

4 Advanced Simplex Procedures

4.1 The Two-Phase Simplex Method

The LP we solved in the previous lecture allowed us to find an initial BFS very easily. In cases where such an obvious candidate for an initial BFS does not exist, we use an additional phase I to find a BFS. In phase II we then proceed as in the previous lecture.

Consider the LP to

$$\begin{aligned} &\text{maximize} && -6x_1 - 3x_2 \\ &\text{subject to} && x_1 + x_2 \geq 1 \\ & && 2x_1 - x_2 \geq 1 \\ & && 3x_2 \leq 2 \\ & && x_1, x_2 \geq 0, \end{aligned}$$

and introduce slack variables to obtain

$$\begin{aligned} &\text{maximize} && -6x_1 - 3x_2 \\ &\text{subject to} && x_1 + x_2 - z_1 = 1 \\ & && 2x_1 - x_2 - z_2 = 1 \\ & && 3x_2 + z_3 = 2 \\ & && x_1, x_2, z_1, z_2, z_3 \geq 0. \end{aligned}$$

Unfortunately, the basic solution with $x_1 = x_2 = 0$, $z_1 = z_2 = -1$, and $z_3 = 2$ is not feasible. We can, however, add an *artificial variable* to the left-hand side of each constraint where the slack variable and the right-hand side have opposite signs, and then minimize the sum of the artificial variables starting from the obvious BFS where the artificial variables are non-zero instead of the corresponding slack variables. In the example, we

$$\begin{aligned} &\text{minimize} && y_1 + y_2 \\ &\text{subject to} && x_1 + x_2 - z_1 + y_1 = 1 \\ & && 2x_1 - x_2 - z_2 + y_2 = 1 \\ & && 3x_2 + z_3 = 2 \\ & && x_1, x_2, z_1, z_2, z_3, y_1, y_2 \geq 0, \end{aligned}$$

and the goal of phase I is to solve this LP starting from the BFS where $x_1 = x_2 = z_1 = z_2 = 0$, $y_1 = y_2 = 1$, and $z_3 = 2$. If the original problem is feasible, we will be able to find a BFS where $y_1 = y_2 = 0$. This automatically gives us an initial BFS for the original problem.

In summary, the two-phase simplex method proceeds as follows:

1. Bring the constraints into equality form. For each constraint in which the slack variable and the right-hand side have opposite signs, or in which there is no slack variable, add a new artificial variable that has the same sign as the right-hand side.
2. Phase I: minimize the sum of the artificial variables, starting from the BFS where the absolute value of the artificial variable for each constraint, or of the slack variable in case there is no artificial variable, is equal to that of the right-hand side.
3. If some artificial variable has a positive value in the optimal solution, the original problem is infeasible; stop.
4. Phase II: solve the original problem, starting from the BFS found in phase I.

While the original objective is not needed for phase I, it is useful to carry it along as an extra row in the tableau, because it will then be in the appropriate form at the beginning of phase II. In the example, phase I therefore starts with the following tableau:

	x_1	x_2	z_1	z_2	z_3	y_1	y_2	
y_1	1	1	-1	0	0	1	0	1
y_2	2	-1	0	-1	0	0	1	1
z_3	0	3	0	0	1	0	0	2
II	-6	-3	0	0	0	0	0	0
I	3	0	-1	-1	0	0	0	2

Note that the objective for phase I is written in terms of the variables that are not in the basis. This can be obtained by first writing it in terms of y_1 and y_2 , such that we have -1 in the columns for y_1 and y_2 and 0 in all other columns because we are *maximizing* $-y_1 - y_2$, and then adding the first and second row to make the entries for all variables in the basis equal to zero.

Phase I now proceeds by pivoting on a_{21} to get

	x_1	x_2	z_1	z_2	z_3	y_1	y_2	
	0	$\frac{3}{2}$	-1	$\frac{1}{2}$	0	1	$-\frac{1}{2}$	$\frac{1}{2}$
	1	$-\frac{1}{2}$	0	$-\frac{1}{2}$	0	0	$\frac{1}{2}$	$\frac{1}{2}$
	0	3	0	0	1	0	0	2
II	0	-6	0	-3	0	0	3	3
I	0	$\frac{3}{2}$	-1	$\frac{1}{2}$	0	0	$-\frac{3}{2}$	$\frac{1}{2}$

and on a_{14} to get

	x_1	x_2	z_1	z_2	z_3	y_1	y_2	
	0	3	-2	1	0	2	-1	1
	1	1	-1	0	0	1	0	1
	0	3	0	0	1	0	0	2
II	0	3	-6	0	0	6	0	6
I	0	0	0	0	0	-1	-1	0

Note that we could have chosen a_{12} as the pivot element in the second step, and would have obtained the same result.

This ends phase I as $y_1 = y_2 = 0$, and we have found a BFS for the original problem with $x_1 = z_2 = 1$, $z_3 = 2$, and $x_2 = z_1 = 0$. After dropping the columns for y_1 and y_2 and the row corresponding to the objective for phase I, the tableau is in the right form for phase II:

	x_1	x_2	z_1	z_2	z_3	
	0	3	-2	1	0	1
	1	1	-1	0	0	1
	0	3	0	0	1	2
	0	3	-6	0	0	6

By pivoting on a_{12} we obtain the following tableau, corresponding to an optimal solution of the original problem with $x_1 = 2/3$, $x_2 = 1/3$, and value -5 :

	x_1	x_2	z_1	z_2	z_3	
	0	1	$-\frac{2}{3}$	$\frac{1}{3}$	0	$\frac{1}{3}$
	1	0	$-\frac{1}{3}$	$-\frac{1}{3}$	0	$\frac{2}{3}$
	0	0	2	-1	1	1
	0	0	-4	-1	0	5

It is worth noting that the problem we have just solved is the dual of the LP in Example 2.3, which we solved in the previous lecture, augmented by the constraint $3x_2 \leq 2$. Ignoring the column and row corresponding to z_3 , the slack variable for this new constraint, the final tableau is essentially the negative of the transpose of the final tableau we obtained in the previous lecture. This makes sense because the additional constraint is not tight in the optimal solution, as we can see from the fact that $z_3 \neq 0$.

4.2 The Dual Simplex Method

The (primal) simplex method maintains feasibility of the primal solution along with complementary slackness and seeks feasibility of the dual solution. Alternatively one

could maintain feasibility of the dual solution and complementary slackness and seek feasibility of the primal solution. This is known as the dual simplex method.

One situation where the dual simplex method can be useful is when an initial feasible solution for the dual is easier to find than one for the primal. Consider the following LP, to which we have already added slack variables z_1 and z_2 :

$$\begin{aligned} \text{minimize} \quad & 2x_1 + 3x_2 + 4x_3 \\ \text{subject to} \quad & x_1 + 2x_2 + x_3 - z_1 = 3 \\ & 2x_1 - x_2 - 3x_3 - z_2 = 4 \\ & x_1, x_2, x_3, z_1, z_2 \geq 0. \end{aligned}$$

The primal simplex algorithm would have to use two phases, since the solution where $z_1 = -3$ and $z_2 = -4$ is not feasible. On the other hand, $c \geq 0$, so the dual solution with $\lambda_1 = \lambda_2 = 0$ satisfies $c^T - \lambda^T A \geq 0$ and is therefore feasible. We obtain the following tableau:

-1	-2	-1	1	0	-3
-2	1	3	0	1	-4
2	3	4	0	0	0

In the dual simplex algorithm the pivot is selected by picking a row i such that $a_{i0} < 0$ and a column $j \in \{j' : a_{ij'} < 0\}$ that minimizes $-a_{0j}/a_{ij}$. Pivoting then works just like in the primal algorithm. In the example we can pivot on a_{21} to obtain

0	$-\frac{5}{2}$	$-\frac{5}{2}$	1	$-\frac{1}{2}$	-1
1	$-\frac{1}{2}$	$-\frac{3}{2}$	0	$-\frac{1}{2}$	2
0	4	7	0	1	-4

and then on a_{12} to obtain

0	1	1	$-\frac{2}{5}$	$\frac{1}{5}$	$\frac{2}{5}$
1	0	-1	$-\frac{1}{5}$	$-\frac{2}{5}$	$\frac{11}{5}$
0	0	3	$\frac{8}{5}$	$\frac{1}{5}$	$-\frac{28}{5}$

We have reached the optimum of $28/5$ with $x_1 = 11/5$, $x_2 = 2/5$, and $x_3 = 0$.

It is worth pointing out that for problems in which all constraints are inequality constraints, the optimal dual solution can also be read off from the final tableau. For problems of this type, the last $n - m$ columns of the extended constraint matrix A correspond to the slack variables and therefore contain values 1 or -1 on the diagonal and 0 everywhere else. For the same reason, the last $n - m$ columns of the vector c^T are 0. The values of the dual variables, each of them with opposite sign of the slack variable in the corresponding constraint, thus appear in the last $n - m$ columns of the vector $(c^T - \lambda^T A)$ in the last row of the final tableau. In our example, we have $\lambda_1 = 8/5$ and $\lambda_2 = 1/5$.

4.3 Gomory's Cutting Plane Method

Another situation where the dual simplex method can be useful is when we need to add constraints to an already solved LP. While such constraints can make the primal solution infeasible, they do not affect feasibility of the dual solution. We can therefore simply add the constraint and continue running the dual LP algorithm from the current solution until the primal solution again becomes feasible. The need to add constraints to an LP for example arises naturally in Gomory's cutting plane approach for solving integer programs (IPs). An IP is a linear program with the additional requirement that variables should be integral.

Assume that for a given IP we have already found an optimal (fractional) solution x^* with basis B , and let a_{ij} denote the entries of the final tableau, i.e., $a_{ij} = (A_B^{-1}A_j)_i$ and $a_{i0} = (A_B^{-1}b)_i$. If x^* is not integral, there has to be a row i such that a_{i0} is not integral, and for every feasible solution x ,

$$x_i + \sum_{j \in N} \lfloor a_{ij} \rfloor x_j \leq x_i + \sum_{j \in N} a_{ij} x_j = a_{i0}.$$

The inequality holds because x is feasible, i.e., $x \geq 0$, the equality follows from the properties of the final tableau. If x is integral, the left-hand side is integral as well, and the inequality must still hold if the right-hand side is rounded down. Thus,

$$x_i + \sum_{j \in N} \lfloor a_{ij} \rfloor x_j \leq \lfloor a_{i0} \rfloor.$$

This inequality is satisfied by every (integral) feasible solution, but not by the current solution x^* , for which $x_i^* = a_{i0}$. It corresponds to a so-called *cutting plane*, a hyperplane that separates the current solution x^* from the feasible set. The idea behind the cutting plane method is to iteratively add cutting planes and solve the resulting linear programs using the dual simplex algorithm. As it turns out, this always leads to an optimal integral solution after a finite number of steps.

Consider again the final tableau on Page 22, and assume that we are now looking for an integral solution. By the first row, and assuming that all variables are integral and non-negative,

$$x_2 + x_3 - 1z_1 + 0z_2 \leq x_2 + x_3 - \frac{2}{5}z_1 + \frac{1}{5}z_2 = \frac{2}{5},$$

and in fact

$$x_2 + x_3 - z_1 \leq 0.$$

If we turn this into an equality constraint using a new slack variable, add it to the tableau, and bring it into the right form by subtracting the first constraint from it, we obtain

0	1	1	$-\frac{2}{5}$	$\frac{1}{5}$	0	$\frac{2}{5}$
1	0	-1	$-\frac{1}{5}$	$-\frac{2}{5}$	0	$\frac{11}{5}$
0	0	0	$-\frac{3}{5}$	$-\frac{1}{5}$	1	$-\frac{2}{5}$
0	0	3	$\frac{8}{5}$	$\frac{1}{5}$	0	$-\frac{28}{5}$

After one more round of the dual simplex algorithm we reach the optimal integral solution with $x_1 = 3$ and $x_2 = x_3 = 0$:

0	1	1	-1	0	1	0
1	0	-1	1	0	-2	3
0	0	0	3	1	-5	2
0	0	3	1	0	1	-6

We will return to IPs, and learn about a different method for solving them that often works better in practice, in a later lecture.