Theoretical Computer Science 516 (2014) 121–126

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Note

Adaptivity in the stochastic blackjack knapsack problem

Asaf Levin *, Aleksander Vainer

Faculty of Industrial Engineering and Management, The Technion, 32000 Haifa, Israel

ARTICLE INFO

Article history:
Received 19 August 2013
Received in revised form 12 November 2013
Accepted 19 November 2013
Communicated by D.-Z. Du

Keywords:
Knapsack problem
Benefit of adaptivity

ABSTRACT

We consider a stochastic variant of the NP-hard 0/1 knapsack problem in which item values are deterministic and item sizes are independent random variables with known, arbitrary distributions. Items are placed in the knapsack sequentially, and the act of placing an item in the knapsack instantiates its size. In every stage of insertion if the subset of the items inserted thus far is feasible, then it has a total value that equals to the sum of values of all items of this subset. Otherwise, if the subset violates the constraint, then its value equals to zero. The goal is to compute a policy for insertion of the items, that maximizes the expected total value of items placed in the knapsack.

We consider both non-adaptive policies (that designate a priori a fixed subset of items to insert) and adaptive policies (that can make dynamic decisions based on the instantiated sizes of items placed in the knapsack thus far). Our work characterizes the benefit of adaptivity. For this purpose we use a measure called the adaptivity gap: the supremum over instances of the ratio between the expected value obtained by an optimal adaptive policy and the expected value obtained by an optimal non-adaptive policy. First we show a tight bound of $\frac{3}{2}$ on the adaptivity gap for the case of inputs consisting of only two items. Then we present a non-adaptive policy with expected value that is at least $(\sqrt{2}-1)^2/2 \approx 1/11.66$ times the expected value of the optimal adaptive policy. Thus the adaptivity gap in this model is at most 11.66. Additionally this non-adaptive policy is computed in polynomial time. Finally, we consider a special case of the model where all sizes are distributed according to Bernoulli distribution with different parameters. For this special case we improve our result and bound the adaptivity gap by 8.

1. Introduction

The DETERMINISTIC KNAPSACK PROBLEM (KP) is a well studied NP-hard problem, where the input consists of a set of n items characterized by non-negative values $v_1, \dots, v_n$ and non-negative sizes $S_1, \dots, S_n$, and the goal is to find a maximum-value subset of these items whose total size is at most one. That is, formally KP is defined as follows

$$\max\left\{\sum_{i=1}^{n} v_i x_i\colon \sum_{i=1}^{n} S_i x_i \leqslant 1,\ x_i \in \{0,1\},\ \forall i\right\}.$$

See [13,11] for surveys of results on the knapsack problem. Despite its theoretical importance, the deterministic knapsack problem fails to capture many realistic scenarios. In many practical applications, the a priori information regarding item sizes is stochastic, so a deterministic model does not fit in these cases.

* Corresponding author.
E-mail addresses: levinas@ie.technion.ac.il (A. Levin), alevai@tx.technion.ac.il (A. Vainer).

0304-3975/$ – see front matter
http://dx.doi.org/10.1016/j.tcs.2013.11.029

Stochastic behavior of optimization problems was studied for many problems that are similar to the knapsack problem (see e.g. [2,14]) and also to the knapsack problem itself (see e.g. [10,15,4,7]). In this work we consider two models defined as follows.

Definition 1. Stochastic Knapsack Problem (SKP) [4] is a stochastic variant of KP where for all j, v_j is deterministic, while S_j is a random variable with a known, arbitrary distribution. Moreover, S_j's are mutually independent and the distributions of distinct items may be different. The goal is to compute a *policy* that maximizes the expected value of items successfully placed in the knapsack, where the first overflowing item contributes no value, and stops the process of inserting items into the knapsack.

Definition 2. Blackjack Knapsack Problem (BKP) is a stochastic variant of KP where for all j, v_j is deterministic, while S_j is a random variable with a known, arbitrary distribution. Moreover, S_j's are mutually independent and the distributions of distinct items may be different. The goal is to compute a *policy* that maximizes the expected value of items placed in the knapsack, where for a given subset of items, its value is the total value of its items if their total size is at most 1, and otherwise the value of this subset is 0.

The motivation for considering BKP is based on the card game Blackjack (also known as twenty-one). It is a comparing card game between a player and a dealer, where every card is characterized by some fixed number of points. After receiving his initial two cards, the player has the option of getting a "hit", i.e., receiving an additional card(s). The goal of the game is to reach 21 points or to reach a score higher than a dealer without exceeding 21. If a player is "busting", i.e., exceeds 21, then he automatically gains 0 points and thus loses the game. This is the same kind of penalty used in the definition of BKP.

The policy (in both models) can decide to terminate after any stage. However, in SKP, there is no need to decide to terminate, since there is no penalty on overflowing the knapsack's capacity. The policies (in both models) could be adaptive or non-adaptive. An adaptive policy can use the instantiated values of the stochastic information revealed by the policy thus far. A non-adaptive policy does not use this kind of information, and thus in SKP a non-adaptive policy is a permutation of (perhaps a subset of) the items and in BKP it is a subset of items. Note that the characteristics of non-adaptive policies are similar to the ones of general policies in static models (see e.g. [17,16,5]).

The implementation of adaptive policies is much more difficult than the implementation of non-adaptive policies, as for example the memory that is needed to encode an adaptive policy need not be finite. Thus, the restriction to non-adaptive policies is useful, and the goal of this work is to examine the extent to which the performance of the system deteriorates due to this constraint. Therefore, we consider the benefit of adaptivity [4]. For this purpose we examine a measure, called the *adaptivity gap*, defined as the supremum ratio between the expected value of an optimal adaptive policy and the expected value of an optimal non-adaptive policy, where the supremum is over all possible instances in our model. Note that by definition, in every instance of SKP or BKP there is an optimal non-adaptive policy (the set of feasible policies is finite in each of these cases), but there might be instances for which there is no optimal adaptive policy. In the sequel, we assume that there is an optimal adaptive policy (that is the supremum is attained by some adaptive policy). However, if this is not the case, we can replace an optimal adaptive policy by an adaptive policy whose expected value is close to the optimal value (up to an additive error of infinitesimally small value of ε). This does not change our results, and thus we prefer to keep the presentation for the case in which the supremum is attained.

Notation. Given an event C, we denote by $P(C)$ the probability of C. For a random variable X, we let $E(X)$ be its expected value (also called its expectation). For a policy $\mathcal{P}$ for BKP, we denote by $val(\mathcal{P})$ the value of the subset of items selected by policy $\mathcal{P}$ (the value as a solution to BKP).

Related work. The study of the benefit of adaptivity in stochastic combinatorial optimization problems started with the seminal work by Dean, Goemans and Vondrák [4]. In this work, the authors consider the SKP. They show that the adaptivity gap in this model is at most 4. More precisely, they construct a non-adaptive policy which approximates the optimal adaptive policy by a factor of 4.

Let $w_i = v_i \cdot P(S_i \leqslant 1)$ be the effective value of item i, $\mu_i = E(\min\{S_i, 1\})$ be the mean truncated size of item i. An inferior $\frac{32}{7}$ bound on the adaptivity gap for SKP is achieved by [4] based on a linear program Φ that allows to bound *ADAPT*, the expected value of the optimal adaptive policy in SKP. We will use the following result (Theorem 3.1 in [4]).

Theorem 3. *ADAPT* $\leqslant \Phi(2)$, *where* $\Phi(t) = \max\{\sum_{i=1}^n w_i x_i,\ \text{s.t.}\ \sum_{i=1}^n \mu_i x_i \leqslant t,\ x_i \in [0,1],\ \forall i\}$.

A generalization of SKP was presented by Gupta et al. [9] where both sizes and values are random and the assumption of independence between sizes and values of every item is omitted. Using a new linear programming formulation, they construct a non-adaptive policy that in expectation achieves at least $\frac{1}{8}$ of the expected value achieved by an optimal adaptive policy. Thus, the adaptivity gap in their model is at most 8.