The Online Knapsack Problem with Departures

Bo Sun
The Chinese University of Hong Kong
Hong Kong, China
bsun@cse.cuhk.edu.hk

Lin Yang*
Nanjing University
Nanjing, China
linyang@nju.edu.cn

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

Adam Wierman
California Institute of Technology
Pasadena, USA
adamw@caltech.edu

John C.S. Lui
The Chinese University of Hong Kong
Hong Kong, China
cslui@cse.cuhk.edu.hk

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

Danny H.K. Tsang
The Hong Kong University of Science and Technology (Guangzhou)
Guangzhou, China
The Hong Kong University of Science and Technology
Hong Kong, China
eetsang@ust.hk

ABSTRACT

The online knapsack problem is a classic online resource allocation problem in networking and operations research. Its basic version studies how to pack online arriving items of different sizes and values into a capacity-limited knapsack. In this paper, we study a general version that includes *item departures*, while also considering *multiple knapsacks* and *multi-dimensional item sizes*. We design a threshold-based online algorithm and prove that the algorithm can achieve order-optimal competitive ratios. Beyond worst-case optimized algorithms, we also propose a data-driven online algorithm that can achieve near-optimal average performance under typical instances while guaranteeing the worst-case performance.

CCS CONCEPTS

• **Theory of computation** → **Online algorithms**; • **Networks** → *Network economics.*

KEYWORDS

online knapsack problems; knapsack with departures; data-driven algorithms; competitive ratio

ACM Reference Format:

Bo Sun, Lin Yang, Mohammad Hajiesmaili, Adam Wierman, John C.S. Lui, Don Towsley, and Danny H.K. Tsang. 2023. The Online Knapsack Problem with Departures. In *Abstract Proceedings of the 2023 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '23 Abstracts), June 19–23, 2023, Orlando, FL, USA.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3578338.3593576

*Corresponding author; the first two authors contribute equally to this work.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
SIGMETRICS '23 Abstracts, June 19–23, 2023, Orlando, FL, USA
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0074-3/23/06.
https://doi.org/10.1145/3578338.3593576

1 PROBLEM STATEMENT

We study the classic online knapsack problem (OKP). In its basic version, there is only one knapsack, and each item is characterized by its value and one-dimensional (scalar) size. The problem is to irrevocably decide whether to admit each item upon its arrival with the goal of maximizing the total values of admitted items while respecting the capacity of the knapsack. The sequence of items can only be revealed one-by-one and may even be adversarial. In this paper, we focus on a novel generalization that includes item departures, while also considering multiple knapsacks and multi-dimensional item sizes.

Consider K knapsacks in a slotted time horizon $[T] = \{1, \dots, T\}$, where each knapsack $k \in [K]$ has capacity $C_k \in \mathbb{R}^+$. A total of N items arrive sequentially and each item n is characterized by its item information $\mathcal{I}_n = \{a_n, \{w_{nk}, v_{nk}, \mathcal{T}_{nk}\}_{k\in[K]}\}$, where a_n is the arrival time, and for each knapsack k, w_{nk} and v_{nk} are the size and value, and $\mathcal{T}_{nk} := \{s_{nk}, \dots, s_{nk} + d_{nk} - 1\}$ is the set of time slots that item n requests to stay in knapsack k from its starting time s_{nk} to its departure time $s_{nk} + d_{nk} - 1$. The set $\mathcal{T}_{nk}$ contains d_{nk} consecutive time slots and we call d_{nk} the duration of the item.

Upon arrival of item n, a decision maker observes its item information $\mathcal{I}_n$ and determines (i) whether to admit this item, and (ii) which knapsack this item should be assigned to if it is admitted. Let $\boldsymbol{x}_n = \{x_{nk}\}_{k\in[K]}$ denote the decision variable, where $x_{nk} \in \{0, 1\}$ indicates whether to admit item n to knapsack k and $\sum_{k\in[K]} x_{nk} = 0$ represents declining the item. The goal is then to design an online algorithm to causally determine $\boldsymbol{x}_n$ based on the item information up to n, i.e., $\{\mathcal{I}_{n'}\}_{n'\le n}$, that maximizes the total value of all admitted items while ensuring the capacities of all knapsacks not to be violated over the time horizon.

Let $\mathcal{I} := \{\mathcal{I}_n\}_{n\in[N]}$ denote an instance of the online knapsack with departures (OKD). Given $\mathcal{I}$, the offline problem is cast as

$$\max_{x_{nk}} \quad \sum\nolimits_{n\in[N]}\sum\nolimits_{k\in[K]} v_{nk}x_{nk}, \tag{1a}$$

$$\text{s.t.} \quad \sum\nolimits_{n\in[N]:t\in\mathcal{T}_{nk}} w_{nk}x_{nk} \le C_k, \forall k\in[K], t\in[T], \tag{1b}$$

$$\sum\nolimits_{k\in[K]} x_{nk} \le 1, \forall n\in[N], \tag{1c}$$

$$x_{nk}\in\{0,1\}, \forall n\in[N], k\in[K]. \tag{1d}$$

Let $\mathtt{OPT}(\mathcal{I})$ and $\mathtt{ALG}(\mathcal{I})$ denote the values obtained by the offline problem (1) and an online algorithm under the instance $\mathcal{I}$. The competitive ratio of the online algorithm is defined as the worst-case performance ratio of the offline and online algorithms, i.e., $\mathtt{CR} = \max_{\mathcal{I}} \mathtt{OPT}(\mathcal{I})/\mathtt{ALG}(\mathcal{I})$. Our goal is to design an online algorithm that can achieve the smallest competitive ratio.

Assumption 1.1 (Value density fluctuation). *The value density of each item n in knapsack k is bounded, i.e., $v_{nk}/(w_{nk}d_{nk}) \in [1,\theta_k], \forall n\in[N]$, where θ_k is the value density ratio.*

Assumption 1.2 (Duration fluctuation). *The duration of each item n in knapsack k is bounded, i.e., $d_{nk}\in[\underline{D}_k,\overline{D}_k], \forall n\in[N]$ and duration ratio is defined as $\alpha_k = \overline{D}_k/\underline{D}_k$.*

Assumption 1.3 (Upper bound of item size). *The size of each item n is upper bounded, i.e., $w_{nk}\le\varepsilon_k\le C_k, \forall k\in[K], n\in[N]$.*

2 ALGORITHMS & MAIN RESULTS

Algorithm 1 Online Algorithms for OKD (OA-OKD)

input: thresholds $\phi = \{\phi_k(\cdot)\}_{k\in[K]}$, capacities $\{C_k\}_{k\in[K]}$;
output: admission and assignment decision $\boldsymbol{x}_n = \{x_{nk}\}_{k\in[K]}$;
initialization: initial knapsack utilization $z_{kt}^{(0)} = 0, \forall k, t$;
for each item n with item information $\mathcal{I}_n$ **do**
  for each knapsack $k\in[K]$ **do**
    call Algorithm 2 to check admissibility $\hat{x}_{nk} = \mathtt{OTA}(\mathcal{I}_n, \phi_k, \{z_{kt}^{(n-1)}\}_{t\in\mathcal{T}_{nk}}, C_k)$;
  end for
  if $\sum_{k\in[K]}\hat{x}_{nk} > 0$ **then**
    admit item n and assign it to knapsack $k' = \arg\max_{k\in[K]:\hat{x}_{nk}=1} v_{nk}$;
    set $x_{nk'} = 1$ and $x_{nk} = 0, \forall k\in[K]\setminus\{k'\}$;
  else
    decline item n and set $x_{nk} = 0, \forall k\in[K]$;
  end if
  update $z_{kt}^{(n)} = z_{kt}^{(n-1)} + w_{nk}x_{nk}, \forall k\in[K], t\in\mathcal{T}_{nk}$.
end for

We propose a simple yet effective online algorithm (OA-OKD) to solve OKD in Algorithm 1. It consists of two parts: decomposing the multiple knapsack problem into the admissibility check of each individual knapsack and admission control of each individual knapsack via an online threshold-based algorithm. The key step is the admission control of items to each knapsack by calling the OTA subroutine in Algorithm 2. To check admissibility, OTA defines a threshold value (line 3) as $\Phi = \sum_{t\in\mathcal{T}} w\phi(z_t)$, where $\phi(z_t)$ can be interpreted as the marginal cost of the unit item if it stays in

Algorithm 2 Online Threshold-based Algorithm (OTA)

1: **input:** item information $\{v, w, \mathcal{T}\}$, threshold function ϕ, utilization $\{z_t\}_{t\in\mathcal{T}}$, capacity C;
2: **output:** admission decision $\hat{x}$;
3: determine a threshold value $\Phi = \sum_{t\in\mathcal{T}} w\phi(z_t)$;
4: **if** $v\ge\Phi$ and $z_t + w\le C, \forall t\in\mathcal{T}$ **then**
5:   item is admissible and set $\hat{x} = 1$;
6: **else**
7:   item is inadmissible and set $\hat{x} = 0$.
8: **end if**

the knapsack in slot t, and is a function of the real-time knapsack utilization z_t. Since OTA is fully parameterized by ϕ, the key design question is how to determine the threshold function ϕ such that OA-OKD is competitive with the offline optimum.

Theorem 2.1. *Under Assumptions 1.1-1.3, there exist parameters $\gamma_k = O(\ln(\alpha_k\theta_k)), \forall k\in[K]$, if the item size is upper bounded by $\varepsilon_k \le C_k\ln 2/\gamma_k, \forall k\in[K]$, and the threshold function is given by $\phi^{\gamma} := \{\phi_k^{\gamma_k}\}_{k\in[K]}$, where*

$$\phi_k^{\gamma_k}(z) = \exp(z\gamma_k/C_k) - 1, z\in[0,C_k], \forall k\in[K], \tag{2}$$

then the competitive ratio of Algorithm 1 is $O(\ln(\alpha\theta))$, where $\theta = \max_{k\in[K]}\theta_k$ and $\alpha = \max_{k\in[K]}\alpha_k$.

Theorem 2.2. *There is no online algorithm that can achieve a competitive ratio smaller than $\Omega(\ln(\alpha\theta))$ for the online multiple one-dimensional knapsacks with departures.*

Combining the upper bound result in Theorem 2.1 and the lower bound result in Theorem 2.2, we conclude that Algorithm 1 achieves an order-optimal competitive ratio for OKD. In the full paper [1], this proposed algorithm can be further extended to consider the multi-dimensional item size and also achieve the order-optimal competitive ratio. In addition, we additionally design a data-driven online algorithm that can adaptively select the parameter γ such that the overall algorithm can work well on practical instances and, in the meanwhile, provide the worst-case guarantees.

ACKNOWLEDGMENTS

Danny H.K. Tsang acknowledges the support received from the Hong Kong Research Grant Council (RGC) General Research Fund (Project 16202619 and Project 16211220) and Guangdong Research Project under Grant 2021JC02X149. Adam Wierman acknowledges the support received from NSF grants (CNS-2146814, CPS-2136197, CNS-2106403, and NGSDI-210564) and the additional support from Amazon AWS. Mohammad Hajiesmaili's research is supported by NSF grants (CNS-2106299, CNS-2102963, CPS-2136199, NGSDI-2105494, and CAREER-2045641). The work of John C.S. Lui is supported in part by the RGC's SRFS2122-4S02. Don Towsley acknowledges the support received from Army Research Laboratory under Cooperative Agreement W911NF-17-2-0196 (IoBT CRA).

REFERENCES

[1] Bo Sun, Lin Yang, Mohammad Hajiesmaili, Adam Wierman, John CS Lui, Don Towsley, and Danny HK Tsang. 2022. The Online Knapsack Problem with Departures. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 6, 3 (2022), 1–32.