

The Online Stochastic Generalized Assignment Problem

Saeed Alaei ^{*}, MohammadTaghi Hajiaghayi ^{*}, and Vahid Liaghat ^{*}

Department of Computer Science, University of Maryland,
College Park, MD 20742
{saeed, hajiagha, vliaghat}@cs.umd.edu

Abstract. We present a $1 - \frac{1}{\sqrt{k}}$ -competitive algorithm for the online stochastic generalized assignment problem under the assumption that no item takes up more than $\frac{1}{k}$ fraction of the capacity of any bin. Items arrive online; each item has a value and a size; upon arrival, an item can be placed in a bin or discarded; the objective is to maximize the total value of the placement. Both value and size of an item may depend on the bin in which the item is placed; the size of an item is revealed only after it has been placed in a bin; distribution information is available about the value and size of each item in advance (not necessarily i.i.d), however items arrive in adversarial order (non-adaptive adversary). We also present an application of our result to subscription-based advertising where each advertiser, if served, requires a given minimum number of impressions (i.e., the “*all or nothing*” model).

1 Introduction

The generalized assignment problem (GAP) and its special cases *multiple knapsack*¹ and *bin packing*² capture several fundamental optimization problems and have many practical applications in computer science, operations research, and related disciplines. The (offline) GAP is defined as follows:

Definition 1 (Generalized Assignment Problem). *There is a set of n items and m bins. Each bin has a hard capacity where the total size of items placed in a bin cannot exceed its capacity. Each item has a value and a size if placed in a bin; both might depend on the bin. The goal is to find a maximum valued assignment of items to the bins which respects the capacities of the bins.*

For example GAP can be viewed as a scheduling problem on parallel machines, where each machine has a capacity (or a maximum load) and each job has a size (or a processing time) and a profit, each possibly dependent on the machine to which it

^{*} Supported in part by NSF CAREER award 1053605, ONR YIP award N000141110662, DARPA/AFRL award FA8650-11-1-7162, a Google Faculty Research Award, and a University of Maryland Research and Scholarship Award (RASA).

¹ In the multiple knapsack problem, we are given a set of items and a set of bins (knapsacks) such that each item j has a profit v_j and a size s_j , and each bin i has a capacity C_i . The goal is to find a subset of items of maximum profit such that they have a feasible packing in the bins.

² In the bin packing problem, given a set of items with different sizes, the goal is to find a packing of items into unit-sized bins that minimizes the number of bins used.

is assigned, and the objective is to find a feasible scheduling which maximizes the total profit. Though multiple knapsack and bin packing have a fully polynomial-time approximation scheme (asymptotic for bin packing) in the offline setting [6], GAP is APX-hard and the best known approximation ratio is $1 - 1/e + \epsilon$ where $\epsilon \approx 10^{-180}$ [10], which improves on a previous $(1 - 1/e)$ -approximation [12].

In this paper we consider the *online stochastic* variant of the problem:

Definition 2 (Online Stochastic Generalized Assignment Problem). *There are n items arriving in an online manner which can be of different types. There are m (static) bins each with a capacity limit on the total size of items that can be placed in it. A type of an item is associated with a value and a size distribution which may depend on the bin to which the item is placed. Stochastic information is known about the type of an item and sizes/values of the types. Upon arrival of an item, the type of that item is revealed. However the realization of the size of the item is revealed only after it has been placed in a bin. The goal is to find a maximum valued assignment of items to the bins. We consider a large-capacity assumption: no item takes up more than $\frac{1}{k}$ fraction of the capacity of any bin.*

We emphasize that there are two sources of uncertainty in our model: the type of an item and the size of the item. The type of an item (which contains a size-distribution) is revealed before making the assignment, however the actual size of an item is revealed after the assignment. To the best of our knowledge, Feldman et al. [11] were the first to consider the generalized assignment problem in an online setting, albeit with deterministic sizes. In the adversarial model where the items and the order of arrivals are chosen by an adversary, there is no competitive algorithm. Consider the simple case of one bin with capacity one and two arriving items each with size one. The value of the first item is 1. The value of the second item would be either $\frac{1}{\epsilon}$ or 0 based on whether we place the first item in the bin. Thus the online profit cannot be more than ϵ factor of the offline profit. Indeed one can show a much stronger hardness result for the adversarial model: no algorithm can be competitive for the two special cases of GAP, namely the adword problem³ and the display ad problem⁴ even under the large-capacity assumption [11,19].

Since no algorithm is competitive for online GAP in the adversarial model, Feldman et al. consider this model with *free disposal*. In the free disposal model, the total size of items placed in a bin may exceed its capacity, however, the profit of the bin is the maximum-valued subset of the items in the bin which does not violate the capacity. Feldman et al. give a $(1 - \frac{1}{e} - \epsilon)$ -competitive primal-dual algorithm for GAP under the free disposal assumption and the additional large-capacity assumption by which the capacity of each bin is at least $O(\frac{1}{\epsilon})$ times larger than the maximum size of a single item. Although the free disposal assumption might be counter-intuitive in time-sensitive applications such as job scheduling where the machine may start doing a job right after the job assignment, it is a very natural assumption in many applications including applications in economics like ad allocation – a buyer does not mind receiving more items.

³ The adword problem is a special case of GAP where the size and the value of placing an item in a bin is the same.

⁴ The display ad problem is a special case of GAP where all sizes are uniform.

Dean, Goemans, and Vondrak [7] consider the (offline) stochastic knapsack problem which is closely related to GAP. In their model, there is only one bin and the value of each item is known. However, the size of each item is drawn from a known distribution only after it is placed in the knapsack. We note that this is an offline setting in the sense that we may choose any order of items for allocation. This model is motivated by job scheduling on a single machine where the actual processing time required for a job is learned only after the completion of the job. Dean et al. give various adaptive and non-adaptive algorithms for their model where the best one has a competitive ratio $\frac{1}{3} - \epsilon$. Recently Bhalgat improved the competitive ratio to $\frac{1}{2} - \epsilon$ [4]. Other variations, such as soft capacity constraints, have also been considered for which we refer the reader to [5,13,16]. Dean et al. [7] also introduce an *ordered model* where items must be considered in a specific order, which can be seen as a version of the the online model with a known order. Dean et al. [7] present a $\frac{1}{9.5}$ -competitive algorithm. In general, the online model can be considered as a more challenging variation of the models proposed by Dean et al, and we show that the large-capacity assumption is enough to overcome this challenge.

To the best of our knowledge, the current variation of the online stochastic GAP has not been considered before. We note that since the distributions are not necessary i.i.d., this model generalizes the well-known *prophet inequalities*.⁵ Even with stochastic information about the arriving queries, no online algorithm can achieve a competitive ratio better than $\frac{1}{2}$ [1,14,17,18]. Consider the simple example from before where the value of the first item is 1 with probability one and the value of the second item is $\frac{1}{\epsilon}$ with probability ϵ , and 0 with probability $1 - \epsilon$. The algorithm can only select one item. No online (randomized) algorithm can achieve a profit more than $\max\{1, \epsilon(\frac{1}{\epsilon})\} = 1$ in expectation. However, the expected profit of the optimum offline assignment is $(1 - \epsilon)1 + \epsilon(\frac{1}{\epsilon}) = 2 - \epsilon$. Therefore without any additional assumption one cannot get a competitive ratio better than $1/2$. We overcome this difficulty by considering the natural large-capacity assumption which arises in many applications such as online advertising.

A summary of the other related work is presented in the full version of the paper.

1.1 Our Contribution

The *loss factor* of an online algorithm is the ratio α such that the profit of the algorithm is at least $1 - \alpha$ fraction of the optimal offline profit. The main result of the paper can be summarized in the following theorem (formally stated in Theorem 4).

Theorem. *For the stochastic generalized assignment problem, there exists a randomized online algorithm (see Definition 6) with the loss factor at most $\frac{1}{\sqrt{k}}$ in expectation.*

The proposed algorithm initially computes an optimal solution for a linear program corresponding to a fractional expected instance. In the online stage, the algorithm tentatively assigns each item upon arrival to one of the bins at random with probabilities

⁵ In the classic prophet inequality problem, given the distribution of a sequence of random variables, an onlooker has to choose from the succession of these values. The onlooker can only choose a certain number of values and cannot choose a past value. The goal is to maximize the total sum of selected numbers.

proportional to the fractional LP solution. This ensures that the expected total size of items assigned tentatively to each bin does not exceed its capacity. However, once a bin becomes full, any item which gets tentatively assigned to that bin will have to be discarded. In general, a straightforward randomized assignment based on the LP solution could be arbitrarily far from optimal; that is because the probability of an item being discarded due to a bin being full could be arbitrarily close to 1 for items that arrive towards the end. To overcome this problem, we incorporate an adaptive threshold based strategy for each bin so that an item tentatively assigned to a bin is only placed in the bin if the remaining capacity of the bin is more than a certain threshold. This ensures the online algorithm discards a tentatively assigned item with a probability at most $\frac{1}{\sqrt{k}}$ of the fractional LP assignment. The thresholds are computed adaptively based on previously observed items.

Indeed by using the fractional solution as a guideline, it is possible to achieve a non-adaptive competitive algorithm. One may scale down the fractional solution by a factor $1 - O(\frac{\log k}{\sqrt{k}})$ and assign the items with the modified probabilities, thus achieving a loss factor $O(\frac{\log k}{\sqrt{k}})$. By using the Chernoff bound it can be shown that the probability of exceeding bin capacities is very small. This simple algorithm gives an asymptotically optimum solution, however there are two drawbacks. The first issue is that the constants in various implementations of this idea are large. Thus unless the value k is very large, this algorithm cannot guarantee a reasonable competitive ratio; in contrast, the loss factor of our algorithm is exactly $\frac{1}{\sqrt{k}}$. On the other hand, in the applications of online GAP such as the Adword problem, the factor $O(\frac{\log k}{\sqrt{k}})$ is the loss factor of the millions of dollars. Therefore our algorithm saves a logarithmic factor in the loss of revenue. Indeed these drawbacks were also the motivation for improving the loss factor in the special cases of online GAP in previous papers [2,3].

The threshold based strategy of the online algorithm is presented in Section 3 in the form of a generalization of the magician's problem of [1]. The original magician's problem can be interpreted as a stochastic knapsack problem with unit size items and a knapsack of size k and such that each item arrives in one of two possible states (e.g. good/bad) with known probabilities; the objective being to maximize – simultaneously for all items – the probability of picking every item that is good. On the other hand, in the generalized variant presented in the current paper, the size of each item can vary according to an arbitrary (but known) distribution; in this version k is an integer lower bound on the ratio of the total size of the knapsack to the maximum possible size of a single item. Although the bound we obtain for the generalized magician is similar to that of [1], they are incomparable for small k ; in particular for $k = 1$, one can easily achieve a $\frac{1}{2}$ -competitive algorithm for the magician's problem with fixed size items, whereas for the generalized version with variable size items, no constant competitive algorithm is possible for $k = 1$.

Recently, Alaei et al. [3,2]⁶ use a combination of expected linear programming approach and dynamic programming to achieve a $1 - \epsilon$ -competitive algorithm for adword and display ad. They use a relatively simple dynamic programming in combination with

⁶ In an independent work, Devanur et al.[8] also consider the expected linear program of a similar problem.

the LP solution to check whether they should assign an item to a bidder or discard it. Using an approach similar to “dual fitting” [15], they demonstrate an analysis of the combination of a dynamic programming approach with an online LP-based algorithm and prove a loss factor $\frac{1}{\sqrt{k+3}}$ for the display ad problem. They use the sand theorem of [1] as a black-box in their analysis to derive proper dual variables in their dual fitting analysis of the algorithm. A dynamic programming approach needs to know the stochastic information about the remaining items, while a threshold-based approach needs to know the past, i.e., they are complements of each other. However, analysis of the dynamic programming even with uniform sizes is involved. Furthermore, it is not easy to generalize the approach of [2] to the model of Goemans et al. [7] where the given sizes show only the expected size of an item.

It is worthwhile to compare the stochastic model of the current paper against other popular models, i.e., random order model⁷ and unknown distribution model⁸. While both the random order and the unknown distribution models require less stochastic information, they both treat items uniformly; hence they are more suitable for applications where items are more symmetric⁹. The model considered in the current paper is more suitable when there is a high degree of distribution asymmetry across the items. In particular, the extra stochastic information allows us to obtain practical bounds even for small values of k whereas in other stochastic models the obtained bounds often become meaningful only asymptotically in k .

The all-or-nothing model. The online algorithm of this paper can be applied to a slightly different model in which each item should still be either fully assigned or discarded, however in case of assignment, unlike GAP, an item can be fractionally split across multiple bins (i.e., the all-or-nothing assignment model). Note that the LP for the expected instance is still the same for the all-or-nothing model, therefore our online algorithm still obtains the same bound in expectation compared to the optimal offline solution. The all-or-nothing model is suitable for *subscription-based advertisement* and *banner advertisement*.

The subscription-based advertisement problem is an example of an offline ad allocation setting motivated by the banner advertisement. In this problem, there is a set of contracts proposed by the advertisers and the goal is to accept the contracts of a subset of the advertisers which maximizes the revenue. The contract proposed by an advertiser specifies a collection of webpages which are relevant to his product, a desired total quantity of impressions on these pages, and a price. Each webpage has an ad inventory with a certain number of banner ads. The problem of selecting a feasible subset of advertisers with the maximum total value does not have any non-trivial approximation. This can be shown by a reduction from the Independent Set problem on a graph; advertisers represent the vertices of the graph and webpages represent the edges of the graph.

⁷ Random arrival model: items are chosen by an adversary, but they arrive in a random order.

⁸ Unknown distribution model: items are chosen i.i.d. from a fixed but unknown distribution.

⁹ At a first glance, the random order model may appear to allow for asymmetry, however note that for any i and j , the i^{th} arriving item and the j^{th} arriving item have the same ex ante distribution in the random order model.

Advertisers desire all the impressions of the relevant webpages. Thus any feasible subset of advertisers would denote an independent set in the graph. This shows maximizing the total value does not have a non-trivial approximation. Different pricing models have also been introduced by Feige et al. [9]. The proof of the following corollary is by a reduction from the all-or-nothing model.

Corollary. *There exists a randomized algorithm for the subscription-based advertisement problem which obtains a loss factor $\frac{1}{\sqrt{k}}$ in expectation where the number of available impressions on each website is at least k times the required impressions of each relevant advertiser.*

Proof. One can show that this is an offline version of the all-or-nothing model where bins denote the webpages and items denote the advertisers. The size of an item is the required number of ads of an advertiser and the value of an item is the proposed price. By Theorem 4, we can achieve at least $1 - \frac{1}{\sqrt{k}}$ fraction of the optimal profit in expectation.

2 Preliminaries

Model. We consider the problem of assigning n items to m bins; items arrive online and stochastic information is known about the size/value of each item; the objective is to maximize the total value of the assignment. The item $i \in [n]$ (arriving at time i) has r_i possible types with each type $t \in [r_i]$ having a probability of p_{it} , a value of $v_{itj} \in \mathbb{R}_+$, and a size of $S_{itj} \in [0, 1]$ if placed in bin j (for each $j \in [m]$); S_{itj} is a random variable which is drawn from a distribution with a CDF of F_{itj} if the item is placed in bin j . Each bin $j \in [m]$ has a capacity of $c_j \in \mathbb{N}_0$ which limits the total size of the items placed in that bin¹⁰. The type of each item is revealed upon arrival and the item must be either placed in a bin or discarded; this decision cannot be changed later. The size of an item is revealed only after it has been placed in a bin, furthermore an item can be placed in a bin only if the bin has at least one unit of capacity left. We assume that n, m, c_j, v_{itj} and F_{itj} are known in advance.

Benchmark. Consider the following linear program in which $\tilde{s}_{itj} = \mathbf{E}_{S_{itj} \sim F_{itj}}[S_{itj}]$ ¹¹.

$$\begin{aligned}
 & \text{maximize} && \sum_i \sum_t \sum_j v_{itj} x_{itj} && (\overline{OPT}) \\
 & \text{subject to} && \sum_i \sum_t \tilde{s}_{itj} x_{itj} \leq c_j, && \forall j \in [m] \\
 & && \sum_j x_{itj} \leq p_{it}, && \forall i \in [n], \forall t \in [r_i] \\
 & && x_{itj} \in [0, 1], &&
 \end{aligned}$$

¹⁰ Our results hold for non-integer capacities, however we assume integer capacities to simplify the exposition.

¹¹ Throughout the rest of this paper, we often omit the range of the sums whenever the range is clear from the context (e.g., \sum_i means $\sum_{i \in [n]}$, and \sum_j means $\sum_{j \in [m]}$, etc).

The optimal value of this linear program, which corresponds to the expected instance, is an upper bound on the expected value of the optimal offline assignment.

Theorem 1. *The optimal value of the linear program (\overline{OPT}) is an upper bound on the expected value of the offline optimal assignment.*

Proof. Let x_{itj}^* denote the ex ante probability that item i is of type t and is assigned to bin j in the optimal offline assignment. It is easy to see that x_{itj}^* is a feasible assignment for the linear program. Furthermore, the expected value of the optimal offline assignment is exactly $\sum_i \sum_t \sum_j v_{itj} x_{itj}^*$ which is equal to the value of the linear program for x_{itj}^* which is itself no more than the optimal value of the linear program. Note that the optimal value of the linear program may be strictly higher since a feasible assignment of the linear program does not necessarily correspond to a feasible offline assignment policy.

Section 4 presents an online adaptive algorithm which obtains a loss factor $\frac{1}{\sqrt{k}}$ w.r.t. the optimal value of the above linear program, where $k = \min_j c_j$. We emphasize that our adaptive algorithm saves a logarithmic factor in the loss of the outcome compared to the non-adaptive methods. Next section presents a stochastic toy problem and its solution which is used in our online algorithm.

3 The Generalized Magician's Problem

We present a generalization of the magician's problem, which was originally introduced in [1]; we also present a near-optimal solution for this generalization.

Definition 3 (The Generalized Magician's Problem). *A magician is presented with a series of boxes one by one, in an online fashion. There is a prize hidden in one of the boxes. The magician has a magic wand that can be used to open the boxes. The wand has k units of mana [20]. If the wand is used on box i and has at least 1 unit of mana, the box opens, but the wand loses a random amount of mana $X_i \in [0, 1]$ drawn from a distribution specified on the box by its cumulative distribution function F_{x_i} (i.e., the magician learns F_{x_i} upon seeing box i). The magician wishes to maximize the probability of obtaining the prize, but unfortunately the sequence of boxes, the distributions written on the boxes, and the box containing the prize have been arranged by a villain; the magician has no prior information (not even the number of the boxes); however, it is guaranteed that $\sum_i \mathbf{E}[X_i] \leq k$, and that the villain has to prepare the sequence of boxes in advance (i.e., cannot make any changes once the process has started).*

The magician could fail to open a box either because (a) he might choose to skip the box, or (b) his wand might run out of mana before getting to the box. Note that once the magician fixes his strategy, the best strategy for the villain is to put the prize in the box which, based on the magician's strategy, has the lowest ex ante probability of being opened. Therefore, in order for the magician to obtain the prize with a probability of at least γ , he has to devise a strategy that guarantees an ex ante probability of at least γ for opening each box. Notice that allowing the prize to be split among multiple boxes does not affect the problem. We present an algorithm parameterized by a probability

$\gamma \in [0, 1]$ which guarantees a minimum ex-ante probability of γ for opening each box while trying to minimize the mana used. We show that for $\gamma \leq 1 - \frac{1}{\sqrt{k}}$ this algorithm never requires more than k units of mana.

Definition 4 (γ -Conservative Magician). *The magician adaptively computes a sequence of thresholds $\theta_1, \theta_2, \dots \in \mathbb{R}_+$ and makes a decision about each box as follows: let W_i denote the amount of mana lost prior to seeing the i^{th} box; the magician makes a decision about box i by comparing W_i against θ_i ; if $W_i < \theta_i$, it opens the box; if $W_i > \theta_i$, it does not open the box; and if $W_i = \theta_i$, it randomizes and opens the box with some probability (to be defined). The magician chooses the smallest threshold θ_i for which $\Pr[W_i \leq \theta_i] \geq \gamma$ where the probability is computed ex ante (i.e., not conditioned on X_1, \dots, X_{i-1}). Note that γ is a parameter that is given. Let $F_{W_i}(w) = \Pr[W_i \leq w]$ denote the ex ante CDF of random variable W_i , and let Y_i be the indicator random variable which is 1 iff the magician opens the box i . Formally, the probability with which the magician should open box i conditioned on W_i is computed as follows¹².*

$$\Pr[Y_i = 1 | W_i] = \begin{cases} 1 & W_i < \theta_i \\ (\gamma - F_{W_i}^-(\theta_i)) / (F_{W_i}(\theta_i) - F_{W_i}^-(\theta_i)) & W_i = \theta_i \\ 0 & W_i > \theta_i \end{cases} \quad (Y)$$

$$\theta_i = \inf\{w | F_{W_i}(w) \geq \gamma\} \quad (\theta)$$

In the above definition, $F_{W_i}^-$ is the left limit of F_{W_i} , i.e., $F_{W_i}^-(w) = \Pr[W_i < w]$.

Note that $F_{W_{i+1}}$ and $F_{W_{i+1}}^-$ are fully determined by F_{W_i} and F_{X_i} and the choice of γ (see Theorem 3). Observe that θ_i is in fact computed before seeing box i itself.

A γ -conservative magician may fail for a choice of γ unless all thresholds θ_i are less than or equal to $k - 1$. The following theorem states a condition on γ that is sufficient to guarantee that $\theta_i \leq k - 1$ for all i .

Theorem 2 (γ -Conservative Magician). *For any $\gamma \leq 1 - \frac{1}{\sqrt{k}}$, a γ -conservative magician with k units of mana opens each box with an ex ante probability of γ exactly.*

Proof. See Section 5.

Definition 5 (γ_k). *We define γ_k to be the largest probability such that for any $k' \geq k$ and any instance of the magician's problem with k' units of mana, the thresholds computed by a γ_k -conservative magician are no more than $k' - 1$. In other words, γ_k is the optimal choice of γ which works for all instances with $k' \geq k$ units of mana. By Theorem 2, we know that γ_k must be¹³ at least $1 - \frac{1}{\sqrt{k}}$.*

Observe that γ_k is a non-decreasing function in k and approaches 1 as $k \rightarrow \infty$. However $\gamma_1 = 0$ which is in contrast to the bound of $\frac{1}{2}$ obtained for $k = 1$ in [1] in which all X_i are Bernoulli random variables. It turns out that when X_i are arbitrary random variables in $[0, 1]$, no algorithm exists for the magician that can guarantee a constant non-zero probability for opening each box.

¹² Assume $W_0 = 0$

¹³ Because for any $k' \geq k$ obviously $1 - \frac{1}{\sqrt{k}} \leq 1 - \frac{1}{\sqrt{k'}}$.

Proposition 1. *For the generalized magician's problem for $k = 1$, no algorithm for the magician (online or offline) can guarantee a constant non-zero probability for opening each box.*

Proof. Suppose there is an algorithm for the magician that is guaranteed to open each box with a probability of at least $\gamma \in (0, 1]$. We construct an instance in which the algorithm fails. Let $n = \lceil \frac{1}{\gamma} \rceil + 1$. Suppose all X_i are (independently) drawn from the distribution specified below.

$$X_i = \begin{cases} \frac{1}{2n} & \text{with prob. } 1 - \frac{1}{2n}, \\ 1 & \text{with prob. } \frac{1}{2n}, \end{cases} \quad \forall i \in [n]$$

As soon as the magician opens a box, the remaining mana will be less than 1, so he will not be able to open any other box, i.e., the magician can open only one box at every instance. Let Y_i denote the indicator random variable which is 1 iff the magician opens box i . Since $\sum_i Y_i \leq 1$, it must be $\sum_i \mathbf{E}[Y_i] \leq 1$. On the other hand, $\mathbf{E}[Y_i] \geq \gamma$ because the magician has guaranteed to open each box with a probability of at least γ . However $\sum_i \mathbf{E}[Y_i] \geq n\gamma > 1$ which is a contradiction. Note that $\sum_i \mathbf{E}[X_i] < 1$ so it satisfies the requirement of Definition 3.

Computation of $F_{W_i}(\cdot)$. For every $i \in [n]$, the equation $W_{i+1} = W_i + Y_i X_i$ relates the distribution of W_{i+1} to those of W_i and X_i ¹⁴. The following lemma shows that the distribution of W_{i+1} is fully determined by the information available to the magician before seeing box $i + 1$.

Theorem 3. *In the algorithm of γ -conservative magician (Definition 4), the choice of γ and the distributions of X_1, \dots, X_i fully determine the distribution of W_{i+1} , for every $i \in [n]$. In particular, $F_{W_{i+1}}$ can be recursively defined as follows.*

$$\begin{aligned} F_{W_{i+1}}(w) &= F_{W_i}(w) - G_i(w) + \mathbf{E}_{X_i \sim F_{X_i}} [G_i(w - X_i)] \quad \forall i \in [n], \forall w \in \mathbb{R}_+ \quad (F_w) \\ G_i(w) &= \min(F_{W_i}(w), \gamma) \quad \forall i \in [n], \forall w \in \mathbb{R}_+ \quad (G) \end{aligned}$$

Proof. See Section 5.

As a corollary of Theorem 3, we show how F_{W_i} can be computed using dynamic programming, assuming X_i can only take discrete values that are proper multiples of some minimum value.

Corollary 1. *If all X_i are proper multiple of $\frac{1}{D}$ for some $D \in \mathbb{N}$, then $F_{W_i}(\cdot)$ can be computed using the following dynamic program.*

$$F_{W_{i+1}}(w) = \begin{cases} F_{W_i}(w) - G_i(w) + \sum_{\ell} \mathbf{Pr}[X_i = \frac{\ell}{D}] G_i(w - \frac{\ell}{D}) & i \geq 1, w \geq 0 \\ 1 & i = 0, w \geq 0 \\ 0 & \text{otherwise.} \end{cases}$$

for all values of $i \in [n]$ and $w \in \mathbb{R}_+$. In particular, the γ -conservative magician makes a decision for each box in time $O(D)$.

¹⁴ Note that the distribution of Y_i is dependent on/determined by W_i .

Note that it is enough to compute F_{w_i} only for proper multiples of $\frac{1}{D}$ because $F_{w_i}(w) = F_{w_i}(\lfloor \frac{Dw}{D} \rfloor)$ for any $w \in \mathbb{R}_+$.

4 The Online Algorithm

We present an online algorithm which obtains at least $1 - \frac{1}{\sqrt{k}}$ -fraction of the optimal value of the linear program (\overline{OPT}). The algorithm uses, as a black box, the solution of the generalized magician's problem.

Definition 6 (Online Stochastic GAP Algorithm).

1. Solve the linear program (\overline{OPT}) and let x be an optimal assignment.
2. For each $j \in [m]$, create a γ -conservative magician (Definition 4) with c_j units of mana for bin j . γ is a parameter that is given.
3. Upon arrival of each item $i \in [n]$, do the following:
 - (a) Let t denote the type of item i .
 - (b) Choose a bin at random such that each bin $j \in [m]$ is chosen with probability $\frac{x_{itj}}{p_{it}}$. Let j^* denote the chosen bin.
 - (c) For each $j \in [m]$, define the random variable X_{ij} as $X_{ij} \leftarrow S_{itj}$ if $j^* = j$, and $X_{ij} \leftarrow 0$ otherwise¹⁵. For each $j \in [m]$, write the CDF of X_{ij} on a box and present it to the magician of bin j . The CDF of X_{ij} is $F_{X_{ij}}(s) = (1 - \sum_{t'} x_{it'j}) + \sum_{t'} x_{it'j} F_{it'j}(s)$.
 - (d) If the magician for bin j^* opened his box in step 3c, then assign item i to bin j^* , otherwise discard the item. For each $j \in [m]$, if the magician of bin j opened his box in step 3c, decrease the mana of that magician by X_{ij} . In particular, $X_{ij} = 0$ for all $j \neq j^*$, and $X_{ij^*} = S_{itj^*}$.

Theorem 4. For any $\gamma \leq \gamma_k$, the online algorithm of Definition 6 obtains in expectation at least a γ -fraction of the expected value of the optimal offline assignment (recall that $\gamma_k \geq 1 - \frac{1}{\sqrt{k}}$).

Proof. By Theorem 1, it is enough to show that the online algorithm obtains in expectation at least a γ -fraction of the optimal value of the linear program (\overline{OPT}). Let x be an optimal assignment for the LP. The contribution of each item $i \in [n]$ to the value of bin $j \in [m]$ in the LP is exactly $\sum_t v_{itj} x_{itj}$. We show that the online algorithm obtains in expectation $\gamma \sum_t v_{itj} x_{itj}$ from each item i and each bin j .

Consider an arbitrary item $i \in [n]$ and an arbitrary bin $j \in [m]$. WLOG, suppose the items are indexed in the order in which they arrive. Observe that

$$\mathbf{E}[X_{ij}] = \sum_t p_{it} \frac{x_{itj}}{p_{it}} \mathbf{E}[S_{itj}] = \sum_t x_{itj} \tilde{s}_{itj}.$$

Consequently, $\sum_i \mathbf{E}[X_{ij}] = \sum_i \sum_t x_{itj} \tilde{s}_{itj} \leq c_j$.

¹⁵ Note that S_{itj} is learned only after item i is placed in bin j which implies that X_{ij} may not be known at this point, however the algorithm does not use X_{ij} until after it is learned.

The last inequality follows from the first set of constraints in the LP of (\overline{OPT}) . Given that $\sum_i \mathbf{E}[X_{ij}] \leq c_j$ and $\gamma \leq \gamma_k \leq \gamma_{c_j}$, Theorem 2 implies that the magician of bin j opens each box with a probability of γ . Therefore, the expected contribution of item i to bin j is exactly $\sum_t \gamma p_{it} \frac{x_{itj}}{p_{it}} v_{itj} = \gamma \sum_t x_{itj} v_{itj}$. Consequently, the online algorithm obtains $\gamma \sum_i \sum_j \sum_t x_{itj} v_{itj}$ in expectation which is at least a γ -fraction of the expected value of the optimal offline assignment. Furthermore, each magician guarantees that the total size of the items assigned to each bin does not exceed the capacity of that bin.

5 Analysis of Generalized γ -Conservative Magician

We present the proof of Theorem 2. We prove the theorem in two parts. In the first part, we show that the thresholds computed by the γ -conservative magician indeed guarantee that each box is opened with an ex-ante probability of γ , assuming there is enough mana. In the second part, we show that for any $\gamma \leq 1 - \frac{1}{\sqrt{k}}$, the thresholds θ_i are less than or equal to $k - 1$, for all i , which implies that the magician never requires more than k units of mana. Below, we repeat the formulation of the threshold based strategy of the magician.

$$\Pr [Y_i = 1 | W_i] = \begin{cases} 1 & W_i < \theta_i \\ (\gamma - F_{w_i}^-(\theta_i)) / (F_{w_i}(\theta_i) - F_{w_i}^-(\theta_i)) & W_i = \theta_i \\ 0 & W_i > \theta_i \end{cases} \quad (Y)$$

$$\theta_i = \inf\{w | F_{w_i}(w) \geq \gamma\} \quad (\theta)$$

Part 1. We show that the thresholds computed by a γ -conservative magician guarantee that each box is opened with an ex ante probability of γ (i.e., $\Pr[Y_i = 1] = \gamma$), assuming there is enough mana.

$$\begin{aligned} \Pr [Y_i \leq w] &= \Pr [Y_i = 1 \cap W_i < \theta_i] + \Pr [Y_i = 1 \cap W_i = \theta_i] \\ &\quad + \Pr [Y_i = 1 \cap W_i > \theta_i] \\ &= \Pr [W_i < \theta_i] + \frac{\gamma - F_{w_i}^-(\theta_i)}{F_{w_i}(\theta_i) - F_{w_i}^-(\theta_i)} \Pr [W_i = \theta_i] = \gamma \end{aligned}$$

Part 2. Assuming $\gamma \leq 1 - \frac{1}{\sqrt{k}}$, we show that the thresholds computed by a γ -conservative magician are no more than $k - 1$ (i.e., $\theta_i \leq k - 1$ for all i). First, we present an interpretation of how $F_{w_i}(\cdot)$ evolves in i in terms of a sand displacement process.

Definition 7 (Sand Displacement Process). Consider one unit of infinitely divisible sand which is initially at position 0 on the real line. The sand is gradually moved to the right and distributed over the real line in n rounds. Let $F_{w_i}(w)$ denote the total amount of sand in the interval $[0, w]$ at the beginning of round $i \in [n]$. At each round i the following happens.

- (I) The leftmost γ -fraction of the sand is selected by first identifying the smallest threshold $\theta_i \in \mathbb{R}_+$ such that $F_{w_i}(\theta_i) \geq \gamma$ and then selecting all the sand in the interval $[0, \theta_i]$ and selecting a fraction of the sand at position θ_i itself such that the total amount of selected sand is equal to γ . Formally, if $G_i(w)$ denotes the total amount of sand selected from $[0, w]$, the selection of sand is such that $G_i(w) = \min(F_{w_i}(w), \gamma)$, for every $w \in \mathbb{R}_+$. In particular, this implies that only a fraction of the sand at position θ_i itself might be selected, however all the sand to the left of position θ_i is selected.
- (II) The selected sand is moved to the right as follows. Consider the given random variable $X_i \in [0, 1]$ and let $F_{X_i}(\cdot)$ denote its CDF. For every point $w \in [0, \theta_i]$ and every distance $\delta \in [0, 1]$, take a fraction proportional to $\Pr[X_i = \delta]$ out of the sand which was selected from position w and move it to position $w + \delta$.

It is easy to see that θ_i and $F_{w_i}(w)$ resulting from the above process are exactly the same as those computed by the γ -conservative magician.

Lemma 1. At the end of the i^{th} round of the sand displacement process, the total amount of sand in the interval $[0, w]$ is given by the following equation.

$$F_{w_{i+1}}(w) = F_{w_i}(w) - G_i(w) + \mathbf{E}_{X_i \sim F_{X_i}}[G_i(w - X_i)] \quad \forall i \in [n], \forall w \in \mathbb{R}_+ \quad (F_w)$$

Proof. According to definition of the sand displacement process, $F_{w_{i+1}}(w)$ can be defined as follows.

$$\begin{aligned} F_{w_{i+1}}(w) &= (F_{w_i}(w) - G_i(w)) + \iint_{\omega + \delta \leq w} dG_i(\omega) dF_{X_i}(\delta) \\ &= F_{w_i}(w) - G_i(w) + \int G_i(\omega - \delta) dF_{X_i}(\delta) \\ &= F_{w_i}(w) - G_i(w) + \mathbf{E}_{X_i \sim F_{X_i}}[G_i(w - X_i)] \end{aligned}$$

Proof (Proof of Theorem 3). The claim follows directly from Lemma 1

Consider a conceptual barrier which is at position $\theta_i + 1$ at the beginning of round i and is moved to position $\theta_{i+1} + 1$ for the next round, for each $i \in [n]$. It is easy to verify (i.e., by induction) that the sand never crosses to the right side of the barrier (i.e., $F_{w_{i+1}}(\theta_i + 1) = 1$). In what follows, the sand theorem implies that the sand remains concentrated near the barrier throughout the process. The barrier theorem implies that the barrier never passes k .

Theorem 5 (Sand). Throughout the sand displacement process (Definition 7), at the beginning of round $i \in [n]$, the following inequality holds.

$$F_{w_i}(w) < \gamma F_{w_i}(w + 1), \quad \forall i \in [n], \forall w \in [0, \theta_i] \quad (F_w\text{-ineq})$$

Furthermore, at the beginning of round $i \in [n]$, the average distance of the sand from the barrier, denoted by d_i , is upper bounded by the following inequalities¹⁶ in which the first inequality is strict except for $i = 1$.

$$d_i \leq (1 - \{\theta_i\}) \frac{1 - \gamma^{\lfloor \theta_i \rfloor + 1}}{1 - \gamma} + \{\theta_i\} \frac{1 - \gamma^{\lceil \theta_i \rceil + 1}}{1 - \gamma} \leq \frac{1 - \gamma^{\lceil \theta_i \rceil + 1}}{1 - \gamma} < \frac{1}{1 - \gamma}, \quad \forall i \in [n] \quad (d)$$

Proof. We start by proving the inequality (F_w -ineq). The proof is by induction on i . The case of $i = 1$ is trivial because all the sand is at position 0 and so $\theta_1 = 0$. Suppose the inequality holds at the beginning of round i for all $w \in [0, \theta_i]$; we show that it holds at the beginning of round $i + 1$ for all $w \in [0, \theta_{i+1}]$. Note that $\theta_i \leq \theta_{i+1} \leq \theta_i + 1$, so there are two possible cases:

Case 1. $w \in [0, \theta_i]$. Observe that $G_i(w) = F_{w_i}(w)$ in this interval, so:

$$\begin{aligned} F_{w_{i+1}}(w) &= F_{w_i}(w) - G_i(w) + \mathbf{E}_{X_i} [G_i(w - X_i)] && \text{by } (F_w). \\ &= \mathbf{E}_{X_i} [F_{w_i}(w - X_i)] && \text{by } G_i(w) = F_{w_i}(w), \text{ for } w \in [0, \theta_i]. \\ &< \mathbf{E}_{X_i} [\gamma F_{w_i}(w - X_i + 1)] && \text{by induction hypothesis.} \\ &= \gamma \mathbf{E}_{X_i} [F_{w_i}(w - X_i + 1) - G_i(w - X_i + 1) + G_i(w - X_i + 1)] \\ &\leq \gamma (F_{w_i}(w + 1) - G_i(w + 1) + \mathbf{E}_{X_i} [G_i(w - X_i + 1)]) \\ & && \text{by monotonicity of } F_{w_i}(\cdot) - G_i(\cdot). \\ &= \gamma F_{w_{i+1}}(w + 1) && \text{by } (F_w). \end{aligned}$$

Case 2. $w \in [\theta_i, \theta_{i+1}]$. We prove the claim by showing that $F_{w_{i+1}}(w) < \gamma$ and $F_{w_{i+1}}(w + 1) = 1$. Observe that $F_{w_{i+1}}(w) < \gamma$ because $w < \theta_{i+1}$ and because of the definition of θ_{i+1} in (θ) . Furthermore, observe that $F_{w_{i+1}}(w + 1) \geq F_{w_{i+1}}(\theta_i + 1) = 1$ both before and after round i all the sand is still contained in the interval $[0, \theta_i + 1]$.

Next, we prove inequality (d) which upper bounds the average distance of the sand from the barrier at the beginning of round $i \in [n]$.

$$\begin{aligned} d_i &= \int_0^{\theta_i + 1} (\theta_i + 1 - w) dF_{w_i}(w) \\ &= \int_0^{\theta_i + 1} F_{w_i}(w) dw && \text{by integration by part.} \\ &= \sum_{\ell=0}^{\lceil \theta_i \rceil} \int_{\theta_i - \ell}^{\theta_i + 1 - \ell} F_{w_i}(w) dw \end{aligned}$$

¹⁶ Note that $\{z\} = z - \lfloor z \rfloor$, for any z .

$$\begin{aligned}
&\leq \sum_{\ell=0}^{\lfloor \theta_i \rfloor} \int_{\theta_i}^{\theta_i+1} \gamma^\ell F_{w_i}(w) dw + \int_{\lfloor \theta_i \rfloor + 1}^{\theta_i+1} \gamma^{\lceil \theta_i \rceil} F_{w_i}(w) dw && \text{by } (F_w\text{-ineq}). \\
&\leq \sum_{\ell=0}^{\lfloor \theta_i \rfloor} \gamma^\ell + \{\theta_i\} \gamma^{\lceil \theta_i \rceil} && \text{by } F_{w_i}(w) \leq 1. \\
&= (1 - \{\theta_i\}) \sum_{\ell=0}^{\lfloor \theta_i \rfloor} \gamma^\ell + \{\theta_i\} \sum_{\ell=0}^{\lceil \theta_i \rceil} \gamma^\ell \\
&= (1 - \{\theta_i\}) \frac{1 - \gamma^{\lfloor \theta_i \rfloor + 1}}{1 - \gamma} + \{\theta_i\} \frac{1 - \gamma^{\lceil \theta_i \rceil + 1}}{1 - \gamma} \leq \frac{1 - \gamma^{\lceil \theta_i \rceil + 1}}{1 - \gamma}
\end{aligned}$$

The last inequality follows because $(1 - \beta)L + \beta H \leq H$ for any $\beta \in [0, 1]$ and any L, H with $L \leq H$. Note that at least one of the first two inequalities is strict except for $i = 1$ which proves the claim.

Theorem 6 (Barrier). *If $\sum_{i=1}^n \mathbf{E}_{X_i \sim F_{X_i}}[X_i] \leq k$ for some $k \in \mathbb{N}$, and $\gamma \leq 1 - \frac{1}{\sqrt{k}}$, then the distance of the barrier from the origin is no more than k throughout the process, i.e., $\theta_i \leq k - 1$ for all $i \in [n]$.*

Proof. At the beginning of round i , let d_i and d'_i denote the average distance of the sand from the barrier and from the origin respectively. Recall that the barrier is defined to be at position $\theta_i + 1$ at the beginning of round i . Observe that $d_i + d'_i = \theta_i + 1$. Furthermore, $d'_{i+1} = d'_i + \gamma \mathbf{E}[X_i]$, i.e., the average distance of the sand from the origin is increased exactly by $\gamma \mathbf{E}[X_i]$ during round i (because the amount of selected sand is exactly γ and the sand selected from every position $w \in [0, \theta_i]$ is moved to the right by an expected distance of $\mathbf{E}[X_i]$). By applying Theorem 5 we get the following inequality for all $i \in [n]$.

$$\begin{aligned}
\theta_i + 1 &= d'_i + d_i < \gamma \sum_{r=1}^{i-1} \mathbf{E}[X_r] + d_i \\
&\leq \gamma k + (1 - \{\theta_i\}) \frac{1 - \gamma^{\lfloor \theta_i \rfloor + 1}}{1 - \gamma} + \{\theta_i\} \frac{1 - \gamma^{\lceil \theta_i \rceil + 1}}{1 - \gamma}
\end{aligned}$$

In order to show that distance of the barrier from the origin is no more than k throughout the process, it is enough to show that the above inequality cannot hold for $\theta_i > k - 1$. In fact it is just enough to show that it cannot hold for $\theta_i = k - 1$; alternatively, it is enough to show that the complement of the above inequality holds for $\theta_i = k - 1$, i.e., $k \geq \gamma k + \frac{1 - \gamma^k}{1 - \gamma}$. To complete the proof, consider the stronger inequality $k \geq \gamma k + \frac{1}{1 - \gamma}$ which is quadratic in γ and can be solved to get a bound of $\gamma \leq 1 - \frac{1}{\sqrt{k}}$.

Theorem 6 implies that a γ -conservative magician requires no more than k units of mana, assuming that $\gamma \leq 1 - \frac{1}{\sqrt{k}}$. That completes the proof of Theorem 2.

References

1. S. Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In *FOCS*, 2011.
2. S. Alaei, M. T. Hajiaghayi, and V. Liaghat. Online prophet-inequality matching with applications to ad allocation. In *EC*, 2012.
3. S. Alaei, M. T. Hajiaghayi, V. Liaghat, D. Pei, and B. Saha. Adcell: Ad allocation in cellular networks. In *ESA*, 2011.
4. A. Bhalgat. A $(2 + \epsilon)$ -approximation algorithm for the stochastic knapsack problem. In *Unpublished Manuscript*, 2012.
5. A. Bhalgat, A. Goel, and S. Khanna. Improved approximation results for stochastic knapsack problems. In *SODA*, 2011.
6. C. Chekuri and S. Khanna. A ptas for the multiple knapsack problem. In *SODA*, 2000.
7. B. C. Dean, M. X. Goemans, and J. Vondrak. Approximating the stochastic knapsack problem: The benefit of adaptivity. In *FOCS*, 2004.
8. N. R. Devanur, K. Jain, B. Sivan, and C. A. Wilkens. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In *EC*, 2011.
9. U. Feige, N. Immorlica, V. Mirrokni, and H. Nazerzadeh. A combinatorial allocation mechanism with penalties for banner advertising. In *WWW*, 2008.
10. U. Feige and J. Vondrak. Approximation algorithms for allocation problems: Improving the factor of $1 - 1/e$. In *FOCS*, 2006.
11. J. Feldman, N. Korula, V. Mirrokni, S. Muthukrishnan, and M. Pál. Online ad assignment with free disposal. In *WINE*, 2009.
12. L. Fleischer, M. X. Goemans, V. S. Mirrokni, and M. Sviridenko. Tight approximation algorithms for maximum general assignment problems. In *SODA*, 2006.
13. A. Goel and P. Indyk. Stochastic load balancing and related problems. In *FOCS*, 1999.
14. M. T. Hajiaghayi, R. D. Kleinberg, and T. Sandholm. Automated online mechanism design and prophet inequalities. In *AAAI*, 2007.
15. K. Jain, M. Mahdian, E. Markakis, A. Saberi, and V. V. Vazirani. Greedy facility location algorithms analyzed using dual fitting with factor-revealing LP. *J. ACM*, 2003.
16. J. Kleinberg, Y. Rabani, and E. Tardos. Allocating bandwidth for bursty connections. In *STOC*, 1997.
17. U. Krengel and L. Sucheston. Semiamarts and finite values. *Bull. Am. Math. Soc.*, 1977.
18. U. Krengel and L. Sucheston. On semiamarts, amarts, and processes with finite value. *Kuelbs, J., ed., Probability on Banach Spaces.*, 1978.
19. A. Mehta, A. Saberi, U. Vazirani, and V. Vazirani. Adwords and generalized online matching. *J. ACM*, 2007.
20. Wikipedia. <http://en.wikipedia.org/wiki/mana>, 2012.