Homework 3 for Columbia B9136

Will Ma

due Feb 17, 11:59pm

Exercise 1. Recall the LP for k-unit OCRS that we saw in class (assuming arrival order and activeness probabilities are known in advance):

$$
\begin{aligned}
\max \quad & c && && \text{(0.1a)}\\
\text{s.t.} \quad & cx_t \le \sum_{\ell=1}^{k} \alpha_t^\ell && \forall t \in [T] && \text{(0.1b)}\\
& \alpha_t^\ell \le x_t \beta_t^\ell && \forall t \in [T], \ell \in [k] && \text{(0.1c)}\\
& \beta_t^\ell = \begin{cases} 1 & t = 1, \ell = k \\ 0 & t = 1, \ell < k \\ \beta_{t-1}^\ell - \alpha_{t-1}^\ell + \alpha_{t-1}^{\ell+1} & t > 1 \end{cases} && \forall t \in [T], \ell \in [k] && \text{(0.1d)}\\
& \alpha_t^\ell \ge 0 && \forall t \in [T], \ell \in [k] && \text{(0.1e)}
\end{aligned}
$$

Now consider the accept/reject problem, starting with k units where agents arrive in order $t = 1, \dots, T$, each with a valuation V_t drawn independently from a *weighted Bernoulli* distribution where $V_t = r_t$ w.p. x_t, and $V_t = 0$ w.p. $1 - x_t$. An adversary sets the values $(r_t)_{t=1}^T \in \mathbb{R}_{\ge 0}^T$ to *minimize* the expected reward of the optimal dynamic programming policy, subject to the constraint that $\sum_{t=1}^T r_t x_t = 1$. (In other words, the adversary normalizes $\textsc{Flu}(\mathbf{F})$ to be 1.)

Show that the optimal objective value for the adversary's problem equals that of LP (0.1), under any values of k, T, and $(x_t)_{t=1}^T \in [0,1]^T$.

Hint: LP duality is your friend. To formulate the adversary's problem as an LP, use a variable J_t^ℓ to denote the value-to-go when agent t arrives with ℓ units remaining, for all $t \in [T]$ and $\ell \in [k]$. By Bellman's equations from dynamic programming, we have $J_t^\ell - J_{t+1}^\ell = x_t[r_t - J_{t+1}^\ell + J_{t+1}^{\ell-1}]^+$. One way to formulate this in the LP is to introduce an auxiliary variable K_t^ℓ with constraints:

$$
\begin{aligned}
K_t^\ell &\ge r_t - J_{t+1}^\ell + J_{t+1}^{\ell-1}\\
K_t^\ell &\ge 0.
\end{aligned}
$$

To get the final result, you may have to manipulate the LP before/after taking the dual.

Exercise 2. We saw in class the example where $k = 2$, $T = 3$, and agents have $(x_1, x_2, x_3) = (1, 1/2, 1/2)$ (and arrive in that order). It can be verified that the instance-optimal value of c is $3/4$.

Consider a modification to this example where now there is a 4th agent arriving at the end, who can only be accepted if *both* units still remain. Show that the optimal value of c on this modified example is $4/9$.

Hint: The optimal value of c does not depend on what x_4 is. You want to design an online policy such that the probability of having both units remaining for agent 4 is exactly $4/9$.

Exercise 3. Your goal in this question is to design an OCRS for knapsack. The general challenge with knapsack is that accepting an agent of size ε can make an agent of size 1 infeasible. Therefore, typical analyses for knapsack partition the agent set $[T]$ into "Large", $T_L := \{t : s_t > 1/2\}$, and "Small", $T_S := \{t : s_t \leq 1/2\}$, and consider either one or the other. Let $w_L := \sum_{t \in T_L} s_t x_t$ and $w_S := \sum_{t \in T_S} s_t x_t$ denote the total expected size of the Large and Small agents respectively, and since $\mathbf{x}$ lies in the knapsack polytope, we have that $w_L + w_S \leq 1$.

We study the following form of algorithm. First, it flips a global coin and w.p. α, it considers only the Large agents; w.p. $1 - \alpha$, it considers only the Small agents. α is to be optimized later.

1. Conditional on the algorithm focusing on Large agents, show that any agent $t \in T_L$ can be accepted with probability $1/(1 + 2w_L)$.

   Hint: Because only one Large agent can fit anyway, the algorithm should just be the single-unit OCRS algorithm. To complete the proof, you will need to use the fact that $s_t > 1/2$ for all $t \in T_L$.

2. Conditional on the algorithm focusing on Small agents, show that any agent $t \in T_S$ can be accepted with probability $1/(1 + 2w_S)$.

   Hint: Although the end result is identical, the argument is different. The idea is still to accept every agent $t \in T_S$ with probability exactly cx_t, for $c = 1/(1 + 2w_S)$. This may require flipping a coin that is calibrated based on the probability that agent t can be feasibly accepted, i.e. the **probability that the space used before time t is at most** $1 - s_t$. You may have to use Markov's inequality at some point.

3. Optimize the parameter α to get a 1/4-selectable OCRS.