Competitive Online Optimization with Multiple Inventories: A Divide-and-Conquer Approach

Qiulin Lin, Yanfang Mo, Junyan Su, and Minghua Chen
School of Data Science, City University of Hong Kong

ABSTRACT

We study an online inventory trading problem where a user seeks to maximize the aggregate revenue of trading multiple inventories over a time horizon. The trading constraints and concave revenue functions are revealed sequentially in time, and the user needs to make irrevocable decisions. The problem has wide applications in various engineering domains. Existing works employ the primal-dual framework to design online algorithms with sub-optimal, albeit near-optimal, competitive ratios (CR). We exploit the problem structure to develop a new divide-and-conquer approach to solve the online multi-inventory problem by solving multiple calibrated single-inventory ones separately and combining their solutions. The approach achieves the optimal CR of $\ln\theta+1$ if $N \leq \ln\theta+1$, where N is the number of inventories and θ represents the revenue function uncertainty; it attains a CR of $1/[1-e^{-1/(\ln\theta+1)}] \in [\ln\theta+1, \ln\theta+2)$ otherwise. The divide-and-conquer approach reveals novel structural insights for the problem, (partially) closes a gap in existing studies, and generalizes to broader settings. For example, it gives an algorithm with a CR within a constant factor to the lower bound for a generalized one-way trading problem with price elasticity with no previous results. When developing the above results, we also extend a recent CR-Pursuit algorithmic framework and introduce an online allocation problem with allowance augmentation, both of which can be of independent interest.

1 INTRODUCTION

We consider an important class of online optimization problem, optimizing the trading or allocation of limited resources, such as inventories, cryptocurrency, budgets, or electric power, across a multi-round decision period with dynamic per-round revenue functions and allocation conditions. In the problem, the online decision maker has N capacity-limited inventories to trade in a decision period of T slots, e.g., airlines selling flight tickets with different classes. At each slot t, there are two types of trading constraints. The first is the total allocation of all inventories bounded by the allowance A_t. The second is that the allocation of each inventory i is bounded by the rate limit $\delta_{i,t}$. The problem has applications in different domains, including ad allocation, bilateral trading, energy storage management, EV charging, etc.

We consider the online optimization problem with multiple inventories (OOIC-M) and formulate it in (1). In OOIC-M, we optimize the inventory trading $\{v_{i,t}\}_{i\in[N],t\in[T]}$ to achieve the maximum total revenue subjecting to the capacity constraint of each inventory (2), the allowance constraint at each slot (3), and the allocation rate limit constraint for each inventory at each slot (4).

$$\text{OOIC-M}: \max \sum_{i\in[N]}\sum_{t\in[T]} g_{i,t}(v_{i,t}) \tag{1}$$

$$\text{s.t.} \sum_t v_{i,t} \leq C_i, \forall i \in [N], \tag{2}$$

$$\sum_i v_{i,t} \leq A_t, \forall t \in [T], \tag{3}$$

$$0 \leq v_{i,t} \leq \delta_{i,t}, \forall t \in [T], i \in [N], \tag{4}$$

We are interested in the online setting that at each slot t, the online decision maker without the information of the decision period T is fed the revenue functions $\{g_{i,v}(\cdot)\}_{i\in[N]}$, the allowance A_t and the rate limits $\{\delta_{i,t}\}_{i\in[N]}$. We need to irrevocably determine the allocation at slot t, i.e., $\{v_{i,t}\}_{i\in[N]}$. After that, if the decision period ends, we stop and know the information of T. Otherwise, we move to the next slot and continue the trading. We consider the following set of revenue functions ($\triangleq \mathcal{G}$),

- $g_{i,t}(\cdot)$ is concave and differentiable with $g(0) = 0$;
- $g'_{i,t}(v_{i,t}) \in [p_{\min}, p_{\max}], \forall v_{i,t} \in [0, \delta_{i,t}]$.

We consider that $p_{\max} \geq p_{\min} > 0$ and denote $\theta = p_{\max}/p_{\min}$. The revenue functions capture the case where the marginal revenue of allocating more inventory is non-increasing in the allocation amount but always between $p_{\min}$ and $p_{\max}$.

We use *Competitive Ratio* (CR) as the performance metric for online algorithms. The CR of an algorithm $\mathcal{A}$ is defined as the worst-case performance ratio between the offline optimal and the online objective under the algorithm, i.e.,

$$\mathcal{CR}(\mathcal{A}) = \sup_{\sigma\in\Sigma} \frac{OPT(\sigma)}{ALG(\sigma)}, \tag{5}$$

where $\sigma \in \Sigma$ denotes an input and Σ represents all possible input; $OPT(\sigma)$ and $ALG(\sigma)$ denote the offline optimal objective and the online objective of $\mathcal{A}$ under input σ, respectively.

2 OUR APPROACH AND SELECTED RESULTS

We propose a divide-and-conquer approach for deriving online algorithms for OOIC-M. The general idea is that we can optimize OOIC-M by first allocating the allowance to the inventories and then separately optimizing the allocation of each inventory. More specifically, we define the following subproblem for each $i \in [N]$,

$$\text{OOIC-S}_i: \max \sum_t \tilde{g}_{i,t}(v_{i,t}) \tag{6}$$

$$s.t. \sum_t v_{i,t} \leq C_i \tag{7}$$

$$0 \leq v_{i,t} \leq a_{i,t}, \forall t, \tag{8}$$

where $a_{i,t}$ is the allocated allowance to user i at slot t. $\tilde{g}_{i,t}(v_{i,t})$ is another algorithmic design space that allows us to exploit the

---

Corresponding author: Minghua Chen. The results were presented at ACM SIGMETRICS/IFIP PERFORMANCE '22 [1], and the full version is published in ACM POMACS [2].

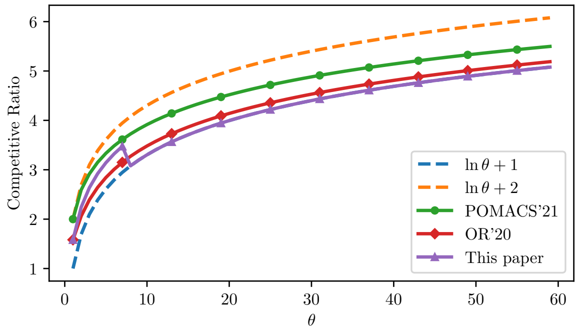

Figure 1: The competitive ratios obtained by OR'20 [5], POMACS'21 [6], and ours. We fix $N = 3$ and vary θ from 1 to 60.

online allowance augmentation. Following this structure, we can design an online algorithm with two main steps at each slot t,

(1) Step-I: Determine the allowance allocation, $\{\hat{a}_{i,t}\}_{i\in[N]}$, irrevocably.
(2) Step-II: Determine the inventory allocation for each online OOIC-S$_i$, $\{\hat{v}_{i,t}\}_{i\in[N]}$, irrevocably.

We note that this divide-and-conquer approach allows us to tackle the two main challenges of the problem separately. First, the revenue functions come online while the allocation across the decision period is coupled due to the capacity constraint for each inventory, which we can handle in Step-II. Second, the online allowance constraints and the rate constraints couple the decisions across the inventories, which we tackle in Step-I.

To achieve a good online performance under the divide-and-conquer structure, we can 1) (Step-I) find allowance allocation such that the sum of the optimal objectives of the subproblems is close to the optimal revenue of the original problem OOIC-M and 2) (Step-II) design inventory allocation to achieve a close-to-optimal online performance in each subproblem. Under such a goal, we note that the problem in Step-II is close-related to the online ad allocation problem with free disposal [3]. And Step-II is reduced to the online optimization problem with a single inventory [4].

Our single-parametric online algorithm A&P(π) with π as a parameter to be specified stands for Allowance Augmented Allocation and Pursuit. In Step-II, following the framework CR-Pursuit [4], we determine the online inventory allocation $\hat{v}_{i,t}$ for each subproblem to maintain the offline-to-online performance ratio being π at all slots. A useful property is that the online solution at most utilizes $1/\pi$ of the allocated allowance/rate limit. Given that, for Step-I, we consider an online allowance augmentation scenario where the online decision maker can utilize π-time allowance and is subject to π-time relaxer rate limit constraints, which generalizes the existing study [3]. When $N \leq \pi$, we can directly allocate the allowance to each subproblem as its rate limit $\delta_{i,t}$. When $N > \pi$, we design a scheme of Allowance Allocation at slot t with Augmentation, $\triangleq$ AAt-A(π). Please refer to the main paper for more details. In particular, when choosing $\pi = \ln\theta + 1$, A&P($\ln\theta + 1$) achieves a close-to-optimal performance guarantee.

Theorem 1. *Our online algorithm A&P($\ln\theta + 1$) achieves the CR,*

$$\mathcal{CR}_1(A\&P(\ln\theta+1)) = \begin{cases} \ln\theta + 1, & N \leq \ln\theta + 1 \\ 1/[1 - e^{-1/(\ln\theta+1)}], & otherwise. \end{cases} \quad (9)$$

Compared with a lower bound $\ln\theta + 1$, the CR we obtain matches the lower bound when $N \leq \ln\theta + 1$ and is within an additive constant of one to the lower bound, otherwise. The problem has also been studied in [6] and [5]. We provide an illustration and comparison of the CRs achieved by [5, 6] and ours in Fig. 1. We note that both studies are threshold-based algorithms following the online primal-and-dual framework. We also discuss their empirical performance and algorithmic behaviors in the main paper.

3 CONCLUDING REMARKS

Our consideration in Step-I generalizes the existing studies on the online maximum allocation problem and the online ad allocation problem with free disposal [3] to the online allowance augmentation scenario. We also discuss how to generalize our approach to different sets of revenue functions that appear in real-world applications, e.g., IoT device information uploading and one-way trading with price elasticity. We are interested in exploring the divide-and-conquer approach under various online optimization problems that involve multiple inventories or participants. Another interesting future direction is adapting and generalizing our approach to various real-world applications [7].

ACKNOWLEDGEMENT

The work presented in this paper was supported in part by a General Research Fund from Research Grants Council, Hong Kong (Project No. 11206821), an InnoHK initiative, The Government of the HKSAR, and Laboratory for AI-Powered Financial Technologies. Qiulin Lin would like to thank Dr. Hanling Yi and Dr. John Z. F. Pang for the discussions at the initial stage of the study.

REFERENCES

[1] Qiulin Lin, Yanfang Mo, Junyan Su, and Minghua Chen. Competitive online optimization with multiple inventories: A divide-and-conquer approach. *SIGMETRICS Perform. Eval. Rev.*, 50(1):83–84, jul 2022.

[2] Qiulin Lin, Yanfang Mo, Junyan Su, and Minghua Chen. Competitive online optimization with multiple inventories: A divide-and-conquer approach. *Proc. ACM Meas. Anal. Comput. Syst.*, 6(2), jun 2022.

[3] Jon Feldman, Nitish Korula, Vahab Mirrokni, S. Muthukrishnan, and Martin Pál. Online ad assignment with free disposal. In Stefano Leonardi, editor, *Internet and Network Economics*, pages 374–385, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.

[4] Qiulin Lin, Hanling Yi, John Pang, Minghua Chen, Adam Wierman, Michael Honig, and Yuanzhang Xiao. Competitive online optimization under inventory constraints. *Proc. ACM Meas. Anal. Comput. Syst.*, 3(1), March 2019.

[5] Will Ma and David Simchi-Levi. Algorithms for online matching, assortment, and pricing with tight weight-dependent competitive ratios. *Operations Research*, 68(6):1787–1803, 2020.

[6] Bo Sun, Ali Zeynali, Tongxin Li, Mohammad Hajiesmaili, Adam Wierman, and Danny H.K. Tsang. Competitive algorithms for the online multiple knapsack problem with application to electric vehicle charging. *Proc. ACM Meas. Anal. Comput. Syst.*, 4(3), December 2020.

[7] Yanfang Mo, Qiulin Lin, Minghua Chen, and S. Joe Qin. Optimal online algorithms for peak-demand reduction maximization with energy storage. In *Proc. ACM e-Energy*, pages 73–83, 2021.