
Projection-Free Online Convex Optimization with Time-Varying Constraints

Dan Garber¹ Ben Kretzu¹

Abstract

We consider the setting of online convex optimization with adversarial time-varying constraints in which actions must be feasible w.r.t. a fixed constraint set, and are also required on average to approximately satisfy additional time-varying constraints. Motivated by scenarios in which the fixed feasible set (hard constraint) is difficult to project on, we consider projection-free algorithms that access this set only through a linear optimization oracle (LOO). We present an algorithm that, on a sequence of length T and using overall T calls to the LOO, guarantees $\tilde{O}(T^{3/4})$ regret w.r.t. the losses and $O(T^{7/8})$ constraints violation (ignoring all quantities except for T). In particular, these bounds hold w.r.t. any interval of the sequence. This algorithm however also requires access to an oracle for minimizing a strongly convex nonsmooth function over a Euclidean ball. We present a more efficient algorithm that does not require the latter optimization oracle but only first-order access to the time-varying constraints, and achieves similar bounds w.r.t. the entire sequence. We extend the latter to the setting of bandit feedback and obtain similar bounds (as a function of T) in expectation.

1. Introduction

We consider a particular setting of the well studied paradigm for sequential prediction known as online convex optimization (OCO) (Hazan, 2019; Shalev-Shwartz et al., 2012). In (standard) OCO, a decision maker is required to iteratively choose an action — some point \mathbf{x}_t on each round t (the total number of rounds T is finite and assumed for simplicity to be known in advanced), which must belong to some fixed

¹Technion - Israel Institute of Technology, Haifa, Israel. Correspondence to: Ben Kretzu <benkretzu@campus.technion.ac.il>, Dan Garber <dangar@technion.ac.il>.

(throughout all rounds) feasible convex set $\mathcal{K} \subset \mathbb{R}^n$ ¹, which we will also assume to be compact (as is often standard). After choosing \mathbf{x}_t , a scalar loss $f_t(\mathbf{x}_t)$ is incurred, where $f_t : \mathbb{R}^n \rightarrow \mathbb{R}$ is some arbitrary convex loss function (assumed for sake of analysis to be chosen adversarially). The standard goal is, on the course of the T rounds, to choose feasible actions $\mathbf{x}_1, \dots, \mathbf{x}_T$ such that the regret, given by the difference $\sum_{t=1}^T f_t(\mathbf{x}_t) - \min_{\mathbf{x} \in \mathcal{K}} \sum_{t=1}^T f_t(\mathbf{x})$, grows (as a function of T) only at a sublinear rate (the slower the better).

In the particular setting considered in this work — *OCO with time-varying constraints*, we assume that besides the *hard and fixed constraint* given by the set \mathcal{K} to which all played points must belong, there are additional *soft and time-varying constraints* given by convex functions $g_1, \dots, g_T \in \mathbb{R}^n \rightarrow \mathbb{R}$, where g_t is revealed at the end of round t and encodes the constraint $g_t(\mathbf{x}) \leq 0$. In this setting the standard goal is two folded: I. to guarantee sublinear regret w.r.t. the loss functions f_1, \dots, f_T against the best action in hindsight in the intersection of the hard constraint and all soft constraints, i.e., to guarantee that

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - \min_{\mathbf{x} \in \mathcal{K}: \mathbf{g}_t(\mathbf{x}) \leq 0 \forall t \in [T]} \sum_{t=1}^T f_t(\mathbf{x}) = o(T), \quad (1)$$

and II. to guarantee that the cumulative violation of the soft constraints is also sublinear in T , i.e., that

$$\sum_{t=1}^T \{g_t^+(\mathbf{x}_t) := \max\{g_t(\mathbf{x}_t), 0\}\} = o(T). \quad (2)$$

Here, $g_t^+(\mathbf{x}) := \max\{g_t(\mathbf{x}), 0\}$ is introduced to prevent a natural undesired phenomenon in which strongly satisfying some constraints can compensate strongly violating others.

This setting was recently studied in (Yi et al., 2022; Guo et al., 2022; Castiglioni et al., 2022; Neely & Yu, 2017; Cao & Liu, 2018). The state-of-the-art bounds with full-information (i.e., $f_t(\cdot), g_t(\cdot)$ become known to the algorithm at the end of each round t) for this setting are $O(\sqrt{T})$ for the regret (as given in (1)) and $O(T^{3/4})$ for the constraint violation (as given in (2)), see for instance (Guo et al., 2022).

However, previous works require as a sub-routine to solve on each iteration a convex optimization problem that is at

¹for ease of presentation we consider the underlying space to be \mathbb{R}^n , however any finite-dimensional Euclidean space is suitable

least as difficult as computing a Euclidean projection onto the feasible set \mathcal{K} . In high-dimensional settings and when the set \mathcal{K} admits a non-trivial structure, this highly limits the applicability of the proposed methods. Motivated by this observation, in this work, and to the best of our knowledge for the first time, we consider so-called *projection-free algorithms* for OCO with time-varying constraints. Concretely, motivated by vast work on projection-free methods for OCO in recent years, e.g., (Hazan & Kale, 2012; Garber & Kretzu, 2022; Hazan & Minasyan, 2020; Garber & Kretzu, 2023; Mhammedi, 2021; 2022), we consider algorithms that only access the feasible set \mathcal{K} through a linear optimization oracle (that is an oracle that given a linear function, can find a minimizer of it over \mathcal{K}) and consider online algorithms that throughout the T rounds make at most T calls to this oracle.

Example I: online minimum-cost capacitated flow. Consider a fixed directed acyclic graph G with n nodes, m edges, source node s and target node e . A decision maker (DM) must route on each round t a unit flow from s to e , i.e., choose some point \mathbf{x}_t in the corresponding *unit flow polytope*² \mathcal{K} . The DM then incurs cost $f_t(\mathbf{x}_t)$. If each edge is associated with a linear cost, we may have $f_t(\mathbf{x}) = \mathbf{f}_t^\top \mathbf{x}$ for some $\mathbf{f}_t \in \mathbb{R}^m$. The DM also needs to respect time-varying edge capacity constraints given by $\mathbf{x}_t(i) \leq \mathbf{c}_t(i)$, $i = 1, \dots, m$, where \mathbf{c}_t has non-negative entries. Here, the constraint function on time t is simply $g_t(\mathbf{x}) = \max_{i \in [m]} \{\mathbf{x}(i) - \mathbf{c}_t(i)\}$. The unit flow polytope is difficult to project on, however linear optimization over it corresponds to a simple weighted shortest path computation, which takes linear time using dynamic programming.

Example II: online semidefinite optimization. Consider a sequence of T semidefinite optimization problems over the bounded positive semidefinite (PSD) cone $\mathcal{K} = \{\mathbf{X} \in \mathbb{S}^n \mid \mathbf{X} \succeq 0, \text{Tr}(\mathbf{X}) \leq \tau\}$, where \mathbb{S}^n denotes the space of $n \times n$ real symmetric matrices and $\tau > 0$. Each instance t in the sequence consist of a convex objective $f_t(\mathbf{X}) : \mathbb{S}^n \rightarrow \mathbb{R}$ and a set of m_t linear inequalities $\text{Tr}(\mathbf{A}_{t,i}^\top \mathbf{X}) \leq b_{t,i}$, $i \in \{1, \dots, m_t\}$. Here the fixed hard constraints are given by the convex bounded PSD cone, and the time-varying soft constraints can be encoded using $g_t(\mathbf{X}) = \max_{i \in [m_t]} \{\text{Tr}(\mathbf{A}_{t,i}^\top \mathbf{X}) - b_{t,i}\}$. Projecting onto the bounded PSD cone requires a full eigen-decomposition of a $n \times n$ symmetric matrix, which in practical implementations requires $O(n^3)$ runtime. Linear optimization over this set amounts to a single leading eigenvector computation whose runtime scales only with n^2 , and even faster when the input matrix is sparse, using fast iterative eigenvector methods.

We believe this setting is well suited to capture important online scenarios with changing constraints: at one extreme, some works as the ones mentioned above, enforce the hard

²this is also the convex-hull of all identifying vectors of paths from node s to node e in the graph G

constraint \mathcal{K} using costly computational procedures (projections) which can be prohibitive, as in the examples above. At the other extreme, for the sake of computational efficiency one can also model membership to \mathcal{K} as a long term soft constraint (this was the motivation of (Mahdavi et al., 2012) which pioneered the study of OCO with long term constraints), however this might result in a too crude approximation of real-world scenarios. For instance, in Example I above it makes sense that even if the capacity constraints could be violated, the decision maker still must choose a feasible flow on each round, or that in Example II the chosen matrix must indeed be PSD. We view our work as an appealing middle ground between these two extremes: distinguishing between hard and soft (varying) constraints, while insisting on tractable procedures for high-dimensional and complex domains.

1.1. Contributions (informally stated)

I. We give an algorithm that using no more than T calls to the LOO of \mathcal{K} throughout all T rounds, guarantees that on each time interval $[s, e]$, $1 \leq s < e \leq T$, the regret w.r.t. the set $\widehat{\mathcal{K}}_{s,e} = \{\mathbf{x} \in \mathcal{K} \mid g_t(\mathbf{x}) \leq 0 \forall s \leq t \leq e\}$ is upper-bounded by $\tilde{O}(T^{3/4})$, and the overall non-negative constraint violation on the interval $\sum_{t=s}^e g_t^+(\mathbf{x}_t)$ is upper-bounded by $O(T^{7/8})$ (ignoring all quantities except for T and n). Note this metric (bounds w.r.t. any interval of the sequence, also known as *adaptive regret* (Hazan & Seshadhri, 2009)) is substantially stronger than the ones in (1), (2), and in particular is interesting even when there is no point $\mathbf{x} \in \mathcal{K}$ that satisfies all constraints $g_t(\mathbf{x}) \leq 0, t \in [T]$. Aside from a LOO for the set \mathcal{K} and a first-order oracle for the observed loss functions, the algorithm also requires access to an oracle for minimizing a strongly convex nonsmooth function over a Euclidean ball (which can be efficiently implemented via standard methods for convex optimization and is independent of the set \mathcal{K}). See Section 3.

II. We give a more efficient algorithm than the latter which does not require an oracle for strongly convex nonsmooth minimization over Euclidean balls, and instead only requires first-order access to the constraints g_1, \dots, g_T . This algorithm provides similar bounds, but only w.r.t. to the entire sequence (i.e., bounds as in (1), (2)). See Section 4.

III. We extend the latter algorithm to the bandit setting in which only the scalar values $f_t(\mathbf{x}_t), g_t(\mathbf{x}_t)$ are observed on each round t , where \mathbf{x}_t is the played point. The algorithm guarantees $\tilde{O}(\sqrt{n}T^{3/4})$ expected regret and expected non-negative constraint violation $\sum_{t=1}^T g_t^+(\mathbf{x}_t) = \tilde{O}(n^{1/4}T^{7/8})$. See Section 5.

Importantly, our regret bounds w.r.t. the loss functions f_1, \dots, f_T , both in the full-information and bandit settings, match up to logarithmic terms the current state-of-the-art

bounds (in terms of the horizon T and dimension n) for projection-free OCO (i.e., without additional time-varying constraints) (Hazan & Kale, 2012; Garber & Kretzu, 2022; 2020), and are thus (up to log factors) the best one can currently expect to achieve in our more challenging setting of projection-free OCO with time-varying constraints.

2. Preliminaries

2.1. Assumptions and notation

As stated, we assume that the feasible set $\mathcal{K} \subset \mathbb{R}^n$, which encodes the hard constraints (must be satisfied on each round t), is convex and compact and we let R be such that $\mathcal{K} \subseteq RB$, where B denotes the origin-centered unit Euclidean ball. We also assume w.l.o.g. that $\mathbf{0} \in \mathcal{K}$. We assume all loss and constraint functions $f_1, g_1, \dots, f_T, g_T$ are convex over RB , and we let G_f and G_g denote upper-bounds on the ℓ_2 norm of the subgradients of each $f_t, t \in [T]$ and $g_t, t \in [T]$ over the ball RB , respectively.

Our online algorithms will consider the T prediction rounds in disjoint blocks (batches) of size K , where K is an integer parameter. We assume throughout for simplicity and w.l.o.g. that T/K is an integer. We also use the notation $m(t) := \lceil t/K \rceil$ to map from some $t \in [T]$ to the corresponding block index $m(t) \in [T/K]$.

2.2. Fast approximately-feasible projections via linear optimization

In order to construct our projection-free algorithms that rely only on solving linear optimization problems over the fixed convex set \mathcal{K} (as opposed to projection or even more complex steps) we rely on the technique recently introduced in (Garber & Kretzu, 2022; 2023) of fast computation of approximately-feasible projections using a LOO.

Definition 2.1 (Approximately-feasible Projection Oracle³). Given a convex and compact set $\mathcal{K} \subset \mathbb{R}^n$, and a tolerance $\epsilon > 0$, we say a function $\mathcal{O}_{AFP}(\mathbf{y}, \epsilon, \mathcal{K})$ is an approximately-feasible projection (AFP) oracle (for the set \mathcal{K} with parameter ϵ), if for any input point $\mathbf{y} \in \mathbb{R}^n$, it returns some $(\mathbf{x}, \tilde{\mathbf{y}}) \in \mathcal{K} \times \mathbb{R}^n$ such that i. for all $\mathbf{z} \in \mathcal{K}$, $\|\tilde{\mathbf{y}} - \mathbf{z}\| \leq \|\mathbf{y} - \mathbf{z}\|$, and ii. $\|\mathbf{x} - \tilde{\mathbf{y}}\|^2 \leq \epsilon$.

Algorithm 1 given below is taken from (Garber & Kretzu, 2023) and implements an AFP oracle for the set \mathcal{K} using only linear optimization steps over \mathcal{K} . It uses as a subroutine the well known Frank-Wolfe method (Algorithm 2). If the point-to-project is far from the set, the output of Algorithm 2 is used to construct a hyperplane separating the point from \mathcal{K} , which is then used to “pull the point closer”

³(Garber & Kretzu, 2023) considers arbitrary matrix-induced norms. Here we only consider the Euclidean case which is equivalent to setting $\mathbf{A} = \mathbf{I}$ in their definition.

to \mathcal{K} . Otherwise, if the point is already sufficiently close to \mathcal{K} (but not necessarily in \mathcal{K}), Algorithm 2 outputs a point in \mathcal{K} sufficiently close to it.

Algorithm 1 Approximately-Feasible Projection Oracle via a Linear Optimization Oracle (see (Garber & Kretzu, 2023))

Data: parameters $\mathbf{y}_1 \in \mathbb{R}^n, \mathbf{x}_0 \in \mathcal{K}, \epsilon$
if $\|\mathbf{x}_0 - \mathbf{y}_1\|^2 \leq 3\epsilon$ **then**
 return $\mathbf{x} \leftarrow \mathbf{x}_0, \mathbf{y} \leftarrow \mathbf{y}_1$
end if
for $i = 1, 2, \dots$ **do**
 $\mathbf{x}_i \leftarrow$ Output of Algorithm 2 when called with tolerance ϵ , feasible point \mathbf{x}_{i-1} , and initial point \mathbf{y}_i
 if $\|\mathbf{x}_i - \mathbf{y}_i\|^2 > 3\epsilon$ **then**
 $\mathbf{y}_{i+1} = \mathbf{y}_i - \frac{2}{3}(\mathbf{y}_i - \mathbf{x}_i)$
 else
 return $\mathbf{x} \leftarrow \mathbf{x}_i, \mathbf{y} \leftarrow \mathbf{y}_i$
 end if
end for

Algorithm 2 Separating Hyperplane via Frank-Wolfe

Data: parameters $\mathbf{x}_1 \in \mathcal{K}, \mathbf{y} \in \mathbb{R}^n, \epsilon$
for $i = 1, 2, \dots$ **do**
 $\mathbf{v}_i \in \operatorname{argmin}_{\mathbf{x} \in \mathcal{K}} \{(\mathbf{x}_i - \mathbf{y})^\top \mathbf{x}\}$ {call to LOO of \mathcal{K} }
 if $(\mathbf{x}_i - \mathbf{y})^\top (\mathbf{x}_i - \mathbf{v}_i) \leq \epsilon$ or $\|\mathbf{x}_i - \mathbf{y}\|^2 \leq 3\epsilon$ **then**
 return $\tilde{\mathbf{x}} \leftarrow \mathbf{x}_i$
 end if
 $\sigma_i = \operatorname{argmin}_{\sigma \in [0,1]} \{\|\mathbf{y} - \mathbf{x}_i - \sigma(\mathbf{v}_i - \mathbf{x}_i)\|^2\}$
 $\mathbf{x}_{i+1} = \mathbf{x}_i + \sigma_i(\mathbf{v}_i - \mathbf{x}_i)$
end for

The following lemma is from (Garber & Kretzu, 2023).

Lemma 2.2. *Let $\mathcal{K} \subset \mathbb{R}^n$ be convex and compact such that $\mathbf{0} \in \mathcal{K} \subseteq RB$ for some $R > 0$. Algorithm 1 guarantees that it returns $(\mathbf{x}, \mathbf{y}) \in \mathcal{K} \times \mathbb{R}^n$ such that the following three conditions hold: I. $\forall \mathbf{z} \in \mathcal{K} : \|\mathbf{y} - \mathbf{z}\|^2 \leq \|\mathbf{y}_1 - \mathbf{z}\|^2$, II. $\|\mathbf{x} - \mathbf{y}\|^2 \leq 3\epsilon$, and III. $\|\mathbf{y}\| \leq \|\mathbf{y}_1\|$. The overall number of calls to the LOO of \mathcal{K} (in Algorithm 2) throughout the run of Algorithm 1 is upper-bounded by $\frac{27R^2}{\epsilon} \max \left\{ 2.25 \log \left(\frac{\|\mathbf{y}_1 - \mathbf{x}_0\|^2}{\epsilon} \right) + 1, 0 \right\}$.*

3. Drift-plus-Penalty-inspired Algorithm with Adaptive Regret Guarantees

Our first online algorithm, Algorithm 3, is inspired by the *drift plus penalty* approach used extensively in the literature on OCO with constraints, see for instance (Guo et al., 2022; Yi et al., 2021; 2022; Yu & Neely, 2020). The algorithm considers the T rounds in disjoint blocks of fixed size K , where K is a parameter of the algorithm. At the end of each block $m \in [T/K]$, the algorithm performs its two key steps: The first step is to find a minimizer of a *drift-plus-penalty*-style objective function (the function h_m defined in

the algorithm), involving the gradients (i.e., linearizations) of the loss functions and the constraints functions (without linearization) observed during the K iterations of the block, over a Euclidean ball enclosing the set \mathcal{K} ⁴. This amounts to minimizing a strongly convex nonsmooth function over a Euclidean ball (which in particular does not involve the set \mathcal{K} and does not require calls to the LOO of \mathcal{K}), and can be efficiently implemented via a variety of standard and efficient convex optimization algorithms. The second step is to feed to computed minimizer into the approximately-feasible projection oracle — Algorithm 1, to obtain the next state of the algorithm.

Algorithm 3 LOO-based Drift-plus-Penalty Method

Data: parameters $T, K, \epsilon, \delta, \alpha$
 $\mathbf{x}_1 = \tilde{\mathbf{y}}_1 \leftarrow$ arbitrary point in \mathcal{K}
for $m = 1, \dots, T/K$ **do**
 for $t = (m-1)K + 1, \dots, mK$ **do**
 Play \mathbf{x}_m and observe $f_t(\cdot), g_t(\cdot)$
 Set $\nabla_t \in \partial f_t(\tilde{\mathbf{y}}_m)$
 end for
 Denote $\mathcal{T}_m = \{(m-1)K + 1, \dots, mK\}$
 $\bar{\nabla}_m \leftarrow \frac{1}{K} \sum_{t \in \mathcal{T}_m} \nabla_t$
 $G_m^+(\mathbf{x}) \leftarrow \delta \sum_{t \in \mathcal{T}_m} g_t^+(\mathbf{x})$
 $\mathbf{y}_{m+1} \leftarrow \operatorname{argmin}_{\mathbf{x} \in R\mathcal{B}} \{h_m(\mathbf{x}) := \bar{\nabla}_m^\top (\mathbf{x} - \tilde{\mathbf{y}}_m) + G_m^+(\mathbf{x}) + \frac{\alpha}{2} \|\mathbf{x} - \tilde{\mathbf{y}}_m\|^2\}$
 $(\mathbf{x}_{m+1}, \tilde{\mathbf{y}}_{m+1}) \leftarrow \mathcal{O}_{AFP}(\mathbf{y}_{m+1}, \mathbf{x}_m, \epsilon, \mathcal{K})$ {Alg. 1}
end for

Theorem 3.1. Consider running Algorithm 3 with parameters $\delta = K = T^{\frac{1}{2}}, \alpha = \frac{G_f}{R} T^{\frac{1}{4}}, \epsilon = 61R^2 T^{-\frac{1}{2}} \log T$. Fix an interval $[s, e], 1 \leq s < e \leq T$. If $\hat{\mathcal{K}}_{s,e} := \{\mathbf{x} \in \mathcal{K} : g_t(\mathbf{x}) \leq 0, \forall t \in [s, e]\} \neq \emptyset$ then the regret on the interval is upper-bounded by

$$\sum_{t=s}^e f_t(\mathbf{x}_{m(t)}) - \min_{\mathbf{x} \in \hat{\mathcal{K}}_{s,e}} \sum_{t=s}^e f_t(\mathbf{x}) = O\left(RG_f T^{\frac{3}{4}} \sqrt{\log T}\right),$$

and the constraints violation on the interval is upper bounded by $\sum_{t=s}^e g_t^+(\mathbf{x}_{m(t)}) = O\left(RG_g T^{\frac{7}{8}}\right)$. The overall number of calls to the LOO of \mathcal{K} is upper-bounded by T .

Note Theorem 3.1 in particular provides meaningful bounds in the important scenario in which no single point in \mathcal{K} satisfies all the soft constraints given by g_1, \dots, g_T .

Relying only on a LOO (with an overall budget of T calls) comes with a price: the bounds in Theorem 3.1 could be

⁴We note that our *drift-plus-penalty*-style objectives are simpler versions of those used in previous works such as (Guo et al., 2022; Yi et al., 2021; 2022; Yu & Neely, 2020) since the penalty multiplying the constraint function G_m^+ is fixed throughout the run and not adaptive.

compared to the $O(\sqrt{T})$ regret and $O(T^{3/4})$ constraint violation achievable by methods with *unrestricted optimizations*, see for instance (Guo et al., 2022). This deterioration in rates is expected and is a known phenomena in projection-free methods for OCO (without soft constraints), e.g., (Hazan & Kale, 2012; Garber & Kretzu, 2022).

For the proof of the theorem we need the following lemma.

Lemma 3.2. Consider the (infeasible) sequence $\{\tilde{\mathbf{y}}_{m(t)}\}_{t \in [T]}$ produced by Algorithm 3. Fix an interval $[s, e], 1 \leq s < e \leq T$. If $\hat{\mathcal{K}}_{s,e} \neq \emptyset$ then the regret of $\{\tilde{\mathbf{y}}_{m(t)}\}_{t \in [T]}$ is upper-bounded by:

$$\sum_{t=s}^e f_t(\tilde{\mathbf{y}}_{m(t)}) - \min_{\mathbf{x} \in \hat{\mathcal{K}}_{s,e}} \sum_{t=s}^e f_t(\mathbf{x}) \leq 2KR(2G_f + R\alpha) + \frac{G_f^2}{2\alpha} T,$$

and for every $c > 0$, the constraints violation of the sequence on the interval is upper-bounded by:

$$\begin{aligned} \sum_{t=s}^e g_t^+(\tilde{\mathbf{y}}_{m(t)}) &\leq \frac{G_g^2}{2c} T + 2cR \left(\frac{G_f T}{\alpha} + RK \right) \\ &\quad + \frac{2\alpha R^2}{\delta} + \left(\frac{G_f}{\alpha} + 4R \right) \frac{G_f T}{2\delta K} + 4KRG_g. \end{aligned}$$

Proof. Fix some $m \in [T/K]$. Note $h_m(\cdot)$ is α -strongly convex. Since $\mathbf{y}_{m+1} \in R\mathcal{B}$ minimizes $h_m(\mathbf{x})$ over $R\mathcal{B}$, it follows that for any $\mathbf{x} \in R\mathcal{B}$, $h_m(\mathbf{x}) \geq h_m(\mathbf{y}_{m+1}) + \frac{\alpha}{2} \|\mathbf{x} - \mathbf{y}_{m+1}\|^2$. Thus, for every $\mathbf{x} \in \mathcal{K} \subseteq R\mathcal{B}$,

$$\begin{aligned} \bar{\nabla}_m^\top (\mathbf{x} - \tilde{\mathbf{y}}_m) + G_m^+(\mathbf{x}) + \frac{\alpha}{2} \left(\|\mathbf{x} - \tilde{\mathbf{y}}_m\|^2 - \|\mathbf{x} - \mathbf{y}_{m+1}\|^2 \right) \\ \geq \bar{\nabla}_m^\top (\mathbf{y}_{m+1} - \tilde{\mathbf{y}}_m) + G_m^+(\mathbf{y}_{m+1}) + \frac{\alpha}{2} \|\mathbf{y}_{m+1} - \tilde{\mathbf{y}}_m\|^2. \end{aligned}$$

Since $\tilde{\mathbf{y}}_{m+1}$ is the output of $\mathcal{O}_{AFP}(\mathbf{y}_{m+1}, \epsilon, \mathcal{K})$, it follows that $\|\mathbf{x} - \tilde{\mathbf{y}}_{m+1}\| \leq \|\mathbf{x} - \mathbf{y}_{m+1}\|$ (Lemma 2.2) and so,

$$\begin{aligned} \bar{\nabla}_m^\top (\mathbf{x} - \tilde{\mathbf{y}}_m) + G_m^+(\mathbf{x}) + \frac{\alpha}{2} \left(\|\mathbf{x} - \tilde{\mathbf{y}}_m\|^2 - \|\mathbf{x} - \tilde{\mathbf{y}}_{m+1}\|^2 \right) \\ \geq \bar{\nabla}_m^\top (\mathbf{y}_{m+1} - \tilde{\mathbf{y}}_m) + G_m^+(\mathbf{y}_{m+1}) + \frac{\alpha}{2} \|\mathbf{y}_{m+1} - \tilde{\mathbf{y}}_m\|^2. \end{aligned} \quad (3)$$

Note that

$$\begin{aligned} \bar{\nabla}_m^\top (\mathbf{y}_{m+1} - \tilde{\mathbf{y}}_m) + \frac{\alpha}{2} \|\mathbf{y}_{m+1} - \tilde{\mathbf{y}}_m\|^2 \\ = \left\| \frac{\bar{\nabla}_m}{\sqrt{2\alpha}} + \sqrt{\frac{\alpha}{2}} (\mathbf{y}_{m+1} - \tilde{\mathbf{y}}_m) \right\|^2 - \frac{\|\bar{\nabla}_m\|^2}{2\alpha} \geq -\frac{\|\bar{\nabla}_m\|^2}{2\alpha}. \end{aligned}$$

Combining the last two inequalities, we have for any $\mathbf{x} \in \mathcal{K}$,

$$\begin{aligned} \bar{\nabla}_m^\top (\mathbf{x} - \tilde{\mathbf{y}}_m) + G_m^+(\mathbf{x}) + \frac{\alpha}{2} \left(\|\mathbf{x} - \tilde{\mathbf{y}}_m\|^2 - \|\mathbf{x} - \tilde{\mathbf{y}}_{m+1}\|^2 \right) \\ \geq G_m^+(\mathbf{y}_{m+1}) - \frac{\|\bar{\nabla}_m\|^2}{2\alpha}, \end{aligned}$$

which by rearranging gives that for every $\mathbf{x} \in \mathcal{K}$,

$$\begin{aligned} \bar{\nabla}_m^\top (\tilde{\mathbf{y}}_m - \mathbf{x}) &\leq \frac{\alpha}{2} \left(\|\mathbf{x} - \tilde{\mathbf{y}}_m\|^2 - \|\mathbf{x} - \tilde{\mathbf{y}}_{m+1}\|^2 \right) \\ &\quad + \frac{\|\bar{\nabla}_m\|^2}{2\alpha} + G_m^+(\mathbf{x}) - G_m^+(\mathbf{y}_{m+1}). \end{aligned} \quad (4)$$

Let us now fix some interval $[s, e]$, $1 \leq s \leq e \leq T$. Let m_s and m_e denotes the smallest block index and the largest block index that are fully contained in the interval $[s, e]$, respectively. Note that $\|\bar{\nabla}_m\| \leq G_f$ for every m . Recall $G_m^+(\cdot) \geq 0$, and $G_m^+(\mathbf{x}) = 0$ for every $\mathbf{x} \in \hat{\mathcal{K}}_{s,e}$ and $m \in [m_s, m_e]$. Summing Eq.(4) over m from m_s to m_e , we have that for every $\mathbf{x} \in \hat{\mathcal{K}}_{s,e}$ it holds that,

$$\sum_{t=m_s}^{m_e} \sum_{t \in \mathcal{T}_m} \nabla_t^\top (\tilde{\mathbf{y}}_m - \mathbf{x}) \leq K \frac{\alpha}{2} \|\mathbf{x} - \tilde{\mathbf{y}}_{m_s}\|^2 + \frac{G_f^2}{2\alpha} T. \quad (5)$$

Since $f_t(\cdot)$ is convex, for every $\mathbf{x} \in \hat{\mathcal{K}}_{s,e}$ it holds that

$$\begin{aligned} \sum_{t=s}^e f_t(\tilde{\mathbf{y}}_{m(t)}) - f_t(\mathbf{x}) &= \sum_{t=s}^{(m_s-1)K} f_t(\tilde{\mathbf{y}}_{m_s-1}) - f_t(\mathbf{x}) \\ &\quad + \sum_{m=m_s}^{m_e} \sum_{t \in \mathcal{T}_m} f_t(\tilde{\mathbf{y}}_m) - f_t(\mathbf{x}) + \sum_{t=m_e K+1}^e f_t(\tilde{\mathbf{y}}_{m_e+1}) - f_t(\mathbf{x}) \\ &\leq \sum_{t=s}^{(m_s-1)K} \|\nabla_t\| \|\tilde{\mathbf{y}}_{m_s-1} - \mathbf{x}\| + \sum_{m=m_s}^{m_e} \sum_{t \in \mathcal{T}_m} \nabla_t^\top (\tilde{\mathbf{y}}_m - \mathbf{x}) \\ &\quad + \sum_{t=m_e K+1}^e \|\nabla_t\| \|\tilde{\mathbf{y}}_{m_e+1} - \mathbf{x}\|. \end{aligned}$$

From Lemma 2.2 it follows that $\tilde{\mathbf{y}}_m \in R\mathcal{B}$ for every m . plugging this, Eq.(5), and the fact that $\forall t \in [T] \|\nabla_t\| \leq G_f$ into the last inequality, we have that for every $\mathbf{x} \in \hat{\mathcal{K}}_{s,e}$,

$$\sum_{t=s}^e f_t(\tilde{\mathbf{y}}_{m(t)}) - f_t(\mathbf{x}) \leq 4KG_f R + 2KR^2\alpha + \frac{G_f^2}{2\alpha} T.$$

Now, we move on to bound the constraint violation on the interval $[s, e]$. We start with some preliminaries. Rearranging Eq.(3), we have for every $\mathbf{x} \in \mathcal{K}$ and $m \in [T/K]$ that,

$$\begin{aligned} \|\mathbf{y}_{m+1} - \tilde{\mathbf{y}}_m\|^2 &\leq \frac{2}{\alpha} \bar{\nabla}_m^\top (\mathbf{x} - \mathbf{y}_{m+1}) \\ &\quad + \frac{2}{\alpha} (G_m^+(\mathbf{x}) - G_m^+(\mathbf{y}_{m+1})) + \|\mathbf{x} - \tilde{\mathbf{y}}_m\|^2 - \|\mathbf{x} - \tilde{\mathbf{y}}_{m+1}\|^2 \\ &\leq \frac{4R}{\alpha} \|\bar{\nabla}_m\| + \frac{2}{\alpha} G_m^+(\mathbf{x}) + \|\mathbf{x} - \tilde{\mathbf{y}}_m\|^2 - \|\mathbf{x} - \tilde{\mathbf{y}}_{m+1}\|^2, \end{aligned}$$

where in the last inequality we have used the facts that $\mathbf{x} \in \mathcal{K} \subseteq R\mathcal{B}$ and $\mathbf{y}_{m+1} \in R\mathcal{B}$, and that $G_m^+(\cdot) \geq 0$. Summing the above inequality over m from m_s to m_e , and

recalling that $\tilde{\mathbf{y}}_m \in R\mathcal{B}$ and $\|\bar{\nabla}_m\| \leq G_f$ for all m , we have that for every $\mathbf{x} \in \mathcal{K}$ it holds that,

$$\sum_{m=m_s}^{m_e} \|\mathbf{y}_{m+1} - \tilde{\mathbf{y}}_m\|^2 \leq \frac{4G_f RT}{\alpha K} + \frac{2}{\alpha} \sum_{m=m_s}^{m_e} G_m^+(\mathbf{x}) + 4R^2. \quad (6)$$

Rearranging Eq.(4), and summing it over m from m_s to m_e , we have that for every $\mathbf{x} \in \mathcal{K}$ it holds that,

$$\begin{aligned} \sum_{m=m_s}^{m_e} G_m^+(\mathbf{y}_{m+1}) &\leq \sum_{m=m_s}^{m_e} \frac{\alpha}{2} \left(\|\mathbf{x} - \tilde{\mathbf{y}}_m\|^2 - \|\mathbf{x} - \tilde{\mathbf{y}}_{m+1}\|^2 \right) \\ &\quad + \sum_{m=m_s}^{m_e} \left(\frac{\|\bar{\nabla}_m\|^2}{2\alpha} + G_m^+(\mathbf{x}) + \bar{\nabla}_m^\top (\mathbf{x} - \tilde{\mathbf{y}}_m) \right) \\ &\leq 2\alpha R^2 + \frac{G_f^2 T}{2\alpha K} + \sum_{m=m_s}^{m_e} G_m^+(\mathbf{x}) + \frac{2G_f RT}{K}. \end{aligned}$$

Recalling $G_m^+(\mathbf{x}) = \delta \sum_{t \in \mathcal{T}_m} g_t^+(\mathbf{x})$, we have $\forall \mathbf{x} \in \mathcal{K}$,

$$\begin{aligned} \sum_{m=m_s}^{m_e} \sum_{t \in \mathcal{T}_m} g_t^+(\mathbf{y}_{m(t)+1}) &\leq \frac{2\alpha R^2}{\delta} + \frac{G_f^2 T}{2\alpha \delta K} \\ &\quad + \frac{1}{\delta} \sum_{m=m_s}^{m_e} G_m^+(\mathbf{x}) + \frac{2G_f RT}{\delta K}. \end{aligned} \quad (7)$$

Since $g_t(\cdot)$ is convex and G_g -Lipschitz over $R\mathcal{B}$ for every t , $g_t^+(\mathbf{x}) = \max\{0, g_t(\mathbf{x})\}$ is also convex and G_g -Lipschitz over $R\mathcal{B}$. Using the inequality $2ab \leq a^2 + b^2$ for every $a, b \in \mathbb{R}$, we have that for every $t \in [T]$ and $c > 0$,

$$\begin{aligned} g_t^+(\tilde{\mathbf{y}}_{m(t)}) - g_t^+(\mathbf{y}_{m(t)+1}) &\leq G_g \|\tilde{\mathbf{y}}_{m(t)} - \mathbf{y}_{m(t)+1}\| \\ &\leq \frac{G_g^2}{2c} + \frac{c}{2} \|\tilde{\mathbf{y}}_{m(t)} - \mathbf{y}_{m(t)+1}\|^2. \end{aligned}$$

Summing the last inequality over t , and combining it with Eq.(6), and Eq.(7), we have that for every $\mathbf{x} \in \hat{\mathcal{K}}_{s,e} \subseteq \mathcal{K}$,

$$\begin{aligned} \sum_{m=m_s}^{m_e} \sum_{t \in \mathcal{T}_m} g_t^+(\tilde{\mathbf{y}}_{m(t)}) &\leq \frac{G_g^2}{2c} T + \frac{cK}{2} \sum_{m=m_s}^{m_e} \|\tilde{\mathbf{y}}_m - \mathbf{y}_{m+1}\|^2 \\ &\quad + \sum_{m=m_s}^{m_e} \sum_{t \in \mathcal{T}_m} g_t^+(\mathbf{y}_{m(t)+1}) \\ &\leq \frac{G_g^2}{2c} T + \frac{2cG_f RT}{\alpha} + 2cR^2 K + \left(\frac{cK}{\alpha} + \frac{1}{\delta} \right) \sum_{m=m_s}^{m_e} G_m^+(\mathbf{x}) \\ &\quad + \frac{2\alpha R^2}{\delta} + \frac{G_f^2 T}{2\alpha \delta K} + \frac{2G_f RT}{\delta K} \\ &= \frac{G_g^2 T}{2c} + \frac{2cG_f RT}{\alpha} + 2cR^2 K + \frac{2\alpha R^2}{\delta} + \frac{G_f^2 T}{2\alpha \delta K} + \frac{2G_f RT}{\delta K}, \end{aligned}$$

where that last equality holds since $G_m^+(\mathbf{x}) = 0$ for every $\mathbf{x} \in \hat{\mathcal{K}}_{s,e}$ and $m \in [m_s, m_e]$.

The proof is concluded by noting that the number of first and last iterates in the interval $[s, e]$ that are not contained in the blocks m_s, \dots, m_e , i.e., the set $[s, e] \setminus \bigcup_{m=m_s}^{m_e} \mathcal{T}_m$, is at most $2K$. These add an additional term which is at most $2K \cdot 2RG_g$ to the RHS of the last inequality by using the fact that for each $t \in [s, e] \setminus \bigcup_{m=m_s}^{m_e} \mathcal{T}_m$ and $\mathbf{x} \in \widehat{\mathcal{K}}_{s,e}$, $g_t^+(\tilde{\mathbf{y}}_{m(t)}) \leq g_t^+(\mathbf{x}) + \|\tilde{\mathbf{y}}_{m(t)} - \mathbf{x}\| \cdot G_g \leq 2RG_g$. \square

Proof of Theorem 3.1. Fix some interval $[s, e]$, $1 \leq s \leq e \leq T$ such that $\widehat{\mathcal{K}}_{s,e} \neq \emptyset$. Applying Lemma 3.2 and Lemma 2.2, we have that for every $\mathbf{x} \in \widehat{\mathcal{K}}_{s,e}$ it holds that,

$$\begin{aligned} & \sum_{t=s}^e f_t(\mathbf{x}_{m(t)}) - f_t(\mathbf{x}) = \\ & \sum_{t=s}^e f_t(\mathbf{x}_{m(t)}) - f_t(\tilde{\mathbf{y}}_{m(t)}) + f_t(\tilde{\mathbf{y}}_{m(t)}) - f_t(\mathbf{x}) \leq \\ & 4KG_fR + 2KR^2\alpha + \frac{G_f^2}{2\alpha}T + \sum_{t=s}^e \|\nabla_t\| \|\mathbf{x}_{m(t)} - \tilde{\mathbf{y}}_{m(t)}\| \\ & \leq 4KG_fR + 2KR^2\alpha + \frac{G_f^2}{2\alpha}T + TG_f\sqrt{3\epsilon}, \end{aligned}$$

where the last inequality also uses the facts that $\tilde{\mathbf{y}}_m \in RB$ and that $\|\nabla_t\| \leq G_f$ for every $t \in [T]$.

plugging-in the values of K, α, ϵ stated in the theorem, we obtain the required bound on the regret.

Now we move on to bound the constraints violation. Since for every $t \in [T]$, $g_t(\cdot)$ is convex and G_g -Lipschitz over RB , it follows that $g_t^+(\mathbf{x})$ is also convex and G_g -Lipschitz over RB . Using Lemma 2.2 we have that for every $t \in [T]$,

$$g_t^+(\mathbf{x}_{m(t)}) - g_t^+(\tilde{\mathbf{y}}_{m(t)}) \leq G_g \|\mathbf{x}_{m(t)} - \tilde{\mathbf{y}}_{m(t)}\| \leq G_g\sqrt{3\epsilon}.$$

Summing over $t \in [s, e]$, and using Lemma 3.2, it holds that

$$\begin{aligned} \sum_{t=s}^e g_t^+(\mathbf{x}_{m(t)}) & \leq \frac{G_g^2}{2c}T + 2cR \left(\frac{G_f T}{\alpha} + RK \right) + \frac{2\alpha R^2}{\delta} \\ & + \left(\frac{G_f}{\alpha} + 4R \right) \frac{G_f T}{2\delta K} + 4KRG_g + G_g\sqrt{3\epsilon}T. \end{aligned}$$

Plugging-in $c = \frac{G_g}{12R}T^{\frac{1}{8}}$ and the values of $K, \alpha, \epsilon, \delta$ listed in the theorem, the overall constraints violation bound stated in the theorem is obtained.

Finally, we move on to bound the overall number of calls to the linear optimization oracle. Using Lemma 2.2, the

overall number of calls to linear optimization oracle is

$$\begin{aligned} N_{calls} & \leq \sum_{m=2}^{T/K} \frac{27R^2}{\epsilon} \left(2.25 \log \left(\frac{\|\tilde{\mathbf{y}}_{m+1} - \mathbf{x}_m\|^2}{\epsilon} \right) + 1 \right) \\ & \leq \sum_{m=2}^{T/K} \frac{27R^2}{\epsilon} \left(2.25 \log \left(\frac{4R^2}{\epsilon} \right) + 2.25 \log(e^{\frac{4}{9}}) \right) \\ & \leq \frac{27R^2 T}{\epsilon K} \left(2.25 \log \left(\frac{7R^2}{\epsilon} \right) \right) \leq \frac{61R^2 T}{\epsilon K} \log \left(\frac{7R^2}{\epsilon} \right). \end{aligned}$$

Plugging-in the values of ϵ, K listed in the theorem, we obtain that $N_{calls} \leq T$, as required. \square

4. Lagrangian-based Primal-Dual Algorithm

Our second algorithm, Algorithm 4, is based on an online first-order primal-dual approach which has also been studied extensively, see for instance (Mahdavi et al., 2012; Cao & Liu, 2018; Jenatton et al., 2016; Sun et al., 2017). This algorithm is more efficient than Algorithm 3 since it does not require any exact minimization of an auxiliary optimization problem involving the constraint functions $g_t, t \in [T]$, and instead only requires first-order access (i.e., to compute subgradients) of $f_t, g_t, t \in [T]$. On the downside, we will only have standard (i.e., non-adaptive) guarantees for this method that hold only w.r.t. the entire sequence of functions.

The algorithm performs online primal-dual (sub)gradient updates w.r.t. the dual-regularized Lagrangian functions

$$\mathcal{L}_t(\mathbf{x}, \lambda) := f_t(\mathbf{x}) + \lambda g_t^+(\mathbf{x}) - \frac{\delta\eta}{2}\lambda^2 \quad t \in [T]$$

and their block aggregations $\mathcal{L}_m(\mathbf{x}, \lambda) := \sum_{t=(m-1)K+1}^{mK} \mathcal{L}_t(\mathbf{x}, \lambda), m \in [T/K]$. Note that the vectors $\nabla_{m,\mathbf{x}}, \nabla_{m,\lambda}$ in Algorithm 4 are indeed the derivatives of $\mathcal{L}_m(\mathbf{x}, \lambda)$ w.r.t. \mathbf{x} and λ , evaluated at $(\mathbf{x}_m, \lambda_m)$.

Algorithm 4 LOO-based Primal-Dual Method

data: parameters $T, K, \epsilon, \delta, \eta$
 $\mathbf{x}_1 = \tilde{\mathbf{y}}_1 \leftarrow$ arbitrary point in $\mathcal{K}, \lambda_1 \leftarrow 0$
for $m = 1, \dots, T/K$ **do**
 for $t = (m-1)K + 1, \dots, mK$ **do**
 Play \mathbf{x}_m and observe $f_t(\cdot), g_t(\cdot)$
 Set $\nabla f_t \in \partial f_t(\mathbf{x}_m)$ and $\nabla g_t^+ \in \partial g_t^+(\mathbf{x}_m)$
 end for
 Denote $\mathcal{T}_m = \{(m-1)K + 1, \dots, mK\}$
 $\nabla_{m,\mathbf{x}} \leftarrow \sum_{t \in \mathcal{T}_m} \nabla f_t + \lambda_m \nabla g_t^+$
 $\nabla_{m,\lambda} \leftarrow \sum_{t \in \mathcal{T}_m} (g_t^+(\mathbf{x}_m) - \delta\eta\lambda_m)$
 $\mathbf{y}_{m+1} \leftarrow \Pi_{RB}[\tilde{\mathbf{y}}_m - \eta \nabla_{m,\mathbf{x}}]$
 $(\mathbf{x}_{m+1}, \tilde{\mathbf{y}}_{m+1}) \leftarrow \mathcal{O}_{AFP}(\mathbf{y}_{m+1}, \mathbf{x}_m, \epsilon, \mathcal{K})$ {Alg. 1}
 $\lambda_{m+1} \leftarrow \Pi_{\mathbb{R}_+}(\lambda_m + \eta \nabla_{m,\lambda})$
end for

Theorem 4.1. Suppose $\widehat{\mathcal{K}} = \{\mathbf{x} \in \mathcal{K} : g_t(\mathbf{x}) \leq 0 \forall t \in [T]\} \neq \emptyset$. For every T sufficiently large, setting $\delta = 32(G_f^2 + G_g R)T^{\frac{1}{2}}\sqrt{\log T}$, $\eta = T^{-\frac{3}{4}}$, $\epsilon = 61R^2T^{-\frac{1}{2}}\log T$, $K = T^{\frac{1}{2}}$ in Algorithm 4 guarantees that the regret is upper bounded by

$$\sum_{t=1}^T f_t(\mathbf{x}_{m(t)}) - \min_{\mathbf{x} \in \widehat{\mathcal{K}}} \sum_{t=1}^T f_t(\mathbf{x}) = O\left(R(G_f + G_g)T^{\frac{3}{4}}\sqrt{\log T} + (G_f^2 + (G_g^2 + 1)R^2)T^{3/4}\right),$$

the constraints violation is upper bounded by

$$\sum_{t=1}^T g_t^+(\mathbf{x}_{m(t)}) = O\left(\sqrt{G_f R(1 + G_g(G_g + R)\sqrt{\log T})}T^{\frac{7}{8}}\right),$$

and the overall number of calls to the LOO is at most T .

The proof of the theorem builds on the following lemma.

Lemma 4.2. Algorithm 4 guarantees that for every $\mathbf{x} \in \widehat{\mathcal{K}}$ and $\lambda \in \mathbb{R}_+$,

$$\begin{aligned} & \sum_{t=1}^T \mathcal{L}_t(\mathbf{x}_{m(t)}, \lambda) - \mathcal{L}_t(\mathbf{x}, \lambda_{m(t)}) \leq \\ & \frac{2R^2}{\eta} + \frac{\lambda^2}{2\eta} + G_f\sqrt{3\epsilon}T + G_f^2\eta KT + 4G_g^2R^2\eta KT \\ & + G_g\sqrt{3\epsilon}K \sum_{m=1}^{T/K} \lambda_m + (\delta^2\eta^3K^2 + G_g^2\eta K^2) \sum_{m=1}^{T/K} \lambda_m^2. \end{aligned}$$

Proof. Using Lemma 2.2 and the facts that $\mathbf{y}_{m+1} = \Pi_{RB}(\tilde{\mathbf{y}}_m - \eta\nabla_{m,\mathbf{x}})$ and $\mathcal{K} \subseteq RB$, for every $\mathbf{x} \in \mathcal{K}$ and $m \in [T/K]$ we have that,

$$\begin{aligned} \|\tilde{\mathbf{y}}_{m+1} - \mathbf{x}\|^2 & \leq \|\mathbf{y}_{m+1} - \mathbf{x}\|^2 \leq \|\tilde{\mathbf{y}}_m - \eta\nabla_{m,\mathbf{x}} - \mathbf{x}\|^2 \\ & = \|\tilde{\mathbf{y}}_m - \mathbf{x}\|^2 - 2\eta\nabla_{m,\mathbf{x}}^\top(\tilde{\mathbf{y}}_m - \mathbf{x}) + \eta^2\|\nabla_{m,\mathbf{x}}\|^2, \end{aligned}$$

and rearranging, we have that

$$\begin{aligned} \nabla_{m,\mathbf{x}}^\top(\tilde{\mathbf{y}}_m - \mathbf{x}) & \leq \frac{1}{2\eta} \left(\|\tilde{\mathbf{y}}_m - \mathbf{x}\|^2 - \|\tilde{\mathbf{y}}_{m+1} - \mathbf{x}\|^2 \right) \\ & \quad + \frac{\eta}{2} \|\nabla_{m,\mathbf{x}}\|^2. \end{aligned} \quad (8)$$

Since $\lambda_{m+1} = \Pi_{\mathbb{R}_+}(\lambda_m + \eta\nabla_{m,\lambda})$, for every $\lambda \in \mathbb{R}_+$ it holds that,

$$\begin{aligned} (\lambda_{m+1} - \lambda)^2 & \leq (\lambda_m + \eta\nabla_{m,\lambda} - \lambda)^2 \\ & = (\lambda_m - \lambda)^2 + 2\eta\nabla_{m,\lambda}(\lambda_m - \lambda) + \eta^2\nabla_{m,\lambda}^2, \end{aligned}$$

and rearranging, we have that

$$\begin{aligned} \nabla_{m,\lambda}(\lambda - \lambda_m) & \leq \frac{1}{2\eta} \left((\lambda_m - \lambda)^2 - (\lambda_{m+1} - \lambda)^2 \right) \\ & \quad + \frac{\eta}{2} \nabla_{m,\lambda}^2. \end{aligned} \quad (9)$$

Since $\mathcal{L}_t(\cdot, \lambda)$ is convex, it holds that

$$\mathcal{L}_t(\mathbf{x}_m, \lambda_m) - \mathcal{L}_t(\mathbf{x}, \lambda_m) \leq \nabla_{\mathbf{x}}\mathcal{L}_t(\mathbf{x}_m, \lambda_m)^\top(\mathbf{x}_m - \mathbf{x}),$$

and since $\mathcal{L}_t(\mathbf{x}_m, \cdot)$ is concave, it holds that

$$\mathcal{L}_t(\mathbf{x}_m, \lambda) - \mathcal{L}_t(\mathbf{x}_m, \lambda_m) \leq \nabla_{\lambda}\mathcal{L}_t(\mathbf{x}_m, \lambda_m)(\lambda - \lambda_m).$$

Combining the last two inequalities we have that,

$$\begin{aligned} \mathcal{L}_t(\mathbf{x}_m, \lambda) - \mathcal{L}_t(\mathbf{x}, \lambda_m) & \leq \nabla_{\mathbf{x}}\mathcal{L}_t(\mathbf{x}_m, \lambda_m)^\top(\mathbf{x}_m - \mathbf{x}) \\ & \quad + \nabla_{\lambda}\mathcal{L}_t(\mathbf{x}_m, \lambda_m)(\lambda - \lambda_m). \end{aligned}$$

Summing the above inequality over $t \in \mathcal{T}_m$ and recalling that $\nabla_{m,\mathbf{x}} = \sum_{t \in \mathcal{T}_m} \nabla_{\mathbf{x}}\mathcal{L}_t(\mathbf{x}_m, \lambda_m)$ and $\nabla_{m,\lambda} = \sum_{t \in \mathcal{T}_m} \nabla_{\lambda}\mathcal{L}_t(\mathbf{x}_m, \lambda_m)$, we have that

$$\begin{aligned} & \sum_{t \in \mathcal{T}_m} \mathcal{L}_t(\mathbf{x}_m, \lambda) - \mathcal{L}_t(\mathbf{x}, \lambda_m) \leq \\ & \nabla_{m,\mathbf{x}}^\top(\mathbf{x}_m - \mathbf{x}) + \nabla_{m,\lambda}(\lambda - \lambda_m) = \\ & \nabla_{m,\mathbf{x}}^\top(\mathbf{x}_m - \tilde{\mathbf{y}}_m) + \nabla_{m,\mathbf{x}}^\top(\tilde{\mathbf{y}}_m - \mathbf{x}) + \nabla_{m,\lambda}(\lambda - \lambda_m). \end{aligned}$$

Combining the last inequality with Eq. (8) and (9), we have that for every $\mathbf{x} \in \mathcal{K}$, $\lambda \in \mathbb{R}_+$ and m , it holds that

$$\begin{aligned} & \sum_{t \in \mathcal{T}_m} \mathcal{L}_t(\mathbf{x}_m, \lambda) - \mathcal{L}_t(\mathbf{x}, \lambda_m) \leq \\ & \nabla_{m,\mathbf{x}}^\top(\mathbf{x}_m - \tilde{\mathbf{y}}_m) + \frac{1}{2\eta} \left(\|\tilde{\mathbf{y}}_m - \mathbf{x}\|^2 - \|\tilde{\mathbf{y}}_{m+1} - \mathbf{x}\|^2 \right) \\ & + \frac{\eta}{2} \|\nabla_{m,\mathbf{x}}\|^2 + \frac{1}{2\eta} \left((\lambda_m - \lambda)^2 - (\lambda_{m+1} - \lambda)^2 \right) + \frac{\eta}{2} \nabla_{m,\lambda}^2. \end{aligned}$$

Summing the above over m , we have for every $\mathbf{x} \in \mathcal{K}$ and $\lambda \in \mathbb{R}_+$ that

$$\begin{aligned} & \sum_{m=1}^{T/K} \sum_{t \in \mathcal{T}_m} \mathcal{L}_t(\mathbf{x}_m, \lambda) - \mathcal{L}_t(\mathbf{x}, \lambda_m) \leq \\ & \sum_{m=1}^{T/K} \|\nabla_{m,\mathbf{x}}\| \|\mathbf{x}_m - \tilde{\mathbf{y}}_m\| + \frac{\eta}{2} \sum_{m=1}^{T/K} \left(\|\nabla_{m,\mathbf{x}}\|^2 + \nabla_{m,\lambda}^2 \right) \\ & + \frac{\|\tilde{\mathbf{y}}_1 - \mathbf{x}\|^2 + (\lambda_1 - \lambda)^2}{2\eta}. \end{aligned}$$

Using Lemma 2.2, it holds that $\|\mathbf{x}_m - \tilde{\mathbf{y}}_m\|^2 \leq 3\epsilon$ for every m . Since $\mathbf{y}_1 \in \mathcal{K}$ and $\lambda_1 = 0$, we have for every $\mathbf{x} \in \mathcal{K}$ and $\lambda \in \mathbb{R}_+$ that

$$\begin{aligned} & \sum_{m=1}^{T/K} \sum_{t \in \mathcal{T}_m} \mathcal{L}_t(\mathbf{x}_m, \lambda) - \mathcal{L}_t(\mathbf{x}, \lambda_m) \leq \sqrt{3\epsilon} \sum_{m=1}^{T/K} \|\nabla_{m,\mathbf{x}}\| \\ & + \frac{4R^2 + \lambda^2}{2\eta} + \frac{\eta}{2} \sum_{m=1}^{T/K} \left(\|\nabla_{m,\mathbf{x}}\|^2 + \nabla_{m,\lambda}^2 \right). \end{aligned} \quad (10)$$

Since for every m

$$\|\nabla_{m,\mathbf{x}}\| = \left\| \sum_{t \in \mathcal{T}_m} \nabla f_t + \lambda_m \nabla g_t^+ \right\| \leq K(G_f + \lambda_m G_g),$$

and $\nabla_{m,\lambda} = \sum_{t \in \mathcal{T}_m} (g_t^+(\mathbf{x}_m) - \delta\eta\lambda_m)$, we obtain

$$\begin{aligned} & \sum_{t=1}^T \mathcal{L}_t(\mathbf{x}_{m(t)}, \lambda) - \mathcal{L}_t(\mathbf{x}, \lambda_{m(t)}) \leq \\ & G_g \sqrt{3\epsilon} K \sum_{m=1}^{T/K} \lambda_m + G_f \sqrt{3\epsilon} T + \frac{4R^2 + \lambda^2}{2\eta} \\ & + \frac{\eta K^2}{2} \sum_{m=1}^{T/K} (G_f + \lambda_m G_g)^2 + \frac{\eta}{2} \sum_{m=1}^{T/K} \left(\sum_{t \in \mathcal{T}_m} (g_t^+(\mathbf{x}_m) - \delta\eta\lambda_m) \right)^2 \\ & \leq \frac{2R^2}{\eta} + \frac{\lambda^2}{2\eta} + G_f \sqrt{3\epsilon} T + G_f^2 \eta K T \\ & + \eta \sum_{m=1}^{T/K} \left(\sum_{t \in \mathcal{T}_m} g_t^+(\mathbf{x}_m) \right)^2 + G_g \sqrt{3\epsilon} K \sum_{m=1}^{T/K} \lambda_m \\ & + (\delta^2 \eta^3 K^2 + G_g^2 \eta K^2) \sum_{m=1}^{T/K} \lambda_m^2, \end{aligned} \quad (11)$$

where in the last inequality we have used the inequality $(a+b)^2 \leq 2a^2 + 2b^2$ for every $a, b \in \mathbb{R}$.

Note that $g_t^+(\mathbf{x}) = 0$ for every $\mathbf{x} \in \widehat{\mathcal{K}}$. Since $g_t(\mathbf{x})$ is convex and G_g -Lipschitz, $g_t^+(\mathbf{x}) = \max\{0, g_t(\mathbf{x})\}$ is also convex and G_g -Lipschitz and thus, for every $\mathbf{x} \in \widehat{\mathcal{K}}$ it holds that $g_t^+(\mathbf{x}_m) \leq g_t^+(\mathbf{x}) + \nabla g_t^{+\top}(\mathbf{x}_m - \mathbf{x}) \leq 2G_g R$. Plugging this into the RHS of (11) yields the lemma. \square

Proof of Theorem 4.1. From Lemma 4.2, for every $\mathbf{x} \in \widehat{\mathcal{K}}$ and $\lambda \in \mathbb{R}_+$ it holds that,

$$\begin{aligned} & \sum_{t=1}^T (f_t(\mathbf{x}_{m(t)}) - f_t(\mathbf{x})) + \sum_{t=1}^T (\lambda g_t^+(\mathbf{x}_{m(t)}) - \lambda_{m(t)} g_t^+(\mathbf{x})) \\ & + \frac{\delta\eta K}{2} \sum_{m=1}^{T/K} \lambda_m^2 - \frac{\delta\eta}{2} T \lambda^2 \leq \\ & \frac{2R^2}{\eta} + \frac{\lambda^2}{2\eta} + G_f \sqrt{3\epsilon} T + (G_f^2 + 4G_g^2 R^2) \eta K T \\ & + G_g \sqrt{3\epsilon} K \sum_{m=1}^{T/K} \lambda_m + (\delta^2 \eta^2 + G_g^2) \eta K^2 \sum_{m=1}^{T/K} \lambda_m^2. \end{aligned}$$

Rearranging, we have that

$$\begin{aligned} & \sum_{t=1}^T (f_t(\mathbf{x}_{m(t)}) - f_t(\mathbf{x})) + \sum_{t=1}^T (\lambda g_t^+(\mathbf{x}_{m(t)}) - \lambda_{m(t)} g_t^+(\mathbf{x})) \\ & - \left(\frac{\delta\eta}{2} T + \frac{1}{2\eta} \right) \lambda^2 \leq \\ & \frac{2R^2}{\eta} + G_f \sqrt{3\epsilon} T + (G_f^2 + 4G_g^2 R^2) \eta K T + G_g \sqrt{3\epsilon} K \sum_{m=1}^{T/K} \lambda_m \\ & + \left(\delta^2 \eta^3 K^2 + G_g^2 \eta K^2 - \frac{\delta\eta K}{2} \right) \sum_{m=1}^{T/K} \lambda_m^2. \end{aligned}$$

$g_t^+(\mathbf{x}) = 0$ for every $\mathbf{x} \in \widehat{\mathcal{K}}$ and $t \in [T]$. Also, by the choice of parameters in the theorem, we have that for any T larger than some constant, $\delta^2 \eta^3 K^2 + \eta K^2 G_g^2 + G_g \sqrt{3\epsilon} K \leq \frac{\eta K \delta}{2}$. Thus, for every $\mathbf{x} \in \widehat{\mathcal{K}}$ and $\lambda \in \mathbb{R}_+$ we have that,

$$\begin{aligned} & \sum_{t=1}^T (f_t(\mathbf{x}_{m(t)}) - f_t(\mathbf{x})) + \lambda \sum_{t=1}^T g_t^+(\mathbf{x}_{m(t)}) \\ & - \left(\frac{\delta\eta}{2} T + \frac{1}{2\eta} \right) \lambda^2 \leq \frac{2R^2}{\eta} + G_f \sqrt{3\epsilon} T \\ & + (G_f^2 + 4G_g^2 R^2) \eta K T + G_g \sqrt{3\epsilon} K \sum_{m=1}^{T/K} (\lambda_m - \lambda_m^2). \end{aligned}$$

Since $z - z^2 \leq \frac{1}{4}$ for every $z \in \mathbb{R}$, it holds that

$$\begin{aligned} & \sum_{t=1}^T (f_t(\mathbf{x}_{m(t)}) - f_t(\mathbf{x})) + \lambda \sum_{t=1}^T g_t^+(\mathbf{x}_{m(t)}) \\ & - \left(\frac{\delta\eta}{2} T + \frac{1}{2\eta} \right) \lambda^2 \leq \\ & \frac{2R^2}{\eta} + (G_f^2 + 4G_g^2 R^2) \eta K T + \sqrt{3\epsilon} \left(G_f + \frac{G_g}{4} \right) T. \end{aligned}$$

Let $\mathbf{x}^* \in \operatorname{argmin}_{\mathbf{x} \in \widehat{\mathcal{K}}} \sum_{t=1}^T f_t(\mathbf{x})$. Note that $\lambda^* = \frac{\sum_{t=1}^T g_t^+(\mathbf{x}_{m(t)})}{\delta\eta T + \eta^{-1}}$ maximizes the term $\lambda \sum_{t=1}^T g_t^+(\mathbf{x}_{m(t)}) - \left(\frac{\delta\eta}{2} T + \frac{1}{2\eta} \right) \lambda^2$ in the last inequality. Plugging-in \mathbf{x}^* and λ^* into the last inequality, we have that

$$\begin{aligned} & \sum_{t=1}^T (f_t(\mathbf{x}_{m(t)}) - f_t(\mathbf{x}^*)) + \frac{1}{2} \frac{\left(\sum_{t=1}^T g_t^+(\mathbf{x}_{m(t)}) \right)^2}{\delta\eta T + \eta^{-1}} \leq \\ & \frac{2R^2}{\eta} + (G_f^2 + 4G_g^2 R^2) \eta K T + \sqrt{3\epsilon} \left(G_f + \frac{G_g}{4} \right) T. \end{aligned} \quad (12)$$

This yields the regret bound

$$\begin{aligned} & \sum_{t=1}^T f_t(\mathbf{x}_{m(t)}) - f_t(\mathbf{x}^*) \leq \\ & \frac{2R^2}{\eta} + (G_f^2 + 4G_g^2 R^2) \eta K T + \sqrt{3\epsilon} (G_f + G_g/4) T. \end{aligned}$$

Plugging-in the choice of parameters listed in the theorem yields the desired regret bound.

Since $f_t(\mathbf{x}^*) - f_t(\mathbf{x}_{m(t)}) \leq 2RG_f$ for all $t \in [T]$, from (12) we also obtain the constraints violation bound

$$\sum_{t=1}^T g_t^+(\mathbf{x}_{m(t)}) \leq \sqrt{\delta\eta T + \eta^{-1}} \left(\frac{4R^2}{\eta} + 2(G_f^2 + 4G_g^2 R^2)\eta KT + 4RG_f T + 2\sqrt{3}\epsilon(G_f + G_g/4)T \right)^{1/2}.$$

Plugging-in our choice of parameters, yields the bound.

Finally, the bound on the overall number of calls to the LOO of \mathcal{K} follows from exactly the same argument as in the proof of Theorem 3.1. In particular we use the same values of K, ϵ which go into the bound. \square

5. Bandit Feedback Setting

In the bandit setting, after each round $t \in [T]$, the algorithm does not get to observe the functions f_t, g_t (or their subgradients), but only their value at the point played on round t . Building on the standard approach pioneered in (Flaxman et al., 2005) (see also (Hazan, 2019; Garber & Kretzu, 2020; 2022)), we extend Algorithm 4 to this setting by replacing the functions f_t, g_t^+ with their μ -smoothed versions: $\hat{f}_t(\mathbf{x}) = \mathbb{E}_{\mathbf{u} \sim U(\mathcal{S}^n)}[f_t(\mathbf{x} + \mu\mathbf{u})]$, $\hat{g}_t^+(\mathbf{x}) = \mathbb{E}_{\mathbf{u} \sim U(\mathcal{S}^n)}[g_t^+(\mathbf{x} + \mu\mathbf{u})]$, where $U(\mathcal{S}^n)$ denotes the uniform distribution over the origin-centered unit sphere. By sampling uniformly around the point to play, unbiased estimators for the gradients of \hat{f}_t, \hat{g}_t^+ ⁵ could be constructed using only the values of the functions at the sampled point.

We make the additional standard assumptions: **I.** $\{f_t, g_t\}_{t \in [T]}$ are chosen *obliviously*, i.e., before the first round and are thus independent of any randomness introduced by the algorithm. **II.** there is $r > 0$ such that $r\mathcal{B} \subseteq \mathcal{K}$.

For a fixed smoothing parameter $\mu \in (0, r]$, we denote a shrinking of the set \mathcal{K} by $\mathcal{K}_{\mu/r} := (1 - \mu/r)\mathcal{K} = \{(1 - \mu/r)\mathbf{x} \mid \mathbf{x} \in \mathcal{K}\}$. It is well known that for all $\mathbf{x} \in \mathcal{K}_{\mu/r}$ and $\mathbf{y} \in r\mathcal{B}$ it holds that $\mathbf{x} + \mathbf{y} \in \mathcal{K}$.

The proof of the following theorem is given in the appendix.

Theorem 5.1. *Suppose $\hat{\mathcal{K}} = \{\mathbf{x} \in \mathcal{K} : g_t(\mathbf{x}) \leq 0 \forall t \in [T]\} \neq \emptyset$. For any T sufficiently large, setting*

$$\delta = 24\sqrt{183} \max\{G_f, G_g, M_f, M_g\}^2 \times \max\left\{\frac{n}{r}, \frac{n}{\sqrt{r}}, \sqrt{n}, \frac{n}{\sqrt{r}R}\right\} T^{\frac{1}{2}} \sqrt{\log(T)},$$

$\eta = \frac{R}{\max\{G_f, G_g, M_f, M_g\}} T^{-\frac{3}{4}}$, $\epsilon = 61R^2 T^{-\frac{1}{2}} \log T$, $K = T^{\frac{1}{2}}$, $\mu = \sqrt{nr} T^{-1/4}$ in Algorithm 5, guarantees that the

⁵ \hat{f}_t, \hat{g}_t^+ are always differentiable, even if f_t, g_t^+ are not

Algorithm 5 LOO-based Bandit Primal-Dual Method

Data: parameters $T, K, \epsilon, \delta, \eta$
 $\mathbf{x}_1 = \hat{\mathbf{y}}_1 \leftarrow$ arbitrary point in $\mathcal{K}_{\mu/r}$, $\lambda_1 \leftarrow 0$
for $m = 1, \dots, T/K$ **do**
 for $t = (m-1)K + 1, \dots, mK$ **do**
 $\mathbf{z}_t \leftarrow \mathbf{x}_m + \mu\mathbf{u}_t$ where $\mathbf{u}_t \sim U(\mathcal{S}^n)$ $\{\mathbf{u}_t$ sampled from uniform dist. over the unit sphere}
 Play \mathbf{z}_t and observe $f_t(\mathbf{z}_t), g_t(\mathbf{z}_t)$
 $\nabla_{t,\mathbf{x}} \leftarrow \frac{n}{\mu} (f_t(\mathbf{z}_t) + \lambda_m g_t^+(\mathbf{z}_t)) \mathbf{u}_t$
 $\nabla_{t,\lambda} \leftarrow g_t^+(\mathbf{z}_t) - \delta\eta$
 end for
 Denote $\mathcal{T}_m = \{(m-1)K + 1, \dots, mK\}$
 $\nabla_{m,\mathbf{x}} \leftarrow \sum_{t \in \mathcal{T}_m} \nabla_{t,\mathbf{x}}$ and $\nabla_{m,\lambda} \leftarrow \sum_{t \in \mathcal{T}_m} \nabla_{t,\lambda}$
 $\mathbf{y}_{m+1} \leftarrow \Pi_{r\mathcal{B}}(\hat{\mathbf{y}}_m - \eta \nabla_{m,\mathbf{x}})$
 $(\mathbf{x}_{m+1}, \hat{\mathbf{y}}_{m+1}) \leftarrow \mathcal{O}_{AFP}(\mathbf{y}_{m+1}, \mathbf{x}_m, \epsilon, \mathcal{K}_{\mu/r})$
 {Apply Algorithm 1 w.r.t. the set $\mathcal{K}_{\mu/r}$ }
 $\lambda_{m+1} \leftarrow \Pi_{\mathbb{R}_+}(\lambda_m + \eta \nabla_{m,\lambda})$
end for

expected regret is upper bounded by

$$\mathbb{E} \left[\sum_{t=1}^T f_t(\mathbf{z}_t) \right] - \min_{\mathbf{x} \in \hat{\mathcal{K}}} \sum_{t=1}^T f_t(\mathbf{x}) = O \left(\sqrt{n} \max\{G_f, G_g, M_f, M_g\} T^{3/4} \left(\max\{R/r, R, 1/\sqrt{r}\} + R(1/\sqrt{n} + 1/\sqrt{r})\sqrt{\log T} \right) \right) = O \left(\sqrt{n} T^{3/4} \sqrt{\log T} \right),$$

the expected constraints violation is upper bounded by

$$\mathbb{E} \left[\sum_{t=1}^T g_t^+(\mathbf{z}_t) \right] = O \left(\sqrt{RG_f \max\{G_f, G_g, M_f, M_g\}} \times \sqrt{\max\left\{\frac{R}{r}, \frac{R}{\sqrt{r}}, \frac{R}{\sqrt{n}}, \frac{1}{\sqrt{r}}, \frac{1}{R}\right\}} \times n^{1/4} T^{7/8} \log(T)^{1/4} \right) = O \left(n^{1/4} T^{7/8} \log(T)^{1/4} \right),$$

and the overall number of calls to the LOO is upper bounded by T .

Acknowledgements

Dan Garber is supported by the ISRAEL SCIENCE FOUNDATION (grant No. 2267/22).

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are potential societal conse-

quences of our work, none which we feel must be specifically highlighted here.

References

- Cao, X. and Liu, K. R. Online convex optimization with time-varying constraints and bandit feedback. *IEEE Transactions on automatic control*, 64(7):2665–2680, 2018.
- Castiglioni, M., Celli, A., Marchesi, A., Romano, G., and Gatti, N. A unifying framework for online optimization with long-term constraints. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 33589–33602. Curran Associates, Inc., 2022.
- Flaxman, A. D., Kalai, A. T., and McMahan, H. B. Online convex optimization in the bandit setting: gradient descent without a gradient. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pp. 385–394, 2005.
- Garber, D. and Kretzu, B. Improved regret bounds for projection-free bandit convex optimization. In *International Conference on Artificial Intelligence and Statistics*, pp. 2196–2206. PMLR, 2020.
- Garber, D. and Kretzu, B. New projection-free algorithms for online convex optimization with adaptive regret guarantees. In Loh, P.-L. and Raginsky, M. (eds.), *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pp. 2326–2359. PMLR, 02–05 Jul 2022.
- Garber, D. and Kretzu, B. Projection-free online exp-concave optimization. In Neu, G. and Rosasco, L. (eds.), *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pp. 1259–1284. PMLR, 12–15 Jul 2023.
- Guo, H., Liu, X., Wei, H., and Ying, L. Online convex optimization with hard constraints: Towards the best of two worlds and beyond. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 36426–36439. Curran Associates, Inc., 2022.
- Hazan, E. Introduction to online convex optimization. *arXiv preprint arXiv:1909.05207*, 2019.
- Hazan, E. and Minasyan, E. Faster projection-free online learning. In *Conference on Learning Theory*, pp. 1877–1893. PMLR, 2020.
- Hazan, E. and Seshadhri, C. Efficient learning algorithms for changing environments. In *Proceedings of the 26th annual international conference on machine learning*, pp. 393–400, 2009.
- Hazan, E. E. and Kale, S. Projection-free online learning. In *29th International Conference on Machine Learning, ICML 2012*, pp. 521–528, 2012.
- Jenatton, R., Huang, J., and Archambeau, C. Adaptive algorithms for online convex optimization with long-term constraints. In Balcan, M. F. and Weinberger, K. Q. (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 402–411, New York, New York, USA, 20–22 Jun 2016. PMLR.
- Mahdavi, M., Jin, R., and Yang, T. Trading regret for efficiency: online convex optimization with long term constraints. *The Journal of Machine Learning Research*, 13(1):2503–2528, 2012.
- Mhammedi, Z. Efficient projection-free online convex optimization with membership oracle. *arXiv preprint arXiv:2111.05818*, 2021.
- Mhammedi, Z. Exploiting the curvature of feasible sets for faster projection-free online learning, 2022. URL <https://arxiv.org/abs/2205.11470>.
- Neely, M. J. and Yu, H. Online convex optimization with time-varying constraints. *arXiv preprint arXiv:1702.04783*, 2017.
- Shalev-Shwartz, S. et al. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2012.
- Sun, W., Dey, D., and Kapoor, A. Safety-aware algorithms for adversarial contextual bandit. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 3280–3288. PMLR, 06–11 Aug 2017. URL <https://proceedings.mlr.press/v70/sun17a.html>.
- Yi, X., Li, X., Yang, T., Xie, L., Chai, T., and Johansson, K. Regret and cumulative constraint violation analysis for online convex optimization with long term constraints. In *International Conference on Machine Learning*, pp. 11998–12008. PMLR, 2021.
- Yi, X., Li, X., Yang, T., Xie, L., Chai, T., and Karl, H. Regret and cumulative constraint violation analysis for distributed online constrained convex optimization. *IEEE Transactions on Automatic Control*, 2022.
- Yu, H. and Neely, M. J. A low complexity algorithm with $O(\sqrt{T})$ regret and $O(1)$ constraint violations for online

convex optimization with long term constraints. *Journal of Machine Learning Research*, 21(1):1–24, 2020.
URL <http://jmlr.org/papers/v21/16-494.html>.

A. Proof of Theorem 5.1

Before we prove the theorem we shall require some preliminary results.

Throughout this section we use the following notation:

$$\begin{aligned}\mathcal{L}_t(\mathbf{x}, \lambda) &:= f_t(\mathbf{x}) + \lambda g_t^+(\mathbf{x}) - \frac{\delta\eta}{2}\lambda^2 \quad \forall t \in [T], \\ \widehat{\mathcal{L}}_t(\mathbf{x}, \lambda) &:= \mathbb{E}_{\mathbf{u} \sim U(\mathcal{S}^n)} [f_t(\mathbf{x} + \mu\mathbf{u}) + \lambda g_t^+(\mathbf{x} + \mu\mathbf{u})] - \frac{\delta\eta}{2}\lambda^2 \quad \forall t \in [T], \\ \mathcal{L}_m(\mathbf{x}, \lambda) &:= \sum_{t \in \mathcal{T}_m} \mathcal{L}_t(\mathbf{x}, \lambda) \quad \text{and} \quad \widehat{\mathcal{L}}_m(\mathbf{x}, \lambda) := \sum_{t \in \mathcal{T}_m} \widehat{\mathcal{L}}_t(\mathbf{x}, \lambda) \quad \forall m \in [T/K],\end{aligned}$$

where we recall that $\mathcal{T}_m = \{(m-1)K + 1, \dots, mK\}$.

In the following Lemma A.1 and Lemma A.2, the notation $\nabla_{m,\mathbf{x}}$ is as defined in Algorithm 5.

Lemma A.1. *For every $m \in [T/K]$ the following holds:*

$$\forall \mathbf{x} \in \mathcal{K}_{\mu/r}, \lambda \in \mathbb{R}_+ : \quad 0 \leq \widehat{\mathcal{L}}_m(\mathbf{x}, \lambda) - \mathcal{L}_m(\mathbf{x}, \lambda) \leq \mu K(G_f + \lambda G_g),$$

$$\nabla_{\mathbf{x}} \widehat{\mathcal{L}}_m(\mathbf{x}_m, \lambda_m) = \mathbb{E}_{\{\mathbf{u}_t\}_{t \in \mathcal{T}_m}} [\nabla_{m,\mathbf{x}}] \quad \text{and} \quad \nabla_{\lambda} \widehat{\mathcal{L}}_m(\mathbf{x}_m, \lambda_m) = \sum_{t \in \mathcal{T}_m} \mathbb{E}_{\mathbf{u} \sim U(\mathcal{S}^n)} [g_t^+(\mathbf{x} + \mu\mathbf{u})] - \delta\eta K\lambda.$$

Proof. The second part of the lemma is a well known result, see for instance (Flaxman et al., 2005) and (Hazan, 2019) (Lemma 6.7). As for the first part, it is also well known that

$$\forall \mathbf{x} \in \mathcal{K}_{\mu/r}, \lambda \in \mathbb{R}_+ : \quad |\widehat{\mathcal{L}}_m(\mathbf{x}, \lambda) - \mathcal{L}_m(\mathbf{x}, \lambda)| \leq \mu K(G_f + \lambda G_g),$$

see (Flaxman et al., 2005) and (Hazan, 2019) (Lemma 2.8). One side of this bound could be improved, as suggested by the lemma, by noticing that for any convex function $h : \mathcal{K} \rightarrow \mathbb{R}$ and any $\mathbf{x} \in \mathcal{K}_{\mu/r}$ we have that,

$$h(\mathbf{x}) - \mathbb{E}_{\mathbf{u} \sim U(\mathcal{S}^n)} [h(\mathbf{x} + \mu\mathbf{u})] = \mathbb{E}_{\mathbf{u} \sim U(\mathcal{S}^n)} [h(\mathbf{x}) - h(\mathbf{x} + \mu\mathbf{u})] \leq \mathbb{E}_{\mathbf{u} \sim U(\mathcal{S}^n)} [-\mu\mathbf{u}^\top \mathbf{g}_{\mathbf{x}}] = -\mu \mathbb{E}_{\mathbf{u} \sim U(\mathcal{S}^n)} [\mathbf{u}]^\top \mathbf{g}_{\mathbf{x}} = 0,$$

where $\mathbf{g}_{\mathbf{x}}$ is some subgradient in $\partial h(\mathbf{x})$. □

Lemma A.2 (Lemma 5 in (Garber & Kretzu, 2020)). *For every $m \in [T/K]$,*

$$\mathbb{E}[\|\nabla_{m,\mathbf{x}}\|^2] \leq \mathbb{E}[\|\nabla_{m,\mathbf{x}}\|^2] \leq K \left(\frac{n(M_f + \lambda_m M_g)}{\mu} \right)^2 + K^2(G_f + \lambda_m G_g)^2.$$

The following lemma is analogous to Lemma 4.2 in the full-information setting.

Lemma A.3. *Algorithm 5 guarantees that for every $\mathbf{x} \in \mathcal{K}_{\mu/r}$ and $\lambda \in \mathbb{R}_+$,*

$$\begin{aligned}\mathbb{E} \left[\sum_{t=1}^T \mathcal{L}_t(\mathbf{x}_{m(t)}, \lambda) - \mathcal{L}_t(\mathbf{x}, \lambda_{m(t)}) \right] &\leq \left(\frac{\sqrt{3\epsilon}\sqrt{K}nM_g}{\mu} + KG_g(\sqrt{3\epsilon} + \mu) \right) \sum_{m=1}^{T/K} \mathbb{E}[\lambda_m] \\ &\quad + \left(\frac{\eta K n^2 M_g^2}{\mu^2} + \eta K^2 G_g^2 + K^2 \delta^2 \eta^3 \right) \sum_{m=1}^{T/K} \mathbb{E}[\lambda_m^2] \\ &\quad + \frac{T}{K} \left(\frac{\sqrt{3\epsilon}\sqrt{K}nM_f}{\mu} + KG_f(\sqrt{3\epsilon} + \mu) \right. \\ &\quad \left. + \eta \frac{K n^2 M_f^2}{\mu^2} + \eta K^2 (G_f^2 + M_g^2) \right) + \frac{4R^2 + \lambda^2}{2\eta}.\end{aligned}$$

Proof. Fix some block index m in Algorithm 5. Using Lemma 2.2 and the facts that $\mathbf{y}_{m+1} = \Pi_{RB}(\tilde{\mathbf{y}}_m - \eta \nabla_{m,\mathbf{x}})$ and $\mathcal{K}_{\mu/r} \subseteq \mathcal{K} \subseteq RB$, we have that for every $\mathbf{x} \in \mathcal{K}_{\mu/r}$ it holds that

$$\|\tilde{\mathbf{y}}_{m+1} - \mathbf{x}\|^2 \leq \|\mathbf{y}_{m+1} - \mathbf{x}\|^2 \leq \|\tilde{\mathbf{y}}_m - \eta \nabla_{m,\mathbf{x}} - \mathbf{x}\|^2 = \|\tilde{\mathbf{y}}_m - \mathbf{x}\|^2 - 2\eta \nabla_{m,\mathbf{x}}^\top (\tilde{\mathbf{y}}_m - \mathbf{x}) + \eta^2 \|\nabla_{m,\mathbf{x}}\|^2,$$

and rearranging, we have that

$$\nabla_{m,\mathbf{x}}^\top (\tilde{\mathbf{y}}_m - \mathbf{x}) \leq \frac{1}{2\eta} \left(\|\tilde{\mathbf{y}}_m - \mathbf{x}\|^2 - \|\tilde{\mathbf{y}}_{m+1} - \mathbf{x}\|^2 \right) + \frac{\eta}{2} \|\nabla_{m,\mathbf{x}}\|^2. \quad (13)$$

Since $\lambda_{m+1} = \Pi_{\mathbb{R}_+}(\lambda_m + \eta \nabla_{m,\lambda})$, for every $\lambda \in \mathbb{R}_+$ it holds that

$$(\lambda_{m+1} - \lambda)^2 \leq (\lambda_m + \eta \nabla_{m,\lambda} - \lambda)^2 = (\lambda_m - \lambda)^2 + 2\eta \nabla_{m,\lambda} (\lambda_m - \lambda) + \eta^2 \nabla_{m,\lambda}^2,$$

and rearranging, we have that

$$\nabla_{m,\lambda} (\lambda - \lambda_m) \leq \frac{1}{2\eta} \left((\lambda_m - \lambda)^2 - (\lambda_{m+1} - \lambda)^2 \right) + \frac{\eta}{2} \nabla_{m,\lambda}^2. \quad (14)$$

Let us now fix some $t \in [T]$. Since $\hat{\mathcal{L}}_t(\cdot, \lambda)$ is convex, for all $\mathbf{x} \in \mathcal{K}_{\mu/r}$ it holds that,

$$\hat{\mathcal{L}}_t(\mathbf{x}_{m(t)}, \lambda_{m(t)}) - \hat{\mathcal{L}}_t(\mathbf{x}, \lambda_{m(t)}) \leq \nabla_{\mathbf{x}} \hat{\mathcal{L}}_t(\mathbf{x}_{m(t)}, \lambda_{m(t)})^\top (\mathbf{x}_{m(t)} - \mathbf{x}),$$

and since $\hat{\mathcal{L}}_t(\mathbf{x}_{m(t)}, \cdot)$ is concave, for all $\lambda \in \mathbb{R}_+$ it holds that,

$$\hat{\mathcal{L}}_t(\mathbf{x}_{m(t)}, \lambda) - \hat{\mathcal{L}}_t(\mathbf{x}_{m(t)}, \lambda_{m(t)}) \leq \nabla_{\lambda} \hat{\mathcal{L}}_t(\mathbf{x}_{m(t)}, \lambda_{m(t)}) (\lambda - \lambda_{m(t)}).$$

Combining the last two inequalities we have that for any $\mathbf{x} \in \mathcal{K}_{\mu/r}$ and $\lambda \in \mathbb{R}^+$ it holds that,

$$\hat{\mathcal{L}}_t(\mathbf{x}_{m(t)}, \lambda) - \hat{\mathcal{L}}_t(\mathbf{x}, \lambda_{m(t)}) \leq \nabla_{\mathbf{x}} \hat{\mathcal{L}}_t(\mathbf{x}_{m(t)}, \lambda_{m(t)})^\top (\mathbf{x}_{m(t)} - \mathbf{x}) + \nabla_{\lambda} \hat{\mathcal{L}}_t(\mathbf{x}_{m(t)}, \lambda_{m(t)}) (\lambda - \lambda_{m(t)}).$$

Let us introduce the notation $\mathbb{E}_m[\cdot] = \mathbb{E}_{\{\mathbf{u}_t\}_{t \in \mathcal{T}_m}}[\cdot]$, where we recall that $\mathcal{T}_m = \{(m-1)K + 1, \dots, mK\}$. That is, $\mathbb{E}_m[\cdot]$ denotes the expectation w.r.t. all randomness introduced by the random vectors $\{\mathbf{u}_t\}_{t \in \mathcal{T}_m}$ on some block m during the run of the algorithm. Let us also recall that according to Lemma A.1, $\mathbb{E}_m[\nabla_{m,\mathbf{x}}] = \nabla_{\mathbf{x}} \hat{\mathcal{L}}_m(\mathbf{x}_m, \lambda_m) = \sum_{t \in \mathcal{T}_m} \nabla_{\mathbf{x}} \hat{\mathcal{L}}_t(\mathbf{x}_m, \lambda_m)$ and $\mathbb{E}_m[\nabla_{m,\lambda}] = \nabla_{\lambda} \hat{\mathcal{L}}_m(\mathbf{x}_m, \lambda_m) = \sum_{t \in \mathcal{T}_m} \nabla_{\lambda} \hat{\mathcal{L}}_t(\mathbf{x}_m, \lambda_m)$.

Summing the above inequality over $t \in \mathcal{T}_m$, we have that for any $\mathbf{x} \in \mathcal{K}_{\mu/r}$ and $\lambda \in \mathbb{R}^+$ it holds that,

$$\begin{aligned} \sum_{t \in \mathcal{T}_m} \hat{\mathcal{L}}_t(\mathbf{x}_m, \lambda) - \hat{\mathcal{L}}_t(\mathbf{x}, \lambda_m) &\leq \mathbb{E}_m[\nabla_{m,\mathbf{x}}]^\top (\mathbf{x}_m - \mathbf{x}) + \mathbb{E}_m[\nabla_{m,\lambda}] (\lambda - \lambda_m) \\ &= \mathbb{E}_m[\nabla_{m,\mathbf{x}}]^\top (\mathbf{x}_m - \tilde{\mathbf{y}}_m) + \mathbb{E}_m[\nabla_{m,\mathbf{x}}]^\top (\tilde{\mathbf{y}}_m - \mathbf{x}) + \mathbb{E}_m[\nabla_{m,\lambda}] (\lambda - \lambda_m). \end{aligned}$$

Combining the last inequality with Eq. (13) and (14), and noting that $\mathbf{x}_m, \mathbf{y}_m, \lambda_m$ are independent of the random vectors $\{\mathbf{u}_t\}_{t \in \mathcal{T}_m}$, we have that for every $\mathbf{x} \in \mathcal{K}_{\mu/r}$ and $\lambda \in \mathbb{R}_+$ it holds that,

$$\begin{aligned} \sum_{t \in \mathcal{T}_m} \hat{\mathcal{L}}_t(\mathbf{x}_m, \lambda) - \hat{\mathcal{L}}_t(\mathbf{x}, \lambda_m) &\leq \mathbb{E}_m[\nabla_{m,\mathbf{x}}^\top (\mathbf{x}_m - \tilde{\mathbf{y}}_m)] \\ &\quad + \frac{1}{2\eta} \mathbb{E}_m \left[\|\tilde{\mathbf{y}}_m - \mathbf{x}\|^2 - \|\tilde{\mathbf{y}}_{m+1} - \mathbf{x}\|^2 \right] + \frac{\eta}{2} \mathbb{E}_m [\|\nabla_{m,\mathbf{x}}\|^2] \\ &\quad + \frac{1}{2\eta} \mathbb{E}_m \left[(\lambda_m - \lambda)^2 - (\lambda_{m+1} - \lambda)^2 \right] + \frac{\eta}{2} \mathbb{E}_m [\nabla_{m,\lambda}^2]. \end{aligned}$$

Summing the above over $m \in \{1, \dots, T/K\}$ and taking expectation w.r.t. all random variables $\mathbf{u}_1, \dots, \mathbf{u}_T$, we have for every $\mathbf{x} \in \mathcal{K}_{\mu/r}$ and $\lambda \in \mathbb{R}_+$ that,

$$\begin{aligned} \mathbb{E} \left[\sum_{m=1}^{T/K} \sum_{t \in \mathcal{T}_m} \hat{\mathcal{L}}_t(\mathbf{x}_m, \lambda) - \hat{\mathcal{L}}_t(\mathbf{x}, \lambda_m) \right] &\leq \sum_{m=1}^{T/K} \mathbb{E} [\|\nabla_{m,\mathbf{x}}\| \|\mathbf{x}_m - \tilde{\mathbf{y}}_m\|] \\ &\quad + \frac{\eta}{2} \sum_{m=1}^{T/K} \left(\mathbb{E} [\|\nabla_{m,\mathbf{x}}\|^2] + \mathbb{E} [\nabla_{m,\lambda}^2] \right) + \frac{\|\tilde{\mathbf{y}}_1 - \mathbf{x}\|^2 + (\lambda_1 - \lambda)^2}{2\eta}. \end{aligned}$$

Using Lemma 2.2 we have that $\|\mathbf{x}_m - \tilde{\mathbf{y}}_m\|^2 \leq 3\epsilon$ for every m . Also, since $\tilde{\mathbf{y}}_1 \in \mathcal{K}_{\mu/r} \subseteq R\mathcal{B}$ and $\lambda_1 = 0$, for every $\mathbf{x} \in \mathcal{K}_{\mu/r}$ and $\lambda \in \mathbb{R}_+$ we have that,

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T \widehat{\mathcal{L}}_t(\mathbf{x}_{m(t)}, \lambda) - \widehat{\mathcal{L}}_t(\mathbf{x}, \lambda_{m(t)}) \right] &\leq \sqrt{3\epsilon} \sum_{m=1}^{T/K} \mathbb{E}[\|\nabla_{m,\mathbf{x}}\|] + \frac{4R^2 + \lambda^2}{2\eta} \\ &\quad + \frac{\eta}{2} \sum_{m=1}^{T/K} \left(\mathbb{E}[\|\nabla_{m,\mathbf{x}}\|^2] + \mathbb{E}[\nabla_{m,\lambda}^2] \right). \end{aligned} \quad (15)$$

Using Lemma A.2 we have that for every $m \in [T/K]$,

$$\mathbb{E}_m[\|\nabla_{m,\mathbf{x}}\|^2] \leq \mathbb{E}_m[\|\nabla_{m,\mathbf{x}}\|^2] \leq K \left(\frac{n(M_f + \lambda_m M_g)}{\mu} \right)^2 + K^2(G_f + \lambda_m G_g)^2.$$

Recalling also that $\nabla_{m,\lambda} = \sum_{t \in \mathcal{T}_m} (g_t^+(\mathbf{z}_t) - \delta\eta\lambda_m)$, we obtain

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T \widehat{\mathcal{L}}_t(\mathbf{x}_{m(t)}, \lambda) - \widehat{\mathcal{L}}_t(\mathbf{x}, \lambda_{m(t)}) \right] &\leq \sqrt{3\epsilon} \sum_{m=1}^{T/K} \mathbb{E} \left[\sqrt{K \left(\frac{n(M_f + \lambda_m M_g)}{\mu} \right)^2 + K^2(G_f + \lambda_m G_g)^2} \right] \\ &\quad + \frac{\eta}{2} \sum_{m=1}^{T/K} \mathbb{E} \left[K \left(\frac{n(M_f + \lambda_m M_g)}{\mu} \right)^2 + K^2(G_f + \lambda_m G_g)^2 \right] \\ &\quad + \frac{\eta}{2} \sum_{m=1}^{T/K} \mathbb{E} \left[\left(\sum_{t \in \mathcal{T}_m} (g_t^+(\mathbf{z}_t) - \delta\eta\lambda_m) \right)^2 \right] + \frac{4R^2 + \lambda^2}{2\eta} \\ &\leq \sqrt{3\epsilon} \sum_{m=1}^{T/K} \mathbb{E} \left[\sqrt{K} \frac{n(M_f + \lambda_m M_g)}{\mu} + K(G_f + \lambda_m G_g) \right] \\ &\quad + \frac{\eta}{2} \sum_{m=1}^{T/K} \mathbb{E} \left[K \left(\frac{n(M_f + \lambda_m M_g)}{\mu} \right)^2 + K^2(G_f + \lambda_m G_g)^2 \right] \\ &\quad + \eta \sum_{m=1}^{T/K} \mathbb{E} \left[\left(\sum_{t \in \mathcal{T}_m} g_t^+(\mathbf{z}_t) \right)^2 + K^2\delta^2\eta^2\lambda_m^2 \right] + \frac{4R^2 + \lambda^2}{2\eta}, \end{aligned} \quad (16)$$

where in the last inequality we have used the facts that $(a+b)^2 \leq 2a^2 + 2b^2$ for any $a, b \in \mathbb{R}$, and $\sqrt{a^2 + b^2} \leq a + b$ for any $a, b \in \mathbb{R}_+$.

Rearranging the RHS in the above using the fact that $g_t^+(\mathbf{z}_t) \leq M_g$ for all $t \in [T]$, and also using again the inequality $(a+b)^2 \leq 2a^2 + 2b^2$ gives,

$$\begin{aligned} \text{RHS of (16)} &\leq \sqrt{3\epsilon} \left(\frac{\sqrt{K}nM_g}{\mu} + KG_g \right) \sum_{m=1}^{T/K} \mathbb{E}[\lambda_m] + \left(\frac{\eta K n^2 M_g^2}{\mu^2} + \eta K^2 G_g^2 + K^2 \delta^2 \eta^3 \right) \sum_{m=1}^{T/K} \mathbb{E}[\lambda_m^2] \\ &\quad + \frac{T}{K} \left(\frac{\sqrt{3\epsilon}\sqrt{K}nM_f}{\mu} + \sqrt{3\epsilon}KG_f + \eta \frac{Kn^2 M_f^2}{\mu^2} + \eta K^2 (G_f^2 + M_g^2) \right) + \frac{4R^2 + \lambda^2}{2\eta}. \end{aligned}$$

Finally, using Lemma A.1, for any $\mathbf{x} \in \mathcal{K}_{\mu/r}$ and $\lambda \in \mathbb{R}_+$ we have that,

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T \mathcal{L}_t(\mathbf{x}_{m(t)}, \lambda) - \mathcal{L}_t(\mathbf{x}, \lambda_{m(t)}) \right] &\leq \left(\frac{\sqrt{3\epsilon}\sqrt{K}nM_g}{\mu} + KG_g(\sqrt{3\epsilon} + \mu) \right) \sum_{m=1}^{T/K} \mathbb{E}[\lambda_m] \\ &\quad + \left(\frac{\eta Kn^2 M_g^2}{\mu^2} + \eta K^2 G_g^2 + K^2 \delta^2 \eta^3 \right) \sum_{m=1}^{T/K} \mathbb{E}[\lambda_m^2] \\ &\quad + \frac{T}{K} \left(\frac{\sqrt{3\epsilon}\sqrt{K}nM_f}{\mu} + KG_f(\sqrt{3\epsilon} + \mu) \right. \\ &\quad \left. + \eta \frac{Kn^2 M_f^2}{\mu^2} + \eta K^2 (G_f^2 + M_g^2) \right) + \frac{4R^2 + \lambda^2}{2\eta}. \end{aligned}$$

□

Proof of Theorem 5.1. From Lemma 4.2 we have for every $\mathbf{x} \in \mathcal{K}_{\mu/r}$ and $\lambda \in \mathbb{R}_+$ that,

$$\begin{aligned} &\mathbb{E} \left[\sum_{t=1}^T (f_t(\mathbf{x}_{m(t)}) - f_t(\mathbf{x})) + \sum_{t=1}^T (\lambda g_t^+(\mathbf{x}_{m(t)}) - \lambda_{m(t)} g_t^+(\mathbf{x})) + \frac{\delta\eta K}{2} \sum_{m=1}^{T/K} \lambda_m^2 - \frac{\delta\eta}{2} T \lambda^2 \right] \\ &\leq \left(\frac{\sqrt{3\epsilon}\sqrt{K}nM_g}{\mu} + KG_g(\sqrt{3\epsilon} + \mu) \right) \sum_{m=1}^{T/K} \mathbb{E}[\lambda_m] + \left(\frac{\eta Kn^2 M_g^2}{\mu^2} + \eta K^2 G_g^2 + K^2 \delta^2 \eta^3 \right) \sum_{m=1}^{T/K} \mathbb{E}[\lambda_m^2] \\ &\quad + \frac{T}{K} \left(\frac{\sqrt{3\epsilon}\sqrt{K}nM_f}{\mu} + \sqrt{3\epsilon}KG_f + \eta \frac{Kn^2 M_f^2}{\mu^2} + \eta K^2 (G_f^2 + M_g^2) + \mu KG_f \right) + \frac{4R^2 + \lambda^2}{2\eta}. \end{aligned}$$

Rearranging, we have that for every $\mathbf{x} \in \mathcal{K}_{\mu/r}$ and $\lambda \in \mathbb{R}_+$ it holds that,

$$\begin{aligned} &\mathbb{E} \left[\sum_{t=1}^T (f_t(\mathbf{x}_{m(t)}) - f_t(\mathbf{x})) + \sum_{t=1}^T \lambda g_t^+(\mathbf{x}_{m(t)}) \right] - \left(\frac{\delta\eta}{2} T + \frac{1}{2\eta} \right) \lambda^2 \\ &\leq \frac{2R^2}{\eta} + \left(\frac{\sqrt{3\epsilon}\sqrt{K}nM_g}{\mu} + KG_g(\sqrt{3\epsilon} + \mu) \right) \sum_{m=1}^{T/K} \mathbb{E}[\lambda_m] \\ &\quad + \left(\frac{\eta Kn^2 M_g^2}{\mu^2} + \eta K^2 G_g^2 + K^2 \delta^2 \eta^3 - \frac{\delta\eta K}{2} \right) \sum_{m=1}^{T/K} \mathbb{E}[\lambda_m^2] \\ &\quad + T \left(\frac{\sqrt{3\epsilon}nM_f}{\mu\sqrt{K}} + \sqrt{3\epsilon}G_f + \eta \frac{n^2 M_f^2}{\mu^2} + \eta K (G_f^2 + M_g^2) + \mu G_f \right) + \sum_{t=1}^T g_t^+(\mathbf{x}) \mathbb{E}[\lambda_{m(t)}]. \end{aligned}$$

Fix some $\mathbf{x}^* \in \widehat{\mathcal{K}}$ such that $\mathbf{x}^* \in \arg \min_{\mathbf{w} \in \widehat{\mathcal{K}}} \sum_{t=1}^T f_t(\mathbf{w})$ and from now on fix $\mathbf{x} = (1 - \mu/r)\mathbf{x}^* \in \mathcal{K}_{\mu/r}$. Note that for all $t \in [T]$, using the convexity of g_t^+ and the fact that $g_t^+(\mathbf{x}^*) = 0$ we have that,

$$g_t^+(\mathbf{x}) = g_t^+((1 - \mu/r)\mathbf{x}^* + (\mu/r)\mathbf{0}) \leq (1 - \mu/r)g_t^+(\mathbf{x}^*) + (\mu/r)g_t^+(\mathbf{0}) \leq \frac{\mu}{r}M_g.$$

This gives,

$$\begin{aligned}
 & \mathbb{E} \left[\sum_{t=1}^T (f_t(\mathbf{x}_{m(t)}) - f_t(\mathbf{x})) + \sum_{t=1}^T \lambda g_t^+(\mathbf{x}_{m(t)}) \right] - \left(\frac{\delta\eta}{2} T + \frac{1}{2\eta} \right) \lambda^2 \\
 & \leq \frac{2R^2}{\eta} + \left(\frac{\sqrt{3\epsilon}\sqrt{K}nM_g}{\mu} + KG_g(\sqrt{3\epsilon} + \mu) + \frac{\mu KM_g}{r} \right) \sum_{m=1}^{T/K} \mathbb{E}[\lambda_m] \\
 & \quad + \left(\frac{\eta K n^2 M_g^2}{\mu^2} + \eta K^2 G_g^2 + K^2 \delta^2 \eta^3 - \frac{\delta\eta K}{2} \right) \sum_{m=1}^{T/K} \mathbb{E}[\lambda_m^2] \\
 & \quad + T \left(\frac{\sqrt{3\epsilon}nM_f}{\mu\sqrt{K}} + \sqrt{3\epsilon}G_f + \eta \frac{n^2 M_f^2}{\mu^2} + \eta K(G_f^2 + M_g^2) + \mu G_f \right).
 \end{aligned}$$

Suppose throughout the rest of the proof that the following condition holds:

$$\begin{aligned}
 \frac{\delta\eta K}{2} & \geq \frac{\eta K n^2 M_g^2}{\mu^2} + \eta K^2 G_g^2 + K^2 \delta^2 \eta^3 + \frac{\sqrt{3\epsilon}\sqrt{K}nM_g}{\mu} + KG_g(\sqrt{3\epsilon} + \mu) + \frac{\mu KM_g}{r} \\
 & = \eta K \left(\frac{n^2 M_g^2}{\mu^2} + KG_g^2 + \eta^2 K \delta^2 \right) + \sqrt{3\epsilon}K \left(\frac{nM_g}{\mu} + \sqrt{K}G_g \right) + K\mu \left(G_g + \frac{M_g}{r} \right).
 \end{aligned}$$

In particular, for the above to hold it suffices that the following holds:

$$\delta \geq 12 \max \left\{ \frac{n^2 M_g^2}{\mu^2}, KG_g^2, K\delta^2 \eta^2, \frac{\sqrt{3\epsilon}nM_g}{\eta\mu\sqrt{K}}, \frac{G_g(\sqrt{3\epsilon} + \mu)}{\eta}, \frac{\mu M_g}{r\eta} \right\}. \quad (17)$$

Then we can write,

$$\begin{aligned}
 & \mathbb{E} \left[\sum_{t=1}^T (f_t(\mathbf{x}_{m(t)}) - f_t(\mathbf{x})) + \sum_{t=1}^T \lambda g_t^+(\mathbf{x}_{m(t)}) \right] - \left(\frac{\delta\eta}{2} T + \frac{1}{2\eta} \right) \lambda^2 \\
 & \leq \frac{2R^2}{\eta} + \left(\frac{\sqrt{3\epsilon}\sqrt{K}nM_g}{\mu} + KG_g(\sqrt{3\epsilon} + \mu) + \frac{\mu KM_g}{r} \right) \sum_{m=1}^{T/K} \mathbb{E}[\lambda_m - \lambda_m^2] \\
 & \quad + T \left(\frac{\sqrt{3\epsilon}nM_f}{\mu\sqrt{K}} + \sqrt{3\epsilon}G_f + \eta \frac{n^2 M_f^2}{\mu^2} + \eta K(G_f^2 + M_g^2) + \mu G_f \right).
 \end{aligned}$$

Since $z - z^2 \leq \frac{1}{4}$ for every $z \in \mathbb{R}$, it holds that

$$\begin{aligned}
 & \mathbb{E} \left[\sum_{t=1}^T (f_t(\mathbf{x}_{m(t)}) - f_t(\mathbf{x})) + \sum_{t=1}^T \lambda g_t^+(\mathbf{x}_{m(t)}) \right] - \left(\frac{\delta\eta}{2} T + \frac{1}{2\eta} \right) \lambda^2 \leq \frac{2R^2}{\eta} \\
 & \quad + \frac{T}{4} \left(\frac{\sqrt{3\epsilon}n(M_g + 4M_f)}{\mu\sqrt{K}} + G_g(\sqrt{3\epsilon} + \mu) + \frac{\mu M_g}{r} + 4\sqrt{3\epsilon}G_f + 4\eta \frac{n^2 M_f^2}{\mu^2} + 4\eta K(G_f^2 + M_g^2) + 4\mu G_f \right). \quad (18)
 \end{aligned}$$

Note that $\lambda^* = \frac{\sum_{t=1}^T \mathbb{E}[g_t^+(\mathbf{x}_{m(t)})]}{\delta\eta T + \eta^{-1}}$ maximizes the term $\lambda \sum_{t=1}^T \mathbb{E}[g_t^+(\mathbf{x}_{m(t)})] - \left(\frac{\delta\eta}{2} T + \frac{1}{2\eta} \right) \lambda^2$ in the above inequality as a function of λ . Plugging-in $\mathbf{x} = (1 - \mu/r)\mathbf{x}^*$ and λ^* into the last inequality, we have that

$$\mathbb{E} \left[\sum_{t=1}^T (f_t(\mathbf{x}_{m(t)}) - f_t((1 - \mu/r)\mathbf{x}^*)) \right] + \frac{1}{2} \frac{\left(\sum_{t=1}^T \mathbb{E}[g_t^+(\mathbf{x}_{m(t)})] \right)^2}{\delta\eta T + \eta^{-1}} \leq \text{RHS of (18)}. \quad (19)$$

Suppose $\eta = c_1 T^{-3/4}$, $K = T^{1/2}$, $\epsilon = 61R^2 T^{-1/2} \log T$, $\delta = c_2 T^{1/2} \sqrt{\log T}$ and $\mu = c_3 T^{-1/4}$ for some constants c_1, c_2, c_3 . This gives

$$\begin{aligned} \text{RHS of (18)} &= O\left(\frac{R^2 T^{3/4}}{c_1} + T\left(\frac{RT^{-1/4} \sqrt{\log T} n(M_g + M_f)}{c_3} + G_g\left(RT^{-1/4} \sqrt{\log T} + c_3 T^{-1/4}\right)\right.\right. \\ &\quad \left.\left. + \frac{c_3 T^{-1/4} M_g}{r} + G_f RT^{-1/4} \sqrt{\log T} + \frac{c_1 T^{-1/4} n^2 M_f^2}{c_3^2} + c_1 T^{-1/4} (G_f^2 + M_g^2) + c_3 G_f T^{-1/4}\right)\right) \\ &= O\left(T^{3/4} \sqrt{\log T} \left(\frac{nR(M_g + M_f)}{c_3} + R(G_g + G_f)\right)\right. \\ &\quad \left. + T^{3/4} \left(\frac{R^2}{c_1} + c_3 \left(G_f + G_g + \frac{M_g}{r}\right) + \frac{c_1 n^2 M_f^2}{c_3^2} + c_1 (G_f^2 + M_g^2)\right)\right). \end{aligned}$$

Setting $c_1 = \frac{R}{\sqrt{n} \max\{G_f, G_g, M_f, M_g\}}$, $c_3 = \sqrt{nr}$, we obtain

$$\begin{aligned} \text{RHS of (18)} &= O\left(T^{3/4} \sqrt{\log T} \left(\sqrt{\frac{n}{r}} R(M_g + M_f) + R(G_g + G_f)\right)\right. \\ &\quad \left. + T^{3/4} \left(\sqrt{n} \max\{R/r, R, 1/\sqrt{r}\} \max\{G_f, G_g, M_f, M_g\}\right)\right) \\ &= O\left(\sqrt{n} \max\{G_f, G_g, M_f, M_g\} T^{3/4} \left(\max\{R/r, R, 1/\sqrt{r}\} + R(1/\sqrt{n} + 1/\sqrt{r}) \sqrt{\log T}\right)\right). \end{aligned}$$

Note that for all $t \in [T]$, using the convexity of f_t , we have that

$$f_t(\mathbf{z}_t) = f_t(\mathbf{x}_{m(t)} + \mu \mathbf{u}_t) \leq f_t(\mathbf{x}_{m(t)}) + \mu G_f, \quad f_t((1 - \mu/r)\mathbf{x}^*) \leq f_t(\mathbf{x}^*) + \frac{\mu G_f}{r}.$$

Plugging these into (19) yields the regret bound:

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T f_t(\mathbf{z}_t) - f_t(\mathbf{x}^*) \right] &\leq \text{RHS of (18)} + T \mu G_f \left(1 + \frac{1}{r}\right) \\ &= O\left(\sqrt{n} \max\{G_f, G_g, M_f, M_g\} T^{3/4} \left(\max\{R/r, R, 1/\sqrt{r}\} + R(1/\sqrt{n} + 1/\sqrt{r}) \sqrt{\log T}\right)\right). \end{aligned}$$

Using the fact that for all $t \in [T]$, $f_t((1 - \mu/r)\mathbf{x}^*) - f_t(\mathbf{x}_{m(t)}) \leq 2RG_f$ and $g_t^+(\mathbf{z}_t) \leq g_t^+(\mathbf{x}_{m(t)}) + \mu G_g$, we obtain from (19) the constraints violation bound:

$$\mathbb{E} \left[\sum_{t=1}^T g_t^+(\mathbf{z}_t) \right] \leq \sqrt{2(\delta \eta T + \eta^{-1}) \times (\text{RHS of (18)} + 2RG_f T)} + \mu G_g T. \quad (20)$$

Let us now get back to the condition in (17) which we assumed holds true. Plugging-in our choices of $\eta, \mu, \epsilon, K, \delta$, it suffices that the constant c_2 satisfies:

$$\begin{aligned} c_2 \geq \frac{12T^{-1/2}}{\sqrt{\log T}} \max \left\{ \frac{nM_g^2 T^{1/2}}{r}, G_g^2 T^{1/2}, \frac{c_2^2 R^2}{n \max\{G_f, G_g, M_f, M_g\}^2} \log T, \right. \\ \frac{\sqrt{183n} \max\{G_f, G_g, M_f, M_g\}^2 T^{1/2} \sqrt{\log T}}{\sqrt{r}}, \\ \frac{\sqrt{n} \max\{G_f, G_g, M_f, M_g\}^2 T^{1/2} (\sqrt{183R} \sqrt{\log T} + \sqrt{nr})}{R}, \\ \left. \frac{nT^{1/2} \max\{G_f, G_g, M_f, M_g\}^2}{\sqrt{r} R} \right\}. \end{aligned}$$

In particular, for any T large enough (so that the second term inside the max in the RHS is not dominant), it suffices to take

$$c_2 = 24\sqrt{183} \max\{G_f, G_g, M_f, M_g\}^2 \max\left\{\frac{n}{r}, \frac{n}{\sqrt{r}}, \sqrt{n}, \frac{n}{\sqrt{r}R}\right\}.$$

Plugging this choice into (20), and noting that the RHS of (18) scales (as a function of T) only as $T^{3/4}\sqrt{\log T}$, we have that for any T large enough,

$$\mathbb{E} \left[\sum_{t=1}^T g_t^+(\mathbf{z}_t) \right] = O \left(\sqrt{RG_f \max\{G_f, G_g, M_f, M_g\} \max\left\{\frac{R}{r}, \frac{R}{\sqrt{r}}, \frac{R}{\sqrt{n}}, \frac{1}{\sqrt{r}}, \frac{1}{R}\right\}} n^{1/4} T^{7/8} \log(T)^{1/4} \right).$$

Finally, since ϵ, K are the same as in Theorem 4.1, we have that the same upper-bound on the total number of calls to the linear optimization oracle applies, which concludes the proof. \square