Lecture 5

Uri Feige

May 15, 2002

1 Linear programming duality

1.1 The diet problem revisited

Recall the diet problem from Lecture 1. There are n foods, m nutrients, and a person (the *buyer*) is required to consume at least b_i units of nutrient i (for $1 \leq i \leq m$). Let a_{ij} denote the amount of nutrient i present in one unit of food j. Let c_i denote the cost of one unit of food item i. One needs to design a diet of minimal cost that supplies at least the required amount of nutrients. This gives the following linear program.

minimize $c^T x$
subject to
$Ax \geq b$
$x \geq 0$

Now assume that some other person (the *seller*) has a way of supplying the nutrients directly, not through food. (For example, the nutrients may be vitamins, and the seller may sell vitamin pills.) The seller wants to charge as much as he can for the nutrients, but still have the buyer come to him to buy nutrients. A plausible constraint in this case is that the price of nutrients is such that it is never cheaper to buy a food in order to get the nutrients in it rather than buy the nutrients directly. If y is the vector of nutrient prices, this gives the constraints $A^T y \leq c$. In addition, we have the nonnegativity constrain $y \geq 0$. Under these constraints the seller wants to set the prices of the nutrients in a way that would maximize the sellers profit (assuming that the buyer does indeed by all his nutrients from the seller). This gives the the following *dual* LP:

maximize $b^T y$
Subject to
$A^T y \leq c$
$y \geq 0$

As one can replace any food by its nutrients and not pay more, one gets *weak duality*, namely, the dual provides a lower bound for the primal. Weak duality goes beyond the diet problem and holds even if A, b, c have some entries that are negative. That is, for every pair of feasible solutions to the primal and dual LPs we have:

$$b^T y \leq (Ax)^T y = x^T A^T y \leq x^T c = c^T x \tag{1}$$

In particular, weak duality implies that if the optimal value of the primal is unbounded then the dual is infeasible, and if the optimal value of the dual is unbounded, then the primal is infeasible.

Assume that there is a pair of solutions x^* and y^* for which the values of the primal and dual LPs are equal, namely $c^T x^* = b^T y^*$. Then necessarily both x^* and y^* are optimal solutions to their respective LPs. In economics, the vector y^* is refereed to as *shadow prices*. These optimal solutions need to satisfy the inequalities of (1) with equality. This gives the following *complementary slackness* conditions:

$$(Ax^* - b)^T y^* = 0 \tag{2}$$

$$(c - A^T y^*)^T x^* = 0 \tag{3}$$

Condition (2) has the following economic interpretation. If a certain nutrient is in surplus in the optimal diet, then its shadow price is free (a *free good*). Condition (3) can be interpreted to say that if a food is overpriced (more expensive than the shadow price of its nutrients) then this food does not appear in the optimal diet.

1.2 Derivation of dual for standard form

Consider a (primal) linear program in standard form.

minimize $c^T x$
suject to
$Ax = b$
$x \geq 0$

Let us derive a lower bound on its value. The lower bound is obtained by removing the constraints $Ax = b$, and replacing them by a charging mechanism for violated constraints. There is a price y_i for one unit of violation of constraint $A_i x = b_i$. The following lower bound is valid for every vector y of prices (because feasible solutions of the primal LP give the same value for this new LP):

minimize (over x**)** $c^T x + y^T(b - Ax)$
suject to
$x \geq 0$

Rearranging the objective function, one gets:

minimize (over x**)** $y^T b + (c^T - y^T A)x$

Unless $c^T - y^T A \geq 0^T$, this lower bound can be made $-\infty$, and hence useless. So we shall restrict attention to y satisfying $y^T A \leq c^T$.

If we now remove the nonnegative $(c^T - y^T A)x$ from the objective function, we still have a lower bound. Moreover, this lower bound no longer depends on x, so we may remove the minimization over x, and the constraints $x \geq 0$. Hence for every y satisfying $y^T A \leq c^T$, $y^T b$ is a lower bound on the primal LP. What is the best (highest) lower bound that we can get by this approach? This is simply the solution to the dual LP:

maximize (over y**)** $b^t y$
subject to
$A^T y \leq c$

Observe that complementary slackness condition (3) can be extracted from the above development. Complementary slackness condition (2) is trivially true for LPs in standard form.

1.3 Strong duality

If the primal is feasible and its optimal solution is bounded, then the lower bound given by the dual is tight. Namely, the value of the dual is equal it this case to value of the primal. We shall show this for LPs in standard form.

The proof we give is based on the stopping condition for the simplex algorithm. Recall that at that point $X_B = B^{-1}b$, and the reduced costs are nonnegative, namely $c_j - c_B^T B^{-1} A_j \geq 0$, or in matrix notation:

$c^T - c_B^T B^{-1} A \geq 0^T$.

Let us "guess" the following solution to the dual: $y^T = c_B^T B^{-1}$. Then

$y^T A = c_B^T B^{-1} A \leq c^T$

satisfying the constraints of the dual. Hence y is a feasible solution to the dual. Its value is:

$b^T y = y^T b = c_B^T B^{-1} b = c_B^T x_B = c^T x$,

implying that the lower bound is tight.

1.4 LPs in general form

The following table explains how to obtain the dual of a primal LP that is in general form. Here A_j denotes a row of matrix A and A^j denotes a column.

$\min c^T x$		$\max b^T y$
$A_i x \geq b_i$	$i \in I^+$	$y_i \geq 0$
$A_i x = b_i$	$i \in I^=$	y_i free
$x_j \geq 0$	$j \in J^+$	$y^T A^j \leq c_j$
x_j free	$j \in J^=$	$y^T A^j = c_j$

Note that the dual of the dual is the primal. (Verifying this is left as homework.)

Weak and strong duality apply also in this case. More specifically, if the optimum to the primal is bounded, then so is the optimum to the dual, and vice versa. If the optimum to one of the LPs is unbounded, then the other is not feasible. It may also happen that neither one of them is feasible, as can be shown by the following example:

minimize $-x_1 + 2x_2$
subject to
$x_1 - x_2 = -1$
$-x_1 + x_2 = 2$

Coincidently, its dual happens to be the same:

maximize $-y_1 + 2y_2$
subject to
$y_1 - y_2 = -1$
$-y_1 + y_2 = 2$

1.5 A geometric view of the dual

Consider the primal LP

minimize $c^T x$
subject to
$Ax \geq b$

Its dual is:

maximize $b^t y$
subject to
$A^T y = c$
$y \geq 0$

Hence in a feasible solution to the dual, c is expressed as a positive combitation of the rows of A. The rows of A are the normals to the hyperplanes defining the primal LP, pointing into the polyhedron. In the optimal solution to the dual, we get more (value for the objective function) by including those normals with highest values of b. Intuitively, those are the constraints that are hardest to satisfy in the primal.

1.6 Farkas' lemma

Farkas' lemma states that given a system of equation $Ax = b$ for which we seek a nonnegative solution (meaning $x \geq 0$), then either the system has a solution, or (and this is an *exclusive or*) there is the following type of explanation for the insolvability of the system: there is some vector y such that $y^T b > 0$, but $y^T A \leq 0$.

Geometrically, this vector y is a normal to a hyperplane that separates b from the columns of A.

To prove Farkas' lemma we use the following LP primal-dual pair (though it would be fair to note that historically, Farkas' lemma predated LP duality).

minimize $0^T x$
subject to
$Ax = b$
$x \geq 0$

and

maximize $b^t y$
subject to
$y^t A \leq 0$

If the primal is feasible, its optimal value is 0, which upper bounds the dual, refuting the second alternative.

If the primal is not feasible, then the dual must be unbounded (because it is feasible), proving the second alternative.

1.7 Some easy consequences of duality

Using duality it is a simple matter to show that the following problem is in $NP \bigcap coNP$:

Given a minimization LP, does it have a feasible solution of value at most k?

A witness for this statement is a feasible solution to the LP with value at most k. A refutation is a solution for the dual of value higher than k. Both solutions have complexity polynomial in the input size.

We can also show that finding a feasible solution to an LP is essentially as hard as finding the optimal solution (when the optimum is bounded). Given an LP, consider a new LP composed of all constraints of the dual and the primal, and the additional constraints that both objective functions must have equal value. Its size is polynomially related to the size of the original LP. A feasible solution to this new LP is an optimal solution to the original LP.