

Teaching Duality in Linear Programming - the Multiplier Approach

Jens Clausen
DIKU, Department of Computer Science
University of Copenhagen
Universitetsparken 1
DK 2100 Copenhagen Ø

March 31, 2003

Abstract

Duality in LP is often introduced through the relation between LP problems modeling different aspects of a planning problem. Though providing a good motivation for the study of duality this approach does not support the general understanding of the interplay between the primal and the dual problem with respect to the variables and constraints.

This paper describes the multiplier approach to teaching duality: Replace the primal LP-problem P with a relaxed problem by including in the objective function the violation of each primal constraint multiplied by an associated multiplier. The relaxed problem is trivial to solve, but the solution provides only a bound for the solution of the primal problem. The new problem is hence to choose the multipliers so that this bound is optimized. This is the dual problem of P .

LP duality is described similarly in the work by A. M. Geoffrion on Lagrangean Relaxation for Integer Programming. However, we suggest here that the approach is used not only in the technical parts of a method for integer programming, but as a general tool in teaching LP.

Key words: Linear programming, Duality.

1 Introduction

Duality is one of the most fundamental concepts in connection with linear programming and provides the basis for better understanding of LP models and their results and for algorithm construction in linear and in integer programming. Duality in LP is in most

textbooks as e.g. [2, 7, 9] introduced using examples building upon the relationship between the primal and the dual problem seen from an economical perspective. The primal problem may e.g. be a blending problem:

Determine the contents of each of a set of available ingredients (e.g. fruit or grain) in a final blend. Each ingredient contains varying amounts of vitamins and minerals, and the final blend has to satisfy certain requirements regarding the total contents of minerals and vitamins. The costs of units of the ingredients are given, and the goal is to minimize the unit cost of the blend.

The dual problem here turns out to be a problem of prices: What is the maximum price one is willing to pay for “artificial” vitamins and minerals to substitute the natural ones in the blend?

After such an introduction the general formulations and results are presented including the statement and proof of the duality theorem:

Theorem 1 (Duality Theorem) *If feasible solutions to both the primal and the dual problem in a pair of dual LP problems exist, then there is an optimum solution to both systems and the optimal values are equal.*

Also accompanying theorems on unboundedness and infeasibility, and the Complementary Slackness Theorem are presented in most textbooks.

While giving a good motivation for studying dual problems this approach has an obvious shortcoming when it comes to explaining duality in general, i.e. in situations, where no natural interpretation of the dual problem in terms of primal parameters exists.

General descriptions of duality are often handled by means of symmetrical dual forms as introduced by von Neumann. Duality is introduced by stating that two LP-problems are dual problems *by definition*. The classical duality theorems are then introduced and proved. The dual of a given LP-problem can then be found by transforming this to a problem of one of the two symmetrical types and deriving its dual through the definition. Though perfectly clear from a formal point of view this approach does not provide any understanding the interplay between signs of variables in one problem and type of constraints in the dual problem. In [1], Appendix II, a presentation of general duality trying to provide this understanding is given, but the presentation is rather complicated.

In the following another approach used by the author when reviewing duality in courses on combinatorial optimization is suggested. The motivating idea is that of problem relaxation: If a problem is difficult to solve, then find a family of easy problems each resembling the original one in the sense that the solution provides information in terms of bounds on the solution of our original problem. Now find that problem among the easy ones, which provides the strongest bounds.

In the LP case we relax the primal problem into one with only non-negativity constraints by including in the objective function the violation of each primal constraint multiplied by

Fruit type	F1	F2	F3	F4
Preservative R	2	3	0	5
Preservative Q	3	0	2	4
Cost pr. ton	13	6	4	12

Table 1. Contents of R and Q and price for different types of fruit.

an associated multiplier. For each choice of multipliers respecting sign conditions derived from the primal constraints, the optimal value of the relaxed problem is a lower bound (in case of primal minimization) on the optimal primal value. The dual problem now turns out to be the problem of maximizing this lower bound.

The advantages of this approach are 1) that the dual problem of any given LP problem can be derived in a natural way without problem transformations and definitions 2) that the primal/dual relationship between variables and constraints and the signs/types of these becomes very clear for all pairs of primal/dual problems and 3) that *Lagrangean Relaxation* in integer programming now becomes a natural extension as described in [3, 4]. A similar approach is sketched in [8, 5].

The reader is assumed to have a good knowledge of basic linear programming. Hence, concepts as Simplex tableau, basic variables, reduced costs etc. will be used without introduction. Also standard transformations between different forms of LP problems are assumed to be known.

The paper is organized as follows: Section 2 contains an example of the approach sketched, Section 3 presents the general formulation of duality through multipliers and the proof of the Duality Theorem, and Section 4 discusses the pros and cons of the approach presented. The main contribution of the paper is not theoretical but pedagogical: the derivation of the dual of a given problem can be presented without problem transformations and definitions, which are hard to motivate to people with no prior knowledge of duality.

2 A blending example

The following example is adapted from [6].

The Pasta Basta company wants to evaluate an ecological production versus a traditional one. One of their products, the Pasta Basta lasagne, has to contain certain preservatives, R and Q, in order to ensure durability. Artificially produced counterparts R' and Q' are usually used in the production - these are bought from a chemical company PresChem, but are undesirable from the ecological point of view. R and Q can alternatively be extracted from fresh fruit, and there are four types of fruit each with their particular content (number of units) of R and Q in one ton of the fruit. These contents and the cost of buying the fruit are specified in Table 1.

Pasta Basta has a market corresponding to daily needs of 7 units of R and 2 units of Q. If the complete production is based on ecologically produced preservatives, which types of fruit and which amounts should be bought in order to supply the necessary preservatives in the cheapest way ?

The problem is obviously an LP-problem:

$$\begin{aligned}
 (P) \quad & \min \quad 13x_1 + 6x_2 + 4x_3 + 12x_4 \\
 & \text{s.t.} \quad 2x_1 + 3x_2 + 5x_4 = 7 \\
 & \quad \quad 3x_1 + 2x_3 + 4x_4 = 2 \\
 & \quad \quad x_1, x_2, x_3, x_4 \geq 0
 \end{aligned}$$

With x_2 and x_3 as initial basic variables the Simplex method solves the problem with the tableau of Table 2 as result.

Natural questions when one knows one solution method for a given type of problem are: Is there an easier way to solve such problems ? Which LP problems are trivial to solve ? Regarding the latter, it is obvious that LP minimization problems with no constraints but non-negativity of x_1, \dots, x_n are trivial to solve:

$$\begin{aligned}
 \min \quad & c_1x_1 + \dots + c_nx_n \\
 \text{s.t.} \quad & x_1, \dots, x_n \geq 0
 \end{aligned}$$

If at least one cost coefficient is negative the value of the objective function is unbounded from below (in the following termed “equal to $-\infty$ ”) with the corresponding variable unbounded (termed “equal to ∞ ”) and all other variables equal to 0, otherwise it is 0 with all variables equal to 0.

The blending problem P of 2 is not of the form just described. However, such a problem may easily be constructed from P : Measure the violation of each of the original constraints by the difference between the right-hand and the left-hand side:

$$\begin{aligned}
 & 7 - (2x_1 + 3x_2 + 5x_4) \\
 & 2 - (3x_1 + 2x_3 + 4x_4)
 \end{aligned}$$

	x_1	x_2	x_3	x_4	
Red. Costs	15/2	0	3	0	-15
x_2	-7/12	1	-5/6	0	3/2
x_4	3/4	0	1/2	1	1/2

Table 2. Optimal Simplex tableau for problem LP

Multiply these by penalty factors y_1 and y_2 and add them to the objective function:

$$(PR(y_1, y_2)) \quad \min_{x_1, \dots, x_4 \geq 0} \left\{ \begin{array}{l} 13x_1 + 6x_2 + 4x_3 + 12x_4 \\ + y_1(7 - 2x_1 - 3x_2 - 5x_4) \\ + y_2(2 - 3x_1 - 2x_3 - 4x_4) \end{array} \right\}$$

We have now constructed a family of relaxed problems, one for each value of y_1, y_2 , which are easy to solve. None of these seem to be the problem we actually want to solve, but the solution of each $PR(y_1, y_2)$ gives some information regarding the solution of P . It is a lower bound for any y_1, y_2 , and the maximum of these lower bound turns out to be equal to the optimal value of P . The idea of replacing a difficult problem by an easier one, for which the optimal solution provides a lower bound for the optimal solution of the original problem, is also the key to understanding Branch-and-Bound methods in integer programming.

Let $\text{opt}(P)$ and $\text{opt}(PR(\cdot))$ denote the optimal values of P and $PR(\cdot)$ resp. Now observe the following points:

1. $\forall y_1, y_2 \in \mathcal{R}: \text{opt}(PR(y_1, y_2)) \leq \text{opt}(P)$
2. $\max_{y_1, y_2 \in \mathcal{R}} (\text{opt}(PR(y_1, y_2))) \leq \text{opt}(P)$

1) states that $\text{opt}(PR(y_1, y_2))$ is a lower bound for $\text{opt}(P)$ for any choice of y_1, y_2 and follows from the fact that for any set of values for x_1, \dots, x_4 satisfying the constraints of P , the values of P and PR are equal since the terms originating in the violation of the constraints vanish. Hence $\text{opt}(PR(y_1, y_2))$ is found by minimizing over a set containing all values of feasible solutions to P implying 1). Since 1) holds for all pairs y_1, y_2 it must also hold for the pair giving the maximum value of $\text{opt}(PR(\cdot))$, which is 2). In the next section we will prove that

$$\max_{y_1, y_2 \in \mathcal{R}} (\text{opt}(PR(y_1, y_2))) = \text{opt}(P)$$

The best bound for our LP problem P is thus obtained by finding optimal multipliers for the relaxed problem. We have here tacitly assumed that P has an optimal solution, i.e. it is neither infeasible nor unbounded - we return to that case in Section 3.

Turning back to the relaxed problem the claim was that it is easily solvable for any *given* y_1, y_2 . We just collect terms to find the coefficients of x_1, \dots, x_4 in $PR(y_1, y_2)$:

$$(PR(y_1, y_2)) \quad \min_{x_1, \dots, x_4 \geq 0} \left\{ \begin{array}{l} (13 - 2y_1 - 3y_2) x_1 \\ + (6 - 3y_1) x_2 \\ + (4 - 2y_2) x_3 \\ + (12 - 5y_1 - 4y_2) x_4 \\ + 7y_1 + 2y_2 \end{array} \right\}$$

Since y_1, y_2 are fixed the term $7y_1 + 2y_2$ in the objective function is a constant. If any coefficient of a variable is less than 0 the value of PR is $-\infty$. The lower bound for $\text{opt}(P)$ provided by such a pair of y -values is of no value. Hence, we concentrate on y -values for which this does not happen. These pairs are exactly those assuring that each coefficient for x_1, \dots, x_4 is non-negative:

$$\begin{aligned} (13 \quad -2y_1 \quad -3y_2) &\geq 0 \Leftrightarrow 2y_1 + 3y_2 \leq 13 \\ (6 \quad -3y_1 \quad \quad) &\geq 0 \Leftrightarrow 3y_1 \leq 6 \\ (4 \quad \quad \quad -2y_2) &\geq 0 \Leftrightarrow 2y_2 \leq 4 \\ (12 \quad -5y_1 \quad -4y_2) &\geq 0 \Leftrightarrow 5y_1 + 4y_2 \leq 12 \end{aligned}$$

If these constraints all hold the optimal solution to $PR(y_1, y_2)$ has x_1, \dots, x_4 all equal to 0 with a value of $7y_1 + 2y_2$ which, since y_1, y_2 are finite, is larger than $-\infty$. Since we want to maximize the lower bound $7y_1 + 2y_2$ on the objective function value of P , we have to solve the following problem to find the optimal multipliers:

$$\begin{aligned} \max \quad & 7y_1 + 2y_2 \\ \text{s.t.} \quad & 2y_1 + 3y_2 \leq 13 \\ & 3y_1 \leq 6 \\ & 2y_2 \leq 4 \\ & 5y_1 + 4y_2 \leq 12 \\ & y_1, y_2 \in \mathcal{R} \end{aligned} \tag{DP}$$

The problem DP resulting from our reformulation is exactly the dual problem of P . It is again a linear programming problem, so nothing is gained with respect to ease of solution - we have no reason to believe that DP is any easier to solve than P . However, the above example indicates that linear programming problems appear in pairs defined on the *same* data with one being a minimization and the other a maximization problem, with variables of one problem corresponding to constraints of the other, and with the type of constraints determining the signs of the corresponding dual variables. Using the multiplier approach we have derived the dual problem DP of our original problem P , and we have through 1) and 2) proved the so-called Weak Duality Theorem - that $\text{opt}(P)$ is greater than or equal to $\text{opt}(DP)$.

In the next section we will discuss the construction in general, the proof of the Duality Theorem as stated in the introduction, and the question of unboundedness/infeasibility of the primal problem. We end this section by deriving DP as frequently done in textbooks on linear programming.

The company PresChem selling the artificially produced counterparts R' and Q' to Pasta Basta at prices r and q is considering to increase these as much as possible well knowing that many consumers of Pasta Basta lasagne do not care about ecology but about prices.

These customers want as cheap a product as possible, and Pasta Basta must also produce a cheaper product to maintain its market share.

If the complete production of lasagne is based on P' and Q', the profit of PresChem is $7r + 2q$. Of course r and q cannot be so large that it is cheaper for Pasta Basta to extract the necessary amount of R and Q from fruit. For example, at the cost of 13, Pasta Basta can extract 2 units of R and 3 units of Q from one ton of F1. Hence

$$2r + 3q \leq 13$$

The other three types of fruit give rise to similar constraints. The prices r and q are normally regarded to be non-negative, but the very unlikely possibility exists that it may pay off to offer Pasta Basta money for each unit of one preservative used in the production provided that the price of the other is large enough. Therefore the prices are allowed to take also negative values. The optimization problem of PresChem is thus exactly *DP*.

3 General formulation of dual LP problems

3.1 Proof of the Duality Theorem

The typical formulation of an LP problem with n nonnegative variables and m equality constraints is

$$\begin{aligned} \min \quad & cx \\ & Ax = b \\ & x \geq 0 \end{aligned}$$

where c is an $1 \times n$ matrix, A is an $m \times n$ matrix and b is an $n \times 1$ matrix of reals. The process just described can be depicted as follows:

$$\begin{aligned}
& \min \left. \begin{array}{l} cx \\ Ax = b \\ x \geq 0 \end{array} \right\} \mapsto \\
& \max_{y \in \mathcal{R}^m} \{ \min_{x \in \mathcal{R}_+^n} \{ cx + y(b - Ax) \} \} \mapsto \\
& \max_{y \in \mathcal{R}^m} \{ \min_{x \in \mathcal{R}_+^n} \{ (c - yA)x + yb \} \} \mapsto \\
& \left\{ \begin{array}{l} \max \quad yb \\ \quad yA \leq c \\ \quad y \text{ free} \end{array} \right.
\end{aligned}$$

The proof of the Duality Theorem proceeds in the traditional way: We find a set of multipliers which satisfy the dual constraints and gives a dual objective function value equal to the optimal primal value.

Assuming that P and DP have feasible solutions implies that P can be neither infeasible nor unbounded. Hence an optimal basis \mathcal{B} and a corresponding optimal basic solution $x_{\mathcal{B}}$ for P exists. The vector $y_{\mathcal{B}} = c_{\mathcal{B}}\mathcal{B}^{-1}$ is called the *dual solution* or the set of *Simplex multipliers* corresponding to \mathcal{B} . The vector satisfies that if the reduced costs of the Simplex tableau is calculated using $y_{\mathcal{B}}$ as π in the general formula

$$\bar{c} = c - \pi A$$

then \bar{c}_i equals 0 for all basic variables and \bar{c}_j is non-negative for all non-basic variables. Hence,

$$y_{\mathcal{B}}A \leq c$$

holds showing that $y_{\mathcal{B}}$ is a feasible dual solution. The value of this solution is $c_{\mathcal{B}}\mathcal{B}^{-1}b$, which is exactly the same as the primal objective value obtained by assigning to the basic variables $x_{\mathcal{B}}$ the values defined by the updated right-hand side $\mathcal{B}^{-1}b$ multiplied by the vector of basic costs $c_{\mathcal{B}}$.

The case in which the problem P has no optimal solution is for all types of primal and dual problems dealt with as follows. Consider first the situation where the objective function is unbounded on the feasible region of the problem. Then any set of multipliers must give rise to a dual solution with value $-\infty$ (resp. $+\infty$ for a maximization problem) since this is the only “lower bound” (“upper bound”) allowing for an unbounded primal objective function. Hence, no set of multipliers satisfy the dual constraints, and the dual feasible set is empty. If maximizing (resp. minimizing) over an empty set returns the value $-\infty$ (resp. $+\infty$), the desired relation between the primal and dual problem with respect to objective function value holds - the optimum values are equal.

Finally, if no primal solution exist we minimize over an empty set - an operation returning the value $+\infty$. In this case the dual problem is either unbounded or infeasible.

3.2 Other types of dual problem pairs

Other possibilities of combinations of constraint types and variable signs of course exist. One frequently occurring type of LP problem is a maximization problem in non-negative variables with less than or equal constraints. The construction of the dual problem is outlined below:

$$\begin{aligned} & \left. \begin{array}{l} \max \quad cx \\ Ax \leq b \\ x \geq 0 \end{array} \right\} \mapsto \\ & \min_{y \in \mathcal{R}_+^m} \{ \max_{x \in \mathcal{R}_+^n} \{ cx + y(b - Ax) \} \} \mapsto \\ & \min_{y \in \mathcal{R}_+^m} \{ \max_{x \in \mathcal{R}_+^n} \{ (c - yA)x + yb \} \} \mapsto \\ & \left\{ \begin{array}{l} \min \quad yb \\ yA \geq c \\ y \geq 0 \end{array} \right. \end{aligned}$$

Note here that the multipliers are restricted to being non-negative, thereby ensuring that for any feasible solution, $\hat{x}_1, \dots, \hat{x}_n$, to the original problem, the relaxed objective function will have a value greater than or equal to that of the original objective function since $b - A\hat{x}$ and hence $y(b - A\hat{x})$ will be non-negative. Therefore the relaxed objective function will be pointwise larger than or equal to the original one on the feasible set of the primal problem, which ensures that an upper bound results for all choices of multipliers. The set of multipliers minimizing this bound must now be determined.

Showing that a set of multipliers exists such that the optimal value of the relaxed problem equals the optimal value of the original problem is slightly more complicated than in the previous case. The reason is that the value of the relaxed objective function no longer is equal to the value of the original one for each feasible point, it is larger than or equal to this.

A standard way is to formulate an LP problem P' equivalent to the given problem P by adding a *slack variable* to each of the inequalities thereby obtaining a problem with equality constraints:

$$\left. \begin{array}{l} \max \quad cx \\ Ax \leq b \\ x \geq 0 \end{array} \right\} = \left\{ \begin{array}{l} \max \quad cx + 0s \\ Ax + Is = b \\ x, s \geq 0 \end{array} \right.$$

Note now that if we derive the dual problem for P' using multipliers we end up with the dual problem of P : Due to the equality constraints, the multipliers are now allowed to take both positive and negative values. The constraints on the multipliers imposed by the identity matrix corresponding to the slack variables are, however,

$$yI \geq 0$$

i.e. exactly the non-negativity constraints imposed on the multipliers by the inequality constraints of P . The proof just given now applies for P' and DP' , and the non-negativity of the optimal multipliers y_B are ensured through the sign of the reduced costs in optimum since these now satisfy

$$\bar{c} = (c \ 0) - y_B(A \ I) \leq 0 \quad \Leftrightarrow \quad y_B A \geq c \quad \wedge \quad y_B \geq 0$$

Since P' and P are equivalent the theorem holds for P and DP as well.

The interplay between the types of primal constraints and the signs of the dual variables is one of the issues of duality, which often creates severe difficulties in the teaching situation. Using the common approach to teaching duality, often no explanation of the interplay is provided. We have previously illustrated this interplay in a number of situations. For the sake of completeness we now state all cases corresponding to a primal minimization problem - the case of primal maximization can be dealt with likewise.

First note that the relaxed primal problems provide *lower bounds*, which we want to *maximize*. Hence the relaxed objective function should be pointwise *less than or equal* to the original one on the feasible set, and the dual problem is a *maximization problem*. Regarding the signs of the dual variables we get the following for the three possible types of primal constraints (A_i denotes the i 'th row of the matrix \mathbf{A}):

$A_i x \leq b_i$ For a feasible x , $b_i - A_i x$ is larger than or equal to 0, and $y_i(b_i - A_i x)$ should be non-positive. Hence, y_i should be non-positive as well.

$A_i x \geq b_i$ For a feasible x , $b_i - A_i x$ is less than or equal to 0, and $y_i(b_i - A_i x)$ should be non-positive. Hence, y_i should be non-negative.

$A_i x = b_i$ For a feasible x , $b_i - A_i x$ is equal to 0, and $y_i(b_i - A_i x)$ should be non-positive. Hence, no sign constraints should be imposed on y_i .

Regarding the types of the dual constraints, which we previously have not explicitly discussed, these are determined by the sign of the coefficients to the variables in the relaxed primal problem in combination with the sign of the variables themselves. The coefficient of x_j is $(c - yA)_j$. Again we have three cases:

$x_j \geq 0$ To avoid unboundedness of the relaxed problem $(c - yA)_j$ must be greater than or equal to 0, i.e. the j 'th dual constraint will be $(yA)_j \leq c_j$.

$x_j \leq 0$ In order not to allow unboundedness of the relaxed problem $(c - yA)_j$ must be less than or equal to 0, i.e. the j 'th dual constraint will be $(yA)_j \geq c_j$.

x_j *free* In order not to allow unboundedness of the relaxed problem $(c - yA)_j$ must be equal to 0 since no sign constraints on x_j are present, i.e. the j 'th dual constraint will be $(yA)_j = c_j$.

3.3 The Dual Problem for Equivalent Primal Problems

In the previous section it was pointed out that the two equivalent problems

$$\left. \begin{array}{l} \max \quad cx \\ Ax \leq b \\ x \geq 0 \end{array} \right\} = \left\{ \begin{array}{l} \max \quad cx + 0s \\ Ax + Is = b \\ x, s \geq 0 \end{array} \right.$$

give rise to exactly the same dual problem. This is true in general. Suppose P is any given minimization problem in variables, which may be non-negative, non-positive or free. Let P' be a minimization problem in standard form, i.e a problem in non-negative variables with equality constraints, constructed from P by means of addition of slack variables to \leq -constraints, subtraction of surplus variables from \geq -constraints, and change of variables. Then the dual problems of P and P' are equal.

We have commented upon the addition of slack variables to \leq -constraints in the preceding section. The subtraction of slack variables are dealt with similarly. A constraint

$$a_{i1}x_1 + \cdots + a_{in}x_n \geq b_i \Leftrightarrow (b_i - a_{i1}x_1 - \cdots - a_{in}x_n) \leq 0$$

gives rise to a multiplier, which must be non-negative in order for the relaxed objective function to provide a lower bound for the original one on the feasible set. If a slack variable is subtracted from the left-hand side of the inequality constraint to obtain an equation

$$a_{i1}x_1 + \cdots + a_{in}x_n - s_i = b_i \Leftrightarrow (b_i - a_{i1}x_1 - \cdots - a_{in}x_n) + s_i = 0$$

the multiplier must now be allowed to vary over \mathcal{R} . A new constraint in the dual problem, however, is introduced by the column of the slack variable, cf. Section 2:

$$-y_i \leq 0 \Leftrightarrow y_i \geq 0,$$

thereby reintroducing the sign constraint for y_i .

If a non-positive variable x_j is substituted by x'_j of opposite sign, all signs in the corresponding column of the Simplex tableau change. For minimization purposes however, a positive sign of the coefficient of a non-positive variable is beneficial, whereas a negative sign of the coefficient of a non-negative variable is preferred. The sign change of the column in combination with the change in preferred sign of the objective function coefficient leaves the dual constraint unchanged.

Finally, if a free variable x_j is substituted by the difference between two non-negative variables x'_j and x''_j two equal columns of opposite sign are introduced. These give rise to two dual constraints, which when taken together result in the same dual equality constraint as obtained directly.

The proof of the Duality Theorem for all types of dual pairs P and DP of LP problems may hence be given as follows: Transform P into a standard problem P' in the well known fashion. P' also has DP as its dual problem. Since the Duality Theorem holds for P' and DP as shown previously and P' is equivalent to P , the theorem also holds for P and DP .

4 Discussion: Pros and Cons of the Approach

The main advantages of teaching duality based on multipliers are in my opinion

- the independence of the problem modeled by the primal model and the introduction of the dual problem, i.e. that no story has to go with the dual problem,
- the possibility to avoid problem transformation and “duality by definition” in the introduction of general duality in linear programming,
- the clarification of the interplay between the sign of variables and the type of the corresponding constraints in the dual pair of problems,
- the early introduction of the idea of getting information about the optimum of an optimization problem through bounding using the solution of an easier problem,
- the possibility of introducing partial dualization by including only some constraint violations in the objective function, and
- the resemblance with duality in non-linear programming, cf. [3].

The only disadvantage in my view is one listed also as an advantage:

- the independence of the introduction of the dual problem and the problem modeled by the primal model

since this may make the initial motivation weaker.

I do not advocate that duality should be taught based solely on the multiplier approach, but rather that it is used as a supplement to the traditional presentation (or vice versa). In my experience, it offers a valuable supplement, which can be used to avoid the situation of frustrated students searching for an intuitive interpretation of the dual problem in cases, where such an interpretation is not natural. The decision on whether to give the traditional presentation of duality or the multiplier approach first of course depends on the particular audience.

Acknowledgement

I am grateful to Professor Susan Powell, London School of Economics, for introducing the possibility of performing Lagrangean Relaxation on an LP problem to me, and to Professor Jakob Krarup, DIKU, for valuable comments and for a pointer to [1].

References

- [1] R. E. Bellman and S. E. Dreyfus, *Applied Dynamic Programming*, Princeton University Press, 1962.
- [2] G. B. Dantzig, *Linear Programming and Extensions*, Princeton University Press, 1963.
- [3] A. M. Geoffrion, *Duality in Non-Linear Programming: A Simplified Applications-Oriented Approach*, *SIAM Review* 13, 11 - 37, 1971.
- [4] A. M. Geoffrion, *Lagrangean Relaxation for Integer Programming*, *Math. Prog. Study* 2 82 - 114, 1974.
- [5] M. X. Goemans, *Linear Programming, Course Notes*, 1994, available from <http://theory.lcs.mit.edu/goemans>.
- [6] J. Krarup and C. Vanderhoeft, *Belgium's Best Beer and Other Stories*, VUB Press (to appear).
- [7] K. G. Murty, *Linear and Combinatorial Programming*, John Wiley, 1976.
- [8] H. P. Williams, *Model Solving in Mathematical Programming*, John Wiley, 1993.
- [9] W. L. Winston, *Operations Research - Applications and Algorithms*, Int. Thomson Publishing, 1994.