Solved exercises for the course of

Foundations of Operations Research

Roberto Cordone

The dual simplex method

Given the following *LP* problem:

$$\begin{aligned} \max z = 5x_1 + 8x_2 & & \\ x_1 + x_2 & \leq & 6 \\ 5x_1 + 9x_2 & \leq & 45 \\ x_1, x_2 & \geq & 0 \end{aligned}$$

1. solve it and write the *tableau* corresponding to the optimal basic solution;
2. add to the problem constraint $-1/4x_3 - 3/4x_4 \leq -1/4$, where x_3 and x_4 are the *slack* variables of the two original constraints, and determine whether the previous optimal solution is still feasible; if not, compute the new optimal solution with the dual simplex method;
3. add to the problem constraint $-1/3x_3 - 2/3x_5 \leq -1/3$ and reoptimize;
4. add to the problem constraint $-1/2x_3 - 1/2x_6 \leq -1/2$ and reoptimize.

Solution

Let us solve the problem graphically. Figure 1 shows the result: the optimal point is $(9/4, 15/4)$, which corresponds to the optimal solution $x_1 = 9/4$, $x_2 = 15/4$, $x_3 = x_4 = 0$.

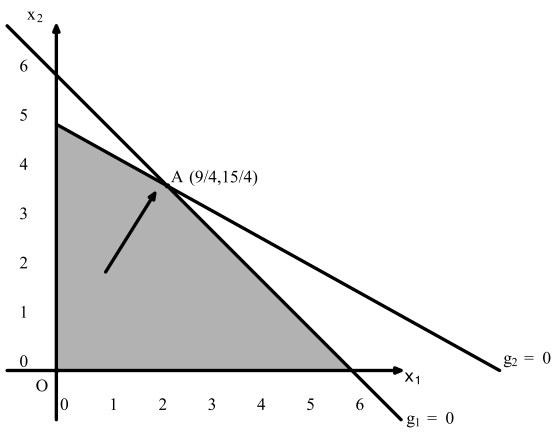

Figura 1: Graphical solution

In order to obtain the full *tableau*, either we apply the algebraic method, inverting the basic submatrix B, or (equivalently) we apply the simplex method, which yields in two iterations the following result:

41 + 1/4	0	0	5/4	3/4
9/4	1	0	9/4	-1/4
15/4	0	1	-5/4	1/4

The dual simplex method

The dual simplex method allows to start from a basic solution with nonnegative reduced costs (in general, unfeasible) to a basic optimal solution, visiting a sequence of basic solutions with nonnegative reduced costs and decreasing (or at least nonincreasing) unfeasibility. In other words, while the primal simplex method preserves solution feasibility and gradually approaches the optimality condition $\bar{c} \geq 0$, the dual simplex method preserves the optimality condition (also called *dual feasibility*) and gradually approaches feasibility.

The dual simplex method provides a much simpler alternative to the two-phase method for the cases in which the starting solution is unfeasible. However, while the two-phase method can always be applied, the dual simplex method can be applied only when all reduced costs are nonnegative. The most common practical cases in which this occurs are:

- when new regulations or unexpected events introduce an additional constraint or make an existing constraint tighter for an *LP* problem whose optimal solution has already been computed, so as to make that solution unfeasible;

- when the continuous relaxation of an *ILP* problem has been solved and the relaxed solution is fractionary (hence, unfeasible for the original problem); in this case, several techniques allow to introduce into the problem one or more constraints (named *cutting planes* or *cuts*), which forbid the relaxed solution, without removing any integer one.

The dual simplex method consists in applying *pivot* operations, just as in the primal simplex, but selecting the *pivot* element with different rules:

1. the *pivot* row is any row with a negative right-hand-side (in general, one selects the *most negative* one)

2. the *pivot* element is the *negative* element $a_{ij} < 0$ with the minimum absolute ratio between the column cost and the element: $\min_{j:a_{ij}<0} c_j/|a_{ij}|$ (notice the absolute value operator applied to a_{ij})

The rule is clearly "dual" with respect to the one used in the primal simplex: the element is negative instead of positive; the row is chosen first, instead of the column; the ratios considered concern columns, instead of rows. The element chosen is still the one with the minimum ratio, and the heuristic rule for the first choice still prefers the most negative value.

An operational explanation of the method is the following one: the *pivot* operation concerns a negative element so that the right-hand-side, divided by a_{ij}, will become positive; the choice of the column is made so as to guarantee that no reduced cost will become negative.

There is, however, a deeper relation to duality. Consider a primal problem in basic canonical form:

$$\begin{aligned} \min z = c_N\, x_N & \\ N\, x_N + x_B \;&=\; b \\ x_N, x_B \;&\geq\; 0 \end{aligned}$$

where B and N are the basic and nonbasic submatrix, respectively. The problem can be equivalently written as:

$$\begin{aligned} \max z = -c_N\, x_N & \\ N\, x_N \;&\leq\; b \\ x_N \;&\geq\; 0 \end{aligned}$$

and its dual is

$$\begin{aligned} \min w = b_N\, y_N & \\ N^T\, y_N \;&\geq\; -c \\ y_N \;&\geq\; 0 \end{aligned}$$

which can be written, replacing $N' y_N - y_B = -c$ with $-N' y_N + y_B = c$, as

$$\begin{aligned} \min w = b_N\, y_N & \\ -N^T\, y_N + y_B \;&=\; c \\ y_N, y_B \;&\geq\; 0 \end{aligned}$$

This shows that the *tableau* can be seen as a simultaneous representation of the primal and the dual problem. This holds for all pairs of basic solutions (primal and dual) which correspond to each other through the complementary slackness conditions.

Now, a *pivot* operation on the primal problem modifies the primal basis, but correspondingly also the dual basis. Conversely, a *pivot* operation on the dual problem modifies the dual basis, but also the primal one. Now, the dual simplex method rule simply corresponds to switching to the dual problem, applying to it the standard *pivot* rule, and switching back to the primal problem:

- the negative elements of the primal correspond to the positive ones of the dual;
- the negative right-hand-sides of the primal correspond to the negative reduced costs of the dual;

- the minimum ratio $\min_{j:a_{ij}<0} c_j/|a_{ij}|$ of the primal corresponds to the minimum ratio between right-hand-side and *pivot* element of the dual.

The *pivot* operation, however, is always the same, working on the rows of the problem (not on the columns), because we are working on the primal problem and visiting primal basic solution: the *tableau* of the dual is used only as a guide: we are performing on the primal problem a move which makes the dual solution approach its optimum. The method works because approaching the dual optimum means approaching also the primal optimum.

41 + 1/4	0	0	5/4	3/4
9/4	1	0	9/4	-1/4
15/4	0	1	-5/4	1/4
0			1	0
0			0	1

First additional constraint

Constraint $-1/4x_3 - 3/4x_4 \leq -1/4$ turns into $-1/4x_3 - 3/4x_4 + x_5 = -1/4$ in order to keep the problem in standard form. This requires to add a new column and row to the original *tableau*.

41 + 1/4	0	0	5/4	3/4	0
9/4	1	0	9/4	-1/4	0
15/4	0	1	-5/4	1/4	0
-1/4	0	0	-1/4	-3/4	1

The dual simplex method requires to select a *pivot* element. The *pivot* row is row 3, since it is the only one with a negative right-hand-side. Column 4 is the one with minimum ratio $c_j/|a_{3j}|$ among those with $a_{3j} < 0$. So, the *pivot* element is a_{34}.

41	0	0	1	0	1
7/3	1	0	7/3	0	-1/3
11/3	0	1	-4/3	0	1/3
1/3	0	0	1/3	1	-4/3

The new optimal solution is $x_1 = 7/3$, $x_2 = 11/3$, $x_3 = x_5 = 0$ and $x_4 = 1/3$. The objective value is obviously worse due to the new constraint; it increased from $-41 - 1/4$ to -41.

Second additional constraint

Constraint $-1/3x_3 - 2/3x_5 \leq -1/3$ turns into $-1/3x_3 - 2/3x_5 + x_6 = -1/3$ in order to keep the problem in standard form. This requires to add a new column and row to the original *tableau*.

41	0	0	1	0	1	0
7/3	1	0	7/3	0	-1/3	0
11/3	0	1	-4/3	0	1/3	0
1/3	0	0	1/3	1	-4/3	0
-1/3	0	0	-1/3	0	-2/3	1

Performing a *pivot* operation on element a_{45}, one obtains

40 + 1/2	0	0	1/2	0	0	3/2
5/2	1	0	5/2	0	0	-1/2
7/2	0	1	-3/2	0	0	1/2
1	0	0	1	1	0	-2
1/2	0	0	1/2	0	1	-3/2

Third additional constraint

Finally, constraint $-1/2x_3 - 1/2x_6 \leq -1/2$ turns into $-1/2x_3 - 1/2x_6 + x_7 = -1/2$.

40 + 1/2	0	0	1/2	0	0	3/2	0
5/2	1	0	5/2	0	0	-1/2	0
7/2	0	1	-3/2	0	0	1/2	0
1	0	0	1	1	0	-2	0
1/2	0	0	1/2	0	1	-3/2	0
-1/2	0	0	-1/2	0	0	-1/2	1

from which, performing a *pivot* operation on a_{53}:

40	0	0	0	0	0	1	1
0	1	0	0	0	0	3	5
5	0	1	0	0	0	2	-3
0	0	0	0	1	0	-3	2
0	0	0	0	0	1	-2	1
1	0	0	1	0	0	1	-2

Any further constraint can be treated in the same way. Notice that in all three examples a single *pivot* operation was always sufficient to achieve feasibility (and optimality). This is not true in general.

As well, in all three examples, the additional constraint could be introduced almost directly in the problem producing a basic canonical form, though an unfeasible one. This is not true in general: the new constraint could actually require the application of further *pivot* operations to introduce the basic canonical form, and this could destroy the optimality condition $\bar{c} \geq 0$, thus forbidding the use of the dual simplex method. In the special case of cutting planes, however, the new constraints usually respect the canonical form, and the dual simplex method can be directly applied.