Online Conversion with Group Fairness Constraints

Hossein Nekouyan Jazi*, Faraz Zargari*
University of Alberta
{nekouyan, fzargari}@ualberta.ca

Bo Sun
University of Waterloo
bo.sun@uwaterloo.ca

Xiaoqi Tan
University of Alberta
xiaoqi.tan@ualberta.ca

ABSTRACT

In this paper, we initiate the study of an online conversion problem that incorporates group fairness guarantees. This problem aims to distribute a resource with fixed capacity to a sequence of buyers based on their offered prices. Each buyer belongs to a distinct group, and the objective is to maximize revenue while ensuring fairness across groups by guaranteeing that each group receives a predetermined quantity of resources. We propose a novel threshold-based online algorithm and prove that it achieves the optimal competitive ratio with fairness guarantees.

1. INTRODUCTION

The online conversion is a classic resource allocation problem in the literature of online algorithms [1]. In this problem, a decision maker aims to sell a fixed amount of divisible resources to a sequence of online arriving buyers. Each buyer offers their conversion price for purchasing per-unit resources upon arrival, and the seller immediately and irrevocably decides the amount of resources to sell without knowing the prices and the number of future buyers. The goal is to maximize the seller's total revenue. Under competitive analysis, this problem and its variants have been well-studied, and it is known that a threshold-based algorithm can achieve optimal competitive performance [2].

Although existing algorithms can achieve optimal performance in maximizing the seller's revenues, there is a concern that specific subgroups of buyers (defined by characteristics such as age, financial status, or education) may be consistently disadvantaged in obtaining the resource due to their limited purchasing power and delayed access to information. This issue becomes particularly critical when the seller is allocating essential resources (e.g., energy, food). Consequently, it is increasingly important to incorporate fairness guarantees across different groups in algorithm design. Some recent research has begun to address this need by integrating fairness guarantees into resource allocation problems, including knapsack problems [3] and online matching [4].

Motivated by the algorithmic challenges posed by the aforementioned societal issues, we introduce and study online conversion with multi-class arrivals and group fairness constraints. In this problem, each buyer belongs to a specific class (or group), and our objective is to ensure fairness by guaranteeing that each class receives a predetermined amount of resources. We propose a novel, modified threshold-based algorithm tailored to this challenge and demonstrate that it achieves optimal competitive performance under the given fairness constraint. Furthermore, our algorithms are promising for the allocation of essential resources in various contexts. For instance, in energy systems, a charging station might distribute energy to electric vehicles across diverse categories, such as private, commercial, and public vehicles. Ensuring group fairness in such scenarios is crucial for equitable resource distribution, which highlights the practical relevance of our proposed approach.

*These authors contributed equally to this work

2. PROBLEM STATEMENT

We formally define online conversion with multi-class arrivals, termed Online Multi-class Conversion (OMcC), as follows: A seller with an initial resource of B is trying to maximize its revenue by trading its resource to a sequence of buyers arriving one at time. Upon the arrival of buyer $t \in \{1, 2, \dots, T\}$, she offers a price p_t; an immediate and irrevocable decision $x_t \in [0, 1]$ must be made regarding the amount of resource to be sold to this buyer. The goal is to maximize the total revenue in the end, i.e., $\sum_t p_t x_t$, subject to the resource constraint $\sum_t x_t \le B$.

In OMcC, each buyer t belongs to a specific class $j_t \in [K]$. We assume that there exists a class-dependent finite support for possible prices within each class.

Assumption 1. *We assume that the prices of buyers in each class $j \in [K]$ are normalized and bounded within $[1, \theta_j]$, i.e., $p_t \in [1, \theta_j]$ holds for all $j \in [K]$ and $t \in [T]$ if $j_t = j$.*

We refer to θ_j as the *fluctuation ratio* of class j. Intuitively, $\theta_j \ge 1$ holds for all $j \in [K]$ and a smaller fluctuation ratio indicates that arrivals within that class tend to be more homogeneous. We also assume w.l.o.g. that $\theta_1 \le \theta_2 \le \cdots \le \theta_K$.

Since the seller aims to maximize her revenue, resources tend to be allocated to buyers with high prices, even if they are all from the same class. To avoid such unfair treatment of different classes, we seek the following fairness guarantee, referred to as *group fairness by quantity* (GFQ) [3].

Definition 1 (Group Fairness by Quantity)**.** *The total amount of resource allocated to buyers from group $j \in [K]$ is at least m_j for all $j \in [K]$.*

It is worth noting that when $m_j = B/K$ holds for all $j \in [K]$, GFQ reduces to the well-known *proportional fairness* (by groups/classes); namely, each class will receive at least $1/K$ fraction of the total resource. The fairness requirement $\boldsymbol{m} := \{m_j\}_{j\in[K]}$ is provided beforehand, and our objective is to study how to design online algorithms for a given $\boldsymbol{m}$.

In the online setting, the sequence of price arrivals does not necessarily follow any pattern and could be adversarially chosen. Following the *worst-case competitive analysis* framework, the performance of an online algorithm is quantified by its *competitive ratio.* Given an arrival instance $\sigma = \{p_1, \cdots, p_T\}$, let us denote by $\mathsf{OPT}(\sigma)$ the optimal revenue achieved in the offline setting when the information of the arrival sequence σ is known beforehand. Mathematically, $\mathsf{OPT}(\sigma)$ can be obtained by solving the following linear program.

$$\begin{aligned}\mathsf{OPT}(\sigma) = \max_{x_t\in[0,1]} & \sum\nolimits_{t\in[T]} p_t \cdot x_t \\ \text{s.t.} & \sum\nolimits_{t\in[T]} x_t \le B, \\ & \sum\nolimits_{t\in[T]} x_t \cdot \mathbf{1}_{\{j_t=j\}} \ge m_j, \forall j \in [K].\end{aligned}$$

Let $\mathsf{ALG}(\sigma)$ denote the revenue achieved by an online algorithm ALG. Our goal is to develop online algorithms that minimize the worst-case competitive ratio, i.e., $\min_{\mathsf{ALG}} \max_{\sigma\in\Omega} \frac{\mathsf{OPT}(\sigma)}{\mathsf{ALG}(\sigma)}$, where Ω represents the family of arrival sequences that satisfy Assumption 1.

3. MAIN RESULTS

3.1 A Threshold-based Algorithm

We propose a threshold-based algorithm in Algorithm 1 for OMcC with group fairness guarantee. Upon receiving the first set of $\lceil m_j \rceil$ buyers from each class $j \in [K]$, Algorithm 1 ensures the corresponding fairness guarantee sought for that class based on Definition 1. As a result, the M-portion of total resource B is reserved to meet the fairness requirements, where $M = \sum_j m_j$. The remaining $(B - M)$-portion of the total resource is then allocated based on the threshold function $\phi(u_t) : [0, B - M] \to [1, \theta_K]$, where u_t represents the utilization level of the algorithm from the $(B - M)$-portion of the resource up to time t.

Algorithm 1: Threshold-based algorithm for OMcC with group fairness guarantee

Input: B; $(m_j, \theta_j), \forall j \in [K]$.
Initialization: $u_0 = 0, f_0^j = 0, \forall j \in [K]$.
while *buyer t arrives* **do**
  Obtain the price and class information of buyer t: p_t and j_t ;
  Set $x_t = 0, y_t = 0$
  if $f_{t-1}^{j_t} < m_{j_t}$ **then**
    $y_t = \min\{1, m_{j_t} - f_{t-1}^{j_t}\}$;
    Update $f_t^{j_t} = f_{t-1}^{j_t} + y_t$;
  end
  Decide the conversion according to:
  if $y_t < 1$ **then**
    if $p_t \ge \phi(u_{t-1})$ **then**

$$x_t = \min\{ \arg\max_{a\in[0,1-y_t]} \{a \cdot p_t - \int_{u_{t-1}}^{u_{t-1}+a} \phi(\eta) d\eta\}, B - M - u_{t-1}\}$$

    end
  end
  Update the cumulative conversion: $u_t = u_{t-1} + x_t$.
  Update the conversion amount of buyer t: $x_t = x_t + y_t$
end

Theorem 1 below shows that Algorithm 1 can achieve the optimal competitive ratio if the threshold function $\phi : [0, B - M] \to [1, \theta_K]$ is strategically designed.

Theorem 1. *There exits a threshold function ϕ such that Algorithm 1 is optimal in the sense that no online algorithm can achieve a smaller competitive ratio under Assumption 1 with GFQ guarantee.*

In the following section, we formally present our design of the optimal threshold functions for the cases when $K = 1$ and $K = 2$, and briefly discuss our general results regarding $K \ge 2$ due to page limit.

3.2 Optimal Threshold Function Design

3.2.1 *Single-Class Case* ($K = 1$)

We start with the basic case that there is only one class of buyers with prices bounded within $[1, \theta]$. In this case, the definition of GFQ basically requires that at least m units of the total resources be allocated to buyers regardless of their prices. Theorem 2 below shows the design of the optimal threshold function for OMcC with $K = 1$.

Theorem 2. *The competitive ratio of Algorithm 1 when $K = 1$ can be determined in the following two cases.*

- *When $m \le \frac{B}{1+\ln\theta}$, let $\alpha_0 := 1 + \ln\theta$, Algorithm 1 is α_0-competitive if the threshold function ϕ is given by*

$$\phi(u) = \begin{cases} 1 & u \in [0, \frac{B}{\alpha_0} - m], \\ e^{\left(\frac{\alpha_0\cdot(u+m)}{B} - 1\right)} & u \in (\frac{B}{\alpha_0} - m, B - m]. \end{cases}$$

- *When $m > \frac{B}{1+\ln\theta}$, let $\alpha_1 := \frac{B}{B-m} \cdot W(\theta(\frac{B}{m} - 1))$, where $W(\cdot)$ is the Lambert function, Algorithm 1 is α_1-competitive if the threshold function ϕ is given by*

$$\phi(u) = \frac{\alpha_1 m}{B} \cdot e^{\left(\frac{\alpha_1\cdot u}{B}\right)}, \quad \forall u \in [0, B - m].$$

Note that when m is small, the best-possible competitive ratio for OMcC with GFQ guarantee is $1 + \ln\theta$. As m approaches B, α_1 converges to θ. This implies that when $m = B$, no online algorithm can effectively reserve any portion of its resource for future buyers, i.e., no online algorithm can outperform the trivial competitive ratio θ.

3.2.2 *Two-Class Case* ($K = 2$)

In this section, we extend the results of the previous section to the two-class arrival case. Recall that we assume w.l.o.g. that $\theta_2 \ge \theta_1$. Theorem 3 below shows the design of the optimal threshold function for OMcC with $K = 2$.

Theorem 3. *The competitive ratio of Algorithm 1 when $K = 2$ can be determined in the following three cases.*

- *When $M \le \frac{B}{\alpha_0}$, where α_0 is defined as follows:*

$$\alpha_0 := 1 + \ln\theta_2 - \frac{m_1}{B} \ln\frac{\theta_2}{\theta_1},$$

Algorithm 1 is α_0-competitive if ϕ is given by

$$\phi(u) = \begin{cases} 1 & u \in [0, \frac{B}{\alpha_0} - M] \\ e^{\left(\frac{\alpha_0(u+M)-B}{B}\right)} & u \in (\frac{B}{\alpha_0} - M, \frac{B}{\alpha_0} - M + \frac{B}{\alpha_0}\ln\theta_1], \\ e^{\left(\frac{\alpha_0(u+M)-B-m_1\ln\theta_1}{B-m_1}\right)} & u \in (\frac{B}{\alpha_0} - M + \frac{B}{\alpha_0}\ln\theta_1, B - M]. \end{cases}$$

- *When $M \in (\frac{B}{\alpha_1}, \frac{B}{\alpha_1}\theta_1]$, where α_1 is defined as follows:*

$$\alpha_1 := \frac{B}{B-M} W\left(\frac{\theta_2(B-M)}{M} e^{\left(-\frac{m_1}{B}\ln\frac{\theta_2}{\theta_1}\right)}\right),$$

Algorithm 1 is α_1-competitive if ϕ is given by

$$\phi(u) = \begin{cases} v^* e^{\left(\frac{\alpha_1 u}{B}\right)} & u \in [0, \frac{B}{\alpha_1}\ln\frac{\theta_1}{v^*}], \\ v^* e^{\left(\frac{\alpha_1 u - m_1 \ln\frac{\theta_1}{v^*}}{B-m_1}\right)} & u \in [\frac{B}{\alpha_1}\ln\frac{\theta_1}{v^*}, B - M], \end{cases}$$

where $v^* = \alpha_1 M/B$.

- *When $M \in (\frac{B}{\alpha_2}\theta_1, B]$, where α_2 is defined as follows:*

$$\alpha_2 := \theta_1 \frac{m_1}{M} + \frac{B-m_1}{B-M} W\left(\frac{\theta_2(B-M)}{M} e^{\left(-\frac{\theta_1 m_1(B-M)}{(B-m_1)M}\right)}\right),$$

Algorithm 1 is α_2-competitive if ϕ is given by

$$\phi(u) = v^* e^{\left(\frac{\alpha_2 u}{B-m_1}\right)}, \qquad \forall u \in [0, B - M],$$

where $v^* = (\alpha_2 M - m_1\theta_1)/(B - m_1)$.

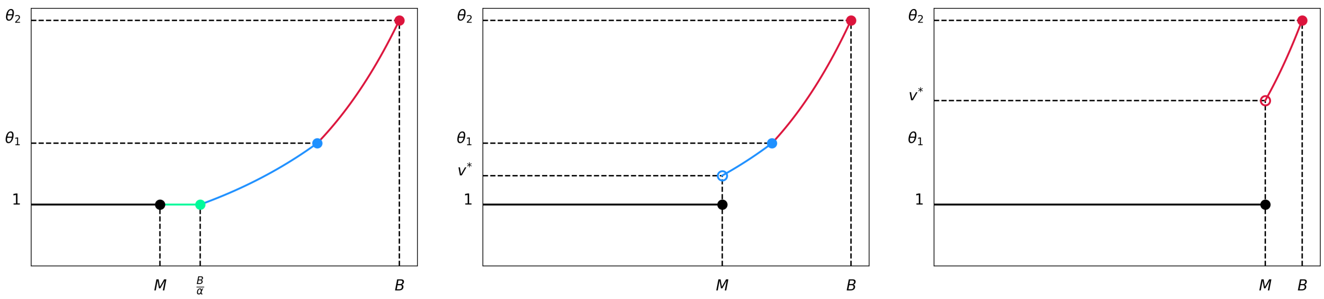

(a) $m_2 = 5$, $\alpha_0 = 2.18$ (b) $m_2 = 35$, $\alpha_1 = 2.26$ (c) $m_2 = 60$, $\alpha_2 = 2.76$

Figure 1: Threshold function dynamics with $\theta_1 = 2$, $\theta_2 = 4$, $B = 100$ and $m_1 = 30$ and changing m_2.

Theorem 3 shows that when M approaches B, α_2 converges to $\theta_1 \frac{m_1}{B} + \theta_2 \frac{B-m_1}{B}$. The intuition is that when $M = B$, due to the fairness requirement, no online algorithm can effectively reserve any portion of its resource for future buyers. Thus, in the worst case, no online algorithm can perform better than B, while the offline optimal algorithm may achieve the maximum revenue $\theta_1 m_1 + \theta_2 (B - m_1)$, leading to the worst-case competitive ratio $\theta_1 \frac{m_1}{B} + \theta_2 \frac{B-m_1}{B}$.

Remark 1. *At first glance, it might appear that the three intervals of M that define the three cases in Theorem 3 are not continuous and do not fully cover the range of $[0, B]$. However, as the value of M approaches the end-point of one interval (e.g., the end-point $\frac{B}{\alpha_1}\theta_1$ of the second interval), the start-point of the next interval (e.g., $\frac{B}{\alpha_2}\theta_1$) also converges to the end-point of the last interval. This observation is illustrated in Figure 2, where α_0 and α_1 converge to the same value as m_2 approaches* 15.91.

3.2.3 *General Cases* ($K \geq 2$)

Our results can be generalized to scenarios with more than two classes, i.e., $K \geq 2$. An informal statement of our most general results is given in Theorem 4 below.

Theorem 4 (Informal)**.** *For* OMcC *with $K \geq 2$ classes, there exists a threshold function ϕ with at most $K + 1$ segments such that Algorithm 1 achieves the optimal competitive ratio. The number of segments and the competitive ratio depend on the sum of the minimum allocation requirements (i.e., $\boldsymbol{m} := \{m_j\}_{j \in [K]}$) prescribed by the GFQ requirement.*

4. NUMERICAL RESULTS

In this section, we focus on OMcC with two classes and aim to numerically illustrate the design of the optimal threshold functions given in Theorem 3. For the M-portion of the resource, where a buyer is accepted irrespective of their price, we set the threshold to be 1. As for the remaining portion, we use the threshold function outlined in Theorem 3. In Figure 1, we fix all the parameters of the problem except the value of m_2. For three different values of m_2, Figure 1 shows the corresponding threshold function that Algorithm 1 uses based on Theorem 3. It can be seen that depending on the value of M, the threshold function may be designed in three different cases. For each case, the threshold function contains at most three segments (illustrated in different colors in Figure 1). Moreover, the threshold of each segment increases w.r.t. the resource utilization (i.e., the portion of the resources that have been allocated).

In Figure 2, we fix all the parameters except the value of m_2 and show how the competitive ratio of Algorithm 1 changes with variations in M. As m_2 increases, the competitive ratio of Algorithm 1, denoted by CR^*, continuously increases. This outcome was foreseeable, since we should allocate a larger share of resources to buyers regardless of their prices to ensure fairness guarantee. Moreover, we can observe that CR^* switches from α_0 to α_1 and α_2 w.r.t. the increase of m_2 (or equivalently, the increase of M).

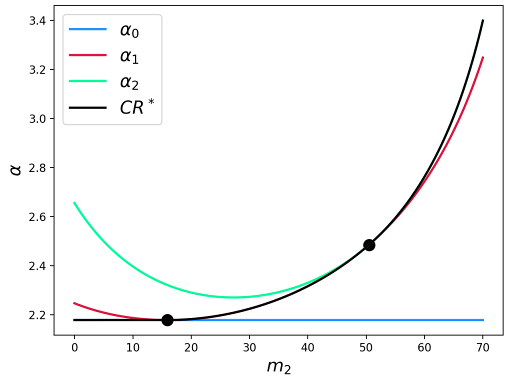

Figure 2: Illustrating the competitive ratio as a function of m_2. CR^* denotes the optimal competitive ratio.

5. CONCLUSIONS AND FUTURE WORK

In this work, we initiated the study of online conversion with multi-class arrivals and group fairness constraints. We first considered the simple one-class case and then extended our results to the two-class case. We also showed that it is possible to extend our results to the general K-class case with $K \geq 2$. As a natural next step, it is interesting to investigate online multi-class conversion with other group fairness constraints. Another interesting direction is to extend our results to consider the trade-off between group fairness and individual fairness within each class.

6. REFERENCES

[1] E. Mohr, I. Ahmad, and G. Schmidt, "Online algorithms for conversion problems: a survey," *Surveys in Operations Research and Management Science*, vol. 19, no. 2, pp. 87–104, 2014.

[2] A. Lechowicz, N. Christianson, B. Sun, N. Bashir, M. Hajiesmaili, A. Wierman, and P. Shenoy, "Online conversion with switching costs: Robust and learning-augmented algorithms," in *ACM SIGMETRICS/IFIP PERFORMANCE*, 2023.

[3] D. Patel, A. Khan, and A. Louis, "Group fairness for knapsack problems," in *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, ser. AAMAS '21, 2021.

[4] H. Hosseini, Z. Huang, A. Igarashi, and N. Shah, "Class fairness in online matching," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 5, 2023, pp. 5673–5680.