

On Solving the Linear Programming Problem Approximately

NIMROD MEGIDDO

ABSTRACT. This paper studies the complexity of some approximate solutions of linear programming problems with real coefficients.

1. Introduction

The general linear programming problem is to maximize a linear function over a set defined by linear inequalities and equations. There are many equivalent ways to represent instances of the linear programming problem. For example, consider the symmetric form

$$\begin{array}{ll} \text{Maximize} & \mathbf{c}^T \mathbf{x} \\ \text{subject to} & \mathbf{A}\mathbf{x} \leq \mathbf{b} \\ & \mathbf{x} \geq \mathbf{0}. \end{array} \quad (\text{Sym}(\mathbf{A}, \mathbf{b}, \mathbf{c}))$$

The dual is then

$$\begin{array}{ll} \text{Minimize} & \mathbf{b}^T \mathbf{y} \\ \text{subject to} & \mathbf{A}^T \mathbf{y} \geq \mathbf{c} \\ & \mathbf{y} \geq \mathbf{0}. \end{array}$$

Intuitively, two representations are equivalent if there is an easy way to transform solutions of one to solutions of the other and vice versa. We first mention some of the well-known equivalences. First, any set of linear inequalities and linear equations can be reduced to a set of linear equations with nonnegativity constraints or to a set of inequality and nonnegativity constraints. Also, any linear programming problem can be reduced to a linear programming problem with a nonempty set of solutions by using artificial variables. Moreover, any linear programming problem can be reduced to a problem of finding a solution to a system of linear inequalities (by combining the constraints of the primal and dual and adding the inequality $\mathbf{c}^T \mathbf{x} \geq \mathbf{b}^T \mathbf{y}$)

1980 *Mathematics Subject Classification* (1985 Revision). Primary 90C05.

© 1990 American Mathematical Society
0271-4132/90 \$1.00 + \$.25 per page

or else concluding that the system has no solution. By the duality theorem, if we put the problem in the combined primal-dual form, every problem can be reduced to a problem that is either infeasible or feasible and bounded (but not unbounded). It is interesting to note that over any ordered field, every linear programming problem can be reduced to one that is both feasible and bounded. This is done as follows. Suppose the problem is given in the symmetric form $\text{Sym}(A, b, c)$. Consider the following:

$$\begin{aligned} \text{Minimize} \quad & c^T x - b^T y + t \\ \text{subject to} \quad & Ax - te \leq b \\ & A^T y + te \geq c \\ & c^T x - b^T y \geq 0 \\ & x, y, t \geq 0, \end{aligned}$$

where e denotes a vector of 1's. It is easy to verify that the optimal value of the latter is zero if and only if the former has an optimal solution. In this case the latter provides optimal solutions for the former and its dual.

The equivalences mentioned above are valid as long as exact computation is feasible. In practice one usually works with finite precision and hence obtains results that are only "approximately true." However, the meaning of the last sentence depends on the particular representation of the practical problem. Indeed, a good approximate solution for one representation of the problem may transform into a very bad approximate solution for another "equivalent" representation of the same problem.

When two people talk about approximate solutions, they often think of different notions of approximation. It is quite likely though that they refer to one of the following:

- (i) A feasible point (i.e., one that satisfies all of the constraints in the exact sense) and is close in a certain metric to an optimal point.
- (ii) A feasible solution whose objective function value is close to the optimal value.
- (iii) A point, not necessarily feasible, close to an optimal solution.
- (iv) A point that approximately satisfies every constraint, and whose objective function value is close to the optimal value.
- (v) A point close to the feasible domain, whose objective function value is close to the optimal value (called the "weak optimization problem" in [7]).
- (vi) A basis where the simplex algorithm (using exact arithmetic) terminates, but the numerical values of variables are only approximate.
- (vii) A basis where the simplex algorithm terminates due to a prescribed tolerance.

The choice of the right definition depends very much on the practical situation. In fact, practical considerations dictate which constraints must be satisfied and which may be approximately satisfied. In other words, the tolerance may be different for different constraints.

It is not known whether the $m \times n$ linear programming problem with real data can be solved in a polynomial number of arithmetic operations and comparisons in terms of m and n . (We will refer to this notion of complexity as strongly polynomial time, even though the usual definition of this concept also requires polynomial time in the usual sense.) Thus, another natural question is whether for any $\varepsilon > 0$, any of the ε -approximation problems can be solved in a polynomial number of operations in terms of m , n , and $-\log \varepsilon$. To deal with this question, we first have to define what we mean by an " ε -approximate" solution. In particular, such a definition should not make the second question trivially equivalent to the first one. Consider, for example, the concept suggested in (i) above. Thus, for the problem Maximize $c^T x$ subject to $Ax \leq b$ (assuming its maximum V^* exists) an ε -approximate solution is a point x such that $Ax \leq b$ and $c^T x \geq V^* - \varepsilon$. This definition is not satisfactory since it is not clear whether an ε -approximation algorithm is required to decide the existence of an x such that $Ax \leq b$, and the boundedness of the function $c^T x$ on the feasible domain. If indeed it is required to decide these questions, then in the worst-case sense this approximation problem is trivially equivalent to the exact problem.

The consequences of the ellipsoid algorithm with respect to approximation problems on convex sets are studied in [7]. It is not clear whether these results can be applied to achieve the type of results we seek here. The reason is that over the real numbers it seems difficult to obtain estimates of the radii of a circumscribing sphere and an inscribed sphere. The main complexity result on convex minimization in [7] (Theorem 2.2.15) assumes that the convex set is given with estimates of such radii. Our main interest here is the question of what is a reasonable sense of approximation when the algorithm fails to classify the instance correctly as feasible, unbounded, etc.

In Section 2 we give some preliminaries and discuss the difficulties involved in classifying the problem. In Section 3 we discuss approximate solutions based on satisfying a termination criterion within some tolerance. In Section 4 we discuss a notion of approximation that is based on solving a perturbed instance exactly. Section 5 gives an analysis of complexity for various notions of complexity.

2. Preliminaries

We pointed out in the introduction that the practical situation usually dictates the right notion of approximate solution. For a theoretical discussion it is often convenient to consider the problem in the symmetric form $\text{Sym}(A, b, c)$. (See Section 1.) Traditionally, an exact algorithm for this problem (for example, the simplex method) is supposed to provide the user with information as follows. It has to classify the problem into one of the following three categories:

- (i) Infeasible. (The domain X defined by $Ax \leq b$ and $x \geq 0$ is empty.)

- (ii) Feasible and bounded. (There exists a maximizer of $\mathbf{c}^T \mathbf{x}$ over X .)
- (iii) Feasible and unbounded. (The function $\mathbf{c}^T \mathbf{x}$ is unbounded over X .)

In case (ii) the algorithm has to provide an optimal solution. The algorithm may also be required in case (iii) to provide a ray contained in X along which $\mathbf{c}^T \mathbf{x}$ tends to infinity. A nice property of the simplex method is that it also solves the dual problem

$$\begin{aligned} &\text{Minimize} && \mathbf{b}^T \mathbf{y} \\ &\text{subject to} && \mathbf{A}^T \mathbf{y} \geq \mathbf{c} \\ &&& \mathbf{y} \geq \mathbf{0}. \end{aligned}$$

Thus, besides providing such a ray in case (iii), the algorithm also provides in case (ii) an optimal solution to the dual, and in case (i) a “certificate” in the form of a ray of a related problem.

In fact, the (exact) simplex method always computes a *basis* that provides the required information. Specifically, it provides a representation of the problem (by a suitable linear transformation of the space) from which the classification and the numerical values of both the primal and the dual variables are transparent. Thus, the simplex method classifies problems into one of *four* categories (even though the commonly used variants do not distinguish between IF and II):

- (i) FF: primal feasible, dual feasible.
- (ii) FI: primal feasible, dual infeasible (that is, unbounded primal).
- (iii) IF: primal infeasible, dual feasible (here the dual is unbounded).
- (iv) II: primal infeasible, dual infeasible.

It is quite common to include this classification in the requirements from an exact algorithm for the general linear programming problem. We refer to it later as the classification problem of linear programs.

An approximation algorithm should be expected sometimes to fail in classifying the input into the categories FF, FI, IF, and II. Interestingly, the existence of a strongly polynomial algorithm for the classification problem implies the existence of one for the problem itself. (See page 445 in [1].)

So far we have discussed the subject of approximation under the assumption that the result should be “close” to the true one. However, a different approach can sometimes be useful. We may allow the algorithm to be totally wrong in a small number of cases. This approach is approximate when the output space of the algorithm is discrete and has no natural metric associated with it. For example, consider the following trivial problem: Given two numbers α , β , recognize whether $\alpha > \beta$ or $\alpha \leq \beta$. Suppose the comparison of α to β can be performed with arbitrary finite precision. Thus, for any given $\varepsilon > 0$, we can recognize either that

$$\alpha \leq \beta + \varepsilon$$

or that

$$\alpha \geq \beta - \varepsilon.$$

The algorithm reports $\alpha \leq \beta$ in the first case and $\alpha \geq \beta$ in the second one. Thus, the algorithm gives the correct answer if

$$|\beta - \alpha| > \varepsilon$$

but may fail otherwise. The grey area is the set of pairs (α, β) such that $|\alpha - \beta| \leq \varepsilon$. Of course, the smaller ε the smaller the grey area. Thus, the measure of the grey area reflects the quality of the approximation.

The problem of the preceding paragraph can be cast as a linear programming problem

$$\begin{array}{ll} \text{Maximize} & x_1 \\ \text{subject to} & \alpha x_1 \leq x_2 \leq \beta x_1. \end{array}$$

Here the point $(0, 0)$ is feasible for any α and β . The problem is unbounded if and only if $\alpha \leq \beta$. This suggests that the grey area approach would be suitable for the classification problem of linear programs.

A general linear programming problem in standard form

$$\begin{array}{ll} \text{Maximize} & \mathbf{c}^T \mathbf{x} \\ \text{subject to} & \mathbf{A}\mathbf{x} = \mathbf{b} \\ & \mathbf{x} \geq \mathbf{0} \end{array}$$

(SF(\mathbf{A} , \mathbf{b} , \mathbf{c}))

is determined by $\mathbf{A} \in R^{m \times n}$, $\mathbf{b} \in R^m$, and $\mathbf{c} \in R^n$. There is a one-to-one correspondence between problems of order $m \times n$ and points of $\Omega = R^{mn+m+n}$. The classification corresponds to a partition of Ω into four sets: FF, FI, IF, and II, as discussed above. For example, IF is the set of triples $(\mathbf{A}, \mathbf{b}, \mathbf{c})$ that determine infeasible primal problems whose dual problems are feasible (and hence unbounded).

Let Ω' denote the union of the boundaries of these four sets. Obviously, an approximation algorithm (for the classification problem) may fail if the input $(\mathbf{A}, \mathbf{b}, \mathbf{c})$ is close to Ω' . For example, if an instance is close to the common boundary of FF and FI, but far from the boundaries of IF and II, then an approximation algorithm is expected to recognize that the problem is feasible, but is expected to fail in deciding whether it is bounded. Interestingly, there are more "pathological" cases. In fact, the intersection of all four boundaries, which we denote by Ω_4 , is not empty. Thus, given an instance close to the intersection of the four boundaries, an approximation algorithm may not be able to recognize anything in terms of the above classification. This observation is obvious in view of the invariance of the classification under multiplication of columns and rows by positive scalars. Thus, the neighborhood of the origin is obviously pathological in this sense. The difficulties with the origin can be avoided by scaling rows and columns. However, it is easy to construct other examples with similar characteristics.

PROPOSITION 2.1. *The instance*

$$\begin{array}{ll} (*) & \begin{array}{ll} \text{Maximize} & x_1 - x_2 \\ \text{subject to} & x_1 - x_2 \leq -1 \\ & -x_1 + x_2 \leq -1 \\ & x_1, x_2 \geq 0. \end{array} \end{array}$$

belongs to Ω_4 .

PROOF. It is easy to see that $(*)$ itself is in IF. Using small perturbations, one can move from $(*)$ to instances in any of the other three classes. More precisely, if only c_1 is slightly increased then we can get instances in II. If only A_{21} is slightly decreased then we get instances in FF. Finally, if only A_{11} is slightly increased then an instance in FI is obtained. \square

We note that most of the numerical difficulties in solving linear programming problems are due to the fact that many such problems are ill posed. It is well known in numerical analysis (see, e.g., [3]) that near-singularities in the matrix A can cause problems. However, in this paper we also discuss intrinsic aspects of approximate solutions that arise even when exact arithmetic is used. For example, if one is interested only in an approximate solution, what should be a good termination criterion? Because of such questions we have to deal with perturbations of the vectors b and c and not only of the matrix A .

Due to the classification aspect of the problem, we clearly cannot always measure the quality of the approximation by the distance between the exact solution and the approximate one (neither in terms of the solution vector nor in terms of the objective function value). Thus a different approach to approximation may be proposed for a general situation, where there is some natural metric on the input space, but there does not seem to exist one for the output space. The following definition is similar to backward analysis of errors in numerical analysis [9].

DEFINITION 2.2. Let $M = (S, d)$ be a metric space and let f be a mapping from S into some set T that does not necessarily have any metric associated with it. A mapping $g: S \rightarrow T$ is called an ε -approximation to f if for every $x \in S$, there exists an $x' \in S$ such that $d(x, x') < \varepsilon$ and $g(x) = f(x')$.

In linear programming the output space (including the classification) does have a metric structure. Besides the classification information, there are also numerical values associated with the variables. One might propose for the linear programming problem the following approach to approximation by posing the following problem:

PROBLEM 2.3. Given the problem $SF(A, b, c)$ and $\varepsilon > 0$, name a class $S \in \{FF, FI, IF, II\}$ and assign numerical values to the variables so that the following condition is satisfied: There exists an instance (A', b', c') in S for which the numerical values are correct within an error of ε , such that $\|(A, b, c) - (A', b', c')\|_\infty < \varepsilon$.

The approach represented by Problem 2.3 takes care of pathological cases where some other approaches fail. Consider, for comparison, a different notion of approximate solution of systems of inequalities reflected in the following problem:

PROBLEM 2.4. Given A, b and $\varepsilon > 0$, either give an x such that $Ax \leq b + \varepsilon e$ or conclude that there is no x such that $Ax \leq b - \varepsilon e$.

The approach represented by Problem 2.4 seems to be a natural generalization of the obvious approximate comparison of two real numbers. Its weakness is apparent in the following example. Consider the problem

$$\begin{aligned}\alpha x_1 - x_2 &\leq -1, \\ -x_1 + x_2 &\leq -1.\end{aligned}$$

If $\alpha = 1$ then, obviously, for every $\varepsilon < 1$ the system

$$\begin{aligned}x_1 - x_2 &\leq -1 + \varepsilon, \\ -x_1 + x_2 &\leq -1 + \varepsilon\end{aligned}$$

is infeasible. However, for any $\alpha \neq 1$ the system is feasible for every ε . Thus, in order to solve Problem 2.4 we have to know whether $\alpha = 1$. This example can easily be generalized so that in order to solve Problem 2.4 one has to know whether a certain matrix is singular. The latter involves some numerical difficulties in practice.

The weakness of the concept of Problem 2.4 is that it considers perturbations of the given problem only in a limited and quite arbitrary way. In Problem 2.3 we allow perturbations in all directions. We note that for certain classes of linear programming problems (e.g., the min-cost flow problem) certain coefficients have the values 1 or 0 throughout the class. In such cases we would allow only perturbations within the subject class. Thus, if the concept represented by Problem 2.3 were to apply to a min-cost flow problem $SF(A, b, c)$ then we would require that $A' = A$.

3. Tolerance-based approximation

In this section we discuss the issue of approximation as it arises in the context of the simplex method. There are two ways to look at the question. First, imagine we run the simplex algorithm using exact arithmetic but have to compute only “ ε -approximate” solutions. Thus, rather than running the algorithm to the end, we seek to apply some stopping rule that guarantees our output to be ε -approximate. The interesting problem is of course to devise such stopping rules for various concepts of approximation. Another way to look at the question is to realize that on a machine we usually have numerical errors, and thus we almost always have to specify some “tolerance” within which we accept our results. It is important to know the implications of using a certain tolerance with regard to the results.

Consider the problem $SF(A, b, c)$. In the Appendix we review some properties of basic solutions and how the simplex method uses them. (See the Appendix for the notation.)

In practice, one usually works with some “tolerance” $\delta > 0$, so that any number $a \leq \delta$ is accepted as nonpositive and any $a \geq -\delta$ is accepted as nonnegative. This suggests another approach to approximation, which may

be called the “tolerance” approach:

DEFINITION 3.1. A basic solution $\mathbf{x} = \mathbf{x}(\mathbf{B}) = (\mathbf{x}_B, \mathbf{x}_N)$ (where $\mathbf{x}_B = \mathbf{B}^{-1}\mathbf{b}$ and $\mathbf{x}_N = \mathbf{0}$) is *optimal with tolerance δ* for $\text{SF}(\mathbf{A}, \mathbf{b}, \mathbf{c})$ if for every j ,

$$x_j \geq -\delta \quad \text{and} \quad c_j \leq \mathbf{c}_B^T \mathbf{B}^{-1} \mathbf{A}_j + \delta.$$

Obviously, such a solution is not necessarily feasible. It is only “approximately feasible” in the sense that the equality constraints are satisfied while the nonnegativity constraints are approximately satisfied. Analogously, the dual vector $\mathbf{y}(\mathbf{B})$ is approximately feasible in the dual problem. Moreover, the vectors $\mathbf{x}(\mathbf{B})$ and $\mathbf{y}(\mathbf{B})$ satisfy the complementary slackness condition

$$x_j(\mathbf{y}^T \mathbf{A}_j - c_j) = 0,$$

which is necessary for optimality. This implies that both vectors yield the same objective function value in their respective problem

$$\mathbf{c}^T \mathbf{x} = \mathbf{c}^T \mathbf{B}^{-1} \mathbf{b} = \mathbf{y}^T \mathbf{b}.$$

Note that the problem may also be feasible and unbounded but that an optimal solution with tolerance δ may still exist. One can also define a notion of an unbounded ray with tolerance δ .

4. Perturbation-based approximation

It is interesting to observe that from a solution that is optimal with small tolerance we can easily obtain an exact solution to an instance that is close to the given one. More precisely, we have the following proposition:

PROPOSITION 4.1. Suppose $\mathbf{x} = \mathbf{x}(\mathbf{B})$ is a basic solution that is optimal with tolerance δ . Let \mathbf{x}' be defined by $x'_j = x_j + \delta$ for j associated with \mathbf{B} and $x'_j = 0$ otherwise. Also, let \mathbf{c}' be defined by $c'_j = c_j$ for j associated with \mathbf{B} and $c'_j = c_j - \delta$ otherwise, and let $\mathbf{b}' = \mathbf{b} + \delta \mathbf{B}\mathbf{e}$. Under these conditions, \mathbf{x}' is an optimal solution for the problem $\text{SF}(\mathbf{A}, \mathbf{b}', \mathbf{c}')$, and $\mathbf{y} = \mathbf{B}^{-T} \mathbf{c}_B$ is optimal for its dual problem.

PROOF. We have

$$\begin{aligned} \mathbf{x}'_B &= \mathbf{B}^{-1} \mathbf{b}' = \mathbf{x}_B + \delta \mathbf{e} \geq \mathbf{0}, \\ \mathbf{A}^T \mathbf{y} &\geq \mathbf{c}' \end{aligned}$$

and \mathbf{x}' and \mathbf{y} satisfy the complementary slackness conditions required in $\text{SF}(\mathbf{A}, \mathbf{b}', \mathbf{c}')$. \square

In simpler words, we have

COROLLARY 4.2. If \mathbf{x} is optimal with tolerance δ for $\text{SF}(\mathbf{A}, \mathbf{b}, \mathbf{c})$ then there exist \mathbf{b}' and \mathbf{c}' and an optimal solution \mathbf{x}' for $\text{SF}(\mathbf{A}, \mathbf{b}', \mathbf{c}')$ such that

$$\|\mathbf{x}' - \mathbf{x}\|_\infty, \|\mathbf{c}' - \mathbf{c}\|_\infty \leq \delta,$$

and

$$\|\mathbf{b}' - \mathbf{b}\|_\infty \leq \min\{\delta \|\mathbf{A}\|_\infty, mT\delta\}$$

where $\|A\|_\infty$ is the usual operator norm, corresponding to the vector supremum norm $\|\cdot\|$, T is the maximum absolute value of any entry in A , and m is the number of rows of A .

A analogous proposition can be proven with respect to the unbounded case:

PROPOSITION 4.3. Suppose $x = x(B)$ is a basic solution such that

$$x_j \geq -\delta$$

and for some k not associated with B ,

$$c_k > c_B^T B^{-1} A_k$$

and

$$B^{-1} A_k \geq -\delta e.$$

Let c' be the same as c except that $c'_k = c_k + \delta c_B^T e$, and let $b' = b + \delta B e$. Also, let A' be the same as A except that $A'_k = A_k + \delta B e$. Under these conditions, the problem $SF(A', b', c')$ is unbounded.

PROOF. In the new problem the basis B certifies unboundedness since

$$\begin{aligned} B^{-1} b' &\geq 0, \\ c'_k &> (c')^T B^{-1} A'_k, \end{aligned}$$

and

$$B^{-1} A'_k \geq 0. \quad \square$$

COROLLARY 4.4. If $SF(A, b, c)$ is concluded within tolerance δ to be unbounded, then there exist A' , b' , and c' such that $SF(A', b', c')$ is unbounded,

$$\|A' - A\|_1 \leq \delta \|A\|_\infty$$

(where $\|A' - A\|_1$ is the operator norm corresponding to the vector norm $\|\cdot\|_1$),

$$\|b' - b\|_\infty \leq \delta \|A\|_\infty,$$

and

$$\|c' - c\|_\infty \leq \delta \|c\|_1.$$

5. Complexity questions

We start this section with yet another variant of an approximation problem. Again, we consider approximation concepts that avoid the difficulties involved in the classification problem. Suppose the exact problem is given in the dual form

$$\begin{aligned} (\text{DF}(A, b, c)) \quad & \text{Minimize } c^T x \\ & \text{subject to } Ax \geq b \end{aligned}$$

where the output has to be one of the following:

- (i) a point x^* that minimizes $c^T x$ subject to $Ax \geq b$,

- (ii) a point \mathbf{x}^* and a scalar $t^* > 0$ that minimize the value of t subject to $\mathbf{A}\mathbf{x} + t\mathbf{e} \geq \mathbf{b}$ (in which case the problem is infeasible), or
- (iii) vectors \mathbf{x} and \mathbf{u} such that $\mathbf{A}\mathbf{x} \geq \mathbf{b}$, $\mathbf{c}^T \mathbf{u} < 0$, and $\mathbf{A}\mathbf{u} \geq \mathbf{0}$ (in which case the problem is unbounded).

The above motivates the definition of the following approximation problem:

PROBLEM 5.1. Denote the optimal value of a given problem $\text{DF}(\mathbf{A}, \mathbf{b}, \mathbf{c})$ by V^* (allowing $V^* = \pm\infty$) and let t^* denote the minimum of t subject to $\mathbf{A}\mathbf{x} + t\mathbf{e} \geq \mathbf{b}$. Given a number $\varepsilon > 0$, output one of the following:

- (i) a point \mathbf{x} such that $\mathbf{c}^T \mathbf{x} \leq V^* + \varepsilon$ and $\mathbf{A}\mathbf{x} \geq \mathbf{b} - \varepsilon\mathbf{e}$,
- (ii) a point \mathbf{x} and a scalar t , $-\varepsilon \leq t \leq t^* + \varepsilon$ such that $\mathbf{A}\mathbf{x} + t\mathbf{e} \geq \mathbf{b} - \varepsilon\mathbf{e}$,
or
- (iii) vectors \mathbf{x} and \mathbf{u} such that $\mathbf{A}\mathbf{x} \geq \mathbf{b} - \varepsilon\mathbf{e}$, $\mathbf{c}^T \mathbf{u} < \varepsilon$, and $\mathbf{A}\mathbf{u} \geq -\varepsilon\mathbf{e}$.

It is interesting to look at the question of the existence of an algorithm for the approximation problem, where the number of operations is expressed in terms of ε as well.

PROPOSITION 5.2. *Over any ordered field, if Problem 5.1 can be solved in $f(m, n, \varepsilon)$ field operations (including comparisons) then it can be solved in $g(m, n) = O(f(m, n, 1))$ operations.*

PROOF. Suppose \mathcal{A} is an algorithm for Problem 5.1 that runs in $f(m, n, \varepsilon)$ field operations. Given \mathbf{A} , \mathbf{b} , \mathbf{c} and $\varepsilon > 0$, let

$$\bar{\mathbf{A}} = \varepsilon^{-1} \mathbf{A}, \quad \bar{\mathbf{b}} = \varepsilon^{-1} \mathbf{b}, \quad \bar{\mathbf{c}} = \varepsilon^{-1} \mathbf{c}.$$

The instance $\text{DF}(\bar{\mathbf{A}}, \bar{\mathbf{b}}, \bar{\mathbf{c}})$ is equivalent to $\text{DF}(\mathbf{A}, \mathbf{b}, \mathbf{c})$. Moreover, a valid output for $\text{DF}(\bar{\mathbf{A}}, \bar{\mathbf{b}}, \bar{\mathbf{c}})$ with precision $\varepsilon = 1$ is also a valid output for $\text{DF}(\mathbf{A}, \mathbf{b}, \mathbf{c})$ with the prescribed precision ε . \square

COROLLARY 5.3. *If there exists a polynomial $f(m, n, \varepsilon)$ such that Problem 5.1 with rational data can be solved in $f(m, n, \varepsilon)$ arithmetic operations, then the exact problem with rational data can be solved in a polynomial number of operations.*

PROOF. For a problem with rational data it is easy to determine a value ε such that an exact solution can be computed from a solution of Problem 5.1 in a polynomial number of operations. Thus, the problem can be scaled so that $\varepsilon = 1$ suffices for determining an exact solution. \square

In view of Proposition 5.2 it is reasonable to ask whether Problem 5.1 can be solved in a polynomial number of operations in terms of m , n , and $\log R/\varepsilon$, where $R = R(\mathbf{A}, \mathbf{b}, \mathbf{c}) > 0$ is some quantity such that for any positive scalar λ

$$R(\lambda \mathbf{A}, \lambda \mathbf{b}, \lambda \mathbf{c}) = \lambda R(\mathbf{A}, \mathbf{b}, \mathbf{c})$$

(e.g., R equals the maximum absolute value of any input coefficients).

Consider first the feasibility problem

$$\begin{array}{ll} \text{Minimize} & t \\ \text{subject to} & \mathbf{Ax} + t\mathbf{e} \geq \mathbf{b} \\ & t \geq -1, \end{array}$$

and the associated approximation problem:

PROBLEM 5.4. Given $\text{FB}(\mathbf{A}, \mathbf{b})$ and $\varepsilon > 0$, find \mathbf{x} and t such that $\mathbf{Ax} + t\mathbf{e} \geq \mathbf{b}$ and $t < t^* + \varepsilon$ (where t^* is the minimum of $\text{FB}(\mathbf{A}, \mathbf{b})$).

Before stating the next proposition, recall that the set of solutions of the system $\mathbf{Ax} \geq \mathbf{b}$ is bounded for every \mathbf{b} if and only if the rows of \mathbf{A} span the space R^n in nonnegative linear combinations.

PROPOSITION 5.5. *If Problem 5.4 is given with real data such that the rows of \mathbf{A} span the space R^n in nonnegative linear combinations, then it can be solved in a polynomial number of arithmetic operations in terms of m , n , and*

$$\rho = \log \left(\frac{b_{\max} + 1}{\varepsilon} \right)$$

(assuming b_{\max} , the maximal b_i , is positive).

PROOF. The point $\mathbf{x}^0 = \mathbf{0}$, $t^0 = b_{\max} + 1$ is in the interior of the feasible domain of $\text{FB}(\mathbf{A}, \mathbf{b})$. By our assumption, the set of optimal solutions of $\text{FB}(\mathbf{A}, \mathbf{b})$ is bounded. Several interior point algorithms are now known (e.g., [2], [10]) that can start from any interior point and reduce the value of the objective function to a value not greater than $t^* + \varepsilon$ in a polynomial number of iterations in terms m , n , and ρ , where each iteration takes a polynomial number of operations in terms of m and n . \square

An obvious consequence of Proposition 5.5 is the following:

COROLLARY 5.6. *Suppose the rows of \mathbf{A} spanned the whole R^n in nonnegative linear combinations. It takes a polynomial number of operations in terms of m , n , and ρ to either compute a vector \mathbf{x} such that $\mathbf{Ax} \geq \mathbf{b}$ or conclude that there is no \mathbf{x} such that $\mathbf{Ax} \geq \mathbf{b} + \varepsilon\mathbf{e}$.*

PROOF. Run any of the polynomial interior point algorithms for a number of iterations that guarantees that $t < t^* + \varepsilon$. As soon as t becomes nonpositive, stop. (The current \mathbf{x} is feasible.) If at the end t is still positive then $t^* > -\varepsilon$, and hence there is no \mathbf{x} such that $\mathbf{Ax} - \varepsilon\mathbf{e} \geq \mathbf{b}$. \square

Another consequence with respect to optimal solutions with tolerance can be stated conveniently when the problem is in the symmetric form $\text{Sym}(\mathbf{A}, \mathbf{b}, \mathbf{c})$:

$$\begin{array}{ll} \text{Maximize} & \mathbf{c}^T \mathbf{x} \\ \text{subject to} & \mathbf{Ax} \leq \mathbf{b} \\ & \mathbf{x} \geq \mathbf{0}, \end{array}$$

whose dual is

$$\begin{array}{ll} \text{Minimize} & \mathbf{b}^T \mathbf{y} \\ \text{subject to} & \mathbf{A}^T \mathbf{y} \geq \mathbf{c} \\ & \mathbf{y} \geq \mathbf{0}. \end{array}$$

Recall that by the duality theorem, a feasible solution x is optimal if and only if there exists a dual feasible solution y such that $c^T x = b^T y$. On the other hand, $c^T x \leq b^T y$ for any pair of feasible x and y . This suggests the following approximation problem:

PROBLEM 5.7. Given A , b , c and $\varepsilon > 0$, either find a pair of vectors x , y such that

$$\begin{aligned} Ax &\leq b + \varepsilon e, \\ A^T y &\geq c - \varepsilon e, \\ b^T y - c^T x &\leq \varepsilon, \\ x &\geq -\varepsilon e, \quad y \geq -\varepsilon e \end{aligned}$$

or conclude that $\text{Sym}(A, b, c)$ does not have an optimal solution.

The notion of approximation presented in Problem 5.7 is very close to the one used in practice, where optimality criteria are applied without knowledge of proximity to the value of an optimal solution.

Note that Problem 5.7 is trivial if $c \leq 0 \leq b$. Thus, assume

$$\gamma = \max\{-b_{\min}, c_{\max}\} > 0$$

and denote

$$\rho^* = \log \left(\frac{\gamma + 1}{\varepsilon} \right).$$

PROPOSITION 5.8. Suppose the rows of the matrix

$$\begin{bmatrix} 0 & A^T \\ -A & 0 \\ c^T & -b^T \\ I & 0 \\ 0 & I \end{bmatrix}$$

span the space R^{m+n} in nonnegative linear combinations. Problem 5.7 can be solved in a polynomial number of operations in terms of m , n , and ρ^* .

PROOF. Consider the problem

$$\begin{aligned} &\text{Minimize } t \\ &\text{subject to } \begin{aligned} Ax - te &\leq b \\ A^T y + te &\geq c \\ b^T y - c^T x - t &\leq 0 \\ x + te, y + te &\geq 0. \end{aligned} \end{aligned}$$

Starting at $t = \gamma + 1$, $x = 0$, and $y = 0$, run a number of iterations that guarantees $t \leq t^* + \varepsilon$. This number is polynomial in m , n , and ρ^* , and ρ^* is trivial to compute. If $t \leq \varepsilon$ then the current x and y solve Problem 5.7. Otherwise, $t^* > 0$, and $\text{Sym}(A, b, c)$ does not have an optimal solution. \square

Note that a solution of Problem 5.7 does not guarantee that $c^T x$ is close to the optimal value V^* when the latter exists. It seems much more difficult to solve the approximation problem in the latter sense. This difficulty can be explained by considering the following practical question, which

arises when one applies the simplex method to the problem in standard form $SF(A, b, c)$. Suppose x and y satisfy

$$\begin{aligned} Ax &= b, \\ x &\geq -\delta e, \\ A^T y &\geq c - \delta e, \\ c^T x &= b^T y. \end{aligned}$$

Assuming the problem has an optimum whose value is V^* , we are interested in finding what should be the value of δ in order to ensure that

$$c^T x \geq V^* - \varepsilon.$$

To answer the question, suppose x^* is an optimal solution and we get

$$c^T x = b^T y = y^T Ax^* \geq c^T x^* - \delta e^T x^*.$$

Thus, we have to estimate the quantity $e^T x^*$.

As is known from the analysis of the ellipsoid algorithm, the coordinates of a basic solution x^* can be bounded as follows: First, x^* satisfies an equation $Bx^* = b$, where B is a nonsingular square submatrix of A . Thus,

$$x_i^* = \frac{\det(B^i)}{\det(B)}$$

where B^i is the matrix obtained from B by substituting b for the i th column. The known analyses do not use this relationship between B^i and B . In the case of integer coefficients an obvious lower bound on the absolute value of a nonzero determinant is 1. Thus, a lower bound can also be obtained in the case of rational coefficients. It is not known whether over the reals the following problem can be solved in a polynomial number of arithmetic operations and comparisons:

PROBLEM 5.9. Given a real matrix $A \in R^{m \times n}$, compute a positive α such that for every nonsingular submatrix $B \in R^{m \times m}$ of A , $|\det(B)| \geq \alpha$.

A tight upper bound on the absolute value of a determinant in terms of the maximum value M of any entry can be obtained over the reals as follows: Suppose the columns of B are v^1, \dots, v^m . Obviously,

$$|\det(B)| \leq \|v^1\| \cdots \|v^m\| \leq (\sqrt{m}M)^m.$$

At least for values of m for which there exist Hadamard matrices (i.e., matrices of orthogonal columns consisting of ± 1 's), there exist matrices with the determinant $(\sqrt{m}M)^m$.

Without exploiting the relationship between B^i and B the only claim we can prove is that

$$x_i^* \leq \frac{(\sqrt{m}M)^m}{\eta(A)}$$

where $\eta(A)$ is any positive lower bound on the absolute value of the determinant of any nonsingular $m \times m$ submatrix of A . We note in passing that

in the case of integral coefficients the bounds on x_i^* cannot be improved dramatically. For example, if

$$B = \begin{bmatrix} 1 & & & & \\ -M & 1 & & & \\ -M & -M & 1 & & \\ \vdots & \vdots & \ddots & \ddots & \\ -M & -M & \cdots & -M & 1 \end{bmatrix}$$

and $\mathbf{b} = (M, \dots, M)^T$ then $x_m^* = \sum_{i=1}^m M^i$. Also, if

$$B = \begin{bmatrix} M & M & M & \cdots & M \\ -1 & M & M & \cdots & M \\ & -1 & M & \cdots & M \\ & & \ddots & \ddots & \vdots \\ & & & -1 & M \end{bmatrix}$$

and $\mathbf{b} = (1, 0, \dots, 0)^T$ then $x_m^* = 1/(M(M+1)^{m-1})$.

Recall that to guarantee an ε -approximation in terms of the function value, we have to choose δ such that

$$\delta \leq \frac{\varepsilon}{\mathbf{e}^T \mathbf{x}^*}.$$

Hence we get the estimate

$$\delta \leq \frac{\varepsilon \eta(\mathcal{A})}{m(\sqrt{m}T)^m}.$$

Another estimate can be derived by

$$\mathbf{e}^T \mathbf{x}^* \leq \sqrt{m} \|\mathbf{x}^*\| = \sqrt{m} \|\mathbf{B}^{-1} \mathbf{b}\| \leq \frac{\sqrt{m} \|\mathbf{b}\|}{\sqrt{\lambda_{\min}(\mathbf{B}\mathbf{B}^T)}}$$

where λ_{\min} denotes the least eigenvalue. Thus, what we would need is a lower bound on the least eigenvalue of any matrix of the form $\mathbf{B}\mathbf{B}^T$ where \mathbf{B} is a basis. Note that our estimates depend on properties of an optimal basis rather than one that supports the approximate solution. So, the optimal basis may be ill conditioned, this fact being unknown to the user, and the terminal basis well conditioned and satisfying the optimality conditions within tolerance.

Obviously, if any $\eta(\mathcal{A})$ is known then we can solve the ε -approximation problem in a polynomial number of operations in terms of m , n , and $-\log \delta$. However, in general η is not known. It is interesting to note that for a practical solution of problems with thousands of variables, even with a sparse and well-structured problem with small integral coefficients, the required δ may be too small to be practical. It is not clear that an approximate solution based on tolerance has a value close to the optimal. In fact, the value may be far from the optimum even though the duality gap is very small, since the vectors \mathbf{x} and \mathbf{y} are only approximately feasible.

The following example illustrates the difficulties described above. Suppose we solve a problem in standard form, and the representation using the current basis is

$$\begin{aligned} &\text{Maximize} && \tilde{\mathbf{c}}_N^T \mathbf{x}_N \\ &\text{subject to} && \mathbf{x}_B + \tilde{\mathbf{N}} \mathbf{x}_N = \mathbf{e} \\ &&& \mathbf{x}_B, \mathbf{x}_N \geq \mathbf{0}, \end{aligned}$$

where $\tilde{\mathbf{N}} \in R^{m \times (n-m)}$ contains the following $m \times m$ submatrix:

$$\mathbf{B}' = \begin{bmatrix} 1 & & & & \\ -2 & 1 & & & \\ & -2 & 1 & & \\ & & \ddots & \ddots & \\ & & & -2 & 1 \end{bmatrix}.$$

Suppose further that the coordinates of $\tilde{\mathbf{c}}_N$ are negative except for those corresponding to the columns of \mathbf{B}' that are all equal to zero except for the last one, which equals 10^{-20} . Suppose $m = 1000$, which is quite common in practice. Assuming 10^{-50} is considered nonpositive within tolerance, the solution $\mathbf{x} = \mathbf{e}$ is accepted as optimal, with an objective function value of 0. However, the basis \mathbf{B}' determines a feasible solution \mathbf{x} where $x_{B'_i} = 2^i - 1$ and the objective function value is greater than 10^{250} . Note that the current basis \mathbf{B} is very well conditioned. Moreover, the underlying matrix is very sparse and well structured.

Appendix

We review some known characteristics of the simplex method for problems in standard form $\text{SF}(\mathbf{A}, \mathbf{b}, \mathbf{c})$. Suppose, for simplicity, we run only “phase II” of the algorithm; i.e., we start from a basic feasible solution and attempt to find an optimal one or conclude that the problem is unbounded. (It is well known that if a problem in standard form has a feasible solution then it has a basic feasible one; the “phase I” problem of finding a basic feasible solution, or concluding that none exists, can be formulated as a problem in standard form with a known basic feasible solution.) Assuming exact arithmetic, the algorithm then terminates with a feasible basis. The termination criterion is stated in terms of signs of certain entries of the “tableau.” As is common, let $\mathbf{B} \in R^{m \times m}$ denote the nonsingular submatrix of \mathbf{A} whose columns constitute the current basis, and let \mathbf{N} denote the matrix consisting of the other columns. Let \mathbf{x}_B and \mathbf{c}_B denote the restrictions of the vectors \mathbf{x} and \mathbf{c} , respectively, to the indices corresponding to the columns of \mathbf{B} . Let \mathbf{x}_N and \mathbf{c}_N denote the complementary restrictions of the vectors. The “tableau” is essentially a representation of the problem in an equivalent form:

$$\begin{aligned} &\text{Maximize} && (\mathbf{c}_N^T - \mathbf{c}_B^T \mathbf{B}^{-1} \mathbf{N}) \mathbf{x}_N \\ &\text{subject to} && \mathbf{x}_B + \mathbf{B}^{-1} \mathbf{N} \mathbf{x}_N = \mathbf{B}^{-1} \mathbf{b} \\ &&& \mathbf{x}_B \geq \mathbf{0}, \mathbf{x}_N \geq \mathbf{0}. \end{aligned}$$

Given a basis B , the corresponding basic primal solution $x = x(B)$ is given by

$$x_B = B^{-1}b, \quad x_N = 0.$$

The basic dual vector $y = y(B)$ associated with B is given by the equation

$$y^T B = c_B^T.$$

The dual problem is

$$\begin{array}{ll} \text{Minimize} & y^T b \\ \text{subject to} & y^T A \geq c^T. \end{array}$$

Thus, $y(B)$ is feasible in the dual problem if and only if

$$c_N^T \leq c_B^T B^{-1}N,$$

in which case y is optimal for the dual.

Assuming that the algorithm (namely, the primal simplex method) works with exact arithmetic, for every B occurring in the process,

$$B^{-1}b \geq 0.$$

The algorithm terminates in one of the following cases:

- (i) $c_N^T \leq c_B^T B^{-1}N$, in which case B is optimal, or
- (ii) there exists a column N_j such that $c_j > c_B^T B^{-1}N_j$ and $B^{-1}N_j \leq 0$, in which case the problem is unbounded.

In either case the basis B is said to be *terminal*. The termination criterion applies to signs of certain entries in the tableau.

REFERENCES

1. V. Chvátal, *Linear programming*, W. H. Freeman, New York, 1983.
2. D. M. Gay, *A variant of Karmarkar's linear programming algorithm for problems in standard form*, Math. Programming **37** (1987), 81–90.
3. G. H. Golub and C. Van Loan, *Matrix computations*, Johns Hopkins University Press, 1983.
4. C. C. Gonzaga, *Polynomial affine algorithms for linear programming*, Report ES-139/88, Dept. of Systems Engineering and Computer Sciences, COPPE, Rio de Janeiro, 1988.
5. M. Grötschel, L. Lovász, and A. Schrijver, *The ellipsoid method and its consequences in combinatorial optimization*, Combinatorica **1** (1981), 169–197; Corrigendum: Combinatorica **4** (1984), 291–295.
6. L. G. Khachiyan, *A polynomial algorithm in linear programming*, Soviet. Math. Dokl. **20** (1979), 191–194.
7. L. Lovász, *An algorithmic theory of numbers, graphs and convexity*, SIAM, Philadelphia, 1986.
8. N. Megiddo, *Combinatorial optimization with rational objective functions*, Math. Oper. Res. **4** (1979), 414–424.
9. J. H. Wilkinson, *Error analysis of floating-point computation*, Numer. Math. **2** (1960), 219–340.
10. Y. Ye and M. Kojima, *Recovering optimal dual solutions in Karmarkar's polynomial algorithm for linear programming*, Math. Programming **39** (1987), 305–317.

IBM RESEARCH DIVISION, ALMADEN RESEARCH CENTER SAN JOSE, CALIFORNIA 95120-6099,
 AND SCHOOL OF MATHEMATICAL SCIENCES, TEL AVIV UNIVERSITY, TEL AVIV, ISRAEL
 E-mail address: MEGIDDO@IBM.COM