The Primal-Dual Simplex Method

A Parametric Method

Math 354, Spring 2006

With the simplex method and the dual simplex method, we can solve any LPP as long as we can find a starting point–a basic solution which is feasible for either the primal or the dual problem. In the two-phase method, the goal of the first phase is to obtain such a starting point for the second phase.

A more unified method is the parametric self-dual simplex method or "primal-dual" method. The idea is to **begin by altering the LPP** to one with an obvious initial tableau which is feasible and dual-feasible[1], so is optimal. The altered problem is obtained from the original problem by increasing or decreasing certain coefficients by a large but unspecified number μ. The focus then shifts to getting rid of the parameter μ without losing optimality. Getting rid of it means making $\mu = 0$. We keep μ in the tableau **symbolically**, i.e., as a variable, and we imagine it taking on various values at various stages of the algorithm. At first, we imagine μ to be very large, and the goal is to reduce μ to $\mu = 0$ without giving up optimality. Then we will have an optimal solution of the original LPP.

We give an example[2] below, and discuss it there in some detail. But first, here is the plan. The LPP translates to Tableau 0, which is neither feasible nor dual-feasible. The parametrized problem begins with Tableau 1, which is feasible and dual-feasible if μ is large enough.

So imagine that μ starts out very large, and then shrinks toward 0, slowly. As μ shrinks, the tableau may suddenly stop being feasible or dual-feasible when μ crosses certain "threshold" values. When this happens, we immediately pause to "fix" the tableau. If we lost dual-feasibility (resp. feasibility) we restore it by the simplex method (resp. the dual simplex method). Having done that, we have a tableau which is feasible and dual-feasible for smaller values of μ than before. It will not be feasible and dual-feasible for very large μ any more, but we don't care because our goal is $\mu = 0$. See Tableau 2 below.

Now let μ resume shrinking toward 0, and whenever the tableau loses feasibility or dual-feasibility, pause to "repair" the tableau. As there are only finitely many possible sets of basic variables, there are only finitely

[1] We continue to use the term "dual-feasible" for a tableau satisfying the optimality criterion.

[2] The example is taken from Robert J. Vanderbei, *Linear Programming, Foundations and Extensions*, 2nd Edition, Springer, 2001. A highly recommended book.

many possible "threshold" values that μ can cross. Therefore, after finitely many repairs, this procedure finally reaches a tableau that is feasible and dual-feasible for $\mu = 0$, i.e., for the original problem, as in Tableaux 3, 4, 5.

Of course there is also the possibility that before reaching $\mu = 0$, we reach a tableau showing that for some $\mu > 0$, the perturbed problem is infeasible or unbounded. Then we stop and can conclude that the original LPP or its dual was infeasible (exercise).

While we are reducing μ we "freeze" the basis until we cross a threshold value of μ. Then, we pause and freeze μ, and use simplex or dual-simplex to change the basis. Then we re-freeze the basis and reduce μ, and so on. Throughout, as μ changes, so does the LPP that the tableau represents.[3]

THE EXAMPLE.

Maximize $z = -2x + 3y$
subject to

$$\begin{aligned} -x + \ y &\le -1 \\ -x - 2y &\le -2 \\ y &\le 1 \\ x, y &\ge 0 \end{aligned}$$

	x	y	u	v	w	
u	-1	1	1	0	0	-1
v	-1	-2	0	1	0	-2
w	0	1	0	0	1	1
	2	-3	0	0	0	0

TABLEAU 0: original LPP

By adding slack variables we get Tableau 0. By duality theory, we want a tableau that is both feasible and dual-feasible. Our Tableau 0 is neither. We "perturb" the original problem, adding μ to each right side, and subtracting μ from each non-slack coefficient of the objective function. Thus μ is **added** to the objective row entries in Tableau 1. Each value of μ corresponds to a different LPP. (Our original problem corresponds to $\mu = 0$.)

	x	y	u	v	w	
u	-1	1	1	0	0	$-1+\mu$
v	-1	-2	0	1	0	$-2+\mu$
w	0	1	0	0	1	$1+\mu$
	$(2+\mu)$	$(-3+\mu)$	0	0	0	0

TABLEAU 1: feasible and dual-feas. for $\mu \geq 3$

[3] This is an example of a general technique sometimes called "deformation", where we introduce a parameter μ into a given problem, making a family of problems, one for each value of μ. One tries to do this in a way that (a) the problem is easy (or easier) for some value(s) of μ, and (b) one can keep track of what happens to the solution as we let $\mu \to 0$.

Tableau 1 is the beginning of our analysis of these problems. Keep in mind that we really only care about $\mu = 0$. If μ is very large, Tableau 1 is clearly feasible and dual-feasible. Instead of specifying a value for μ, we observe that as long as $\mu \geq 3$, this tableau is feasible and dual-feasible. We will reduce the perturbation factor μ to 0, gradually, repairing any loss of feasibility or dual-feasibility as soon as it happens.

You should try to spot the threshold values of μ from the tableaux, and use the text to check yourself. What is the threshold for Tableau 1? We just watch for values of μ where one of $2+\mu$, $-3+\mu$, $-1+\mu$, $-2+\mu$, or $1+\mu$ changes sign. As μ descends, the first such threshold is $\mu = 3$. For values of μ slightly less than $\mu = 3$, Tableau 1 has lost dual-feasibility: the $-3+\mu$ in the objective row becomes negative. However, the tableau is still feasible for such values of μ. Therefore, thinking of μ as slightly less than 3, we perform primal simplex iterations. The only choice for entering variable is y. The ratios $(-1+\mu)/1$ and $(1+\mu)/1$ compete to determine the departing variable. The first of these is the smaller, and so u departs. Pivoting on the 1 in the u-row, y-column brings us to Tableau 2.

For Tableau 2 to be feasible and dual-feas., the conditions are that $-1+2\mu$, $3-\mu$, $-1+\mu$, and $-4+3\mu$ must all be nonnegative. These conditions all hold only for μ between $\frac{4}{3}$ and 3. Next threshold: $\mu = \frac{4}{3}$.

It is convenient to ignore the lower right corner until the final tableau.

	x	y	u	v	w	
y	-1	1	1	0	0	$-1+\mu$
v	-3	0	2	1	0	$-4+3\mu$
w	1	0	-1	0	1	2
	$(2\mu-1)$	0	$(3-\mu)$	0	0	

TABLEAU 2: feas. and dual-feas. if $\frac{4}{3} \leq \mu \leq 3$

Just below $\mu = 4/3$ Tableau 2 is infeasible in the v-row, but remains dual-feasible.

So the dual simplex method tells us to let v depart and x enter. After pivoting on the -3, we have reached Tableau 3.

	x	y	u	v	w	
y	0	1	$1/3$	$-1/3$	0	$1/3$
x	1	0	$-2/3$	$-1/3$	0	$\frac{4}{3}-\mu$
w	0	0	$-1/3$	$1/3$	1	$\frac{2}{3}+\mu$
	0	0	$(7+\mu)/3$	$(2\mu-1)/3$	0	

TABLEAU 3: feas. and dual-feas. if $\frac{1}{2} \leq \mu \leq \frac{4}{3}$

As μ drops below $4/3$, feasibility holds all the way to $\mu = 0$ in Tableau 3, but dual-feasibility fails if $2\mu - 1 < 0$. The next threshold is $\mu = 1/2$.

	x	y	u	v	w	
y	0	1	0	0	1	$1+\mu$
x	1	0	-1	0	1	2
v	0	0	-1	1	3	$2+3\mu$
	0	0	$2+\mu$	0	$1-2\mu$	

TABLEAU 4: feasible and dual-feasible if $0 \le \mu \le 1/2$

	x	y	u	v	w	
y	0	1	0	0	1	1
x	1	0	-1	0	1	2
v	0	0	-1	1	3	2
	0	0	2	0	1	-1

TABLEAU 5: optimal for the original LPP

Next imagine μ slightly below 1/2 and make a (primal) pivot in Tableau 3, with v entering, w departing. This leads to Tableau 4.

Tableau 4 is feasible and dual-feasible even for $\mu = 0$, so we have achieved our objective. Explicitly, set $\mu = 0$ to get Tableau 5, which is the final tableau for the original LPP. Now, compute the optimal value of z. Our optimal solution is $x = 2$, $y = 1$, $z = -1$.

(Shortcuts for hand calculations. It is not really necessary to go to the trouble of writing down all of Tableau 5 since it's implicit in Tableau 4, except for the value of z. Also, we actually could have set $\mu = 0$ in Tableau 3. The result isn't dual-feasible, but it's at least feasible, and then we can use primal simplex method.)

Exercises.

1. Solve some LPP's by the primal-dual method.

2. In the primal-dual method, assuming that $\mu \ge 0$, any feasible solution of the original LPP is also feasible for the perturbed problem. Explain why this is true. (The same is true for dual-feasibility, too.)

3. In The Example, we were always able to restore feasibility and dual-feasibility. But it is possible in principle that in the course of the primal-dual method, the simplex or dual-simplex method might lead to the unboundedness or infeasibility criterion for the perturbed problem. What would that mean about the original LPP and its dual? (**Hint.** Use the previous exercise.)

4. Another conceivable difficulty is that **both** feasibility and dual-feasibility might be lost at the **same** threshold value of μ. Once both are lost, we can't restore either one, at least by the simplex or dual-simplex methods. Can you think of a slightly different way to deform the original LPP that prevents this unpleasant possibility from ever occurring?

5. In The Example, as we crossed each threshold value of μ, it took just one iteration of simplex or dual-simplex to restore feasibility and dual-feasibility. Do you think that this always necessarily true? If so, why? If several iterations might be needed, how do you tell when to freeze the basis and start reducing μ again?