

Combining Regularization with Look-Ahead for Competitive Online Convex Optimization

Ming Shi, *Member, IEEE*, Xiaojun Lin, *Fellow, IEEE*, and Lei Jiao, *Member, IEEE*

Abstract—There has been significant interest in leveraging limited look-ahead to achieve low competitive ratios for online convex optimization (OCO). However, existing online algorithms (such as Averaging Fixed Horizon Control (AFHC)) that can leverage look-ahead to reduce the competitive ratios still produce competitive ratios that grow unbounded as the coefficient ratio (i.e., the maximum ratio of the switching-cost coefficient and the service-cost coefficient) increases. On the other hand, the regularization method can attain a competitive ratio that remains bounded when the coefficient ratio is large, but it does not benefit from look-ahead. In this paper, we propose a new algorithm, called Regularization with Look-Ahead (RLA), that can get the best of both AFHC and the regularization method, i.e., its competitive ratio decreases with the look-ahead window size when the coefficient ratio is small, and remains bounded when the coefficient ratio is large. Moreover, we provide a matching lower bound for the competitive ratios of all online algorithms with look-ahead, which differs from the achievable competitive ratio of RLA within a factor that only depends on the problem size. Further, the competitive analysis of RLA involves a non-trivial generalization of online primal-dual analysis to the case with look-ahead.

Index Terms—Competitive analysis, look-ahead, online convex optimization, regularization, switching costs.

I. INTRODUCTION

ONLINE convex optimization (OCO) problem with switching costs has many applications in the context of networking [2]–[6], cloud or edge computing [7]–[12], cyber-physical systems [13]–[16], machine learning [17]–[21] and beyond [22]–[24]. Typically, a decision maker and the adversary (or environment) interact sequentially over time. At each time t , after receiving the current input, the decision maker must make a decision. This decision incurs a service cost (that is a function of the current decision) and a switching cost (that depends on the difference between the current decision and the previous decision). In competitive OCO, the goal is to design online algorithms with low competitive ratios. The competitive ratio is defined as, over all possible input sequences, the worst-case ratio between the total cost of an

online algorithm and that of the optimal offline algorithm, who knows the entire input sequence in advance [25, p. 3].

As a more concrete example (based on which we will conduct a case study in Sec. VIII), the importance of online decisions and switching costs can be seen in serverless computing [26], [27]. In serverless computing, customers can dynamically invoke serverless functions on demand. Thus, from the service provider’s point of view (who needs to manage, i.e., start/stop, the actual instances executing these serverless functions), the decision must be made in an online manner without knowing the future demands. Further, whenever the number of active instances is smaller than the requested number, there will be a cold-start delay, which degrades end-user experience and thus corresponds to switching costs. Therefore, the online decisions of the service provider must balance the service costs (of running the actual instances) with the switching costs. Similar examples could also be found in network functions virtualization (NFV) [5], online geographical load balancing [2] and dynamic right-sizing in data centers [3].

In the literature, many online algorithms with guaranteed competitive ratios have been provided for OCO. For example, [3], [13]–[15] provide online algorithms with constant competitive ratios for some limited settings, e.g., 1-dimensional OCO problems. However, for more general settings and under no future information, the competitive ratios of existing online algorithms [23], [28]–[30] depend on problem parameters and can usually be quite large. This is not surprising because, when there is absolutely no future information, it would be difficult to choose one online decision that is good for all possible future inputs.

To overcome this difficulty, a recent line of work has focused on how to utilize limited look-ahead information to improve the competitive ratios of online algorithms [2], [9], [31]–[33]. Here, look-ahead means that, at each time t , the decision maker knows not only the current input, but also the inputs of the immediately following K time-slots (i.e., a look-ahead window of size K). Intuitively, as K increases, the competitive ratios of online algorithms should become smaller. The Averaging Fixed Horizon Control (AFHC) algorithm, which was proposed in [2], achieves exactly that. Specifically, assume that the service cost for each decision variable $x_n(t)$ is linear, i.e., $c_n(t)x_n(t)$, and the switching cost for $x_n(t)$ is in the form of $w_n|x_n(t) - x_n(t-1)|$, where $c_n(t)$ and w_n are the service-cost and switching-cost coefficients, respectively. Then, the competitive ratio of AFHC is $1 + \max_{\{n,t\}} \frac{w_n}{c_n(t)(K+1)}$.

This work has been partially supported by the U.S. National Science Foundation (through Grants CNS-2113893, CNS-2225950, CNS-2047719 and CNS-2225949). A preliminary version of this work has appeared in IEEE International Conference on Computer Communications (IEEE INFOCOM), May 10-13, 2021 [1].

Ming Shi and Xiaojun Lin are with the Elmore Family School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907, USA (e-mail: sming@purdue.edu, linx@ecn.purdue.edu).

Lei Jiao is with the Department of Computer Science, University of Oregon, Eugene, OR 97403, USA (e-mail: jiao@cs.uoregon.edu).

In the rest of this paper, we define the “**coefficient ratio**” r_{co} to be the maximum ratio of the switching-cost and service-cost coefficients, i.e., $r_{\text{co}} \triangleq \max_{\{n,t\}} \frac{w_n}{c_n(t)}$. Thus, for any fixed coefficient ratio, the competitive ratio of AFHC decreases with the look-ahead window size K .

However, what remains unsatisfactory is that the competitive ratio of AFHC still grows with the coefficient ratio. In other words, regardless of the size K of the look-ahead window, as the coefficient ratio increases (e.g., some service-cost coefficients $c_n(t)$ may be very close to 0), the competitive ratio of AFHC will go to infinity. In a similar manner, the competitive ratio of a related algorithm in [33] could also be arbitrarily large when the coefficient ratio increases.

The above performance degradation when the coefficient ratio is large leaves much to be desired. Indeed, even with no look-ahead information, the regularization method [23] can achieve a competitive ratio that is independent of the coefficient ratio r_{co} . Of course, the downside of the regularization method of [23] is that it cannot leverage look-ahead. Therefore, it would be much more desirable if we can get the best of both worlds, i.e., achieve a competitive ratio that both decreases with K when r_{co} is small (similar to AFHC), and remains bounded when r_{co} is large (similar to the regularization method). Our previous work [34] claimed to achieve this by providing a $(1 + \frac{1}{K})$ -competitive online algorithm. Unfortunately, there appears to be an error in the proof so that the claimed competitive ratio does not hold. (Indeed, as we show in Sec. III in this paper, no algorithms can achieve a competitive ratio that low.) To the best of our knowledge, it remains an open question how to combine the strengths of both AFHC and the regularization method.

In this paper, we present new results that answer this open question. We first focus on a more restrictive setting, where the service cost is linear in the decision variables and the feasible decisions are chosen from a convex set formed by fractional covering constraints (see (1) for the specific form). While we begin with this model for simplicity and ease of exposition, it still captures the key features of practical problems [12], [22], [23], [25], [28], [29], [35], [36] (i.e., the allocated resources must meet the incoming demand).

Under this simplified model, our first contribution is to provide a lower bound on the competitive ratio for all online algorithms. Specifically, we show that, there exists instances such that the competitive ratio cannot be lower than $1 + \frac{\log_2 N}{2[1 + \frac{1}{r_{\text{co}}}((K+1)\log_2 N + 1)]}$, where N is the total number of the decision variables. *To the best of our knowledge, this is the first such lower bound in the literature for OCO problems with look-ahead.* This lower bound reveals several important insights. First, it is larger than $1 + \frac{1}{K}$ when r_{co} is large, indicating that the competitive results reported in [34] were incorrect. Second, it reveals how the coefficient ratio r_{co} affects the fundamental limit that online decisions can benefit from look-ahead. Specifically, if the size of the look-ahead window K is much larger than the coefficient ratio r_{co} , the lower bound will be driven to 1 as K increases (similar to AFHC). On the other hand, if the size of the look-ahead window is much smaller than the coefficient ratio, the lower bound will not be

close to 1. However, unlike AFHC, even when r_{co} approaches infinity, the lower bound remains at $1 + \frac{1}{2} \log_2 N$. This suggests that one may indeed design online algorithms that can get the best of both AFHC and the regularization method.

Inspired by the lower bound, our second important contribution is to provide a new online algorithm, called Regularization with Look-Ahead (RLA), whose competitive ratio matches with the lower bound up to a factor that only depends on the problem size N and is independent of the coefficient ratio r_{co} . Specifically, let $\eta \triangleq \ln(\frac{N+\epsilon}{\epsilon})$, where ϵ is a positive value chosen by RLA. We show that, when $\lceil r_{\text{co}} \rceil < K + 1$, the competitive ratio of RLA is $1 + \frac{3\eta(1+\epsilon)\lceil r_{\text{co}} \rceil}{K+1}$, which approaches 1 as the look-ahead window size K decreases. When $\lceil r_{\text{co}} \rceil \geq K + 1$, the competitive ratio of RLA is $1 + 2\eta(1 + \epsilon)$, which remains upper-bounded even when the coefficient ratio r_{co} increases to infinity. We can show that the competitive ratio of RLA differs from the lower bound by a factor $\max\left\{36\eta(1 + \epsilon), \frac{4\eta(1+\epsilon)[\frac{3}{2} + \log_2 N]}{\log_2 N}\right\}$. *To the best of our knowledge, RLA is the first such online algorithm in the literature that can get the best of both AFHC and the regularization method, i.e., achieve a competitive ratio that both decreases with K when the coefficient ratio is small, and remains upper-bounded when the coefficient ratio is large.*

Such an improved competitive ratio of RLA is achieved by carefully modifying the objective function that RLA optimizes in each episode of $K + 1$ time-slots (see Section IV). Note that within each such episode, AFHC [2] directly optimizes the total cost. However, as shown in the counter-example in [34], simply optimizing the total cost may produce poor decisions at the end of the episode, leading to poor competitive ratios. Instead, RLA replaces the switching cost in the first time-slot of each episode by two specially-chosen regularization terms at the beginning and the end of the episode. These two regularization terms avoid poor decisions at the boundary between episodes, so that the switching costs will not be excessively high. These regularization terms were inspired by that of [23], but are different because we need to leverage look-ahead. *To the best of our knowledge, this way of adding regularization terms for problems with look-ahead is also new.*

The competitive ratio of RLA is shown via an online primal-dual analysis [28]. However, there arise two new technical difficulties. First, we need to verify that the online dual variables from different episodes are feasible for the offline dual optimization problem. Second, we need to carefully bound the gap between the online primal cost and the online dual cost induced by the two regularization terms. We resolve these difficulties by providing a new competitive analysis, which extends the primal-dual analysis [28] to the case with look-ahead. *This analysis is also a key contribution of this paper and of independent interest.*

Furthermore, while the above results are stated for OCO problems with fractional covering constraints, we show in Sec. VII that these results can be extended to more general demand-supply balance constraints and capacity constraints, which are more useful for computing and networking applications.

Our work is also related to regret minimization for OCO

problems with constraints [37], [38]. In particular, [37] shows that one cannot simultaneously obtain sub-linear regret in both the objective and the constraint violation. However, our study of competitive OCO is different as the competitive ratio focuses on the *relative ratio* to the cost of the best offline *dynamic* decision, while [37], [38] focus on the *absolute difference* from the cost of the best *static* decision. Thus, even if sub-linear regret is not attainable, it is still possible to attain a low competitive ratio.

Finally, our work is related to the convex body/function chasing (CBC) problem [39]–[41]. Although their results can be applied to our problem, their competitive ratios are looser as they do not exploit the particular structure of our problem. For example, the dependency of our competitive ratio on the problem size is $O(\ln N)$, which is significantly smaller than the $O(N)$ dependency attained in [41]. Moreover, our results show how the size K of the look-ahead window and the coefficient ratio r_{co} affect the competitive ratios, which are not captured by existing work for CBC.

A preliminary version of the results was published in IEEE INFOCOM 2021 [1]. This journal version substantially enhances the conference version by (i) adding a case study on serverless computing (in Sec. VIII), (ii) providing new results on tightening the competitive ratio of our RLA algorithm when $K = 0$ (in Sec. VI), and (iii) including key proofs of our main results (in Sec. III and appendices).

II. PROBLEM FORMULATION

A. OCO with Switching Costs

The decision maker and the adversary (or environment) interact in \mathcal{T} time-slots. At each time $t = 1, \dots, \mathcal{T}$, first a feasible convex set $\mathbb{X}(t)$ and service-cost coefficients $\vec{C}(t) = [c_n(t), n = 1, \dots, N]^T \in \mathbb{R}_+^{N \times 1}$ are revealed, where $[\cdot]^T$ denotes the transpose of a vector, \mathbb{R}_+ represents the set of non-negative real numbers. For now, we restrict the set $\mathbb{X}(t)$ to be a polyhedron formed by fractional covering constraints, i.e.,

$$\sum_{n \in S_m(t)} x_n(t) \geq 1, \text{ for all } m = 1, \dots, M(t), \quad (1)$$

where $S_m(t)$ is a subset of $\{1, 2, \dots, N\}$ and could change over time. The number $M(t)$ of such constraints at each time t could also change over time. The fractional covering constraints have been widely used to model many important practical problems [22], [25], [35], [36], [42], [43]. Although the right-hand-side of (1) must be 1, which simplifies our exposition, such constraints capture the essential feature of practical constraints that the amount of resource allocated must be no smaller than the incoming demand. Further, note that there is no upper-bound constraint on the decision variable $x_n(t)$. In Sec. VII, we will extend our results to the case with more general constraints.

After receiving the input $\mathbb{X}(t)$ and $\vec{C}(t)$, the decision maker must choose a decision $\vec{X}(t) = [x_n(t), n = 1, \dots, N]^T \in \mathbb{R}_+^{N \times 1}$ from the convex set $\mathbb{X}(t)$. Then, it incurs a service cost $\langle \vec{C}(t), \vec{X}(t) \rangle$ for the current decision $\vec{X}(t)$ and a switching

cost $\langle \vec{W}, [\vec{X}(t) - \vec{X}(t-1)]^+ \rangle$ for the increment¹ of $\vec{X}(t)$ from the last decision $\vec{X}(t-1)$, where $\vec{W} = [w_n, n = 1, \dots, N]^T \in \mathbb{R}_+^{N \times 1}$ is the switching-cost coefficient. We assume that the coefficient ratio $r_{\text{co}} \triangleq \max_{\{n,t\}} \frac{w_n}{c_n(t)}$ satisfies $r_{\text{co}} \geq 1$.

In an offline setting, at time $t = 1$, the current and all the future inputs $\mathbb{X}(1 : \mathcal{T})$ and $\vec{C}(1 : \mathcal{T})$ are known. Thus, the optimal offline solution can be obtained by solving a standard convex optimization problem as follows,

$$\min_{\vec{X}(1:\mathcal{T})} \sum_{t=1}^{\mathcal{T}} \left\{ \vec{C}^T(t) \vec{X}(t) + \vec{W}^T [\vec{X}(t) - \vec{X}(t-1)]^+ \right\} \quad (2a)$$

$$\text{sub. to: } \vec{X}(t) \geq 0, \text{ for all } t \in [1, \mathcal{T}], \quad (2b)$$

$$\sum_{n \in S_m(t)} x_n(t) \geq 1, \text{ for all } m \in [1, M(t)], t \in [1, \mathcal{T}], \quad (2c)$$

where $[a, b]$ denotes the set $\{a, a+1, \dots, b\}$. As typically in many OCO problems [2], [5], [10], [23], we assume $\vec{X}(0) = 0$. For ease of exposition, we use $\vec{X}(t_1 : t_2)$ to collect $\vec{X}(t)$ from time $t = t_1$ to t_2 , i.e., $\vec{X}(t_1 : t_2) \triangleq \{\vec{X}(t), \text{ for all } t \in [t_1, t_2]\}$. Define $\vec{C}(t_1 : t_2)$ and $\mathbb{X}(t_1 : t_2)$ similarly.

B. Look-Ahead Model and Performance Metric

A recent line of work has focused on how to use look-ahead to improve competitive online algorithms [2], [9], [32], [33], [44]. Let the look-ahead window size be $K \geq 1$. Then, at each time t , the decision maker not only knows the exact input $(\mathbb{X}(t), \vec{C}(t))$, but also knows the near-term future $(\mathbb{X}(t+1 : t+K), \vec{C}(t+1 : t+K))$. Note that at time t the decision maker still does not know the future inputs beyond time $t+K$.

For an online algorithm π , let $\vec{X}^\pi(t)$ be the decision at time t . Then, its cost from time $t = t_1$ to t_2 is given as follows,

$$\text{Cost}^\pi(t_1 : t_2) \triangleq \sum_{t=t_1}^{t_2} \vec{C}^T(t) \vec{X}^\pi(t) + \sum_{t=t_1}^{t_2} \vec{W}^T [\vec{X}^\pi(t) - \vec{X}^\pi(t-1)]^+ \quad (3)$$

Let $\vec{X}^{\text{OPT}}(1:\mathcal{T})$ be the optimal offline solution to the optimization problem (2), whose total cost is $\text{Cost}^{\text{OPT}}(1 : \mathcal{T})$. Different from the *offline* setting, in an *online* setting, the decision maker only knows the current input $(\mathbb{X}(t), \vec{C}(t))$ and the look-ahead information $(\mathbb{X}(t+1 : t+K), \vec{C}(t+1 : t+K))$. Moreover, the decision $\vec{X}(t)$ made at each time is irrevocable. Then, the competitive ratio of the *online* algorithm π is defined as,

$$\text{CR}^\pi \triangleq \max_{\{\text{all possible } (\mathbb{X}(1:\mathcal{T}), \vec{C}(1:\mathcal{T}))\}} \frac{\text{Cost}^\pi(1 : \mathcal{T})}{\text{Cost}^{\text{OPT}}(1 : \mathcal{T})}, \quad (4)$$

i.e., the worst-case ratio of its total cost to that of the optimal offline solution, over all possible inputs.

¹Note that, as shown in [29], our results assuming this type of the switching cost also imply a competitive ratio for the case when the switching cost penalizes the absolute difference $|\vec{X}(t) - \vec{X}(t-1)|$ [5], [14].

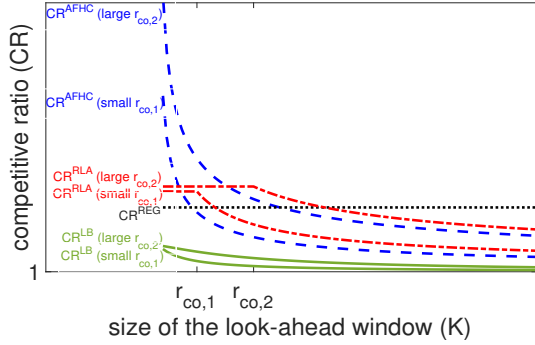


Fig. 1: Compare the lower bound of the competitive ratio (CR^{LB}) and the competitive ratios of AFHC (CR^{AFHC}), the regularization method (CR^{REG}) and RLA (CR^{RLA}).

III. A LOWER BOUND

Although OCO with look-ahead has been extensively studied, e.g., in [2], [9], [44], most existing results in the literature focus on achievable competitive ratios, but do not provide lower bounds on the competitive ratio. Such lower bounds are important because they can reveal the fundamental limit that one can hope to reach with online decisions. Note that the lower bounds in [28] and [44] are for different settings (ℓ_2 -norm switching costs and online packing problems). Further, they do not consider look-ahead. Next, we provide a new lower bound for our OCO formulation, which reveals how the relationship between the coefficient ratio r_{co} and the size K of the look-ahead window will affect the competitive ratio. (Note that Theorem 1 holds not only for $K \geq 1$, but also for $K = 0$, i.e., the case without look-ahead, which we will discuss in Sec. VI.)

Theorem 1. *Consider the OCO problem in Sec. II-A. With a look-ahead window of size $K \geq 0$, the competitive ratio of any online algorithm is lower-bounded by*

$$\text{CR}^{\text{LB}} = 1 + \frac{\log_2 N}{2 \left[1 + \frac{1}{r_{\text{co}}} ((K+1) \log_2 N + 1) \right]}. \quad (5)$$

Theorem 1 reveals important insights on how the competitive ratio is impacted by the look-ahead window size K relative to the coefficient ratio r_{co} .

(i) The lower bound CR^{LB} in (5) is always increasing in r_{co} and decreasing in K . Further, we have

$$\text{CR}^{\text{LB}} \leq 1 + \frac{r_{\text{co}}}{2(K+1)}. \quad (6)$$

Note that the right-hand-side is close to the competitive ratio of AFHC [2].

(ii) When the look-ahead window size K is large, in particular when $K+1 > r_{\text{co}}$, CR^{LB} will not be far away from (6) and the competitive ratio of AFHC. Indeed, we have

$$\text{CR}^{\text{LB}} > 1 + \frac{\log_2 N}{6 \frac{1}{r_{\text{co}}} (K+1) \log_2 N} = 1 + \frac{r_{\text{co}}}{6(K+1)}, \quad (7)$$

where the first inequality is because $(K+1) \log_2 N \geq 1$ and $\frac{1}{r_{\text{co}}} (K+1) \log_2 N \geq 1$. This behavior is illustrated by the two

solid curves in Fig. 1 (for two coefficient ratios $r_{\text{co},1} < r_{\text{co},2}$), which decrease to 1 as K increases beyond $r_{\text{co},1}$ and $r_{\text{co},2}$. Notice that this is also the range where AFHC [2] will produce a low competitive ratio (see the dashed curves in Fig. 1). In contrast, the competitive ratio of the regularization method (REG) of [23] does not decrease with K (see the dotted line in Fig. 1).

(iii) When the look-ahead window size K is small, e.g., when $K+1 \leq r_{\text{co}}$, (5) could be quite far away from (6) and the competitive ratio of AFHC. Specifically, for small K , the competitive ratio of AFHC increases to infinity when the coefficient ratio increases, which can be seen in Fig. 1 by comparing the two dashed curves at small K . In contrast, the lower bound CR^{LB} and the competitive ratio of the regularization method CR^{REG} are upper-bounded by a function of the problem size N . Indeed, even when r_{co} increases to infinity, the lower bound in (5) still satisfies that

$$\text{CR}^{\text{LB}} \leq 1 + \frac{1}{2} \log_2 N, \quad (8)$$

which suggests room for improvement for AFHC. Please see below for the proof of Theorem 1.

Proof of Theorem 1: Lower bound instance: We first present the problem instance leading to the lower bound in (5). Let $c_n(t) = c > 0$ and $w_n = w > 0$ for all n and t . Moreover, let the total number of decision variables be $N = 2^\alpha$, where α is a positive integer. Consider a total of $\mathcal{T} = (K+1)\alpha + 1$ time-slots, which is divided into $\alpha + 1$ episodes. Specifically, each of the first α episodes contains $K+1$ consecutive time-slots, while the last episode contains the last time-slot.

Our key idea of the proof is to let the adversary reveal new inputs based on the decisions of the online algorithm, so that the online algorithm has to switch at least once in each episode. Specifically, there is only $M(t) = 1$ constraint with $S_1(t)$ for every episode. In the first episode, i.e., for any $t \in [1, K+1]$, the constraint is $\sum_{n=1}^N x_n(t) \geq 1$, i.e., $S_1(t) = [1, N]$.

The constraint in the second episode is based on the decision $\bar{X}^\pi(1)$. (Note that the online algorithm must choose $\bar{X}^\pi(1)$ without knowing the constraint in the second episode.) (i) If $\sum_{n=1}^{N/2} x_n^\pi(1) \leq \sum_{n=N/2+1}^N x_n^\pi(1)$, the adversary chooses $S_1(t) = [1, \frac{N}{2}]$, and the constraint becomes $\sum_{n=1}^{N/2} x_n(t) \geq 1$ in the second episode, i.e., for all $t \in [K+2, 2K+2]$. (ii) Otherwise, the adversary chooses $S_1(t) = [\frac{N}{2} + 1, N]$, and the constraint becomes $\sum_{n=N/2+1}^N x_n(t) \geq 1$ in the second episode. Intuitively, assuming that the online algorithm π does not over-provision in the first time-slot (i.e., $\sum_{n=1}^N x_n^\pi(1) = 1$), we must have $\sum_{n \in S_1(t)} x_n^\pi(1) \leq 1/2$ for $S_1(t)$ chosen in the second episode. This choice of $S_1(t)$ then forces the online algorithm π to increase $x_n^\pi(t)$ for $n \in S_1(t)$ during the second episode (in order to meet the new constraint), and therefore the online algorithm has to incur a large switching cost as we will show below.

In a similar way, the constraint in the i -th episode ($2 \leq i \leq \alpha + 1$) will always be on the half of the previous constraint set, for which the decision variables at the beginning of the $(i-1)$ -th episode add up to a smaller sum. Following these

steps, at the last time $t = (K+1)\alpha + 1$, the constraint set will reduce to a singleton $S_1(t) = \{\tilde{n}\}$ for some $\tilde{n} \in \{1, \dots, N\}$.

Total cost of the optimal offline solution: The offline solution can simply choose, for all time-slots, $x_n^{\text{OFF}}(1 : \mathcal{T}) = 1$ for $n = \tilde{n}$, and $x_n^{\text{OFF}}(1 : \mathcal{T}) = 0$ for $n \neq \tilde{n}$. It only incurs a switching cost of w at time $t = 1$. Thus, the optimal offline cost is upper-bounded as follows,

$$\text{Cost}^{\text{OPT}}(1 : \mathcal{T}) \leq w + c((K+1)\alpha + 1). \quad (9)$$

Total cost of any online algorithm π : First, at each time $t \in [1, \mathcal{T}]$, to satisfy the constraint, at least a service cost of c is incurred. Next, we show that the total switching cost of any online algorithm π is at least $\frac{1}{2}w\alpha + w$. To see this, consider any decision variable x_n that last saw a constraint in episode $i_n \leq \alpha$, whose first time-slot is $t'(i_n) \triangleq (K+1)(i_n - 1) + 1$. It must be because the decision variable x_n is one of those that are in the constraint in episode i_n , but are excluded from the constraint in episode $i_n + 1$. Let $S'(i_n)$ be the set of all such decision variables in episode i_n . Because (i) in episode i_n the constraint must be met, and (ii) the adversary chooses the half of the decision variables whose sum are smaller to form the constraint in episode $i_n + 1$, we must have $\sum_{n \in S'(i_n)} x_n^\pi(t'(i_n)) \geq \frac{1}{2}$. Across α episodes, there are α such sets $S'(i_n)$, which are non-overlapping. Finally, in the last time-slot, the decision $x_{\tilde{n}}^\pi(\mathcal{T}) \geq 1$. Together, we have $\sum_{n=1}^N x_n^\pi(t'(i_n)) \geq \frac{\alpha}{2} + 1$. Finally, note that the total switching cost associated with $x_n(\cdot)$ is at least $w_n x_n(t'(i_n))$. Therefore, the total cost of any online algorithm π is lower-bounded as follows,

$$\text{Cost}^\pi(1 : \mathcal{T}) \geq c((K+1)\alpha + 1) + w + \frac{\alpha w}{2}. \quad (10)$$

The result then follows by dividing the right-hand-side of (10) by the right-hand-side of (9). \square

IV. REGULARIZATION WITH LOOK-AHEAD (RLA)

Inspired by Fig. 1, a nature question is then: can we develop an online algorithm that gets the best of both AFHC and the regularization method? In this section, we present a new online algorithm, called Regularization with Look-Ahead (RLA), which achieves exactly that, i.e., a competitive ratio that not only remains upper-bounded when r_{co} is large, but also decreases with K when r_{co} is small.

Specifically, let τ be an integer from 0 to K . RLA runs $K+1$ versions of a subroutine, called Regularization-Fixed Horizon Control (R-FHC), indexed by τ . We denote the τ -th version of R-FHC by $\text{R-FHC}^{(\tau)}$. $\text{R-FHC}^{(\tau)}$ divides the time horizon into episodes. Each episode starts from time $t^{(\tau)}$ to $t^{(\tau)} + K$, where $t^{(\tau)} = \tau + (K+1)u$ and $u = -1, 0, \dots, \left\lceil \frac{\mathcal{T}}{K+1} \right\rceil$. Recall that at time $t^{(\tau)}$, the inputs $(\mathbb{X}(t^{(\tau)} : t^{(\tau)} + K), \tilde{C}(t^{(\tau)} : t^{(\tau)} + K))$ at the current time and in the look-ahead window have been revealed. $\text{R-FHC}^{(\tau)}$ then computes the solution to the following problem,

$$\min_{\bar{x}(t^{(\tau)} : t^{(\tau)} + K)} \left\{ \sum_{s=t^{(\tau)}}^{t^{(\tau)}+K} \sum_{n=1}^N c_n(s) x_n(s) \right\} \quad (11a)$$

$$+ \sum_{n=1}^N \frac{w_n}{\eta} x_n(t^{(\tau)}) \ln \left(\frac{1 + \frac{\epsilon}{N}}{x_n^{\text{R-FHC}^{(\tau)}}(t^{(\tau)} - 1) + \frac{\epsilon}{N}} \right) \quad (11b)$$

$$+ \sum_{s=t^{(\tau)}+1}^{t^{(\tau)}+K} \sum_{n=1}^N w_n [x_n(s) - x_n(s-1)]^+ \quad (11c)$$

$$+ \sum_{n=1}^N \frac{w_n}{\eta} \left[\left(x_n(t^{(\tau)} + K) + \frac{\epsilon}{N} \right) \cdot \ln \left(\frac{x_n(t^{(\tau)} + K) + \frac{\epsilon}{N}}{1 + \frac{\epsilon}{N}} \right) - x_n(t^{(\tau)} + K) \right] \quad (11d)$$

$$\text{sub. to: } \sum_{n \in S_m(s)} x_n(s) \geq 1, \text{ for all } m \in [1, M(s)],$$

$$s \in [t^{(\tau)}, t^{(\tau)} + K], \quad (11e)$$

$$x_n(s) \geq 0, \text{ for all } n \in [1, N], s \in [t^{(\tau)}, t^{(\tau)} + K], \quad (11f)$$

where $\eta = \ln \left(\frac{N+\epsilon}{\epsilon} \right)$, $\epsilon > 0$ and the decision $x_n^{\text{R-FHC}^{(\tau)}}(t^{(\tau)} - 1)$ were given by the solution of the previous episode of $\text{R-FHC}^{(\tau)}$ from time $t^{(\tau)} - K - 1$ to $t^{(\tau)} - 1$.

According to (11), in each episode from time $t^{(\tau)}$ to $t^{(\tau)} + K$, RLA does not simply optimize the corresponding service costs and switching costs, but instead adds additional regularization terms similar to [23]. Next, we explain the intuition behind these regularization terms. Note that since there is no look-ahead in [23], it suffices to use a single regularization term that is based directly on the current decision $x_n(t)$ and the last decision $x_n(t-1)$, i.e.,

$$\frac{w_n}{\eta} \left[x_n(t) + \frac{\epsilon}{N} \right] \ln \left(\frac{x_n(t) + \frac{\epsilon}{N}}{x_n(t-1) + \frac{\epsilon}{N}} \right) - x_n(t) + x_n(t-1). \quad (12)$$

Using this regularization term for every time t produces the online dual variables equal to $\theta_n(t+1) = \frac{w_n}{\eta} \ln \left(\frac{1 + \frac{\epsilon}{N}}{x_n(t) + \frac{\epsilon}{N}} \right)$, which is crucial for the online primal-dual proof in [23]. Unfortunately, for RLA we cannot use a term like (12) at *every time*. The reason is that, when there is look ahead (i.e., $K > 0$), the decisions *inside* the look-ahead window of size K should be as close to the offline optimal solution as possible. Using the term (12) at those time-slots distorts the objective function too much that it will lead to sub-optimal decisions. Due to this reason, we cannot use (12) directly in our algorithm for $K > 0$. Instead, in our RLA algorithm, we can think of each look-ahead window as a single super-time-slot. Naturally, the regularization term should only involve decisions around time t and time $t+K$, but its impact to other time-slots inside the look-ahead window should be as little as possible. This motivates us to construct our regularization term as the sum of two terms: the part related to the decision $x_n(t+K)$ at the end of the look-ahead window (which is similar to the role of $x_n(t)$ in (12)) and the part related to the previous decision $x_n(t-1)$. Moreover, in order to still guarantee dual feasibility, we need the dual variables $\theta_n(t)$ at these two boundaries of each look-ahead window to be of a similar form as in [23]. Thus, we replace the switching cost in the first time-slot $t^{(\tau)}$ of the current episode by the regularization term (11b), and adds

Algorithm 1 Regularization with Look-Ahead (RLA)**Parameters:** $\epsilon > 0$ and $\eta = \ln\left(\frac{N+\epsilon}{\epsilon}\right)$.**for** $t = -K + 1 : \mathcal{T}$ **do** Step 1: $\tau \leftarrow t \bmod (K + 1)$ and $t^{(\tau)} \leftarrow t$. Step 2: Solve (11) to get $\vec{X}^{\text{R-FHC}^{(\tau)}}(t^{(\tau)} : t^{(\tau)} + K)$. (If $t^{(\tau)} \leq 0$, remove (11b). If $t^{(\tau)} \geq \mathcal{T} - K$, remove (11d).)

Step 3:

if $1 \leq t \leq \mathcal{T}$ **then**

$$\vec{X}^{\text{RLA}}(t) = \frac{1}{K+1} \sum_{\tau=0}^K \vec{X}^{\text{R-FHC}^{(\tau)}}(t). \quad (13)$$

end if**end for**

another regularization term (11d) for the decision variables in the last time-slot $t^{(\tau)} + K$ of the current episode.

Similar to [23], the regularization term (11d) makes the objective function strictly convex in $x_n(t^{(\tau)} + K)$, and thus discourages it from taking extreme values. More specifically, without (11d), it is possible that the decision in the last time-slot goes down to zero if the associated service-cost coefficient is high or if there is no constraint. However, if the next input at time $t^{(\tau)} + K + 1$ requires the next decision to be high, the algorithm will incur a high switching cost. In contrast, (11d) is decreasing and strictly convex in $x_n(t^{(\tau)} + K)$, so it discourages the decision in the last time-slot $t^{(\tau)} + K$ to be too low. When combined with the regularization term (11b), they together ensure that the switching cost at the boundary between two episodes is not too high (see details in our analysis in Sec. V). Thus, unlike AFHC, the competitive ratio of RLA can be upper-bounded even if r_{co} is large. Readers familiar with [23] will recognize that, when the size of the look-ahead window $K = 0$, these two regularization terms combined reduce to the original regularization term in [23]. However, our formulation of the regularization terms for $K \geq 1$ is new and has not been reported in the literature.

Finally, at each time $t \in [1, \mathcal{T}]$, RLA takes the average of $\vec{X}^{\text{R-FHC}^{(\tau)}}(t)$ for all τ as the final decision $\vec{X}^{\text{RLA}}(t)$ at time t . As K increases, since R-FHC^(τ) optimizes the real service costs and switching costs in the middle of each episode, more and more decision variables are close to optimal. Thus, by taking the average of all versions of R-FHC^(τ), the performance of RLA should improve with K . The details of RLA are given in Algorithm 1. Note that for any version of R-FHC^(τ) whose first episode starts at time $t^{(\tau)} \leq 0$, (11b) can be removed. Similarly, for any version of R-FHC^(τ) whose last episode ends at time $t^{(\tau)} + K \geq \mathcal{T}$, (11d) can be removed.

V. COMPETITIVE ANALYSIS

Theorem 2 below provides the theoretical competitive ratio of RLA. Recall that $\eta = \ln\left(\frac{N+\epsilon}{\epsilon}\right)$ and $r_{\text{co}} \geq 1$.

Theorem 2. Consider the OCO problem introduced in Sec. II-A. With a look-ahead window of size $K \geq 1$, the

competitive ratio of RLA is upper-bounded as follows,

$$\text{CR}^{\text{RLA}} \leq 1 + \frac{3\eta(1+\epsilon) \lceil r_{\text{co}} \rceil}{K+1}, \text{ if } \lceil r_{\text{co}} \rceil < K+1; \quad (14a)$$

$$\text{CR}^{\text{RLA}} \leq 1 + 2\eta(1+\epsilon), \text{ if } \lceil r_{\text{co}} \rceil \geq K+1. \quad (14b)$$

It is easy to see that the competitive ratio of RLA in (14) matches the lower bound (5) within a factor that only depends on the problem size N (see the two dash-dot curves in Fig. 1). Specifically, (i) when $r_{\text{co}} \leq \lceil r_{\text{co}} \rceil \leq K+1$, both (14a) and (14b) differ from (7) (and thus (5)) by at most $36\eta(1+\epsilon)$. Note that CR^{RLA} decreases to 1 as K increases. (ii) When $r_{\text{co}} \geq K+1$, we can show that the lower bound (5) is larger than $1 + \frac{\log_2 N}{2 \lceil \frac{3}{2} + \log_2 N \rceil}$. Thus, the gap between (14b) and (5) is at most $\frac{4\eta(1+\epsilon) \lceil \frac{3}{2} + \log_2 N \rceil}{\log_2 N}$. Further, when $r_{\text{co}} \geq (K+1) \log_2 N$, the gap between (14b) and (5) is at most $\frac{10\eta(1+\epsilon)}{\log_2 N}$, which is upper-bounded by a constant $10(1+\epsilon) \ln\left(\frac{2+\epsilon}{\epsilon}\right)$ for all $N \geq 2$. Note that in all cases (even when r_{co} increases to infinity), CR^{RLA} is upper-bounded. Therefore, RLA gets the best of both AFHC and the regularization method. To the best of our knowledge, RLA is the first algorithm in the literature that can utilize look-ahead to attain a competitive ratio that matches the lower bound (5).

The rest of this section is devoted to the proof of Theorem 2. We first give the high-level idea, starting from a typical online primal-dual analysis [28]. For the offline problem (2), by introducing an auxiliary variable $y_n(t)$ for the switching term $[x_n(t) - x_n(t-1)]^+$, together with a new constraint

$$y_n(t) \geq x_n(t) - x_n(t-1), \text{ for all } n \in [1, N], \quad (15)$$

we can get an equivalent formulation of the offline optimization problem (2). Then, let $\vec{\beta}(t) = [\beta_m(t), m = 1, \dots, M(t)]^T$ and $\vec{\theta}(t) = [\theta_n(t), n = 1, \dots, N]^T$ be the Lagrange multipliers for constraints (2c) and (15), respectively. We have the offline dual optimization problem as follows,

$$\max_{\{\vec{\beta}(1:\mathcal{T}), \vec{\theta}(1:\mathcal{T})\}} \sum_{t=1}^{\mathcal{T}} \sum_{m=1}^{M(t)} \beta_m(t) \quad (16a)$$

$$\text{sub. to: } c_n(t) - \sum_{m:n \in S_m(t)} \beta_m(t) + \theta_n(t) - \theta_n(t+1) \geq 0, \quad (16b)$$

$$\text{for all } n \in [1, N], t \in [1, \mathcal{T}],$$

$$w_n - \theta_n(t) \geq 0, \text{ for all } n \in [1, N], t \in [1, \mathcal{T}], \quad (16c)$$

$$\beta_m(t) \geq 0, \text{ for all } m \in [1, M(t)], t \in [1, \mathcal{T}], \quad (16d)$$

$$\theta_n(t) \geq 0, \text{ for all } n \in [1, N], t \in [1, \mathcal{T}]. \quad (16e)$$

Let $\beta_m^{\text{OPT}}(t)$ and $\theta_n^{\text{OPT}}(t)$ be the optimal solution to (16). Then, the optimal offline dual cost is,

$$D^{\text{OPT}}(1:\mathcal{T}) \triangleq \sum_{t=1}^{\mathcal{T}} \sum_{m=1}^{M(t)} \beta_m^{\text{OPT}}(t). \quad (17)$$

Let $D^{\text{RLA}}(1:\mathcal{T})$ be the total dual cost of RLA. Then, we can prove the competitive performance of RLA by establishing the following inequalities,

$$\text{Cost}^{\text{RLA}}(1:\mathcal{T}) \stackrel{(a)}{\leq} \text{CR} \cdot D^{\text{RLA}}(1:\mathcal{T})$$

$$\stackrel{(b)}{\leq} \text{CR} \cdot D^{\text{OPT}}(1 : \mathcal{T}) \stackrel{(c)}{\leq} \text{CR} \cdot \text{Cost}^{\text{OPT}}(1 : \mathcal{T}). \quad (18)$$

In (18), step (c) simply follows from standard duality [45, p. 225]. Step (b) is established by showing that RLA produces a set of online dual variables that are also feasible for the offline dual optimization problem (16). Since (16) is a maximization problem, step (b) then holds. Finally, step (a) is related to the regularization terms (11b) and (11d) added to the objective function of R-FHC, which leads to a gap between $\text{Cost}^{\text{RLA}}(1 : \mathcal{T})$ and $D^{\text{RLA}}(1 : \mathcal{T})$. This gap needs to be carefully bounded to establish (a). Below, we will address (b) and (a).

Step-1 (Checking the dual feasibility): We now focus on one version τ of R-FHC. For simplicity, in the rest of this section, we use (τ) instead of R-FHC $^{(\tau)}$ in the superscript, e.g., use $\vec{X}^{(\tau)}(t)$ to denote $\vec{X}^{\text{R-FHC}^{(\tau)}}(t)$. We now show that the decisions produced by all episodes of R-FHC $^{(\tau)}$ generate a feasible set of dual variables for (16). Focus on one episode from time $t^{(\tau)}$ to $t^{(\tau)} + K$. As in (16), we introduce the variable $y_n(t)$ and the constraint (15) to (11). We can then form the dual problem of the equivalent form of (11). As in (16), we let $\beta_m^{(\tau)}(t)$ and $\theta_n^{(\tau)}(t)$ be the corresponding online dual solution of (11). However, note that the objective function of (11) does not contain the switching cost of the first time-slot $t^{(\tau)}$. Therefore, we are still missing the dual variables $\theta_n^{(\tau)}(t^{(\tau)})$. To remediate this, for all $n \in [1, N]$, we let

$$\theta_n^{(\tau)}(t^{(\tau)}) \triangleq \frac{w_n}{\eta} \ln \left(\frac{1 + \frac{\epsilon}{N}}{x_n^{(\tau)}(t^{(\tau)} - 1) + \frac{\epsilon}{N}} \right). \quad (19)$$

Lemma 1 below shows that we have constructed a feasible dual solution for the offline dual optimization problem (16).

Lemma 1. *The $\vec{\beta}^{(\tau)}(1 : \mathcal{T})$ and $\vec{\theta}^{(\tau)}(1 : \mathcal{T})$ constructed above from (19) and the online dual solution of R-FHC $^{(\tau)}$ are feasible for the offline dual optimization problem (16).*

Lemma 1 can be proved by verifying that the Karush-Kuhn-Tucker (KKT) conditions [45, p. 243] of (11) satisfies the dual constraints (16b)-(16e). (16c) to (16e) are easy to verify, so is (16b) for $t = t^{(\tau)} + 1$ to $t^{(\tau)} + K - 1$, because the KKT conditions for (11) in those time-slots are exactly the same as that of (16). Thus, it only remains to verify (16b) at time $t = t^{(\tau)}$ and $t = t^{(\tau)} + K$. At time $t^{(\tau)}$, by examining the KKT conditions for (11), we have,

$$c_n(t^{(\tau)}) - \sum_{m:n \in S_m(t^{(\tau)})} \beta_m^{(\tau)}(t^{(\tau)}) + \frac{w_n}{\eta} \ln \left(\frac{1 + \frac{\epsilon}{N}}{x_n^{(\tau)}(t^{(\tau)} - 1) + \frac{\epsilon}{N}} \right) - \theta_n^{(\tau)}(t^{(\tau)} + 1) \geq 0.$$

Using (19), (16b) at time $t = t^{(\tau)}$ is verified. We can verify (16b) at time $t^{(\tau)} + K$ similarly. Lemma 1 then follows. Please see Appendix A for the proof of Lemma 1.

Step-2 (Quantifying the gap between the online primal cost and the online dual cost): As before, we focus on one episode (from time $t^{(\tau)}$ to $t^{(\tau)} + K$) of version τ of R-FHC.

We define the primal cost $\text{Cost}^{(\tau)}(t^{(\tau)} : t^{(\tau)} + K)$ as in (3) and the online dual cost

$$D^{(\tau)}(t^{(\tau)} : t^{(\tau)} + K) \triangleq \sum_{t=t^{(\tau)}}^{t^{(\tau)}+K} \sum_{m=1}^M \beta_m^{(\tau)}(t). \quad (20)$$

However, note that (11) contains additional terms (11b) and (11d) in the primal objective function. Thus, there will be some gap between $\text{Cost}^{(\tau)}(t^{(\tau)} : t^{(\tau)} + K)$ and $D^{(\tau)}(t^{(\tau)} : t^{(\tau)} + K)$. Lemma 2 below captures this gap. Define the tail-terms as

$$\Omega_n^{(\tau)}(t^{(\tau)}) \triangleq w_n \left[x_n^{(\tau)}(t^{(\tau)}) - x_n^{(\tau)}(t^{(\tau)} - 1) \right]^+, \quad (21)$$

$$\phi_n^{(\tau)}(t^{(\tau)}) \triangleq -\frac{w_n}{\eta} x_n^{(\tau)}(t^{(\tau)}) \ln \left(\frac{1 + \frac{\epsilon}{N}}{x_n^{(\tau)}(t^{(\tau)} - 1) + \frac{\epsilon}{N}} \right), \quad (22)$$

$$\psi_n^{(\tau)}(t^{(\tau)}) \triangleq \frac{w_n}{\eta} x_n^{(\tau)}(t^{(\tau)} + K) \ln \left(\frac{1 + \frac{\epsilon}{N}}{x_n^{(\tau)}(t^{(\tau)} + K) + \frac{\epsilon}{N}} \right). \quad (23)$$

Lemma 2. *For each version τ of R-FHC, we have,*

$$\text{Cost}^{(\tau)}(t^{(\tau)} : t^{(\tau)} + K) \leq D^{(\tau)}(t^{(\tau)} : t^{(\tau)} + K) + \sum_{n=1}^N \Omega_n^{(\tau)}(t^{(\tau)}) + \sum_{n=1}^N \phi_n^{(\tau)}(t^{(\tau)}) + \sum_{n=1}^N \psi_n^{(\tau)}(t^{(\tau)}). \quad (24)$$

Lemma 2 captures the gap between the online primal cost and the online dual cost of each version τ of R-FHC. In (24), the first tail-term $\Omega_n^{(\tau)}(t^{(\tau)})$ is because R-FHC $^{(\tau)}$ does not optimize over the real switching cost $w_n [x_n(t^{(\tau)}) - x_n(t^{(\tau)} - 1)]^+$ in the first time-slot. The second and third tail-terms, $\phi_n^{(\tau)}(t^{(\tau)})$ and $\psi_n^{(\tau)}(t^{(\tau)})$, are because of the regularization terms (11b) and (11d) added to the primal objective function in the first time-slot and the last time-slot. Lemma 2 can be shown via the duality theorem [45, p.225]. Please see Appendix B for the proof of Lemma 2.

Recall that, to establish step (a) in (18), the main difficulty is to bound the gap due to the tail-terms in Lemma 2. We resolve this difficulty by designing two important steps as follows. The ideas we propose in these two steps are novel and critical for online primal-dual analysis. This non-trivial generalization of online primal-dual analysis to the case with look-ahead is of independent interest.

Step 2-1 (Bounding the tail-terms): Next, we show in Lemma 3 that, with a factor that will appear in the final competitive ratio, the tail-terms (21)-(23) from the same version τ of R-FHC are actually bounded by a carefully-chosen portion of the online dual costs. We let $\Delta = \min\{K, \lceil r_{\text{co}} \rceil - 1\}$.

Lemma 3. *For each version τ of R-FHC, the following holds,*

$$(i) \sum_{u=0}^{\lceil \frac{\mathcal{T}}{K+1} \rceil} \sum_{t^{(\tau)}=\tau+(K+1)u} \sum_{n=1}^N \Omega_n^{(\tau)}(t^{(\tau)}) \leq \eta(1 + \epsilon) \times \sum_{u=0}^{\lceil \frac{\mathcal{T}}{K+1} \rceil} \sum_{t^{(\tau)}=\tau+(K+1)u} D^{(\tau)}(t^{(\tau)} : t^{(\tau)} + \Delta), \quad (25)$$

$$\begin{aligned}
 (ii) \quad & \sum_{u=-1}^{\lceil \frac{\tau}{K+1} \rceil} \sum_{t^{(\tau)}=\tau+(K+1)u}^{\tau} \sum_{n=1}^N \left[\phi_n^{(\tau)}(t^{(\tau)}) + \psi_n^{(\tau)}(t^{(\tau)}) \right] \\
 & \leq \eta(1+\epsilon) \sum_{u=-1}^{\lceil \frac{\tau}{K+1} \rceil} \sum_{\substack{t^{(\tau)}=\tau \\ +(K+1)u}} D^{(\tau)}(t^{(\tau)}+K-\Delta : t^{(\tau)}+K), \quad (26)
 \end{aligned}$$

where $D^{(\tau)}(t) = 0$ for all $t \leq 0$ and $t > \mathcal{T}$.

To interpret (25), the tail-term $\Omega_n^{(\tau)}(t^{(\tau)})$ are bounded by the right-hand-side of (25), which corresponds to a partial sum of online dual costs over sub-intervals of length $\Delta + 1$ at the beginning of each episode. (Note that when $\lceil r_{\text{co}} \rceil$ is large, $\Delta = K$ and thus this sub-interval will contain the whole episode.) Expression (26) has a similar interpretation, while the partial sum is over sub-intervals at the end of each episode. Due to page limits, we only provide a sketch of proof of Lemma 3 below. Please see our technical report [46] for the complete proof.

Sketch of Proof of Lemma 3: We focus on the proof of (25), and (26) follows along a similar line. Consider any $t^{(\tau)}$ and n such that $\Omega_n^{(\tau)}(t^{(\tau)}) > 0$, i.e., $x_n^{(\tau)}(t^{(\tau)}) > x_n^{(\tau)}(t^{(\tau)} - 1)$. First, since $a - b \leq a \ln\left(\frac{a}{b}\right)$ for all $a, b > 0$ and $x_n(t) \leq 1$, we can show that each $\Omega_n^{(\tau)}(t^{(\tau)})/\eta$ is upper-bounded by

$$\frac{w_n}{\eta} \left[x_n^{(\tau)}(t^{(\tau)}) + \frac{\epsilon}{N} \right] \ln \left(\frac{1 + \frac{\epsilon}{N}}{x_n^{(\tau)}(t^{(\tau)} - 1) + \frac{\epsilon}{N}} \right). \quad (27)$$

Let $\hat{\beta}_n^{(\tau)}(t) = \sum_{m:n \in \mathcal{S}_m(t)} \beta_m^{(\tau)}(t)$. Consider any $t' > t^{(\tau)}$ such that $x_n^{(\tau)}(t) > 0$ for all $t \in [t^{(\tau)}, t']$. Using KKT conditions of (11), we can show that (27) is equal to

$$\begin{aligned}
 & \sum_{t=t^{(\tau)}}^{t'} \left[x_n^{(\tau)}(t) + \frac{\epsilon}{N} \right] \hat{\beta}_n^{(\tau)}(t) + \left[x_n^{(\tau)}(t') + \frac{\epsilon}{N} \right] \theta_n^{(\tau)}(t' + 1) \\
 & - \sum_{t=t^{(\tau)}}^{t'} c_n(t) \left[x_n^{(\tau)}(t) + \frac{\epsilon}{N} \right] - \sum_{t=t^{(\tau)}+1}^{t'} w_n y_n^{(\tau)}(t). \quad (28)
 \end{aligned}$$

Next, we show that

$$\frac{\Omega_n^{(\tau)}(t^{(\tau)})}{\eta} \leq (28) \leq \sum_{t=t^{(\tau)}}^{t^{(\tau)}+\Delta} \left[x_n^{(\tau)}(t) + \frac{\epsilon}{N} \right] \hat{\beta}_n^{(\tau)}(t) \quad (29)$$

by considering the following two cases. (i) If there exists a time-slot $t < t^{(\tau)} + \Delta$, such that $x_n^{(\tau)}(t + 1) < x_n^{(\tau)}(t)$, we take t' as the first such t after $t^{(\tau)}$. Then, we must have $\theta_n^{(\tau)}(t' + 1) = 0$ (from complementary slackness) and (29) follows. (ii) If no such time-slot t exists, we let $t' = t^{(\tau)} + \Delta$. There are two sub-cases. (ii-a) If $\lceil r_{\text{co}} \rceil - 1 < K$, then we consider the last three terms in (28). Since $x_n^{(\tau)}(t') - \sum_{t=t^{(\tau)}+1}^{t'} y_n^{(\tau)}(t) = x_n^{(\tau)}(t^{(\tau)})$ (because $x_n^{(\tau)}(t)$ does not decrease before time t') and $\theta_n^{(\tau)}(t' + 1) \leq w_n$, the second and fourth term in (28) can be upper-bounded by $w_n \left[x_n^{(\tau)}(t^{(\tau)}) + \frac{\epsilon}{N} \right]$. Then, since $x_n^{(\tau)}(t) \geq x_n^{(\tau)}(t^{(\tau)})$ for all $t \in [t^{(\tau)}, t']$ and $\sum_{t=t^{(\tau)}}^{t'} c_n(t) \geq \frac{w_n}{r_{\text{co}}} (\Delta + 1) \geq w_n$, the last

three terms in (28) are upper-bounded by 0, and (29) then follows. (ii-b) If $\lceil r_{\text{co}} \rceil - 1 \geq K$, we can show that

$$\begin{aligned}
 \frac{\Omega_n^{(\tau)}(t^{(\tau)})}{\eta} & \leq \frac{w_n}{\eta} \left[x_n^{(\tau)}(t^{(\tau)}) + \frac{\epsilon}{N} \right] \ln \left(\frac{1 + \frac{\epsilon}{N}}{x_n^{(\tau)}(t^{(\tau)} - 1) + \frac{\epsilon}{N}} \right) \\
 & - \frac{w_n}{\eta} \left[x_n^{(\tau)}(t^{(\tau)}) + \frac{\epsilon}{N} \right] \ln \left(\frac{1 + \frac{\epsilon}{N}}{x_n^{(\tau)}(t^{(\tau)}) + \frac{\epsilon}{N}} \right). \quad (30)
 \end{aligned}$$

(29) can then be verified similarly by combining (28) and (30).

Finally, (25) follows by taking the sum of (29) over all n and all episodes, and applying complementary slackness (i.e., $\sum_{n=1}^N x_n^{(\tau)}(t) \sum_{m:n \in \mathcal{S}_m(t)} \beta_m^{(\tau)}(t) = \sum_{m=1}^{M(t)} \beta_m^{(\tau)}(t)$). \square

Step 2-2 (Bounding the portions of the online dual costs): Lemma 4 below connects the online dual cost on the right-hand-side of (25) and (26) to the optimal offline dual cost, which follows from standard duality [45, p. 225]. Please see our technical report [46] for the complete proof of Lemma 4.

Lemma 4. *In any interval from time $t = t_0$ to t_1 , we have*

$$\begin{aligned}
 D^{(\tau)}(t_0 : t_1) & \leq D^{\text{OPT}}(t_0 : t_1) - \sum_{n=1}^N \theta_n^{\text{OPT}}(t_0) x_n^{\text{OPT}}(t_0 - 1) \\
 & + \sum_{n=1}^N \theta_n^{\text{OPT}}(t_1 + 1) x_n^{\text{OPT}}(t_1) + \sum_{n=1}^N \theta_n^{(\tau)}(t_0) x_n^{\text{OPT}}(t_0 - 1) \\
 & - \sum_{n=1}^N \theta_n^{(\tau)}(t_1 + 1) x_n^{\text{OPT}}(t_1), \quad (31)
 \end{aligned}$$

where $x_n^{\text{OPT}}(t)$ and $\theta_n^{\text{OPT}}(t)$ are optimal offline primal and dual solutions, respectively, and $x_n^{(\tau)}(t)$ and $\theta_n^{(\tau)}(t)$ are online primal and dual solutions, respectively.

We can now prove Theorem 2.

Proof of Theorem 2: The total cost of RLA can be calculated as in (3), where the decision $\vec{X}^{\text{RLA}}(t)$ is calculated as in (13). Then, applying Jensen's Inequality, we have that

$$\text{Cost}^{\text{RLA}}(1 : \mathcal{T}) \leq \frac{1}{K+1} \sum_{\tau=0}^K \text{Cost}^{(\tau)}(1 : \mathcal{T}). \quad (32)$$

Then, applying Lemma 2 to (32), we have that the total cost $\text{Cost}^{\text{RLA}}(1 : \mathcal{T})$ of RLA is upper-bounded by

$$\begin{aligned}
 & \frac{1}{K+1} \sum_{\tau=0}^K \sum_{u=-1}^{\lceil \frac{\tau}{K+1} \rceil} \sum_{t^{(\tau)}=\tau+(K+1)u}^{\tau} \left\{ D^{(\tau)}(t^{(\tau)} : t^{(\tau)} + K) \right. \\
 & \left. + \sum_{n=1}^N \Omega_n^{(\tau)}(t^{(\tau)}) + \sum_{n=1}^N \phi_n^{(\tau)}(t^{(\tau)}) + \sum_{n=1}^N \psi_n^{(\tau)}(t^{(\tau)}) \right\}. \quad (33)
 \end{aligned}$$

According to Lemma 1, the online dual costs in (33) add up to $\frac{1}{K+1} \sum_{\tau=0}^K D^{(\tau)}(1 : \mathcal{T}) \leq D^{\text{OPT}}(1 : \mathcal{T})$. It only remains to bound the three tail-terms in (33). We divide into two cases, i.e., $\lceil r_{\text{co}} \rceil < K + 1$ and $\lceil r_{\text{co}} \rceil \geq K + 1$.

i. When $\lceil r_{\text{co}} \rceil < K + 1$, we have $\Delta = \lceil r_{\text{co}} \rceil - 1$. According to Lemma 3, the sum of the tail-terms in (33) can be upper-bounded by

$$\sum_{\tau=0}^K \sum_{u=-1}^{\lceil \frac{\tau}{K+1} \rceil} \sum_{t^{(\tau)}=\tau+(K+1)u} \left\{ D^{(\tau)}(t^{(\tau)} : t^{(\tau)} + \lceil r_{\text{co}} \rceil - 1) + D^{(\tau)}(t^{(\tau)} + K - \lceil r_{\text{co}} \rceil + 1 : t^{(\tau)} + K) \right\} \cdot \eta(1 + \epsilon). \quad (34)$$

Applying Lemma 4 to (34), we can replace $D^{(\tau)}$ by D^{OPT} , with additional tail-terms as shown in (31). When we sum these tail-terms over τ and $t^{(\tau)}$, note that the sum of the tail-terms $-\sum_{n=1}^N \theta_n^{\text{OPT}}(t_0)x_n^{\text{OPT}}(t_0-1)$ and $\sum_{n=1}^N \theta_n^{\text{OPT}}(t_1+1)x_n^{\text{OPT}}(t_1)$ get cancelled across all versions and episodes, and thus can be upper-bounded by 0. The tail-term $-\sum_{n=1}^N \theta_n^{(\tau)}(t_1+1)x_n^{\text{OPT}}(t_1)$ is upper-bounded by 0. Moreover, since the tail-term $\theta_n^{(\tau)}(t_0)x_n^{\text{OPT}}(t_0-1) \leq w_n x_n^{\text{OPT}}(t_0-1)$, the sum of the tail-terms $\sum_{n=1}^N \theta_n^{(\tau)}(t_0)x_n^{\text{OPT}}(t_0-1)$ over all versions and episodes can be upper-bounded by $\max_{\{n,t\}} \frac{w_n}{c_n(t)} \cdot \text{Cost}^{\text{OPT}}(1 : \mathcal{T}) \leq \lceil r_{\text{co}} \rceil \text{Cost}^{\text{OPT}}(1 : \mathcal{T})$. Together, the total cost of RLA is upper-bounded as follows,

$$\begin{aligned} \text{Cost}^{\text{RLA}}(1 : \mathcal{T}) &\leq D^{\text{OPT}}(1 : \mathcal{T}) + \frac{\eta(1 + \epsilon)}{K + 1} \\ &\quad \cdot \{2 \lceil r_{\text{co}} \rceil D^{\text{OPT}}(1 : \mathcal{T}) + \lceil r_{\text{co}} \rceil \text{Cost}^{\text{OPT}}(1 : \mathcal{T})\} \\ &\leq \left\{ 1 + \frac{3\eta(1 + \epsilon) \lceil r_{\text{co}} \rceil}{K + 1} \right\} \text{Cost}^{\text{OPT}}(1 : \mathcal{T}). \end{aligned} \quad (35)$$

This shows (14a).

ii. When $\lceil r_{\text{co}} \rceil \geq K + 1$, we have $\Delta = K$. Similar to the first case, by applying Lemma 3 and Lemma 4, we can show that the total cost of RLA is upper-bounded as follows,

$$\begin{aligned} \text{Cost}^{\text{RLA}}(1 : \mathcal{T}) &\leq D^{\text{OPT}}(1 : \mathcal{T}) + \frac{\eta(1 + \epsilon)}{K + 1} \\ &\quad \cdot \sum_{\tau=0}^K \sum_{u=-1}^{\lceil \frac{\tau}{K+1} \rceil} \sum_{t^{(\tau)}=\tau+(K+1)u} 2D^{(\tau)}(t^{(\tau)} : t^{(\tau)} + K) \\ &\leq \{1 + 2\eta(1 + \epsilon)\} \text{Cost}^{\text{OPT}}(1 : \mathcal{T}). \end{aligned} \quad (36)$$

(14b) then follows. \square

VI. TIGHTENING THE COMPETITIVE RATIO WHEN $K = 0$

Readers may notice that the gap between the competitive ratio of RLA in (14) and the lower bound in (5) grows with η , which is of the order of $\Theta(\ln N)$ when the problem size is large. It would be of interest to see whether this dependency on N can be eliminated. In this section, we show that when $K = 0$, i.e., without look-ahead, this gap can be further tightened to a small constant factor that is independent of not only the coefficient ratio r_{co} , but also the problem size N . Note that as a side product of this result, we also tighten the known competitive ratio of the regularization method [23] and provide a new matching lower bound on the competitive ratio without

look-ahead. We leave for future work whether such tightening of the gap can be attained for $K \geq 1$.

First, according to Theorem 1, the lower bound of the competitive ratio when $K = 0$ is

$$\text{CR}_{K=0}^{\text{LB}} = 1 + \frac{\log_2 N}{2 \left[1 + \frac{1}{r_{\text{co}}} (\log_2 N + 1) \right]}. \quad (37)$$

(37) suggests that, when $K = 0$, the competitive ratio should remain upper-bounded for any value of the coefficient ratio r_{co} and the problem size N . In particular, for any value of the problem size N , the lower bound (37) remains upper-bounded as follows,

$$\text{CR}_{K=0}^{\text{LB}} \leq 1 + \frac{1}{2} r_{\text{co}}, \text{ for all } N \geq 2. \quad (38)$$

In contrast, the competitive ratio $1 + \ln(1 + \frac{N}{\epsilon})(1 + \epsilon)$ of the regularization method obtained in [23] could increase to infinity as N increases. This significant gap thus motivates us to further study whether the competitive ratio of the regularization method can also be improved.

Indeed in Theorem 3 below, we show that the competitive ratio of RLA in the case without look-ahead (i.e., $K = 0$) also remains upper-bounded for any value of N . Recall that $\eta = \ln(\frac{N+\epsilon}{\epsilon})$ and $\epsilon > 0$ are parameters of RLA.

Theorem 3. *Consider the OCO problem introduced in Sec. II-A. When there is no look-ahead, i.e., $K = 0$, the competitive ratio of RLA is upper-bounded as follows,*

$$\text{CR}_{K=0}^{\text{RLA}} \leq 1 + \frac{\eta(1 + \epsilon)}{1 + \frac{\eta}{r_{\text{co}}}}. \quad (39)$$

Note that for any value of the problem size N , the competitive ratio of RLA in (39) keeps upper-bounded as follows,

$$\text{CR}_{K=0}^{\text{RLA}} \leq 1 + r_{\text{co}}(1 + \epsilon), \text{ for all } N \geq 2. \quad (40)$$

Comparing (38) and (40), we can see that the competitive ratio of RLA in (39) matches the lower bound (37) within a small constant factor that is independent of not only the coefficient ratio r_{co} , but also the problem size N . To the best of our knowledge, this is the first such result in the literature for OCO problems. Note that when $K = 0$, our RLA algorithm reduces back to the regularization method in [23]. Therefore, as a side product, we have also tightened the competitive ratio of the regularization method to match the lower bound in (37).

Next, we provide a sketch of proof of Theorem 3 below. Please see our technical report [46] for the complete proof.

Sketch of Proof of Theorem 3: First of all, by letting $K = 0$ in our proofs of Lemma 1 and Lemma 2 in the appendices, it is not difficult to show that Lemma 1 and Lemma 2 hold for $K = 0$. That is, (i) the online dual solution of RLA when $K = 0$ is feasible for the offline dual optimization problem (16), and (ii) the primal online cost at each time can be upper-bounded as in (24) with $K = 0$.

Moreover, when $K = 0$, the sum of the tail-terms $\phi_n^{\text{RLA}}(t)$ and $\psi_n^{\text{RLA}}(t)$ in (24) can be upper-bounded as follows,

$$\sum_{t=1}^{\mathcal{T}} \sum_{n=1}^N \{ \phi_n^{\text{RLA}}(t) + \psi_n^{\text{RLA}}(t) \}$$

$$\begin{aligned}
&= \sum_{t=1}^{\mathcal{T}} \sum_{n=1}^N -\frac{w_n}{\eta} x_n^{\text{RLA}}(t) \ln \left(\frac{x_n^{\text{RLA}}(t) + \frac{\epsilon}{N}}{x_n^{\text{RLA}}(t-1) + \frac{\epsilon}{N}} \right) \\
&= \sum_{t=1}^{\mathcal{T}} \sum_{n=1}^N -\frac{w_n}{\eta} \left[x_n^{\text{RLA}}(t) + \frac{\epsilon}{N} \right] \ln \left(\frac{x_n^{\text{RLA}}(t) + \frac{\epsilon}{N}}{x_n^{\text{RLA}}(t-1) + \frac{\epsilon}{N}} \right) \\
&\quad + \sum_{t=1}^{\mathcal{T}} \sum_{n=1}^N \frac{w_n}{\eta} \frac{\epsilon}{N} \ln \left(\frac{x_n^{\text{RLA}}(t) + \frac{\epsilon}{N}}{x_n^{\text{RLA}}(t-1) + \frac{\epsilon}{N}} \right),
\end{aligned}
\tag{44}$$

where the first equality is according to the definition of the tail-terms in (22) and (23). Next, because of $a \ln \left(\frac{a}{b} \right) \geq a - b$ for any $a, b > 0$ and the telescoping sum, we have

$$\begin{aligned}
&\sum_{t=1}^{\mathcal{T}} \sum_{n=1}^N \{ \phi_n^{\text{RLA}}(t) + \psi_n^{\text{RLA}}(t) \} \\
&\leq \sum_{n=1}^N \frac{w_n}{\eta} \left\{ -x_n^{\text{RLA}}(\mathcal{T}) + \frac{\epsilon}{N} \ln \left(\frac{x_n^{\text{RLA}}(\mathcal{T}) + \frac{\epsilon}{N}}{\frac{\epsilon}{N}} \right) \right\}.
\end{aligned}$$

Then, because $-a + \frac{\epsilon}{N} \ln \left(\frac{a + \frac{\epsilon}{N}}{\frac{\epsilon}{N}} \right) \leq 0$ for any $a \geq 0$ and $\epsilon > 0$, we have

$$\sum_{t=1}^{\mathcal{T}} \sum_{n=1}^N \{ \phi_n^{\text{RLA}}(t) + \psi_n^{\text{RLA}}(t) \} \leq 0. \tag{41}$$

Finally, to prove (39), we only need to upper-bound the sum of the remaining tail-term $\Omega_n^{\text{RLA}}(t)$ as follows,

$$\sum_{t=1}^{\mathcal{T}} \sum_{n=1}^N \Omega_n^{\text{RLA}}(t) \leq \frac{\eta(1 + \epsilon)}{1 + \frac{\eta}{r_{co}}} D^{\text{RLA}}(1 : \mathcal{T}). \tag{42}$$

When $x_n^{\text{RLA}}(t) \leq x_n^{\text{RLA}}(t-1)$, $\Omega_n^{\text{RLA}}(t) \leq 0$. Then, (42) is obviously true. Thus in the following, we only need to focus on the case when $x_n^{\text{RLA}}(t) > x_n^{\text{RLA}}(t-1)$, which implies that $\Omega_n^{\text{RLA}}(t) = w_n [x_n^{\text{RLA}}(t) - x_n^{\text{RLA}}(t-1)]$. We consider the following two sub-cases. Sub-case 1: if $\sum_{m:n \in S_m(t)} \beta_m^{\text{RLA}}(t) \leq c_n(t) + \frac{w_n}{\eta}$, we have

$$\begin{aligned}
&w_n [x_n^{\text{RLA}}(t) - x_n^{\text{RLA}}(t-1)] \\
&\leq w_n \left(x_n^{\text{RLA}}(t) + \frac{\epsilon}{N} \right) \frac{\eta}{w_n} \left(\sum_{m:n \in S_m(t)} \beta_m^{\text{RLA}}(t) - c_n(t) \right) \\
&\leq \left(x_n^{\text{RLA}}(t) + \frac{\epsilon}{N} \right) \sum_{m:n \in S_m(t)} \beta_m^{\text{RLA}}(t) \eta \left(1 - \frac{c_n(t)}{c_n(t) + \frac{w_n}{\eta}} \right) \\
&= \frac{\eta}{1 + \frac{\eta c_n(t)}{w_n}} \left(x_n^{\text{RLA}}(t) + \frac{\epsilon}{N} \right) \sum_{m:n \in S_m(t)} \beta_m^{\text{RLA}}(t), \tag{43}
\end{aligned}$$

where the first inequality is because $a - b \leq a \ln \left(\frac{a}{b} \right)$ for any $a, b > 0$ and the optimality condition of the KKT conditions (see (58b) with $K = 0$), and the second inequality is because of the condition $\sum_{m:n \in S_m(t)} \beta_m^{\text{RLA}}(t) \leq c_n(t) + \frac{w_n}{\eta}$ of this case. Sub-case 2: if $\sum_{m:n \in S_m(t)} \beta_m^{\text{RLA}}(t) > c_n(t) + \frac{w_n}{\eta}$, we have

$$\begin{aligned}
&w_n [x_n^{\text{RLA}}(t) - x_n^{\text{RLA}}(t-1)] \\
&= w_n \left(x_n^{\text{RLA}}(t) + \frac{\epsilon}{N} \right) \left[1 - e^{-\frac{\eta}{w_n} \left(\sum_{m:n \in S_m(t)} \beta_m^{\text{RLA}}(t) - c_n(t) \right)} \right]
\end{aligned}$$

where the first equality is according to the optimality condition of the KKT conditions (see (58b) with $K = 0$), and the first inequality is because $1 - e^{-x} \leq 1$ for any $x \geq 0$ and the condition $\sum_{m:n \in S_m(t)} \beta_m^{\text{RLA}}(t) > c_n(t) + \frac{w_n}{\eta}$ of this case. Finally, by taking the sum of (43) and (44) over all n and t , according to the complementary slackness of the KKT conditions (see (56a) with $K = 0$), we will have (42). This concludes the proof. \square

Remark 1. Notice that when $K \geq 1$, (41) and (42) may not hold. This is exactly where the difficulty lies to tighten the competitive ratio of RLA when $K \geq 1$. However, according to our numerical results in Sec. VIII-C, we conjecture that the competitive ratio of RLA when $K \geq 1$ is also upper-bounded for any value of N . Thus, we believe the true competitive ratio of RLA when $K \geq 1$ may also match the lower bound (5) by a constant factor independent of not only r_{co} and K , but also N . This requires a new competitive analysis method, which we leave for future work.

VII. GENERALIZATION

The fractional covering constraint in (1) corresponds to a demand $a_m(t)$ that is either 1 (when the constraint is present) or 0 (when the constraint is not present). Further, the coefficients on the left-hand-side of (1) must always be 1. Both assumptions are restrictive in practice. In this section, we will extend our results to the more general case, where the decision variables must meet constraints of the type,

$$\sum_{n \in S_m(t)} b_{mn}(t) x_n(t) \geq a_m(t), \text{ for all } m \in [1, M(t)], \tag{45}$$

where $b_{mn}(t)$ and $a_m(t)$ can be any positive integers as in [12], [23], [47]. Moreover, we allow capacity constraints that each decision variable must be upper-bounded, i.e.,

$$x_n(t) \leq X_n^{\text{cap}}, \text{ for all } n \in [1, N], \tag{46}$$

where X_n^{cap} are positive integers. (We do not consider constraints such that the sum of some decision variables needs to be upper-bounded, which will be a subject for future work.)

For this type of OCO problem, with minor modifications, the Regularization with Look-Ahead (RLA) algorithm still works. Specifically, we only need to change $1 + \frac{\epsilon}{N}$ term in the two regularization terms (11b) and (11d) to $X_n^{\text{cap}} + \frac{\epsilon}{N}$, and change η to be $\eta_n \triangleq \ln \left(\frac{X_n^{\text{cap}} + \frac{\epsilon}{N}}{\frac{\epsilon}{N}} \right)$ for each n . Thus, at each time $t^{(\tau)} \in [-K + 1, \mathcal{T}]$, R-FHC $^{(\tau)}$ now calculates the solution to the following problem,

$$\begin{aligned}
&\min_{\bar{X}(t^{(\tau)}:t^{(\tau)}+K)} \left\{ \sum_{t=t^{(\tau)}}^{t^{(\tau)}+K} \sum_{n=1}^N c_n(t) x_n(t) \right. \\
&\quad \left. + \sum_{n=1}^N \frac{w_n}{\eta_n} x_n(t^{(\tau)}) \ln \left(\frac{X_n^{\text{cap}} + \frac{\epsilon}{N}}{x_n^{\text{R-FHC}^{(\tau)}}(t^{(\tau)} - 1) + \frac{\epsilon}{N}} \right) \right\}
\end{aligned}$$

$$\begin{aligned}
& + \sum_{t=t^{(\tau)}+1}^{t^{(\tau)}+K} \sum_{n=1}^N w_n [x_n(t) - x_n(t-1)]^+ \\
& + \sum_{n=1}^N \frac{w_n}{\eta_n} \left[\left(x_n(t^{(\tau)} + K) + \frac{\epsilon}{N} \right) \right. \\
& \quad \left. \cdot \ln \left(\frac{x_n(t^{(\tau)} + K) + \frac{\epsilon}{N}}{X_n^{\text{cap}} + \frac{\epsilon}{N}} \right) - x_n(t^{(\tau)} + K) \right] \} \quad (47a)
\end{aligned}$$

$$\text{sub. to: (11f), (45), (46), for all } t \in [t^{(\tau)}, t^{(\tau)} + K]. \quad (47b)$$

In the analysis, we similarly change $\theta_n^{(\tau)}(t^{(\tau)})$ in (19) to $\frac{w_n}{\eta_n} \ln \left(\frac{X_n^{\text{cap}} + \frac{\epsilon}{N}}{x_n^{(\tau)}(t^{(\tau)}-1) + \frac{\epsilon}{N}} \right)$, which ensures that the online dual variables satisfy the dual constraints. The rest of the analysis then follows the same line, by changing $1 + \frac{\epsilon}{N}$ to $X_n^{\text{cap}} + \frac{\epsilon}{N}$ and by using the knapsack cover (KC) inequalities [48]. Finally, in Theorem 4, we provide the competitive ratio of RLA for this case. Please see our technical report [46] for the complete proof of Theorem 4.

Theorem 4. *Given a look-ahead window of size $K \geq 1$, for the OCO problem with constraints (45) and (46), the competitive ratio of Regularization with Look-Ahead (RLA) is, (with $\eta \triangleq \max_n \eta_n$ and $\bar{B} \triangleq \max_{\{m,n,t\}} b_{mn}(t)$)*

$$\text{CR}^{RLA} = \begin{cases} 1 + \frac{3\eta(1+\epsilon\bar{B})\lceil r_{co} \rceil}{K+1}, & \text{if } \lceil r_{co} \rceil < K+1; \\ 1 + 2\eta(1+\epsilon\bar{B}), & \text{if } \lceil r_{co} \rceil \geq K+1. \end{cases} \quad (48)$$

VIII. NUMERICAL RESULTS

In this section, we demonstrate our theoretic results using numerical experiments. We will mainly focus on the more general OCO problem formulation in Sec. VII with general demand-supply balance constraints (45). Please see our IEEE INFOCOM 2021 paper [1] for numerical results for the less-general OCO problem introduced in Sec. II with fractional covering constraints (1). We choose $\epsilon = 1$ for both Regularization with Look-Ahead (RLA) that we propose in Sec. VII and the regularization method (REG) that was proposed in [23]. First, we show the impact of the coefficient ratio r_{co} and the look-ahead window size K on the empirical competitive ratios (ECRs) of RLA, AFHC [2] and REG [23]. Second, we show the impact of the problem size N on the gap between the ECR of RLA and the lower bound (5).

A. The Simulation Setting for Serverless Computing

Background: We perform a case study on serverless computing [26], [27]. Serverless computing, e.g., Microsoft Azure Serverless Computing and Amazon Lambda, has been a prominent way for customers to deploy applications without the need of worrying about the management of the infrastructure. With serverless computing, customers can dynamically invoke serverless functions on demand, but the service provider has to manage (i.e., starting/stopping) the actual instances executing these functions. Specifically, suppose that there are N functions. At each time t , the customer requests the number $a_n(t)$ of instances needed for function n . Thus, the number $x_n(t)$ of active instances of function n must be no smaller than $a_n(t)$,

which corresponds to the demand-supply constraint (45). Let $x_n(t-1)$ denote the number of instances that are already active. If $a_n(t) > x_n(t-1)$, a cold-start delay [26], [27] will be incurred to start new instances, which corresponds to the switching cost w_n . To avoid such cold-start delay, some instances may be kept active by the service provider even when the demand $a_n(t)$ goes down, which then incurs higher service costs. Thus, the service provider can use the online algorithms in this paper to balance the service cost and the switching cost.

To allow more flexibility in dynamically managing active instances, one possibility is to use the concept of composite functions [49]. A composite function corresponds to an instance that loads the code of multiple functions in memory, and therefore can easily shift the processing across these functions (e.g., by adjusting CPU allocation) without significant switching costs. As an example, let us consider a composite function n_{12} for the non-composite functions n_1 and n_2 . Recall that without composite function n_{12} , the number of instances needs to satisfy, for all time t ,

$$x_{n_1}(t) \geq a_{n_1}(t) \text{ and } x_{n_2}(t) \geq a_{n_2}(t).$$

Thus, if $a_{n_1}(t)$ increases, but $a_{n_2}(t)$ decreases by the same amount, a switching cost on $x_{n_1}(t)$ may be incurred. In contrast, with composite function n_{12} , these demand-supply balance constraints become, for all time t ,


$$\begin{aligned}
x_{n_1}(t) + x_{n_{12}}(t) & \geq a_{n_1}(t), x_{n_2}(t) + x_{n_{12}}(t) \geq a_{n_2}(t), \\
\text{and } x_{n_1}(t) + x_{n_2}(t) + x_{n_{12}}(t) & \geq a_{n_1}(t) + a_{n_2}(t).
\end{aligned}$$

Then, in the same scenario where $a_{n_1}(t)$ increases and $a_{n_2}(t)$ decreases by the same amount, the use of $x_{n_{12}}(t)$ of composite functions may eliminate the need of changing $x_{n_1}(t)$ or $x_{n_2}(t)$. Our goal is therefore to evaluate the performance of our proposed online algorithm for serverless computing with and without composite functions.

Simulation setups: For our simulation, we use the Microsoft's Azure serverless-function traces [26]. Each datum in the trace represents the number of invocations of a function in one minute. There are $\mathcal{T} = 1440$ time-slots, i.e., 1440 minutes for a one-day trace. We discard those functions whose invocation numbers are very small (e.g., 0) or do not change for most time-slots, since they do not contribute much to the performance comparison of the online decisions. We then randomly pick 20 non-composite serverless functions whose variances of the invocation numbers are not too small, e.g., larger than 100. Then, we randomly divide them into 10 pairs. We assign a composite function to each pair of non-composite functions. Therefore, there are 30 functions in total, i.e., $N = 30$.

Further, we generate the cost coefficient as follows. The service-cost coefficient $c_n(t)$ of each non-composite function is randomly generated in the range $[1, 2]$, which represents the resource (e.g., CPU and memory) costs [26]. Then, for each pair of non-composite functions n_1 and n_2 , the service-cost coefficient $c_{n_{12}}(t)$ of the composite function is set to be


$$\begin{aligned}
c_{n_{12}}(t) & = \max \{c_{n_1}(t), c_{n_2}(t)\} \\
& \quad + 0.1 \cdot [c_{n_1}(t) + c_{n_2}(t) - \max \{c_{n_1}(t), c_{n_2}(t)\}].
\end{aligned}$$



(a) ECR vs. look-ahead.

(b) ECR vs. coefficient ratio.


Fig. 2: Compare the ECRs of RLA, AFHC and REG.



(a) Cost vs. look-ahead.

(b) Cost vs. coefficient ratio.

Fig. 3: The value of using the composite function.


 (a) $\lceil r_{co} \rceil < K + 1$

 (b) $\lceil r_{co} \rceil \geq K + 1$

 Fig. 4: Impact of N .

In this way, the service cost of the composite function is higher than both of the two non-composite functions, which models the inherent overhead due to the use of composite functions. To simulate different values of the coefficient ratio r_{co} , the switching-cost coefficient w_n of each non-composite function n is randomly generated in the range $[0.9r_{co}, r_{co}]$, which models the performance loss due to cold-start [26]. Then, for each pair of non-composite functions n_1 and n_2 , the switching-cost coefficient of the composite function is set to be $w_{n_{12}} = \max\{w_{n_1}, w_{n_2}\}$. That is, the cold-start cost of the composite function is equal to the larger cold-start cost of the two non-composite functions.

B. The Impact of r_{co} and K on the ECRs

The numerical results are shown in Fig. 2 and Fig. 3. First, in Fig. 2a, we fix $r_{co} = 50$ and vary K from 1 to 100. We can see that, as the look-ahead window size K increases, the ECRs of RLA and AFHC decrease quickly to a value close to 1 and become much smaller than the ECR of REG. Second, note that the relation between the switching cost w_n and the cold-start time could be affected by various practical factors, e.g., the platform provider and how much the customers dislike the cold-start delay. As a result, the coefficient ratio r_{co} could vary significantly across different scenarios. Therefore, in Fig. 2b, we fix $K = 10$ and vary r_{co} from 5 to 500. We can see that, as the coefficient ratio r_{co} increases, the ECR of AFHC increases to be very large. In contrast, the ECRs of RLA and REG remain at a low value. Fig. 2a and Fig. 2b confirm our analytic results that the competitive ratio of RLA not only decreases with K when r_{co} is small, but also remains upper-bounded for any large value of r_{co} . Moreover, to show the value of using the composite function, we compare in Fig. 3 the total costs of RLA for the case with and without composite functions, labeled as “RLA-composite” and “RLA-no-com”, respectively. We can see that when using the composite functions, the total

costs of serverless computing are indeed lower than the case when no composite function is used.

C. The Impact of N on the ECRs

Recall that in Sec. VI, we tighten the competitive ratio (39) of RLA when $K = 0$, so that the gap from the lower bound (37) remains upper-bounded for any value of problem size N . However, when $K \geq 1$, due to the analytical difficulty that we mentioned in Remark 1, the gap between the proved competitive ratio (14) of RLA and the lower bound (5) increases with $\Theta(\log_2 N)$. Interestingly, according to our simulation results in Fig. 4, we find that the gap between the empirical competitive ratios (ECRs) of RLA and the lower bound (5) remains upper-bounded for any value of N . In the simulation, we use three different traces: the lower-bound trace we designed in Sec. III, an *i.i.d.* random trace and the Microsoft’s Azure Virtual-Machine (VM) trace [50]. In Fig. 4, we show the gaps between the ECRs of RLA and the lower bound (5) for different values of the problem size N . Specifically, we evaluate the gap by dividing the ECRs of RLA by the lower bound (5). To show the impact of N , we change N from 4 to 128 so that $\log_2 N$ increases linearly from 2 to 6. Remember that the theoretical competitive ratio of RLA in (14) depends on the relation between K and r_{co} . To simulate the case when $\lceil r_{co} \rceil < K + 1$, we let $r_{co} = 5$ and $K = 20$ (please see Fig. 4a). To simulate the case when $\lceil r_{co} \rceil \geq K + 1$, we let $r_{co} = 20$ and $K = 5$ (please see Fig. 4b). Fig. 4 shows that the gap does not increase much when $\log_2 N$ increases. Thus, we conjecture that the true competitive ratio of RLA may remain upper-bounded for any value of the problem size N .

IX. CONCLUSION AND FUTURE WORK

In this paper, we study competitive online convex optimization (OCO) with look-ahead. We develop a new online algorithm RLA that can utilize look-ahead to achieve a competitive ratio that not only remains bounded when the coefficient ratio is large, but also decreases with the size of the look-ahead window when the coefficient ratio is small. In this way, the new online algorithm gets the best of both AFHC [2] and the regularization method [23]. To prove the competitive ratio of RLA, we extend the online primal-dual method analysis [28] to the case with look-ahead, which is of independent interest. We also provide a lower bound of the competitive ratio, which matches with the competitive ratio of RLA up to a factor that

only depends on the problem size N . Finally, we generalize RLA to OCO problems with more general constraints.

There are several directions of future work. First, from the experiment results for $K \geq 1$, we observe that the empirical competitive ratio of RLA is only a constant factor (independent of the problem size N) away from the lower bound. Thus, we will study ways to tighten the competitive ratio of RLA when $K \geq 1$. Second, we have not allowed constraints of the form that the sum of some decision variables is upper-bounded. Third, it would be interesting to study whether regularization helps for online maximization problems with packing constraints and study the case with convex service costs. We note that the regularization method in [23] has not been extended to maximization problems or convex service costs either. Thus, it would be of interest to study such extensions. Fourth, it may also be of interest to study how regularization may help to improve the regret, instead of the competitive ratio.

Acknowledgements: We wish to thank Prof. Minghua Chen at City University of Hong Kong and Ms. Jianing Zhai at Chinese University of Hong Kong for pointing out the error in the results in our ACM SIGMETRICS 2019 paper [34].

APPENDIX A PROOF OF LEMMA 1

Proof. To prove Lemma 1, we need to prove, together with the dual variables $\bar{\theta}^{(\tau)}(t^{(\tau)})$ constructed in (19), the online dual variables $\bar{\beta}^{(\tau)}(t)$ and $\bar{\theta}^{(\tau)}(t)$ from each version τ of R-FHC satisfy the constraints (16b)-(16e). We consider one episode from time $t^{(\tau)}$ to $t^{(\tau)} + K$. The proof is similar in all other episodes. (Please see our technical report [46] for the complete proof.)

First, according to the KKT conditions of (11), we have the following inequalities,

$$c_n(t^{(\tau)}) - \sum_{m:n \in S_m(t)} \beta_m^{(\tau)}(t^{(\tau)}) + \frac{w_n}{\eta} \ln \left(\frac{1 + \frac{\epsilon}{N}}{x_n^{(\tau)}(t^{(\tau)} - 1) + \frac{\epsilon}{N}} \right) - \theta_n^{(\tau)}(t^{(\tau)} + 1) \geq 0, \quad \text{for all } n \in [1, N], \quad (49)$$

$$c_n(t) - \sum_{m:n \in S_m(t)} \beta_m^{(\tau)}(t) + \theta_n^{(\tau)}(t) - \theta_n^{(\tau)}(t + 1) \geq 0, \quad \text{for all } n \in [1, N], t \in [t^{(\tau)} + 1, t^{(\tau)} + K - 1], \quad (50)$$

$$c_n(t^{(\tau)} + K) - \sum_{m:n \in S_m(t)} \beta_m^{(\tau)}(t^{(\tau)} + K) + \theta_n^{(\tau)}(t^{(\tau)} + K) - \frac{w_n}{\eta} \ln \left(\frac{1 + \frac{\epsilon}{N}}{x_n^{(\tau)}(t^{(\tau)} + K) + \frac{\epsilon}{N}} \right) \geq 0, \quad \text{for all } n \in [1, N], \quad (51)$$

$$w_n - \theta_n^{(\tau)}(t) \geq 0, \quad \text{for all } n \in [1, N], t \in [t^{(\tau)} + 1, t^{(\tau)} + K], \quad (52)$$

$$\beta_m^{(\tau)}(t) \geq 0, \quad \text{for all } m \in [1, S_m(t)], t \in [t^{(\tau)}, t^{(\tau)} + K], \quad (53)$$

$$\theta_n^{(\tau)}(t) \geq 0, \quad \text{for all } n \in [1, N], t \in [t^{(\tau)} + 1, t^{(\tau)} + K]. \quad (54)$$

Thus, constraint (16b) from time $t^{(\tau)} + 1$ to $t^{(\tau)} + K - 1$, constraint (16c) from time $t^{(\tau)} + 1$ to $t^{(\tau)} + K$, constraint (16d) from time $t^{(\tau)}$ to $t^{(\tau)} + K$, and constraint (16e) from time $t^{(\tau)} + 1$ to $t^{(\tau)} + K$ are satisfied.

Moreover, according to (19), we know $\theta_n^{(\tau)}(t^{(\tau)}) = \frac{w_n}{\eta} \ln \left(\frac{1 + \frac{\epsilon}{N}}{x_n^{(\tau)}(t^{(\tau)} - 1) + \frac{\epsilon}{N}} \right)$ and $\theta_n^{(\tau)}(t^{(\tau)} + K + 1) = \frac{w_n}{\eta} \ln \left(\frac{1 + \frac{\epsilon}{N}}{x_n^{(\tau)}(t^{(\tau)} + K) + \frac{\epsilon}{N}} \right)$. Thus, according to (49) and (51), we have that constraint (16b) at time $t^{(\tau)}$ and $t^{(\tau)} + K$, constraint (16c) at time $t^{(\tau)}$, and constraint (16e) at time $t^{(\tau)}$ are satisfied.

Hence, together with the dual variables $\bar{\theta}^{(\tau)}(t^{(\tau)})$ constructed in (19), the online dual variables $\bar{\beta}^{(\tau)}(t)$ and $\bar{\theta}^{(\tau)}(t)$ from each version τ of R-FHC satisfy the constraints (16b)-(16e). Lemma 1 then follows. \square

APPENDIX B SKETCH OF PROOF OF LEMMA 2

Due to page limits, we only provide a sketch of proof of Lemma 2 in this section. Please see our technical report [46] for the complete proof.

Proof. First, for each version τ of R-FHC, the total cost from time $t^{(\tau)}$ to $t^{(\tau)} + K$ is equal to

$$\begin{aligned} & \text{Cost}^{(\tau)}(t^{(\tau)} : t^{(\tau)} + K) \\ &= \sum_{t=t^{(\tau)}}^{t^{(\tau)}+K} \sum_{n=1}^N c_n(t) x_n^{(\tau)}(t) + \sum_{t=t^{(\tau)}+1}^{t^{(\tau)}+K} \sum_{n=1}^N w_n y_n^{(\tau)}(t) \\ & \quad + \sum_{n=1}^N w_n \left[x_n^{(\tau)}(t^{(\tau)}) - x_n^{(\tau)}(t^{(\tau)} - 1) \right]^+. \end{aligned} \quad (55)$$

Then, notice that the complementary slackness and the optimality condition of the KKT conditions of (11) implies that

$$\sum_{t=t^{(\tau)}}^{t^{(\tau)}+K} \sum_{m=1}^M \beta_m^{(\tau)}(t) \left[1 - \sum_{n \in S_m(t)} x_n^{(\tau)}(t) \right] = 0, \quad (56a)$$

$$\sum_{t=t^{(\tau)}+1}^{t^{(\tau)}+K} \sum_{n=1}^N \theta_n^{(\tau)}(t) \left[x_n^{(\tau)}(t) - x_n^{(\tau)}(t-1) - y_n^{(\tau)}(t) \right] = 0. \quad (56b)$$

By taking the sum of the right-hand-side of (55) and the left-hand-side of the two equations in (56), together with (19), we have that the total cost is equal to

$$\begin{aligned} & \text{Cost}^{(\tau)}(t^{(\tau)} : t^{(\tau)} + K) \\ &= \sum_{t=t^{(\tau)}}^{t^{(\tau)}+K} \sum_{m=1}^M \beta_m^{(\tau)}(t) + \sum_{t=t^{(\tau)}}^{t^{(\tau)}+K} \sum_{n=1}^N x_n^{(\tau)}(t) \left[c_n(t) \right. \\ & \quad \left. - \sum_{m:n \in S_m(t)} \beta_m^{(\tau)}(t) + \theta_n^{(\tau)}(t) - \theta_n^{(\tau)}(t + 1) \right] \\ & \quad + \sum_{t=t^{(\tau)}+1}^{t^{(\tau)}+K} \sum_{n=1}^N y_n^{(\tau)}(t) \left[w_n - \theta_n^{(\tau)}(t) \right] \end{aligned}$$

$$\begin{aligned}
& + \sum_{n=1}^N w_n \left[x_n^{(\tau)}(t^{(\tau)}) - x_n^{(\tau)}(t^{(\tau)} - 1) \right]^+ \\
& - \sum_{n=1}^N \frac{w_n}{\eta} x_n^{(\tau)}(t^{(\tau)}) \ln \left(\frac{1 + \frac{\epsilon}{N}}{x_n^{(\tau)}(t^{(\tau)} - 1) + \frac{\epsilon}{N}} \right) \\
& + \sum_{n=1}^N \frac{w_n}{\eta} x_n^{(\tau)}(t^{(\tau)} + K) \ln \left(\frac{1 + \frac{\epsilon}{N}}{x_n^{(\tau)}(t^{(\tau)} + K) + \frac{\epsilon}{N}} \right). \tag{57}
\end{aligned}$$

Finally, notice that the complementary slackness and the optimality condition of the KKT conditions of (11) implies that

$$\sum_{t=t^{(\tau)}+1}^{t^{(\tau)}+K} \sum_{n=1}^N y_n^{(\tau)}(t) \left[w_n - \theta_n^{(\tau)}(t) \right] = 0, \tag{58a}$$

$$\begin{aligned}
& \sum_{t=t^{(\tau)}}^{t^{(\tau)}+K} \sum_{n=1}^N x_n^{(\tau)}(t) \left[c_n(t) - \sum_{m:n \in S_m(t)} \beta_m^{(\tau)}(t) \right. \\
& \left. + \theta_n^{(\tau)}(t) - \theta_n^{(\tau)}(t+1) \right] = 0. \tag{58b}
\end{aligned}$$

Lemma 2 then follows by combining (57) and (58). \square

REFERENCES

- [1] M. Shi, X. Lin, and L. Jiao, "Combining regularization with look-ahead for competitive online convex optimization," in *IEEE Conference on Computer Communications (INFOCOM)*, 2021, pp. 1–10.
- [2] M. Lin, Z. Liu, A. Wierman, and L. L. Andrew, "Online algorithms for geographical load balancing," in *International Green Computing Conference (IGCC)*. IEEE, 2012, pp. 1–10.
- [3] M. Lin, A. Wierman, L. L. Andrew, and E. Thereska, "Dynamic right-sizing for power-proportional data centers," *IEEE/ACM Transactions on Networking (TON)*, vol. 21, no. 5, pp. 1378–1391, 2013.
- [4] Z. Liu, A. Wierman, Y. Chen, B. Razon, and N. Chen, "Data center demand response: Avoiding the coincident peak via workload shifting and local generation," *Performance Evaluation*, vol. 70, no. 10, pp. 770–791, 2013.
- [5] M. Shi, X. Lin, and S. Fahmy, "Competitive online convex optimization with switching costs and ramp constraints," *IEEE/ACM Transactions on Networking*, vol. 29, no. 2, pp. 876–889, 2021.
- [6] Y. Jia, C. Wu, Z. Li, F. Le, and A. Liu, "Online scaling of NFV service chains across geo-distributed datacenters," *IEEE/ACM Transactions on Networking (TON)*, vol. 26, no. 2, pp. 699–710, 2018.
- [7] Z. Liu, I. Liu, S. Low, and A. Wierman, "Pricing data center demand response," *ACM SIGMETRICS Performance Evaluation Review*, vol. 42, no. 1, pp. 111–123, 2014.
- [8] H. Wang, J. Huang, X. Lin, and H. Mohsenian-Rad, "Exploring smart grid and data center interactions for electric power load balancing," *ACM SIGMETRICS Performance Evaluation Review*, vol. 41, no. 3, pp. 89–94, 2014.
- [9] N. Chen, A. Agarwal, A. Wierman, S. Barman, and L. L. Andrew, "Online convex optimization using predictions," *ACM SIGMETRICS Performance Evaluation Review*, vol. 43, no. 1, pp. 191–204, 2015.
- [10] N. Chen, J. Comden, Z. Liu, A. Gandhi, and A. Wierman, "Using predictions in online optimization: Looking forward with an eye on the past," *ACM SIGMETRICS Performance Evaluation Review*, vol. 44, no. 1, pp. 193–206, 2016.
- [11] L. Jiao, A. M. Tulino, J. Llorca, Y. Jin, and A. Sala, "Smoothed online resource allocation in multi-tier distributed cloud networks," *IEEE/ACM Transactions on Networking (TON)*, vol. 25, no. 4, pp. 2556–2570, 2017.
- [12] L. Jiao, L. Pu, L. Wang, X. Lin, and J. Li, "Multiple granularity online control of cloudlet networks for edge computing," in *15th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*, 2018, pp. 1–9.
- [13] L. Lu, J. Tu, C.-K. Chau, M. Chen, and X. Lin, "Online energy generation scheduling for microgrids with intermittent energy sources and co-generation," *ACM SIGMETRICS Performance Evaluation Review*, vol. 41, no. 1, pp. 53–66, 2013.
- [14] N. Bansal, A. Gupta, R. Krishnaswamy, K. Pruhs, K. Schewior, and C. Stein, "A 2-competitive algorithm for online convex optimization with switching costs," in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2015)*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2015, pp. 96–109.
- [15] M. H. Hajiesmaili, C.-K. Chau, M. Chen, and L. Huang, "Online microgrid energy generation scheduling revisited: The benefits of randomization and interval prediction," in *Proceedings of the 7th International Conference on Future Energy Systems*. ACM, 2016, pp. 1–11.
- [16] S.-J. Kim and G. B. Giannakis, "An online convex optimization approach to real-time energy pricing for demand response," *IEEE Transactions on Smart Grid*, vol. 8, no. 6, pp. 2784–2793, 2017.
- [17] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *Proceedings of the 20th International Conference on Machine Learning (ICML)*. AAAI Press, 2003, pp. 928–935.
- [18] E. Hazan, A. Agarwal, and S. Kale, "Logarithmic regret algorithms for online convex optimization," *Machine Learning*, vol. 69, no. 2-3, pp. 169–192, 2007.
- [19] L. Xiao, "Dual averaging methods for regularized stochastic learning and online optimization," *Journal of Machine Learning Research*, vol. 11, no. Oct, pp. 2543–2596, 2010.
- [20] S. Shalev-Shwartz, "Online learning and online convex optimization," *Foundations and Trends® in Machine Learning*, vol. 4, no. 2, pp. 107–194, 2012.
- [21] M. Shi, Y. Liang, and N. Shroff, "Near-optimal adversarial reinforcement learning with switching costs," in *The Eleventh International Conference on Learning Representations (ICLR)*, 2023.
- [22] N. Buchbinder and J. Naor, "Online primal-dual algorithms for covering and packing problems," in *European Symposium on Algorithms*. Springer, 2005, pp. 689–701.
- [23] N. Buchbinder, S. Chen, and J. S. Naor, "Competitive analysis via regularization," in *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, 2014, pp. 436–444.
- [24] M. Shi, X. Lin, and L. Jiao, "Power-of-2-arms for bandit learning with switching costs," in *Proceedings of the Twenty-Third International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing*, 2022, pp. 131–140.
- [25] A. Borodin and R. El-Yaniv, *Online Computation and Competitive Analysis*. Cambridge University Press, 2005.
- [26] M. Shahrhad, R. Fonseca, Í. Goiri, G. Chaudhry, P. Batum, J. Cooke, E. Laureano, C. Tresness, M. Russinovich, and R. Bianchini, "Serverless in the wild: Characterizing and optimizing the serverless workload at a large cloud provider," in *2020 USENIX Annual Technical Conference (USENIX ATC 20)*, 2020, pp. 205–218.
- [27] I. Baldini, P. Castro, K. Chang, P. Cheng, S. Fink, V. Ishakian, N. Mitchell, V. Muthusamy, R. Rabbah, A. Slominski *et al.*, "Serverless computing: Current trends and open problems," *Research Advances in Cloud Computing*, p. 1, 2017.
- [28] N. Buchbinder and J. S. Naor, "The design of competitive online algorithms via a primal–dual approach," *Foundations and Trends® in Theoretical Computer Science*, vol. 3, no. 2–3, pp. 93–263, 2009.
- [29] N. Buchbinder, S. Chen, J. S. Naor, and O. Shamir, "Unified algorithms for online learning and competitive analysis," in *Proceedings of the 25th Annual Conference on Learning Theory*. Proceedings of Machine Learning Research, 2012, pp. 23:5.1–5.18.
- [30] N. Chen, G. Goel, and A. Wierman, "Smoothed online convex optimization in high dimensions via online balanced descent," in *Proceedings of the 31st Conference On Learning Theory*, 2018, pp. 1574–1594.
- [31] J. B. Rawlings and D. Q. Mayne, *Model Predictive Control: Theory and Design*. Nob Hill Pub. Madison, Wisconsin, 2009.
- [32] E. F. Camacho and C. B. Alba, *Model Predictive Control*. Springer Science & Business Media, 2013.
- [33] Y. Li, G. Qu, and N. Li, "Online optimization with predictions and switching costs: Fast algorithms and the fundamental limit," *arXiv preprint arXiv:1801.07780*, 2018.
- [34] M. Shi, X. Lin, and L. Jiao, "On the value of look-ahead in competitive online convex optimization," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 3, no. 2, pp. 1–42, 2019.
- [35] A. Borodin, N. Linial, and M. E. Saks, "An optimal on-line algorithm for metrical task system," *Journal of the ACM (JACM)*, vol. 39, no. 4, pp. 745–763, 1992.

- [36] N. Alon, B. Awerbuch, and Y. Azar, "The online set cover problem," in *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, 2003, pp. 100–105.
- [37] S. Mannor, J. N. Tsitsiklis, and J. Y. Yu, "Online learning with sample path constraints," *Journal of Machine Learning Research*, vol. 10, no. 3, 2009.
- [38] J. Abernethy, P. L. Bartlett, A. Rakhlin, and A. Tewari, "Optimal strategies and minimax lower bounds for online convex games," in *21st Annual Conference on Learning Theory (COLT)*, 2008.
- [39] J. Friedman and N. Linial, "On convex body chasing," *Discrete & Computational Geometry*, vol. 9, no. 3, pp. 293–321, 1993.
- [40] S. Bubeck, B. Klartag, Y. T. Lee, Y. Li, and M. Sellke, "Chasing nested convex bodies nearly optimally," in *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2020, pp. 1496–1508.
- [41] M. Sellke, "Chasing convex bodies optimally," in *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2020, pp. 1509–1518.
- [42] D. D. Sleator and R. E. Tarjan, "Amortized efficiency of list update and paging rules," *Communications of the ACM*, vol. 28, no. 2, pp. 202–208, 1985.
- [43] N. Bansal, N. Buchbinder, A. Madry, and J. Naor, "A polylogarithmic-competitive algorithm for the k-server problem," *Journal of the ACM (JACM)*, vol. 62, no. 5, pp. 1–49, 2015.
- [44] Y. Lin, G. Goel, and A. Wierman, "Online optimization with predictions and non-convex losses," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 4, no. 1, pp. 1–32, 2020.
- [45] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [46] M. Shi, X. Lin and L. Jiao, "Combining regularization with look-ahead for competitive online convex optimization," Purdue University, Tech. Rep., 2021. Available at <https://engineering.purdue.edu/~7elinx/papers.html>.
- [47] L. Jiao, R. Zhou, X. Lin, and X. Chen, "Online scheduling of traffic diversion and cloud scrubbing with uncertainty in current inputs," in *ACM MOBIHOC*, 2019.
- [48] R. D. Carr, L. K. Fleischer, V. J. Leung, and C. A. Phillips, "Strengthening integrality gaps for capacitated network design and covering problems," in *Proceedings of the 11th annual ACM-SIAM symposium on Discrete algorithms*, 2000, pp. 106–115.
- [49] A. Suresh, G. Somashekar, A. Varadarajan, V. R. Kakarla, H. Upadhyay, and A. Gandhi, "Ensure: Efficient scheduling and autonomous resource management in serverless environments," in *2020 IEEE International Conference on Autonomic Computing and Self-Organizing Systems (AC-SOS)*. IEEE, 2020, pp. 1–10.
- [50] E. Cortez, A. Bonde, A. Muzio, M. Russinovich, M. Fontoura, and R. Bianchini, "Resource central: Understanding and predicting workloads for improved resource management in large cloud platforms," in *Proceedings of the 26th Symposium on Operating Systems Principles*, 2017, pp. 153–167.




Xiaojun Lin (S'02 M'05 SM'12 F'17) received his B.S. from Zhongshan University, Guangzhou, China, in 1994, and his M.S. and Ph.D. degrees from Purdue University, West Lafayette, IN, in 2000 and 2005, respectively. He joined the School of Electrical and Computer Engineering at Purdue University in 2005, and became a Professor of ECE in 2017. Since June 2023, he joined the Department of Information Engineering, The Chinese University of Hong Kong, as a Professor and Global STEM Scholar.

Dr. Lin's research interests are in the analysis, control and optimization of large and complex networked systems, including both communication networks and power grid. He received the NSF CAREER award in 2007. He received 2005 best paper of the year award from Journal of Communications and Networks, IEEE INFOCOM 2008 best paper award, ACM MobiHoc 2021 best paper award, and ACM e-Energy 2022 best paper award. He was the Workshop co-chair for IEEE GLOBECOM 2007, the Panel co-chair for WICON 2008, the TPC co-chair for ACM MobiHoc 2009, the Mini-Conference co-chair for IEEE INFOCOM 2012, and the General co-chair for ACM e-Energy 2019. He has served as an Area Editor for (Elsevier) Computer Networks Journal, an Associate Editor for IEEE/ACM Transactions on Networking, and a Guest Editor for (Elsevier) Ad Hoc Networks journal.



Lei Jiao received the Ph.D. degree in computer science from the University of Göttingen, Germany. He is currently a faculty member at the Department of Computer Science, University of Oregon, USA. Previously he worked at Nokia Bell Labs in Ireland. He is interested in the mathematics of optimization, control, learning, and mechanism design applied to computer and telecommunication systems, networks, and services. He is a recipient of the NSF CAREER award. He publishes papers in journals such as JSAC, ToN, TPDS, and TMC and in conferences

such as INFOCOM, MOBIHOC, ICNP, ICDCS, SECON, and IPDPS. He also received the Best Paper Awards of IEEE LANMAN 2013 and IEEE CNS 2019, and the 2016 Alcatel-Lucent Bell Labs UK and Ireland Recognition Award. He has been on the program committees of INFOCOM, MOBIHOC, ICDCS, TheWebConf, and IWQoS, and served as the program chair of multiple workshops with INFOCOM and ICDCS.



Ming Shi received his B.S. degree from Tianjin University, Tianjin, China, in 2015, and his Ph.D. degree from Purdue University, West Lafayette, IN, USA, in 2022. In June 2022, he joined the Department of Electrical and Computer Engineering and the AI-EDGE Institute, The Ohio State University, as a Post-Doctoral Scholar. Dr. Shi's research interests are in the analysis, e.g., the regret and competitive ratio, of learning algorithms and online optimization algorithms, with applications in Edge AI, networking and cloud computing.