



Online Resource Allocation under Horizon Uncertainty

Santiago Balseiro
Columbia University
New York, USA
srb2155@columbia.edu

Christian Kroer
Columbia University
New York, USA
christian.kroer@columbia.edu

Rachitesh Kumar
Columbia University
New York, USA
rk3068@columbia.edu

ABSTRACT

We study stochastic online resource allocation: a decision maker needs to allocate limited resources to stochastically-generated sequentially-arriving requests in order to maximize reward. At each time step, requests are drawn independently from a distribution that is unknown to the decision maker. Online resource allocation and its special cases have been studied extensively in the past, but prior results crucially and universally rely on the strong assumption that the total number of requests (the horizon) is known to the decision maker in advance. In many applications, such as revenue management and online advertising, the number of requests can vary widely because of fluctuations in demand or user traffic intensity. In this work, we develop online algorithms that are robust to horizon uncertainty. In sharp contrast to the known-horizon setting, no algorithm can achieve even a constant asymptotic competitive ratio that is independent of the horizon uncertainty. We introduce a novel generalization of dual mirror descent which allows the decision maker to specify a schedule of time-varying target consumption rates, and prove corresponding performance guarantees. We go on to give a fast algorithm for computing a schedule of target consumption rates that leads to near-optimal performance in the unknown-horizon setting. In particular, our competitive ratio attains the optimal rate of growth (up to logarithmic factors) as the horizon uncertainty grows large. Finally, we also provide a way to incorporate machine-learned predictions about the horizon which interpolates between the known and unknown horizon settings.

CCS CONCEPTS

• Theory of computation → Online algorithms.

KEYWORDS

online resource allocation; traffic spikes; dual mirror descent

ACM Reference Format:

Santiago Balseiro, Christian Kroer, and Rachitesh Kumar. 2023. Online Resource Allocation under Horizon Uncertainty. 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/10.1145/3578338.3593559>

Online resource allocation is a general framework that includes as special cases various fundamental problems like network revenue

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/10.1145/3578338.3593559>




Figure 1: Suppose the horizon T (total number of requests) is unknown but it is constrained to lie in some uncertainty window $[\tau_1, \tau_2]$. Above plot shows (i) our upper bound on the best-possible competitive ratio which scales as $\tilde{O}(\ln(\tau_2/\tau_1)^{-1})$, (ii) the (asymptotic) competitive ratio of our algorithm which scales as $\Omega(\ln(\tau_2/\tau_1)^{-1})$, and (iii) an upper bound on the competitive ratio of algorithms that are optimal for the known-horizon setting when used with some proxy horizon $T^* \in [\tau_1, \tau_2]$, which scales as $(\tau_2/\tau_1)^{-1/2}$. Even for small values of τ_2/τ_1 , our algorithm significantly outperforms previous ones.

management, online advertising, online linear/convex programming, bidding in repeated auctions, and assortment optimization under inventory constraints. It captures any setting in which a decision-maker endowed with a limited amount of resources faces sequentially-arriving requests, each of which consumes a certain amount of resources and generates a reward. At each time step, the decision maker observes the request and then takes an action, with the overarching aim of maximizing cumulative reward subject to resource constraints. In this work, we focus on the setting in which requests are generated by some stationary distribution unknown to the decision maker. We impose very mild assumptions on this distribution, and in particular do not require the requests to satisfy common convexity assumptions. To the best of our knowledge, nearly all of the previous works on stochastic online resource allocation assume that the total number of requests (horizon) is known to the decision maker. Importantly, this assumption is vital for previous algorithms and performance guarantees because it allows them to compute a *per-period resource budget* (total amount of resources divided by total number of requests) and use it as the target consumption in each time step. However, if the total number of requests is not known to the decision maker, one can no longer compute this quantity and these previous works fail to offer any guidance. Juxtapose this with a world in which viral trends are becoming ever more common, causing online advertising platforms, retailers and service providers to routinely experience traffic spikes.

These spikes inject uncertainty into the system and make it difficult to accurately predict the total number of requests that will arrive. In fact, these spikes present lucrative opportunities for the advertiser/retailer, which makes addressing the uncertainty even more pertinent [4]. Moreover, it is usually difficult to predict these spikes, e.g. a news story breaks about COVID-travel bans being lifted, which results in a sudden and large uptick in the number of advertising opportunities for an airline. In fact, search-traffic spikes might be so large that they cause websites to crash¹. This motivates us to relax the previously-ubiquitous known-horizon assumption and address this omission in the literature by developing algorithms which are robust to horizon uncertainty. We use as our benchmark the hindsight optimal allocation that can be computed with full knowledge of all requests and the time horizon.

Impossibility Results. If no assumptions are made about the horizon, it has been shown in the context of matching [3] and prophet inequalities [1] that no algorithm can guarantee a positive constant fraction of the hindsight optimal reward. Contrast this with the online convex optimization literature, where one can easily obtain a low-regret algorithm (i.e., asymptotic competitive ratio of one) for the unknown horizon setting by applying the doubling trick to a low-regret algorithm for the known-horizon setting². In this paper, we prove a new impossibility result for the setting when the horizon T is constrained to lie in an uncertainty window $[\tau_1, \tau_2]$ that is known to the decision maker, which is the mildest possible assumption that renders the problem interesting. In the uncertainty-window setting, we show that no online algorithm can achieve a greater than $\tilde{O}(\ln(\tau_2/\tau_1)^{-1})$ fraction of the hindsight optimal reward. This upper bound holds even when (i) there is only 1 type of resource, (ii) the decision maker receives the same request at each time step, (iii) this request is known to the decision maker ahead of time, (iv) the request has a smooth concave reward function and linear resource consumption, (v) τ_1 is arbitrarily large, and (vi) the initial resource endowment $B = \Theta(\tau_1)$ scales with the horizon. In particular, unlike the known-horizon setting, vanishing regret is impossible to achieve under horizon uncertainty, leading us to focus on developing algorithms with a good asymptotic competitive ratio (fraction of the hindsight optimal reward).

Variable Target Dual Mirror Descent. Dual mirror descent is a natural algorithm for the known-horizon setting introduced by Balseiro et al. [2], who build on a long line of primal-dual algorithms for online allocation problems. It maintains a price (i.e., dual variable) for each resource and then dynamically updates them with the goal of consuming the per-period resource budget at each step—if the resource is being over-consumed, increase its price; and vice-versa. As stated earlier, this approach fails if the horizon is not known because the per-period budget cannot be computed ahead of time. A natural approach to handle horizon uncertainty is to use dual mirror descent with some proxy horizon $T^* \in [\tau_1, \tau_2]$ in the hopes of getting good performance for all $T \in [\tau_1, \tau_2]$. We

show that this approach can be extremely suboptimal, not just for dual mirror descent but for any algorithm which is optimal for the known-horizon setting. Thus, the unknown-horizon setting calls for new algorithms. Our main insight is that, even though one cannot compute the per-period resource budget and target its consumption, it is possible to compute a time-varying sequence of target consumptions which, if consumed at those rates, perform well no matter what the horizon turns out to be. To achieve this, we develop *Variable Target Dual Mirror Descent*, which takes a sequence of target consumptions as input and dynamically updates the prices to hit those targets. One of our primary technical contributions is generalizing the analysis of dual mirror descent to develop a fundamental bound that allows for general target consumption sequences. We leverage this bound to show that there exists a simple time-varying target consumption sequence which can be described in closed form and achieves a near-optimal $\Omega(\ln(\tau_2/\tau_1)^{-1})$ asymptotic competitive ratio when deployed with our *Variable Target Dual Mirror Descent* algorithm, matching the upper bound up to logarithmic factors.

Optimizing the Target Sequence and Incorporating Predictions. Variable Target Dual Mirror Descent reduces the complex problem of finding an algorithm which maximizes the competitive ratio to the much simpler problem of finding the optimal target consumption sequence. We develop another algorithm to solve the latter efficiently, leading to substantial gains over previous algorithms even for small values of τ_2/τ_1 (see Figure 1). Importantly, our algorithm does not require one to solve computationally-expensive linear programs (LPs), which can be desirable in time-sensitive applications. We then use the Algorithms-with-Predictions framework to study incorporating (potentially inaccurate) predictions about the horizon with the goal of performing well if the prediction comes true, while also ensuring a good competitive ratio no matter what the horizon turns out to be. We show that the problem of computing the optimal target consumption sequence for the goal of optimally incorporating predictions can also be solved efficiently by our algorithm. In particular, it allows the decision maker to account for the level of confidence she has in the predictions, and smoothly interpolate between the known-horizon and uncertainty-window settings.

Acknowledgement. Christian Kroer was supported by the Office of Naval Research award N00014-22-1-2530, and the National Science Foundation award IIS-2147361.

Full Version. The full paper can be found at <https://arxiv.org/abs/2206.13606>

REFERENCES

- [1] Reza Alijani, Siddhartha Banerjee, Sreenivas Gollapudi, Kamesh Munagala, and Kangning Wang. 2020. Predict and match: Prophet inequalities with uncertain supply. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 4, 1 (2020), 1–23.
- [2] Santiago Balseiro, Haihao Lu, and Vahab Mirrokni. 2020. The best of many worlds: Dual mirror descent for online allocation problems. *arXiv preprint arXiv:2011.10124* (2020).
- [3] Brian Brubach, Nathaniel Grammel, Will Ma, and Aravind Srinivasan. 2019. Online Matching Frameworks under Stochastic Rewards, Product Ranking, and Unknown Patience. *arXiv preprint arXiv:1907.03963* (2019).
- [4] Hossein Esfandiari, Nitish Korula, and Vahab Mirrokni. 2015. Online allocation with traffic spikes: Mixing adversarial and stochastic models. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. 169–186.

¹<https://developers.google.com/search/blog/2012/02/preparing-your-site-for-traffic-spike>

²Specifically, the doubling trick involves repeatedly running the low-regret algorithm for the known-horizon setting with increasing lengths of horizons. This is possible because the online convex optimization problem decouples across time into independent subproblems. Such decomposition is not possible in our problem because of the resource constraints: resources that have been consumed in the past restrict future actions of the algorithm.