

Think Eternally: Improved Algorithms for the Temp Secretary Problem and Extensions*

Thomas Kesselheim^{†1} and Andreas Tönnis^{‡2}

- 1 Max-Planck-Institut für Informatik and Saarland University, Saarbrücken, Germany
thomas.kesselheim@mpi-inf.mpg.de
- 2 Department of Computer Science, RWTH Aachen University, Germany
toennis@cs.rwth-aachen.de

Abstract

The *Temp Secretary Problem* was recently introduced by Fiat et al. [11]. It is a generalization of the Secretary Problem, in which commitments are temporary for a fixed duration. We present a simple online algorithm with improved performance guarantees for cases already considered by Fiat et al. and give competitive ratios for new generalizations of the problem. In the classical setting, where candidates have identical contract durations $\gamma \ll 1$ and we are allowed to hire up to B candidates simultaneously, our algorithm is $(\frac{1}{2} - O(\sqrt{\gamma}))$ -competitive. For large B , the bound improves to $1 - O(\frac{1}{\sqrt{B}}) - O(\sqrt{\gamma})$.

Furthermore we generalize the problem from cardinality constraints towards general packing constraints. We achieve a competitive ratio of $1 - O\left(\sqrt{\frac{(1+\log d + \log B)}{B}}\right) - O(\sqrt{\gamma})$, where d is the sparsity of the constraint matrix and B is generalized to the capacity ratio of linear constraints. Additionally we extend the problem towards arbitrary hiring durations.

Our algorithmic approach is a relaxation that aggregates all temporal constraints into a non-temporal constraint. Then we apply a linear scaling algorithm that, on every arrival, computes a tentative solution on the input that is known up to this point. This tentative solution uses the non-temporal, relaxed constraints scaled down linearly by the amount of time that has already passed.

1998 ACM Subject Classification F.1.2 Models of Computation/Online computation

Keywords and phrases Secretary Problem, Online Algorithms, Scheduling Problems

Digital Object Identifier 10.4230/LIPIcs.ESA.2016.54

1 Introduction

Online resource allocation problems have a notion of time: Choices have to be made at some point in time without knowing the future input. Each decision may make a future one infeasible. The standard example of such a setting is the secretary problem where candidates of different value arrive over time. After each arrival, the algorithm has to decide whether to permanently accept or reject this candidate. Every decision is final. That is, once rejected a candidate will never come back again. Once a candidate is accepted, no other candidate can be accepted anymore.

* A full version of the paper is available at <http://arxiv.org/abs/1606.06926>.

[†] Supported in part by the DFG through Cluster of Excellence MMCI.

[‡] Supported by the DFG GRK/1298 “AlgoSyn”.



In many practical applications, however, commitments are not eternal but affect only a finite time horizon. They may limit options for the upcoming days but not for the rest of the year or even longer. Nevertheless, even with such an assumption, traditional worst-case competitive analysis is typically too strong a benchmark. It is trivial to see that for the respective version of the secretary problem no algorithm achieves a bounded competitive ratio.

Therefore, we consider a partly stochastic model introduced by Fiat et al. [11]. First an adversary chooses which items will arrive. However, it does not determine the arrival times, which are instead drawn from a probability distribution, typically the uniform distribution on $[0, 1]$. In more detail, in the *temp secretary problem*, an adversary defines values of items v_1, \dots, v_n . Afterwards, arrival times τ_j are drawn independently uniformly from $[0, 1]$. As time proceeds, the values and arrival times are revealed to the algorithm. Upon each arrival, the algorithm has to decide whether to accept or to reject the respective item. Each item is accepted for a duration of γ , which is assumed to be much smaller than 1. At any point in time t at most B items may overlap, that is, during time $t - \gamma$ and t at most B items may be accepted.

The objective is to maximize $\sum_{j \in \text{ALG}} v_j$, where $\text{ALG} \subseteq [n]$ denotes the selection by the algorithm. By OPT we denote the optimal selection $\text{OPT} \subseteq [n]$, which maximizes $\sum_{j \in \text{OPT}} v_j$. As the arrival times τ_1, \dots, τ_n are random, both ALG and OPT are random variables. We evaluate the performance of an algorithm by its competitive ratio, defined as $\mathbf{E} \left[\sum_{j \in \text{ALG}} v_j \right] / \mathbf{E} \left[\sum_{j \in \text{OPT}} v_j \right]$.

1.1 Our Contribution

We introduce a new algorithmic approach to online packing problems with temporal constraints. As key idea we consider a relaxation to OPT by removing the temporal constraints and exchanging them with global ones. In the special case of the temp secretary problem, we exploit that for every realization of the arrival dates τ_1, \dots, τ_n the optimal offline solution OPT never contains more than $B \lceil 1/\gamma \rceil$ elements. Therefore, we exchange the constraints by only requiring $B \lceil 1/\gamma \rceil$ items to be picked throughout the process. An online solution to this relaxation can be found using algorithms for online linear packing problems. It then remains to derive a solution to the original constraints.

For the temp secretary problem, this approach allows us to derive a light-weight, easy to state algorithm. We show it to be $\frac{1}{2} (1 - O(\sqrt{\gamma}))$ -competitive for all values of B . Furthermore, for large values of B , a different analysis shows a better competitive ratio of $1 - O(1/\sqrt{B}) - O(\sqrt{\gamma})$. The previous best results for this setting were $\frac{1}{2} \left(1 - O\left(\sqrt{\gamma \ln(1/\gamma)}\right) \right)$ for $B = 1$ and $1 - O\left(\sqrt{(\ln B)/B}\right) - O\left(\sqrt{\gamma \ln(1/\gamma)}\right)$ for large values of B , both by Fiat et al. [11]. Note that $1/2$ is known to be an asymptotic upper bound to the competitive ratio for $B = 1$ [11].

We also generalize the cardinality constraint in the temp secretary problem to arbitrary linear constraints. This enables us to capture more general combinatorial problems, like multiple knapsack constraints that have to be fulfilled simultaneously. For example, we could model scenarios in which the algorithm has to select production orders online in such a way that none of the involved machines is overloaded. Our algorithm is $1 - O\left(\sqrt{(1 + \log d + \log B)/B}\right) - O(\sqrt{\gamma})$ -competitive, where d denotes the maximum number of constraints a single item is contained in. By B we denote the *capacity ratio*, which is defined to be the minimum ratio of a constraint's capacity and the usage of a single item. For non-timed constraints, there are lower bounds in the order of $1 - O\left(\sqrt{\log m/B}\right)$, where m is the number of constraints and $d = m$ [2, 6].

Our algorithm also has a natural generalization to settings with items of different lengths. For the temp secretary problem, we show a competitive ratio of $\frac{1}{4} - \Theta(\sqrt{\gamma})$.

The main technical contribution are bounds on the probability that tentative selections made by the algorithm are actually feasible. In related work, it is usually enough to pretend all previous tentative choices were actually feasible. As these can be considered independent, a concentration bound can be applied. These techniques are apparently not strong enough here and we have to bound the actual commitments. We do so by analyzing coupled random variables that provide an upper bound on the random process. For the case of large B , this analysis is based on a symmetric random walk, representing arrivals and departures of items.

1.2 Related Work

Secretary problems have gained a lot of attention over the last decade, even though the most famous variant was already introduced and solved in the 1960s [12, 21, 7].

The most famous combinatorial generalization is the matroid secretary problem, introduced by Babaioff et al [4]. As of now, the big question of whether there is a constant competitive algorithm for the matroid secretary problem is still open. The best known algorithms for the problem are $O(\log \log \rho)$ -competitive [10, 20]. Constant competitive algorithms are known for most special cases, e.g. there is a $1/2e$ -competitive algorithm for graphical matroids [19], a $1/9.6$ -competitive algorithm for laminar matroids [23] and there is an optimal $1/e$ -competitive algorithm for transversal matroids [16]. For k -uniform matroids, the problem is also known as multiple-choice secretary problem and was solved by Kleinberg, who gave a $(1 - O(1/\sqrt{k}))$ -competitive algorithm and showed that this is optimal [18].

Furthermore, online models with random arrival order have been used for online packing problems. The knapsack secretary problem was introduced by Babaioff et al. [3] and the currently best known competitive ratio is $1/8.1$ [17]. This problem was generalized towards general packing linear programs with special attention on the case with large capacities. There are several known algorithms [1, 13, 17] that feature a competitive ratio of $1 - O\left(\sqrt{\log m/B}\right)$, where m is the total number of constraints and B is a lower bound on the capacities of the constraints. These results match the lower bound by Agrawal et al.[2] and Devanur et al. [6] for the random order and i.i.d. model respectively. Note that the result in [17] is stronger in case of sparse matrices: If the maximal number of non-zero entries in any column is bounded by d , the guarantee only depends on d rather than m .

Another important way of generalizing the secretary problem is the submodular secretary problem introduced by Bateni et al. [5]. The problem generalizes the multiple-choice secretary problem towards submodular objective functions. The currently best known competitive ratio is $\frac{e-1}{e^2+e}$ by Feldmann et al. [9]. For submodular, transversal matroids the best known algorithm is $1/95$ -competitive [23] and for linear packing constraints the best algorithm is known to be $\Omega(1/m)$ -competitive [5].

The temp secretary problem that we consider and generalize in this paper was introduced in 2015 by Fiat et al. [11]. It introduces temporal constraints to the field of online algorithms with random order in a way that had only been considered before in the worst-case model for online interval scheduling [22, 25]. Fiat et al. give an algorithm that is inspired by Kleinberg's algorithm for the multiple-choice secretary problem. Their algorithm iteratively refines the sample $\log n$ times, while our algorithm updates the sample in every round. Both algorithms are closely related, but the one presented here can be described much more compact and allows for a more simple analysis.

Fiat et al. achieve a competitive ratio of $1/(1+k\gamma) \left(1 - 5/\sqrt{k} - 7.4\sqrt{\gamma \ln(1/\gamma)}\right)$ for the case where at most one candidate can be hired simultaneously and the sum of hires cannot

Algorithm 1: Scaling Algorithm for length γ and capacity B

```

for every arriving item  $j$  do
    Set  $t := \tau_j$ ; // arrival time of  $j$ 
    Let  $S^{(t)}$  be the  $\lfloor tB/\gamma \rfloor$  highest-valued items  $j'$  with  $\tau_{j'} \leq t$ ;
    if  $j \in S^{(t)}$  then // if among best items
        if  $S \cup \{j\}$  is a feasible schedule then // and if feasible
            Set  $S := S \cup \{j\}$ ; // then select  $j$ 

```

exceed the budget k . To compare this result to ours, consider the unconstrained budget case $k = 1/\gamma$. In this case, they achieve a competitive ratio of $\frac{1}{2} \left(1 - O\left(\sqrt{\gamma \ln(1/\gamma)}\right)\right)$. Additionally, they show a lower bound $\frac{1+\gamma}{2}$ for this case, thus both algorithms, ours and theirs, are asymptotically tight for $\gamma \rightarrow 0$. For up to B concurrent hires, their algorithm is $1 - \Theta\left(\sqrt{(\ln B)/B}\right) - \Theta\left(\sqrt{\gamma \ln(1/\gamma)}\right)$ -competitive. Additionally, they describe a black-box procedure that transforms any algorithm for a combinatorial secretary problem into an algorithm for the respective combinatorial temp secretary problem. This transformation loses a factor of $1/2$ in the competitive ratio, but also works for general arrival distributions as long as all items have an identical duration.

2 The Temp Secretary Problem

As our first result, we present a simplified and improved algorithm for the temp secretary problem. Here, an adversary chooses a value v_j for each of the n items and after values have been determined arrival times τ_1, \dots, τ_n are drawn independently uniformly at random from $[0, 1]$. Each item when selected stays active for γ time. At any point in time, at most B elements may be active simultaneously.

The optimal selection $\text{OPT} \subseteq [n]$ is a random variable that depends on the arrival times. However, pointwise we have $|\text{OPT}| \leq B \lceil 1/\gamma \rceil$ for any realization of the arrival times τ_1, \dots, τ_n . Therefore, the expected value of OPT can be upper-bounded by the value of the $B \lceil 1/\gamma \rceil$ highest-valued elements, which we denote by $\text{OPT}^* \subseteq [n]$.

Algorithm 1 is inspired by online approximation algorithms for OPT^* , particularly [17]. If item j arrives at time t , then we determine whether it is among the $\lfloor tB/\gamma \rfloor$ highest-valued items seen so far, called $S^{(t)}$. In this case, we call it tentatively selected. If it is also feasible to accept j , we do so. Otherwise, we reject j .

Note that in expectation at time t we have seen a t fraction of OPT^* , so approximately tB/γ items from OPT^* . The set $S^{(t)}$ approximates this set by including the best $\lfloor tB/\gamma \rfloor$ up to this point.

We give two performance bounds for this algorithm. First, in Section 2.2, we show that it is $\frac{1}{2} (1 - O(\sqrt{\gamma}))$ -competitive for all values of B . Afterwards, in Section 2.3, we perform a different analysis for large values of B giving a competitive ratio of $1 - O(1/\sqrt{B}) - O(\sqrt{\gamma})$.

2.1 Analysis Preliminaries

The analyses for both cases follow a similar pattern. First, we analyze the expected value of the set $S^{(t)}$ and thereby the expected value of the tentative selection. Then, we bound the probability that such a tentative selection is feasible.

For analysis purposes, we discretize time into N uniform intervals of length $1/N$, which we call rounds. If $N \gg n$, then the probability that two items fall into the same interval is negligible. Also, if N is large enough, we can effectively assume that all items arrive at times which are multiples of $1/N$ because the value of $\lfloor tB/\gamma \rfloor$ stays constant in almost all intervals. From time to time, it will be helpful to fill up rounds in which no actual item arrives with dummy items that do not have value but also do not use space. The order of these items and dummy items is then a uniformly drawn permutation. Furthermore, to avoid cumbersome notation, we assume that γN and $\sqrt{\gamma}N$ are integer. We overload notation and write $S^{(\ell)}$ instead of $S^{(\ell/N)}$.

To discuss the probability of feasibility, we introduce 0/1 random variables $(C_\ell)_{\ell \in [N]}$ and $(F_\ell)_{\ell \in [N]}$. We set $C_\ell = 1$ if and only if a tentative selection is made in round ℓ . Furthermore, let $F_\ell = 1$ for every round ℓ in which it would be feasible to actually select an item. Finally, let V_ℓ denote the value of the item tentatively selected in round ℓ if any, otherwise set $V_\ell = 0$. So formally $V_\ell = v_j C_\ell$ if item i arrives in round ℓ . The value achieved by the algorithm is given as $\sum_{j \in \text{ALG}} v_j = \sum_{\ell=1}^N V_\ell F_\ell = \sum_{\ell=1}^N V_\ell C_\ell F_\ell$.

Observe that the value of a single random variable C_ℓ is already fully determined by the set $S^{(\ell)}$ and which of these items arrives in round ℓ . Neither the mutual order in rounds $1, \dots, \ell - 1$ nor in $\ell + 1, \dots, N$ matters. Furthermore, conditioned on any set $S^{(\ell)}$ and any order in rounds $\ell + 1, \dots, N$, the probability of $C_\ell = 1$ is at most $\frac{|S^{(\ell)}|}{\ell} \leq \frac{B}{\gamma N}$. As a consequence, for every sequence of values $a_{\ell'} \in \{0, 1\}$ for $\ell < \ell' \leq N$, we have

$$\Pr [C_\ell = 1 \mid C_{\ell'} = a_{\ell'} \forall \ell < \ell' \leq N] \leq \frac{B}{\gamma N} .$$

Note that there are still complicated dependencies among these random variables. For example, at most n of them can be 1. Therefore, based on the above observation, we define coupled random variables that dominate the actual ones but are easier to deal with. We introduce random variables $(\tilde{C}_\ell)_{\ell \in [N]}$ such that pointwise $\tilde{C}_\ell \geq C_\ell$ for which the above relation holds with equality. To define these formally, we iterate over the rounds from large to small index. Conditioned on any value of the variables $\tilde{C}_{\ell+1} = a_{\ell+1}, \dots, \tilde{C}_N = a_N$, the probability $p = \Pr [C_\ell = 1 \mid \tilde{C}_{\ell'} = a_{\ell'} \forall \ell < \ell' \leq N]$ is always at most $\frac{B}{\gamma N}$. Now let $\tilde{C}_\ell = 1$ whenever $C_\ell = 1$ and additionally $\tilde{C}_\ell = 1$ with probability $\frac{\frac{B}{\gamma N} - p}{1 - p}$ if $C_\ell = 0$. This guarantees

$$\Pr [\tilde{C}_\ell = 1 \mid \tilde{C}_{\ell'} = a_{\ell'} \forall \ell < \ell' \leq N] = \frac{B}{\gamma N} ,$$

and therefore, by induction on all subsets of $[N]$, the random variables $\tilde{C}_1, \dots, \tilde{C}_N$ are independent and identically distributed. Note that $\tilde{C}_\ell = 1$ whenever $C_\ell = 1$. So therefore also $\sum_{j \in \text{ALG}} v_j = \sum_{\ell=1}^n V_\ell \tilde{C}_\ell F_\ell$.

Next, we can bound the expected value of the items contained in the set $S^{(\ell)}$.

► **Lemma 1.** For $\ell \geq 2\sqrt{\frac{\gamma}{B}}N$

$$\mathbf{E} \left[\sum_{j \in S^{(\ell)}} v_j \right] \geq \left(1 - 9 \sqrt{\frac{1}{\frac{\ell}{N} \cdot \frac{B}{\gamma}}} \right) \frac{1 - \frac{1}{2} \sqrt{\frac{\gamma}{B}} \frac{\ell}{N}}{1 + \gamma} \sum_{j \in \text{OPT}^*} v_j .$$

The general idea is as follows. In round ℓ , we have seen an $\frac{\ell}{N}$ -fraction of OPT^* in expectation, so ignoring rounding these are $\frac{\ell}{N} \frac{B}{\gamma}$ items in expectation. This approximately matches the size of $S^{(\ell)}$. The set $S^{(\ell)}$ has a slightly smaller value due to variance and rounding. The first two factors compensate these effects.

For the proof, we show that the non-temporal relaxation described here corresponds to a packing linear program and then we apply a result in [17]. For the detailed proof, please refer to the full version.

2.2 General Analysis

We are now ready to analyze the algorithm.

► **Theorem 2.** *Algorithm 1 is at least $\frac{1}{2} \left(1 - \frac{7}{2}\sqrt{\gamma} - \frac{37}{2}\sqrt{\frac{\gamma}{B}} - \gamma\right)$ -competitive for the temp secretary problem with duration γ for all items.*

The key ingredient to the analysis is a bound on the probability that a tentative selection is feasible.

► **Lemma 3.** *For all L , we have*

$$\mathbf{E} \left[\sum_{\ell=L}^{L+\sqrt{\gamma}N-1} \tilde{C}_\ell F_\ell \right] \geq \left(\frac{1}{2} - \sqrt{\gamma} - \frac{1}{4\sqrt{\gamma}N} \right) \mathbf{E} \left[\sum_{\ell=L}^{L+\sqrt{\gamma}N-1} \tilde{C}_\ell \right].$$

Proof. The crux when proving this lemma is that that F_ℓ depends on \tilde{C}_ℓ in a non-trivial way. Therefore, we instead introduce variables \tilde{F}_ℓ defined as follows. We set $\tilde{F}_\ell = 1$ for $\ell < L$ and $\tilde{F}_\ell = \max\{0, 1 - \frac{1}{B} \sum_{i=1}^{\gamma N} \tilde{C}_{\ell-i} \tilde{F}_{\ell-i}\}$ for $\ell \geq L$. The motivation behind this definition is that $F_\ell = 1$ if and only if it is feasible to select an item in round ℓ . So therefore, $F_\ell = 1$ if and only if $\sum_{i=1}^{\gamma N} C_{\ell-i} F_{\ell-i} < B$, implying $F_\ell \geq \max\{0, 1 - \frac{1}{B} \sum_{i=1}^{\gamma N} C_{\ell-i} F_{\ell-i}\}$. The definition of \tilde{F}_ℓ captures this last bound in a pessimistic way. Note that due to the independence of $(\tilde{C}_\ell)_{\ell \in [N]}$, \tilde{C}_ℓ and \tilde{F}_ℓ are now independent.

We first show that pointwise $\sum_{\ell=L}^{L+k} \tilde{C}_\ell F_\ell \geq \sum_{\ell=L}^{L+k} \tilde{C}_\ell \tilde{F}_\ell$ for all $k \in \mathbb{Z}$ by induction on k . Observe that the statement is trivial for $k < 0$ because then both sums are empty. So, let us consider $k \geq 0$ for the induction step. If $\tilde{C}_{L+k} = 0$ or $\tilde{F}_{L+k} = 0$, then also the statement follows trivially from the induction hypothesis. The only interesting case is $\tilde{C}_{L+k} = 1$ and $\tilde{F}_{L+k} > 0$. In this case, we get

$$\begin{aligned} \sum_{\ell=L}^{L+k} \tilde{C}_\ell \tilde{F}_\ell &= \tilde{F}_{L+k} + \sum_{\ell=L}^{L+k-1} \tilde{C}_\ell \tilde{F}_\ell = 1 - \frac{1}{B} \sum_{i=1}^{\gamma N} \tilde{C}_{L+k-i} \tilde{F}_{L+k-i} + \sum_{\ell=L}^{L+k-1} \tilde{C}_\ell \tilde{F}_\ell \\ &= 1 - \frac{1}{B} \sum_{\ell=L+k-\gamma N}^{L-1} \tilde{C}_\ell \tilde{F}_\ell + \left(1 - \frac{1}{B}\right) \sum_{\ell=L}^{L+k-1} \tilde{C}_\ell \tilde{F}_\ell + \frac{1}{B} \sum_{\ell=L}^{L+k-\gamma N-1} \tilde{C}_\ell \tilde{F}_\ell. \end{aligned}$$

At this point, we can apply the induction hypothesis, which states that in the second and third sum we can replace all occurrences of \tilde{F}_ℓ by F_ℓ to get a lower bound. Furthermore, $1 = \tilde{F}_\ell \geq F_\ell$ for $\ell \leq L-1$. So, we can do the same in the first sum. Therefore

$$\begin{aligned} \sum_{\ell=L}^{L+k} \tilde{C}_\ell \tilde{F}_\ell &\leq 1 - \frac{1}{B} \sum_{\ell=L+k-\gamma N}^{L-1} \tilde{C}_\ell F_\ell + \left(1 - \frac{1}{B}\right) \sum_{\ell=L}^{L+k-1} \tilde{C}_\ell F_\ell + \frac{1}{B} \sum_{\ell=L}^{L+k-\gamma N-1} \tilde{C}_\ell F_\ell \\ &= 1 - \frac{1}{B} \sum_{i=1}^{\gamma N} \tilde{C}_{L+k-i} F_{L+k-i} + \sum_{\ell=L}^{L+k-1} \tilde{C}_\ell F_\ell. \end{aligned}$$

Now, we use that $F_{L+k} = 1$ if and only if less than B items have been selected in rounds $L+k-\gamma N$ to $L+k-1$. This gives us $F_{L+k} \geq 1 - \frac{1}{B} \sum_{i=1}^{\gamma N} C_{L+k-i} F_{L+k-i} \geq$

$1 - \frac{1}{B} \sum_{i=1}^{\gamma N} \tilde{C}_{L+k-i} F_{L+k-i}$. We immediately get

$$\sum_{\ell=L}^{L+k} \tilde{C}_\ell \tilde{F}_\ell \leq \tilde{C}_{L+k} F_{L+k} + \sum_{\ell=L}^{L+k-1} \tilde{C}_\ell F_\ell .$$

Now, it only remains to bound $\mathbf{E} \left[\sum_{\ell=L}^{L+\sqrt{\gamma}N-1} \tilde{C}_\ell \tilde{F}_\ell \right] = \sum_{\ell=L}^{L+\sqrt{\gamma}N-1} \mathbf{E} [\tilde{C}_\ell] \mathbf{E} [\tilde{F}_\ell]$. By definition $\mathbf{E} [\tilde{C}_\ell] = \frac{B}{\gamma N}$ for every ℓ , it is enough to show that $\frac{1}{\sqrt{\gamma}N} \sum_{\ell=L}^{L+\sqrt{\gamma}N-1} \mathbf{E} [\tilde{F}_\ell] \geq \left(\frac{1}{2} - \sqrt{\gamma} - \frac{1}{4\sqrt{\gamma}N} \right)$.

Define $a_\ell = \mathbf{E} [\tilde{F}_\ell]$. We have $a_\ell \geq 1 - \frac{1}{B} \sum_{i=1}^{\gamma N} \mathbf{E} [\tilde{C}_{\ell-i}] a_{\ell-i} = 1 - \frac{1}{\gamma N} \sum_{i=1}^{\gamma N} a_{\ell-i}$ for $\ell \geq L$ and $a_\ell = 1$ for $\ell < L$. Averaging over the rounds $L, \dots, L + \sqrt{\gamma}N - 1$ we get

$$\begin{aligned} \frac{1}{\sqrt{\gamma}N} \sum_{\ell=L}^{L+\sqrt{\gamma}N-1} a_\ell &\geq \frac{1}{\sqrt{\gamma}N} \sum_{\ell=L}^{L+\sqrt{\gamma}N-1} \left(1 - \frac{1}{\gamma N} \sum_{i=1}^{\gamma N} a_{\ell-i} \right) \\ &= \frac{1}{\sqrt{\gamma}N} \left(\sqrt{\gamma}N - \frac{1}{\gamma N} \sum_{\ell=L}^{L+\sqrt{\gamma}N-1} \sum_{i=1}^{\gamma N} a_{\ell-i} \right) . \end{aligned}$$

Here, we change the order of summation and split the inner sum into two parts which we bound separately

$$\frac{1}{\sqrt{\gamma}N} \sum_{\ell=L}^{L+\sqrt{\gamma}N-1} a_\ell \geq \frac{1}{\sqrt{\gamma}N} \left(\sqrt{\gamma}N - \frac{1}{\gamma N} \sum_{i=1}^{\gamma N} \left(\sum_{\ell=L}^{L-1+i} a_{\ell-i} + \sum_{\ell=L+i}^{L+\sqrt{\gamma}N-1} a_{\ell-i} \right) \right) .$$

Since $a_{\ell-i} \leq 1$, we bound the first sum with $\sum_{\ell=L}^{L-1+i} a_{\ell-i} \leq i$. Furthermore, as $a_\ell \geq 0$ for all ℓ , we can pad the second sum so that $\sum_{\ell=L+i}^{L+\sqrt{\gamma}N-1} a_{\ell-i} \leq \sum_{\ell=L}^{L+\sqrt{\gamma}N-1} a_\ell$, thus we have $\frac{1}{\gamma N} \sum_{i=1}^{\gamma N} \sum_{\ell=L+i}^{L+\sqrt{\gamma}N-1} a_{\ell-i} \leq \sum_{\ell=L}^{L+\sqrt{\gamma}N-1} a_\ell$.

We use both bounds and get

$$\frac{1}{\sqrt{\gamma}N} \sum_{\ell=L}^{L+\sqrt{\gamma}N-1} a_\ell \geq \frac{1}{\sqrt{\gamma}N} \left(\sqrt{\gamma}N - \frac{1}{\gamma N} \sum_{i=1}^{\gamma N} i - \sum_{\ell=L}^{L+\sqrt{\gamma}N-1} a_\ell \right) .$$

This implies

$$\frac{1}{\sqrt{\gamma}N} \sum_{\ell=L}^{L+\sqrt{\gamma}N-1} a_\ell = \frac{1}{2} - \frac{\gamma N + 1}{4\sqrt{\gamma}N} = \frac{1}{2} - \frac{\sqrt{\gamma}}{4} - \frac{1}{4\sqrt{\gamma}N} . \quad \blacktriangleleft$$

Now we have all parts required to prove the theorem.

Proof of Theorem 2. If in round ℓ an item is tentatively selected, let V_ℓ denote its value. Otherwise set $V_\ell = 0$. Our task is to bound the sum of $\mathbf{E} [V_\ell F_\ell] = \mathbf{E} [V_\ell \tilde{C}_\ell F_\ell]$.

Fixing which items come in rounds $1, \dots, \ell$ fixes the set $S^{(\ell)}$. As any order among these items and the respective dummy items is still equally likely, the item coming in round ℓ can be considered being drawn uniformly at random. This way, by Lemma 1, we have for all $\ell \geq 2\sqrt{\frac{\gamma}{B}}N$

$$\mathbf{E} [V_\ell] \geq \frac{1}{\ell} \mathbf{E} \left[\sum_{j \in S^{(\ell)}} v_j \right] \geq \left(1 - 9\sqrt{\frac{1}{\frac{\ell}{N} \cdot \frac{B}{\gamma}}} \right) \frac{1}{N} \frac{1 - \frac{1}{2}\sqrt{\frac{\gamma}{B}}}{1 + \gamma} \sum_{j \in \text{OPT}^*} v_j .$$

Next, observe that V_ℓ given that $\tilde{C}_\ell = 1$ is independent of F_ℓ . This is due to the fact that the algorithm is comparison based. For this reason, the course of events in rounds $1, \dots, \ell - 1$ is independent of the identity of the item from $S^{(\ell)}$ that actually arrives in round ℓ . Therefore the events leading up to round ℓ are identical, although they might involve different items. Also exploiting that $V_\ell = 0$ if $\tilde{C}_\ell = 0$, we get

$$\begin{aligned} \mathbf{E}[V_\ell \tilde{C}_\ell F_\ell] &= \Pr[\tilde{C}_\ell = 1, F_\ell = 1] \mathbf{E}[V_\ell \mid \tilde{C}_\ell = 1, F_\ell = 1] \\ &= \Pr[\tilde{C}_\ell = 1, F_\ell = 1] \mathbf{E}[V_\ell \mid \tilde{C}_\ell = 1] = \frac{\Pr[\tilde{C}_\ell = 1, F_\ell = 1]}{\Pr[\tilde{C}_\ell = 1]} \mathbf{E}[V_\ell]. \end{aligned}$$

To get the bound, we split the input sequence into blocks of length $\sqrt{\gamma}N$ and apply Lemma 3 on each of these blocks.

$$\begin{aligned} \mathbf{E}\left[\sum_{j \in \text{ALG}} v_j\right] &\geq \sum_{k=2}^{\lfloor \frac{1}{\sqrt{\gamma}} \rfloor - 1} \sum_{\ell=k\sqrt{\gamma}N+1}^{(k+1)\sqrt{\gamma}N} \mathbf{E}[V_\ell \tilde{C}_\ell F_\ell] \\ &\geq \sum_{k=2}^{\lfloor \frac{1}{\sqrt{\gamma}} \rfloor - 1} \sum_{\ell=k\sqrt{\gamma}N+1}^{(k+1)\sqrt{\gamma}N} \frac{\Pr[\tilde{C}_\ell = 1, F_\ell = 1]}{\Pr[\tilde{C}_\ell = 1]} \left(1 - 9\sqrt{\frac{1}{\frac{\ell}{N} \cdot \frac{B}{\gamma}}}\right) \frac{1}{N} \frac{1 - \frac{1}{2}\sqrt{\frac{\gamma}{B}}}{1 + \gamma} \sum_{j \in \text{OPT}^*} v_j \\ &\geq \left(\frac{1}{2} - \frac{7}{4}\sqrt{\gamma} - \frac{37}{4}\sqrt{\frac{\gamma}{B}} - \frac{\gamma}{2} - \frac{1}{4\sqrt{\gamma}N}\right) \sum_{j \in \text{OPT}^*} v_j. \end{aligned}$$

Details on the calculations can be found in the full version. \blacktriangleleft

2.3 Improved Analysis for Large Capacities

For the same algorithm, we can show a better competitive ratio if B is large, converging to 1 asymptotically.

► **Theorem 4.** *Algorithm 1 is at least $\left(1 - \frac{4}{\sqrt{B}} - \frac{41}{2}\sqrt{\frac{\gamma}{B}} - 3\gamma - O\left(\frac{1}{B}\right)\right)$ -competitive for the temp secretary problem with duration γ for all items.*

The proof of this theorem is very similar to the proof of Theorem 2. Again, we will bound the number of rounds in which $\tilde{C}_\ell = 1$ but $F_\ell = 0$. The main difference is that we consider blocks of γN rounds each. In this case, the duration of an item corresponds to the length of the block. Therefore, no more than B items can be feasibly selected in any block.

If B items are selected at the beginning of such a block, then it is feasible to select one item for every item that times out. We use this concept and construct a symmetric random walk. The maximum of this random walk upper bounds the number of failure events, in which the algorithm performs a tentative selection before sufficiently many previous items have timed out.

In contrast to Lemma 3, the expected ratio of failure events to successful selections of item is decreasing in B and therefore our competitive ratio in Theorem 4 tends to 1 as B increases.

► **Lemma 5.** *The expected number of rounds $\ell \in \{L, \dots, L + \gamma N - 1\}$ in which $\tilde{C}_\ell = 1$ but $F_\ell = 0$ is $|\{\ell \mid \tilde{C}_\ell = 1 \wedge F_\ell = 0\}| \leq \sqrt{B} + \frac{\pi^2}{6}\sqrt{B} + 2\sqrt{\frac{4B}{3\pi}} + O(1) \leq 4\sqrt{B} + O(1)$.*

Proof. We claim that the number of rounds $\ell \in \{L, \dots, L + \gamma N - 1\}$ for which $\tilde{C}_\ell = 1$ but $F_\ell = 0$ is bounded by

$$Q := \max_{L \leq \ell < L + \gamma N} \left(\sum_{r=L}^{\ell} \tilde{C}_r - \sum_{r=L-\gamma N}^{\ell-\gamma N} \tilde{C}_r \right) + \left| \sum_{r=L-\gamma N}^{L-1} \tilde{C}_r - B \right|.$$

To show this, we define an alternative sequence $(\tilde{C}'_r)_{r \in [N]}$ by setting $\tilde{C}_r = \tilde{C}'_r$ for every r except for the first Q occurrences of $\tilde{C}_r = 1$ after L , where we set $\tilde{C}'_r = 0$. Consider an $\ell \geq L$ such that $\tilde{C}'_\ell = 1$. Now observe that

$$\sum_{r=\ell-\gamma N}^{\ell} \tilde{C}'_r = \sum_{r=\ell-\gamma N}^{\ell} \tilde{C}_r - Q \leq \sum_{r=L-\gamma N}^{L-1} \tilde{C}_r - \left| \sum_{r=L-\gamma N}^{L-1} \tilde{C}_r - B \right| \leq B .$$

This implies that $\sum_{r=L}^{L+\gamma N-1} \tilde{C}_r F_r \geq \sum_{r=L}^{L+\gamma N-1} \tilde{C}'_r$ because every case of $\tilde{C}_r = 1$ but $F_r = 0$ can be matched to a case where $\tilde{C}_r = 1$ but $\tilde{C}'_r = 0$. So, to show the lemma, it only remains to show that $\mathbf{E}[Q] = O(\sqrt{B})$.

First, observe that $\sum_{r=L-\gamma N}^{L-1} \tilde{C}_r$ is drawn from a binomial distribution with γN trials and probability $\frac{B}{\gamma N}$. Therefore, its expectation is $\mu = B$ and its standard deviation is $\sigma = \sqrt{B - \frac{B^2}{\gamma N}} \leq \sqrt{B}$. Thus by Chebyshev inequality, we get

$$\mathbf{E} \left[\left| \sum_{r=L-\gamma N}^{L-1} \tilde{C}_r - B \right| \right] \leq \sigma + \sum_{k=1}^{\infty} \mathbf{E} \left[\left| \sum_{r=L-\gamma N}^{L-1} \tilde{C}_r - B \right| \geq k\sigma \right] \sigma \leq \sigma + \sum_{k=1}^{\infty} \frac{\sigma}{k^2} = \sqrt{B} + \frac{\pi^2}{6} \sqrt{B} .$$

So, it only remains to show that

$$\mathbf{E} \left[\max_{L \leq \ell < L+\gamma N} \left(\sum_{r=L}^{\ell} \tilde{C}_r - \sum_{r=L-\gamma N}^{\ell-\gamma N} \tilde{C}_r \right) \right] \leq 2\sqrt{\frac{4B}{3\pi}} + O(1) .$$

Let $\tilde{C}''_\ell \in \{-1, 0, 1\}$ be a random variable with $\tilde{C}''_\ell = \tilde{C}_\ell - \tilde{C}_{\ell-\gamma N}$. Each \tilde{C}''_ℓ takes the values 1 and -1 with probability $p = \frac{B}{\gamma N}(1 - \frac{B}{\gamma N})$ each and 0 with the remaining probability $1-2p = \left(1 - \frac{2B}{\gamma N} + 2(\frac{B}{\gamma N})^2\right)$. Furthermore, because the \tilde{C}_ℓ random variables are independent, $\tilde{C}''_L, \dots, \tilde{C}''_{L+\gamma N-1}$ also are.

We interpret this random process on the \tilde{C}''_ℓ as a random walk of length γN that moves up or down with probability p and stays the same with probability $1 - 2p$. In the next part of the proof, we are going to show that the maximal deviation of this random walk is in $\Theta(\sqrt{B})$. To this end, we condition our random walk on the number of zeros that occur. The remaining random walk is symmetric, thus results from the literature apply.

It has been shown that, for a symmetric random walk that starts in position 0 and does k steps, the expected final position is $\mathbf{E}[S_k] = \sqrt{\frac{2k}{3\pi}} + O(k^{-\frac{1}{2}})$ [14]. Furthermore, it is well known that the expected maximal deviation during such a random walk is $\mathbf{E}[M_k] \leq 2\mathbf{E}[S_k]$. Now, let K be the number of times $\tilde{C}''_r \neq 0$ for $r \in \{L, \dots, L + \gamma N - 1\}$. Then we have

$$\begin{aligned} \mathbf{E} \left[\max_{L \leq \ell < L+\gamma N} \left(\sum_{r=L}^{\ell} \tilde{C}_r - \sum_{r=L-\gamma N}^{\ell-\gamma N} \tilde{C}_r \right) \right] &\leq \sum_{k=0}^{\gamma N} \mathbf{E}[M_k] \cdot \mathbf{Pr}[K = k] \\ &\leq \sum_{k=0}^{\gamma N} 2\mathbf{E}[S_k] \cdot \mathbf{Pr}[K = k] \leq \mathbf{E} \left[2\sqrt{\frac{2K}{3\pi}} + O(K^{-\frac{1}{2}}) \right] \leq 2\sqrt{\frac{2\mathbf{E}[K]}{3\pi}} + O(1) \\ &= 2\sqrt{\frac{4B}{3\pi}} + O(1) . \end{aligned}$$

We use the same proof structure as in the previous proof of Theorem 2, but now we replace Lemma 3 with Lemma 5.

Proof of Theorem 4. Again, if in round ℓ an item is tentatively selected, let V_ℓ denote its value. Otherwise set $V_\ell = 0$. By the same arguments as in the proof of Theorem 2, we have

$$\mathbf{E} [V_\ell \tilde{C}_\ell F_\ell] = \frac{\Pr [\tilde{C}_\ell = 1, F_\ell = 1]}{\Pr [\tilde{C}_\ell = 1]} \mathbf{E} [V_\ell] .$$

To get the bound, we split the input sequence into blocks of length γN and ignore the blocks in which there is a round for which Lemma 1 does not hold.

$$\begin{aligned} \mathbf{E} \left[\sum_{j \in \text{ALG}} v_j \right] &\geq \sum_{k=\lfloor \frac{1}{2\sqrt{\frac{1}{\gamma B}}} \rfloor}^{\lfloor \frac{1}{\gamma} \rfloor - 1} \sum_{\ell=k\gamma N+1}^{(k+1)\gamma N} \mathbf{E} [V_\ell \tilde{C}_\ell F_\ell] \\ &\geq \sum_{k=\lfloor \frac{1}{2\sqrt{\frac{1}{\gamma B}}} \rfloor}^{\lfloor \frac{1}{\gamma} \rfloor - 1} \sum_{\ell=k\gamma N+1}^{(k+1)\gamma N} \frac{\Pr [\tilde{C}_\ell = 1, F_\ell = 1]}{\Pr [\tilde{C}_\ell = 1]} \left(1 - 9\sqrt{\frac{1}{\frac{\ell}{N} \cdot \frac{B}{\gamma}}} \right) \cdot \frac{1}{N} \frac{1 - \frac{1}{2}\sqrt{\frac{\gamma}{B}}}{1 + \gamma} \sum_{j \in \text{OPT}^*} v_j \\ &\geq \left(1 - \frac{4}{\sqrt{B}} - \frac{41}{2} \sqrt{\frac{\gamma}{B}} - 3\gamma - O\left(\frac{1}{B}\right) \right) \sum_{j \in \text{OPT}^*} v_j . \end{aligned}$$

The missing details on the calculations are included in the full version. ◀

3 The Temp Secretary Problem with Packing Constraints

Next, we turn to the temp secretary with general linear packing constraints. This generalizes the timed cardinality constraint of the temp secretary problem towards multiple knapsack constraints. Therefore we can model, e.g., production capacities of different types. Now, the problem is not about selecting a set of best candidates, but a set of contracts with different resource demands such that the value of the selected contracts is maximized and all demands are fulfilled at any point in time.

We assume that the items, or possible contracts, that arrive over time are variables of a packing LP that have to be set immediately and irrevocably at time of arrival. In more detail, we assume that an adversary defines an $n \times m$ constraint matrix A , a capacity vector b , and an objective function vector v . Again, for each variable x_j an arrival time is drawn independently uniformly at random from $[0, 1]$. We now have to find an assignment $\hat{x}_j \in \{0, 1\}$ for all $j \in [n]$, with the property that for every $t \in [0, 1]$, the set of variables that arrive between t and $t + \gamma$ solve the packing LP. That is, for every t , we need $Ax' \leq b$, where $x'_j = \hat{x}_j$ if $t_j \in [t, t + \gamma]$ and 0 otherwise. The objective is to maximize $v^T \hat{x}$. So, \hat{x} represents the aggregate vector of which items are selected whenever they are present in the system. Therefore, the temp secretary problem can be captured by having only one constraint with all coefficients being 1 and the capacity being the cardinality bound. As the output of our algorithm will be integral, this algorithm also solved the temp secretary problem.

We assume that the matrix A is sparse. That is, each column of A contains at most d non-zero entries. Furthermore, we assume that the capacities in each constraint are large compared to the respective coefficients in A . Our bound will depend on the capacity ratio B , defined as $B = \min_{i \in [m]} \frac{b_i}{\max_{j \in [n]} a_{i,j}}$. When modeling the temp secretary problem as above, this capacity ratio coincides with the capacity bound.

Again, we use a relaxation without temporal constraints like the one in Section 2. In this case, it reads $\max v^T x$ s.t. $Ax \leq \lceil \frac{1}{\gamma} \rceil b$, $0 \leq x_j \leq 1$ for all $j \in [n]$. While our algorithm chooses online which items to select and which to reject subject to the temporal constraints,

Algorithm 2: Scaling Algorithm for packing constraints

```

for every arriving item  $j \in [n]$  do
  Set  $t := \tau_j$ ; // arrival time of  $j$ 
  Let  $x^{(t)}$  be an optimal solution to the LP  $\max v^T x$  s.t.  $Ax \leq t(1 - \epsilon)\frac{b}{\gamma}$ ,  $0 \leq x_{j'} \leq 1$ 
  for  $j' \in U_{\leq t}$ ; // optimal offline solution
   $\hat{x}_{j'}^{(t)} = \begin{cases} 1, & \text{with prob. } x_{j'}^{(t)} \text{ if } j' = j; \\ 0, & \text{otherwise;} \end{cases}$  // randomized rounding
  if  $\hat{x} + \hat{x}^{(t)}$  is feasible with respect to temporal constraints then
  | Set  $\hat{x} := \hat{x} + \hat{x}^{(t)}$ ; // make allocation to  $j$  permanent

```

our point of comparison is the best fractional solution to this LP relaxation. We use an algorithm similar to the one presented in [17] to solve the relaxed problem. We scale down the aggregated constraints slightly and solve the resulting linear packing problem. Next, we use randomized rounding to transform the fractional solution into our integral, tentative solution. At this point, the algorithm behaves exactly like the one in Section 2. If the item that just arrived is part of the tentative solution and it is feasible to select it, then the algorithm adds it to the online solution. In contrast to the analysis in [17], we have to show that temporal constraints are likely fulfilled although the algorithm only operates on the relaxed ones.

To define the algorithm, let $U_{\leq t}$ be the set of variables $U_{\leq t} \subseteq [n]$ that arrive before time t . Let $\epsilon = \sqrt{6 \frac{(1 + \log d + \log B)}{B}}$.

This algorithm can also be extended to the case in which there are not only timed constraints but also global ones. Additionally, it can be applied when multiple variables arrive at a time like in [17]. We omit these generalizations because the techniques are identical to the ones presented here, but correct notation gets a lot more involved.

► **Theorem 6.** *Algorithm 2 is $\frac{1}{1+\gamma} - O\left(\sqrt{\frac{1+\log d+\log B}{B}}\right)$ -competitive for the temp secretary problem with linear packing constraints.*

Please note that this algorithm is invariant to scaling constraints and therefore we can assume without loss of generality that $\max_{j \in [n]} a_{i,j} \leq 1$ and $b_i \geq B$ for every constraint $i \in [m]$.

For the proof, we will first bound the expected consumption of the tentative solution of the algorithm. Then we will use a Chernoff bound to bound the probability that last *if*-clause of the algorithm is violated for a single constraint. Finally, we will aggregate the probabilities for all constraints in the relaxation.

We discretize time, like in Section 2, with arbitrary precision such that every discrete time interval only contains a single item of the input. We have $N \gg n$ rounds, spanning time $\frac{1}{N}$ each. For each of the $N - n$ rounds in which no variable arrives, we introduce a dummy variable with all coefficients zero. These N variables can now be considered being assigned to the rounds by a uniformly drawn permutation. In the proofs, we overload notation and write $x^{(\ell)} = x^{(t)}$ and $U_{\leq \ell} = U_{\leq t}$ if t lies within round ℓ .

► **Lemma 7.** *Let $U_{< \ell} \subseteq [n]$ be the set of items that arrive before round ℓ . Then, conditioned on this set, the sum of previous tentative allocations violates any constraint $i \in [m]$ with*

probability at most

$$\Pr \left[\left(\sum_{\ell'=\ell-\gamma N}^{\ell-1} A\hat{x}^{(\ell')} \right)_i > b_i - 1 \mid U_{<\ell} \right] \leq \frac{1}{dB}$$

$$\text{if } \epsilon = \sqrt{6 \frac{1+\log d+\log B}{B}} \leq \frac{1}{2}.$$

We describe the capacity used by the tentative selection through random variables $X_{\ell'} = (A\hat{x}^{(\ell')})_i$. These random variables are not independent but 1-correlated as defined by Panconesi and Srinivasan [24] and this allows us to apply a Chernoff bound to proof the lemma. Details can be found in the full version.

Proof of Theorem 6. We can assume without loss of generality that $\sqrt{6 \frac{1+\log d+\log B}{B}} \leq \frac{1}{2}$ because otherwise the theorem statement follows trivially.

First, we bound the value of the tentative allocation performed in round $\ell \geq 2\sqrt{\frac{1+\ln d}{B}}N$ using a result in [17]. Compared to [17] our non-temporal relaxation is scaled down by an additional factor of $1 - \epsilon$ this gives us

$$\mathbf{E} \left[v^T x^{(\ell)} \right] \geq \left(1 - 9 \sqrt{\frac{1}{\frac{\ell}{N} \cdot (1 - \epsilon) \frac{B}{\gamma}}} \right) \frac{\ell}{N} (1 - \epsilon) \cdot \max_{x: Ax \leq b/\gamma, 0 \leq x_j \leq 1 \text{ for all } j \in [m]} v^T x$$

Letting x^* denote the optimal solution to the relaxation, we also have

$$\max_{x: Ax \leq b/\gamma, 0 \leq x_j \leq 1 \text{ for all } j \in [m]} v^T x \geq \frac{\frac{1}{\gamma}}{\lceil \frac{1}{\gamma} \rceil} v^T x^* \geq \frac{1}{1 + \gamma} v^T x^* .$$

So, this implies

$$\mathbf{E} \left[v^T x^{(\ell)} \right] \geq \left(1 - 9 \sqrt{\frac{1}{\frac{\ell}{N} \cdot (1 - \epsilon) \frac{B}{\gamma}}} \right) \frac{\ell}{N} \frac{1 - \epsilon}{1 + \gamma} v^T x^* .$$

Observe that this outcome only depends on the set $U_{\leq \ell}$ but not the order in this set. Therefore, the variable that arrives in round ℓ can be considered being drawn uniformly from $U_{\leq \ell}$. This way, we get

$$\mathbf{E} \left[v^T \hat{x}^{(\ell)} \mid U_{\leq \ell} \right] = \frac{1}{\ell} \mathbf{E} \left[v^T x^{(\ell)} \mid U_{\leq \ell} \right]$$

Note that these outcomes only depend on the sets $U_{<\ell}$ and $U_{\leq \ell}$ but not the order within $U_{<\ell}$.

To bound the probability of feasibility, we will use Lemma 7, conditioning on the set $U_{<\ell}$ and $U_{\leq \ell}$. Let j be the index of the variable arriving in round ℓ . Taking a union bound over all $\leq d$ constraints in which variable j has non-zero coefficients, we get

$$\begin{aligned} & \Pr [\text{it is feasible to set } \hat{x}_j = 1 \mid U_{<\ell}, U_{\leq \ell}] \\ & \geq 1 - d \cdot \sum_{i: a_{i,j} > 0} \Pr \left[\left(\sum_{\ell'=\ell-\gamma N}^{\ell-1} A\hat{x}^{(\ell')} \right)_i > b_i - 1 \mid U_{<\ell} \right] \geq \left(1 - \frac{1}{B} \right) . \end{aligned}$$

Overall, the expected value of the allocation performed in round ℓ is at least

$$\mathbf{E} \left[v^T \hat{x}^{(\ell)} \right] \geq \left(1 - 9 \sqrt{\frac{1}{\frac{\ell}{N} \cdot (1 - \epsilon) \frac{B}{\gamma}}} \right) \frac{1}{N} \frac{1 - \epsilon}{1 + \gamma} \left(1 - \frac{1}{B} \right) v^T x^* .$$

Algorithm 3: Scaling Algorithm for different lengths and capacity B

```

for every arriving item  $j$  do
  Set  $t := \tau_j$ ,  $U_{\leq t} :=$  items arrived so far;
  Let  $S^{(t)} := \text{GREEDYROUNDUP}(U_{\leq t}, \alpha B)$ ;
  if  $j \in S^{(t)}$  then // if among best items
    if  $S \cup \{j\}$  is a feasible schedule then // and if feasible
      Set  $S := S \cup \{j\}$ ; // then select  $j$ 

```

Taking the sum of all these bounds, we get

$$\begin{aligned}
 \mathbf{E} [v^T \hat{x}] &\geq \sum_{\ell=2\sqrt{\frac{1+\ln d}{B}}N}^N \left(1 - 9\sqrt{\frac{1}{\frac{\ell}{N} \cdot (1-\epsilon)\frac{B}{\gamma}}}\right) \frac{1}{N} \cdot \frac{1-\epsilon}{1+\gamma} \left(1 - \frac{1}{B}\right) v^T x^* \\
 &\geq \frac{1 - O(\epsilon)}{1+\gamma} v^T x^* .
 \end{aligned}$$

4 The Temp Secretary Problem with Different Lengths

We generalize the Temp Secretary Problem towards different item durations $\lambda_j \leq \gamma$ for all items $j \in [n]$. To define the relaxation, we use the fact that pointwise $\sum_{j \in \text{OPT}} \lambda_j \leq B(1+\gamma)$. This is due to the fact that an item j selected at time $\tau_j = 1$ will be active until time $1 + \lambda_j$. Therefore, let OPT^* denote the optimal solution to this knapsack problem with profits v_j . Due to the knapsack nature of the problem, the algorithm cannot be purely comparison-based anymore. Instead, whenever an item arrives, we compute an approximate knapsack solution and tentatively select the item if it is included in this solution. It will be crucial that these solutions only slowly change when adding or removing items. This is why, there is no obvious generalization of our algorithm and analysis to general packing LPs. In more detail, for $U' \subseteq [n]$ and $\Lambda > 0$ we define $\text{GREEDYROUNDUP}(U', \Lambda)$ as the set $U'' \subseteq U'$ which we get by ordering the items in U' in non-increasing order $\frac{v_j}{\lambda_j}$ and taking the minimal prefix such that $\sum_{j \in U''} \lambda_j \geq \Lambda$.

► **Theorem 8.** *With $\alpha = 1/2$, Algorithm 3 is at least $\frac{1}{4} - \Theta(\sqrt{\gamma})$ -competitive for the temp secretary problem with arbitrary durations $\lambda_j \leq \gamma$.*

We use the same proof structure like in Theorem 2. Here Lemma 9 gives a bound on the expected value of the tentative solution. The proof of this lemma is very similar to the proof of Lemma 1 and can be found in the full version.

► **Lemma 9.** *For $\ell \geq 2\sqrt{\frac{\gamma}{B}}N$*

$$\mathbf{E} \left[\sum_{j \in S^{(\ell)}} v_j \right] \geq \alpha \left(1 - 9\sqrt{\frac{1}{\frac{\ell}{N} \frac{B}{\gamma}}}\right) \frac{\ell}{N} \frac{1}{1+\gamma} \sum_{j \in \text{OPT}^*} v_j .$$

The main difference to previous proofs in this paper is the fact that we cannot bound the probability of a selection for a fixed round because it depends on the lengths of the items that arrive up to this round. Fortunately, we can bound the probability that a selection would be feasible by the following lemma.

► **Lemma 10.** *Conditioning on any set of arrivals in rounds $1, \dots, \ell - 1$, the probability that an item can be feasibly selected in round ℓ is at least $1 - \alpha - \frac{3\gamma N}{\ell - \gamma N}$.*

As a key idea, we show that the probability of a selection does not significantly change between rounds.

► **Lemma 11.** *Let $H_{i,\ell}$ be the 0/1 indicator if an item with duration at least i rounds is tentatively selected in rounds ℓ . It holds that $\Pr[H_{i,\ell-i} = 1] \leq \Pr[H_{i,\ell} = 1] + \frac{\alpha B + 1}{\ell - i}$.*

Proof. First, we observe that an item j is tentatively selected in round ℓ' if it is contained in $S^{(\ell')}$ and arrives in round ℓ' . As the set $S^{(\ell')}$ only depends on the items arriving in rounds $1, \dots, \ell'$ but not on their order, the probability of j being tentatively selected is exactly $\frac{1}{\ell'} \Pr[j \in S^{(\ell')}]$. Therefore, to prove this lemma, we will compare the sets $S^{(\ell-i)}$ and $S^{(\ell)}$.

Instead of making statements about the set $S^{(\ell-i)}$ directly, we instead use a set \tilde{S} defined as $\tilde{S} := \text{GREEDYROUNDUP}(U_{\leq \ell-i}, \alpha \frac{\ell}{N} B)$. Note that by definition of GREEDYROUNDUP , this set \tilde{S} is a superset of $S^{(\ell-i)} = \text{GREEDYROUNDUP}(U_{\leq \ell-i}, \alpha \frac{\ell-i}{N} B)$.

Let now $\mathcal{E}_{i,\ell-i}^*$ be the event that an item of length at least i rounds arrives in round $\ell - i$ that is contained in $\tilde{S} \cap S^{(\ell)}$, and let $\mathcal{E}_{i,\ell-i}^-$ be the event that an item of length at least i rounds arrives in round $\ell - i$ that is contained in $\tilde{S} \setminus S^{(\ell)}$. As $S^{(\ell-i)}$ is contained in \tilde{S} , we have $\Pr[H_{i,\ell-i} = 1] \leq \Pr[\mathcal{E}_{i,\ell-i}^*] + \Pr[\mathcal{E}_{i,\ell-i}^-]$.

To bound the first probability, we use that every item in $S^{(\ell)}$ has already arrived in round $\ell - i$ with probability $\frac{\ell-i}{\ell}$. Every such item arrives exactly in round $\ell - i$ with probability $\frac{1}{\ell-i}$. Therefore we have $\Pr[\mathcal{E}_{i,\ell-i}^*] = \Pr[H_{i,\ell} = 1]$.

For the second probability, observe that \tilde{S} contains all elements of $S^{(\ell)}$ that have arrived by round $\ell - i$, i.e., $\tilde{S} \supseteq S^{(\ell)} \cap U_{\leq \ell-i}$. Also if $\sum_{j \in S^{(\ell)}} \lambda_j < \alpha \frac{\ell}{N} B - \gamma$, it has to be $S^{(\ell)} = U_{\leq \ell}$. Therefore, $\tilde{S} = S^{(\ell)}$ and nothing has to be shown. This implies that $\mathbf{E}[\sum_{j \in \tilde{S} \cap S^{(\ell)}} \lambda_j] \geq \frac{\ell-i}{\ell} \alpha \frac{\ell}{N} B = \alpha \frac{\ell-i}{N} B$ because the probability for each item in $S^{(\ell)}$ to arrive until round $\ell - i$ is $\frac{\ell-i}{\ell}$. Therefore, the λ_j values for all but one item in $\tilde{S} \setminus S^{(\ell)}$ add up to at most $\alpha B \frac{i}{N}$.

Let K denote the number of items in $\tilde{S} \setminus S^{(\ell)}$ of length at least i blocks. On the one hand, K is bounded by $\mathbf{E}[K] \leq \alpha B + 1$ due to the above considerations. On the other hand, conditioning on K , we can bound the probability of $\Pr[\mathcal{E}_{i,\ell-i}^-]$ by $\Pr[\mathcal{E}_{i,\ell-i}^- \mid K = k] \leq \frac{k}{\ell-i}$.

Taking the expectation over K , we get $\Pr[\mathcal{E}_{i,\ell-i}^-] \leq \frac{\alpha B + 1}{\ell - i}$. ◀

Proof of Lemma 10. Using Lemma 11, we now have

$$\begin{aligned} \Pr[\text{round } \ell \text{ would be feasible} \mid U_{\leq \ell-1}] &\geq 1 - \frac{1}{B} \sum_{i=1}^{\gamma N} \mathbf{E}[H_{i,\ell-i} \mid U_{\leq \ell-1}] \\ &\geq 1 - \frac{1}{B} \sum_{i=1}^{\gamma N} \mathbf{E}[H_{i,\ell-1} \mid U_{\leq \ell-1}] - \frac{\alpha \gamma N + 1}{\ell - \gamma N}. \end{aligned}$$

Observe that $\sum_{i=1}^{\gamma N} \mathbf{E}[H_{i,\ell-1} \mid U_{\leq \ell-1}]$ is exactly the length in rounds of the tentative selection in round $\ell - 1$, counting no tentative selection as 0. This length is bounded by $\frac{1}{\ell-1} \mathbf{E}[\sum_{j \in S^{(\ell-1)}} \lambda_j N \mid U_{\leq \ell-1}]$. This sum can be bounded pointwise by $\sum_{j \in S^{(\ell-1)}} \lambda_j \leq \alpha \frac{\ell-1}{N} B + \gamma$. So we have $\sum_{i=1}^{\gamma N} \mathbf{E}[H_{i,\ell-1} \mid U_{\leq \ell-1}] \leq \alpha + \frac{\gamma N}{\ell-1}$.

This means $\Pr[\text{round } \ell \text{ would be feasible} \mid U_{\leq \ell-1}] \geq 1 - \alpha - \frac{(1+\alpha)\gamma N + 1}{\ell - \gamma N} \geq 1 - \alpha - \frac{3\gamma N}{\ell - \gamma N}$. ◀

The remaining proof follows essentially the pattern of the previous sections.

Proof of Theorem 8. We have already shown that

$$\Pr[\text{round } \ell \text{ would be feasible} \mid U_{\leq \ell-1}] \geq 1 - \alpha - \frac{3\gamma N}{\ell - \gamma N}.$$

By combining this bound with Lemma 9, we get that the expected value gained from round $\ell \geq 2\sqrt{\gamma}N$ is at least

$$\frac{1}{2\ell} \cdot \left(1 - 9\sqrt{\frac{\gamma}{\frac{\ell}{N}B}}\right) \frac{\ell}{N} \frac{\sum_{j \in \text{OPT}^*} v_j}{1 + \gamma} \cdot \left(\frac{1}{2} - \frac{3\gamma N}{\ell - \gamma N}\right).$$

Similar calculations like in the proof of Theorem 2 and $\alpha = \frac{1}{2}$ give a competitive ratio of

$$\begin{aligned} \mathbf{E} \left[\sum_{j \in \text{ALG}} v_j \right] &\geq \sum_{k=2}^{\frac{1}{\sqrt{\gamma}}-1} \sum_{\ell=k\sqrt{\gamma}N}^{(k+1)\sqrt{\gamma}N-1} \frac{1}{2} \left(1 - 9\sqrt{\frac{\gamma}{\frac{\ell}{N}B}}\right) \frac{1}{N} \frac{\sum_{j \in \text{OPT}^*} v_j}{1 + \gamma} \cdot \left(\frac{1}{2} - \frac{3\gamma N}{\ell - \gamma N}\right) \\ &\geq \left(\frac{1}{4} - 5\sqrt{\gamma} - \frac{3}{2}\gamma \ln\left(\frac{1}{\sqrt{\gamma}}\right)\right) \cdot \frac{1}{1 + \gamma} \sum_{j \in \text{OPT}^*} v_j. \end{aligned}$$

The missing details are included in the full version. ◀

5 Future Work

The area of online problems with stochastic arrivals and temporal constraints leaves many open research directions. First of all, only very few impossibility results are known in this model. It seems natural that bounds in the related random order model should generalize in some way. For example in the temp secretary problem with large capacities, it seems plausible that the competitive ratio cannot be better than $1 - \Omega\left(\sqrt{\frac{1}{B}}\right)$ like in [18], independent of γ . A small γ increases the overall capacity, but in every step the algorithm is still tightly restricted by the timed capacity B .

Furthermore, the results in this paper apply if the arrival times are each drawn independently uniformly from $[0, 1]$. This condition can be relaxed in several ways. Firstly, other distributions than the uniform one are of important interest. Although the algorithms in this paper do admit a reasonable generalization using quantiles, our analyses as they are do not extend. Secondly, it is also very interesting to weaken the assumption that arrival times are independent and identically distributed. There is only little work in related models [15, 8] and none for this particular setting.

Other fruitful directions could be the extension to other feasibility structures, such as (special classes of) matroid, and to other objective functions, such as submodular ones. Finally, it might also be interesting to let the algorithm decide the contract starting/finishing dates or its duration.

References

- 1 Shipra Agrawal and Nikhil R. Devanur. Fast algorithms for online stochastic convex programming. In *Proc. 26th Symp. Discr. Algorithms (SODA)*, pages 1405–1424, 2015. doi:10.1137/1.9781611973730.93.

- 2 Shipra Agrawal, Zizhuo Wang, and Yinyu Ye. A dynamic near-optimal algorithm for online linear programming. *Operations Research*, 62(4):876–890, 2014. doi:10.1287/opre.2014.1289.
- 3 Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. A knapsack secretary problem with applications. In *Proc. 10th Intl. Workshop Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pages 16–28, 2007. doi:10.1007/978-3-540-74208-1_2.
- 4 Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. Matroids, secretary problems, and online mechanisms. In *Proc. 18th Symp. Discr. Algorithms (SODA)*, pages 434–443, 2007. URL: <http://dl.acm.org/citation.cfm?id=1283383.1283429>.
- 5 MohammadHossein Bateni, Mohammad Taghi Hajiaghayi, and Morteza Zadimoghaddam. Submodular secretary problem and extensions. *ACM Trans. Algorithms*, 9(4):32, 2013. doi:10.1145/2500121.
- 6 Nikhil R. Devanur, Kamal Jain, Balasubramanian Sivan, and Christopher A. Wilkens. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In *Proc. 12th Conf. Econom. Comput. (EC)*, pages 29–38, 2011. doi:10.1145/1993574.1993581.
- 7 Eugene B. Dynkin. The optimum choice of the instant for stopping a markov process. In *Sov. Math. Dokl*, volume 4, pages 627–629, 1963.
- 8 Hossein Esfandiari, Nitish Korula, and Vahab S. Mirrokni. Online allocation with traffic spikes: Mixing adversarial and stochastic models. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC'15, Portland, OR, USA, June 15-19, 2015*, pages 169–186, 2015. doi:10.1145/2764468.2764536.
- 9 Moran Feldman, Joseph Naor, and Roy Schwartz. Improved competitive ratios for submodular secretary problems (extended abstract). In *Proc. 14th Intl. Workshop Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pages 218–229, 2011. doi:10.1007/978-3-642-22935-0_19.
- 10 Moran Feldman, Ola Svensson, and Rico Zenklusen. A simple $O(\log \log(\text{rank}))$ -competitive algorithm for the matroid secretary problem. In *Proc. 26th Symp. Discr. Algorithms (SODA)*, pages 1189–1201, 2015. doi:10.1137/1.9781611973730.79.
- 11 Amos Fiat, Ilya Gorelik, Haim Kaplan, and Slava Novgorodov. The temp secretary problem. In *Proc. 23rd European Symp. Algorithms (ESA)*, pages 631–642, 2015. doi:10.1007/978-3-662-48350-3_53.
- 12 M. Gardner. *Scientific american*, 1960.
- 13 Anupam Gupta and Marco Molinaro. How experts can solve LPs online. In *Proc. 22nd European Symp. Algorithms (ESA)*, pages 517–529, 2014. doi:10.1007/978-3-662-44777-2_43.
- 14 Edward G. Coffman Jr., Philippe Flajolet, Leopold Flatto, and Micha Hofri. The maximum of a random walk and its application to rectangle packing. *Probability in the Engineering and Informational Sciences*, 12(3):373–386, 1998. doi:10.1017/S0269964800005258.
- 15 Thomas Kesselheim, Robert D. Kleinberg, and Rad Niazadeh. Secretary problems with non-uniform arrival order. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 879–888, 2015. doi:10.1145/2746539.2746602.
- 16 Thomas Kesselheim, Klaus Radke, Andreas Tönnis, and Berthold Vöcking. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In *Proc. 21st European Symp. Algorithms (ESA)*, pages 589–600, 2013. doi:10.1007/978-3-642-40450-4_50.

- 17 Thomas Kesselheim, Klaus Radke, Andreas Tönnis, and Berthold Vöcking. Primal beats dual on online packing lps in the random-order model. In *Proc. 46th Symp. Theory of Computing (STOC)*, pages 303–312, 2014. doi:10.1145/2591796.2591810.
- 18 Robert D. Kleinberg. A multiple-choice secretary algorithm with applications to online auctions. In *Proc. 16th Symp. Discr. Algorithms (SODA)*, pages 630–631, 2005. URL: <http://dl.acm.org/citation.cfm?id=1070432.1070519>.
- 19 Nitish Korula and Martin Pál. Algorithms for secretary problems on graphs and hypergraphs. In *Proc. 36th Intl. Coll. Autom. Lang. Program. (ICALP)*, pages 508–520, 2009. doi:10.1007/978-3-642-02930-1_42.
- 20 Oded Lachish. $O(\log \log \text{rank})$ competitive ratio for the matroid secretary problem. In *Proc. 55th Symp. Foundations of Computer Science (FOCS)*, pages 326–335, 2014. doi:10.1109/FOCS.2014.42.
- 21 Denis V Lindley. Dynamic programming and decision theory. *Applied Statistics*, pages 39–51, 1961.
- 22 Richard J. Lipton and Andrew Tomkins. Online interval scheduling. In *Proc. 5th Symp. Discr. Algorithms (SODA)*, pages 302–311, 1994. URL: <http://dl.acm.org/citation.cfm?id=314464.314506>.
- 23 Tengyu Ma, Bo Tang, and Yajun Wang. The simulated greedy algorithm for several submodular matroid secretary problems. In *Proc. 30th Symp. Theoret. Aspects of Computer Science (STACS)*, pages 478–489, 2013. doi:10.4230/LIPIcs.STACS.2013.478.
- 24 Alessandro Panconesi and Aravind Srinivasan. Randomized distributed edge coloring via an extension of the chernoff-hoeffding bounds. *SIAM J. Comput.*, 26(2):350–368, 1997. doi:10.1137/S0097539793250767.
- 25 Gerhard J. Woeginger. On-line scheduling of jobs with fixed start and end times. *Theoret. Comput. Sci.*, 130(1):5–16, 1994. doi:10.1016/0304-3975(94)90150-3.