Competitive Algorithms for Online Multidimensional Knapsack Problems

Lin Yang
University of Massachusetts Amherst
Amherst, MA, USA
linyang@cs.umass.edu

Ali Zeynali
University of Massachusetts Amherst
Amherst, MA, USA
azeynali@umass.edu

Mohammad H. Hajiesmaili
University of Massachusetts Amherst
Amherst, MA, USA
hajiesmaili@cs.umass.edu

Ramesh K. Sitaraman
University of Massachusetts Amherst
& Akamai Technologies
Amherst, MA, USA
ramesh@cs.umass.edu

Don Towsley
University of Massachusetts Amherst
Amherst, MA, USA
towsley@cs.umass.edu

ABSTRACT

In this work[1], we study the online multidimensional knapsack problem (called OMdKP) in which there is a knapsack whose capacity is represented in m dimensions, each dimension could have a different capacity. Then, n items with different scalar profit values and m-dimensional weights arrive in an online manner and the goal is to admit or decline items upon their arrival such that the total profit obtained by admitted items is maximized and the capacity of knapsack across all dimensions is respected. This is a natural generalization of the classic single-dimension knapsack problem with several relevant applications such as in virtual machine allocation, job scheduling, and all-or-nothing flow maximization over a graph. We develop an online algorithm for OMdKP that uses an exponential reservation function to make online admission decisions. Our competitive analysis shows that the proposed online algorithm achieves the competitive ratio of $O(\log(\theta\alpha))$, where α is the ratio between the aggregate knapsack capacity and the minimum capacity over a single dimension and θ is the ratio between the maximum and minimum item unit values. We also show that the competitive ratio of our algorithm with exponential reservation function matches the lower bound up to a constant factor.

CCS CONCEPTS

• **Theory of computation** → **Online algorithms**.

KEYWORDS

Online Algorithms, multidimensional knapsack, competitive ratio

ACM Reference Format:
Lin Yang, Ali Zeynali, Mohammad H. Hajiesmaili, Ramesh K. Sitaraman, and Don Towsley. 2022. Competitive Algorithms for Online Multidimensional Knapsack Problems. In *Abstract Proceedings of the 2022 ACM SIGMETRICS/IFIP PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS/PERFORMANCE '22 Abstracts), June 6–10, 2022, Mumbai, India.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3489048.3522627

[1]The full version of this abstract is available in [5].

SIGMETRICS/PERFORMANCE '22 Abstracts, June 6–10, 2022, Mumbai, India.

ACM ISBN 978-1-4503-9141-2/22/06.
https://doi.org/10.1145/3489048.3522627

1 INTRODUCTION

The online knapsack problem [3] (OKP) is a classical online optimization problem that has application in a variety of domains such as cloud and edge computing [7], online admission control [2] and online routing of virtual switches [4]. In the basic version of OKP, an online algorithm must make irrevocable decisions about which items with different values and weights to pack into a capacity-limited knapsack without knowing what items will arrive in the future. The goal of the algorithm is to maximize the aggregate value of admitted items while respecting the capacity of the knapsack. This problem has been tackled using the competitive online algorithm framework [1], and there are algorithms that achieve a competitive ratio of $O(\log\theta)$ for the basic version, where θ is a value fluctuation ratio between the most and the least valuable items. Further, it has been shown that the competitive ratio has a lower bound of $\Omega(\log\theta)$ [8], hence the algorithms in [8] are optimal since their competitive ratio is tight.

In this work, we study a different extension of OKP, the online multidimensional knapsack problem (OMdKP), in which there is a single knapsack whose capacity is represented by an m-dimensional vector, and the weights (or sizes) of online items are m-dimensional. The goal of an online decision maker is to pack the most valuable items so that the capacity of the knapsack in each dimension is not exceeded. Note that this problem differs from the online multiple knapsack problem in which fixed-weight items are packed into one or multiple knapsacks. In other words, in the online multiple knapsack problem decisions must be made about both *admission* and *allocation* whereas in OMdKP, only an admission decision is required.

Our Contributions. In this work, we develop online algorithms for OMdKP and analyze their competitive ratios. Specifically, we design a reservation policy explicitly accounting for item weights across different dimensions, which is an exponential function of knapsack utilization. The main contribution of this work is providing the first order-optimal algorithm for OMdKP under the commonly-made "infinitesimal" assumption that item weights are "small enough"' as compared to the capacity of the knapsack.

2 PROBLEM STATEMENT

We consider a knapsack whose capacities along m dimensions is represented by vector $\mathbf{C} = [C_1, \dots, C_j, \dots, C_m]$, where C_j represents the capacity of dimension $j \in [m] = \{1, \dots, m\}$ and C is the aggregate capacity over all dimensions, i.e., $C = \sum_{j\in[m]} C_j$. Without loss of generality, we assume, $C_1 \le C_2 \le \dots \le C_m$. Items arrive in an online fashion, each with a different value and weights. Specifically, in round $i \in [n] = \{1, \dots, n\}$, item i arrives with value $v_i \ge 0$, and a weight vector $\mathbf{w_i} = [w_{i,1}, \dots, w_{i,j}, \dots, w_{i,m}]$, where $w_{i,j} \ge 0$ is the size of item i in dimension j of the knapsack. Given item values and weights along with the capacity vector of the knapsack, the offline version of OMdKP can be formulated as

$$[\texttt{OMdKP}] \quad \max \sum\nolimits_{i\in[n]} v_i x_i, \quad \text{s.t.,} \sum\nolimits_{i\in[n]} w_{i,j} x_i \le C_j, \forall j \in [m], \tag{1}$$

where x_i's are the optimization variables and $x_i = 1$ if item i is admitted and $x_i = 0$, otherwise. We consider both integral and fractional versions of the problem. In the fractional version, $x_i \in [0, 1], \forall i \in [n]$, and $x_i \in \{0, 1\}, \forall i \in [n]$ for the integral version. We are interested in an online setting in which items arrive one-by-one and an online algorithm has to immediately decide whether to admit the incoming item without knowing the future and in the absence of a stochastic modeling. We present our main results for the integral version of OMdKP.

3 AN EXPONENTIAL RESERVATION POLICY (EXPRP) FOR OMDKP

We propose ExpRP that uses the normalized utilization and the reservation function for evaluating the admission cost. The definition of the normalized cost $z_{i,j}$ is

$$z_{i,j} = \left\lfloor \frac{u_{i,j}}{C_j} \log\left(\theta\alpha_j\right) \right\rfloor, j \in [m], \tag{2}$$

that represents the normalized utilization of each dimension after arrival of the i-th item. Then, the new item i is admitted if there is enough capacity and the following inequality holds.

$$v_i \ge \sum_{j=1}^{m} \left(2^{z_{i-1,j}} - 1\right) w_{i,j}. \tag{3}$$

We note that one may attain a different competitive ratio by specifying a different constant instead of 2 in the reservation function (3). For simplicity, we just choose 2 as the base of the reservation function. It is worth noting that following the recent approach in designing data-driven online algorithms [6], this coefficient could be changed to a parameter that could be learned in real-time to improve practical performance.

Given enough available space for admission, ExpRP makes an admission decision based on Equation (3), in which factor $(2^{z_{i-1,j}} - 1)$ represents the scarcity of dimension j. The larger the variable $z_{i-1,j}$, the larger the scarcity factor. By multiplying the scarcity and the weight $w_{i,j}$ in the same dimension, ExpRP evaluates the cost in each dimension, and admits item i only if its value is larger than or equal to the aggregate cost over all dimensions. Compared to LinRP, the scarcity factor in ExpRP increase exponentially in $z_{i,j}$ and ExpRP admits the item only when its value is at least equal to the sum of the costs over all dimensions. In general, ExpRP algorithm is more conservative in its admission decisions when there are bottleneck dimensions and thus tends to reserve the remaining capacity for higher-valued items.

3.1 Main Results

We present a lower bound for any competitive algorithm providing a feasible solution for OMdKP (in Theorem 3.1) followed by the competitive results for ExpRP (in Theorem 3.2). Note that ε serves as an upper bound of ratios between single-dimension size of items and the capacity of the knapsack. This valid range for this parameter is explicitly characterized to guarantee the competitive ratios.

THEOREM 3.1. *(Lower Bound on Competitive Ratio for OMdKP) The competitive ratio of any online algorithm providing a feasible solution to OMdKP is $\Omega(\log\theta\alpha)$.*

THEOREM 3.2. *With $\varepsilon < \min\{1/3, 1/(2\log(\theta\alpha))\}$, the competitive ratio of ExpRP satisfies*

$$\texttt{CR(ExpRP)} \le \max\left\{12, \frac{4\log(\theta\alpha)}{1 - 2\varepsilon\log(\theta\alpha)}\right\} + 1.$$

Remark. When $\varepsilon \to 0$, the competitive ratio of ExpRP satisfies

$$\texttt{CR(ExpRP)} \le \max\{12, 4\log\theta\alpha\} + 1.$$

Hence, ExpRP is $O(\log(\theta\alpha))$-competitive. Comparing the result in theorems 3.1 and 3.2 shows that ExpRP achieves the optimal competitive ratio up to a constant factor.

4 ACKNOWLEDGMENTS

This research is supported by NSF CAREER 2045641, CNS 2106299, CNS 2102963, CNS 1908298, CNS 1763617, CNS 1901137, the Army Research Laboratory under Cooperative Agreement W911NF-17-2-0196 (IoBT CRA), and the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement W911NF-16-3-0001.

REFERENCES

[1] Allan Borodin and Ran El-Yaniv. 2005. *Online computation and competitive analysis.* Cambridge University Press.
[2] Niv Buchbinder and Joseph Naor. 2009. *The design of competitive online algorithms via a primal-dual approach.* Now Publishers Inc.
[3] Deeparnab Chakrabarty, Yunhong Zhou, and Rajan Lukose. 2008. Online knapsack problems. In *Workshop on internet and network economics (WINE).*
[4] S. Paris, J. Leguay, L. Maggi, M. Draief, and S. Chouvardas. 2016. Online experts for admission control in SDN. In *NOMS 2016 - 2016 IEEE/IFIP Network Operations and Management Symposium.* 1003–1004.
[5] Lin Yang, Ali Zeynali, Mohammad H Hajiesmaili, Ramesh K Sitaraman, and Don Towsley. 2021. Competitive Algorithms for Online Multidimensional Knapsack Problems. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 5, 3 (2021), 1–30.
[6] Ali Zeynali, Bo Sun, Mohammad Hajiesmaili, and Adam Wierman. 2021. Data-driven Competitive Algorithms for Online Knapsack and Set Cover. In *Proc. of AAAI.*
[7] Zizhan Zheng and Ness B Shroff. 2016. Online multi-resource allocation for deadline sensitive jobs with partial values in the cloud. In *IEEE INFOCOM 2016-The 35th Annual IEEE International Conference on Computer Communications.* IEEE, 1–9.
[8] Yunhong Zhou, Deeparnab Chakrabarty, and Rajan Lukose. 2008. Budget constrained bidding in keyword auctions and online knapsack problems. In *International Workshop on Internet and Network Economics.* Springer, 566–576.