Lecture 10: The Primal Dual Schema

Vazirani Chapters 15 and 18

We just saw how to use *Linear Programming Duality* to design approximation algorithms. The technique

- Writes an integer LP describing the problem
- Relaxes the Integer LP
- Writes the *dual* of the relaxed LP
- Finds a feasible solution to dual LP
- Uses the fact that *every* dual solution is a lower bound to *optimal (minimum)* primal solution to derive bound.

The technique relied upon us being able to cleverly construct a feasible dual solution that yielded good bounds.

In this section we will see a technique that allows us to play off the dependence between primal and dual to find such feasible solutions.

COMPLEMENTARY SLACKNESS

Recall that given primal

$$\begin{array}{lll} \text{minimize} & \sum_{j=1}^{n} c_j x_j & \\ \text{subject to} & \sum_{j=1}^{n} a_{ij} x_j \geq b_i, & i = 1, \ldots, m \\ & x_j \geq 0, & j = 1, \ldots, n \end{array}$$

the dual is

$$\begin{array}{lll} \text{maximize} & \sum_{i=1}^{m} b_i y_i & \\ \text{subject to} & \sum_{i=1}^{m} a_{ij} y_i \leq c_j, & j = 1, \ldots, n \\ & y_i \geq 0, & i = 1, \ldots, m \end{array}$$

Theorem (Complementary Slackness): Let x and y respectively be primal and dual feasible solutions. Then x and y are *both* optimal if and only if all of the following conditions are satisfied.

Primal Complementary Slackness conditions

$$\forall 1 \leq j \leq n: \text{ either } x_j = 0 \text{ or } \sum_{i=1}^{m} a_{ij} y_i = c_j$$

Dual Complementary Slackness conditions

$$\forall 1 \leq i \leq m: \text{ either } y_i = 0 \text{ or } \sum_{j=1}^{n} a_{ij} x_j = b_i$$

Our main tool is

Theorem (RELAXED Complementary Slackness):
Let x and y respectively be primal and dual feasible solutions. Suppose further that for some $\alpha > 1$, x and y satisfy all of

Primal Complementary Slackness conditions

$$\forall 1 \le j \le n: \text{ either } x_j = 0 \text{ or } \sum_{i=1}^{m} a_{ij} y_i = c_j$$

RELAXED Dual Complementary Slackness conditions

$$\forall 1 \le i \le m: \text{ either } y_i = 0 \text{ or } \sum_{j=1}^{n} a_{ij} x_j \le \alpha b_i$$

Then

$$\sum_{j=1}^{n} c_j x_j \le \alpha \cdot \sum_{i=1}^{m} b_i y_i$$

Proof:

$$\sum_{j=1}^{n} c_j x_j = \sum_{j=1}^{n} \left(\sum_{i=1}^{m} a_{ij} y_i \right) x_j = \sum_{i=1}^{m} \left(\sum_{j=1}^{n} a_{ij} x_j \right) y_i \le \alpha \sum_{i=1}^{m} b_i y_i.$$

Given such an x, y we immediately know that x is within α of OPT, the minimum cost optimum solution.

The *PRIMAL DUAL SCHEMA* is to start with some with some primal infeasible dual feasible solution and to iterate.

During each iteration it improves the feasibility of the primal and the optimality of the dual (always keeping the dual solution feasible).

At the end it will produce both a feasible primal and feasible dual solution that satisfy the relaxed complimentary slackness conditions, giving an α-approximation algorithm for the primal problem.

The cost of the dual solution will then be a lower bound on the cost of OPT.

Our first example will again be set cover.
The *frequency* of element e will be
$frequency(e) = \{F \in \mathcal{F} \,:\, e \in F\}$

Suppose $f = \max_{e\in U} \text{frequency}(e)$ is the max number of sets an element can appear in. We will design a f-approximation algorithm.

Recall the set-cover LP formulations:
The integer LP will be

$$\begin{array}{|ll|}\hline \textbf{Minimize } \Sigma_{F\in\mathcal{F}}\, C(F)x_F & \\ \text{subject to conditions} & \\ \forall e \in U, & \Sigma_{e\in F}\, x_F \geq 1 \\ \forall F \in \mathcal{F} & x_F \in \{0,1\} \\ \hline \end{array}$$

The relaxation of the LP is

$$\begin{array}{|ll|}\hline \textbf{Minimize } \Sigma_{F\in\mathcal{F}}\, C(F)x_F & \\ \text{subject to conditions} & \\ \forall e \in U, & \Sigma_{e\in F}\, x_F \geq 1 \\ \forall F \in \mathcal{F}, & x_F \geq 0 \\ \hline \end{array}$$

The *dual* of the relaxed LP is then

$$\begin{array}{|ll|}\hline \textbf{Maximize } \Sigma_{e\in U}\, y_e & \\ \text{subject to conditions} & \\ \forall F \in \mathcal{F}, & \Sigma_{e\in F}\, y_e \leq C(F) \\ \forall e \in U, & y_e \geq 0 \\ \hline \end{array}$$

Primal:

Minimize $\Sigma_{F\in\mathcal{F}} C(F)x_F$
subject to conditions
$\forall e \in U, \qquad \Sigma_{e\in F}\, x_F \geq 1$
$\forall F \in \mathcal{F}, \qquad x_F \geq 0$

Dual:

Maximize $\Sigma_{e\in U}\, y_e$
subject to conditions
$\forall F \in \mathcal{F}, \qquad \Sigma_{e\in F}\, y_e \leq C(F)$
$\forall e \in U, \qquad y_e \geq 0$

Our schema will be to start with all of the $x_F = 0$, $y_e = 0$, and then iteratively change some of the x_F to 1 while also changing the y_e (but keeping y feasible). Setting $x_F = 1$ means that we put F in the cover. At the end we will have constructed feasible solutions for both the primal and dual that satisfy the relaxed complementary slackness conditions.

Primal Complementary Slackness Conditions:

$$\forall F \in \mathcal{F} : x_F \neq 0 \Rightarrow \sum_{e:e\in F} y_e = C(F)$$

Relaxed Dual Complementary Slackness Conditions:
Relaxed with $\alpha = f$.

$$\forall e : y_e \neq 0 \Rightarrow \sum_{F:e\in F} x_F \leq f \cdot 1.$$

Primal C.S: $\forall F \in \mathcal{F} : x_F \neq 0 \Rightarrow \Sigma_{e:e\in F}\, y_e = C(F)$
F is *tight* if $\Sigma_{e:e\in F}\, y_e = C(F)$.

Pick only tight sets for the cover

Relaxed Dual C.S: $\forall e : y_e \neq 0 \Rightarrow \Sigma_{F:e\in F}\, x_F \leq f$

Cover every x at most f times.

Primal-Dual Set-Cover
1. Set $\forall F$, $x_F = 0$, $\forall e$, $y_e = 0$.
2. Until all elements are covered do
Pick an uncovered element e, and increase y_e until some set becomes tight.
Add all newly tight sets to the cover. by setting $x_F = 1$ for those sets.
3. Output the cover

In the algorithm an element e is *covered* at a given step if, at that time, there is an F in the current cover s.t. $e \in F$.

<u>Theorem:</u> The algorithm generates a feasible pair x, y that satisfies the relaxed complementary slackness conditions. The algorithm is therefore a f-approximation algorithm.

<u>Proof Sketch:</u> Algorithm starts with feasible y and x that satisfies the primal complementary slackness conditions. At every step, changing y_e keeps y feasible and setting the new $x_F = 1$ keeps the the primal c.s. conditions satisfied. At the end every e is covered so the primal setting has become a feasible solution.

Minimum Multicut

Let $G = (V, E)$ be an undirected weighted graph with weights $c_e > 0$ for all edges $e \in E$.

Let $\{(s_1, t_1), (s_2, t_2), \ldots, (s_k, t_k)\}$ be k specified pairs of vertices. A *multicut* is a set of edges whose removal separates all of the pairs. The problem is to find a minimum weight multicut in G.

If we could solve this problem in polynomial time then we could also solve the *multiway cut* problem, seen before, in polynomial time. Since multiway-cut was NP-hard this problem is also NP-Hard.

In this lesson we will see how to use the Primal-Dual Schema to design a 2 approximation algorithm for the special case when G is a tree. This case can be polynomially reduced to *minimal vertex cover* (see Vazirani, Exercise 18.1) so it is also NP-Hard.

If G is a tree then, for each (s_i, t_i) there is a unique path connecting s_i to t_i. The minimum multicut removes at least one edge from this path.

For each $e \in E$ let $d_e \in \{0, 1\}$ be a variable such that

$$d_e = 1 \text{ iff } e \text{ is in the multicut.}$$

Let p_i denote the set of edges on the unique path connecting s_i and t_i.

The integer LP corresponding to minimum multicut is

Minimize $\Sigma_{e \in E}\, c_e d_e$
subject to conditions
$\forall i \in \{1, \ldots, k\}, \quad \Sigma_{e \in p_i}\, d_e \geq 1$
$\forall e \in E, \qquad d_e \in \{0, 1\}$

The relaxation of the LP is

Minimize $\Sigma_{e \in E}\, c_e d_e$
subject to conditions
$\forall i \in \{1, \ldots, k\}, \quad \Sigma_{e \in p_i}\, d_e \geq 1$
$\forall e \in E, \qquad d_e \geq 0$

We now introduce a variable f_i corresponding to (s_i, t_i). The *dual* of the relaxed LP is then

Maximize $\Sigma_{i=1}^{k}\, f_i$
subject to conditions
$\forall e \in E, \qquad \Sigma_{i: e \in p_i}\, f_i \leq c_e$
$\forall\, i \in \{1, \ldots, k\}, \quad f_i \geq 0$

Maximize $\Sigma_{i=1}^{k} f_i$ subject to conditions $\forall e \in E, \qquad \Sigma_{i:e\in p_i} f_i \le c_e$ $\forall\, i \in \{1,\ldots,k\}, \quad f_i \ge 0$

This dual can be thought of as describing the *multicommodity flow* problem. In this problem there are k different commodities with the ith commodity needing to be shipped from s_i to t_i. The object is to maximize the total amount shipped. The constraint is that the sum of the flows routed through any particular edge is at most c_e.

The *maximum integer multicommodity flow* problem is the multicommodity flow problem with the further restrictions that the f_i are all integers. Note that in this problem we may assume that the c_i are all integers as well; if they are not, we can round them down to $\lfloor c_i \rfloor$ without changing the maximum.

We will use the primal dual schema to derive an algorithm that simultaneously finds a multicut and an integer multicommodity flow that are within a factor of two of each other. This will give a 2-approximation algorithm for the minimum multicut problem *and* a 1/2 "approximation algorithm" for the maximum integer multicommodity flow one.

Primal:

Minimize $\Sigma_{e \in E} c_e d_e$
subject to conditions
$\forall i \in \{1, \ldots, k\}, \quad \Sigma_{e \in p_i} d_e \geq 1$
$\forall e \in E, \qquad d_e \geq 0$

Dual:

Maximize $\Sigma_{i=1}^{k} f_i$
subject to conditions
$\forall e \in E, \qquad \Sigma_{i: e \in p_i} f_i \leq c_e$
$\forall\, i \in \{1, \ldots, k\}, \quad f_i \geq 0$

Edge $e \in E$ will be *saturated* if total flow through e is c_e. The complimentary slackness conditions will then be:

Primal: $\forall e \in E$, $d_e \neq 0, \Rightarrow \Sigma_{i: e \in p_i} f_i = c_e$.
This means *any edge picked in the multicut must be saturated*

Relaxed Dual: $\forall i$, $f_i \neq 0 \Rightarrow \Sigma_{e \in p_i} d_e \leq 2$.
This means *at most two edges can be picked from a path carrying non-zero flow*

We start by rooting the tree G at an arbitrary vertex
Depth of $v \in V$ will be length of path from v to the root.
For $u, v \in V$ let $lca(u, v)$ be the *lowest common ancestor* of u and v.
Let e_1, e_2 be two edges on the same path from a vertex to the root. If e_1 occurs before e_2 on this path, e_1 is *deeper* than e_2.

Algorithm starts with an empty multicut (satisfies primal c.s. conditions) and empty flow (feasible). It then iteratively improves feasibility of primal solution and optimality of dual solution. The edges in the multicut so far will be kept in a list D.

During an iteration it picks the deepest unprocessed vertex so far, v and greedily routes integral flow between pairs that have v as their lca. When no more flow can be routed between these pairs at least one edge has been saturated. All saturated edges are added to D.

After all vertices have been processed D will be a multicut. D might contain extra edges, though. The algorithm ends by working through the edges of D in reverse order in which they were added; if after the removal of an edge D remains a multicut, the edge is removed from D.

Minimum Multicut/ Integer Multicommodity Flow (Trees)
1. $\forall i \; f_i = 0, \quad D = \emptyset$.
2. For all vertices v, in non-increasing order of depth, do:
   For every pair (s_i, t_i) with $lca(s_i, t_i) = v$
      Greedily route integral flow from s_i to t_i
   Add to D all edges e that were saturated in this iteration.
3. Let $e_1, e_2, \ldots, e_l$ be edges in D ordered by insertion time.
4. For $j = l$ downto 1 do.
   If $D - \{e_j\}$ is a multicut in G remove e_j from D.
5. Output flow and multicut D.

First note that the algorithm outputs a legal flow since it starts with the empty (legal) flow and at every step maintains legality.

Next note that at the end of the algorithm D must contain a multicut since at least one edge on the unique (s_i, t_i) path must have been added to D (WHY).
Now construct the integral 0/1 solution $d_e = 1$ iff $e \in D$.

Constructed primal and dual solutions are therefore both feasible. Furthermore, since every edge picked in the multicut was saturated, the solutions satisfy primal c.s. conditions.
$\forall e \in E, \; d_e \neq 0, \Rightarrow \; \Sigma_{i:e \in p_i} f_i = c_e$.

Minimum Multicut/ Integer Multicommodity Flow (Trees)
1. $\forall i \; f_i = 0, \quad D = \emptyset$.
2. For all vertices v, in non-increasing order of depth, do:
   For every pair (s_i, t_i) with $lca(s_i, t_i) = v$
      Greedily route integral flow from s_i to t_i
   Add to D all edges e that were saturated in this iteration.
3. Let $e_1, e_2, \ldots, e_l$ be edges in D ordered by insertion time.
4. For $j = l$ downto 1 do.
   If $D - \{e_j\}$ is a multicut in G remove e_j from D.
5. Output flow and multicut D.

Constructed primal and dual solutions are both feasible and satisfy the primal c.s. conditions.

Lemma: Let (s_i, t_i) be a pair with non-zero flow and let $lca(s_i, t_i) = v$. Then at most one edge is picked in the multicut from each of the two paths s_i to v and v to t_i.

This lemma implies that, for each (s_i, t_i), at most two edges from D are on the path connecting s_i to t_i. Thus the relaxed dual conditions

$$\forall i, \; f_i \neq 0 \Rightarrow \sum_{e: e \in p_i} d_e \leq 2$$

are all satisfied as well.

Combining everything: since d_e and f_i are feasible solutions that satisfy the relaxed c.s. conditions with $\alpha = 2$ we have that

$$\sum_i f_i \leq \sum_e c_e d_e \leq 2 \cdot \sum_i f_i$$

and D has weight within a factor of two optimal solution of a minimum multicut.

This can also be read as

$$\frac{1}{2} \sum_e c_e d_e \leq \sum_i f_i \leq \sum_e c_e d_e$$

implying that the solution is also a $\frac{1}{2}$ "approximation" algorithm for integral multicommodity flow as well.

Lemma: Let (s_i, t_i) be a pair with non-zero flow and let $lca(s_i, t_i) = v$. then at most one edge is picked in the multicut from each of the two paths s_i to v and v to t_i.

Proof: We prove for the s_i to v path. Proof for v to t_i path is the same.

Suppose at the end of the algorithm there are two edges $e, e' \in D$ from the same s_i to v path with e being deeper. By definition e' must remain in D all through step 4.

Consider the moment in step 4 when e is being tested. e is not thrown away so there must be a pair (s_j, t_j) such that e is the only edge on the path between s_j and t_j. Let $u = lca(s_j, t_j)$. Since e' is not on path between s_j and t_j (why) u must be deeper than e' and therefore deeper than v.

This in turn implies that after u was processed D must have contained an edge from the path between s_j and t_j. Call this edge e''.

Now, since non-zero flow was routed from s_i to t_i, e *must* have been added to D during or after the iteration that processed v. Since v is an ancestor of u, e is added after e''. But then e'' must be in D when e is being tested, contradicting fact that at this time e is only edge of D on the path between s_j and t_j.