

Still Simpler Way of Introducing Interior-Point method for Linear Programming

Sanjeev Saxena*
Dept. of Computer Science and Engineering,
Indian Institute of Technology,
Kanpur, INDIA-208 016

November 29, 2014

Abstract

Linear Programming is now included in Algorithm undergraduate and postgraduate courses for Computer Science majors. It is possible to teach interior-point methods directly with just minimal knowledge of Algebra and Matrices.

1 Introduction

Terlaky [3] and Lesaja [1] have suggested simple ways to teach interior-point methods. In this paper a still simpler way is being suggested. Most material required to teach interior-point methods is available in popular text books [2, 4]. However, these books assume knowledge of Calculus, which is not really required. In this paper, it is suggested if appropriate material is selected from these books then it becomes very easy to teach interior-point methods as the first or only method for Linear programming in Computer Science Courses.

Canonical Linear Programming Problem is to

minimise cx^T subject to $Ax = b$ and $x \geq 0$.

Here A is an $n \times m$ matrix, b and c are n -dimensional and x is an m -dimensional vector.

REMARK 1 maximise cx^T is equivalent to minimise $-cx^T$.

REMARK 2 Constraints of type $\alpha_1x_1 + \dots + \alpha_nx_n \leq \beta$ can be replaced by $\alpha_1x_1 + \dots + \alpha_nx_n + \gamma = \beta$ with a new (slack) variable $\gamma \geq 0$. Similarly constraints of type $\alpha_1x_1 + \dots + \alpha_nx_n \geq \beta$ can be replaced by $\alpha_1x_1 + \dots + \alpha_nx_n - \gamma = \beta$ with (surplus) variable $\gamma \geq 0$.

Thus, we assume that there are n constraints and m variables, with $m > n$ (more variables and fewer constraints)— basically slack or surplus are added or subtracted to convert inequalities into equalities.

*E-mail: ssax@iitk.ac.in

We first use pivoting to make first term of all but the first equation as zero. Basically, we multiply i^{th} equation by $-\frac{a_{i1}}{a_{11}}$ and subtract the first equation. In similar way we make first two terms of all but the first two equations as zero— multiply i^{th} equation (for $i \neq 2$) by $-\frac{a_{2i}}{a_{21}}$ and subtract the second equation. And so on. In case, if in any equation all coefficients become zero, we drop those equations. As a result, in the end all remaining equations will be linearly independent. Or the resulting matrix will have full row rank.

REMARK We may have to interchange two columns (interchange two variables), in case, for example, if a diagonal term of an equation becomes zero.

From convexity, it is sufficient to obtain a locally optimal solution, as local optimality will imply global optimality.

We consider another problem, the “dual problem” which is

maximise by^T subject to $A^T y + s = c$, with slack variables $s \geq 0$ and variables y are unconstrained.

Claim 1 $by^T \leq cx^T$. The equality will hold if and only if, $s_i x_i = 0$ for all i s.

REMARK Thus if value of both primal and dual are the same, then both are optimal.

Proof: $s = c - A^T y$, or $x^T s = x^T c - x^T (A^T y) = c^T x - (x^T A^T) y = c^T x - b^T y$. As $x, s \geq 0$, we have $x^T s \geq 0$ or $c^T x \geq b^T y$.

Equality will hold if $x^T s = 0$ or $\sum_i s_i x_i = 0$ but as $s_i, x_i \geq 0$, we want each term (product) $s_i x_i = 0$. ■

Thus, if we are able to find a solution of following equations (last one is not linear, else, an inversion of matrix would have been sufficient), we will be getting optimal solutions of both the original and the dual problems.

$$Ax = b, A^T y + s = c, x_i s_i = 0$$

subject to $x \geq 0, s \geq 0$.

We will relax the last condition to get something like (duality gap):

$$x_i s_i \approx \mu$$

with parameter $\mu \geq 0$. Thus, we will be solving (the exact last equation will be derived in the next section):

$$Ax = b, A^T y + s = c, x_i s_i \approx \mu$$

subject to $x \geq 0, s \geq 0$.

REMARK Thus, $by^T - cx^T \approx m\mu$. If μ is very small, then in the case of rationals, the solution will be exact.

Use of Newton Raphson Method

We will use the Newton-Raphson method[2]. Let us choose the next values as $x + h, y + k, s + f$. Then we want:

1. $A(x + h) = b$ or $Ax + Ah = b$ but as $Ax = b$, we get $Ah = 0$.
2. $A^T(y + k) + (s + f) = c$, from $A^T y + s = c$, we get $A^T k + f = c - A^T y - s = 0$ or $A^T k + f = 0$
3. $(x_i + h_i)(s_i + f_i) \approx \mu$ or $x_i s_i + h_i s_i + f_i x_i + h_i f_i \approx \mu$. Or approximately, $x_i s_i + h_i s_i + f_i x_i = \mu$ (neglecting the non-linear $h_i f_i$ term). Thus, the equation we will be solving is

$$h_i s_i + f_i x_i = \mu - x_i s_i$$

Thus, we have a system of linear equations for h_i, k_i, f_i . We next show that these can be solved by “inverting” a matrix.

But first observe that from the third equation,

Observation 1 $(x_i + h_i)(s_i + f_i) = \mu + h_i f_i$

Theorem 1 *Following equations have a unique solution:*

1. $Ah = 0$
2. $A^T k + f = 0$
3. $h_i s_i + f_i x_i = \mu - x_i s_i$

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. in X the diagonal entries will be x_1, x_2, \dots, x_m). We will also use e to denote a column vector of all ones (usually of length m).

Then in the new notation, the last equation is:

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

Let us look at this equation in more detail.

$$\begin{aligned} Sh + Xf &= \mu e - XSe \\ h + S^{-1}Xf &= S^{-1}\mu e - S^{-1}XSe \text{ pre-multiply by } S^{-1} \\ h + S^{-1}Xf &= \mu S^{-1}e - X \cancel{S^{-1}} e \text{ diagonal matrices commute} \\ h + S^{-1}Xf &= \mu S^{-1}e - x \text{ as } Xe = x \\ Ah + AS^{-1}Xf &= \mu AS^{-1}e - Ax \text{ pre-multiply by } A \\ AS^{-1}Xf &= \mu AS^{-1}e - b \text{ but } Ax = b \text{ and } Ah = 0 \\ -AS^{-1}XA^T k &= \mu AS^{-1}e - b \text{ using } f = -A^T k \\ b - AS^{-1}e &= (AS^{-1}XA^T)k \end{aligned}$$

As XS^{-1} is diagonal with positive items and as A has full rank, if $W = \sqrt{XS^{-1}}$ (each diagonal term is $\sqrt{x_i/s_i}$) then $AS^{-1}XA^T = (AW)(AW)^T = AW^2A^T$ is invertible (see appendix). The last equation can thus be used to get the value of matrix k after inverting the matrix $AS^{-1}XA^T$, or

$$k = (AS^{-1}XA^T)^{-1} (b - AS^{-1}e)$$

Then we can find f from $f = -A^T k$.

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, the above system has a unique solution. ■

Claim 2 $\sum_i h_i f_i = 0$ or equivalently $h^T f = f^T h = 0$

Proof: As $A^T k + f = 0$, we get $h^T A^T k + h^T f = 0$ but $h^T A^T = (Ah)^T = 0$, hence $h^T f = 0$ follows. ■

Invariants in each Iteration

We will maintain following invariants:

1. $Ax^T = b$, with $x > 0$ (strict inequality)
2. $A^T y + s = c$ with $s > 0$ (strict inequality)
3. If μ is the ‘‘approximate duality gap’’ then $\sigma \leq \frac{2}{3} < \sqrt{3} - 1$ where $\sigma^2 = \sum_i \left(\frac{x_i s_i}{\mu} - 1 \right)^2$.

At end of this iteration we want duality gap $\mu' \leq (1 - \delta)\mu$. We will see that δ can be chosen as $\delta = \Theta\left(\frac{1}{\sqrt{m}}\right)$.

We first show that strict inequality invariants hold (in σ' we have $x + h, s + f$ and same μ):

Fact 1 If $\sigma' < 1$ then $x + h > 0$ and $s + f > 0$

Proof: We first show that the product $(x_i + h_i)(s_i + f_i)$ is term-wise positive. From Observation 1, $(x_i + h_i)(s_i + f_i) = \mu + h_i s_i$.

From $\sigma' < 1$ we get $\sigma'^2 < 1$. But (using Observation 1):

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

As the sum is at most one, it follows that each term of the summation must be less than one, or $\left| \frac{h_i s_i}{\mu} \right| < 1$ or $-\mu < h_i s_i < \mu$. In particular $\mu + h_i s_i > 0$.

Thus the product $(x_i + h_i)(s_i + f_i)$ is term-wise positive.

Assume for contradiction that both $x_i + h_i < 0$ and $s_i + f_i < 0$. But as $s_i > 0$ and $h_i > 0$, we have $s_i(x_i + h_i) + x_i(s_i + f_i) < 0$, or $\mu + x_i s_i < 0$. Which is impossible as μ, x_i, s_i are all non-negative, a contradiction. ■

We have to still show that the ‘‘approximate duality gap’’ μ decreases as desired.

Let us define three new variables:

$$H_i = h_i \sqrt{\frac{s_i}{x_i \mu}}$$

$$F_i = f_i \sqrt{\frac{x_i}{s_i \mu}}$$

Observe that $\sum_i H_i F_i = \sum \frac{h_i f_i}{\mu} = 0$ (see Claim 2).

$$\begin{aligned} H_i + F_i &= h_i \sqrt{\frac{s_i}{x_i \mu}} + f_i \sqrt{\frac{x_i}{s_i \mu}} \\ &= \sqrt{\frac{1}{x_i s_i \mu}} (h_i s_i + f_i x_i) \\ &= \sqrt{\frac{1}{x_i s_i \mu}} (\mu - x_i s_i) \\ &= \sqrt{\frac{\mu}{x_i s_i}} \left(1 - \frac{x_i s_i}{\mu}\right) \\ &= -\sqrt{\frac{\mu}{x_i s_i}} \left(-1 + \frac{x_i s_i}{\mu}\right) \end{aligned}$$

From, the proof of Fact 1, we also observe that $\sigma'^2 = \sum_i \left(\frac{h_i f_i}{\mu}\right)^2$, or $\sigma'^2 = \sum_i (H_i F_i)^2$

And finally

$$\begin{aligned} \sigma'^2 &= \sum_i (H_i F_i)^2 \\ &\leq \sum_i \frac{(H_i^2 + F_i^2)^2}{4} \text{ using AM-GM inequality-see Appendix} \\ &\leq \frac{1}{4} \left(\sum_i (H_i^2 + F_i^2) \right)^2 \text{ more positive terms} \\ &\leq \frac{1}{4} \left(\sum_i (H_i + F_i)^2 \right)^2 \text{ from Claim 2} \\ &= \frac{1}{4} \left(\sum_i \frac{\mu}{x_i s_i} \left(\frac{x_i s_i}{\mu} - 1 \right)^2 \right)^2 \\ &\leq \left(\max \frac{\mu}{x_i s_i} \right)^2 \frac{1}{4} \left(\sum_i \left(\frac{x_i s_i}{\mu} - 1 \right)^2 \right)^2 \\ &\leq \frac{\sigma^4}{4} \left(\max \frac{\mu}{x_i s_i} \right)^2 \end{aligned}$$

As $\sigma^2 = \sum \left(\frac{x_i s_i}{\mu} - 1 \right)^2$, each individual term is at most σ^2 or

$$\left| \frac{x_i s_i}{\mu} - 1 \right| \leq \sigma$$

Thus,

$$\begin{aligned} -\sigma &\leq \frac{x_i s_i}{\mu} - 1 \leq \sigma \\ 1 - \sigma &\leq \frac{x_i s_i}{\mu} \leq 1 + \sigma \end{aligned}$$

In particular $\frac{\mu}{x_i s_i} \leq \frac{1}{1-\sigma}$ or

$$\max \frac{\mu}{x_i s_i} \leq \frac{1}{1-\sigma}$$

Thus, $\sigma'^2 \leq \left(\frac{1}{1-\sigma}\right)^2 \frac{\sigma^4}{4}$ or $\sigma' \leq \frac{1}{2} \frac{\sigma^2}{1-\sigma}$

We summarise our observations as:

Observation 2

$$\sigma' \leq \frac{1}{2} \frac{\sigma^2}{1-\sigma}$$

For $\sigma' < 1$, $\frac{\sigma^2}{1-\sigma} < 2$ or $\sigma^2 < 2 - 2\sigma$ or $\sigma^2 + 2\sigma - 2 < 0$ or $\sigma = \frac{-2 \pm \sqrt{2^2 - 4 \times 1 \times (-2)}}{2} = \frac{-2 + \sqrt{12}}{2} = -1 + \sqrt{3} = \sqrt{3} - 1$.

REMARK Thus $\sigma \leq \frac{2}{3}$ is more than enough.

Let us finally try to get bounds on δ (and hence μ). Let us assume $\mu' = \mu(1-\delta)$ then if σ'' corresponds to $x+h, s+f$ and μ' we have

$$\begin{aligned} \sigma''^2 &= \sum_i \left(\frac{(x_i + h_i)(s_i + f_i)}{\mu(1-\delta)} - 1 \right)^2 \\ &= \sum_i \left(\frac{(x_i + h_i)(s_i + f_i) - \mu(1-\delta)}{\mu(1-\delta)} \right)^2 \\ &= \sum_i \left(\frac{(x_i + h_i)(s_i + f_i) - \mu}{\mu(1-\delta)} + \frac{\delta}{1-\delta} \right)^2 \\ &= \sum_i \left(\frac{h_i f_i}{\mu(1-\delta)} + \frac{\delta}{1-\delta} \right)^2 \quad \text{From Observation 1} \\ &= \frac{1}{(1-\delta)^2} \sum_i \left(\frac{h_i f_i}{\mu} + \delta \right)^2 \\ &= \frac{1}{(1-\delta)^2} \left(\sum_i \left(\frac{h_i f_i}{\mu} \right)^2 + m\delta^2 + \frac{2\delta}{\mu} \sum h_i f_i \right) \\ &= \frac{1}{(1-\delta)^2} \left(\sum_i \left(\frac{h_i f_i}{\mu} \right)^2 + m\delta^2 \right) \quad \text{From Claim 2} \\ &= \frac{1}{(1-\delta)^2} (\sigma'^2 + m\delta^2) \end{aligned}$$

Thus observe that

Observation 3

$$\sigma'' = \frac{1}{1-\delta} \sqrt{\sigma'^2 + m\delta^2}$$

We want to choose δ such that $\sigma'' \leq \frac{2}{3}$. As $\sigma'' = \frac{1}{1-\delta} \sqrt{\sigma'^2 + m\delta^2} \leq \frac{1}{1-\delta} \sqrt{m\delta^2} = \frac{\delta\sqrt{m}}{1-\delta}$

We want $\frac{\delta\sqrt{m}}{1-\delta} < \frac{2}{3}$ or $\frac{\delta}{1-\delta} < \frac{2}{3\sqrt{m}}$. We can thus choose $\delta = \frac{1}{4\sqrt{m}}$.

Summary

Let us assume that initial duality gap is μ_0 and final duality gap is μ_f , as after each iteration, $\mu' \leq (1 - \delta)\mu$, thus after r iterations, $\mu_f \leq (1 - \delta)^r \mu_0$, or

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

or

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

As $1 - \sigma \leq \frac{x_i s_i}{\mu} \leq 1 + \sigma$, we have (in last inequality we use $\sigma < \frac{2}{3}$).

$$\mu(1 - \sigma) \leq x_i s_i \leq \mu(1 + \sigma) < \frac{5}{3}\mu$$

Thus, when μ becomes very small, even the products $x_i s_i$ s will be very small. The above method will give a polynomial time algorithm even if $\mu_0 = 2^{m^2}$ and $\mu_f = \frac{1}{2^{m^2}}$.

Initial Solution: Phase 1 Method

Consider a starting solution x_0 s.t. all components of x_0 are strictly positive (say all $x_i = \frac{1}{n}$). Define a vector $\rho = b - Ax_0$.

Consider the following problem[4]:

minimise x_{n+1} (a new variable) subject to

$$Ax + x_{n+1}\rho = b$$

with $x \geq 0, x_{n+1} \geq 0$.

Clearly, $(x_0, 1)^T$ is a feasible solution.

We can solve this problem by the previous method. If the optimal solution has $x_{n+1} > 0$ then the original problem is unfeasible.

REMARK We are ignoring the case when some x_i (for $i \leq n$) turns out to be zero.

Then using the condition $x_i s_i = \mu$, we fix s_i . And finally from $A^T y + s = c$ we get $A^T y = c - s$ or $(AA^T)y = A(c - s)$ but as (AA^T) is invertible (see appendix), we get $y = (AA^T)^{-1}A(c - s)$.

REMARK Variables y are unconstrained. They can either be positive or negative or zero.

Acknowledgement

This work was inspired by an informal lecture given by Nisheeth Vishnoi at IIT Kanpur. I also wish to thank students of CS602 (2014-2015 batch) for their helpful comments and questions when I was teaching this material.

References

- [1] Goran Lesaja, Introducing Interior-Point Methods for Introductory Operations Research Courses and/or Linear Programming Courses, The Open Operational Research Journal, 2009, 3, 1-12.
- [2] C.Roos, T Terlaky and J-P Vial, Interior Point Methods for Linear Optimization, 2nd Ed, 2006, Springer.
- [3] Tamas Terlaky, An easy way to teach interior-point methods, European J of Operational Research, 130 (2001), 1-19
- [4] R.J.Vanderbei, Linear Programming: Foundations and Extension, 1st Ed: Kluwer Academic Publishers, 1997 (3rd Ed: Springer).

Appendix: Result from Algebra

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

As A is $n \times m$ matrix, A^T will be $m \times n$ matrix. The product AA^T will be an $n \times n$ square matrix. Let y^T be an $n \times 1$ matrix (or y is a row-vector of size n).

Consider the equation $(AA^T)y^T = 0$. Pre-multiplying by y we get $yAA^T y^T = 0$ or $(yA)(yA)^T = 0$ or the dot product $\langle yA, yA \rangle = 0$ which, for real vectors (matrices) means, that each term of yA is (individually) zero, or y is identically zero.

Thus, the matrix AA^T has rank n 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 entries of X are x_1, x_2, \dots, x_n then the matrix AX will have rows as $x_1 A_1, x_2 A_2, \dots, x_n A_n$ (i.e., i^{th} row of A gets scaled by x_i). If rows of AX are not independent then there are β s (not all zero) such that:
 $\beta_1 x_1 A_1 + \beta_2 x_2 A_2 + \dots + \beta_n x_n A_n = 0$, or there are α s (not all zero) such that:
 $\alpha_1 A_1 + \alpha_2 A_2 + \dots + \alpha_n A_n = 0$ with $\alpha_i = \beta_i x_i$.

Arithmetic-Geometric Inequality follows from:

$$(x + y)^2 \geq (x + y)^2 - (x - y)^2 = 4xy$$

or taking square roots (of the first and the last) $x + y \geq 2\sqrt{xy}$