LP-Duality
("Approximation Algorithms" by V. Vazirani, Chapter 12)

- Well-characterized problems, min-max relations, approximate certificates
- LP problems in the standard form, primal and dual linear programs
- LP-duality theorem, complementary slackness conditions
- Min-max relations and LP-duality: max-flow min-cut example
- LP-based approximate algorithms design: LP-rounding and Primal-dual schema
- Integrality gap
- Example: Maximum Matching and Minimum Vertex Cover: LP-duality point of view

Considering the decision version of an optimization problem:

- Does it have Yes certificates? (Is it in NP?)
- Does it have No certificates? (Is it in co-NP?)

Problems that have both Yes and No certificates are said to be *well-characterized*.

Well-characterized problems are in NP $\cap$ co-NP, which gives a hope of finding polynomial time algorithms for those.

Min-max relations often help prove that a problem is well-characterized. [And many min-max relations are actually special cases of the LP-duality.]

<u>Example</u>

König-Egervary theorem:

In any bipartite graph, $\max_{\text{matching } M} |M| = \min_{\text{v.cover } U} |U|$

gives us No certificates for the questions

- If G is bipartite, is $\min_{\text{v.cover } U} |U| \leq k$?
- If G is bipartite, is $\max_{\text{matching } M} |M| \geq k$?

In general, $\max_{\text{matching } M} |M| = \min_{\text{v.cover } U} |U|$ is not true. [Example: the Petersen graph.]

Moreover, the Minimum vertex cover is NP-hard, and under the NP $\neq$ co-NP assumption, it doesn't have No certificates. However, it does have *approximate* No certificates, that is, certificates for sufficiently small values of k.

Since $\max_{\text{matching } M} |M| \leq \min_{\text{v.cover } U} |U| \leq 2 \max_{\text{matching } M} |M|$,

if $k < \min_{\text{v.cover } U} |U| / 2$, we have No certificates for the "is $\min_{\text{v.cover } U} |U| \leq k$" question. We say we have factor 2 No certificates for the Minimum vertex cover.

In this case, our approximate No certificates are polynomial time computable. For the shortest lattice vector problem, we have factor n approximate No certificates, but we don't know how to compute those in polynomial time.

Primal and dual linear programs, the LP-duality theorem

Minimization and maximization linear programs in the standard form:

minimize $\quad \sum^{n}_{j=1} c_j x_j$

subject to $\quad \sum^{n}_{j=1} a_{ij} x_j \geq b_i, \quad i = 1, \ldots, m$

$\quad x_j \geq 0, \qquad\qquad j = 1, \ldots, n$

maximize $\quad \sum^{n}_{j=1} c_j x_j$

subject to $\quad \sum^{n}_{j=1} a_{ij} x_j \leq b_i, \quad i = 1, \ldots, m$

$\quad x_j \geq 0, \qquad\qquad j = 1, \ldots, n$

The function being optimized is called the *objective function*. Any solution that satisfies all the constraints is called a *feasible solution*.

LP problems are well-characterized. Existence of Yes certificates is obvious. LP-duality provides us with No certificates.

Let's start with a minimization LP problem. We call it the *primal program*.

To get a No certificate for the question "is the minimum of the objective function $\leq a$?" we need to lower bound the objective function. One way of doing that is via carefully constructed linear combination of the constraints. That is, we introduce new variables $y_i \geq 0$ and use them as the linear combination coefficients:

$y_1 \sum^n_{j=1} a_{1j}\, x_j + \ldots + y_m \sum^n_{j=1} a_{mj}\, x_j \;\geq\; \sum^m_{i=1} y_i\, b_i$, as the coefficients are non-negative.

What do we need to take care of here? That the objective function is no less than the right-hand side above. Since $x_j \geq 0$, we need the following conditions:

$\sum^m_{i=1} a_{ij}\, y_i \;\leq\; c_j,\; j = 1, \ldots, n$

And to get the best possible lower bound, we want $\sum^m_{i=1} y_i\, b_i$ to be as large as possible.
The result is an LP maximization problem, and we call it the *dual program*.

A straightforward argument shows that dual to the dual program is the primal program.

Obviously, any feasible solution to the dual program gives a lower bound on the primal program's objective function. Conversely, any feasible solution to the primal program gives an upper bound on the dual program's objective function.
That is, for any two feasible solutions $\{x_j\}$ and $\{y_i\}$ we have

$$\sum^{n}_{j=1} c_j x_j \geq \sum^{m}_{i=1} y_i b_i ,$$

and this is the **Weak LP-duality theorem**, which in many cases suffices for constructing approximation algorithms. In fact, a stronger min-max claim holds.

LP-duality theorem

The primal program has finite optimum iff the dual has finite optimum. For any two optimum solutions $\{x_j\}$ and $\{y_i\}$, the values of the primal and dual objective functions are equal.

Thus, the LP problem is well-characterized.

From the obvious chain of inequalities for two feasible solutions $\{x_j\}$ and $\{y_i\}$

$$\sum^{n}_{j=1} c_j\, x_j \;\geq\; \sum^{n}_{j=1} \left(\sum^{m}_{i=1} a_{ij}\, y_i\right) x_j \;=\; \sum^{m}_{i=1} \left(\sum^{n}_{j=1} a_{ij}\, x_j\right) y_i \;\geq\; \sum^{m}_{i=1} y_i\, b_i\,,$$

we see that the solutions are both optimum iff the two inequalities are, in fact, equalities.

This gives us the **Complementary slackness conditions** of optimality:

for each $1 \leq j \leq n$: either $x_j = 0$ or $\sum^{m}_{i=1} a_{ij}\, y_i = c_j$

and

for each $1 \leq i \leq m$: either $y_i = 0$ or $\sum^{n}_{j=1} a_{ij}\, x_j = b_i$

LP-duality and max-flow min-cut theorem

We'll now consider an example where LP-duality brings an exact min-max relation. Here's our setting:

$G = (V, E)$ is a directed graph with two distinguished nodes, source s and sink t, and positive edge capacities $c: E \to \mathbf{R}^+$.

Flow $f: E \to \mathbf{R}^+$ must satisfy the following conditions:

- for any edge e, $f(e) \le c(e)$ [capacity constraints]
- for any node v, other than s and t, the total flow into v equals the total flow out of v [flow conservation]

s-t cut is defined by a partition of V into two sets X and X^c, so that s is in X and t is in X^c, and consists of all the edges going from X to X^c. The capacity of the cut, $c(X, X^c)$, is the sum of capacities of all its edges.

Max-flow problem:
Find a flow with the maximum "out-value" (= total-out – total-in) at s.

Min-cut problem:
Find an s-t cut with the minimum capacity.

Capacity of any s-t cut is an upper bound on the value of any feasible flow. Using LP-duality, we can show that it's possible to find a flow and an s-t cut with the equal values, which are actually the optimum values.

To formulate the max-flow problem as a linear program, we introduce a fictitious edge of infinite capacity from t to s, which lets us require flow conservation at s and t. Our primal problem is then:

maximize f_{ts}

subject to $f_{ij} \le c_{ij}$, for any edge (i, j) from E

$\sum_{j:\,(j,\,i)} f_{ji} - \sum_{j:\,(i,\,j)} f_{ij} \le 0$, for any node i from V

$f_{ij} \ge 0$, for any edge (i, j) from E

[Since the second type constraints must hold at each node and their total sum is 0, they actually must be satisfied with equality at each node, thus, they're equivalent to the flow conservation conditions.]

We'll now construct the dual program for our max-flow LP. Because we have two types of the constraints in the primal, we'll introduce two sets of the variables $\{d_{ij}\}_{(i,\, j) \text{ from } E}$ and $\{p_i\}_{i \text{ from } V}$ for the dual. Then we get:

minimize $\sum_{(i,\, j)} c_{ij}\, d_{ij}$

subject to $d_{ij} - p_i + p_j \geq 0$, for any edge $(i, j) \neq (t, s)$ from E

$p_s - p_t \geq 1$

$d_{ij} \geq 0$, for any edge (i, j) from E

$p_i \geq 0$, for any node i from V

The LP-duality theorem tells that the optimum for the primal program (the max flow) is equal to the optimum for the dual.
Now, what if we require all the variables from the dual program to be 0/1 variables?

Claim A: The restriction to 0/1 turns the dual program into the minimum s-t cut problem.

Claim B: The optimum to the dual program is attained on 0/1 values of the variables.

Taken together, these bring us the max-flow min-cut theorem.

Proving Claim A is easy. Claim B requires some work.

The first step is to notice that the optimum for the unconstrained version is the same as for the version when the variables are restricted to [0, 1] interval.

Then, we can show that the optimum for the version with the variables restricted to [0, 1] interval is actually attained at an integral point. [The dual convex polyhedron has integral vertices.]

We call the dual program the *LP-relaxation* of the min s-t cut integer program. Feasible solutions to the dual program are called *fractional solutions* to the integer program (fractional s-t cuts for the min s-t cut integer program).

Finally, we can derive certain properties of the optimum solutions to our max-flow and min-cut problems using the complementary slackness conditions:

- edges from X to X^c in the optimum cut must be saturated by the optimum flow
- edges from X^c to X must carry no flow.

LP-relaxation and LP-duality in algorithm design

Many combinatorial optimizations problems can be stated as integer programs. LP-relaxations of those programs provide a natural way of bounding the optimum solution cost, which is often a key step in the design of exact (when relaxations are exact) or approximation algorithms.

Two basic techniques are *LP-rounding* and *primal-dual schema*.

LP-rounding: solve the LP-relaxation and convert the fractional solution obtained into an integral solution in such a way that the cost doesn't change much. The approximation guarantee is established by comparing the cost of the integral and fractional solutions.

Primal-dual schema: an integral solution to the primal program and a feasible solution to the dual program are constricted iteratively. The approximation guarantee is established by comparing the cost of the two solutions.

Method of dual fitting is used in analyzing combinatorially obtained approximation algorithms. Will see the applications in the later chapters.

Integrality gap, approximation guarantee, and rounding factor

For a minimization problem Π and its LP-relaxation, let LP-OPT(I) denote the cost of an optimal fractional solution to instance I, and OPT(I) denote the cost of an optimal integral solution. We define the integrality gap of the relaxation to be

$$\sup_I \ \mathrm{OPT}(I) \,/\, \text{LP-OPT}(I)$$

[For a maximization problem, the integrality gap is the infimum of the above ratio.]

If the LP-relaxation has an integral optimal solution, then the integrality gap is 1, and we call the relaxation *exact*.

Given an optimal fractional solution to instance I, we can turn it into an integral one, with the cost $A(I)$, by rounding. Then the *rounding factor* for instance I is $A(I)$ / LP-OPT(I). Thus,

rounding factor of I = (approximation factor of I) x (integrality gap of I)

So, (uniformly) minimizing the rounding factor, we can improve the approximation guarantee. Of course, the rounding factor can never be less than the integrality gap.

Maximum Matching and Minimum Vertex Cover: LP-duality point of view

These two problems present us an example of duality and non-exact relaxation.

Maximum matching (no edge weights) can be stated in this way:

maximize $\sum_{(i,j)} x_{ij}$

subject to $\sum_{(i,j)} x_{ij} \leq 1$, for any node i from V

x_{ij} is in $\{0, 1\}$, for any edge (i, j) from E

Minimum vertex cover (no vertex weights) can be stated in this way:

minimize $\sum_i y_i$

subject to $y_i + y_j \geq 1$, for any edge (i, j) from E

y_i is in $\{0, 1\}$, for any vertex i from V

Their LP-relaxations are dual. The same applies to the "weighted" versions, if we choose the vertex and edge weights in a consistent way.

The integrality gap of LP-relaxation for Minimum Vertex Cover

Taking K_n, we see that the integrality gap can approach 2.

We can show that the gap does not exceed 2.

Given a feasible fractional solution $\{y_i\}$, we choose this rounding:

$z_i = 1$ if $y_i \geq 0.5$

$z_i = 0$ if $y_i < 0.5$

This clearly gives us a valid solution.

Now, the integrality gap $\leq$ sup $(\sum_i z_i \ / \ \sum_i y_i)$

And it's easy to see that $\sum_i z_i \leq 2 \sum_i y_i$.