The Temp Secretary Problem

Amos Fiat[1], Ilia Gorelik[1], Haim Kaplan[1], Slava Novgorodov[1]

Tel Aviv University, Israel,
fiat@tau.ac.il, iliagore@post.tau.ac.il, haimk@tau.ac.il,
slavanov@post.tau.ac.il
Research supported by The Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11), and by ISF grant no. 822/10.

Abstract. We consider a generalization of the secretary problem where contracts are temporary, and for a fixed duration γ. This models online hiring of temporary employees, or online auctions for re-usable resources. The problem is related to the question of finding a large independent set in a random unit interval graph.

1 Introduction

This paper deals with a variant of the secretary model, where contracts are temporary. *E.g.*, employees are hired for short-term contracts, or re-usable resources are rented out repeatedly, etc. If an item is chosen, it "exists" for a fixed length of time and then disappears.

Motivation for this problem are web sites such as Airbnb and oDesk. Airbnb offers short term rentals in competition with classic hotels. A homeowner posts a rental price and customers either accept it or not. oDesk is a venture capitalizing on freelance employees. A firm seeking short term freelance employees offers a salary and performs interviews of such employees before choosing one of them.

We consider an online setting where items have values determined by an adversary, ("no information" as in the standard model [15]), combined with stochastic arrival times that come from a prior known distribution (in contrast to the random permutation assumption and as done in [21,7,16]). Unlike much of the previous work on online auctions with stochastic arrival/departure timing ([18]), we do not consider the issue of incentive compatibility with respect to timing, and assume that arrival time cannot be misrepresented.

The temp secretary problem can be viewed

1. As a problem related to hiring temporary workers of varying quality subject to workplace capacity constraints. There is some known prior $F(x) = \int_0^x f(z)dz$ on the arrival times of job seekers, some maximal capacity, d, on the number of such workers that can be employed simultaneously, and a bound k on the total number than can be hired over time. If hired, workers cannot be fired before their contract is up.
2. Alternately, one can view the temp secretary problem as dealing with social welfare maximization in the context of rentals. Customers arrive according

to some distribution. A firm with capacity d can rent out up to d boats simultaneously, possibly constrained to no more than k rentals overall. The firm publishes a rental price, which may change over time *after* a customer is serviced. A customer will choose to rent if her value for the service is at least the current posted price. Such a mechanism is inherently dominant strategy truthful, with the caveat that we make the common assumption that customers reveal their true values in any case.

We give two algorithms, both of which are quite simple and offer posted prices for rental that vary over time. Assuming that the time of arrival cannot be manipulated, this means that our algorithms are dominant strategy incentive compatible.

For rental duration γ, capacity $d = 1$, no budget restrictions, and arrival times from an arbitrary prior, the *time-slice algorithm* gives a $\frac{1}{2e}$ competitive ratio. For arbitrary d the competitive ratio of the time-slice algorithm is at least $(1/2) \cdot (1 - 5/\sqrt{d})$. This can be generalized to more complex settings. The time slice algorithm divides time into slices of length γ. It randomly decides if to work on even or odd slices. Within each slice it uses a variant of some other secretary problem (*E.g.*, [26], [2], [24]) except that it keeps track of the cumulative distribution function rather than the number of secretaries.

The more technically challenging *Charter algorithm* is strongly motivated by the k-secretary algorithm of [24]. For capacity d, employment period γ, and budget $d \leq k \leq d/\gamma$ (the only relevant values), the Charter algorithm does the following:

– Recursively run the algorithm with parameters $\gamma, \lfloor k/2 \rfloor$ on all bids that arrive during the period $[0, 1/2)$.
– Take the bid of rank $\lceil k/2 \rceil$ that appeared during the period $[0, 1/2)$, if such rank exists and set a threshold T to be it's value. If no such rank exists set the threshold T to be zero.
– Greedily accept all items that appear during the period $[1/2, 1)$ that have value at least T — subject to not exceeding capacity (d) or budget (k) constraints.

For $d = 1$ the competitive ratio of the Charter algorithm is at least

$$\frac{1}{1+k\gamma}\left(1 - \frac{5}{\sqrt{k}} - 7.4\sqrt{\gamma \ln(1/\gamma)}\right).$$

Two special cases of interest are $k = 1/\gamma$ (no budget restriction), in which case the expression above is at least $\frac{1}{2}\left(1 - 12.4\sqrt{\gamma \ln(1/\gamma)}\right)$. We also show an upper bound of $1/2 + \gamma/2$ for $\gamma > 0$. As γ approaches zero the two bounds converge to 1/2. Another case of interest is when k is fixed and γ approaches zero in which this becomes the guarantee given by Kleinberg's k-secretary algorithm.

For arbitrary d the competitive ratio of the Charter algorithm is at least

$$1 - \Theta\left(\frac{\sqrt{\ln d}}{\sqrt{d}}\right) - \Theta(\gamma \log(1/\gamma)).$$

We remark that neither the time slice algorithm nor the Charter algorithm requires prior knowledge of n, the number of items due to arrive.

At the core of the analysis of the Charter algorithm we prove a bound on the expected size of the maximum independent set of a random unit interval graph. In this random graph model we draw n intervals, each of length γ, by drawing their left endpoints uniformly in the interval $[0,1)$. We prove that the expected size of a maximum independent set in such a graph is about $n/(1+n\gamma)$. We say that a set of length γ segments that do not overlap is γ-independent. Similarly, a capacity d γ-independent set allows no more than d segments overlapping at any point.

Note that if $\gamma = 1/n$ then this expected size is about $1/2$. This is intuitively the right bound as each interval in the maximum independent set rules out on average one other interval from being in the maximum independent set.

We show that a random unit interval graph with n vertices has a capacity d γ-independent subset of expected size at least $\min(n, d/\gamma)(1-\Theta(\sqrt{\ln d}/\sqrt{d}))$. We also show that when $n = d/\gamma$ the expected size of the maximum capacity d γ-independent subset is no more than $n(1-\Theta(1/\sqrt{d}))$. These results may be of independent interest.

Related work: Worst case competitive analysis of interval scheduling has a long history, e.g., [30,28]. This is the problem of choosing a set of non-overlapping intervals with various target functions, typically, the sum of values.

[19] introduce the question of auctions for reusable goods. They consider a worst case mechanism design setting. Their main goal is addressing the issue of time incentive compatibility, for some restricted set of misrepresentations.

The secretary problem is arguably due to Johannes Kepler (1571-1630), and has a great many variants, a survey by [15] contains some 70 references. The "permutation" model is that items arrive in some random order, all $n!$ permutations equally likely. Maximizing the probability that the best item is chosen, when the items appear in random order, only comparisons can be made, and the number of items is known in advance, was solved by [27] and by [12]. A great many other variants are described in ([15,11]), differing in the number of items to be chosen, the target function to be maximized, taking discounting into account, etc.

An alternative to the random permutation model is the stochastic arrival model, introduced by Karlin [21] in a "full information" (known distribution on values) setting. Bruss [7] subsequently studied the stochastic arrival model in a no-information model (nothing is known about the distribution of values). Recently, [13] made use of the stochastic arrival model as a tool for the analysis of algorithms in the permutation model.

Much of the recent interest in the secretary problem is due to it's connection to incentive compatible auctions and posted prices [18,24,2,3,1,10].

Most directly relevant to this paper is the k-secretary algorithm by R. Kleinberg [24]. Constrained to picking no more than k secretaries, the total value of the secretaries picked by this algorithm is at least a $(1-\frac{5}{\sqrt{k}})$ of the value of the best k secretaries.

Babaioff *et al.* [2] introduced the *knapsack secretary problem* in which every secretary has some weight and a value, and one seeks to maximize the sum of values subject to a upper bound on the total weight. They give a $1/(10e)$ competitive algorithm for this problem. (Note that if weights are one then this becomes the k-secretary problem). The Matroid secretary problem, introduced by Babaioff et al. [4], constrains the set of secretaries picked to be an independent set in some underlying Matroid. Subsequent results for arbitrary Matriods are given in [8,26,14].

Another generalization of the secretary problem is the online maximum bipartite matching problem. See [25,22]. Secretary models with full information or partial information (priors on values) appear in [5] and [29]. This was in the context of submodular procurement auctions ([5]) and budget feasible procurement ([29]). Other papers considering a stochastic setting include [23,17].

In our analysis, we give a detailed and quite technical lower bound on the size of the maximum independent set in a random unit interval graph (produced by the greedy algorithm). Independent sets in other random interval graph models were previously studied in [20,9,6].

2 Formal Statement of Problems Considered

Each item x has a value $v(x)$, we assume that for all $x \neq y$, $v(x) \neq v(y)$ by consistent tie breaking, and we say that $x > y$ iff $v(x) > v(y)$. Given a set of items X, define $v(X) = \sum_{x \in X} v(x)$ and $T_k(X) = \max_{T \subseteq X, |T| \leq k} v(T)$.

Given a set X and a density distribution function f defined on $[0, 1)$, let $\theta_f : X \mapsto [0, 1)$ be a random mapping where $\theta_f(x)$ is drawn independently from the distribution f. The function θ_f is called a *stochastic arrival* function, and we interpret $\theta_f(x)$, $x \in X$, to be the time at which item x arrives. For the special case in which f is uniform we refer to θ_f as θ.

In the problems we consider, the items arrive in increasing order of θ_f. If $\theta_f(x) = \theta_f(y)$ the relative order of arrival of x and y is arbitrary. An online algorithm may select an item only upon arrival. If an item x was selected, we say that the online algorithm *holds* x for γ time following $\theta_f(x)$.

An online algorithm A for the temp secretary problem may hold at most one item at any time and may select at most k items in total. We refer to k as the *budget* of A. The goal of the algorithm is to maximize the expected total value of the items that it selects. We denote by $A(X, \theta_f)$ the set of items chosen by algorithm A on items in X appearing according to stochastic arrival function θ_f.

The set of the arrival times of the items selected by an algorithm for the temp secretary problem is said to be γ-*independent*. Formally, a set $S \subset [0, 1)$ is said to be γ-*independent* if for all $t_1, t_2 \in S$, $t_1 \neq t_2$ we have that $|t_1 - t_2| \geq \gamma$.

Given $\gamma > 0$, a budget k, a set X of items, and a mapping $\theta_f : X \mapsto [0, 1)$ we define $\mathrm{Opt}(X, \theta_f)$ to be a γ-independent set S, $|S| \leq k$, that maximizes the sum of values.

Given rental period $\gamma > 0$, distribution f, and budget k, the competitive ratio of an online algorithm A is defined to be

$$\inf_X \frac{\mathrm{E}_{\theta_f : X \mapsto [0,1)}[v(A(X, \theta_f))]}{\mathrm{E}_{\theta_f : X \mapsto [0,1)}[v(\mathrm{Opt}(X, \theta_f))]}. \tag{1}$$

The competitive ratio of the temp secretary problem is the supremum over all algorithms A of the competitive ratio of A.

Note that when $\gamma \to 0$, the the temp secretary problem reduces to Kleinberg's k-secretary problem.

We extend the γ-temp secretary problem by allowing the algorithm to hold at most d items at any time. Another extension we consider is the *knapsack temp secretary problem* where each item has a weight and we require the set held by the algorithm at any time to be of total weight at most W. Also, we define the *Matroid temp secretary problem* where one restricts the set of items held by the algorithm at any time to be an independent set in some Matroid M.

More generally, one can define a temp secretary problem with respect to some arbitrary predicate P that holds on the set of items held by an online algorithm at all times t. This framework includes all of the variants above. The optimal solution with respect to P is also well defined.

3 The time-slice Algorithm.

In this section we describe a simple time slicing technique. This gives a reduction from temp secretary problems, with arbitrary known prior distribution on arrival times, to the "usual" continuous setting where secretaries arrive over time, do not depart if hired, and the distribution on arrival times is uniform. The reduction is valid for many variants of the temp secretary problem, including the Matroid secretary problem, and the knapsack secretary problem. We remark that although the Matriod and Knapsack algorithms are stated in the random permutation model, they can be replaced with analogous algorithms in the continuous time model and can therefore be used in our context.

We demonstrate this technique by applying it to the classical secretary problem (hire the best secretary). We obtain an algorithm which we call $Slice_\gamma$ for the temp secretary problem with arbitrary prior distribution on arrival times that is $O(1)$ competitive.

Consider the $1/2\gamma$ time intervals (*i.e.* slices) $I_j = [2\gamma j,\ 2\gamma(j+1))$, $0 \leq j \leq 1/(2\gamma) - 1$. We split every such interval into two, $I_j^\ell = [2\gamma j,\ 2\gamma j + \gamma)$, $I_j^r = [2\gamma j + \gamma,\ 2\gamma(j+1))$.[1]

Initially, we flip a fair coin and with probability 1/2 decide to pick points only from the left halves (I_j^ℓ's) or only from the right halves (I_j^r's). In each such interval we pick at most one item by running the following modification of the continuous time secretary algorithm.

[1] For simplicity we assume that $1/(2\gamma)$ is an integer.

The continuous time secretary algorithm [13] observes the items arriving before time $1/e$, sets the largest value of an observed item as a threshold, and then chooses the first item (that arrives following time $1/e$) of value greater than the threshold. The modified continuous time secretary algorithm observes items as long as the cumulative distribution function of the current time is less than $1/e$, then it sets the largest value of an observed item as a threshold compute a threshold, and choose the next item of value larger than the threshold.

It is clear that any two points picked by this algorithm have arrival times separated by at least γ.

Theorem 1. *The algorithm $Slice_\gamma$ is $1/(2e)$ competitive.*

Proof. The analysis is as follows. Fix the mapping of items to each of the left intervals I_j^ℓ's and to each of the right intervals I_j^r's (leaving free the assignment of items to specific arrival times within their the intervals they are assigned to). Let OPT^ℓ (OPT^r) be the sum of the items of maximum value over all intervals I_j^ℓ (I_j^r). Let OPT be the average optimal value conditioned on this mapping of items to intervals. Clearly,

$$OPT^\ell + OPT^r \geq OPT. \tag{2}$$

For any interval I_j's (I_j^ℓ's) $Slice_\gamma$ gain at least $1/e$ over the top value in the interval conditioned on the event that $Slice_\gamma$ doesn't ignore this interval, this happens with probability $1/2$. Therefore the expected sum of values achieved by $Slice_\gamma$ is at least

$$\frac{1}{2} \cdot \frac{1}{e} OPT^\ell + \frac{1}{2} \cdot \frac{1}{e} OPT^r \ . \tag{3}$$

Substitution (2) into (3) we get the lemma. □

Appropriately choosing times (rather than number of elements) as a function of the prior distribution allows us to do the same for other variants of the secretary problem, the Knapsack (achieving a competitive ratio of $\frac{1}{2} \cdot \frac{1}{10e}$, see [2]) and Matriod ($O(\ln \ln \rho)$ when ρ is the rank of the Matroid, see [26,14]).

4 Improved results for the temp secretary problem for the uniform arrival distribution

In this section we give an improved algorithm, referred as the charter algorithm $C_{k,\gamma}$, for the temp secretary problem with uniform arrival times and capacity 1 (at most one secretary can be hired at any time).

As it is never the case that more than $1/\gamma$ items can be selected, setting $k = \lceil 1/\gamma \rceil$ effectively removes the budget constraint. Note that $C_{k,0}$ is Kleinberg's algorithm for the k-secretary problem, with some missing details added to the description.

To analyze the charter algorithm we establish a lower bound on the expected size of the maximum γ-independent subset of a set of uniformly random points in $[0, 1)$. We apply this lower bound to the subset of the items that Kleinberg's algorithm selects.

4.1 The temp secretary algorithm, $C_{k,\gamma}$: a competitive ratio of $1/(1+k\gamma)$.

This charter algorithm, $C_{k,\gamma}$ gets parameters k (the maximal number of rentals allowed) and γ (the rental period) as is described in detail in Algorithm 1. As the entire period is normalized to $[0,1)$, having $k > \lceil 1/\gamma \rceil$ is irrelevant. Thus, we assume that $k \leq \lceil 1/\gamma \rceil$.[2]

We show that $C_{k,\gamma}(X)$ gains in expectation about $1/(1+k\gamma)$ of the top k values of X, which implies that the competitive ratio (see definition (1)) of $C_{k,\gamma}$ is at least about $1/(1+k\gamma)$.

Note that for $k = \lceil 1/\gamma \rceil$, $C_{k,\gamma}$ has a competitive ratio close to $1/2$, while for $\gamma = 0$, $C_{k,\gamma}$ has a competitive ratio close to 1.

It is easy to see that $C_{k,\gamma}$, chooses a γ-independent set of size at most k.

The main theorem of this paper is the following generalization of Kleinberg's k-secretary problem:

Theorem 2. *For any set of items $S = \{x_i\}_{i=1}^n$, $0 < \gamma \leq \gamma^* = 0.003176$ and any positive integer $k \leq 1/\gamma$:*

$$\mathrm{E}_{\theta:S\mapsto[0,1]}[v(C_{k,\gamma}(S,\theta))] \geq \frac{1}{1+\gamma k}(1-\beta(\gamma,k))T_k(S), \tag{4}$$

where $\beta(\gamma,k) = 7.4\sqrt{\gamma \ln(1/\gamma)} + \frac{5}{\sqrt{k}}$, and the expectation is taken oven all uniform mappings of S to the interval $[0,1)$. (Note that the right hand side of Equation (4) is negative for $\gamma^* < \gamma \leq 0.5$.)

4.2 Outline of the proof of Theorem 2

We prove Theorem 2 by induction on k. For $k \leq 25$ the theorem holds vacuously.

The profit, $p^{[0,1/2)}$, on those items that arrive during the time interval $[0,1/2)$ is given by the inductive hypothesis[3]. However, the inductive hypothesis gives this profit, $p^{[0,1/2)}$, in terms of the top $\lfloor k/2 \rfloor$ elements that arrive before time $1/2$, and not in terms of $T_k(X)$, the value of the top k items overall. Thus, we need to relate $p^{[0,1/2)}$ to $T_k(X)$. In the full version of this paper we show that $p^{[0,1/2)}$ is about $1/2$ of $T_k(X)$.

Let $Z_{>T}$ be the set of items that arrive in the time interval $[1/2,1)$ and have value greater than the threshold T. From $Z_{>T}$ we greedily pick a γ-independent subset[4]. It is easy to see that this set is in fact a maximal γ-independent subset.

To bound the expected profit from the items in $Z_{>T}$ we first bound the size of the maximal γ-independent set amongst these items. To do so we use the following general theorem (see also Section 6 and the full version of this paper).

---

[2] To simplify the presentation we shall assume the in sequel that $k \leq 1/\gamma$.

[3] This profit, $p^{[0,1/2)}$ is $\mathrm{E}_{\theta:S\mapsto[0,1)}\left[v(C_{k,\gamma}^{[0,1/2)}(S,\theta))\right]$, where $C_{k,\gamma}^{[0,1/2)}(S,\theta)$ the set of items chosen by the algorithm during the time period $[0,1/2)$.

[4] modulo the caveat that the arrival time of the 1st item chosen from the 2nd half must be at least γ after the arrival time of the last item chosen in the 1st half.

ALGORITHM 1: The Charter Algorithm $C_{k,\gamma}$.

if $k = 1$ **then**

  /* Use the ''continuous secretary'' algorithm [13]: */

  Let x be the largest item to arrive by time $1/e$ (if no item arrives by time $1/e$ — let x be the absolute zero, an item smaller than all other items).

  $C_{k,\gamma}$ accepts the first item y, $y > x$, that arrives after time $1/e$ (if any)

else

  /* Process the items scheduled during the time interval $[0, 1/2)$ */

  Initiate a recursive copy of the algorithm, $C' = C_{\lfloor k/2 \rfloor, 2\gamma}$.

  $x \leftarrow$ next element // If no further items arrive $x \leftarrow \emptyset$

  while $x \neq \emptyset$ *AND* $\theta(x) < 1/2$ **do**

    Simulate C' with input x and modified schedule $\theta'(x) = 2\theta(x)$.

    if C' *accepts* x **then**

      $C_{k,\gamma}$ accepts x

    $x \leftarrow$ next element // If no further items arrive, $x \leftarrow \emptyset$

  /* Determine threshold T */

  Sort the items that arrived during the time interval $[0, 1/2)$: $y_1 > y_2 > \cdots > y_m$ (with consistent tie breaking).

  Let $\tau = \lceil k/2 \rceil$.

  if $m < \tau$ **then**

    set T to be the absolute zero

  else

    set $T \leftarrow y_\tau$.

  /* Process the items scheduled during the time interval $[1/2, 1)$ */

  do

    if $x > T$ *AND* ($\theta(x) \geq \theta(x') + \gamma$ *where* x' *is the last item accepted by* $C_{k,\gamma}$
      OR no items have been previously accepted) **then**

      $C_{k,\gamma}$ accepts x

    $x \leftarrow$ next element // If no further items arrive, $x \leftarrow \emptyset$

  until $x = \emptyset$ *OR* k *items have already been accepted*

Theorem 3. *Let $Z = \{z_1, z_2, \ldots, z_n\}$ be a set of independently uniform samples, z_i, from the real interval $[0, 1)$. For $0 \leq \gamma \leq 1$,*

$$\mathrm{E}_Z[m(Z,\gamma)] \geq \frac{1-\alpha(\gamma)}{\gamma + 1/n} = \frac{(1-\alpha(\gamma))n}{1+n\gamma}, \quad \text{where } \alpha(\gamma) = 3\sqrt{\gamma \ln(1/\gamma)}, \tag{5}$$

where $m(Z,\gamma)$ denotes the size of the largest γ-independent subset of Z.

We apply Theorem 3 to the items in $Z_{>T}$. We can apply this theorem since arrival times of items in $Z_{>T}$ are uniformly distributed in the 2nd half. Specifically, we give a lower bound on the expected profit of the algorithm from the items in the 2nd half as follows:

1. Condition on the size of $Z_{>T}$.

2. Subsequently, condition on the set of arrival times $\{\theta_1, \theta_2, \ldots, \theta_{|Z_{>T}|}\}$ of the items in $Z_{>T}$ but *not* on which item in $Z_{>T}$ arrives when. This conditioning fixes the γ-independent set selected greedily by the algorithm.
3. We take the expectation over all bijections θ whose image on the domain $Z_{>T}$ is the set $\{\theta_1, \theta_2, \ldots, \theta_{|Z_{>T}|}\}$. The expected profit (over the set $Z_{>T}$ and over these bijections) is "approximately"

$$\frac{\text{Size of maximal } \gamma\text{-independent set from } Z_{>T}}{|Z_{>T}|} \cdot \sum_{z \in Z_{>T}} v(z). \tag{6}$$

The "approximately" is because of some technical difficulties:
   - We cannot ignore the last item amongst those arriving prior to time 1/2. If one such item was chosen at some time $1/2 - \gamma < t < 1/2$ then arrivals during the period $[1/2, t + \gamma)$ cannot be chosen.
   - We cannot choose more than k items in total, if the algorithm choose λ items from the time interval $[0, 1/2)$, it cannot choose more than $k - \lambda$ items from the time interval $[1/2, 1)$, but $k - \lambda$ may be smaller than the size of the γ-independent set from $Z_{>T}$.
4. To get an unconditional lower bound we average Equation (6) over the possible sizes of the γ-independent set as given by Theorem 3.

5 Upper bound for the temp secretary problem with uniform arrival times and with no budget restriction

Theorem 4. *For the temp secretary problem where item arrival times are taken from the uniform distribution, for any $\gamma \in (0, 1)$, any online algorithm (potentially randomized) has a competitive ratio $\leq 1/2 + \gamma/2$.*

Proof. Let A denote the algorithm. Consider the following two inputs:

1. The set S of n-1 items of value 1.
2. The set $S' = S \cup \{x_n\}$ where $v(x_n) = \infty$.

Note that these inputs are not of the same size (which is ok as the number of items is unknown to the algorithm).

Condition the mapping $\theta : S \mapsto [0, 1)$ (but not the mapping of x_n). If A accepts an item x at time $\theta(x)$ we say that the segment $[x, x + \gamma)$ is *covered*. For a fixed θ let $g(\theta)$ be the expected fraction of $[0, 1)$ which is not covered when running A on the set S with arrival times θ. This expectation is over the coin tosses of A. Let G be $\mathrm{E}_{\theta:S\mapsto[0,1)}[g(\theta)]$.

The number of items that A picks on the input S with arrival time θ is at most $\frac{1-g(\theta)}{\gamma} + 1$. Taking expectation over all mappings $\theta : S \mapsto [0, 1)$ we get that the value gained by A is at most $(1 - G)/\gamma + 1$.

As $n \to \infty$ the optimal solution consists of $\lceil 1/\gamma \rceil$ items of total value $\lceil 1/\gamma \rceil$. Therefore the competitive ratio of A is at most

$$\frac{(1 - G)/\gamma + 1}{1/\gamma} = 1 - G + \gamma. \tag{7}$$

Note that $g(\theta)$ is exactly the probability that A picks x_n on the input $S \cup \{x_n\}$ (this probability is over the mapping of x_n to $[0, 1)$ conditioned upon the arrival times of all the items in $S \subset S'$). Therefore the competitive ratio of A on the input S' is

$$\mathrm{E}[g(\theta)] = G \ . \tag{8}$$

Therefore the competitive ratio of A is no more than the minimum of the two upper bounds (7) and (8)

$$\min(G,\ 1 - G + \gamma) \leq 1/2 + \gamma/2 \ .$$

□

6 About Theorem 3: A Lower bound on the expected size of the maximum γ-independent subset

Recall the definition of Z and $m(Z, \gamma)$ from Theorem 3.

Define the random variable X_i, $1 \leq i \leq n$ to be the i'th smallest point in Z. Define the random variable C_i to be the number of points from Z that lie in the interval $[X_i, X_i + \gamma)$. Note that at most one of these points can belong to a γ-independent set.

The *greedy algorithm* constructs a maximal γ-independent set by traversing points of Z from the small to large and adding a point whenever possible. Let I_i be a random variable with binary values where $I_i = 1$ iff X_i was chosen by the greedy algorithm. It follows from the definition that $\sum_i I_i$ gives the size of the maximal independent set, $m(Z, \gamma)$, and that $\sum_i I_i C_i = n$.

Note that $\mathrm{E}[C_i] \leq 1 + n\gamma$, one for the point X_i itself, and $n\gamma$ as the expected number of uniformly random points that fall into an interval of length γ. If C_i and I_i were independent random variables, it would follow that

$$\mathrm{E}\Big[\sum I_i C_i\Big] \leq (1 + n\gamma) \sum \mathrm{Prob}[\, I_i = 1\,],$$

and, thus,

$$m(Z, \gamma) = \sum I_1 \geq n/(1 + n\gamma).$$

Unfortunately, C_i and I_i are not independent, and the full proof of Theorem 3, that deals with such dependencies, appears in the full version of this paper (on the Archive).

7 Discussion and Open Problems

We've introduced online optimization over temporal items under stochastic inputs subject to conditions of two different types:

- "Vertical" constraints: Predicates on the set of items held at all times t. In this class, we've considered conditions such as no more than d simultaneous items held at any time, items held at any time of total weight $\leq W$, items held at any time must be independent in some Matroid.

– “Horizontal” constraints: Predicates on the set of items over all times. Here, we’ve considered the condition that no more than k employees be hired over time.

One could imagine much more complex settings where the problem is defined by arbitrary constraints of the first type above, and arbitrary constraints of the 2nd type. For example, consider using knapsack constraints in both dimensions. The knapsack constraint for any time t can be viewed as the daily budget for salaries. The knapsack constraint over all times can be viewed as the total budget for salaries. Many other natural constraints suggest themselves.

It seems plausible that the time slice algorithm can be improved, at least in some cases, by making use of information revealed over time, as done by the Charter algorithm.

References

1. Moshe Babaioff, Michael Dinitz, Anupam Gupta, Nicole Immorlica, and Kunal Talwar. Secretary problems: Weights and discounts. In SODA ’09, pages 1245–1254.
2. Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. A knapsack secretary problem with applications. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 4627 of *Lecture Notes in Computer Science*, pages 16–28, 2007.
3. Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. Online auctions and generalized secretary problems. *SIGecom Exch.*, 7(2):7:1–7:11, June 2008.
4. Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. Matroids, secretary problems, and online mechanisms. In SODA 2007, pages 434–443. SIAM.
5. Ashwinkumar Badanidiyuru, Robert Kleinberg, and Yaron Singer. Learning on a budget: Posted price mechanisms for online procurement. In EC ’12, pages 128–145, 2012.
6. S. Boucheron and W. Fernandez de la Vega. On the independence number of random interval graphs. *Combinatorics, Probability and Computing*, 10:385–396, 9 2001.
7. F. Thomas Bruss. A unified approach to a class of best choice problems with an unknown number of options. *Ann. Probab.*, 12(3):882–889, 08 1984.
8. Sourav Chakraborty and Oded Lachish. Improved competitive ratio for the matroid secretary problem. In SODA ’12, pages 1702–1712. SIAM, 2012.
9. Jr. Coffman, E.G., Bjorn Poonen, and Peter Winkler. Packing random intervals. *Probability Theory and Related Fields*, 102(1):105–121, 1995.
10. Nikhil R Devanur and Thomas P Hayes. The adwords problem: online keyword matching with budgeted bidders under random permutations. In EC’9, pages 71–78.
11. Michael Dinitz. Recent advances on the matroid secretary problem. *SIGACT News*, 44(2):126–142, June 2013.
12. E.B. Dynkin. The optimal choice of the stopping moment for a markov process. *Dokl. Akad. Nauk. SSSR*, pages 238–240, 1963.

13. Moran Feldman, Joseph Naor, and Roy Schwartz. Improved competitive ratios for submodular secretary problems (extended abstract). In *APPROX-RANDOM*, volume 6845 of *LNCS*, pages 218–229. Springer, 2011.
14. Moran Feldman, Ola Svensson, and Rico Zenklusen. A simple $O(\log\log(\text{rank}))$-competitive algorithm for the matroid secretary problem. In SODA 2015, pp 1189–1201. SIAM, 2015.
15. P. R. Freeman. The secretary problem and its extensions: A review. *International Statistical Review / Revue Internationale de Statistique*, 51(2):189–206, 1983.
16. Guillermo Gallego and Garrett van Ryzin. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Manage. Sci.*, 40(8):999–1020, August 1994.
17. Oliver Göbel, Martin Hoefer, Thomas Kesselheim, Thomas Schleiden, and Berthold Vöcking. Online independent set beyond the worst-case: Secretaries, prophets, and periods. In *ICALP*, volume 8573 of *LNCS*, pages 508–519. Springer, 2014.
18. Mohammad Taghi Hajiaghayi, Robert Kleinberg, and David C. Parkes. Adaptive limited-supply online auctions. In EC '04, pages 71–80, 2004.
19. Mohammad Taghi Hajiaghayi, Robert D. Kleinberg, Mohammad Mahdian, and David C. Parkes. Online auctions with re-usable goods. In EC '05, pages 165–174, 2005.
20. Joyce Justicz, Edward R. Scheinerman, and Peter M. Winkler. Random intervals. *The American Mathematical Monthly*, 97(10):pp. 881–889, 1990.
21. Sam Karlin. Stochastic models and optimal policy for selling an asset. pages 148–158, 1962.
22. Thomas Kesselheim, Klaus Radke, Andreas Tönnis, and Berthold Vöcking. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In *ESA'13*, volume 8125 of *LNCS*, pages 589–600. Springer, 2013.
23. Thomas Kesselheim, Andreas Tönnis, Klaus Radke, and Berthold Vöcking. Primal beats dual on online packing lps in the random-order model. In *STOC'14*, pages 303–312. ACM, 2014.
24. Robert D. Kleinberg. A multiple-choice secretary algorithm with applications to online auctions. In *SODA'05*, pages 630–631. SIAM, 2005.
25. Nitish Korula and Martin Pál. Algorithms for secretary problems on graphs and hypergraphs. In *ICALP'09*, volume 5556 of *LNCS*, pages 508–520. Springer, 2009.
26. Oded Lachish. O(log log rank) competitive-ratio for the matroid secretary problem. *CoRR*, abs/1403.7343, 2014.
27. D.V. Lindley. Dynamic programming and decision theory. *Appl. Statist.*, pages 39–52, 1961.
28. R. Lipton. Online interval scheduling. In SODA '94, pages 302–311, 1994.
29. Yaron Singer and Manas Mittal. Pricing mechanisms for crowdsourcing markets. In WWW '13, pages 1157–1166, 2013.
30. Gerhard J. Woeginger. On-line scheduling of jobs with fixed start and end times. *Theoretical Computer Science*, 130(1):5 – 16, 1994.