Optimization **Summer 2010**

Lecture 8 — May 5

Lecturer: Julián Mestre

8.1 Two-player zero-sum games

We consider the following mathematical abstraction of a game played by two players. Each player has a set of possible *strategies* that she can choose to play. We denote by $\mathcal{S}_i$ the set of strategies of player $i = 1, 2$. The outcome of the game is determined by the *pay-off* matrix $\mathbf{D} \in \mathcal{R}^{|S_1| \times |S_2|}$. Suppose that the players select $s_1 \in \mathcal{S}_1$ and $s_2 \in \mathcal{S}_2$ respectively. Then if $D(s_1, s_2) > 0$ we can think of player 1 getting $D(s_1, s_2)$ euros from player 2, and if $D(s_1, s_2) < 0$ we can think of player 1 paying $|D(s_1, s_2)|$ euros to player 2. This type of games are call zero-sum because the sum of the earnings of player 1 and player 2 is always zero.

For a concrete example, consider the children game of rock-paper-scissors. The strategy of the each player is the same, {ROCK, PAPER, SCISSOR}, and the pay-off matrix is as follows:

	ROCK	PAPER	SCISSORS
ROCK	0	-1	1
PAPER	1	0	-1
SCISSORS	-1	1	0

Clearly if a player has to announce her strategy before the other payer, there is no way she can win. An interesting question to consider is whether she could do better by using a randomized strategy? Let us try to make this question more precise. Suppose we are given a pay-off matrix $\mathbf{D} \in \mathcal{R}^{n \times m}$. A mixed strategy for player 1 is a vector $\mathbf{x} \in \mathcal{R}^n$ such that $x_i \geq 0$ for all i and $\sum_i x_i = 1$, where x_i represents the probability that payer 1 choose her ith strategy. Similarly, a mixed strategy for player 2 is a vector $\mathbf{y} \in \mathcal{R}^m$ such that $y_j \geq 0$ and $\sum_j y_j = 1$. To differentiate these from regular strategies, we call the latter *pure strategies*. For fixed mixed strategies $\mathbf{x}$ and $\mathbf{y}$ the expected pay-off is

$$\text{expected pay-off of } \mathbf{x} \text{ and } \mathbf{y} = \sum_{i,j} x_i \, D_{ij} \, y_j = \mathbf{x}' \, \mathbf{D} \, \mathbf{y}.$$

Therefore, if player 1 reveals her strategy before player 2, she should try to solve the following problem:

$$\max_{\mathbf{x}} \min_{\mathbf{y}} \; \mathbf{x}' \, \mathbf{A} \, \mathbf{y},$$

whereas if player 2 reveals her strategy before player 2, she should try to solve the following problem;

$$\min_{\mathbf{y}} \max_{\mathbf{x}} \; \mathbf{x}' \, \mathbf{A} \, \mathbf{y}.$$

Clearly the former should always be smaller than the latter. In fact, this relation also holds for pure strategies; notice, however, that the relation can be strict as the rock-paper-scissors game shows. Surprisingly, in the case of mixed strategies order does not matter. This is called von Neumann's Minimax Theorem and it states that for mixed strategies in both cases we get the same expected pay-off.

Theorem 8.1. *Let* $\mathbf{D}$ *be a pay-off matrix of a two-player zero-sum game. Then*

$$\max_{\mathbf{x}} \min_{\mathbf{y}} \ \mathbf{x}' \, \mathbf{A} \, \mathbf{y} = \min_{\mathbf{y}} \max_{\mathbf{x}} \ \mathbf{x}' \, \mathbf{A} \, \mathbf{y},$$

where the minima and maxima range over all possible mixed strategies of the players.

8.2 Yao's minimax principle

An unexpected application of zero-sum games is that it allows us to prove lower bounds on the computational complexity of randomized algorithms for some problems.

Let $\mathcal{I}_n$ be the set of instances of size n of some computational problem and let $\mathcal{A}$ be the set of deterministic algorithms for the problem. We restrict our attention to correct algorithms, that is, each $A \in \mathcal{A}$ returns the correct answer on all $I \in \mathcal{I}$. We are interested in the complexity of this problem, where the complexity measure can be time, space, or any other resource. Let $D(I, A)$ be the complexity measure of algorithm $A \in \mathcal{A}$ on instance $I \in \mathcal{I}$. The worst-case complexity of algorithm A is

$$\max_{I \in \mathcal{I}} D(I, A).$$

We can interpret a randomized algorithm as a probability distribution over deterministic algorithms. Let $\mathbf{y}$ be the probability distribution induced by some randomized algorithm. The worst-case expected complexity of $\mathbf{y}$ is

$$\max_{I \in \mathcal{I}} \ \mathbb{E}_{A \sim \mathbf{y}}[D(I, A)].$$

Yao's principle states that for a given computational problem, the worst-case expected complexity of any randomized algorithm cannot be better than expected cost of any deterministic algorithm over probability distribution over instances.

Theorem 8.2. *Let* $\mathbf{y}$ *be the probability distribution induced by some randomized algorithm and let* $\mathbf{x}$ *be some probability distribution over instances, then*

$$\min_{\mathbf{y}} \max_{I \in \mathcal{I}} \ \mathbb{E}_{A \sim \mathbf{y}}[D(I, A)] \geq \max_{\mathbf{x}} \min_{A \in \mathcal{A}} \ \mathbb{E}_{I \sim \mathbf{x}}[D(I, A)].$$