HW2 Randomized algorithms & Random structures

MPRI 1.24

Wed. Dec. 2, 2015 - **Due on Wed. Dec. 9, 2015**

You are asked to complete the exercise marked with a [★] and to send me your solutions at:

nicolas.schabanel@cnrs.fr

(or drop it in my mail box at the 4th floor of Sophie Germain) on **Wed. Dec. 9, 2015**.

■ **Exercise 1 (Yao's principle).** We are running an hard-drive and we want to minimize its energy consumption. When the hard drive is on, it consumes x per unit of time; when it is off, it cost y to turn back on. Request to access to the hard drive arrives at unknown dates $t_1 < t_2 < \cdots$ and we want to minimize the overall cost.

► ***Question 1.1)*** *Show that we can restrict the study to minimize the energy consumption during each time interval $[t_i, t_{i+1}]$ independently and thus restrict ourselves to the following situation: at time $t = 0$, the drive is on and we want to serve a request arriving at an unknown beforehand date t for a minimum cost.*

► ***Question 1.2)*** *Show that the only deterministic algorithms to consider are A_d, for $d \geqslant 0$: if the request arrives before $t = d$, serves it; otherwise, turn off the drive at time d and turn it back on when the request arrives.*

We denote by I_t the instance where the request arrive at time t and by $\mathrm{OPT}(I_t)$ the optimal energy consumption to serve a request arriving at time t.

► ***Question 1.3)*** *Show that:*

$$\mathrm{cost}(A_d(I_t)) = \begin{cases} xt & \text{if } t \leqslant d \\ xd + y & \text{if } t > d \end{cases} \quad \text{and } \mathrm{OPT}(I_t) = \min(xt, y).$$

We define the *competitive ratio* of A_d as: $C_d = \max_{t \geqslant 0} \frac{\mathrm{cost}(A_d(I_t))}{OPT(I_t)}$.

► ***Question 1.4)*** *Show that the optimal deterministic competitive ratio is $C_{y/x} = 2$.*

We turn now to randomized algorithms. Any randomized algorithm can be seen as a distribution $p : \mathbb{R} \to \mathbb{R}$ over deterministic algorithms: A_p is the algorithm that runs A_d "with probability $p(\delta)$" (more precisely, A_p runs algorithm A_δ with $\delta \geqslant d$ with probability $\int_d^\infty p(\delta) d\delta$). Then, $\mathbb{E}[\mathrm{cost}(A_p(I_t))] = \int_0^\infty \mathrm{cost}(A_\delta(I_t)) p(\delta) d\delta$ and the randomized competitive ratio is $C_p = \max_{t \geqslant 0} \frac{\mathbb{E}_p[\mathrm{cost}(A_p(I_t))]}{\mathrm{OPT}(I_t)}$. Our goal is to find an optimal distribution $p^* = \arg\min_{p:\mathbb{R}_+ \to \mathbb{R}_+ : \int_0^\infty p(\delta) d\delta = 1} \max_{t \geqslant 0} \frac{\mathbb{E}_p[\mathrm{cost}(A_p(I_t))]}{\min(xt, y)}$. Note that we assume that the adversary chossing the instance t does not known our random choice, otherwise we would be back to the deterministic case.

Zero-sum games. *Zero-sum games* are defined by a *cost matrix* $C = (c_{d,t})_{d=1..D,t=1..T}$. At each round, the line player chooses a line d and the column player chooses (independently) a column t and the line player pays $c_{d,t}$ to the column player (each player earns what the other looses). Now, an optimal deterministic strategy for the line player is a line $d^* \in \arg\min_d \max_t c_{d,t}$ whereas an optimal deterministic strategy for the column player is a column $t^* \in \arg\max_t \min_d c_{d,t}$. It turns out that optimal deterministic strategies do not always exist.

Let us focus on randomized strategies: assume that the line player chooses a line $d \in \{1, \ldots, D\}$ with probability p_d and that the column player chooses a column $t \in \{1, \ldots, R\}$ with probability q_t. Let us see p and q as line and column vectors of dimension D and T respectively.

► ***Question* 1.5*)*** *Show that the expected cost paid by the line player to column player is:* $p\,Cq = \sum_{d=1}^{D} \sum_{t=1}^{T} p_d\, c_{d,t}\, q_t$.

An optimal randomized strategy p^* for the line player must minimize the worst cost to the worst possible distribution of the column player:

$$p^* = \operatorname*{arg\,min}_{p \geqslant 0: ||p||_1 = 1} \max_{q \geqslant 0: ||q||_1 = 1} pCq$$

whereas an optimal randomized strategy q^* for the column player must maximize the cost to the best possible distribution of the line player:

$$q^* = \operatorname*{arg\,max}_{q \geqslant 0: ||q||_1 = 1} \min_{p \geqslant 0: ||p||_1 = 1} pCq$$

Von Neumann showed that there always exists an optimal randomized strategy for both players, and the cost of these two strategies always match! (the minimax theorem is in fact equivalent to the duality theorem in linear programming):

$$\min_{p \geqslant 0: ||p||_1 = 1} \max_{q \geqslant 0: ||q||_1 = 1} pCq = \max_{q \geqslant 0: ||q||_1 = 1} \min_{p \geqslant 0: ||p||_1 = 1} pCq$$

► ***Question* 1.6*)*** *Show the easy inequalities:*

$$\min_p \max_q pCq = \min_p \max_t p\, C_t \geqslant \max_q \min_d C_d\, q = \max_q \min_p p\, Cq$$

where C_d and C_t denote respectively the d-th row and the t-th column of the cost matrix.

Yao's principle. Yao remarked that one can see the interaction between a randomized algorithm and an adversary as a *zero-sum game* where at each round, the algorithm chooses a time d and the adversary chooses (simultaneously) a time t and the algorithm pays $c_{d,t} = \frac{\mathrm{cost}(A_d(I_t))}{\mathrm{OPT}(I_t)}$ to the adversary. It follows from von Neumann's minimax extended to our continuous setting:

$$\min_{p \in \mathfrak{D}} \max_{t \geqslant 0} \mathbb{E}_p \left[\frac{\mathrm{cost}(A_p(I_t))}{\mathrm{OPT}(I_t)} \right] = \max_{q \in \mathfrak{D}} \min_{d \geqslant 0} \mathbb{E}_q \left[\frac{\mathrm{cost}(A_d(I_q))}{\mathrm{OPT}(I_q)} \right]$$

where $\mathfrak{D} = \{f : \mathbb{R}_+ \to \mathbb{R}_+ : \int_0^\infty f(x)dx = 1\}$. That is to say *"the expected cost of the best randomized algorithm on the worst instance is equal to the expected cost of the best deterministic algorithm for the worst distribution of instances"*. The weak version proved in Question 1.6 says that *"the expected cost of any randomized algorithm on the worst instance is always at least the expected cost of the best deterministic algorithm for the worst distribution of instances"*.

We can now use this principle to prove that some strategies are optimal: if we can find two probability distribution p^* and q^* for d and t respectively for which the worst competitive ratio of A_{p^*} is equal to the best competitive ratio of a deterministic algorithm for distribution q^*, then A_{p^*} is an optimal randomized algorithm.

Rescale so that $x = y = 1$ and consider the two distributions:

$$p^*(d) = \begin{cases} \frac{e^d}{e-1} & \text{if } 0 \leqslant d \leqslant 1 \\ 0 & \text{otherwise} \end{cases}$$

and

$$q^*(t) = \begin{cases} \frac{e}{e-1} t e^{-t} & \text{if } 0 \leqslant t \leqslant 1 \\ 0 & \text{if } t > 1 \text{ and } t \neq 2 \end{cases} \quad \text{with a probability mass } \Pr\{ t = 2 \} = \frac{1}{e-1}.$$

► ***Question* 1.7*)*** *Show that the expected competitive ratio of the randomized algorithm A_{p^*} is independent of t and that the expected competitive ratio for distribution of instance q^* is independent of the deterministic algorithm and conclude that A_{p^*} is optimal. What is the best expected competitive ratio achievable by a randomized algorithm?*

It follows that $\max_t C_{p^*,t} = \frac{e}{e-1} = \min_d C_{d,q^*}$ which implies by Yao's principle that $\frac{e}{e-1}$ is the best expected competitive ratio achievable by a randomized algorithm and it is obtained by the randomized algorithm A_{p^*} which chooses the maximum waiting time d according to law p^*.