio@ed.ac.uk

the very first interior point iterations and the construction of a starting basis. Section 5 presents the crossover method for recovering a vertex solution from the final basis. Computational results on a representative set of LP models are discussed in Section 6. This test set is used throughout the paper to provide statistics of IPX runs and to illustrate the effect of algorithmic decisions.

The following notations and conventions are used. When $\mathbf{d}$ is a vector, D is the diagonal matrix with entries d_j on the diagonal. Index sets are understood to be ordered and are denoted by calligraphic letters; $\mathcal{J}_k$ is the k -th index and $|\mathcal{J}|$ the number of indices in the set. When D is a diagonal matrix, $D_{\mathcal{B}}$ is the principal submatrix indexed by $\mathcal{B}$. For the rectangular matrix A , $A_{\mathcal{B}}$ is composed of the corresponding columns of A . A basis $\mathcal{B}$ is an index set such that $A_{\mathcal{B}}$ is square and nonsingular. Associated with $\mathcal{B}$ is the nonbasic set $\mathcal{N}$ with the obvious definition.

IPX internally stores the LP model in the following primal-dual form

$$\begin{aligned} & \underset{\mathbf{x}, \mathbf{x}^l, \mathbf{x}^u}{\text{minimize}} && \mathbf{c}^T \mathbf{x} \\ & \text{subject to} && A\mathbf{x} = \mathbf{b}, \\ & && \mathbf{x} - \mathbf{x}^l = \mathbf{l}, \\ & && \mathbf{x} + \mathbf{x}^u = \mathbf{u}, \\ & && \mathbf{x}^l, \mathbf{x}^u \geq \mathbf{0}, \end{aligned} \tag{1a}$$

and

$$\begin{aligned} & \underset{\mathbf{y}, \mathbf{z}^l, \mathbf{z}^u}{\text{maximize}} && \mathbf{b}^T \mathbf{y} + \mathbf{l}^T \mathbf{z}^l - \mathbf{u}^T \mathbf{z}^u \\ & \text{subject to} && A^T \mathbf{y} + \mathbf{z}^l - \mathbf{z}^u = \mathbf{c}, \\ & && \mathbf{z}^l, \mathbf{z}^u \geq \mathbf{0}, \end{aligned} \tag{1b}$$

where A is an $m \times (n + m)$ matrix whose rightmost m columns form the identity matrix. The first n and the last m components of $\mathbf{x}$ are termed “structural” and “slack” variables, respectively, although the objective coefficients of slack variables do not need to be zero. Entries of $-\mathbf{l}$ and $\mathbf{u}$ are allowed to be infinity, in which case the corresponding dual variable is fixed at zero and its term is dropped from the dual objective. We define the index sets

$$\mathcal{L} = \{j \mid l_j > -\infty\}, \quad \mathcal{U} = \{j \mid u_j < \infty\}.$$

A variable x_j is termed “free” if it has no finite bound and “fixed” if its lower and upper bounds are equal. After preprocessing only structural variables can be free and only slack variables can be fixed.

IPX accepts user input in the more convenient form

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \bar{\mathbf{c}}^T \mathbf{x} \\ & \text{subject to} && \bar{A}\mathbf{x} \{=, \leq, \geq\} \bar{\mathbf{b}}, \\ & && \bar{\mathbf{l}} \leq \mathbf{x} \leq \bar{\mathbf{u}}, \end{aligned}$$

where $\bar{A}$ is an $\bar{m} \times \bar{n}$ matrix. The user model is transformed into computational form with optional dualization. If the input becomes the primal problem, then $A = \begin{bmatrix} \bar{A} & I_{\bar{m}} \end{bmatrix}$, where $I_{\bar{m}}$ denotes the identity matrix of dimension $\bar{m}$; and if it becomes the dual problem, then $A = \begin{bmatrix} \bar{A}^T & (-I_{\bar{n}})_{\mathcal{J}} & I_{\bar{n}} \end{bmatrix}$, where $\mathcal{J}$ is the set of variables with finite lower and upper bounds. The reason for dualization is that the linear algebra implementation is optimized for matrices A with (many) more structural columns than rows. If the user does not make an explicit choice, the model is dualized if $\bar{m} > 2\bar{n}$.

2 Inexact Interior Point Method

A primal-dual interior point method (IPM) simultaneously solves (1a) and (1b) by generating a sequence of iterates for $(\mathbf{x}, \mathbf{x}^l, \mathbf{x}^u, \mathbf{y}, \mathbf{z}^l, \mathbf{z}^u)$ that keeps $(\mathbf{x}^l, \mathbf{x}^u, \mathbf{z}_{\mathcal{L}}^l, \mathbf{z}_{\mathcal{U}}^u)$ positive while driving the residuals

$$\mathbf{r}^b = \mathbf{b} - A\mathbf{x}, \quad (2a)$$

$$\mathbf{r}^l = \mathbf{l} - \mathbf{x} + \mathbf{x}^l, \quad (2b)$$

$$\mathbf{r}^u = \mathbf{u} - \mathbf{x} - \mathbf{x}^u, \quad (2c)$$

$$\mathbf{r}^c = \mathbf{c} - A^T \mathbf{y} - \mathbf{z}^l + \mathbf{z}^u \quad (2d)$$

and the complementarity measure

$$\mu = \frac{1}{n_\mu} \left(\sum_{j \in \mathcal{L}} x_j^l z_j^l + \sum_{j \in \mathcal{U}} x_j^u z_j^u \right)$$

to zero. Here $n_\mu = |\mathcal{L}| + |\mathcal{U}|$ is the number of finite bounds. If the iterate is feasible, i. e. (2a)–(2d) are all zero, then $n_\mu \mu = f_p - f_d$ is the gap between the primal and dual objective values. The usual requirement is to obtain 8-digit accuracy in the objective, leading to the termination criterion

$$|f_p - f_d| \leq 10^{-8} (1 + 0.5|f_p + f_d|). \quad (3)$$

By the behaviour of the algorithm, primal and dual feasibility are practically always satisfied when the objective criterion is reached.

The most time consuming part of each interior point iteration is computing the step direction by solving a small number of linear systems of the form

$$\begin{bmatrix} A^T & & I & -I \\ A & & & \\ I & -I & & \\ I & & I & \\ & Z^l & X^l & \\ & Z^u & X^u & \end{bmatrix} \begin{pmatrix} \Delta \mathbf{x} \\ \Delta \mathbf{y} \\ \Delta \mathbf{x}^l \\ \Delta \mathbf{x}^u \\ \Delta \mathbf{z}^l \\ \Delta \mathbf{z}^u \end{pmatrix} = \begin{pmatrix} \mathbf{r}^c \\ \mathbf{r}^b \\ \mathbf{r}^l \\ \mathbf{r}^u \\ \mathbf{s}^l \\ \mathbf{s}^u \end{pmatrix} \quad (4)$$

for certain vectors $\mathbf{s}^l$ and $\mathbf{s}^u$. IPX implements a version of Mehrotra's method [15], which composes the step from a predictor and a corrector direction. Mehrotra's method is widely considered to be the most efficient method using two linear system solves per iteration. For implementations based on Cholesky factorization it is common practice to compute more than one corrector direction to reduce the iteration count of the IPM [7]. We do not consider this technique here because iterative linear algebra does not offer large savings when solving multiple systems with the same matrix.

After eliminating $\Delta \mathbf{x}^l$, $\Delta \mathbf{x}^u$, $\Delta \mathbf{z}^l$ and $\Delta \mathbf{z}^u$ from (4), the linear system becomes

$$\begin{bmatrix} G & A^T \\ A & 0 \end{bmatrix} \begin{pmatrix} \Delta \mathbf{x} \\ -\Delta \mathbf{y} \end{pmatrix} = \begin{pmatrix} \mathbf{r}^a \\ \mathbf{r}^b \end{pmatrix} \quad (5)$$

for a certain vector $\mathbf{r}^a$ and G the diagonal matrix with entries $g_j = z_j^l/x_j^l + z_j^u/x_j^u$. Note that g_j is zero if x_j is a free variable. We define diagonal matrices D and H with entries

$$d_j = \begin{cases} (g_j)^{-1/2} & \text{if } g_j \neq 0, \\ 1 & \text{if } g_j = 0, \end{cases} \quad h_j = \begin{cases} 1 & \text{if } g_j \neq 0, \\ 0 & \text{if } g_j = 0, \end{cases}$$

so that (5) can be rescaled to

$$\begin{bmatrix} H & DA^T \\ AD & 0 \end{bmatrix} \begin{pmatrix} D^{-1}\Delta\mathbf{x} \\ -\Delta\mathbf{y} \end{pmatrix} = \begin{pmatrix} D\mathbf{r}^a \\ \mathbf{r}^b \end{pmatrix}. \quad (6)$$

The d_j are called “scaling factors” in interior point terminology and the 2×2 matrices are said to have “KKT form”. The KKT system can either be solved as it is written or can be further reduced by eliminating (parts of) $\Delta\mathbf{x}$. In any case, an iterative linear solver only provides an approximate solution to (5) or (6) and the accuracy of the solution needs to be controlled by some means. The iterative methods implemented in IPX compute a solution of the form

$$\begin{bmatrix} H & DA^T \\ AD & 0 \end{bmatrix} \begin{pmatrix} D^{-1}\Delta\mathbf{x} \\ -\Delta\mathbf{y} \end{pmatrix} = \begin{pmatrix} D\mathbf{r}^a \\ \mathbf{r}^b \end{pmatrix} + \begin{pmatrix} \boldsymbol{\delta} \\ \mathbf{0} \end{pmatrix},$$

where a residual $\boldsymbol{\delta}$ occurs only in the first block equation. As termination criterion for the linear solver it is required that

$$\|\boldsymbol{\delta}\|_\infty \leq \tau := \varepsilon_\tau \sqrt{\mu}, \quad (7)$$

where ε_τ is a problem independent parameter. The criterion is motivated theoretically because it maintains boundedness of $\|D^{-1}\Delta\mathbf{x}\|$ in terms of $\sqrt{\mu}$, which is key to the convergence of the IPM (see, for example, [22, Chapter 6]). Choosing a smaller tolerance increases the computational cost per linear system but reduces the number of interior point iterations as the step directions become more accurate. The effect on the total computation time is quite moderate, however. On our test set varying ε_τ between its default value 0.3 and within the range $[0.05, 0.5]$ changed the geometric mean of the IPX runtimes by up to 8%.

Given an approximate solution to the KKT system, the order in which the remaining step components are recovered determines to which of the block equations in (4) the residual propagates. Solving the dual feasibility equations exactly has shown to be necessary for the robustness of the IPM. We therefore compute

$$\begin{aligned} \Delta\mathbf{x}^l &= -\mathbf{r}^l + \Delta\mathbf{x}, \\ \Delta\mathbf{x}^u &= \mathbf{r}^u - \Delta\mathbf{x}, \\ \Delta z_j^l &= (s_j^l - z_j^l \Delta x_j^l) / x_j^l && \text{for all } j \text{ for which } z_j^l / x_j^l < z_j^u / x_j^u, \\ \Delta z_j^u &= (s_j^u - z_j^u \Delta x_j^u) / x_j^u && \text{for all } j \text{ for which } z_j^l / x_j^l \geq z_j^u / x_j^u, \\ \Delta z_j^l &= r_j^c - (A^T \Delta\mathbf{y})_j + \Delta z_j^u && \text{for all } j \text{ for which } z_j^l / x_j^l \geq z_j^u / x_j^u, \\ \Delta z_j^u &= -r_j^c + (A^T \Delta\mathbf{y})_j + \Delta z_j^l && \text{for all } j \text{ for which } z_j^l / x_j^l < z_j^u / x_j^u. \end{aligned}$$

By this choice the first four block equations in (4) are solved exactly and the residual is shifted between the last two blocks. If a variable has two finite bounds and is active at its lower bound in the solution, the residual eventually occurs in the equation $z_j^l \Delta x_j^l + x_j^l \Delta z_j^l = s_j^l$. Since in this case $x_j^l \rightarrow 0$ but $x_j^u \rightarrow (x_j^u)^* > 0$, adjusting Δz_j^l to satisfy dual feasibility causes a smaller residual in the right-hand side to (4) than adjusting Δz_j^u .

3 Basis Preconditioning

In advanced interior point iterations the KKT matrix usually becomes very ill conditioned due to a wide spread of the diagonal entries of G . This ill conditioning is a major obstacle for iterative linear algebra and requires effective preconditioning to obtain an acceptable iteration count of the linear solver. The situation is further complicated because LP problems arise from a variety of applications and therefore have different numerical properties. The

motivation for IPX was to develop a general-purpose LP solver based on iterative linear algebra, which is not tailored to a particular class of problems.

We shall see that preconditioning the linear systems by a careful choice of basis matrix yields a robust method with good overall efficiency. Different variants of basis preconditioning were described in [2, 3, 16]. In this paper a new variant specific for IPMs is presented that takes advantage of free variables while reducing the KKT system to a smaller, positive definite one.

3.1 Preconditioned CR Method

Assume first that the LP model does not contain free variables. Then G has a zero-free diagonal, $D = G^{-1/2}$ and the KKT system (5) can be reduced to positive definite normal equations

$$AD^2A^T\Delta\mathbf{y} = \mathbf{r}^b - AD^2\mathbf{r}^a =: \mathbf{r}.$$

Basis preconditioning transforms this system using a scaled basis matrix $A_{\mathcal{B}}D_{\mathcal{B}}$ into

$$(A_{\mathcal{B}}D_{\mathcal{B}})^{-1}AD^2A^T(A_{\mathcal{B}}D_{\mathcal{B}})^{-T}\Delta\mathbf{u} = (A_{\mathcal{B}}D_{\mathcal{B}})^{-1}\mathbf{r}. \quad (8)$$

In our setting the transformation is carried out explicitly, meaning that the iterative solver computes a solution for $\Delta\mathbf{u}$ rather than $\Delta\mathbf{y}$.

IPX uses the Conjugate Residual (CR) method [18, Algorithm 6.20] for solving positive definite linear systems. The CR method is a Krylov subspace method originally proposed in [13] for the symmetric indefinite case. Compared to the Conjugate Gradient method it requires one more vector update per iteration, but has shown to reach the termination criterion in fewer iterations and smaller computation time.

Matrix-vector products with the transformed normal matrix

$$C := (A_{\mathcal{B}}D_{\mathcal{B}})^{-1}AD^2A^T(A_{\mathcal{B}}D_{\mathcal{B}})^{-T} = I_m + (A_{\mathcal{B}}D_{\mathcal{B}})^{-1}A_{\mathcal{N}}D_{\mathcal{N}}^2A_{\mathcal{N}}^T(A_{\mathcal{B}}D_{\mathcal{B}})^{-T} \quad (9)$$

are implemented through a matrix-vector product with $A_{\mathcal{N}}D_{\mathcal{N}}^2A_{\mathcal{N}}^T$ and two linear system solves with a factorization of $A_{\mathcal{B}}D_{\mathcal{B}}$. Neither $A_{\mathcal{N}}D_{\mathcal{N}}^2A_{\mathcal{N}}^T$ nor C are formed explicitly.

After computing an approximate solution for $\Delta\mathbf{u}$, a solution to the KKT system is recovered from

$$\begin{aligned} \Delta\mathbf{y} &= (A_{\mathcal{B}}D_{\mathcal{B}})^{-T}\Delta\mathbf{u}, \\ \Delta\mathbf{x}_{\mathcal{N}} &= D_{\mathcal{N}}^2(\mathbf{r}_{\mathcal{N}}^a + A_{\mathcal{N}}^T\Delta\mathbf{y}), \\ \Delta\mathbf{x}_{\mathcal{B}} &= A_{\mathcal{B}}^{-1}(\mathbf{r}^b - A_{\mathcal{N}}\Delta\mathbf{x}_{\mathcal{N}}). \end{aligned}$$

By definition of $\Delta\mathbf{x}$, a residual originating from the iterative method occurs only in the basic components of the first block equation in (6). That residual is

$$\begin{aligned} \delta_{\mathcal{B}} &= D_{\mathcal{B}}^{-1}\Delta\mathbf{x}_{\mathcal{B}} - D_{\mathcal{B}}A_{\mathcal{B}}^T\Delta\mathbf{y} - D_{\mathcal{B}}\mathbf{r}_{\mathcal{B}}^a \\ &= D_{\mathcal{B}}^{-1}A_{\mathcal{B}}^{-1}(\mathbf{r}^b - A_{\mathcal{N}}\Delta\mathbf{x}_{\mathcal{N}}) - \Delta\mathbf{u} - D_{\mathcal{B}}\mathbf{r}_{\mathcal{B}}^a \\ &= D_{\mathcal{B}}^{-1}A_{\mathcal{B}}^{-1}\mathbf{r}^b - D_{\mathcal{B}}^{-1}A_{\mathcal{B}}^{-1}A_{\mathcal{N}}D_{\mathcal{N}}^2(\mathbf{r}_{\mathcal{N}}^a + A_{\mathcal{N}}^T\Delta\mathbf{y}) - \Delta\mathbf{u} - D_{\mathcal{B}}\mathbf{r}_{\mathcal{B}}^a \\ &= (A_{\mathcal{B}}D_{\mathcal{B}})^{-1}\mathbf{r}^b - (A_{\mathcal{B}}D_{\mathcal{B}})^{-1}AD^2\mathbf{r}^a - (I_m + (A_{\mathcal{B}}D_{\mathcal{B}})^{-1}A_{\mathcal{N}}D_{\mathcal{N}}^2A_{\mathcal{N}}^T(A_{\mathcal{B}}D_{\mathcal{B}})^{-T})\Delta\mathbf{u} \\ &= (A_{\mathcal{B}}D_{\mathcal{B}})^{-1}\mathbf{r} - C\Delta\mathbf{u}. \end{aligned}$$

Therefore, to satisfy the accuracy requirement (7), the CR method is terminated when the infinity norm of the residual in (8) becomes smaller than τ .

When G has zeros on the diagonal the KKT system can still be reduced to normal equations if all free variables are basic. This requirement is satisfied by the initial basis in

IPX and is maintained throughout; hence let $\mathcal{B} = \mathcal{B}_0 \cup \mathcal{B}_1$, where $\mathcal{B}_0$ are the indices of all free variables. Then (6) can be permuted to

$$\begin{bmatrix} I_{\mathcal{N}} & & D_{\mathcal{N}} A_{\mathcal{N}}^T \\ & \begin{bmatrix} I_{\mathcal{B}_1} & \\ & 0 \end{bmatrix} & D_{\mathcal{B}} A_{\mathcal{B}}^T \\ A_{\mathcal{N}} D_{\mathcal{N}} & A_{\mathcal{B}} D_{\mathcal{B}} & 0 \end{bmatrix} \begin{pmatrix} D_{\mathcal{N}}^{-1} \Delta \mathbf{x}_{\mathcal{N}} \\ D_{\mathcal{B}}^{-1} \Delta \mathbf{x}_{\mathcal{B}} \\ -\Delta \mathbf{y} \end{pmatrix} = \begin{pmatrix} D_{\mathcal{N}} \mathbf{r}_{\mathcal{N}}^a \\ D_{\mathcal{B}} \mathbf{r}_{\mathcal{B}}^a \\ \mathbf{r}^b \end{pmatrix}.$$

Transforming this system using the scaled basis matrix yields

$$\begin{bmatrix} I_{\mathcal{N}} & & D_{\mathcal{N}} A_{\mathcal{N}}^T A_{\mathcal{B}}^{-T} D_{\mathcal{B}}^{-1} \\ & \begin{bmatrix} I_{\mathcal{B}_1} & \\ & 0 \end{bmatrix} & I_{\mathcal{B}} \\ D_{\mathcal{B}}^{-1} A_{\mathcal{B}}^{-1} A_{\mathcal{N}} D_{\mathcal{N}} & I_{\mathcal{B}} & 0 \end{bmatrix} \begin{pmatrix} D_{\mathcal{N}}^{-1} \Delta \mathbf{x}_{\mathcal{N}} \\ D_{\mathcal{B}}^{-1} \Delta \mathbf{x}_{\mathcal{B}} \\ -\Delta \mathbf{u} \end{pmatrix} = \begin{pmatrix} D_{\mathcal{N}} \mathbf{r}_{\mathcal{N}}^a \\ D_{\mathcal{B}} \mathbf{r}_{\mathcal{B}}^a \\ D_{\mathcal{B}}^{-1} A_{\mathcal{B}}^{-1} \mathbf{r}^b \end{pmatrix},$$

where $\Delta \mathbf{u} = D_{\mathcal{B}} A_{\mathcal{B}}^T \Delta \mathbf{y}$ as before. Notice that the components of $\Delta \mathbf{u}$ correspond to basic variables. It follows from the second block equation and $D_{\mathcal{B}_0}$ being the identity matrix that the components of $-\Delta \mathbf{u}$ corresponding to free variables are given by $\mathbf{r}_{\mathcal{B}_0}^a$ and can be eliminated. Also, the components of $\Delta \mathbf{x}_{\mathcal{B}}$ corresponding to free variables can be removed from the linear system and can be computed from the third block equation once the remaining components of $D^{-1} \Delta \mathbf{x}$ are known. Let $P = \begin{bmatrix} I_{\mathcal{B}_1} & 0 \end{bmatrix}$ have m columns. The resulting linear system is

$$\begin{bmatrix} I_{\mathcal{N}} & & D_{\mathcal{N}} A_{\mathcal{N}}^T A_{\mathcal{B}}^{-T} D_{\mathcal{B}}^{-1} P^T \\ & I_{\mathcal{B}_1} & I_{\mathcal{B}_1} \\ P D_{\mathcal{B}}^{-1} A_{\mathcal{B}}^{-1} A_{\mathcal{N}} D_{\mathcal{N}} & I_{\mathcal{B}_1} & 0 \end{bmatrix} \begin{pmatrix} D_{\mathcal{N}}^{-1} \Delta \mathbf{x}_{\mathcal{N}} \\ D_{\mathcal{B}_1}^{-1} \Delta \mathbf{x}_{\mathcal{B}_1} \\ -P \Delta \mathbf{u} \end{pmatrix} = \begin{pmatrix} D_{\mathcal{N}} \mathbf{r}_{\mathcal{N}}^a \\ D_{\mathcal{B}_1} \mathbf{r}_{\mathcal{B}_1}^a \\ P D_{\mathcal{B}}^{-1} A_{\mathcal{B}}^{-1} \mathbf{r}^b \end{pmatrix} - \begin{pmatrix} \boldsymbol{\xi} \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix},$$

where

$$\boldsymbol{\xi} = D_{\mathcal{N}} A_{\mathcal{N}}^T A_{\mathcal{B}}^{-T} D_{\mathcal{B}}^{-1} \begin{pmatrix} \mathbf{0} \\ \mathbf{r}_{\mathcal{B}_0}^a \end{pmatrix} = D_{\mathcal{N}} A_{\mathcal{N}}^T A_{\mathcal{B}}^{-T} \begin{pmatrix} \mathbf{0} \\ \mathbf{r}_{\mathcal{B}_0}^a \end{pmatrix}.$$

Because the upper left 2×2 block now has a zero-free diagonal, the system can be reduced to normal equations

$$\begin{aligned} (I_{\mathcal{B}_1} + P D_{\mathcal{B}}^{-1} A_{\mathcal{B}}^{-1} A_{\mathcal{N}} D_{\mathcal{N}}^2 A_{\mathcal{N}}^T A_{\mathcal{B}}^{-T} D_{\mathcal{B}}^{-1} P^T) (P \Delta \mathbf{u}) = \\ - P D_{\mathcal{B}}^{-1} A_{\mathcal{B}}^{-1} A_{\mathcal{N}} D_{\mathcal{N}} (D_{\mathcal{N}} \mathbf{r}_{\mathcal{N}}^a - \boldsymbol{\xi}) - D_{\mathcal{B}_1} \mathbf{r}_{\mathcal{B}_1}^a + P D_{\mathcal{B}}^{-1} A_{\mathcal{B}}^{-1} \mathbf{r}^b \end{aligned} \quad (10)$$

and can be solved for $P \Delta \mathbf{u}$. The solution for $\Delta \mathbf{u}$ is completed by inserting the components corresponding to free variables given by $-\mathbf{r}_{\mathcal{B}_0}^a$.

The dimension of (10) is m minus the number of free variables. In the implementation the CR method always operates on vectors of dimension m , in which components corresponding to free variables are initialized to zero and are reset to zero after each matrix-vector product with C defined in (9). The only actual change to the code was to replace the right-hand side of (8) by that of (10), scattered into a full size vector. The result is a concise handling of free variables, which implicitly reduces the dimension of the linear systems.

3.2 Maintaining a Basis Matrix

The aim of basis preconditioning is to make the transformed normal matrix (9) well conditioned, since then we expect fast convergence of the linear solver. Clearly, the condition

number of C depends on the choice of the basis. While finding an optimal basis in that sense seems to be impossible without enumerating all bases [3], a criterion that can be targeted in practice is to find a basis $\mathcal{B}$ that satisfies¹

$$\max_{p,q} |D_{\mathcal{B}}^{-1} A_{\mathcal{B}}^{-1} A_{\mathcal{N}} D_{\mathcal{N}}|_{pq} \leq \rho \quad (11)$$

for a parameter $\rho \geq 1$. Such a basis, called “ ρ -maximum volume basis”, exists for any matrix AD of full row rank and bounds the spectrum of C in the interval $[1, 1 + \rho^2 mn]$, see [3, 20]. Compared to the previous approaches for basis preconditioning in IPMs [2, 16], the advantage of the maximum volume concept is that it does not rely on having m large and n small scaling factors. Due to degeneracy the latter assumption is rarely satisfied in practice.

A ρ -maximum volume basis is obtained through pivot operations in the scaled tableau matrix $D_{\mathcal{B}}^{-1} A_{\mathcal{B}}^{-1} A_{\mathcal{N}} D_{\mathcal{N}}$ starting from an arbitrary basis [8, 12]. A generic pivot scheme is given in Algorithm 1. In each iteration any entry that is larger than ρ in absolute value can be chosen as pivot. Because each basis update increases $|\det(A_{\mathcal{B}} D_{\mathcal{B}})|$ (see [8]), the algorithm terminates in a finite number of iterations.

Algorithm 1 Maxvolume

Input: basis $\mathcal{B}$, parameter $\rho \geq 1$.

- 1: **loop**
 - 2: Choose (p, q) such that $|D_{\mathcal{B}}^{-1} A_{\mathcal{B}}^{-1} A_{\mathcal{N}} D_{\mathcal{N}}|_{pq} > \rho$. If no such (p, q) exists, then stop.
 - 3: $\mathcal{B}_p = \mathcal{N}_q$
 - 4: **end loop**
-

It will be seen that the number of pivot operations for updating the basis from one interior point iteration to the next is quite small in practice. The difficulty in implementing Algorithm 1 is finding a pivot element in line 2. Because the tableau matrix is only available implicitly through operations with $A_{\mathcal{B}}^{-1}$, searching systematically for an entry that is larger than ρ can be expensive. In particular, testing if (11) is satisfied would require to compute all tableau entries, which is impractical for large-scale problems.

To make the approach practical we relax the target and only *attempt* to find a ρ -maximum volume basis. The idea is to perform pivot operations in the scaled tableau matrix as long as large entries are readily found, and to terminate the method when the pivot search becomes costly. The pseudocode of the update procedure is given in Algorithm 2.

Assume first that the parameter *nslices* is 1. Then w_j is the sum of entries in column j of the scaled tableau matrix if $j \in \mathcal{N}$. The algorithm chooses the index with maximum such weight as candidate and computes the corresponding column of the scaled tableau matrix. If its maximum entry is larger than ρ in absolute value, the basis is updated. In this case the tableau row needs to be computed as well to update the column sums (line 15). After computing *maxskip*+1 candidate columns without finding a pivot, the algorithm terminates.

When *nslices* > 1, then in each iteration of the outer loop only a subset of rows of the tableau matrix contributes to the column weights. This slicing of the tableau matrix decreases the chance that when a column has large positive and negative entries that these cancel out in the column sum; and it makes the algorithm adaptive to the number of rows of A .

IPX runs Algorithm Maxvolume Heuristic at the beginning of each interior point iteration with the new scaling matrix D and the basis from the previous iteration as starting basis. Free and fixed variables are excluded from the pivot search and remain basic and nonbasic,

¹For a matrix A the expression $|A|$ is meant componentwise.

Algorithm 2 Maxvolume Heuristic

Input: basis $\mathcal{B}$, parameters $(\rho, nslices, maxskip)$ such that $\rho \geq 1$, $1 \leq nslices \leq m$, $maxskip \geq 0$.

```
1: for  $s = 0$  to  $nslices - 1$  do
2:   Let  $\mathbf{u} \in \mathbb{R}^m$ ,  $\mathbf{w} \in \mathbb{R}^{n+m}$ .
3:   for  $i = 1$  to  $m$  do                                     // set row mask
4:      $u_i = 1$  if  $\text{mod}(i, nslices) = s$ 
5:      $u_i = 0$  if  $\text{mod}(i, nslices) \neq s$ 
6:   end for
7:    $\mathbf{w}_{\mathcal{N}}^T = \mathbf{u}^T D_{\mathcal{B}}^{-1} A_{\mathcal{B}}^{-1} A_{\mathcal{N}} D_{\mathcal{N}}$ ,  $\mathbf{w}_{\mathcal{B}} = \mathbf{0}$            // initialize column weights
8:    $skipped = 0$                                            // count "skipped" columns
9:   while  $skipped \leq maxskip$  do
10:     $j = \arg \max_k |w_k|$                                    // choose candidate column
11:    Compute  $\mathbf{v} = D_{\mathcal{B}}^{-1} A_{\mathcal{B}}^{-1} A_j d_j$ .
12:    if  $\|\mathbf{v}\|_{\infty} > \rho$  then
13:       $p = \arg \max_i |v_i|$ 
14:       $\mathcal{B}_p = j$ 
15:      Update  $\mathbf{w}$  .
16:    else                                                 // no pivot found in column  $j$ 
17:       $skipped = skipped + 1$ 
18:      Flag column  $j$  to be excluded from pivot search.
19:    end if
20:  end while
21: end for
```

respectively. The default parameters are $nslices = 5 + \lfloor m/10000 \rfloor$ and $maxskip = 10$, but they have shown little effect on the number of basis updates and the quality of the basis preconditioner.

3.3 Empirical Tests

To investigate the effect of the parameter ρ we consider two LP models, **pds-20** and **nug20**, which have similar dimensions but very different characteristics. The operation counts and computation times for updating the basis and running the iterative solver are reported in Table 1. IPM iterations prior to the start of basis preconditioning are included in the total IPM time but not in the figures for the CR method. These iterations and the starting bases were the same for all values of ρ .

In both examples a smaller ρ systematically reduced the iteration count and runtime of the CR method, but increased the number of updates and the time for finding the basis. A larger set of models confirmed that values between 2.0 and 4.0 lead to similar and close-to-optimal total runtime if the correlations are the same as in Table 1. Therefore IPX uses a problem independent default of $\rho = 2.0$.

The behaviour is not always that clear, however. It turned out that sparsity in the tableau matrix also has large effect on the number of CR iterations. While A is almost always a sparse matrix in real life, for some models $A_{\mathcal{B}}^{-1}$ and $A_{\mathcal{B}}^{-1} A_{\mathcal{N}}$ are sparse as well for all relevant bases; such models are called “hypersparse” in simplex community [4, 9]. On the one hand hypersparsity enables orders of magnitude speedup for finding the basis because operations with $A_{\mathcal{B}}^{-1}$, if implemented properly, take far less than order of m time. On the other hand a sparse tableau matrix has shown to lead to fewer CR iterations. This is not surprising because the maximum eigenvalue of C is bounded by $1 + \|D_{\mathcal{B}}^{-1} A_{\mathcal{B}}^{-1} A_{\mathcal{N}} D_{\mathcal{N}}\|_F^2$ (see [20]) and the Frobenius norm becomes small when the tableau is sparse. Both effects are

name (dimension)	ρ	Maxvolume Heuristic		CR method		IPM total
		updates	time	iter	time	time
pds-20 ($m = 12081$, $n = 81163$)	1.1	2793	1.00	4359	1.81	4.68
	2.0	1964	0.90	5129	2.15	4.91
	4.0	1697	0.85	6467	2.77	5.49
	10.0	1548	0.81	10029	4.34	6.98
nug20 ($m = 14098$, $n = 72546$)	1.1	8871	91.52	11567	34.49	133.24
	2.0	6238	66.83	14928	41.45	115.64
	4.0	5091	55.62	19652	53.39	116.04
	10.0	4385	48.80	31485	87.21	143.11

Table 1: Linear algebra operations on two example models. Times are reported in seconds.

seen in Table 1 by comparing the hypersparse **pds-20** to the not-hypersparse **nug20**. When sparsity in the tableau matrix varies significantly between different bases, it can happen that a smaller ρ leads to more CR iterations. We observed such a behaviour on a number of LP models. It is not obvious how to choose a basis that preserves sparsity as far as possible while targeting the maximum volume criterion. This remains a topic for further investigation.

A second comparison with the two example models demonstrates that the heuristically found basis does not degrade effectiveness of the preconditioner compared to a (true) ρ -maximum volume basis. In Figure 1 the CR iteration counts per interior point iteration are plotted for either basis. Finding the maximum volume basis was a very expensive task for **nug20** and is not an option in practice. For both models and all values of ρ the curves closely resemble each other, showing that the heuristically found basis yields an equally good preconditioner. The curves also illustrate the typical behaviour of the CR iteration counts during the course of the IPM: one or more peaks usually occur after the switch to basis preconditioning, and the iterations level out toward the end of the interior point solve.

4 Initial Iterations and Crash Basis

IPX does not use basis preconditioning from the beginning of the interior point solve. At the early iterations the KKT matrix is usually well conditioned by itself and a simpler (and less expensive) preconditioner is equally effective. This is exploited by performing a few “initial iterations” of the IPM in which the CR method is applied with a variant of diagonal scaling. When the method does not converge within $\min(500, 10 + m/20)$ iterations, a starting basis is constructed and the interior point iteration is repeated with basis preconditioning. This section looks at these techniques in more detail.

4.1 Initial Iterations

The KKT matrix in (5) is nonsingular if and only if A has full row rank and the columns of A corresponding to zeros on the diagonal of G are linearly independent. While the first requirement is satisfied after adding slack variables to all constraints, the latter requirement might not be. To overcome the problem at the beginning, zeros on the diagonal of G are replaced by the regularization value $\min\{\{g_j | g_j \neq 0\}, \mu\}$. Regularization can significantly change the solution to the linear system, but using it in the initial iterations has shown to be unproblematic for the convergence of the IPM. The regularization value is chosen to yield

Figure 1: LP models `pds-20` (above) and `nug20` (below) with ρ -maximum volume basis (solid line) and heuristically computed basis (dashed line).

a well conditioned KKT matrix as long as the remaining diagonal entries of G are balanced and μ is sufficiently large. Including μ in the definition guarantees that regularization is eventually reduced to zero even if none of the remaining g_j becomes small (i.e. if all non-free variables are at a bound in the solution).

An additional benefit of regularization is that the KKT system can always be reduced to normal equations

$$AG^{-1}A^T\Delta\mathbf{y} = \mathbf{r}^b - AG^{-1}\mathbf{r}^a = \mathbf{r}. \quad (12)$$

From an approximate solution for $\Delta\mathbf{y}$, the solution components $\Delta\mathbf{x}$ to the KKT system are recovered from

$$\Delta\mathbf{x}_{\mathcal{N}} = G_{\mathcal{N}}^{-1}(\mathbf{r}^a + A^T\Delta\mathbf{y})_{\mathcal{N}}, \quad \Delta\mathbf{x}_{\mathcal{B}} = \mathbf{r}^b - A_{\mathcal{N}}\Delta\mathbf{x}_{\mathcal{N}},$$

where G is the regularized matrix and $\mathcal{B} = \{n+1, \dots, n+m\}$ is the slack basis. It is easily verified that the accuracy requirement (7) is satisfied if

$$\left\| G_{\mathcal{B}}^{1/2}(\mathbf{r} - AG^{-1}A^T\Delta\mathbf{y}) \right\|_{\infty} \leq \tau.$$

In the initial iterations IPX solves (12) by the CR method using a symmetric positive definite matrix M as preconditioner. The method iterates on $\Delta\mathbf{y}$ and requires one operation with M^{-1} per iteration. In an early version of the code diagonal preconditioning was used, where $M = \text{diag}(AG^{-1}A^T)$ was obtained by dropping all off-diagonal entries. The method has been refined for matrices A with “dense” columns; i.e. columns whose nonzero count is much higher than the average. Let A_s and A_d be the sparse and dense part of A , respectively, and G_d and G_s be the corresponding diagonal blocks of G . Then IPX uses

$$M = \text{diag}(A_s G_s^{-1} A_s^T) + A_d G_d^{-1} A_d^T \quad (13)$$

as preconditioner. By treating the second summand as a low-rank update, an inverse representation of M is obtained from the Sherman-Morrison-Woodbury formula and can be computed through the Cholesky factorization of a dense matrix of dimension equal to the number of columns in A_d . For LP models with a small number of dense columns, using (13) is little more expensive than diagonal scaling but often more effective because $A_s G_s^{-1} A_s^T$ is better approximated by its diagonal than $AG^{-1}A^T$. IPX classifies the maximum number of columns as dense such that each column in A_d has more than 40 nonzeros and more than 10 times the number of nonzeros of any column in A_s . If this yields more than 1000 dense columns, then no columns are treated as dense.

4.2 Crash Basis

At the switch to basis preconditioning a starting basis must be determined, given the current interior point iterate and its associated scaling matrix D . To reduce the number of basis updates in the first run of Algorithm Maxvolume Heuristic and to make the linear algebra operations fast, the basis should have the following (often competing) properties:

- The basis matrix is well conditioned.
- The basis matrix and ideally the tableau matrix are sparse.
- The basis is close to a maximum volume basis for the current D .
- All free variables are basic.
- All fixed variables are nonbasic.

Making all free variables basic is required for the reduction of the KKT system to normal equations (Section 3.1) and is always achievable. If the columns corresponding to free variables are linearly dependent, then the model is either dual infeasible or some of the variables are redundant and can be fixed at an arbitrary value. The analogue requirement is to make all fixed variables nonbasic, which allows them to be removed from the optimization entirely. Again, this is always achievable, for if a fixed slack variable cannot be replaced in the basis, then either is the corresponding row of A redundant and can be removed, or the model is primal infeasible.

Finding efficiently a basis that satisfies the last requirement is already nontrivial if most slack variables are fixed. In particular, the obvious method by computing an LU factorization of the structural part of A would be unacceptably expensive for many large-scale problems. Instead, IPX “crashes” a starting basis through the following steps:

Initial guess A set $\mathcal{J}$ of m column indices is constructed as follows:

- (1) If the LP model contains free variables, the first step is computing an incomplete left-looking LU factorization of the corresponding columns of A . In the left-looking method L starts out to be the identity matrix and its strictly lower triangular part is computed one column at a time. Let j be the next free variable, $\hat{\mathbf{a}} = L^{-1}A_j$ and i be such that $|\hat{a}_i|$ is maximum among all entries of $|\hat{\mathbf{a}}|$ whose row has not been pivotal. If $|\hat{a}_i| > 10^{-3}$, variable j is added to $\mathcal{J}$ and the next column of L is composed from the entries of $\hat{\mathbf{a}}/\hat{a}_i$ that are nonzeros in A_j and have not been pivotal. (Restricting the column of L to the nonzero pattern of A_j makes the factorization incomplete.) After processing all free variables, L is discarded but the index set $\mathcal{I}$ of pivot rows is kept.
- (2) In the second and third step columns of A corresponding to fixed and free variables are treated as vacant; i.e. they are treated as being structurally zero in A and AD . The second step adds singleton columns to $\mathcal{J}$ if their entry in AD is sufficiently large. For each $i \notin \mathcal{I}$ the maximum entry in row i of $|AD|$ and the maximum entry that lies in a singleton column are determined. If the singleton entry is nonzero and not smaller than 0.5 times the maximum of the row, its column is added to $\mathcal{J}$ and i is added to $\mathcal{I}$.
- (3) Let A_{33} be the submatrix of A composed of rows $i \notin \mathcal{I}$ and columns $j \notin \mathcal{J}$. The third step extends $\mathcal{I}$ and $\mathcal{J}$ by choosing a structurally independent subset of the columns of A_{33} . Processing in decreasing order of the d_j , the next column from A_{33} is tested for being structurally dependent on the columns from A_{33} already chosen, by computing an alternating augmenting path [6]. If the candidate column is independent, its index is added to $\mathcal{J}$ and the row that was newly matched by the augmenting path is added to $\mathcal{I}$. The method stops at the latest when $|\mathcal{J}| = m$ or when all columns of A_{33} have been processed. It is stopped before if too many candidate columns were structurally dependent, ensuring that the computation time is small.
- (4) Finally $\mathcal{J}$ is completed by adding slack variables for $i \notin \mathcal{I}$.

Initial factorization The task is to find a well conditioned basis matrix that contains as many columns of $\mathcal{J}$ as possible. A right-looking Markowitz LU factorization of $A_{\mathcal{J}}$ is computed with the usual columnwise threshold pivoting. If during the course of the factorization all entries in a column of the active submatrix become smaller than 10^{-3} , the column is immediately removed without choosing a pivot. After completion, the LU factors are padded with unit columns whose row has not been pivotal. Accordingly a preliminary basis $\mathcal{B}$ is obtained composed of indices of $\mathcal{J}$ and indices of slack variables.

Basis repair The resulting basis matrix A_B would be nonsingular in exact arithmetic. It is well known, however, that A_B can have tiny singular values even if no small pivots occur in the LU factorization, and such cases actually happen in practice. It is therefore essential to control the condition number of A_B and to repair numerical singularities if necessary. Because the model is scaled during preprocessing so that the maximum entry of A is bounded by a moderate number, a high condition number of A_B can only be caused by large entries in A_B^{-1} . As described in [11] a rook search is performed for estimating the maximum absolute entry of A_B^{-1} along with its position (p, i) in the matrix. If the entry is larger than 10^5 , then B_p is replaced by the i -th slack variable and the rook search is repeated. During the basis repair A_B typically requires frequent refactorization due to numerical instability in the LU update. If the basis matrix is refactorized, columns are dropped from the active submatrix and replaced by unit columns as in the initial factorization.

Handling free and fixed variables The obtained basis may contain fixed slack variables and may not contain all free variables. These “defects” are corrected by basis updates. For each free variable that is nonbasic a column of the tableau matrix is computed and the variable is pivoted into the basis if it can replace a non-free basic variable. If not, the model is either declared dual infeasible or the free variable is fixed at zero. Likewise, for each fixed (slack) variable in the basis a row of the tableau matrix is computed and searched for an exchange column. If the tableau row is numerically zero, the model is either declared primal infeasible or the constraint corresponding to the slack variable is removed. After correcting all defects, the remaining fixed variables (which are nonbasic now) are excluded from the IPM solve.

The described procedure has been developed through extensive testing and works efficiently on most real-world problems. Steps (i)–(iii) of the initial guess are designed to be fast and to yield a column subset of close-to full rank. Depending on the characteristics of the LP model, each of the three steps may add the majority of columns to $\mathcal{J}$. Computing the starting basis typically accounts for less than 5% of the total IPX runtime. When the computation time is significant, the last step of the procedure often dominates due to the number of pivots for fixed variables. Processing these variables as described is advantageous later, however, because otherwise fixed variables can lead to small step sizes in the IPM and dependent equality constraints can cause dual variables to blow up.

The crash bases are surprisingly close to the bases at the end of the interior point solve. In Table 2 the number of basis updates in relation to m is reported. Note that a value 0.25, for example, means that at least 75% of the starting basis must be final. It is seen that updating the basis during the interior point solve never exhibits the combinatorial nature that can occur in the simplex method.

5 Crossover

A basic solution to (1) is an optimal solution with an associated basis $\mathcal{B}$ such that $z_{\mathcal{B}}^l$ and $z_{\mathcal{B}}^u$ are zero and $x_{\mathcal{N}}$ is at a bound (or zero if the variable is free). In this case $\mathcal{B}$ is said to be an optimal LP basis. Basic solutions are required in a number of applications, most prominently for classical integer programming. Because the basis from the preconditioner need not become optimal close to the solution, the basis-preconditioned IPM requires a crossover step for recovering a basic solution, as any other interior point solver.

For the crossover it is convenient to combine the dual slack variables into $z = z^l - z^u$.

basis updates / m	instances
0.000–0.010	24
0.010–0.025	8
0.025–0.050	6
0.050–0.100	16
0.100–0.250	46
0.250–0.500	22
0.500–1.000	38
1.000–3.350	5

Table 2: Number of basis updates in the IPM starting from crash basis (165 models in total).

A primal-dual solution $(\mathbf{x}, \mathbf{y}, \mathbf{z})$ then satisfies

$$A\mathbf{x} = \mathbf{b}, \quad A^T\mathbf{y} + \mathbf{z} = \mathbf{c}, \quad (14a)$$

$$\mathbf{l} \leq \mathbf{x} \leq \mathbf{u}, \quad (14b)$$

$$z_j \leq 0 \text{ if } x_j > l_j \quad \text{for all } j, \quad (14c)$$

$$z_j \geq 0 \text{ if } x_j < u_j \quad \text{for all } j. \quad (14d)$$

Let us first assume that $(\mathbf{x}, \mathbf{y}, \mathbf{z})$ from the final interior point iteration is an exact primal-dual solution and let $\mathcal{B}$ be any basis. The variables in $\mathcal{N}^+ = \{j \in \mathcal{N} \mid l_j < x_j < u_j\}$ and $\mathcal{B}^+ = \{j \in \mathcal{B} \mid z_j \neq 0\}$ are called primal and dual “superbasic”, respectively. The crossover method removes superbasic variables by two push phases.

The dual push phase manipulates $(\mathbf{y}, \mathbf{z})$ and $\mathcal{B}$ whilst satisfying (14c), (14d) and

$$A_{\mathcal{B}}^T \mathbf{y} + \mathbf{z}_{\mathcal{B}} = \mathbf{c}_{\mathcal{B}},$$

$$A_{\mathcal{N}}^T \mathbf{y} + \mathbf{z}_{\mathcal{N}} = \mathbf{c}_{\mathcal{N}}.$$

In each iteration an $i \in \mathcal{B}^+$ is chosen and z_i is moved toward zero until the adjustment to $(\mathbf{y}, \mathbf{z}_{\mathcal{N}})$ would violate the sign condition for $\mathbf{z}$. The push is complete if z_i reaches zero. Otherwise a basis update is applied exchanging $i \in \mathcal{B}$ with the blocking index $j \in \mathcal{N}$. Because in the latter case z_j has become zero, each iteration reduces the number of dual superbasic variables by 1.

The primal push phase manipulates $\mathbf{x}$ and $\mathcal{B}$ whilst satisfying (14b) and

$$A_{\mathcal{B}} \mathbf{x}_{\mathcal{B}} + A_{\mathcal{N}} \mathbf{x}_{\mathcal{N}} = \mathbf{b}.$$

In each iteration a $j \in \mathcal{N}^+ \cap (\mathcal{L} \cup \mathcal{U})$ is chosen and x_j is moved toward a bound until the adjustment to $\mathbf{x}_{\mathcal{B}}$ would violate the bound constraints. The push is complete if x_j reaches its bound. Otherwise a basis update is applied exchanging $j \in \mathcal{N}$ with the blocking index $i \in \mathcal{B}$. Because in the latter case x_i has been moved to a bound, each iteration reduces the number of primal superbasic variables by 1.

As pointed out in [5], the push phases are equivalent to Megiddo’s algorithm [14] if the initial basis contains a maximum number of variables for which $l_j < x_j < u_j$ and a minimum number for which $z_j \neq 0$. In this case the analysis in [14] proves that each push maintains complementarity of $(\mathbf{x}, \mathbf{y}, \mathbf{z})$ so that the final iterate is indeed a basic solution. It is easy to see, however, that each dual push as stated above maintains complementarity regardless of the initial basis; if moving z_i ($i \in \mathcal{B}^+$) toward zero makes a z_j ($j \in \mathcal{N}^+$) nonzero, the step is blocked immediately because (14c) or (14d) would be violated. Furthermore, because there

exist no dual superbasic variables after the dual push phase completed, the primal phase can move any basic variable x_i without violating complementarity. Hence the above algorithm can be run from any starting basis.

In practice the final iterate from the IPM is neither exactly feasible nor complementary and care must be taken with small pivot elements to prevent the basis matrix from becoming too ill conditioned for finite precision arithmetic. These issues are addressed by the IPX implementation as follows. Before executing the push phase, the final IPM iterate is dropped to an $(\mathbf{x}, \mathbf{y}, \mathbf{z})$ satisfying (14b)–(14d). For each variable either x_j is set to a bound or z_j is set to zero, depending on which perturbation is smaller. The dropping usually increases the residuals in (14a). Throughout the push operations the conditions (14b)–(14d) are maintained exactly by truncating the update to a variable if necessary. A blocking variable is eligible as pivot only if the pivot element is larger than 10^{-5} in absolute value. Among all candidates, the exchange variable is chosen using the two-pass ratio test from the simplex method [10], which allows larger pivot elements by exploiting a feasibility tolerance for (14b)–(14d) with default value 10^{-7} . The combination of the last two techniques was necessary to make the implementation robust against failure due to a numerically singular basis.

IPX naturally uses the final basis from the preconditioner as starting basis for crossover. The number of pushes is determined from the beginning by the number of superbasic variables, whereas the number of pivot operations depends on the order in which the superbasic variables are processed. IPX proceeds in the primal and dual push phase in decreasing and increasing order, respectively, of the scaling factors from the final interior point iteration. We have not found a systematic difference in the number of pivot operations compared to other orderings and find this the most natural choice that makes the computations invariant to a permutation of the variables.

At the end of the push phases $(\mathbf{x}, \mathbf{y}, \mathbf{z})$ is a basic solution to the LP model with perturbed right-hand side and objective, and $\mathcal{B}$ may or may not be an optimal basis for the original problem. If it is not, a reoptimization with the simplex method is required. The main reason for a non-optimal basis is that the pairwise complementarity products $x_j^l z_j^l$ and $x_j^u z_j^u$ can be quite large when the IPM reaches the termination criterion (3), resulting in a large perturbation. To reduce the number of simplex iterations in the clean-up, IPX uses a more stringent termination criterion for the IPM when crossover is requested. In addition to (3) it is required that

$$\max_j \|\delta x_j A_j\|_\infty \leq 10^{-8}(1 + \|(\mathbf{b}, \mathbf{l}_L, \mathbf{u}_U)\|_\infty), \quad (15a)$$

$$\max_j |\delta z_j| \leq 10^{-8}(1 + \|\mathbf{c}\|_\infty), \quad (15b)$$

where δx_j and δz_j are the perturbations to drop the iterate to complementarity. (For each j either δx_j or δz_j is zero.) On 60% of the models in our test set the stricter condition forced at least one extra interior point iteration, and except for some badly scaled models no more than 6 iterations were needed. Hence the additional computation time for the IPM is moderate, in particular because the final iterations are typically fast with basis preconditioning. Using the stricter termination criterion for the IPM, the basis after the push phases was optimal in 150 out of 165 cases.

A conventional IPM implementation would not be able to achieve (15) reliably because the linear systems would become too ill conditioned in the final iterations to be solved to sufficient accuracy. IPX dynamically eliminates variables from the optimization process when the primal is close to its bound or the dual is close to zero. This technique requires a basis to be implemented correctly and allows the IPM to be run to (arbitrary) high accuracy. Details can be found in [19].

$\min(\bar{m}, \bar{n})$	instances
1,036–2,499	16
2,500–4,999	17
5,000–9,999	27
10,000–24,999	41
25,000–49,999	31
50,000–99,999	11
100,000–249,999	19
250,000–499,999	4
500,000–999,999	3
1,000,000–1,439,571	1

Table 3: Dimensions of 170 test problems after presolve.

6 Results

IPX is written in C++ and comprises about 10,000 lines of code for the core algorithm. The factorization of basis matrices and its update is separated from the interior point code and can be provided by an external package. By default the authors’ BASICLU package is used, which implements a Markowitz *LU* factorization and a variant of the Forrest-Tomlin update [21]. Both packages are available from <https://www.maths.ed.ac.uk/ERGO/software.html>.

In this section the configuration is benchmarked on a diverse set of LP models. For comparison the interior point solver from the commercial software Gurobi [1] version 7.0 is used, which represents a state-of-the-art implementation based on Cholesky factorization. The computations were run on a desktop computer with an Intel i5-6500 processor with 4 cores and 8GB physical memory. While it is clear that the Cholesky factorization would have benefitted from hardware with higher floating point capability more than the iterative solver, this setup was chosen because it is often used by practitioners.

For the test set all LP models from the sources listed in Appendix A were collected. For the mixed-integer problems the canonical LP relaxation was built, and a presolved version was generated for each model by the Gurobi LP presolve. Instances for which the resulting $\bar{m} \times \bar{n}$ constraint matrix was such that $\min(\bar{m}, \bar{n}) \leq 1000$ were removed, because for these models the normal equations (possibly after dualization) can be solved efficiently by dense linear algebra routines.

Both solvers were then applied with default parameters and a time limit of 36,000 seconds. IPX was run on the presolved models, whereas Gurobi was given the original models and its presolve time was subtracted from its total runtime. Infeasible and unbounded models, as well as models for which the Cholesky factorization required more than the 8GB physical memory were removed from the set. Models were also removed if they were solved by both methods within 1 second. The test set was finally cleaned by keeping only 1 or 2 instances of the same model (for example, from the 12 **pds** instances only **pds-60** and **pds-100** were kept). The resulting set contained 170 LP models that are listed with their solution times in Appendix A. An overview of their dimensions is given in Table 3. While there is a lack of truly large instances, the set represents a wide range of medium to large-scale models.

IPX does not provide a simplex implementation for cleaning up the basic solution after the crossover push phases. For the study the final basis was used as starting basis for the

subset	instances	IPX/Gurobi	IPX faster	Gurobi faster
>1s	164	3.67	22	142
>10s	89	3.94	16	73
>100s	41	6.29	6	35

Table 4: Runtime comparison on 164 LP models that were solved by both methods.

Gurobi primal or dual simplex, depending on which infeasibility was smaller. In 150 cases Gurobi decided the initial solution to be optimal within its default tolerances. In 15 cases a simplex run was necessary and the time was added to the total IPX runtime.

All models in the test set were solved by either IPX or Gurobi to basic solution. The Gurobi crossover reached time limit on 1 instance (**nug30**), whereas the IPX interior point method failed on 5 instances:

- On **stormg2.1000**, **ns2122603** and **ns1688926** the IPM stopped after no progress was achieved over a number of iterations. The issue seems to be solvable by a refined IPM implementation (for example using centrality correctors or a more conservative choice of step sizes). The latter two models are questionable numerically, however, due to a wide range of entries in the problem data.
- On **cont1.1** and **cont11.1** the initial LU factorization ran out of memory. It turned out that the large fill-in was caused by the default pivot tolerance of 0.0625 being too small. For the related but smaller instances **cont1** and **cont11** IPX detected the initial LU factorization to be unstable and tightened the pivot tolerance to 0.3. In the repeated factorization the fill-in decreased by about a factor 4. After setting the initial LU pivot tolerance to 0.3, **cont1.1** was solved to an optimal basic solution in 3155 seconds (Gurobi required 1951 seconds); **cont11.1** reached time limit after 3 interior point iterations with basis preconditioning.

Table 4 compares the runtimes on the 164 models that were solved by both codes with default parameters. The column “IPX/Gurobi” shows the geometric mean of the runtime ratios, a value >1.0 meaning that IPX was by that factor slower. The subset “>10s” consists of the models for which at least one solver required more than 10 seconds. The last subset should be considered with care because 41 instances are insufficient to draw a conclusion.

IPX solved the vast majority of test problems and its average runtime on the medium-size instances was in the same order of magnitude as that of Gurobi. Hence the new approach proved to be a general-purpose LP solver. For large instances, say $m \geq 100,000$, the irregular memory access of the iterative solver and update scheme became decisive. Here the Cholesky factorization often performed better as long as the required number of floating point operations was not too high.

The breakdown of the total IPX runtime into different parts of the algorithm is illustrated in Figure 2 for the 87 models that IPX solved successfully but took longer than 10 seconds. Computing the starting basis was inexpensive except for **cont1**, where it accounted for one third of the total time. Here the issue was the large fill-in in the first LU factorization before tightening the pivot tolerance. On average 15% of the time for running the linear solver was spent in the initial IPM iterations (not shown separately). Taking geometric means, preparing the preconditioner and running the iterative solver accounted for 22% and 45% of the total time, respectively. Crossover took 2% on average, but dominated the IPM runtime on 7 models.

Figure 2: Fraction of total runtime spent for computing the crash basis (black bar at the bottom), for updating the basis during the interior point solve (dark grey), for running the linear solver (light grey) and for crossover (black bar at the top). The 87 LP models are ordered on the x-axis by total runtime.

7 Conclusions

The results show that a robust implementation of an interior point solver based on iterative linear algebra is possible. The key idea of the approach presented in this paper is the maximum volume criterion for choosing a basis matrix for preconditioning. We proposed a heuristical method for maintaining a “good” basis at reasonable cost. Although the approach is considerably slower than a state-of-the-art Cholesky factorization on average, it is faster for relevant problems and might therefore be used as an alternative when appropriate.

Maintaining a basis matrix in the IPM enables further options to improve numerical stability. We described the treatment of free variables, otherwise an issue in IPM implementation, and mentioned the ability to remove close-to-converged variables from the optimization process. The latter option allows running the IPM to high accuracy, which is required for a clean crossover. These techniques can also be used to supplement direct linear algebra, for example by performing some extra IPM iterations with basis preconditioning prior to crossover.

A Test Set and Solution Times

The LP models have been obtained from the following sources:

- (1) J. Castro: <http://www-eio.upc.es/~jcastro/>;
(a) huge CTA instances, (b) integrated refinery problems, (c) L1.zip, (d) Linf.zip
- (2) J. A. J. Hall: <http://www.maths.ed.ac.uk/hall/PublicLP/>
- (3) Kennington collection: <http://www.netlib.org/lp/data/kennington/index.html>
- (4) C. Mészáros: <http://old.sztaki.hu/~meszaros/publicftp/lptestset/>;
(a) misc, (b) New, (c) problematic, (d) stochlp
- (5) MIPLIB 2010: <http://miplib.zib.de/download/miplib2010-complete.tgz>
- (6) H. Mittelmann: <http://plato.asu.edu/ftp/lptestset/>;
(a) fome, (b) misc, (c) nug, (d) pds, (e) rail
- (7) Netlib collection: <http://www.netlib.org/lp/data/index.html>

The table provides the computation times in seconds, excluding presolve, obtained on an Intel i5-6500 CPU (4 cores, 3.2GHz, 6MB L3 cache) and 8GB of physical memory.

src	name	IPX				Gurobi		
		total	IPM	push	clean-up	total	IPM	crossover
1a	L1.sixm250obs	180.13	176.50	3.63		43.51	42.11	1.40
1a	L1.sixm500obs	553.82	506.01	47.81		471.78	468.72	3.06
1b	srd120	329.48	327.37	2.11		496.50	484.87	11.63
1b	srd240	1883.34	1877.22	6.12		3615.52	3577.24	38.28
1c	L1.bts4	5.70	5.56	0.14		1.43	1.27	0.16
1c	L1.five20b	47.36	47.21	0.15		6.50	3.74	2.76
1d	Linf.bts4	11.14	10.80r	0.26	0.08	49.29	48.72	0.57
1d	Linf.five20b	243.51	64.63r	16.31	162.57	25.09	10.79	14.30
2	dcp2	5.70	5.68r	0.02		0.74	0.59	0.15
3	cre-b	2.30	2.27	0.03		0.43	0.38	0.05
3	ken-18	5.54	5.46	0.08		1.26	1.11	0.15
3	osa-60	5.66	5.59	0.07		1.07	0.99	0.08
4a	bas1lp	4.00	3.84	0.16		1.18	1.03	0.15
4a	baxter	2.45	2.39	0.06		1.94	1.89	0.05
4a	co9	2.11	2.08	0.03		0.43	0.35	0.08
4a	dbic1	58.43	46.69	11.74		11.97	2.79	9.18
4a	dbir1	17.32	16.32	1.00		1.62	1.42	0.20
4a	e18	12.57	12.33	0.24		63.23	63.07	0.16
4a	ex3sta1	3.45	3.43	0.02		0.44	0.21	0.23
4a	jendrec1	0.91	0.91	0.00		1.13	0.71r	0.42
4a	lp11	21.91	21.44	0.47		4.10	1.79	2.31
4a	mod2	6.11	6.07	0.04		1.47	1.24	0.23
4a	model10	1.56	1.54	0.01	0.01	0.31	0.24	0.07
4a	nemsemm1	5.78	5.77	0.01		0.63	0.58	0.05
4a	nl	1.03	1.02	0.01		0.40	0.37	0.03
4a	nsct1	7.66	7.50	0.16		3.52	3.41	0.11
4a	p010	4.25	4.24	0.01		0.15	0.13	0.02
4a	rat7a	2.04	2.03	0.01		0.53	0.34	0.19
4a	route	2.37	2.29	0.08		0.44	0.39	0.05

src	name	IPX				Gurobi		
		total	IPM	push	clean-up	total	IPM	crossover
4a	stat96v1	21.68	21.16r	0.39	0.13	20.75	4.94r	15.81
4a	stat96v3	1423.94	470.08r	13.97	939.89	8063.47	5.22r	8058.25
4a	ulevimin	3.90	3.87	0.03		0.60	0.53	0.07
4a	world	6.74	6.70	0.04		1.70	1.35	0.35
4b	degme	2496.77	2495.89	0.88		166.39	79.77	86.62
4b	karted	499.65	499.42	0.23		144.93	27.22	117.71
4b	tp-6	2820.88	2820.03	0.85		65.97	41.95	24.02
4b	ts-palko	314.70	314.50	0.20		79.19	8.60	70.59
4c	gen4	12.28	11.80r	0.48		1.82	0.50	1.32
4c	l30	1.79	1.59	0.15	0.05	4.03	0.81	3.22
4d	fxm3_16	7.24	7.18	0.06		0.95	0.76	0.19
4d	pltexpa4_6	3.03	2.98	0.05		1.27	1.19	0.08
4d	scfxm1-2r-256	8.14	8.05	0.09		1.56	1.04	0.52
4d	stormg2-125	28.26	28.05	0.21		2.85	2.47	0.38
4d	stormg2_1000		f			29.31	24.14	5.17
5	30_70_45_095_100	2.44	0.96	1.48		1.22	0.32	0.90
5	app1-2	4.45	4.39	0.06		1.04	0.87	0.17
5	atlanta-ip	10.16	10.00	0.16		4.83	4.49	0.34
5	bab3	120.10	119.76	0.34		6.23	5.89	0.34
5	bley_xl1	10.42	7.82	2.60		13.47	13.22	0.25
5	buildingenergy	262.46	260.49	1.97		6.14	5.66	0.48
5	circ10-3	1.24	0.97	0.27		0.56	0.38	0.18
5	core4872-1529	4.24	4.14	0.10		1.58	1.27	0.31
5	dano3mip	1.68	1.66	0.02		1.09	0.94	0.15
5	datt256	1218.29	158.08	1060.21		718.31	0.99	717.32
5	dc11	5.91	5.85	0.06		1.08	1.00	0.08
5	dolom1	2.38	2.35	0.03		0.58	0.54	0.04
5	ds-big	196.95	195.40	1.55		6.02	5.02	1.00
5	ex10	38.71	12.89	25.82		121.79	114.36	7.43
5	f2000	3.59	2.33	1.26		2.77	0.87	1.90
5	germanrr	1.98	1.92	0.03	0.03	0.32	0.29	0.03
5	gmut-75-50	5.97	5.92	0.02	0.03	0.77	0.68	0.09
5	in	5381.13	4622.91	758.22		101.16	82.71	18.45
5	ivu06-big	2067.81	2064.87	2.94		34.32	30.73	3.59
5	ivu52	69.17	68.36	0.81		3.50	3.19	0.31
5	map06	11.29	11.20	0.09		3.48	3.18	0.30
5	mining	772.01	769.99	2.02		46.61	36.77	9.84
5	momentum3	25.74	25.58	0.16		22.06	18.04	4.02
5	msc98-ip	4.13	3.62	0.27	0.24	3.01	2.60	0.41
5	mspp16	65.30	64.97	0.33		57.42	52.71	4.71
5	mzzv11	2.29	2.12	0.17		1.56	1.49r	0.07
5	n15-3	18.96	18.62	0.34		2.41	2.07	0.34
5	n3seq24	27.02	26.51	0.51		4.09	3.74	0.35
5	nb10tb	92.78	78.21	0.58	13.99	53.61	41.88	11.73
5	neos-1140050	6.81	6.78	0.03		3.93	3.29	0.64
5	neos-1429212	34.07	31.44	2.63		3.28	1.88	1.40
5	neos-1605075	1.28	0.94	0.34		1.01	0.86	0.15
5	neos-476283	23.75	23.27	0.48		20.02	19.75	0.27

src	name	IPX				Gurobi		
		total	IPM	push	clean-up	total	IPM	crossover
5	neos-506428	7.51	4.38	3.13		1.27	0.96	0.31
5	neos-520729	12.09	11.76	0.33		0.43	0.35	0.08
5	neos-631710	220.24	3.41	216.83		8.65	1.19	7.46
5	neos-738098	2.63	1.49	1.14		0.62	0.48	0.14
5	neos-799711	1.44	1.31	0.13		0.50	0.40	0.10
5	neos-824661	2.04	1.73	0.31		0.35	0.20	0.15
5	neos-826694	1.16	0.81	0.35		0.18	0.08	0.10
5	neos-933638	1.94	1.05	0.89		0.61	0.36	0.25
5	neos-941313	8.44	6.87	1.57		1.10	0.37	0.73
5	neos-948126	1.64	1.12	0.52		0.40	0.21	0.19
5	neos-957389	4.42	4.31	0.11		0.58	0.51	0.07
5	neos-984165	1.70	1.21	0.49		0.41	0.22	0.19
5	neos6	1.09	1.06	0.03		0.45	0.42	0.03
5	neos808444	2.80	1.19	1.61		0.64	0.39	0.25
5	net12	2.84	2.81	0.03		4.41	4.36	0.05
5	netdiversion	110.27	105.20	5.07		14.44	4.16	10.28
5	npmv07	23.06	22.90	0.16		3.73	3.50	0.23
5	ns1111636	5.32	5.10	0.22		0.50	0.38	0.12
5	ns1116954	6.96	2.10	4.86		12.17	11.72	0.45
5	ns1631475	4.62	4.37	0.25		1.04	0.55	0.49
5	ns1644855	49.55	47.95	1.60		131.62	131.15	0.47
5	ns1663818	1105.00	1077.75 r	27.25		208.25	197.32 r	10.93
5	ns1685374	9.89	9.87	0.02		11.95	1.56	10.39
5	ns1696083	4.94	4.86	0.08		1.70	1.61	0.09
5	ns1758913	26.17	19.57	6.60		20.98	3.89	17.09
5	ns1853823	492.59	473.30	19.29		32.71	20.82 r	11.89
5	ns1854840	136.53	124.13	12.40		6.23	4.67	1.56
5	ns1904248	6.60	1.67	4.93		0.84	0.66	0.18
5	ns1905797	1.57	1.45	0.12		11.45	11.31	0.14
5	ns2017839	20.15	20.06	0.06	0.03	7.30	3.20 r	4.10
5	ns2118727	42.02	41.91	0.11		13.81	13.51	0.30
5	ns2122603		f			4.45	1.55	2.90
5	ns2124243	4.78	4.37	0.41		0.61	0.31	0.30
5	ns2137859	11.63	11.42	0.21		2.66	2.13	0.53
5	ns894244	3.11	2.64	0.47		0.83	0.63	0.20
5	ns930473	7.18	6.80	0.38		1.13	0.75	0.38
5	nsr8k	16.80	16.44	0.36		4.66	3.93	0.73
5	ofi	135.49	134.65	0.75	0.09	28.99	28.46	0.53
5	opm2-z11-s8	19.45	10.88	8.57		22.40	15.04	7.36
5	opm2-z12-s7	44.36	20.42	23.94		37.79	23.23	14.56
5	pb-simp-nonunif	1.31	1.04	0.27		1.27	0.74	0.53
5	rail02	70.95	62.32	8.63		45.21	42.02	3.19
5	rail03	399.19	318.68	80.51		55.43	48.65	6.78
5	ramos3	4.56	0.98	3.58		4.44	0.33	4.11
5	reblock420	1.95	1.18	0.77		0.82	0.57	0.25
5	rmatr100-p5	3.87	3.83	0.04		9.12	9.10	0.02
5	rmatr200-p5	56.49	55.73	0.76		1.29	1.22	0.07
5	rmine14	83.03	80.43	2.60		48.61	20.80	27.81

src	name	IPX				Gurobi		
		total	IPM	push	clean-up	total	IPM	crossover
5	rmine21	2707.83	2580.18	127.65		1619.99	694.80	925.19
5	rocII-9-11	6.59	6.25	0.34		1.09	0.74	0.35
5	satellites3-40-fs	22.04	14.67	7.37		6.18	3.08	3.10
5	satellites3-40	34.42	25.38	9.04		61.16	56.60	4.56
5	sct1	4.18	4.12	0.06		1.31	1.23	0.08
5	shs1023	173.49	170.75	2.74		14.42	12.80	1.62
5	siena1	3.86	3.78	0.08		1.12	1.02	0.10
5	sing161	281.89	273.55	8.34		13.59	12.08	1.51
5	sing359	268.72	258.44	10.28		11.87	10.05	1.82
5	sp97ar	1.44	1.43	0.01		0.23	0.20	0.03
5	splan1	5242.17	4989.83	250.40	1.94	1775.73	1587.59	188.14
5	stockholm	5.59	5.49	0.10		0.73	0.63	0.10
5	stp3d	196.49	163.36	33.13		11.23	9.40	1.83
5	tanglegram1	7.35	5.82	1.53		7.74	0.30	7.44
5	triptim3	14.20	13.25	0.95		5.30	4.82	0.48
5	uc-case3	7.19	6.96	0.23		1.06	0.90	0.16
5	unitcal_7	12.49	12.06	0.43		1.11	0.90	0.21
5	van	2.73	2.72	0.01		7.48	7.33	0.15
5	vpphard	8.21	7.19	1.02		3.70	3.40	0.30
5	vpphard2	26.22	25.84	0.38		36.63	36.30	0.33
5	wnq-n100-mw99-14	20.27	20.04	0.23		39.36	39.02	0.34
6a	fome13	40.48	36.03	4.45		5.85	4.62	1.23
6a	fome21	15.51	15.00	0.51		3.47	3.10	0.37
6b	cont1	88.48	88.39	0.09		18.82	2.00	16.82
6b	cont11	668.59	641.68	4.80	22.11	249.96	1.86	248.10
6b	cont11_1		f			3256.25	32.06	3224.19
6b	cont1_1		f			1951.28	50.34	1900.94
6b	neos	118.83	118.63	0.20		8.11	7.76	0.35
6b	neos1	7.90	7.79	0.11		1.43	1.32	0.11
6b	neos3	60.89	27.39	33.50		38.56	11.39	27.17
6b	ns1687037	145.85	145.75	0.10		24.51	15.43	9.08
6b	ns1688926		f			28.94	22.52r	6.42
6b	sgpf5y6	3.42	3.32	0.10		1.70	1.41	0.29
6b	watson_1	85.67	85.17	0.50		4.51	3.26	1.25
6b	watson_2	193.78	192.68	1.10		6.49	3.49	3.00
6c	nug08-3rd	573.80	253.75	320.05		91.23	18.80	72.43
6c	nug20	192.26	176.12	16.14		435.45	37.49	397.96
6c	nug30	11473.21	5125.17	6336.99	11.05		1400.47	t
6d	pds-100	110.60	104.93	5.67		30.88	28.14	2.74
6d	pds-60	39.51	37.67	1.84		16.52	15.40	1.12
6e	rail2586	259.01	258.58	0.43		15.32	12.66	2.66
6e	rail4284	370.69	369.90	0.79		32.89	29.68	3.21
7	df1001	2.16	2.03	0.13		0.64	0.53	0.11
7	pilot87	1.62	1.62	0.00		0.43	0.25	0.18
7	qap15	15.41	13.28	2.13		4.26	2.78	1.48

f: failed, t: time limit, r: IPM solution reported not optimal

References

- [1] Gurobi Optimization. <http://www.gurobi.com>. Accessed: Jan 24, 2017.
- [2] G. Al-Jeiroudi, J. Gondzio, and J. A. J. Hall. Preconditioning indefinite systems in interior point methods for large scale linear optimisation. *Optim. Methods Softw.*, 23(3), 2008.
- [3] M. Arioli and I. S. Duff. Preconditioning linear least-squares problems by identifying a basis matrix. *SIAM J. Sci. Comput.*, 37(5):S544–S561, 2015.
- [4] R. E. Bixby, M. Fenelon, Z. Gu, E. Rothberg, and R. Wunderling. MIP: theory and practice—closing the gap. In *System modelling and optimization (Cambridge, 1999)*, pages 19–49. Kluwer Acad. Publ., Boston, MA, 2000.
- [5] R. E. Bixby and M. J. Saltzman. Recovering an optimal LP basis from an interior point solution. *Oper. Res. Lett.*, 15(4):169–178, 1994.
- [6] I. S. Duff. On algorithms for obtaining a maximum transversal. *ACM Trans. Math. Softw.*, 7(3):315–330, September 1981.
- [7] J. Gondzio. Multiple centrality corrections in a primal-dual method for linear programming. *Comput. Optim. Appl.*, 6(2), 1996.
- [8] S. A. Goreinov, I. V. Oseledets, D. V. Savostyanov, E. E. Tyrtyshnikov, and N. L. Zamarashkin. How to find a good submatrix. In *Matrix methods: theory, algorithms and applications*. World Sci. Publ., Hackensack, NJ, 2010.
- [9] J. A. J. Hall and K. I. M. McKinnon. Hyper-sparsity in the revised simplex method and how to exploit it. *Comput. Optim. Appl.*, 32(3):259–283, 2005.
- [10] P. M. J. Harris. Pivot selection methods of the Devex LP code. *Math. Programming*, 5, 1973.
- [11] N. J. Higham and S. D. Relton. Estimating the largest elements of a matrix. *SIAM J. Sci. Comput.*, 38(5), 2016.
- [12] D. E. Knuth. Semioptimal bases for linear dependencies. *Linear and Multilinear Algebra*, 17(1), 1985.
- [13] D. G. Luenberger. The conjugate residual method for constrained minimization problems. *SIAM J. Numer. Anal.*, 7:390–398, 1970.
- [14] N. Megiddo. On finding primal- and dual-optimal bases. *ORSA J. Comput.*, 3(1), 1991.
- [15] S. Mehrotra. On the implementation of a primal-dual interior point method. *SIAM J. Optim.*, 2(4):575–601, 1992.
- [16] A. R. L. Oliveira and D. C. Sorensen. A new class of preconditioners for large-scale linear systems from interior point methods for linear programming. *Linear Algebra Appl.*, 394:1–24, 2005.
- [17] C.-T. Pan. On the existence and computation of rank-revealing LU factorizations. *Linear Algebra Appl.*, 316, 2000.
- [18] Y. Saad. *Iterative methods for sparse linear systems*. Society for Industrial and Applied Mathematics, Philadelphia, PA, second edition, 2003.

- [19] L. Schork. *Basis Preconditioning in Interior Point Methods*. PhD thesis, University of Edinburgh, submitted 2018.
- [20] L. Schork and J. Gondzio. Maintaining a basis matrix in the linear programming interior point method. Technical Report ERGO-17-009, University of Edinburgh, 2017.
- [21] L. Schork and J. Gondzio. Permuting spiked matrices to triangular form and its application to the Forrest-Tomlin update. Technical Report ERGO-17-002, University of Edinburgh, 2017.
- [22] S. J. Wright. *Primal-dual interior-point methods*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1997.