Combining Regularization with Look-Ahead for Competitive Online Convex Optimization

Ming Shi, Xiaojun Lin
School of Electrical and Computer Engineering
Purdue University, West Lafayette, IN, USA
{sming@purdue.edu,linx@ecn.purdue.edu}

Lei Jiao
Department of Computer and Information Science
University of Oregon, Eugene, Oregon, USA
jiao@cs.uoregon.edu

IEEE INFOCOM 2021 - IEEE Conference on Computer Communications | 978-1-6654-0325-2/21/$31.00 ©2021 IEEE | DOI: 10.1109/INFOCOM42981.2021.9488766

***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. We also 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 by a factor that only depends on the problem size. The competitive analysis of RLA involves a non-trivial generalization of online primal-dual analysis to the case with look-ahead.**

***Index Terms*—online convex optimization, competitive analysis, look-ahead, regularization.**

I. Introduction

Online convex optimization (OCO) problem with switching costs has many applications in the context of networking [1]–[5], cloud or edge computing [6]–[11], cyber-physical systems [12]–[15], machine learning [16]–[19] and beyond [20], [21]. 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 for the total cost [22, p. 3].

In the literature, many online algorithms with guaranteed competitive ratios have been provided for OCO. For example, [2], [12]–[14] 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 [21], [23]–[25] 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 [1], [8], [26]–[28]. 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 [1], 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 $1+\max\limits_{\{n,t\}}\frac{w_n}{c_n(t)(K+1)}$. 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., $r_{\text{co}} \triangleq \max\limits_{\{n,t\}}\frac{w_n}{c_n(t)}$. 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 [28] 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 [21] can achieve a competitive ratio that is independent of the coefficient ratio r_{co}. Of course, the downside of the regularization method of [21] 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 [29] claimed to achieve this by providing a $\left(1+\frac{1}{K}\right)$-competitive online algorithm. Unfortunately, there appears to be an error in the proof so that the claimed competitive ratio does not hold [30]. (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 [11], [20]–[24], [31], [32] (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 $1+\frac{\log_2 N}{2\left[1+\frac{1}{r_{\text{co}}}((K+1)\log_2 N+1)\right]}$, 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 [29] 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 $\eta \triangleq \ln\left(\frac{N+\epsilon}{\epsilon}\right)$, 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 [1] directly optimizes the total cost. However, as shown in the counterexample in [29], 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 [21], 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 [23]. 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 [23] 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. VI 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 [33], [34]. In particular, [33] shows that one cannot simultaneously obtain sublinear 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 [33], [34] focus on the *absolute difference* from the cost of the best *static* decision. Thus, even if sublinear regret is not attainable, it is still possible to attain a low competitive ratio.

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 $t=1,...,\mathcal{T}$, first a feasible convex set $\mathbb{X}(t)$ and service-cost coefficients $\vec{C}(t)=[c_n(t), n=1,...,N]^{\text{T}} \in \mathbb{R}_+^{N\times 1}$ are revealed, where $[\cdot]^{\text{T}}$ 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.,

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

where $S_m(t)$ is a subset of $\{1,2,...,N\}$ 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 [20], [22], [31], [32], [35], [36]. 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. VI, 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 $\vec{X}(t) = [x_n(t), n = 1, ..., N]^{\mathrm{T}} \in \mathbb{R}_+^{N\times 1}$ 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 increment[1] of $\vec{X}(t)$ from the last decision $\vec{X}(t-1)$, where $\vec{W} = [w_n, n = 1, ..., N]^{\mathrm{T}} \in \mathbb{R}_+^{N\times 1}$ 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,

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

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

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

where $[a, b]$ denotes the set $\{a, a + 1, ..., b\}$. As typically in many OCO problems [1], [4], [9], [21], 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 [1], [8], [27], [28], [37]. 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,

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

[1]Note that, as shown in [24], 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 $|\vec{X}(t) - \vec{X}(t-1)|$ [4], [13].

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,

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

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

III. A Lower Bound

Although OCO with look-ahead has been extensively studied, e.g., in [1], [8], [37], 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 [23] and [37] 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.

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

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

Please see Appendix A for the proof. 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,

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

Note that the right-hand-side is close to the competitive ratio of AFHC [1].

(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,

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

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 [1] will produce a low competitive ratio (see the dashed curves in Fig. 1). In

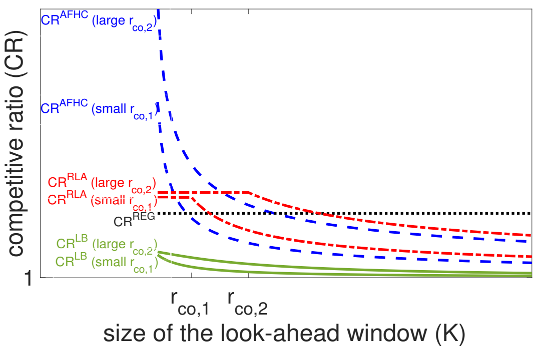

Fig. 1: Compare the lower bound of the competitive ratio ($\mathrm{CR}^{\mathrm{LB}}$) and the competitive ratio of AFHC ($\mathrm{CR}^{\mathrm{AFHC}}$), the regularization method ($\mathrm{CR}^{\mathrm{REG}}$) and RLA ($\mathrm{CR}^{\mathrm{RLA}}$).

contrast, the competitive ratio of the regularization method (REG) of [21] 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 $K+1 \le r_{\mathrm{co}}$, (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 $\mathrm{CR}^{\mathrm{LB}}$ and the competitive ratio of the regularization method $\mathrm{CR}^{\mathrm{REG}}$ 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,

$$\mathrm{CR}^{\mathrm{LB}} \le 1 + \frac{1}{2}\log_2 N, \tag{8}$$

which suggests room for improvement for AFHC.

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 R-FHC$^{(\tau)}$. 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 $u = -1, 0, ..., \left\lceil \frac{\mathcal{T}}{K+1} \right\rceil$. Recall that at time $t^{(\tau)}$, the inputs $\left(\mathbb{X}(t^{(\tau)} : t^{(\tau)}+K), \vec{C}(t^{(\tau)} : t^{(\tau)}+K)\right)$ at the current time and in the look-ahead window have been revealed. R-FHC$^{(\tau)}$ then computes the solution to the following problem,

$$\min_{\vec{X}(t^{(\tau)}:t^{(\tau)}+K)} \left\{ \sum_{s=t^{(\tau)}}^{t^{(\tau)}+K} \sum_{n=1}^{N} c_n(s) x_n(s) \right. \tag{9a}$$

$$+ \sum_{n=1}^{N} \frac{w_n}{\eta} x_n(t^{(\tau)}) \ln \left(\frac{1+\frac{\epsilon}{N}}{x_n^{\text{R-FHC}^{(\tau)}}(t^{(\tau)}-1)+\frac{\epsilon}{N}} \right) \tag{9b}$$

$$+ \sum_{s=t^{(\tau)}+1}^{t^{(\tau)}+K} \sum_{n=1}^{N} w_n \left[x_n(s) - x_n(s-1) \right]^+ \tag{9c}$$

$$+ \sum_{n=1}^{N} \frac{w_n}{\eta} \left[\left(x_n(t^{(\tau)}+K) + \frac{\epsilon}{N} \right) \right.$$

$$\left. \left. \cdot \ln \left(\frac{x_n(t^{(\tau)}+K)+\frac{\epsilon}{N}}{1+\frac{\epsilon}{N}} \right) - x_n(t^{(\tau)}+K) \right] \right\} \tag{9d}$$

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

$$s \in [t^{(\tau)}, t^{(\tau)}+K], \tag{9e}$$

$$x_n(s) \ge 0, \text{ for all } n \in [1, N], s \in [t^{(\tau)}, t^{(\tau)}+K], \tag{9f}$$

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 R-FHC$^{(\tau)}$ from time $t^{(\tau)}-K-1$ to $t^{(\tau)}-1$.

According to (9), in each episode from time $t^{(\tau)}$ to $t^{(\tau)}+K$, RLA does not simply optimize the corresponding service costs and switching costs. Instead, it replaces the switching cost in the first time-slot $t^{(\tau)}$ of the current episode by the regularization term (9b), and adds another regularization term (9d) for the decision variables in the last time-slot $t^{(\tau)}+K$ of the current episode. Similar to [21], the regularization term (9d) makes the objective function strictly convex in $x_n(t^{(\tau)}+K)$, and thus discourages it from taking extreme values. More specifically, without (9d), 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, (9d) 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 (9b), 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 [21] 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 [21]. However, our formulation of the regularization terms for $K \ge 1$ 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 $\vec{X}^{\mathrm{RLA}}(t)$ at time t. As K increases, since R-FHC$^{(\tau)}$ 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$^{(\tau)}$, 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$^{(\tau)}$ whose first episode starts at time $t^{(\tau)} \le 0$, (9b) can be removed.

Algorithm 1 Regularization with Look-Ahead (RLA)

Parameters: $\epsilon > 0$ and $\eta = \ln\left(\frac{N+\epsilon}{\epsilon}\right)$.
FOR $t = -K+1 : \mathcal{T}$
Step 1: $\tau \leftarrow t \bmod (K+1)$ and $t^{(\tau)} \leftarrow t$.
Step 2: Solve (9) to get $\vec{X}^{\text{R-FHC}^{(\tau)}}(t^{(\tau)} : t^{(\tau)} + K)$. (If $t^{(\tau)} \leq 0$, remove (9b). If $t^{(\tau)} \geq \mathcal{T} - K$, remove (9d).)
Step 3: **if** $1 \leq t \leq \mathcal{T}$, **then** let

$$\vec{X}^{\text{RLA}}(t) = \frac{1}{K+1}\sum_{\tau=0}^{K}\vec{X}^{\text{R-FHC}^{(\tau)}}(t). \quad (10)$$

end if
END

Similarly, for any version of R-FHC$^{(\tau)}$ whose last episode ends at time $t^{(\tau)} + K \geq \mathcal{T}$, (9d) 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,*

$$\text{CR}^{RLA} = 1 + \frac{3\eta(1+\epsilon)\lceil r_{co}\rceil}{K+1}, \ \textit{if } \lceil r_{co}\rceil < K+1; \quad (11a)$$

$$\text{CR}^{RLA} = 1 + 2\eta(1+\epsilon), \ \textit{if } \lceil r_{co}\rceil \geq K+1. \quad (11b)$$

It is easy to see that the competitive ratio of RLA in (11) 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 (11a) and (11b) 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 $1 + \frac{\log_2 N}{2\left[\frac{3}{2} + \log_2 N\right]}$. Thus, the gap between (11b) and (5) is at most $\frac{4\eta(1+\epsilon)[\frac{3}{2}+\log_2 N]}{\log_2 N}$. Further, when $r_{\text{co}} \geq (K+1)\log_2 N$, the gap between (11b) 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 [23]. 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

$$y_n(t) \geq x_n(t) - x_n(t-1), \text{ for all } n \in [1, N], \quad (12)$$

we can get an equivalent formulation of the offline optimization problem (2). Then, let $\vec{\beta}(t) = [\beta_m(t), m = 1, ..., M(t)]^{\text{T}}$ and $\vec{\theta}(t) = [\theta_n(t), n = 1, ..., N]^{\text{T}}$ be the Lagrange multipliers for constraints (2c) and (12), respectively. We have the offline dual optimization problem as follows,

$$\max_{\{\vec{\beta}(1:\mathcal{T}), \vec{\theta}(1:\mathcal{T})\}} \sum_{t=1}^{\mathcal{T}}\sum_{m=1}^{M(t)}\beta_m(t) \quad (13a)$$

$$\text{sub. to: } c_n(t) - \sum_{m:n\in S_m(t)}\beta_m(t) + \theta_n(t) - \theta_n(t+1) \geq 0,$$
$$\text{for all } n \in [1,N],\ t \in [1,\mathcal{T}], \quad (13b)$$

$$w_n - \theta_n(t) \geq 0, \text{ for all } n \in [1,N],\ t \in [1,\mathcal{T}], \quad (13c)$$

$$\beta_m(t) \geq 0, \text{ for all } m \in [1, M(t)], t \in [1,\mathcal{T}], \quad (13d)$$

$$\theta_n(t) \geq 0, \text{ for all } n \in [1,N],\ t \in [1,\mathcal{T}]. \quad (13e)$$

Let $\beta_m^{\text{OPT}}(t)$ and $\theta_n^{\text{OPT}}(t)$ be the optimal solution to (13). Then, the optimal offline dual cost is,

$$D^{\text{OPT}}(1:\mathcal{T}) \triangleq \sum_{t=1}^{\mathcal{T}}\sum_{m=1}^{M(t)}\beta_m^{\text{OPT}}(t). \quad (14)$$

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,

$$\text{Cost}^{\text{RLA}}(1:\mathcal{T}) \overset{(a)}{\leq} \text{CR}\cdot D^{\text{RLA}}(1:\mathcal{T})$$
$$\overset{(b)}{\leq} \text{CR}\cdot D^{\text{OPT}}(1:\mathcal{T}) \overset{(c)}{\leq} \text{CR}\cdot\text{Cost}^{\text{OPT}}(1:\mathcal{T}). \quad (15)$$

In (15), step (c) simply follows from standard duality [38, 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 (13). Since (13) is a maximization problem, step (b) then holds. Finally, step (a) is related to the regularization terms (9b) and (9d) 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 (13). Focus on one episode from time $t^{(\tau)}$ to $t^{(\tau)}+K$. As in (13), we introduce the variable $y_n(t)$ and the constraint (12) to (9). We can then form the dual problem of the equivalent form of (9). As in (13), we let $\beta_m^{(\tau)}(t)$ and $\theta_n^{(\tau)}(t)$ be the corresponding online dual solution of (9). However, note that the objective function of (9) 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

$$\theta_n^{(\tau)}(t^{(\tau)}) \triangleq \frac{w_n}{\eta}\ln\left(\frac{1+\frac{\epsilon}{N}}{x_n^{(\tau)}(t^{(\tau)}-1)+\frac{\epsilon}{N}}\right). \quad (16)$$

Lemma 3 below shows that we have constructed a feasible dual solution for the offline dual optimization problem (13). Due to page limits, we omit the detailed proofs of Lemma 3 and several other results below. Please see our technical report [39] for their complete proofs.

Lemma 3. *The $\vec{\beta}^{(\tau)}(1:\mathcal{T})$ and $\vec{\theta}^{(\tau)}(1:\mathcal{T})$ constructed above from (16) and the online dual solution of R-FHC$^{(\tau)}$ are feasible for the offline dual optimization problem (13).*

Lemma 3 can be proved by verifying that the Karush-Kuhn-Tucker (KKT) conditions [38, p. 243] of (9) satisfies the dual constraints (13b)-(13e). (13c) to (13e) are easy to verify, so is (13b) for $t = t^{(\tau)} + 1$ to $t^{(\tau)} + K - 1$, because the KKT conditions for (9) in those time-slots are exactly the same as that of (13). Thus, it only remains to verify (13b) at time $t = t^{(\tau)}$ and $t = t^{(\tau)} + K$. At time $t^{(\tau)}$, by examining the KKT conditions for (9), 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 (16), (13b) at time $t = t^{(\tau)}$ is verified. We can verify (13b) at time $t^{(\tau)} + K$ similarly. Lemma 3 then follows.

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

$$D^{(\tau)}(t^{(\tau)} : t^{(\tau)} + K) \triangleq \sum_{t=t^{(\tau)}}^{t^{(\tau)}+K} \sum_{m=1}^{M(t)} \beta_m^{(\tau)}(t). \tag{17}$$

However, note that (9) contains additional terms (9b) and (9d) 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 4 below captures this gap. Define the tail-terms as

$$\Omega_n^{(\tau)}(t^{(\tau)}) \triangleq w_n \left[x_n^{(\tau)}(t^{(\tau)}) - x_n^{(\tau)}(t^{(\tau)}-1)\right]^+, \tag{18}$$

$$\phi_n^{(\tau)}(t^{(\tau)}) \triangleq -\frac{w_n}{\eta} x_n^{(\tau)}(t^{(\tau)}) \ln\left(\frac{1+\frac{\epsilon}{N}}{x_n^{(\tau)}(t^{(\tau)}-1)+\frac{\epsilon}{N}}\right), \tag{19}$$

$$\psi_n^{(\tau)}(t^{(\tau)}) \triangleq \frac{w_n}{\eta} x_n^{(\tau)}(t^{(\tau)}+K) \ln\left(\frac{1+\frac{\epsilon}{N}}{x_n(t^{(\tau)}+K)+\frac{\epsilon}{N}}\right). \tag{20}$$

Lemma 4. *For each version τ of R-FHC, we have,*

$$\textit{Cost}^{(\tau)}(t^{(\tau)} : t^{(\tau)} + K) \leq D^{(\tau)}(t^{(\tau)} : t^{(\tau)} + K) + \sum_{n=1}^{N} \Omega_n^{(\tau)}(t^{(\tau)}) + \sum_{n=1}^{N} \phi_n^{(\tau)}(t^{(\tau)}) + \sum_{n=1}^{N} \psi_n^{(\tau)}(t^{(\tau)}). \tag{21}$$

Lemma 4 captures the gap between the online primal cost and the online dual cost. In (21), the first tail-term $\Omega_n^{(\tau)}(t^{(\tau)})$ is because R-FHC$^{(\tau)}$ does not optimize over the real switching cost $w_n\left[x_n(t^{(\tau)}) - x_n(t^{(\tau)}-1)\right]^+$ 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 (9b) and (9d) added to the primal objective function in the first time-slot and the last time-slot. Lemma 4 can be shown via the duality theorem [38, p.225]. Recall that, to establish step (a) in (15), the main difficulty is to bound this gap, which we divide into the following two sub-steps.

Step 2-1 (Bounding the tail-terms): Next, we show in Lemma 5 that, with a factor that will appear in the final competitive ratio, the tail-terms (18)-(20) 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 5. *For each version τ of R-FHC, the following holds,*

$$(i) \sum_{u=0}^{\lceil \frac{\mathcal{T}}{K+1} \rceil} \sum_{t^{(\tau)}=\tau+(K+1)u} \sum_{n=1}^{N} \Omega_n^{(\tau)}(t^{(\tau)}) \leq \eta(1+\epsilon) \times \sum_{u=0}^{\lceil \frac{\mathcal{T}}{K+1} \rceil} \sum_{t^{(\tau)}=\tau+(K+1)u} D^{(\tau)}(t^{(\tau)} : t^{(\tau)} + \Delta), \tag{22}$$

$$(ii) \sum_{u=-1}^{\lceil \frac{\mathcal{T}}{K+1} \rceil} \sum_{t^{(\tau)}=\tau+(K+1)u} \sum_{n=1}^{N} \left[\phi_n^{(\tau)}(t^{(\tau)}) + \psi_n^{(\tau)}(t^{(\tau)})\right] \leq \eta(1+\epsilon) \sum_{u=-1}^{\lceil \frac{\mathcal{T}}{K+1} \rceil} \sum_{\substack{t^{(\tau)}=\tau \\ +(K+1)u}} D^{(\tau)}(t^{(\tau)}+K-\Delta : t^{(\tau)}+K), \tag{23}$$

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

To interpret (22), the tail-term $\Omega_n^{(\tau)}(t^{(\tau)})$ are bounded by the right-hand-side of (22), 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 (23) has a similar interpretation, while the partial sum is over sub-intervals at the end of each episode.

Sketch of Proof of Lemma 5: We focus on the proof of (22), and (23) 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

$$\frac{w_n}{\eta}[x_n^{(\tau)}(t^{(\tau)}) + \frac{\epsilon}{N}] \ln\left(\frac{1+\frac{\epsilon}{N}}{x_n^{(\tau)}(t^{(\tau)}-1)+\frac{\epsilon}{N}}\right). \tag{24}$$

Let $\hat{\beta}_n^{(\tau)}(t) = \sum_{m:n\in S_m(t)} \beta_m^{(\tau)}(t)$. Consider any $t' > t^{(\tau)}$ such that $x_n^{(\tau)}(t) > 0$ for all $t \in [t^{(\tau)}, t']$. Using KKT conditions of (9), we can show that (24) is equal to

$$\sum_{t=t^{(\tau)}}^{t'} [x_n^{(\tau)}(t) + \frac{\epsilon}{N}]\hat{\beta}_n^{(\tau)}(t) + [x_n^{(\tau)}(t') + \frac{\epsilon}{N}]\theta_n^{(\tau)}(t'+1) - \sum_{t=t^{(\tau)}}^{t'} c_n(t)[x_n^{(\tau)}(t) + \frac{\epsilon}{N}] - \sum_{t=t^{(\tau)}+1}^{t'} w_n y_n^{(\tau)}(t). \tag{25}$$

Next, we show that

$$\frac{\Omega_n^{(\tau)}(t^{(\tau)})}{\eta} \leq (25) \leq \sum_{t=t^{(\tau)}}^{t^{(\tau)}+\Delta} [x_n^{(\tau)}(t) + \frac{\epsilon}{N}]\hat{\beta}_n^{(\tau)}(t) \tag{26}$$

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 (26) 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 (25). 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 $\theta_n^{(\tau)}(t'+1) \le w_n$, the second and fourth term in (25) can be upper-bounded by $w_n[x_n^{(\tau)}(t^{(\tau)}) + \frac{\epsilon}{N}]$. Then, since $x_n^{(\tau)}(t) \ge x_n^{(\tau)}(t^{(\tau)})$ for all $t \in [t^{(\tau)}, t']$ and $\sum_{t=t^{(\tau)}}^{t'} c_n(t) \ge \frac{w_n}{r_{\text{co}}}(\Delta+1) \ge w_n$, the last three terms in (25) are upper-bounded by 0, and (26) then follows. (ii-b) If $\lceil r_{\text{co}} \rceil - 1 \ge K$, we can show that,

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

(26) can then be verified similarly by combining (25) and (27).

Finally, (22) follows by taking the sum of (26) over all n and all episodes, and applying complementary slackness (i.e., $\sum_{n=1}^{N} x_n^{(\tau)}(t) \sum_{m:n\in S_m(t)} \beta_m^{(\tau)}(t) = \sum_{m=1}^{M(t)} \beta_m^{(\tau)}(t)$). □

Step 2-2 (Bounding the portions of the online dual costs): Lemma 6 below connects the online dual cost on the right-hand-side of (22) and (23) to the optimal offline dual cost, which follows from standard duality [38, p. 225].

Lemma 6. *In any interval from time $t = t_0$ to t_1, we have*

$$D^{(\tau)}(t_0 : t_1) \le D^{OPT}(t_0 : t_1) - \sum_{n=1}^{N} \theta_n^{OPT}(t_0) x_n^{OPT}(t_0 - 1) + \sum_{n=1}^{N} \theta_n^{OPT}(t_1+1) x_n^{OPT}(t_1) + \sum_{n=1}^{N} \theta_n^{(\tau)}(t_0) x_n^{OPT}(t_0-1) - \sum_{n=1}^{N} \theta_n^{(\tau)}(t_1+1) x_n^{OPT}(t_1), \tag{28}$$

where $x_n^{OPT}(t)$ and $\theta_n^{OPT}(t)$ 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 (10). Then, applying Jensen's Inequality, we have that,

$$\text{Cost}^{\text{RLA}}(1:\mathcal{T}) \le \frac{1}{K+1} \sum_{\tau=0}^{K} \text{Cost}^{(\tau)}(1:\mathcal{T}). \tag{29}$$

Then, applying Lemma 4 to (29), we have that the total cost $\text{Cost}^{\text{RLA}}(1:\mathcal{T})$ of RLA is upper-bounded by,

$$\frac{1}{K+1} \sum_{\tau=0}^{K} \sum_{u=-1}^{\lceil \frac{\mathcal{T}}{K+1} \rceil} \sum_{t^{(\tau)}=\tau+(K+1)u} \Bigg\{ D^{(\tau)}(t^{(\tau)} : t^{(\tau)} + K) + \sum_{n=1}^{N} \Omega_n^{(\tau)}(t^{(\tau)}) + \sum_{n=1}^{N} \phi_n^{(\tau)}(t^{(\tau)}) + \sum_{n=1}^{N} \psi_n^{(\tau)}(t^{(\tau)}) \Bigg\}. \tag{30}$$

According to Lemma 3, the online dual costs in (30) add up to $\frac{1}{K+1}\sum_{\tau=0}^{K} D^{(\tau)}(1:\mathcal{T}) \le D^{\text{OPT}}(1:\mathcal{T})$. It only remains to bound the three tail-terms in (30). We divide into two cases, i.e., $\lceil r_{\text{co}} \rceil < K+1$ and $\lceil r_{\text{co}} \rceil \ge K+1$.

i. When $\lceil r_{\text{co}} \rceil < K+1$, we have $\Delta = \lceil r_{\text{co}} \rceil - 1$. According to Lemma 5, the sum of the tail-terms in (30) can be upper-bounded by

$$\sum_{\tau=0}^{K} \sum_{u=-1}^{\lceil \frac{\mathcal{T}}{K+1} \rceil} \sum_{t^{(\tau)}=\tau+(K+1)u} \Bigg\{ D^{(\tau)}(t^{(\tau)} : t^{(\tau)} + \lceil r_{\text{co}} \rceil - 1) + D^{(\tau)}(t^{(\tau)} + K - \lceil r_{\text{co}} \rceil + 1 : t^{(\tau)} + K) \Bigg\} \cdot \eta(1+\epsilon). \tag{31}$$

Applying Lemma 6 to (31), we can replace $D^{(\tau)}$ by D^{OPT}, with additional tail-terms as shown in (28). 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 $\theta_n^{(\tau)}(t_0) x_n^{\text{OPT}}(t_0 - 1) \le w_n x_n^{\text{OPT}}(t_0 - 1)$, 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 $\max_{\{n,t\}} \frac{w_n}{c_n(t)} \cdot \text{Cost}^{\text{OPT}}(1 : \mathcal{T}) \le \lceil r_{\text{co}} \rceil \text{Cost}^{\text{OPT}}(1:\mathcal{T})$. Together, the total cost of RLA is upper-bounded by,

$$\begin{aligned} \text{Cost}^{\text{RLA}}(1:\mathcal{T}) &\le D^{\text{OPT}}(1:\mathcal{T}) + \frac{\eta(1+\epsilon)}{K+1} \cdot \left\{2 \lceil r_{\text{co}} \rceil D^{\text{OPT}}(1:\mathcal{T}) + \lceil r_{\text{co}} \rceil \text{Cost}^{\text{OPT}}(1:\mathcal{T})\right\} \\ &\le \left\{1 + \frac{3\eta(1+\epsilon)\lceil r_{\text{co}} \rceil}{K+1}\right\} \text{Cost}^{\text{OPT}}(1:\mathcal{T}). \end{aligned} \tag{32}$$

This shows (11a).

ii. When $\lceil r_{\text{co}} \rceil \ge K+1$, we have $\Delta = K$. Similar to the first case, by applying Lemma 5 and Lemma 6, we can show that the total cost of RLA is upper-bounded by,

$$\begin{aligned} \text{Cost}^{\text{RLA}}(1:\mathcal{T}) &\le D^{\text{OPT}}(1:\mathcal{T}) + \frac{\eta(1+\epsilon)}{K+1} \cdot \sum_{\tau=0}^{K} \sum_{u=-1}^{\lceil \frac{\mathcal{T}}{K+1} \rceil} \sum_{t^{(\tau)}=\tau+(K+1)u} 2D^{(\tau)}(t^{(\tau)} : t^{(\tau)} + K) \\ &\le \{1 + 2\eta(1+\epsilon)\} \text{Cost}^{\text{OPT}}(1:\mathcal{T}). \end{aligned} \tag{33}$$

(11b) then follows. □

VI. 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 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,

$$\sum_{n\in S_m(t)} b_{mn}(t)x_n(t) \geq a_m(t), \text{ for all } m \in [1, M(t)], \quad (34)$$

where $b_{mn}(t)$ and $a_m(t)$ can be any positive integers as in [11], [21], [40]. Moreover, we allow capacity constraints that each decision variable must be upper-bounded, i.e.,

$$x_n(t) \leq X_n^{\text{cap}}, \text{ for all } n \in [1, N], \quad (35)$$

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 (9b) and (9d) 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,

$$\min_{\vec{X}(t^{(\tau)}:t^{(\tau)}+K)} \left\{ \sum_{t=t^{(\tau)}}^{t^{(\tau)}+K} \sum_{n=1}^{N} c_n(t)x_n(t) \right.$$
$$+ \sum_{n=1}^{N} \frac{w_n}{\eta_n} x_n(t^{(\tau)}) \ln\left(\frac{X_n^{\text{cap}} + \frac{\epsilon}{N}}{x_n^{\text{R-FHC}^{(\tau)}}(t^{(\tau)}-1) + \frac{\epsilon}{N}}\right)$$
$$+ \sum_{t=t^{(\tau)}+1}^{t^{(\tau)}+K} \sum_{n=1}^{N} w_n \left[x_n(t) - x_n(t-1)\right]^+$$
$$+ \sum_{n=1}^{N} \frac{w_n}{\eta_n} \left[\left(x_n(t^{(\tau)}+K) + \frac{\epsilon}{N}\right) \right.$$
$$\left.\left. \cdot \ln\left(\frac{x_n(t^{(\tau)}+K) + \frac{\epsilon}{N}}{X_n^{\text{cap}} + \frac{\epsilon}{N}}\right) - x_n(t^{(\tau)}+K) \right] \right\} \quad (36a)$$

$$\text{sub. to: } (9f), (34), (35), \text{ for all } t \in [t^{(\tau)}, t^{(\tau)}+K]. \quad (36b)$$

In the analysis, we similarly change $\theta_n^{(\tau)}(t^{(\tau)})$ in (16) 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 [41]. Finally, in Theorem 7, we provide the competitive ratio of RLA for this case.

Theorem 7. *Given a look-ahead window of size $K \geq 1$, for the OCO problem with constraints (34) and (35), the competitive ratio of Regularization with Look-Ahead (RLA) is, (with $\eta \triangleq \max_n \eta_n$ and $\bar{B} \triangleq \max_{\{m,n,t\}} b_{mn}(t)$)*

$$\text{CR}^{RLA} = \begin{cases} 1 + \frac{3\eta(1+\epsilon\bar{B})\lceil r_{co}\rceil}{K+1}, & \text{if } \lceil r_{co}\rceil < K+1; \\ 1 + 2\eta(1+\epsilon\bar{B}), & \text{if } \lceil r_{co}\rceil \geq K+1. \end{cases} \quad (37)$$

VII. Numerical Results

A. Data and Settings

We generate a synthetic setting using demand from the Google cluster-usage traces [42]. We focus on 100 machines in the trace with the lowest *machine-ids*. Each machine corresponds to a decision variable in our OCO problem (and thus $N = 100$). We then generate the synthetic constraints as follows. Suppose the m-th lowest *machine-ids* is i. Then, the m-th constraint corresponds to a set $S_m(t)$ in (1) that contains all machines with *machine-ids* in $[i, 3i]$. Thus, such a constraint models the situation where any machine in $S_m(t)$ can be used to meet a certain aggregate demand to the group.

We consider one week of the trace, and sample the values of CPU usages in the *instance-event* table every 1 hour to generate the synthetic demand as follows. First, for OCO problem with fractional covering constraints (1), the constraint corresponding to $S_m(t)$ is present if there exists any CPU usage to the machines in $S_m(t)$. Second, for OCO problems with demand-supply balance constraint (34) and capacity constraint (35), the demand $a_m(t)$ for the constraint corresponds to $S_m(t)$ is the sum of CPU usage over all machines in $S_m(t)$, multiplied by 1000 and rounded to the nearest integers (so that $a_m(t)$ is an integer). We take the value of $b_{mn}(t)$ as the maximum CPU speed on the n-th machine in $S_m(t)$. We take the capacity X_n^{cap} as the value of *capacity* of the n-th machine in the *machine-event* table.

Finally, the service-cost coefficient $c_n(t)$ is randomly generated in $[1, 10]$. For Fig. 2a and Fig. 3a, we fix $K = 10$ hours and vary r_{co}. To simulate the setting with each value of r_{co}, we generate the switching-cost coefficient w_n randomly in $[0.7r_{\text{co}}, r_{\text{co}}]$. For Fig. 2b and Fig. 3b, we fix $r_{\text{co}} = 15$ and vary K. Correspondingly, we generate the switching-cost coefficient w_n randomly in $[5, 15]$, which produces $r_{\text{co}} = 15$.

Similar to the notation of the competitive ratio, we use ECR^{RLA}, ECR^{AFHC} and ECR^{REG} to denote the "empirical competitive ratios" (ECRs) of RLA, AFHC and the regularization method (REG), respectively.

B. Evaluation Results

In Fig. 2a and Fig. 2b, we compare the empirical competitive ratios (ECRs) of RLA, AFHC and REG for the OCO problem with fractional covering constraints (1). Fig. 2a shows that, as the coefficient ratio increases, the ECR of AFHC increases to be very large. In contrast, the ECRs of RLA and the regularization method remain at a low value even for large coefficient ratio r_{co}. In particular, the ECR of RLA is 1.891 even when $r_{\text{co}} = 400$. Furthermore, Fig. 2b shows that, as the look-ahead window size K increases, the ECRs of RLA and AFHC decrease quickly to a value close to 1. In particular, when $K = 50$, the ECR of RLA is about 1.032, which is much smaller than the ECR of the regularization method.

In Fig. 3a and Fig. 3b, we compare the empirical competitive ratios (ECRs) of RLA, AFHC and REG for the more general OCO problem in Sec. VI. The conclusions are similar to that from Fig. 2a and Fig. 2b: the ECRs of RLA not only

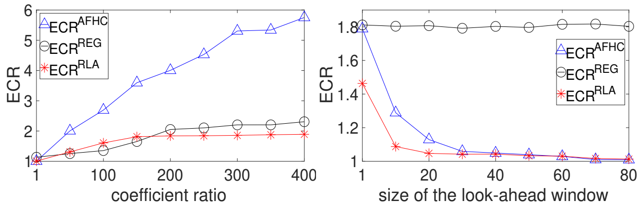

(a) ECR vs. coefficient ratio. (b) ECR vs. look-ahead.

Fig. 2: Compare the ECRs for OCO introduced in Sec. II-A.

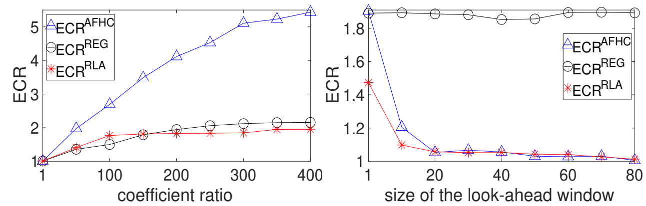

(a) ECR vs. coefficient ratio. (b) ECR vs. look-ahead.

Fig. 3: Compare the ECRs for OCO introduced in Sec. VI.

remain at a low value even for large r_{co} (Fig. 3a), but also decrease quickly to a value close to 1 as K increases (Fig. 3b).

VIII. Conclusion

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 [1] and the regularization method [21]. To prove the competitive ratio of RLA, we extend the online primal-dual method analysis [23] 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 additional experiment results (not reported), we observe that the actual competitive ratio of RLA is only a constant factor away from the lower bound, independent of the problem size. Thus, we will study ways to tighten the competitive ratio of RLA. Second, we have not allowed constraints of the form that the sum of some decision variables is upper-bounded. We note that the regularization method in [21] has a similar limitation. We will study how to generalize our results in this direction.

Acknowledgements: This work has been partially supported by the National Science Foundation (through grants CNS-1717493 and CNS-1703014) and the Ripple Faculty Fellowship.

Appendix A: 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, each of length $K + 1$.

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 one constraint for every episode. In the first episode, the constraint is $\sum_{n=1}^{N} x_n(t) \geq 1$, i.e., $S_1(t) = [1, N]$, for all $t \in [1, K+1]$.

The constraint in the second episode is based on the decision $\vec{X}(1)$. (Note that the decision maker must choose $\vec{X}^\pi(1)$ 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 $\sum_{n=1}^{N/2} x_n(t) \geq 1$ in the second episode. (ii) Otherwise, the adversary chooses $S_1(t) = [\frac{N}{2}+1, N]$ and $\sum_{n=N/2+1}^{N} x_n(t) \geq 1$ in the second episode.

In a similar way, the constraint in the i-th episode ($i \geq 2$) 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 $\tilde{n} \in \{1, ..., N\}$.

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 by

$$\text{Cost}^{\text{OPT}}(1:\mathcal{T}) \leq w + c\left((K+1)\alpha + 1\right). \quad (38)$$

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 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 by,

$$\text{Cost}^\pi(1:\mathcal{T}) \geq c((K+1)\alpha + 1) + w + \frac{\alpha w}{2}. \quad (39)$$

The result then follows by dividing the right-hand-side of (39) by the right-hand-side of (38).

REFERENCES

[1] 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.

[2] M. Lin, A. Wierman, L. L. H. 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.

[3] 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.

[4] M. Shi, X. Lin, S. Fahmy, and D.-H. Shin, "Competitive online convex optimization with switching costs and ramp constraints," in *IEEE Conference on Computer Communications (INFOCOM)*, 2018, pp. 1835–1843.

[5] 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.

[6] 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.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

[11] 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.

[12] 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.

[13] 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.

[14] 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.

[15] 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.

[16] 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.

[17] 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.

[18] 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.

[19] S. Shalev-Shwartz, "Online learning and online convex optimization," *Foundations and Trends® in Machine Learning*, vol. 4, no. 2, pp. 107–194, 2012.

[20] N. Buchbinder and J. Naor, "Online primal-dual algorithms for covering and packing problems," in *European Symposium on Algorithms*. Springer, 2005, pp. 689–701.

[21] 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.

[22] A. Borodin and R. El-Yaniv, *Online Computation and Competitive Analysis*. Cambridge University Press, 2005.

[23] 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.

[24] 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.

[25] 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.

[26] J. B. Rawlings and D. Q. Mayne, *Model Predictive Control: Theory and Design*. Nob Hill Pub. Madison, Wisconsin, 2009.

[27] E. F. Camacho and C. B. Alba, *Model Predictive Control*. Springer Science & Business Media, 2013.

[28] 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.

[29] 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.

[30] M. Shi, X. Lin and L. Jiao, "Corrigendum to 'On the Value of Look-Ahead in Competitive Online Convex Optimization'," 2020. Available at https://engineering.purdue.edu/%7elinx/paper/pomacs-correction20.html.

[31] 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.

[32] 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.

[33] 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.

[34] 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.

[35] 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.

[36] 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.

[37] 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.

[38] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.

[39] 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.

[40] 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.

[41] 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.

[42] A. Verma, L. Pedrosa, M. Korupolu, D. Oppenheimer, E. Tune, and J. Wilkes, "Large-scale cluster management at google with borg," in *Proceedings of the Tenth European Conference on Computer Systems*, 2015, pp. 1–17.