

How to use Farkas' lemma to say something important about linear infeasible problems

MOSEK Technical report: TR-2011-1.

Erling D. Andersen*

12-September-2011

Abstract

When formulating a linear optimization problem it is easy accidentally to create an infeasible problem. In such a case it is obvious to ask: How can the infeasibility be identified and, possibly, repaired.

A common, but dangerous, practice is to use the solution reported by the optimization algorithm (read: *optimization software*) to figure out which constraints are causing the infeasibility. This practice used to make some sense 20 years or more ago when linear optimization problem was exclusively optimized with the phase 1 and phase 2 primal simplex algorithm. However, nowadays it is more common to employ either the dual simplex algorithm or an interior-point algorithm. In both cases the primal solution is somewhat arbitrary and perhaps not much better than the all-zeros solution. Therefore, what kind information should a user of linear optimization require and expect from an linear optimization algorithm? In this note we will argue that the Farkas' certificate of infeasibility is the answer.

1 Introduction

The linear optimization problem

$$\begin{aligned} & \text{minimize} && x_1 \\ & \text{subject to} && x_1 \leq 1, \\ & && x_1 \geq 2, \end{aligned} \tag{1}$$

is clearly primal infeasible, i.e. the problem does not have a solution. There are several possible ways to repair the problem. For instance, one of the constraints may be removed, or the right-hand side of the constraints may be changed appropriately. In this simple case it is easy to discover the infeasibility and figure out a repair. However, frequently infeasible problems are much larger and hence locating the infeasibility “by hand” can be almost impossible.

*MOSEK ApS, Fruebjergvej 3, Box 16, 2100 Copenhagen, Denmark. Email: e.d.andersen@mosek.com

2 The Farkas' certificate

If somebody says that (1) is feasible, then all that is needed to certify the claim is a feasible solution x . Moreover, it is simple to check the feasibility claim using the feasible solution. Similarly, a certificate of infeasibility should be simple and easy to verify. Fortunately, a certificate of infeasibility exists and is specified by the well-known *Farkas' lemma*. Indeed, Farkas' Lemma states that the linear optimization problem

$$\begin{aligned} & \text{minimize} && c^T x \\ & \text{subject to} && Ax = b, \\ & && x \geq 0 \end{aligned} \tag{2}$$

is infeasible if and only if there exists a y such that

$$\begin{aligned} b^T y &> 0, \\ A^T y &\leq 0. \end{aligned} \tag{3}$$

In other words any y satisfying (3) is a certificate of the primal infeasibility. Note it is easy to verify that a Farkas' certificate y^* is valid because it corresponds to checking the conditions

$$b^T y^* > 0 \tag{4}$$

and

$$A^T y^* \leq 0 \tag{5}$$

which only requires simple linear algebra.

It may seem strange that y^* is a certificate of infeasibility but it is easy to prove that it is the case. Consider a vector y^* that satisfies (4) and (5). If we assume that (2) is feasible which implies that there exists an x^* such that

$$Ax^* = b$$

and

$$x^* \geq 0.$$

Since $x^* \geq 0$ and $A^T y^* \leq 0$ then $(y^*)^T Ax^* \leq 0$ implying that

$$\begin{aligned} 0 &\geq (y^*)^T Ax^* \\ &= (y^*)^T b \\ &= b^T (y^*)^T \\ &> 0, \end{aligned}$$

which is a contradiction. Therefore, if an infeasibility certificate exists, then (2) cannot be feasible since we would otherwise have a contradiction.

A simple generalization of Farkas' Lemma is that

$$\begin{aligned} & \text{minimize} && c^T x \\ & \text{subject to} && A_1 x = b_1, \\ & && A_2 x \leq b_2, \\ & && A_3 x \geq b_3, \\ & && x \geq 0 \end{aligned} \tag{6}$$

is infeasible if and only if there exists a (y_1, y_2, y_3) such that

$$\begin{aligned} b_1^T y_1 + b_2^T y_2 + b_3^T y_3 &> 0, \\ A_1^T y_1 + A_2^T y_2 + A_3^T y_3 &\leq 0, \\ y_2 &\leq 0, \\ y_3 &\geq 0. \end{aligned} \tag{7}$$

Using this generalization the Farkas' certificate for the example (1) is

$$\begin{aligned} y_1 + 2y_2 &> 0 \\ y_1 &\leq 0, \\ y_2 &\geq 0 \end{aligned} \tag{8}$$

and hence $y_1 = -1$ and $y_2 = 1$ is a valid certificate.

Observe that an infeasibility certificate is not unique because if it is multiplied by any strictly positive number then it is still a certificate. Moreover, the problem

$$\begin{aligned} \text{minimize} \quad & x_1 \\ \text{subject to} \quad & x_1 \leq 1, \\ & x_1 \geq 2, \\ & x_2 \leq 1, \\ & x_2 \geq 10, \end{aligned} \tag{9}$$

illustrates that a certificate may only pinpoint one of the infeasibilities i.e. both $y = (-1, 2, 0, 0)$ and $y = (0, 0, -1, 2)$ are certificates.

The infeasibility certificate is a property of the optimization problem rather than of the algorithm, therefore it is reasonable to request an infeasibility certificate from an algorithm whenever it claims a problem is infeasible. Furthermore, since the properties of an infeasibility certificate is algorithm-independent, decisions based on infeasibility certificates will be similarly algorithm independent.

Note that an infeasibility certificate can always be reported in place of the dual solution because the infeasibility certificate has the same dimensions as the dual solution.

3 How to use the Farkas' certificate

Not only can the Farkas' certificate be used to certify that a problem is infeasible, it can also be used to pinpoint the cause infeasibility. Usually, if a problem is infeasible we would like to repair it, or at least know which part of the problem is causing the infeasibility. Take the example (1): We may think that changing the problem to

$$\begin{aligned} \text{minimize} \quad & x_1 \\ \text{subject to} \quad & x_1 \leq 1, \\ & x_1 \geq 1.5, \end{aligned} \tag{10}$$

will remove the infeasibility (we have changed the right-hand side for the second constraint from 2 to 1.5). However, the revised Farkas' conditions are

$$\begin{aligned} y_1 + 1.5y_2 &> 0, \\ y_1 &\leq 0, \\ y_2 &\geq 0 \end{aligned} \tag{11}$$

and it is seen that the old certificate $y_1 = -1$ and $y_2 = 1$ is still a valid certificate. Note that if we had changed the right-hand side of the second constraint to 1 then Farkas' certificate no longer would be valid. Therefore, in general when repairing an infeasible problem it should be changed at least so much that the certificate of infeasibility is no longer valid because otherwise the problem stays infeasible.

Another **important** observation is that if $y_i = 0$ then the i th constraint is not involved in the pinpointed infeasibility since if the i th constraint is removed from the problem and y_i is removed from the vector y , then the compressed y is still an infeasibility certificate. Indeed, assume y is an infeasibility certificate to (2) and define

$$\mathcal{I} := \{i : y_i \neq 0\}. \tag{12}$$

Next, consider the relaxed problem

$$\begin{aligned} A_{\mathcal{I},:}x &= b_{\mathcal{I}} \\ x &\geq 0 \end{aligned}$$

where all constraints for which y_i is zero have been removed. It is easy to verify that

$$\begin{aligned} b_{\mathcal{I}}^T y_{\mathcal{I}} &> 0 \\ A_{\mathcal{I},:}^T y_{\mathcal{I}} &\leq 0, \end{aligned} \tag{13}$$

and hence that $y_{\mathcal{I}}$ is a Farkas' certificate for the relaxed problem.

In general, it can be hoped that the set \mathcal{I} contains only few elements, which would imply that we have located a small set of constraints causing the infeasibility, and hence making a repair easy.

We have implicitly made the assumption that the constraint $x \geq 0$ is not causing the infeasibility. However, it can be verified that

$$\mathcal{J} := \{j : A^T y < 0\}$$

then

$$\begin{aligned} A_{\mathcal{J},:}x &= b_{\mathcal{J}} \\ x_{\mathcal{J}} &\geq 0 \end{aligned}$$

is still infeasible. Hence, potentially many of the simple inequality constraints is irrelevant for the infeasibility.

4 An infeasibility report

One way to exploit the infeasibility certificates is to generate to generate an infeasibility report looking something like this:

Equality 5 is important for the infeasibility. $y[5]$ is 23.67.
 Equality 15 is important for the infeasibility. $y[15]$ is 1.0.

based on the sets \mathcal{I} and \mathcal{J} . Here we have assumed that $\mathcal{I} = \{5, 15\}$.

Now, the definition of \mathcal{I} in (12) works well in theory, but in practice something like

$$\mathcal{I} := \left\{ i : |y_i| > \varepsilon \left(\left\| \begin{bmatrix} b_i \\ A_{i,:}^T \end{bmatrix} \right\|_{\infty} \right)^{-1} \right\}. \quad (14)$$

is better, where ε is a small constant, typically $\varepsilon \in [1.0e - 8, 1.0e - 5]$. The definition of the set \mathcal{I} in (14) is independent of the scaling of the constraints which is an attractive property.

5 The Farkas' certificate and column generation

A common practice when solving an optimization problem with many more variables than constraints is to employ column generation. The idea is solve the so-called restricted problem

$$\begin{aligned} & \text{minimize} && c_{\mathcal{J}}^T x_{\mathcal{J}} \\ & \text{subject to} && A_{:, \mathcal{J}} x_{\mathcal{J}} = b, \\ & && x_{\mathcal{J}} \geq 0 \end{aligned} \quad (15)$$

where \mathcal{J} contains a small subset of all the variables. It is well-known that if the restricted problem has an optimal solution and

$$c_j - A_{:,j}^T y \geq 0,$$

where y is an optimal dual solution to the restricted problem, then the optimal solution to the full problem has been located. Otherwise, one or more of the variables for which it holds

$$c_j - A_{:,j}^T y < 0$$

are added to the restricted problem. This process is iterated and will terminate in a finite number iterations if the number variables in the full problem is finite.

The process requires that the restricted problem is feasible but that requirement can be relaxed as follows. Assume that the restricted problem is infeasible and an infeasibility certificate is obtained, then variables for which it holds that

$$0c_j - A_{:,j}^T y < 0$$

should be added to the restricted problem. Such variables will invalidate the infeasibility certificate and hence move the restricted problem closer to a feasible problem. Hence, in column generation there is no need to assume that the initial restricted problem is feasible since a slight modification of the column generation procedure makes it possible to build a feasible restricted problem.

6 Dual infeasibility

Since the dual problem of a linear problem is itself a linear problem, the ideas presented in this note can also be applied to the dual problem. Dual infeasibility is therefore not discussed separately here.

7 Discussion

For an infeasible linear optimization problem there is no well-defined solution to report. This follows from the fact that the constraints of the problems has no solution. In addition, even if such a solution exists, it would most likely not be a certificate of the infeasible status, and it would be hard to verify the infeasible status from such a solution. This implies that users of optimization for linear problems should be very careful about using the reported infeasible solution for anything.

8 Conclusion

This note has argued that if a linear optimization problem is infeasible then the right question to ask is: “Oh yeah? Can I see the infeasibility certificate?”