Recitation Notes: LP Duality and Zero-Sum Games

1 Lecture 16: Linear Programming Duality

- Standard primal/dual pair:

$$\min\{c^\top x : Ax = b,\ x \geq 0\} \qquad \longleftrightarrow \qquad \max\{b^\top y : A^\top y \leq c\}.$$

- Weak duality: for any primal-feasible x and dual-feasible y,

$$b^\top y \leq c^\top x.$$

  Equality certifies optimality.
- Primal constraints correspond to dual variables, and primal variables correspond to dual constraints.
- Many problems fit into LP form:
  - shortest path as a flow LP,
  - max independent set as an LP relaxation,
  - max with $\leq$ constraints can be converted using slack variables.
- Max-flow/min-cut comes from LP duality:
  - max-flow can be written as an LP,
  - its dual looks like a cut problem,
  - this gives another proof idea for max-flow $=$ min-cut.

2 Lecture 17: Zero-Sum Games

- A zero-sum game has payoff function $p(a,b)$: player 1 wants to maximize it, player 2 wants to minimize it.
- Mixed strategies are distributions over pure strategies, and the expected payoff is

$$x^\top P y.$$

- Main theorem about zero-sum games is that mixed strategies are optimal.

Theorem 1: Pure Response lemma / Minimax Theorem

To evaluate the minimum expected payoff of a mixed strategy for player 1, it suffices to minimize only over pure strategies for player 2.

This holds because once player 1 fixes a mixed strategy x, the other player has no reason to randomize, as their payoffs would be an average of the payoffs. They can instead pick the column with the best payoff every time.

- The minimax theorem follows from LP duality:

$$\max_{x\in\Delta_A}\min_{b\in B}\sum_a x(a)p(a,b) = \min_{y\in\Delta_B}\max_{a\in A}\sum_b y(b)p(a,b).$$

- Example: in rock-paper-scissors, both players play uniformly, and the game value is 0.

3 Zero-Sum Game Example

Problem: Matrix game

Let

$$A = \{a_{i,j}\}_{i,j=1}^{n}$$

be an $n \times n$ diagonal matrix with positive diagonal entries

$$d_1, d_2, \ldots, d_n$$

and all off-diagonal entries equal to 0. In other words,

$$A = \begin{pmatrix} d_1 & 0 & \cdots & 0 \\ 0 & d_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & d_n \end{pmatrix}.$$

Consider the zero-sum game in which the row player receives payoff $A_{i,j}$ when row i and column j are chosen. Determine the value V of the game as a function of the nonzero entries $d_1, \ldots, d_n$.

Use a guess-and-bound approach:

1. Analyze the row player's uniform mixed strategy, assigning probability $1/n$ to each row. What is the best response of the column player, and what lower bound does this imply for the game value?
2. Give a stronger mixed strategy for the row player and prove the corresponding lower bound V.
3. Give a matching strategy for the column player and prove the corresponding upper bound V, thereby identifying the value of the game.

Solution For symmetry we may assume that $d_1 \geq d_2 \geq \cdots \geq d_n$.

1. The column player would just plays strategy n since it has the lowest payoff and the column player is minimizing payoff. This implies a lower bound of d_n/n.

2. Say the row player players $(x_1, \ldots, x_n)$ where $x_1 + \cdots + x_n = 1$ and $x_i \geq 0$. The value of the game would then be $\min_{i \in [n]} x_i d_i$.

   Let
   $$\lambda = \left(\sum_{i=1}^{n} \frac{1}{d_i} \right)^{-1}.$$
   The strategy $x_i = \lambda / d_i$ is valid because the d_is, and in turn λ, are non-negative, and
   $$\sum_{i=1}^{n} x_i = \sum_i \frac{1}{d_i} \cdot \lambda = \lambda \cdot \left(\sum_i \frac{1}{d_i} \right) = 1.$$
   The game value of playing it is then λ since each $d_i x_i = d_i(\lambda / d_i) = \lambda$.

3. The column player plays the same strategy. The game value is the same.