Available online at www.sciencedirect.com

ScienceDirect

journal homepage: www.elsevier.com/locate/cosrev

Review Article

A still simpler way of introducing interior-point method for linear programming

Kurt Mehlhorn[a,*], *Sanjeev Saxena*[b]

[a] *Max-Planck-Institut für Informatik, Saarbrücken, Germany*
[b] *Computer Science and Engineering, Indian Institute of Technology, Kanpur, 208 016, India*

A R T I C L E I N F O

Article history:
Received 13 June 2016
Accepted 23 July 2016
Published online 12 August 2016

Keywords:
Linear programming
Interior point method
Simple presentation

A B S T R A C T

Linear programming is now included in algorithm undergraduate and postgraduate courses for computer science majors. We give a self-contained treatment of an interior-point method which is particularly tailored to the typical mathematical background of CS students. In particular, only limited knowledge of linear algebra and calculus is assumed.

Contents

1. Introduction 2
2. Iterative improvement: Use of the Newton– Raphson method 3
3. Invariants in each iteration 4
4. Initial solution 6
5. Extracting an optimal solution 8
6. Complexity 9
7. The proofs of Facts 1–3 9
Acknowledgments 10
Appendix. Result from algebra 10
References 11

* *Corresponding author.*
E-mail addresses: mehlhorn@mpi-inf.mpg.de (K. Mehlhorn), ssax@iitk.ac.in (S. Saxena).

http://dx.doi.org/10.1016/j.cosrev.2016.07.001

1. Introduction

Terlaky [1] and Lesaja [2] have suggested simple ways to teach interior-point methods. In this paper, we suggest an alternative and maybe still simpler way which is particularly tailored to the typical mathematical background of CS students. In particular, only limited knowledge of linear algebra and calculus is assumed. We have selected most of the material from popular textbooks [3–8] to assemble a self-contained presentation of an interior point method—little of this material is new.

The canonical *linear programming problem* is to

$$\text{minimize } c^T x \text{ subject to } Ax = b \text{ and } x \geq 0. \tag{1}$$

Here, A is an $m \times n$ matrix, c and x are n-dimensional, and b is an m-dimensional vector. A *feasible solution* is any vector x with $Ax = b$ and $x \geq 0$. The problem is *feasible* if there is a feasible solution, and *infeasible* otherwise. A feasible problem is *unbounded* (or more precisely the corresponding objective function is unbounded) if for every real z, there is a feasible x with $c^T x \leq z$, and *bounded* otherwise.

In our presentation, we first assume that feasible solutions to the primal and the corresponding dual LP satisfying a certain set of properties (properties (I1)–(I3) in Section 3) are available. We then show how to iteratively improve these solutions in Sections 2 and 3. In each iteration the gap between the primal and the dual objective value is reduced by a factor $1 - O(1/\sqrt{n})$, where n is the number of variables. The iterative improvement scheme leads to solutions that are arbitrarily close to optimality. In Sections 4 and 5 we discuss how to find the appropriate initial solutions and how to extract an optimal solution from a sufficiently good solution by rounding. Either or both these sections may be skipped in a first course.

Remark 1. It is easy to deal with maximization instead of minimization and with inequality constraints. Indeed, maximize $c^T x$ is equivalent to minimize $-c^T x$. Constraints of type $\alpha_1 x_1 + \cdots + \alpha_n x_n \leq \beta$ can be replaced by $\alpha_1 x_1 + \ldots + \alpha_n x_n + \gamma = \beta$ with a new (slack) variable $\gamma \geq 0$. Similarly, constraints of type $\alpha_1 x_1 + \cdots + \alpha_n x_n \geq \beta$ can be replaced by $\alpha_1 x_1 + \cdots + \alpha_n x_n - \gamma = \beta$ with a (surplus) variable $\gamma \geq 0$.

We consider another problem, the *dual problem*, which is

$$\text{maximize } b^T y, \text{ subject to } A^T y + s = c, \text{ with variables}$$
$$s \geq 0 \text{ and unconstrained variables } y. \tag{2}$$

The vector y has m components and the vector s has n components. We will call the original problem the *primal problem*.

Claim 1 (*Weak Duality*). *If x is a solution of $Ax = b$ with $x \geq 0$ and (y, s) is a solution of $A^T y + s = c$ with $s \geq 0$, then*

1. $x^T s = c^T x - b^T y$, *and*
2. $b^T y \leq c^T x$, *with equality if and only if $s_i x_i = 0$ for all* is.

Proof. We multiply $s = c - A^T y$ with x^T from the left and obtain

$$\begin{aligned} x^T s &= x^T c - x^T(A^T y) = c^T x - (x^T A^T) y \\ &= c^T x - (Ax)^T y = c^T x - b^T y. \end{aligned}$$

As $x, s \geq 0$, we have $x^T s \geq 0$, and hence, $c^T x \geq b^T y$.

Equality will hold if $x^T s = 0$, or equivalently, $\sum_i s_i x_i = 0$. Since $s_i, x_i \geq 0$, $\sum_i s_i x_i = 0$ if and only if $s_i x_i = 0$ for all i. ∎

If x is a feasible solution of the primal and (y, s) is a feasible solution of the dual, the difference $c^T x - b^T y$ is called the *objective value gap* of the solution pair. Thus, if the objective values of a primal feasible and a dual feasible solution are the same, then both solutions are optimal. Actually, from the Strong Duality Theorem, if both primal and dual solutions are optimal, then the equality will hold. We will prove the Strong Duality Theorem in Section 5 (Theorem 2).

If the primal and the dual are both feasible, neither of them can be unbounded as by Claim 1, the objective value of all dual feasible solutions are less than or equal to the objective values of any primal feasible solution. As a consequence: If the primal and the dual are feasible, both are bounded. If the primal is unbounded, the dual is infeasible, and if the dual is unbounded, the primal is infeasible. It may happen that both problems are infeasible. It is also true, that if the primal is feasible and bounded, the dual is feasible and bounded, and vice versa. This is a consequence of strong duality.

We will proceed under the assumption that the primal as well as the dual problem are bounded and feasible. This allows us to concentrate on the core of the interior point method, the iterative improvement scheme. We come back to this point in Section 4.

Claim 1 implies, that if we are able to find a solution to the following system of equations and inequalities

$$Ax = b, \; A^T y + s = c, \; x_i s_i = 0 \text{ for all } i, \; x \geq 0, \; s \geq 0,$$

we will get optimal solutions of both the original primal and the dual problem. Notice that the constraints $x_i s_i = 0$ are nonlinear and hence it is not clear whether we have made a step towards the solution of our problem. The idea is now to relax the conditions $x_i s_i = 0$ to the conditions $x_i s_i \approx \mu$ (with the exact form of this equation derived in the next section), where $\mu \geq 0$ is a parameter. We obtain

$$(P_\mu) \quad Ax = b, \; A^T y + s = c, \; x_i s_i \approx \mu \text{ for all } i, \; x > 0, \; s > 0.$$

We will show:

1. (initial solution) For a suitable μ, it is easy to find a solution to the problem P_μ. This will be the subject of Section 4.
2. (iterative improvement) Given a solution (x, y, μ) to P_μ, one can find a solution (x', y', s') to $P_{\mu'}$, where μ' is substantially smaller than μ. This will be the subject of Sections 2 and 3. Applying this step repeatedly, we can make μ arbitrarily small.
3. (final rounding) Given a solution (x, y, μ) to P_μ for sufficiently small μ, one can extract an exact solution for the primal and the dual problem. This will be the subject of Section 5.

For the iterative improvement, it is important that $x > 0$ and $s > 0$. For this reason, we replace the constraints $x \geq 0$ and $s \geq 0$ by $x > 0$ and $s > 0$ when defining problem P_μ (see Fig. 1).

Note that $x_i s_i \approx \mu$ for all i implies $b^T y - c^T x \approx n\mu$ by Claim 1. Thus, repeated application of iterative improvement will make the gap between the primal and dual objective values arbitrarily small.

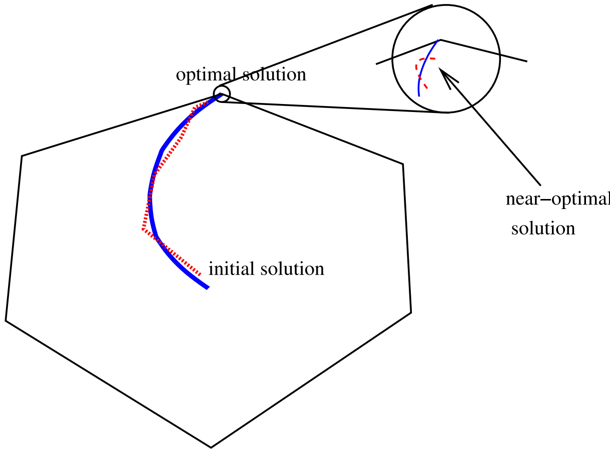

Fig. 1 – The interior of the polygon comprises all points (x, y, s) satisfying $Ax = b$ and $A^T y + s = c$, $x > 0$, and $s > 0$. The blue (bold) line consists of all points in this polygon with $x_i s_i = \mu$ for all i and some $\mu > 0$. These points trace a line inside the polygon that ends in an optimal point. The optimal solution lies on the boundary of the polygon (in the figure, the optimal point is a vertex of the polygon) and satisfies $x_i s_i = 0$ for all i. The red (dashed) line illustrates the steps of the algorithm. It follows the blue (bold) line in discrete steps. The close-up shows the situation near the optimal solution. The algorithm stops tracing the blue curve and rounds to the near-optimal red solution obtained at this point of time to an optimal solution. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Throughout the paper we assume that the rows of A are linearly independent and that $n > m$, i.e., we have more variables than constraints.[1]

[1] Indeed, we can use Gaussian elimination to remove superfluous constraints and to make the rows of A independent. Assume first that A contains a row i in which all entries are equal to zero. If b_i is also zero, we simply delete the row. If b_i is nonzero, the system of equations has no solution, and we declare the problem infeasible and stop. Now, every row of A contains a nonzero entry, in particular, the first row. We may assume that a_{11} is nonzero. Otherwise, we interchange two columns. We multiply the ith equation by $-\frac{a_{11}}{a_{i1}}$ and subtract the first equation. In this way, the first entry of all equations but the first becomes zero. If any row of A becomes equal to the all zero vector, we either delete the equation or declare the problem infeasible. We now proceed in the same way with the second equation. We first make sure that a_{22} is nonzero by interchanging columns if necessary. Then we multiply the ith equation (for $i > 2$) by $-\frac{a_{22}}{a_{21}}$ and subtract the second equation. And so on. In the end, all remaining equations will be linearly independent. Equivalently, the resulting matrix will have full row-rank.

We now have m constraints in n variables with $n \geq m$. If $n = m$, the system $Ax = b$ has a unique solution (recalling that A has full row-rank and is hence invertible). We check whether this solution is non-negative. If so, we have solved the problem. Otherwise, we declare the problem infeasible. So, we may from now on assume $n > m$ (more variables than constraints).

2. Iterative improvement: Use of the Newton–Raphson method

This section and the next follow Roos et al. [3] (see also Vishnoi [9]).

Let us assume that we have a solution (x, y, s) to

$Ax = b$ and $A^T y + s = c$ and $x > 0$ and $s > 0$.

We will use the Newton–Raphson Method [3] to get a "better" solution. Let us choose the next values as $x' = x + h$, $y' = y + k$, and $s' = s + f$. We can think of the steps h, k, and f as small values. Then we want, ignoring the positivity constraints for x' and s' for the moment:

1. $Ax' = A(x + h) = b$, or equivalently, $Ax + Ah = b$. Since $Ax = b$, this is tantamount to $Ah = 0$.
2. $A^T y' + s' = A^T(y + k) + (s + f) = c$. Since $A^T y + s = c$, we get $A^T k + f = c - A^T y - s = 0$.
3. $x'_i s'_i = (x_i + h_i)(s_i + f_i) \approx \mu'$, or equivalently, $x_i s_i + h_i s_i + f_i x_i + h_i f_i \approx \mu'$. We drop the quadratic term $h_i f_i$ (if the steps h_i and f_i are small, the quadratic term $h_i f_i$ will be very small) and turn the approximate equality into an equality, i.e., we require $x_i s_i + h_i s_i + f_i x_i = \mu'$ for all i.

Thus, we have a system of linear equations for h_i, k_i, f_i, namely,

$$\text{system (S)} \qquad \begin{aligned} Ah &= 0 \\ A^T k + f &= 0 \\ h_i s_i + f_i x_i &= \mu' - x_i s_i \quad \text{for all } i. \end{aligned}$$

We show in Theorem 1 that system (S) can be solved by "inverting" a matrix. Note that there are n variables h_i, m variables k_j, and n variables f_i for a total of $2n + m$ unknowns. Also note that $Ah = 0$ constitutes m equations, $A^T k + f = 0$ constitutes n equations, and $h_i s_i + f_i x_i = \mu' - x_i s_i$ for all i comprises n equations. So we have $2n + m$ equations and the same number of unknowns. Also note that the x_i and s_i are *not* variables in this system, but fixed values.

Before we show that the system has a unique solution, we make some simple observations. From the third group of equations, we conclude

Claim 2. *$(x_i + h_i)(s_i + f_i) = \mu' + h_i f_i$, and $(x + h)^T(s + f) = n\mu' + h^T f$.*

Proof. From the third group of equations, we obtain

$(x_i + h_i)(s_i + f_i) = x_i s_i + h_i s_i + f_i x_i + h_i f_i = \mu' + h_i f_i$.

Summation over i yields

$$\begin{aligned} (x + h)^T(s + f) &= \sum_i (x_i + h_i)(s_i + f_i) \\ &= \sum_i (\mu' + h_i f_i) = n\mu' + h^T f. \quad \blacksquare \end{aligned}$$

Claim 3. *$h^T f = f^T h = \sum_i h_i f_i = 0$, i.e., the vectors h and f are orthogonal to each other.*

Proof. Multiplying $A^T k + f = 0$ by h^T from the left, we obtain $h^T A^T k + h^T f = 0$. Since $h^T A^T = (Ah)^T = 0$, the equality $h^T f = 0$ follows. $\blacksquare$

Claim 4. *$c^T(x + h) - b^T(y + k) = (x + h)^T(s + f) = n\mu'$.*

$$\begin{aligned}
Sh + Xf &= \mu' e - XSe & & \\
h + S^{-1}Xf &= S^{-1}\mu' e - S^{-1}XSe & &\text{pre-multiply by } S^{-1} \\
h + S^{-1}Xf &= \mu' S^{-1}e - XS^{-1}Se & &\text{diagonal matrices commute} \\
h + S^{-1}Xf &= \mu' S^{-1}e - x & &\text{as } Xe = x \\
Ah + AS^{-1}Xf &= \mu' AS^{-1}e - Ax & &\text{pre-multiply by } A \\
AS^{-1}Xf &= \mu' AS^{-1}e - b & &\text{since } Ax = b \text{ and } Ah = 0 \\
-AS^{-1}XA^{T}k &= \mu' AS^{-1}e - b & &\text{using } f = -A^{T}k \\
b - \mu' AS^{-1}e &= (AS^{-1}XA^{T})k & &
\end{aligned}$$

Box I.

Proof. From Claims 2 and 3, $(x+h)^T(s+f) = n\mu' + h^Tf = n\mu'$. Also, applying Claim 1 to the primal solution $x' = x+h$ and to the dual solution $(y', s') = (y+k, s+f)$ yields $c^T(x+h) - b^T(y+k) = (x+h)^T(s+f)$. ■

Note that $n\mu'$ is the objective value gap of the updated solution.

Theorem 1. *The system (S) has a unique solution.*

Proof. We will follow Vanderbei [4] and use capital letters (e.g. X) in this proof (only) to denote a diagonal matrix with entries of the corresponding row vector (e.g. X has the diagonal entries $x_1, x_2, \ldots, x_n$). We will also use e to denote a column vector of all ones (usually of length n).

Then, in the new notation, the last group of equations becomes

$Sh + Xf = \mu' e - XSe.$

Let us look at the equation in Box I in more detail.

As XS^{-1} is diagonal with positive items, the matrix $W = \sqrt{XS^{-1}}$ is well-defined. Note that the diagonal terms are $\sqrt{x_i/s_i}$; since $x > 0$ and $s > 0$, we have $x_i/s_i > 0$ for all i. Thus, $AS^{-1}XA^T = AW^2A^T = (AW)(AW)^T$. Since A has full rank, $(AW)(AW)^T$, and hence $AS^{-1}XA^T$, is invertible (see Appendix). Thus,

$$k = (AS^{-1}XA^T)^{-1}\left(b - \mu' AS^{-1}e\right).$$

Then, we can find f from $f = -A^Tk$. And to get h, we use the equation: $h + S^{-1}Xf = \mu' S^{-1}e - x$, i.e.,

$$h = -XS^{-1}f + \mu' S^{-1}e - x.$$

Thus, system (S) has a unique solution. ■

What have we achieved at this point? Given feasible solutions (x, y, s) to the primal and the dual problem, we can compute a solution $(x', y', s') = (x+h, y+k, s+f)$ to $Ax' = b$ and $A^Ty' + s' = c$ that also satisfies $h^Tf = 0$ and $x'^Ts = n\mu'$ for any prescribed parameter μ'. Why do we not simply choose $\mu' = 0$ and be done? It is because we have ignored that we want $x' > 0$ and $s' > 0$. We will attend to these constraints in the next section.

3. Invariants in each iteration

Recall that we want to construct solutions (x, y, s) to P_μ for smaller and smaller values of μ. The solution to P_μ will satisfy the following invariants. The first two invariants state that x is a positive solution to the primal and (y, s) is a solution to the dual with positive s. The third invariant formalizes the condition $x_is_i \approx \mu$ for all i.

(I1) (primal feasibility) $Ax = b$ with $x > 0$ (strict inequality).
(I2) (dual feasibility) $A^Ty + s = c$ with $s > 0$ (strict inequality).
(I3) $\sigma^2 := \sum_i \left(\frac{x_is_i}{\mu} - 1\right)^2 \le \frac{1}{4}$.

Remark 2. Even though the variance of x_is_i is $\frac{1}{n}\sum_i (x_is_i - \mu)^2$, we still use the notation σ^2.

We need to show

$$x' > 0 \text{ and } s' > 0 \text{ and } \sigma'^2 := \sum_i \left(\frac{x_i's_i'}{\mu'} - 1\right)^2 \le \frac{1}{4}.$$

We will do so for $\mu' = (1-\delta)\mu$ and $\delta = \Theta\left(\frac{1}{\sqrt{n}}\right)$. Claim 2 gives us an alternative expression for σ'^2, namely,

$$\sigma'^2 = \sum_i \left(\frac{(x_i+h_i)(s_i+f_i)}{\mu'} - 1\right)^2 = \sum_i \left(\frac{h_if_i}{\mu'}\right)^2. \quad (3)$$

We first show that the positivity invariants hold if σ' is less than one.

Claim 5. *If $\sigma' < 1$, then $x' > 0$, and $s' > 0$.*

Proof. We first show that if $\sigma' < 1$ then each product $x_i's_i' = (x_i+h_i)(s_i+f_i) = \mu' + h_if_i$ is positive. From $\sigma' < 1$, we get $\sigma'^2 < 1$. Since $\sigma'^2 = \sum_i (h_if_i/\mu')^2$, each term of the summation must be less than one, and hence, $-\mu' < h_if_i < \mu'$. In particular, $\mu' + h_if_i > 0$ for every i. Thus, each product $(x_i + h)(s_i + f)$ is positive.

Assume for the sake of a contradiction that both $x_i + h_i < 0$ and $s_i + f_i < 0$. But as $s_i > 0$ and $x_i > 0$, this implies $s_i(x_i + h_i) + x_i(s_i + f_i) < 0$, or equivalently, $\mu' + x_is_i < 0$, which is impossible because μ', x_i, s_i are all non-negative. This is a contradiction. ■

We next show $\sigma' \le 1/2$. We first establish

Claim 6. $\frac{\mu}{x_is_i} \le \frac{1}{1-\sigma}$ *for all i and* $\sum_i \left|1 - \frac{x_is_i}{\mu}\right| \le \sqrt{n}\cdot\sigma$.

$$
\begin{aligned}
\sigma'^2 &= \sum_i (H_iF_i)^2 = \frac{1}{4}\left(\sum_i (H_i^2+F_i^2)^2 - \sum_i (H_i^2-F_i^2)^2\right) \\
&\le \frac{1}{4}\sum_i (H_i^2+F_i^2)^2 && \text{since } \sum_i (H_i^2-F_i^2)^2 \ge 0 \\
&\le \frac{1}{4}\left(\sum_i (H_i^2+F_i^2)\right)^2 && \text{more positive terms} \\
&= \frac{1}{4}\left(\sum_i (H_i+F_i)^2\right)^2 && \text{since } H^TF=0 \\
&= \frac{1}{4}\left(\sum_i \frac{\mu}{x_is_i(1-\delta)}\left(-\delta+1-\frac{x_is_i}{\mu}\right)^2\right)^2 && \text{by (4)} \\
&\le \frac{1}{4(1-\delta)^2(1-\sigma)^2}\left(\sum_i \left(-\delta+1-\frac{x_is_i}{\mu}\right)^2\right)^2 && \text{since } \mu/(x_is_i) \le 1/(1-\sigma) \\
&\le \frac{1}{4(1-\delta)^2(1-\sigma)^2}\left(n\delta^2 - 2\delta\sum_i\left(1-\frac{x_is_i}{\mu}\right) + \sum_i\left(1-\frac{x_is_i}{\mu}\right)^2\right)^2 && \text{remove inner square} \\
&\le \frac{1}{4(1-\delta)^2(1-\sigma)^2}\left(n\delta^2 + 2\delta\sum_i\left|1-\frac{x_is_i}{\mu}\right| + \sum_i\left(1-\frac{x_is_i}{\mu}\right)^2\right)^2 \\
&\le \frac{1}{4(1-\delta)^2(1-\sigma)^2}\left(n\delta^2 + 2\delta\sqrt{n}\cdot\sigma + \sigma^2\right)^2 && \text{by Claim 6} \\
&= \frac{1}{4(1-\delta)^2(1-\sigma)^2}\left((\sqrt{n}\delta+\sigma)^2\right)^2, && \text{forming inner square}
\end{aligned}
$$

Box II.

Proof. As $\sigma^2 = \sum_i (1 - x_is_i/\mu)^2$, each individual term in the sum is at most σ^2. Thus, $|1 - x_is_i/\mu| \le \sigma$, and hence, $x_is_i/\mu \ge 1-\sigma$, and further, $\mu/x_is_i \le 1/(1-\sigma)$.

For the second claim, we have to work harder. Consider any n reals z_1 to z_n. Then $(\sum_i |z_i|)^2 \le n\sum_i z_i^2$; this is the frequently used inequality between the one-norm and the two-norm of a vector.[2] We apply the inequality with $z_i = 1 - x_is_i/\mu$ and obtain the second claim. ■

Let us define two new quantities

$$H_i = h_i\sqrt{\frac{s_i}{x_i\mu'}} \quad \text{and} \quad F_i = f_i\sqrt{\frac{x_i}{s_i\mu'}}.$$

Observe that $\sum_i H_iF_i = \sum \frac{h_if_i}{\mu'} = 0$ (from Claim 3) and $\sum_i (H_iF_i)^2 = \sum_i \left(\frac{h_if_i}{\mu'}\right)^2 = \sigma'^2$. Also,

$$
\begin{aligned}
H_i + F_i &= \sqrt{\frac{1}{x_is_i\mu'}}\,(h_is_i + f_ix_i) = \sqrt{\frac{1}{x_is_i\mu'}}\,(\mu' - \mu + \mu - x_is_i) \\
&= \sqrt{\frac{\mu}{x_is_i}\frac{\mu}{\mu'}}\left(\frac{\mu'}{\mu} - 1 + 1 - \frac{x_is_i}{\mu}\right) \\
&= \sqrt{\frac{\mu}{x_is_i(1-\delta)}}\left(-\delta + 1 - \frac{x_is_i}{\mu}\right).
\end{aligned} \tag{4}
$$

Finally, we have the equation in Box II,

and hence,

$$\sigma' \le \frac{(\sqrt{n}\delta+\sigma)^2}{2(1-\sigma)(1-\delta)} \le \frac{(\sqrt{n}\delta+1/2)^2}{2(1-1/2)(1-\delta)} \overset{!}{\le} \frac{1}{2}, \tag{5}$$

where the second inequality holds since the bound for σ' is increasing in σ, and $\sigma \le 1/2$. We need to choose δ such that the last inequality holds. This is why we put an exclamation mark on top of the $\le$-sign. Setting $\delta = c/\sqrt{n}$ for some to be determined constant c yields the requirement

$$\frac{(c+1/2)^2}{(1-\delta)} \overset{!}{\le} \frac{1}{2}, \quad \text{or equivalently,} \quad (2c+1)^2 \overset{!}{\le} 2\left(1-\frac{c}{\sqrt{n}}\right).$$

This holds true for $c = 1/8$ and all $n \ge 1$. Thus, $\delta = 1/(8\sqrt{n})$.

Remark 3. Why do we require $\sigma \le 1/2$ in the invariant? Let us formulate the bound as $\sigma \le \sigma_0$ for some to be determined σ_0. Then, the inequality (5) becomes

$$\frac{(\sqrt{n}\delta+\sigma_0)^2}{2(1-\sigma_0)(1-\delta)} \overset{!}{\le} \sigma_0.$$

We want this to hold for $\delta = \frac{c}{\sqrt{n}}$ and some $c > 0$. In order for the inequality to hold for $c = 0$, we need $\sigma_0 \le 2(1-\sigma_0)$, or equivalently, $\sigma_0 \le 2/3$. Since we want it to hold for some positive c, we need to choose a smaller σ_0; $1/2$ is a nice number smaller than $2/3$.

An alternative proof for invariant (I3) (provided by Andreas Karrenbauer) Andreas Karrenbauer derived an alternative proof for invariant (I3) that avoids introduction of the quantities H and F and is more compact than the above.

Lemma 1. *Assume $\delta \le 1/6$. Then $\sigma \le \delta$ implies $\sigma' \le \delta$.*

[2] Indeed,

$$
\begin{aligned}
n\sum_i z_i^2 - \left(\sum_i z_i\right)^2 &= n\sum_i z_i^2 - \sum_i z_i^2 - 2\sum_{i<j} z_iz_j \\
&= (n-1)\sum_i z_i^2 - 2\sum_{i<j} z_iz_j = \sum_{i<j}(z_i-z_j)^2 \ge 0.
\end{aligned}
$$

$$
\begin{aligned}
\|\sigma'\|_1 &\le \sum_i \left|\frac{h_if_i}{\mu'}\right| \\
&\le \sum_i \frac{1}{2x_is_i\mu'}\left[(\mu'-x_is_i)^2+(h_is_i)^2+(x_if_i)^2\right] && \text{by (6)} \\
&= \sum_i \frac{(\mu'-x_is_i)^2}{x_is_i\mu'} && \text{by (7)} \\
&\le \sum_i \frac{\left(\frac{\mu'}{\mu}-\frac{x_is_i}{\mu}\right)^2}{(1-\delta)(1-\tau)} && \text{as } x_is_i \ge (1-\delta)\mu \text{ and } \mu'=(1-\tau)\mu \\
&= \sum_i \frac{\left(\tau-\left(\frac{x_is_i}{\mu}-1\right)\right)^2}{(1-\delta)(1-\tau)} \\
&= \frac{n\tau^2-2\tau\sum\left(\frac{x_is_i}{\mu}-1\right)+\sigma^2}{(1-\delta)(1-\tau)} \\
&\le \frac{n\tau^2+2\tau\|\sigma\|_1+\sigma^2}{(1-\delta)(1-\tau)} \\
&\le \frac{(\sqrt{n}\tau+\delta)^2}{(1-\delta)(1-\tau)} && \text{since } \|\sigma\|_1 \le \sqrt{n}\delta \text{ (Claim 6)} \\
&\le \frac{4\delta^2}{(1-\delta)(1-\delta/\sqrt{n})} && \text{for the choice } \tau=\delta/\sqrt{n} \\
&\le \delta && \text{for } \delta \le 1/6
\end{aligned}
$$

Box III.

Proof. As $\sigma^2 = \sum\left(\frac{x_is_i}{\mu}-1\right)^2 \le \delta^2$, each individual term must be bounded by δ^2. Thus, $\sigma \le \delta$ implies $\left|\frac{x_is_i}{\mu}-1\right| \le \delta$, or $-\delta \le \frac{x_is_i}{\mu}-1$ or $x_is_i \ge (1-\delta)\mu$.

We define

$$\|\sigma'\|_1 = \sum_i \left|\frac{(x_i+h_i)(s_i+f_i)}{\mu'}-1\right|.$$

Then from the definition of μ' and triangle inequality,

$$\|\sigma'\|_1 = \sum_i \left|\frac{(x_i+h_i)(s_i+f_i)}{\mu'}-1\right| \le \sum_i \left|\frac{x_is_i+x_if_i+h_is_i}{\mu'}-1\right| + \sum_i \left|\frac{h_if_i}{\mu'}\right| = \sum_i \left|\frac{h_if_i}{\mu'}\right|.$$

Again from $x_if_i+s_ih_i = \mu'-x_is_i$, we obtain (by squaring)

$$h_if_i = \frac{1}{2x_is_i}\left[(\mu'-x_is_i)^2-(h_is_i)^2-(x_if_i)^2\right]. \tag{6}$$

Summing over i and using the fact that from Claim 3, $f^Th=0$ we obtain

$$\sum_i(\mu'-x_is_i)^2 = \sum_i\left((h_is_i)^2+(x_if_i)^2\right). \tag{7}$$

Assume that $\mu'=(1-\tau)\mu$ for a τ to be fixed later. Then we have the equation in Box III.

The claim follows as the two norm is always less than the one norm,[3] $\sigma' = \|\sigma'\|_2 \le \|\sigma'\|_1$. ■

[3] If $\alpha=(\alpha_1,\ldots,\alpha_n)$ then $\left(\sum|\alpha_i|\right)^2 = \sum|\alpha_i|^2+2\sum_{i<j}|\alpha_i||\alpha_i| \ge \sum|\alpha_i|^2 = \sum\alpha_i^2$.

4. Initial solution

This section follows Bertsimas and Tsitsiklis [5, p 430]; see also Karloff [6, p 128–129]. We have to deal with three problems:

1. how to make sure that we are dealing with a bounded problem
2. how to make sure that the problem is feasible and if the problem is feasible, then how to find an initial solution
3. how to guarantee condition (I3) for the initial solution.

A standard solution for the second problem is the *big M method*. Let $x_0 \ge 0$ be an arbitrary nonnegative column vector of length n. We introduce a new variable $z \ge 0$, change $Ax=b$ into $Ax+(b-Ax_0)z=b$ and the objective into "minimize c^Tx+Mz", where M is a big number. Note that $x=x_0$ and $z=1$ is a feasible solution to the modified problem. We solve the modified problem. If $z^*=0$ in an optimal solution, we have also found an optimal solution to the original problem. If $z^*>0$ in an optimal solution and M was chosen big enough, the original problem is infeasible.

Remark 4. There are several other methods of dealing with the problem of getting a starting solution. These include self-dual method [1,4] and the infeasible interior point method [10,11].

We assume for the remainder of the presentation that A, b, and c are integral and that U is an integer with $U \ge \left|a_{ij}\right|, \left|b_i\right|, \left|c_j\right|$ *for all i and j.*

We need the following Fact which we will prove in Section 7.

$$\text{minimize } c^T x + M x_{n+2}, \text{ subject to } \begin{array}{ccccccc} Ax & & & + & \rho x_{n+2} & = & d \\ e^T x & + & x_{n+1} & + & x_{n+2} & = & n+2 \\ x \geq 0 & & x_{n+1} \geq 0 & & x_{n+2} \geq 0 & , & \end{array} \tag{8}$$

where $d = \frac{1}{W} b$, $\rho = d - Ae$.

Box IV.

Fact 1. *Let $W = (mU)^m$. If (1) is feasible, there is a feasible solution with all coordinates bounded by W. If, in addition the problem is bounded, there is an optimal solution with this property.*

We now give the details. We add the constraint $e^T x + z \leq (n+2)W$. If the problem was feasible, it will stay feasible. If the problem was bounded, the additional constraint does not change the optimal objective value. If the problem was unbounded, the additional constraint makes it bounded. Using an additional slack variable x_{n+1} we get the equality $e^T x + x_{n+1} + z = (n+2)W$. If we use "normalized variables" $x_i' = \frac{x_i}{W}$, drop the primes and use x_{n+2} for z, we obtain the following auxiliary primal problem (see Eq. (8) in Box IV). We show later in this section, that M can be chosen as $M = 4nU/R$, where $R = \frac{1}{W^2} \cdot \frac{1}{2n((m+1)U)^{3(m+1)}}$. In matrix form, the auxiliary primal is

$$A' \begin{pmatrix} x \\ x_{n+1} \\ x_{n+2} \end{pmatrix} = b', \quad \text{where} \quad A' = \begin{pmatrix} A & 0 & \rho \\ e^T & 1 & 1 \end{pmatrix}$$

$$\text{and } b' = \begin{pmatrix} d \\ n+2 \end{pmatrix}.$$

We make the following observations.

– As $x_i = 1$ for $1 \leq i \leq n+2$ is a feasible solution, (8) is feasible. The feasible region is a polytope contained in the cube defined by $0 \leq x_i \leq n+2$ for all i. The following Fact is shown in Section 7.

Fact 2. *The nonzero coordinates of the vertices of this polytope are at least R.*

– As $0 \leq x_i \leq n+2$ and $c_i \geq -U$ for all i, the objective value is at least $-U(n+2)$ Thus, (8) is bounded.

– If x is a feasible solution to (1) with $x_i \leq W$ for $1 \leq i \leq n$ then $(\frac{1}{W}x, (n+2) - \frac{1}{W}e^T x, 0)$ is a feasible solution to (8) with objective value $\frac{1}{W}c^T x$.

In particular, if (1) is feasible, then (8) has a solution with objective value less than or equal to nU. This follows from $x_i/W \leq 1$ and $c_i \leq U$ for $1 \leq i \leq n$.

– We next show that if (8) has an optimal solution $(x^*, x^*_{n+1}, x^*_{n+2})$ with $x^*_{n+2} = 0$ then (1) is feasible. Indeed, $AWx^* = WAx^* = Wd = b$ and hence Wx^* is feasible for (1). If, in addition, (1) is bounded, Wx^* is an optimal solution of (1). Note that if (1) is bounded, it has an optimal solution x with $x_i \leq W$ by Fact 1. This solution induces a solution of (8) with objective value $\frac{1}{W}c^T x$ by the preceding item. The optimality of $(x^*, x^*_{n+1}, x^*_{n+2})$ implies $c^T x^* \geq \frac{1}{W}c^T x$.

– We finally show that if (8) has an optimal solution with $x^*_{n+2} > 0$, (1) is infeasible. Indeed, then there must be an optimal vertex solution of (8). For this vertex, $x^*_{n+2} \geq R$. The objective value of this solution is at least $M \cdot R - (n+2)U = 2nU$. On the other hand, if (1) is feasible, (8) has a solution with objective value at most nU. Any value of M for which $M \cdot R - (n+2)U > nU$ would work for this argument. $M = 4U/R$ is one such value. This explains the choice of M.

We summarize: Our original problem is feasible if and only if $x^*_{n+2} = 0$ in every optimal solution to (8) if and only if $x^*_{n+2} = 0$ in some optimal solution to (8). Moreover, if $x^*_{n+2} = 0$, and (1) is bounded, $\frac{1}{W}x^*$ is an optimal solution of (1).

Remark 5. By the above, our original problem is feasible if and only if $x^*_{n+2} = 0$ in an optimal solution to (8). So we can distinguish feasible and infeasible problems. How can we distinguish bounded and unbounded problems? Note that the primal is unbounded if it is feasible and the problem "minimize 0 subject to $c^T x = -1$, $Ax = 0$, and $x \geq 0$" is feasible. So the test for unboundedness reduces to two feasibility tests.

The dual problem (with new dual variables y_{m+1}, s_{n+1} and s_{n+2}) is Eq. (9) given in Box V.

Which initial solution should we choose? Recall that we also need to satisfy (I3) for some choice of μ, i.e., $\sum_{1 \leq i \leq n+2} (x_i s_i/\mu - 1)^2 \leq 1/4$. Also, recall that we set x_i to 1 for all i. As $x_{n+1} = 1$, we choose $s_{n+1} = \mu/x_{n+1} = \mu$. Then, from the last equation, $y_{m+1} = -s_{n+1} = -\mu$. The simplest choice for y is $y = 0$. Then, from the first equation, $s = c + e\mu$, and from the second equation $s_{n+2} = M - y_{m+1} = M + \mu$. Observe that all slack variables are positive (provided μ is large enough). For this choice,

$$\frac{x_i s_i}{\mu} - 1 = \frac{c_i}{\mu} \qquad \text{for } i \leq n$$
$$\frac{x_{n+1} s_{n+1}}{\mu} - 1 = 0$$
$$\frac{x_{n+2} s_{n+2}}{\mu} - 1 = \frac{M}{\mu}.$$

Thus, $\sigma^2 = \left(M^2 + \sum c_i^2\right)/\mu^2$. We can make $\sigma^2 \leq 1/4$ by choosing

$$\mu^2 = 4\left(M^2 + \sum c_i^2\right). \tag{10}$$

Summary: Let us summarize what we have achieved.

– For the auxiliary primal problem and its dual, we have constructed solutions $(x^{(0)}, y^{(0)}, s^{(0)})$ that satisfy the invariants for $\mu^{(0)} = 2\left(M^2 + \sum c_i^2\right)^{1/2}$.
– From the initial solution, we can construct a sequence of solutions $(x^{(t)}, y^{(t)}, s^{(t)})$ and corresponding $\mu^{(t)}$ such that
  – $x^{(t)}$ is a solution to the auxiliary primal,
  – $(y^{(t)}, s^{(t)})$ is a solution to its dual,
  – $\mu^{(t)} = (1-\delta) \cdot \mu^{(t-1)} = (1-\delta)^t \cdot \mu^{(0)}$, and $\sum_j \left(x_j^{(t)} s_j^{(t)}/\mu^{(t)} - 1\right)^2 \leq 1/4$.

  For $t \geq 1$, the difference between the primal and the dual objective value is exactly $(n+2)\mu^{(t)}$ (Claim 4). The gap decreases by a factor $1 - \delta = 1 - 1/(8\sqrt{n+2})$ in each iteration, and hence, can be made arbitrarily small.

$$\text{maximize } d^T y + (n+2)y_{m+1}, \text{ subject to} \quad \begin{aligned} A^T y + e y_{m+1} + s &= c, \\ \rho^T y + y_{m+1} + s_{n+2} &= M \\ y_{m+1} + s_{n+1} &= 0 \end{aligned} \tag{9}$$

with slack variables $s \geq 0, s_{n+1} \geq 0, s_{n+2} \geq 0$ and unconstrained variables y.

Box V.

In the next section, we will exploit this fact and show how to extract the optimal solution. Before doing so, we show the existence of an optimal solution.

Remark 6. Existence of an Optimal Solution: This paragraph requires some knowledge of calculus, namely continuity and accumulation point. Our sequence $(x^{(t)}, y^{(t)}, s^{(t)})$ has an accumulation point (this is clear for the sequence of x^i since the x-variables all lie between 0 and $n+2$ and we ask the reader to accept it for the others). Then there is a converging subsequence. Let (x^*, y^*, s^*) be its limit point. Then x^* and (y^*, s^*) are feasible solutions of the artificial primal and its dual respectively, and $x_i^* s_i^* = 0$ for all i by continuity.

5. Extracting an optimal solution

We will show how to round an approximate solution for the auxiliary problems for a sufficiently small μ to an optimal solution. This section is similar to [7, Theorem 5.3] and to the approach in [3, Section 3.3]. See also [12]. The auxiliary problem has $m+1$ constraints in $n+2$ variables. The auxiliary dual problem has $n+2$ constraints in $m+1+n+2$ variables. We use x to denote the variables of the auxiliary primal including x_{n+1} and x_{n+2}, and y and s for the variable vectors of the dual (including the additional variables). Moreover, we use A for the entire constraint matrix and b for the full right hand side. So A is $(m+1) \times (n+2)$, b is a $(m+2)$-vector and c is a $(n+2)$-vector.

Consider an iterate (x, y, s, μ). We will first show that $x_i \geq x_i^*/(4(n+2))$ and $s_i \geq s_i^*/(4(n+2))$ for all optimal solutions x^* and (y^*, s^*) (Lemma 2), i.e., if $x_i^* > 0$ ($s_i^* > 0$) for some i, then x_i (s_i) cannot become arbitrarily small. However, since $x_i s_i \leq 2\mu$ always and μ decreases exponentially, at least one of x_i or s_i has to become arbitrarily small. We use this observation to conclude that if x_i is sufficiently small (Lemma 3 quantifies what sufficiently small means) then $x_i^* = 0$ in every optimal primal solution. Similarly, if s_i is sufficiently small, then $s_i^* = 0$ in every optimal dual solution.

Let N be the set of indices for which we can conclude $x_i^* = 0$ and let B be the set of indices for which we can conclude $s_i^* = 0$. We show $B \cup N = \{1, \ldots, n\}$ and $B \cap N = \emptyset$. We split our last iterate $\bar{x}$ into two parts $\bar{x}_B$ and $\bar{x}_N$ accordingly, round the N-part to zero and recompute the B-part. Since the coordinates in the N-part are tiny, this has little effect on the B-part and hence the solution stays feasible. It stays optimal because of complementary slackness.

Lemma 2. *Let (x, y, s, μ) satisfy (I1)–(I3).*

1. *For all $i \in \{1, \ldots, n\}$: $x_i \geq x_i^*/(4(n+2))$ for every optimal solution x^* of the auxiliary primal.*
2. *For all $i \in \{1, \ldots, n\}$: $s_i \geq s_i^*/(4(n+2))$ for every optimal solution (y^*, s^*) of the auxiliary dual.*

Proof. By (I1) and (I2), x is a feasible solution of the auxiliary primal and (y, s) a feasible solution of the auxiliary dual. By (I3), we have $\sigma^2 = \sum_i (\frac{x_i s_i}{\mu} - 1)^2 \leq \frac{1}{4}$. Thus, $(\frac{x_i s_i}{\mu} - 1)^2 \leq \frac{1}{4}$, and hence, $\mu/2 \leq x_i s_i \leq 3\mu/2 < 2\mu$ for all i. Further, $x^T s = \sum_i x_i s_i < 2(n+2)\mu$.

Let x^* be any optimal solution of the primal. Then $c^T x \geq c^T x^*$. We apply Claim 1 first to the solution pair x and (y, s) and then to the pair x^* and (y, s) to obtain

$$x^T s = c^T x - b^T y \geq c^T x^* - b^T y = (x^*)^T s.$$

Consider any $i \in \{1, \ldots, n+2\}$ and assume $x_i < x_i^*/(4(n+2))$. Since $x_i s_i \geq \mu/2$, we have $s_i \geq \mu/(2x_i) > 2(n+2)\mu/x_i^*$, and hence

$$(x^*)^T s \geq x_i^* s_i > 2(n+2)\mu \geq x^T s \geq (x^*)^T s,$$

a contradiction.

Let (y^*, s^*) be any optimal solution of the dual. Then $b^T y^* \geq b^T y$. We apply Claim 1 first to the solution pair x and (y, s) and then to the pair x and (y^*, s^*) to obtain

$$x^T s = c^T x - b^T y \geq c^T x - b^T y^* = x^T s^*.$$

Consider any $i \in \{1, \ldots, n+2\}$ and assume $s_i < s_i^*/(4(n+2))$. Since $x_i s_i \geq \mu/2$, we have $x_i \geq \mu/(2s_i) > 2(n+2)\mu/s_i^*$, and hence

$$x^T s^* \geq x_i s_i^* > 2(n+2)\mu \geq x^T s \geq x^T s^*,$$

a contradiction. ■

The preceding Lemma implies strong duality, one of the cornerstones of linear programming theory.

Theorem 2 (*Strong Duality*). *For each i, either $x_i^* = 0$ in every optimal solution or $s_i^* = 0$ in every optimal solution. Thus, $c^T x^* - b^T y^* = (x^*)^T s^* = 0$.*

Proof. Let x^* and (y^*, s^*) be any pair of optimal solutions. Assume that there is an i such that $x_i^* s_i^* > 0$. Let (x, y, s, μ) satisfy the invariants (I1)–(I3). Then $x_i \geq x_i^*/(4(n+2))$ and $s_i \geq s_i^*(4(n+2))$ by Lemma 2. Thus $2\mu > x_i s_i \geq x_i^* s_i^*/(16(n+2)^2)$. For $\mu < x_i^* s_i^*/(32(n+2)^2)$, this is a contradiction. ■

Remark 7. We leave it to the reader to derive strong duality for the original primal and dual from this.

By the Strict Complementarity Theorem (see e.g. [8, pp 77–78] or [7, pp 20–21]), there are optimal solutions x^* and (y^*, s^*) in which $x_i^* > 0$ or $s_i^* > 0$ for every i. A Quantitative version of strict complementarity is next stated in Fact 3 (the proof is in Section 7).

Fact 3. *Let $Q = R/(n+2)$. Then there are optimal solutions x^* and (y^*, s^*) such that for all i either $x_i^* \geq Q$ and $s_i^* = 0$ or $s_i^* \geq Q$ and $x_i^* = 0$.*

The Rounding Procedure: Throughout this section x^* and (y^*, s^*) denote optimal solutions as in Fact 3. We run the iterative improvement algorithm until

$$\mu < \mu_f := R \cdot Q/(64(n+2)^2((m+1)U))^{m+2}. \tag{11}$$

Let $(\bar{x}, \bar{y}, \bar{s}, \bar{\mu})$ be the last iterate. Let

$$B = \{i \mid \bar{s}_i < Q/(4(n+2))\} \quad \text{and} \quad N = \{i \mid \bar{x}_i < Q/(4(n+2))\}.$$

Lemma 3. *$B \cup N = \{1, \ldots, n\}$, $B \cap N = \emptyset$, $x^*_i = 0$ and $\bar{x}_i < 8\bar{\mu}/Q$ for every $i \in N$ and $s^*_i = 0$ and $\bar{s}_i < 8\bar{\mu}/Q$ for every $i \in B$.*

Proof. Since $x_i s_i < 2\mu$ and $\mu \le Q^2/(32n^2)$, we have either $\bar{x}_i < Q/(4(n+2))$ or $\bar{s}_i < Q/(4(n+2))$. Thus $B \cup N = \{1, \ldots, n\}$. Since $\bar{x}_i \ge x^*_i/(4(n+2))$ and $\bar{x}_i \ge s^*_i/(4(n+2))$ and either $x^* \ge Q$ or $s^*_i \ge Q$, we have $B \cap N = \emptyset$. Consider any $i \in B$. Then $\bar{s}_i < Q/(4(n+2))$ and hence $s^*_i < Q$. Thus $s^*_i = 0$. Similarly, $i \in N$ implies $x^*_i = 0$. Finally, since $\bar{x}_i \bar{s}_i < 2\bar{\mu}$, we either have $\bar{x}_i \ge Q/(4(n+2))$ and $\bar{s}_i < 8\bar{\mu}/Q$ or $\bar{s}_i \ge Q/(4(n+2))$ and $\bar{x}_i < 8\bar{\mu}/Q$. ■

We split the variables x into x_B and x_N and the matrix A into A_B and A_N. Then our primal constraint system (ignoring the non-negativity constraints) becomes

$$A_B x_B + A_N x_N = b.$$

(x^*_B, x^*_N) and $(\bar{x}_B, \bar{x}_N)$ are solutions of this system, and $x^*_N = 0$ by Lemma 3. Thus $A_B x^*_B = b$.

Let us concentrate on the equation $A_B x_B = b$. If it has a unique solution, call it $\hat{x}_B$, then $\hat{x}_B = x^*_B$. We can find $\hat{x}_B$ by Gaussian elimination and $(\hat{x}_B, 0)$ will be the optimal solution and we are done.

What can we do if $A_B x_B = b$ has an entire solution set? Then the rank of the matrix A_B is smaller than the cardinality of B. Let $B_1 \subseteq B$ be such that the rank of the matrix A_{B_1} is equal to the cardinality of B_1 and let $B_2 = B \setminus B_1$. We can find B_1 by Gaussian elimination. Then our system becomes

$$A_{B_1} x_{B_1} + A_{B_2} x_{B_2} + A_N x_N = b.$$

For every choice of x_{B_2} and x_N this system has a unique solution[4] for x_{B_1}. Let $\hat{x}_{B_1}$ be the solution of

$$A_{B_1} \hat{x}_{B_1} + A_{B_2} \bar{x}_{B_2} = b \quad (x_N \text{ is set to zero and } x_{B_2} \text{ is set to } \bar{x}_{B_2}).$$

Subtracting this equation from $A_{B_1} \bar{x}_{B_1} + A_{B_2} \bar{x}_{B_2} + A_N \bar{x}_N = b$ yields

$$A_{B_1}(\bar{x}_{B_1} - \hat{x}_{B_1}) + A_N \bar{x}_N = 0.$$

The coordinates of $\bar{x}_N$ are bounded by $8\bar{\mu}/Q$ and hence the coordinates of $A_N \bar{x}_N$ are bounded by $8(n+2)U\bar{\mu}/Q = R/(8(n+2)((m+1)U)^{m+1})$ in absolute value. By the remark after Lemma 4 of Section 7, all coordinates of $\bar{x}_{B_1} - \hat{x}_{B_1}$ are bounded by $((m+1)U)^{m+1}$ times this number in absolute value, i.e., are bounded by $R/(8(n+2))$ in absolute value. Since $\bar{x}_i \ge R/(4(n+2))$ for every $i \in N$, we have $\hat{x}_{B_1} \ge 0$. Thus $\tilde{x} = (\hat{x}_{B_1}, \bar{x}_{B_2}, 0)$ is a feasible solution of (8). Since $\tilde{x}^T s^* = \sum_{i \in B} \tilde{x}_i s^*_i + \sum_{i \in N} \tilde{x}_i s^*_i = 0 + 0 = 0$, $\tilde{x}$ is an optimal solution to (8).

[4] Let $m' \le m$ be the rank of A_B. By row operations and permutation of columns, we can transform the system $A_B x_B + A_N x_N = b$ into

$$\begin{array}{ccccccc} I x_{B_1} & + & A'_{B_2} x_{B_2} & + & A'_N x_N & = & b' \\ 0 & + & 0 & + & A''_N x_N & = & b'', \end{array}$$

where I is a $m' \times m'$ identity matrix, A'_{B_2}, A'_N, and b' have m' rows, and A''_N and b'' have $m - m'$ rows. Since (x^*_B, x^*_N) is a solution to this system and $x^*_N = 0$, we have $b'' = 0$. Since $(\bar{x}_B, \bar{x}_N)$ is a solution to this system, we have further $A''_N \bar{x}_N = 0$. Thus for every choice of x_{B_2} and x_N this system has a unique solution for x_{B_1}.

6. Complexity

Let us assume that the initial value of μ is μ_0 and that we want to decrease μ to μ_f. Since every iteration decreases μ by the factor $(1-\delta)$, we have $\mu = (1-\delta)^r \mu_0$ after r iterations. The smallest r such that $(1-\delta)^r \le \mu_f$ is given by

$$\ln \frac{\mu_0}{\mu_f} = -r \ln(1-\delta) \approx -r(-\delta),$$

or equivalently,

$$r = O\left(\frac{1}{\delta} \log \frac{\mu_0}{\mu_f}\right) = O\left(\sqrt{n} \log \frac{\mu_0}{\mu_f}\right).$$

In (10), we defined

$$\mu_0^2 = 4\left(M^2 + \sum c_i^2\right) \le 4\left(\frac{16n^2U^2}{R^2} + nU^2\right) \le 68\frac{n^2U^2}{R^2}.$$

In (11), we defined μ_f. Thus, the number of iterations will be

$$\begin{aligned} r &= O\left(\sqrt{n}\log\frac{\mu_0}{\mu_f}\right) = O\left(\sqrt{n}\log\frac{n^2U^2/R^2}{RQ/(64(n+2)^2((m+1)U)^{m+2})}\right) \\ &= O\left(\sqrt{n}(\log n + m\log(mU)) + \log\frac{1}{R}\right) \\ &= O\left(\sqrt{n}(\log n + m\left(\log(mU)\right))\right), \end{aligned}$$

as $\log \frac{1}{R} = O(\log n + m(\log(mU)))$.

7. The proofs of Facts 1–3

In the previous sections, we used upper bounds on the components of an optimal solution and lower bounds on the nonzero components of an optimal solution. In this section, we derive these bounds. In this section, we assume more knowledge of linear algebra, namely, determinants and Cramer's rule, and some knowledge of geometry. Unless stated otherwise, we assume that all entries of A and b are integers bounded by U in absolute value.

The determinant of a $k \times k$ matrix G is a sum of $k!$ terms, namely,

$$\det G = \sum_\pi sign(\pi) \cdot g_{1\pi(1)} g_{2\pi(2)} \cdots g_{k\pi(k)}.$$

The summation is over all permutations π of k elements, $sign(\pi) \in \{-1, 1\}$, and the product corresponding to a permutation π selects the $\pi(i)$th element in row i for each i. Each product is at most U^k. As there are $k!$ summands, we have $|\det G| \le k!U^k \le (kU)^k$; see [5, pp 373–374], [6, p 75] or [8, pp 43–44].

Cramer's rule states that the solution of the equation $Gz = g$ (for a $k \times k$ non-singular matrix G) is $z_i = (\det G_i)/\det G$, where G_i is obtained by replacing the ith column of G by g.

Lemma 4. *Let $Gz = g$ be a linear system in k variables with a unique solution. Let z^* be the solution of the system. If all entries of G and g are integers bounded by U in absolute value then $\left|z^*_i\right| \le (kU)^k$ for all i and $z^*_i \neq 0$ implies $\left|z^*_i\right| \ge 1/(kU)^k$.*

Proof. Since the system has a unique solution there is a subsystem $G'z = g'$ consisting of k equations such that G' is non-singular and $G'z^* = g'$. Then $z^*_i = (\det G'_i)/\det G'$, where G'_i is obtained from G' by replacing the ith column of G by g'. Since all entries of G and g are integral, $\det G'$ is at least one in absolute value, $\det G'_i$ is at least one in absolute value if nonzero, and $\det G'_i \le (kU)^k$. The bounds follow. ■

If the entries of the right-hand side g are bounded by U' instead of U, the upper bound becomes $k^k U^{k-1} U'$.

Lemma 5. *Assume that (1) is feasible. Let x be a feasible solution with the maximum number of zero coordinates (equivalently the minimum number of nonzero coordinates).*[5] *Let B be the set of indices for which $x_i \neq 0$, and let A_B be the submatrix of A formed by the columns indexed by B. Then $A_B z = b$ has a unique solution, where the dimension of z is equal to the number of columns of A_B.*

If, in addition (1) is bounded, the same claim holds for an optimal solution with a maximum number of zero coordinates.

Proof. Let x_B be the restriction of x to the indices in B. Then $A_B x_B = b$. Assume there is a second solution x'_B of $A_B z = b$ with $x'_B \neq x_B$. Then all points $z(\lambda) = x_B + \lambda(x'_B - x_B)$, $\lambda \in \mathbb{R}$, satisfy $A_B z = b$. These points form a line. Consider the intersection z^* closest to x_B of this line with one of the coordinate planes $z_i = 0$; if there are several with the same distance choose one of them. Then $z^* \ge 0$ because we consider an intersection closest to x_B and $z^*_i = 0$ for at least one $i \in B$. Thus z^* is a feasible solution to (1) with one more zero coordinate, a contradiction to the definition of x.

If (1) is bounded, there is an optimal solution. Let x be an optimum solution with a maximum number of zero coordinates. Define x_B, x'_B, and $z(\lambda)$ as above. Since $x_B > 0$, the $z(\lambda)$ is feasible for small enough $|\lambda|$. Also $c_B^T z(\lambda) = c_B^T x_B + \lambda(c_B^T x'_B - c_B^T x_B)$; here c_B is the restriction of c to the indices in B. Since λ may be positive or negative, we must have $c_B^T x'_B = c_B^T x_B$ and hence $z(\lambda)$ is feasible and optimal as long as $z(\lambda) \ge 0$. The proof is now completed as in the preceding paragraph. ■

We can now give the proof of Facts 1, 2, and 3.

Proof (Fact 1)**.** Consider a feasible (optimal) solution x of (1) with a maximum number of zero coordinates. Then x is of the form $x = (x_B, x_N)$ with $x_N = 0$ and x_B being the unique solution to the system $A_B x_B = b$. Thus the coordinates of x_B are bounded by $(mU)^m$. ■

Proof (Fact 2)**.** Let x^* be an optimal vertex of the artificial primal (8). How small can a nonzero coordinate of x^* be? The constraint system is

$$\begin{array}{lcccccc} Ax & & & + & (\frac{1}{W}b - Ae)x_{n+2} & = & \frac{1}{W}b \\ e^T x & + & x_{n+1} & + & x_{n+2} & = & (n+2). \end{array}$$

Let B be the index set of the nonzero coordinates of x^*. Then x^*_B is the solution to a subsystem formed by $|B|$ columns of the above and this subsystem has a unique solution. For $i \in B$, $x^*_i = \det G_i / \det G$, where G is a nonsingular square matrix and G_i is obtained from G by replacing the ith column by the corresponding entries of the right hand side. In the system above, the entries in the column corresponding to x_{n+2} are bounded by $(n+1)U$, and all other entries are bounded by U. Since any product in the determinant formula for G can contain only one value of the column for x_{n+2}, we have $\left|\det G\right| \le (m+1)!(n+1)U^{m+1}$. Consider next $\det G_i$. We need to lower bound $\left|\det G_i\right|$. The matrix G_i may contain two columns with fractional values. If we multiply these columns with W, we obtain an integer matrix. Thus, $\left|\det G_i\right| \ge 1/W^2$ if nonzero. Thus

$$x^*_i \ge \frac{1}{W^2} \cdot \frac{1}{2n\left((m+1)U\right)^{m+1}} \ge \frac{1}{2n\left((m+1)U\right)^{3(m+1)}}. \quad ■ \tag{12}$$

Proof (Fact 3)**.** We prove the fact for the auxiliary primal. Let $\mathcal{O}$ be a smallest set of optimal vertices with the property that if for some i there is an optimal solution with $x^*_i > 0$, then $\mathcal{O}$ contains an optimal vertex with this property. Then $|\mathcal{O}| \le n+2$. Let $x^{**} = \frac{1}{|\mathcal{O}|}\sum_{x^* \in \mathcal{O}} x^*$ be the center of gravity of the vertices in $\mathcal{O}$. Then $x^{**}_i \ge x^*_i/(n+2)$ for every $x^* \in \mathcal{O}$. Thus $Q = R/(n+2)$ works. ■

Beyond the integral case If the entries of A and b are rational numbers, we write the entries in each column (or row) with a common denominator. Pulling them out brings us back to the integral case. For example,

$$\begin{vmatrix} 2/3 & 4/5 \\ 1/3 & 6/5 \end{vmatrix} = \frac{1}{15}\begin{vmatrix} 2 & 4 \\ 1 & 6 \end{vmatrix}.$$

Thus, if the determinant is nonzero, it is at least 1/15.

Acknowledgments

The first author thanks Andreas Karrenbauer and Ruben Becker for intensive discussions and Andreas Karrenbauer for teaching from an earlier draft of this note. The work of the second author was inspired by an informal lecture given by Nisheeth Vishnoi at IIT Kanpur. The second author also thanks the students of CS602 (2014–15 and 2015–2016 batches) for their helpful comments and questions. Thanks also to Romesh Saigal for very prompt replies to queries.

The authors also thank Andreas Karrenbauer for providing the proof of Lemma 1.

Appendix. Result from algebra

Assume that A is $m \times n$ matrix and the rank of A is m, with $m < n$. Then, all m rows of A are linearly independent. Or, $\alpha_1 A_1 + \alpha_2 A_2 + \cdots + \alpha_m A_m = 0$ (0 here being a row vector of size n) has only one solution $\alpha_i = 0$. Thus, if x is any $m \times 1$ matrix (a column vector of size m), then $x^T A = 0$ implies $x = 0$. Note that $(x^T A)^T = A^T x$. Thus, $A^T x = 0$ implies $x = 0$.

As A is $m \times n$ matrix, A^T will be $n \times m$ matrix. The product AA^T will be an $m \times m$ square matrix.

Consider the equation $(AA^T)x = 0$. Pre-multiplying by x^T we get $x^T AA^T x = 0$ or $(A^T x)^T (A^T x) = 0$. Now, $(A^T x)^T (A^T x)$ is the squared length of the vector $A^T x$. If a vector has length zero,

[5] Consider minimize 0 subject to $x_1 + x_2 = 1$, $x_1 \ge 0$ and $x_2 \ge 0$. The feasible solutions $(0, 1)$ and $(1, 0)$ have one nonzero coordinate. The feasible solutions (x_1, x_2) with $x_1 > 0$ and $x_2 > 0$ and $x_1 + x_2 = 1$ have two nonzero coordinates.

all its coordinates must be zero. Thus, $A^T x = 0$, and hence, $x = 0$ by the preceding paragraph.

Thus, the matrix AA^T has rank m and is invertible.

Also observe that if X is a diagonal matrix (with all diagonal entries non-zero) and if A has full row-rank, then AX will also have full row-rank. Basically, if the entries of X are $x_1, x_2, \ldots, x_m$ then the matrix AX will have rows as $x_1A_1, x_2A_2, \ldots, x_mA_m$ (i.e., ith row of A gets scaled by x_i). If rows of AX are not independent, then there are βs (not all zero) such that $\beta_1 x_1 A_1 + \beta_2 x_2 A_2 + \cdots + \beta_m x_m A_m = 0$, or there are αs (not all zero) such that $\alpha_1 A_1 + \alpha_2 A_2 + \cdots + \alpha_m A_m = 0$ with $\alpha_i = \beta_i x_i$.

REFERENCES

[1] T. Terlaky, An easy way to teach interior-point methods, Eur. J. Oper. Res. 130 (2001) 1–19.

[2] G. Lesaja, Introducing interior-point methods for introductory operations research courses and/or linear programming courses, Open. Oper. Res. J. 3 (2009) 1–12.

[3] C. Roos, T. Terlaky, J.-P. Vial, Interior Point Methods for Linear Optimization, second ed., Springer, 2006.

[4] R.J. Vanderbei, Linear Programming: Foundations and Extension, first ed., Kluwer Academic Publishers, 1997, (third ed: Springer).

[5] D. Bertsimas, J.N. Tsitsiklis, Introduction to Linear Optimization, Athena Scientific, 1997.

[6] H. Karloff, Linear Programming, Birkhauser, 1991.

[7] Y. Ye, Interior Point Algorithms, Wiley, 1997.

[8] R. Saigal, Linear Programming, in: A Modern Integrated Analysis, Kluwer, 1995.

[9] N.K. Vishnoi, Convex Optimization Notes, http://theory.epfl.ch/vishnoi/Nisheeth-VishnoiFall2014-ConvexOptimization.pdf.

[10] S. Mizuno, M. Kojima, M.J. Todd, Infeasible-interior-point primal–dual potential-reduction algorithms for linear programming, SIAM J. Optim. 5 (1995) 52–67.

[11] Y. Zhang, On the convergence of a class of infeasible interior-point methods for the horizontal linear complementarity problem, SIAM J. Optim. 5 (1994) 208–227.

[12] H.J. Greenberg, The use of the optimal partition in a linear programming solution for postoptimal analysis, Oper. Res. Lett. 15 (1994) 179–185.