CMSC 858F: Algorithmic Game Theory
Fall 2010
Introduction to Algorithmic Game Theory

Instructor: Mohammad T. Hajiaghayi
Scribe: Hyoungtae Cho

October 13, 2010

1 Overview

In this lecture, we introduce *the classic secretary problem* and its generalization *the multiple-choice secretary problem*, and then discuss *the submodular multiple-choice secretary problem*. At the end, we consider *online auctions* for dynamic environments with a power-point slide[2].

2 Secretary Problem

The *Secretary Problem* is traditionally formulated as *marriage(fiancee) problem* or *house selling problem*[1]. It has been richly studied in *optimal stopping theory*, in which we design the mechanism of choosing a time to take a particular action under uncertainty so as to maximize an objective function. The problem is of much interest due to its close connections with online auctions.

2.1 The classical secretary problem

Setting There are N job applicants arriving online one by one(sequentially) in random order, assuming that the total number of applicants is known. After interviewing an applicant i, we can decide its value v_i relative to all $1, 2, \ldots, i\text{-}1$ applicants who have been already interviewed and we must make an irrevocable decision about whether or not to hire the applicant. The goal is to hire the best applicant.

Random Permutation Model Values and thus Ranks of applicants can be decided by the adversary but the order in which they appear is uniformly at random(i.e., same chance of appearance). The motivation of this model is when the values are coming from an unknown distribution which is practical.

A trivial solution is to randomly choose some element $j \in [1, N]$ to achieve N-competitive ratio(i.e., its probability of success is $1/N$) in the worst case. But we can do better by applying the stopping rule policy.

Optimal Algorithm

- Interview the first k applicants without accepting any
- Let max_v be the score(value) of best element
- Among the rest of candidates, choose the first person whose score beats max_v

Indeed, Dynkin showed the optimal stopping rule for this problem in 1963 and proved that the probability of choosing the best applicant with a cutoff k, *P(k)* approaches $\frac{1}{e}$, as N goes to infinity.

Proof: We can obtain $P(k)$ by summing over all possible j, $k+1 \leq j \leq N$. $P(k)=\sum_{j=k+1}^{N} P(j\text{th applicant is best}) \cdot P(\text{the second best is among the first } k \text{ elements}) = \sum_{j=k+1}^{N} \frac{1}{N}(\frac{k}{j-1}) = \frac{k}{N}\sum_{j=k+1}^{N}(\frac{1}{j-1})$.

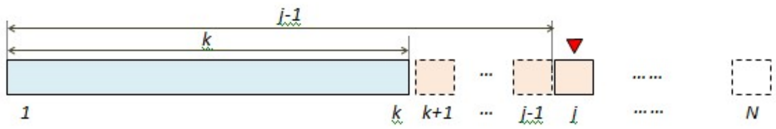

Figure 1: the probability of choosing the best applicant at jth element

As N$\rightarrow \infty$, let x be $\lim_{N\rightarrow\infty} \frac{k}{N}$, t be $\frac{j}{N}$, and dt be $\frac{1}{N}$. The sum can be approximated by the integral, $P(k)=\frac{k}{N}\sum_{j=k+1}^{N}(\frac{1}{j-1}) \approx P(x)=x\int_{x}^{1}\frac{1}{t}dt = -x\ln x$.
Then, we find $k_{optimal}$ by taking the derivation of $P(k)$ with respect to x, setting it to 0, and solving for x. Since $P'(k_{optimal}) = -\ln x - 1 = 0 \Rightarrow x = \frac{1}{e}$, thus $P(k_{optimal}) = P(\frac{1}{e}) = -(\frac{1}{e}) \cdot \ln\frac{1}{e} = \frac{1}{e}$.
In sum, we conclude that the optimal cutoff tends to $\frac{N}{e}$ as N increases and the best applicant is selected with probability $\frac{1}{e}$ which is 0.368(refer to p424~427 in [5]). ■

Scribe: Hyoungtae Cho
Lecture 7 **Date:** 10/13/2010

2.2 Extended Secretary Problem

In addition to the original secretary problem, *marriage(fiancee) problem*, a variety of its extended problems are considered:

- Hiring Secretaries
- Dynamic Auction Market
- Variation of House Selling Problem
- The One-Armed Bandit Problem
- Detecting a change point
- The Burglar Problem

The secretary problem for *unknown number* of applicants in which its optimal solution is much harder has been also studied in several papers. In these days, one of promising applications in this problem is online auction with setting k secretaries instead of just *1* (again in *a random permutation model*). There are three different objectives to design online market:

1. Maximize the probability of selecting best k
2. Minimize the expected sum of ranks of selected k
3. Maximize the expected sum of values of selected k

We focus mainly on the third objective and its generalization in the next section, though the other two can be modeled as well.

2.3 Submodular Secretary Problem

Definition 1 *A function f is* monotone *if and only if f(A)≤f(B) for A⊆B⊆S.*

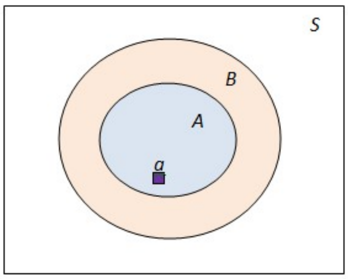

Figure 2: *submodularity*, the marginal profit(return) of each item such as a is non-increasing

Scribe: Hyoungtae Cho
Lecture 7 **Date:** 10/13/2010

Definition 2 *A function f on the subsets of a universe set x is* submodular *if for any pair of subsets of X like A and B such that f(A)+f(B)≥f(A∪B)+f(A∩B) (cf., it is subadditive if f(A)+f(B)≥f(A∪B)). An equivalent characterization is that the marginal profit of each item should be non-increasing and economies of scale, i.e., f(A∪ {a})-f(A)≥f(B∪ {a})-f(B) if A⊆B⊆S and a ∈S\B.*

We say that a function f is *submodular* if adding an element, a to a set B causes a smaller marginal improvement than adding the same element to a subset of B, namely A (refer to *the Figure 2*). So the benefit of adding elements decreases as the set to which they are being added grows.

We want to hire k secretaries such that each subset of secretaries has a value for us which is *not necessarily linear* and this function is *submodular* in many practical applications. If f(A)=$\sum_{a\in A} V(a)$, i.e., linear case, the problem has a (1-$\theta(\frac{1}{\sqrt{k}})$)-competitive algorithm where its competitive ratio is approximated by 1 as k goes to ∞. Also, the case that f(A)=$\max_{a\in A} V(a)$ has been considered in *the classical secretary problem.*

Do we have a c-competitive algorithm for *the (non-monotone) submodular secretary problem* which is a little bit harder? (cf., the ratio can be as bad as $\Omega(\sqrt{n})$ for the subadditive secretary problem)

Theorem 1 *There is an* $8e^2$*-competitive algorithm for the non-monotone submodular secretary problem.*

The algorithm is simple and intuitive. We partition the input stream into k equal part $S_1, S_2, \ldots, S_k$ and choose exactly one secretary in each part. Let T_i be the set of i secretaries we choose in the first i parts. In part S_{i+1}, we try to choose a secretary $x \in S_{i+1}$ that maximizes; Δ_i=f($x\cup T_i$)-f(T_i). We prove the expected value of f(T_k) is within a constant factor of the optimum solution. Let R=$\{a_{i_1}, a_{i_2}, \ldots, a_{i_k}\}$ be the optimal solution. Note that the set $\{i_1, i_2, \ldots, i_k\}$ is a uniformly random subset of $\{1, 2, \ldots, n\}$. Define $X = \{S_i : |S_i \cap R| \neq \phi\}$. For each $S_i \in$X, we pick one element, say s_i, of $S_i \cap R$ randomly. These selected items form a set called $R' = \{s_1, s_2, \ldots, s_{|X|}\}$. Since our algorithm approximates such a set, we study the value of such random samples of R in the following intuitive lemmas. We first show that restricting ourselves to picking at most one element from each segment does not prevent us from picking many elements from the optimal solution (i.e., R).

Lemma 1 *The expected value of the number of items in* R' *is at least* $k(1 - \frac{1}{e})$.

Proof: By the standard methods like balls and bins(refer to p451~452 in [5]), we can show at most $\frac{1}{e}$ fractions of S_i have empty intersection with R[4]. Hence, the expected number of such sets is at most $\frac{k}{e}$. Therefore the expected value of $|R'|$ is at least $k(1 - \frac{1}{e})$. ■

Lemma 2 *For a random subset A of R, the expected value of f*(A) *is at least* $\frac{|A|}{k}$*f*(R).

Proof: The proof can be done by randomness and submodularity. The crux of our proof is that the local optimization steps(i.e., trying to make the best move in each segment) indeed lead to a globally approximate solution. ∎

Lemma 3 *The expected value of* $f(T_k)$ *is at least* $O(\frac{|R'|}{ek}f(R))$*.*

Proof: Assume s_1 is the first element of R' to appear in our ordering. Define $\Delta_j = f(T_j) - f(T_{j-1})$. Note that due to monotonicity $\Delta_j \geq 0$ and thus $E[\Delta_j] \geq 0$ with probability $\frac{1}{e}$, we choose the element in S_j which maximizes the value of $f(T_j)$ given that the set T_{j-1} is fixed. We get factor $\frac{|R'|}{k}$ since we can relate the value of $f(T_k)$ with $O(f(R'))$ which is $\frac{|R'|}{k}$ by *Lemma 2*. ∎

Theorem 2 *The expected value of the output of our algorithm is at least* $O(\frac{1-\frac{1}{e}}{e}f(R))$ *(indeed, the exact number is* $\frac{1-\frac{1}{e}}{7}f(R)$*).*

Proof: Using *Lemma 1* that the expected value of $R' \geq k(1-\frac{1}{e})$. The details of proof are found in [4]. ∎

For non-monotone case, this algorithm does not work, since if we select some item for the ith set, it may hurt us after later step. However slight changes works. Let us consider that we divide the input stream into two equal parts and let $0 \leq x \leq 1$ be a uniformly random variable. If $x \leq \frac{1}{2}$, we ran the algorithm on the first part(and thus ignore the second part!). Otherwise we run the algorithm on the second part and ignore the first part. One can prove that the sum of the solutions in two parts is a constant times the optimal solution, OPT and thus we get at least half of it[4]. Next we consider the applications of secretary problem in online auctions.

3 Online Auctions for Dynamic Environments

Online auctions consider a setting in which agents arrive dynamically and require that an allocation and payment decision is made before they depart. As a motivating example, let us consider *Last-Minute Tickets*[2].

3.1 Example: Last-Minute Tickets

While a traditional auction require that all bids should be received before the auction deadline and then determine the winners, bidders in an online auction have different time constraints and decision horizons, which can be suboptimal.

In this example, the auctioneer is selling one ticket in each hour to the highest bidder at the second highest bid price and only competition makes a price. The value- and time-strategies for three agents in the auction are shown in *the Figure 3*, and if they are truthfully reported then the first bidder wins a ticket for $80 at 11am and the second bidder does for $60 at 12pm, which is

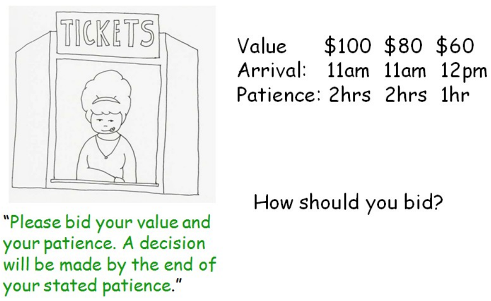

Figure 3: *Last-Minute Tickets Auction*

denoted by $\langle 1, \$80\rangle$, $\langle 2, \$60\rangle$. Would changing a value- or time- strategy affect their incentives?

1. Bidder 1 reduce bid price for a ticket from \$100 to \$65
   $\Rightarrow \langle 2, \$65\rangle$, $\langle 1, \$60\rangle$

2. Bidder 1 delay bid until 12pm with 1hr patience
   $\Rightarrow \langle 2, \$0\rangle$, $\langle 1, \$60\rangle$

3.2 Basic Setting

Setting For k identical goods, each bidder wants at most one copy of the good and the seller derives no utility from retaining copies of the good. Let θ denote the set of agent types, in which an agent i reports its type $\theta_i = (a_i, d_i, w_i)$ that consists of arrival time a_i, departure time d_i and value for the one item w_i.

We make the standard assumption that agent i has a *quasi-linear utility function* $u_i = w_i - \mathrm{price}$.

Quasi-linear utility For an agent i and an allocation x, let $q_i = 1$ if the set $x_i \cap [a_i, d_i]$ contains an interval of length at least 1, otherwise $q_i = 0$.

This online auction mechanism consists of an *allocation rule*, f and a *payment rule*, p. When the vector of reported types is $(\theta_1, \theta_2, \ldots, \theta_n)$, $f(\theta_1, \theta_2, \ldots, \theta_n)$ indicates the allocation which is chosen, and $p(\theta_1, \theta_2, \ldots, \theta_n)$ indicates the amount that agent i must pay.

allocation rule f: $\Theta^n \longmapsto$ Schedules(set of allocations)
payment rule p: $\Theta^n \longmapsto \mathfrak{R}^n$

In addition, we are interested in mechanisms with *dominant-strategy* equilibrium, such that every agent i has a single optimal strategy whatever the strategies and types of other agents. It will require mechanisms to time- and value-SP(*strategyproof* or *truthful*) by reporting its true value $\langle a_i, d_i, w_i \rangle$ immediately upon arrival as a dominant strategy equilibrium (i.e., no benefit otherwise). Recall that we consider *limited-supply (k, 1)* online auction in which an auctioneer has k identical goods to sell in any period before time horizon, T. We assume that agents have no value for receiving an item outside of their arrival-departure interval[3].

We adopt the benchmark offline mechanisms for the purpose of competitive analysis.

- *Efficiency* benchmark is the highest value in this case (in general EFF(v)= $\sum_{i \le k} v^{(i)}$ where $v^{(i)}$ is ith highest value). This is the total value achieved in an offline Vickrey auction, allocating k highest bidders and breaking ties arbitrarily.

- *Revenue* benchmark is Vickrey price, the second highest value in this case (in general $F^{(2)}(v) = \max_{2 \le l \le k} l \cdot v^{(l)}$ where again $v^{(i)}$ is ith highest value). This is the maximal revenue achievable by a fixed-price in which the number of items sold in between 2 and k.

3.3 Adaptive Limited-Supply Auction

First, we consider *the online selection problem*, in which we remove incentives and specialize to the case of disjoint arrival-departure intervals.

Figure 4: The online selection problem

It can be reduced to the secretary problem that we discussed in the previous chapter. To recall it,

- We interview n jobs applicant in random order and want to maximize probability of selecting best applicant

- Knowing *relative ordering* with regard to applicants already interviewed, we must decide to hire or pass

By the secretary algorithm using the stopping rule suggested by Dynkin, we picks the maximum element with probability approaching $\frac{1}{e}$ as $n \rightarrow \infty$.

- We observe the first $\frac{n}{e}$ elements, and set a threshold equal to the maximum seen so far

- We stop(pick) the next time this threshold is reached or exceeded.

Although auction model provides the seller with the numerical values of their bids, not just their relative ordering, it seems plausible that there is no way to capitalize on this numerical information, since we are making no assumptions at all about the distribution of bids.

However, the secretary algorithm is not truthful, since it is not time-SP. Early agents have an incentive to hide until after time $t = \lfloor \frac{n}{e} \rfloor$, when the $\frac{n}{e}$-th agent appears. So change the mechanism:

1. At time τ, denoting arrival j = $\frac{n}{e}$ let p$\geq$q be the top two bids yet received
2. If any agent bidding p has not yet departed, sell to that agent (breaking ties randomly) at price q
3. Else, sell to the next agent whose bid is at least p

Let us consider two examples of this *Adaptive Limited-Supply Auction*.

Figure 5: Adaptive Limited-Supply Auction

Assume that we know the total number of bidders, n=6 and j=$\lfloor \frac{n}{e} \rfloor$=2. In the first example, agent 1 who is *the highest bidder* with \$5(p) bid at time τ wins at price \$2(q), the second highest bid. In the second example, since the highest bidder, agent 1 has departed, sell to the next agent whose bid is at least \$5(p), namely agent 3 in this case, at price p.

In order to analyze the *truthfulness* of this mechanism, we look into as follows:

- If agent i wins, the price charged to her does not depend on her reported valuation

- Possibility agent i wins is (weakly) increasing in w_i, hence *no incentive to understate* w_i
- Reporting $w'_i \geq w_i$ cannot increase the possibility that agent i wins at a price $\leq w_i$, hence *no incentive to overstate* w_i
- Price facing agent i is never influenced by d_i, so *no incentive to misstate* d_i.

To prove truthfulness, only arrival time, a_i which is a bit tricky, is left. We claim that it is alway better to report the true arrival time a_i if possible, given two arrival times $a_i < a'_i$.

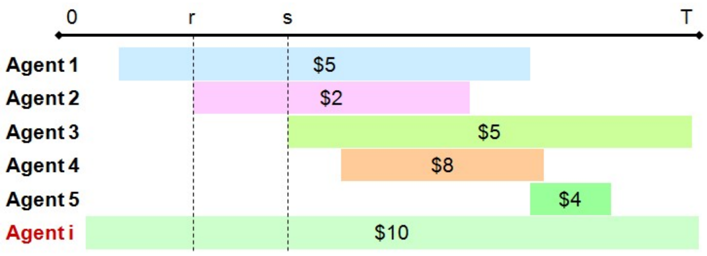

Figure 6: *Arrival time* analysis for truthfulness

Let r,s be the ($\lfloor \frac{n}{e} \rfloor$-1)-th and $\lfloor \frac{n}{e} \rfloor$-th arrival times excluding agent i (say $\frac{n}{e}$ in this case). Stating true arrival a_i, agent i is offered price \$5 on transition while agent *2* defines transition. Stating arrival time a'_i in (a_i,r] does not change anything. If agent i states arrival time a'_i in (r, s), then it influences the transition time τ but not the pricing. Finally, stating arrival time $a'_i \geq$ s influences the transition, but price is not improved (still \$5).

We consider two cases for the competitiveness of the mechanism as follows:
Item sells at time τ: In this case, winner is the highest-bidding agent among first $\frac{n}{e}$. With probability $\frac{1}{e}$, this is also the highest bidder among all n agents.
Item does not sell at time τ: In this case, the auction picks the same outcome as the secretary problem, whose success probability is $\frac{1}{e}$.

For both cases, we have a constant efficiency- and revenue- competitiveness; the competitive ratio for *efficiency* is *e+O(1)*, and the competitive ratio for *revenue* w.r.t. Vickery is e^2*+O(1)*.

In all of this mechanism, we have a *learning phase* to set a fixed price for the rest of the auction, followed by an *accepting phase* to exploit information while making sure truthfulness. It can be extended multiple-choice secretary problem and online auctions with reusable-goods(e.g., Wi-Fi use in Starbucks) and model-based online auction (in which we know the history of valuation in the market).

4 Announcement

- Next class will be the talk on "Algorithms for Online Stochastic Ad Allocation" by *Vahab Mirrokni* from Google Research in CSIC 3117 at 4~5pm, followed by the normal class in our classroom at 5:20~7pm.
- The project report with 5min. presentation will be prepared for class on Oct 27.

References

[1] Hajiaghayi, Mohammad T. *Algorithmic Game Theory: Lecture 7– Handwritten notes.* Oct 13, 2010. University of Maryland.

[2] Hajiaghayi, Mohammad T. *Algorithmic Game Theory: Lecture 7– PowerPoint slides.* Oct 13, 2010. University of Maryland.

[3] Hajiaghayi, Mohammad T. et al. *Adaptive Limited-Supply Online Auctions.* 2004.

[4] Bateni, MohammadHossein et al. *Submodular Secretary Problem and Extensions.* Feb 1, 2010. MIT.

[5] Nisan, Noam et al. *Algorithmic Game Theory.* New York: Cambridge University Press, 2007.