Online Reusable Resource Allocation with Multi-class Arrivals

In today's world, many businesses involve not simply selling products, but instead renting them out for short periods. Sharing economy, including car renting and cloud computing, has become deeply integrated into our daily life. In addition, applications including hotel capacity management, appointment booking involve allocating reusable resources. A key challenge in these settings is the volatility and unpredictability of customers' arrivals, for example due to pandemics and military conflicts.

We study an adversarial online reusable resource allocation problem. There is a single type of resource pool with C units for serving $M(M \geq 2)$ classes of customers. Each class m is associated with a price willing to pay per unit time $r^{(m)}$ and a usage duration $D^{(m)}$. We assume that $0 < r^{(m)} \leq r^{(m+1)}$, $0 < D^{(m)} \leq D^{(m+1)}$ and $r^{(m)}D^{(m)} < r^{(m+1)}D^{(m+1)}$ for all $m \in \{1, \dots, M-1\}$. The time horizon is partitioned so that one customer arrives at each time period. This is without loss of generality since we model empty arrival as class 0 with $r^{(0)} = D^{(0)} = 0$.

At each time step t, one customer arrives with $r_t D_t$ being the price willing to pay for the whole usage. The agent observes the customer type and then decide to accept or to reject the customer. If the decision is to accept, then one unit of resource will be allocated and occupied from time t to $t + D_t - 1$, and a revenue of $r_t D_t$ is earned. If the decision is to reject, then no unit is allocated and no revenue is earned. The future arrival sequence is adversarially chosen, which means that the arrival could follow some changing distributions which are unknown to the agent. The objective is to maximize the total revenue, subject to the constraint that the number of units being occupied at each time step should not exceed the total capacity. Our model is an adversarial variant to the stochastic model by Levi and Radovanovi (2010). When we set the usage durations to be infinite in our model, the specialized model coincides with Ball and Queyranne (2009). Our model can also be cast as a pricing variant to the online assortment problem (see Feng et al. 2019 for example). We quantify the performance of an online algorithm with its Competitive Ratio (CR), which is the minimum ratio between the revenue achieved by the algorithm and the revenue achieved by the optimal offline algorithm who knows the full sequence from the beginning, where the minimum is taken over all possible instances.

We propose an original online policy, the Protection Level Policy for Reusable Resources (PLP-RR). At each time step t, PLP-RR sets a rejection price R_t, and only accept customers with a price willing to pay strictly higher than the rejection price. The rejection price depends on quantities $\{I_t^{(m)}\}_{m=0}^{M-1}$, where $I_t^{(m)} = |\{s \in \{1, \dots t-1\}: R_s = r^{(m)}D^{(m)}, s + D_s - 1 \geq t\}|$. To interpret, $I_t^{(m)}$ is the number of resource units that are (1) allocated under the rejection price $r^{(m)}D^{(m)}$, (2) still being occupied at time t. PLP-RR takes into inputs a set of parameters $\{C^{(m)}\}_{m=0}^{M-1}$, and rejection price R_t is set as $r^{(m)}D^{(m)}$, where m is the smallest integer that satisfies $I_t^{(m)} < C^{(m)}$. The intuition of the algorithm is that the total capacity is partitioned into M classes such that there are $C^{(m-1)}$ number of class m units for all $m \in \{1, \dots, M\}$. Class m units can only be allocated to customers of class $m, \dots, M$. That is, class 1 units can be allocated to all non-empty customers, class 2 units can be allocated to a customer of class 2 or above. The definition of class m units is the units to be allocated with rejection price being $r^{(m-1)}D^{(m-1)}$. As a result of the algorithm, the parameters $\{C^{(m)}\}_{m=0}^{M-1}$ serves as booking limits in the sense that at any time t, the number of customers who has a price willing to pay at most $r^{(m+1)}D^{(m+1)}$, and still occupying a unit of resource at time t is at most $\sum_{n=0}^{m} C^{(n)}$.

We prove that PLP-RR achieves asymptotic optimality with large capacity when the usage duration of any class divides that of any higher class, which is what we call divisible durations. The loss factor, which approaches to zero when capacity is large, is due to the rounding the theoretically optimal booking limits to obtain a set of integer booking limits. PLP-RR achieves optimality if the theoretically optimal booking limits are a set of integers and therefore no rounding is needed. The case where all classes of customers having the same usage duration also satisfies the duration divisibility condition. We highlight that such loss factors are not unique to our result, and similar loss factors appear in Ball and Queyranne (2009) for non-reusable resources, which is a special case in our model. When the durations are general, PLP-RR achieves an asymptotic competitive ratio that is within a factor of half of optimality. To be exact, the factor ratio of optimality depends on the durations of the classes and is lower bounded by half.

We use a primal-dual approach to prove the competitive ratio, where we construct a linear program whose dual is an upper bound to the offline optimum. To be exact, we propose an original Auxiliary Algorithm (AA) where a feasible dual solution is constructed, and the ratio between primal objective value which is the total revenue, and dual objective value achieved by the dual variables constructed in AA are lower bounded by the competitive ratio. Our primal-dual approach is fundamentally different from traditional primal-dual online algorithms from Buchbinder and Naor (2009) in the sense that traditional primal-dual algorithms involve an online update of the dual variables, which is used to guide the online decisions. However, our analysis only sets the values of the dual variables in AA after the whole online process. The dual is only used for analytical purpose but not for online decisions. We are hopeful that such a novel primal-dual approach might find applications in other online problems.

We then show that similar methodology can be extended to the advance booking problem, where usage only starts after a period of time since the customer arrival. We use s_t to denote the starting time of usage for customer arrived at time t. Therefore, the requested usage period for the customer at time t will be from time s_t to $s_t + D_t - 1$. We propose the Protection Level Policy for Advance Booking (PLP-AB), in which the rejection price depends on $\{I_{t,\tau}^{(m)}\}_{\tau\in[s_t,s_t+D_t-1],m\in[0,M-1]}$, where $I_{t,\tau}^{(m)}$ denotes the number of units allocated until time $t-1$, that is still being occupied at time τ, and with the rejection price being $r^{(m)}D^{(m)}$ when allocated. To be exact, the customer will only be accepted if there is available unit with a class number at most that of the customer, for all the time steps in the usage period requested by the customer. This is different from PLP-RR where only the inventory time t is needed to make the decision.

We show that PLP-AB achieves an asymptotic competitive ratio that is within a factor of half of that for resource allocation with large capacity. This is proved using a similar primal-dual approach where an Auxiliary Algorithm is proposed to construct a feasible dual solution that upper bounds the offline optimum.

We conduct numerical experiments to evaluate the empirical performance of both PLP-RR and PLP-AB. We compare the algorithms with the First-Come-First-Serve (FCFS) policy under adversarial arrival with slight stochasticity and completely stochastic arrival patterns. We have three important observations. First, both algorithms have better empirical performance than competitive ratio even with slight stochasticity in the arrival pattern. Second, both algorithms perform dominate the FCFS policy under both arrival patterns for resource allocation and advance booking problem respectively. Lastly, both algorithms are more resilient than the FCFS policy when the arrival pattern changes from stochastic to adversarial, since much less performance degradations are seen.

References

Ball, Queyranne (2009) Toward robust revenue management: Competitive analysis of online booking. Operations Research 57(4):950–963.

Levi R, Radovanovi´c A (2010) Provably near-optimal LP-based policies for revenue management in systems with reusable resources. Operations Research 58(2):503–507.

Niv Buchbinder, Joseph Naor (2009) The Design of Competitive Online Algorithms via a Primal-Dual Approach. Found. Trends Theor. Comput. Sci. 3(2-3): 93-263