

On Dynamic Regret and Constraint Violations in Constrained Online Convex Optimization

Rahul Vaze

School of Technology and Computer Science
Tata Institute of Fundamental Research
Mumbai, India, rahul.vaze@gmail.com

Abstract—A constrained version of the online convex optimization (OCO) problem is considered. With slotted time, for each slot, first an action is chosen. Subsequently the loss function and the constraint violation penalty evaluated at the chosen action point is revealed. For each slot, both the loss function as well as the function defining the constraint set is assumed to be smooth and strongly convex. In addition, once an action is chosen, local information about a feasible set within a small neighborhood of the current action is also revealed. An algorithm is allowed to compute at most one gradient at its point of choice given the described feedback to choose the next action. The goal of an algorithm is to simultaneously minimize the dynamic regret (loss incurred compared to the oracle’s loss) and the constraint violation penalty (penalty accrued compared to the oracle’s penalty). We propose an algorithm that follows projected gradient descent over a suitably chosen set around the current action. We show that both the dynamic regret and the constraint violation is order-wise bounded by the *path-length*, the sum of the distances between the consecutive optimal actions. Moreover, we show that the derived bounds are the best possible.

I. INTRODUCTION

Online convex optimization (OCO) has been a very attractive research problem for the last two decades, because of its versatility in modelling rich optimization problems. With OCO, at each time t , an online algorithm selects an action a_t , after which the loss incurred $f_t(a_t)$ is revealed. Knowing all f_t ’s, $1 \leq t \leq T$ ahead of time, an optimal offline algorithm chooses action $x^* = \arg \min_x \sum_{t=1}^T f_t(x)$, and the *static* regret of an online algorithm is defined as $R_s = \max_{f_t, t=1, \dots, T} \sum_{t=1}^T f_t(a_t) - \sum_{t=1}^T f_t(x^*)$, i.e., an adversary can choose the functions f_t . The name static comes from the fact that the optimal offline algorithm is constrained to use a single action.

Large body of work is known for static regret [1]–[5], where if the functions f_t are convex, the optimal regret is $\Theta(\sqrt{T})$, while if functions f_t are strongly convex then algorithms with regret at most $O(\log T)$ are known. When functions f_t are chosen by an adversary, but the arrival order is randomized, algorithms with better guarantees are also known [3].

Natural generalization of the static regret is the dynamic regret [6], [7], where the regret for a comparator sequence $\mathbf{u} = (u_1, \dots, u_T)$ is defined as

$$R_d(\mathbf{u}) = \max_{f_t, t=1, \dots, T} \sum_{t=1}^T f_t(a_t) - \sum_{t=1}^T f_t(u_t). \quad (1)$$

For this general dynamic regret definition, sub-linear (in T) regret is not always possible, unless some structure is enforced

on the comparator sequence. For example, for a sequence \mathbf{u} , defining $V_{\mathbf{u}} = \sum_{t=2}^T \|u_t - u_{t-1}\|$, the online gradient descent (OGD) algorithm was shown to achieve dynamic regret $R_d(\mathbf{u}) = O(\sqrt{T}(1 + V_{\mathbf{u}}))$ [6], which has been improved to $O(\sqrt{T}(1 + V_{\mathbf{u}}))$ in [8], matching the lower bound $\Omega(\sqrt{T}(1 + V_{\mathbf{u}}))$ [8].

A special case of (1) that is popularly studied [9]–[16] is by restricting $\mathbf{u} = \mathbf{x}^* = (x_1^*, \dots, x_T^*)$, where $x_t^* = \arg \min_x f_t(x)$, the sequence of local optimizers. Moreover, in this case, once the action a_t is chosen, the only information available is $\nabla f_t(a_t)$. For this case, the best known bound on the dynamic regret has been shown to be $O(V_{\mathbf{x}^*})$ [12] using the online projected gradient descent (OPGD) algorithm, when functions f_t are strongly convex, and smooth. Under the special case that the minimizers x_t^* lie in the interior of the feasible set, $O(V_{\mathbf{x}^*})$ regret can be achieved when functions f_t are just convex, and smooth [11]. For strongly convex and smooth functions, defining $S_{\mathbf{x}^*} = \sum_{t=2}^T \|x_t^* - x_{t-1}^*\|^2$, [13] showed that $O(\min\{V_{\mathbf{x}^*}, S_{\mathbf{x}^*}\})$ is also achievable as long as at each time t , gradient of f_t at as many points required is available. Most recently, [16] improved this guarantee to $O(\min\{V_{\mathbf{x}^*}, S_{\mathbf{x}^*}, FV\})$, where

$$FV = \sum_{t=2}^T \sup_{x \in \mathcal{X}} \|f_t(x) - f_{t-1}(x)\| \quad (2)$$

is the maximum function variation over the feasible set \mathcal{X} in successive time slots.

In this paper, we focus on the constrained version of the OCO that has been considered more recently in [17]–[19], where at each time t , the objective is to minimize the loss function f_t subject to a constraint $g_t(x) \leq 0$. For applications of constrained OCO, we refer to prior work [17]–[19]. Similar to the unconstrained OCO, even with constraints, the typical information structure is that information about f_t, g_t is revealed after the action a_t is chosen. The optimizer at time t is $x_t^* = \arg \min_{x \in \mathcal{X}: g_t(x) \leq 0} f_t(x)$, and the objective is to choose actions to minimize the dynamic regret (1) with $\mathbf{u} = \mathbf{x}^*$, where the constraint is already absorbed in the definition of x_t^* , and the maximization in (1) is with respect to both f_t and g_t . In addition to the dynamic regret, in this constrained OCO, an additional metric of interest is the constraint violation penalty,

which can be defined in many different ways. For example,

$$P_g(\mathbf{x}^*) = \max_{f_t, g_t, t=1, \dots, T} \sum_{t=1}^T \|g_t(x_t^*) - g_t(a_t)\|,$$

which measures the gap between the function g_t evaluated at the optimal point and the chosen action or

$$P'_g(\mathbf{x}^*) = \max_{f_t, g_t, t=1, \dots, T} \sum_{t=1}^T g_t(a_t),$$

which just counts the overall constraint violation. We use $P_g(\mathbf{x}^*)$ rather than $P'_g(\mathbf{x}^*)$ since it is a stronger measure as $P_g(\mathbf{x}^*) \geq P'_g(\mathbf{x}^*)$ on account of $g_t(x_t^*) \leq 0$.

In prior work, starting from [17], where functions f_t, g_t are assumed to be convex, Lipschitz and smooth, an algorithm has been proposed that achieves $R_d(\mathbf{x}^*) \leq O(V_{\mathbf{x}^*} \sqrt{T})$ while $P'_g(\mathbf{x}^*) = O(T^{1/2})$, which was improved in [18], to get $R_d(\mathbf{x}^*) \leq O(\sqrt{TV_{\mathbf{x}^*}})$ while $P'_g(\mathbf{x}^*) = O(V_{\mathbf{x}^*}^{1/4} T^{3/4})$, and most recently in [19], an algorithm based on the drift plus penalty method has regret $R_d(\mathbf{x}^*) \leq O(\max\{\sqrt{TV_{\mathbf{x}^*}}, V_g\})$ while $P'_g(\mathbf{x}^*) = O(\sqrt{T}, V_g)$, or $R_d(\mathbf{x}^*) \leq O(\sqrt{TV_{\mathbf{x}^*}})$ while $P'_g(\mathbf{x}^*) = O(T^{3/4}, V_g)$, where V_g is as defined in (2) with $f = g$.

However, notably [19] considers the full information setting, where once a_t is chosen, full functions f_t and g_t are revealed, and hence x_t^* can be computed. Clearly, obtaining this information is highly imposing. Moreover, [19] also needs to know the diameter D of the feasible set. In comparison, the result of [18] requires the knowledge of $V_{\mathbf{x}^*}$ instead of individual x_t^* , which is relatively less demanding, however, still very difficult to obtain in practice, as well as the knowledge of T and D .

In this paper, we consider an alternate information structure that is less imposing than considered in [18], [19]. The full feasible set at time t is $\chi_t = \{x \in \chi : g_t(x) < 0\}$. We assume that once the current action a_t is chosen, for a fixed constant $\text{dist} > 0$ that is independent of T , a subset of χ_t , set $\chi_t(a_t) = \{x : g_t(x) \leq 0\} \cap \mathcal{B}(a_t, \text{dist})$ is made available, where $\mathcal{B}(x, r)$ is a ball with radius r centered at x . Set $\chi_t(a_t)$ captures the feasible set in the neighborhood of the current action. With full information, e.g., in [19], $\text{dist} = \infty$. We will show that our results hold for any $\text{dist} > 0$.

With this new information structure, we consider the problem of simultaneously minimizing the dynamic regret and constraint violation penalty when f_t, g_t are strongly convex, Lipschitz and smooth. Generalizing the results when f_t, g_t are only convex, is part of ongoing work.

Towards this end, we propose an algorithm that uses the projected gradient descent (PGD) algorithm [12] as a black box, and depending on the chosen action a_t being feasible $g_t(a_t) < 0$, on the boundary $g_t(a_t) = 0$ or infeasible $g_t(a_t) > 0$, executes PGD over a suitably chosen subset that may or may not be contained in the feasible region of g_t . The main concept that the algorithm relies on is the property of the PGD algorithm [12] when executed over a convex set I and starting point a_t , is that the next action a_{t+1} satisfies

$$\|x_I^* - a_{t+1}\| \leq c \|x_I^* - a_t\|, \quad (3)$$

for a constant $c < 1$, where $x_I^* = \min_{x \in I} f(x)$ when f is strongly convex and smooth.

If the whole feasible region $\chi_t = \{x \in \chi : g_t(x) < 0\}$ was known, then using $I = \chi_t$, (3) will imply that the algorithm is making ‘quick’ progress towards the optimal point x_t^* . Unfortunately only local information about the feasible region χ_t is known. In particular, only $\chi_t(a_t) = \chi_t \cap \mathcal{B}(a_t, \text{dist})$ is available for a constant dist . Thus, we proceed in two steps. We identify a small region I_t at time t around a_t that is contained in χ_t and use (3) to claim that we are making progress towards the optimal point in this subset I_t (which could be far away from the global optimal). Next, exploiting the strong convexity and the smoothness of the functions, we extend the same claim to the optimal point x_t^* which need not be in I_t .

Since we have only local information about g_t around a_t , it can happen that the size of I_t is arbitrarily small or I_t is empty in case $g_t(a_t) > 0$ (current choice is infeasible). For both these cases, we show that the algorithm makes progress of a finite distance towards the optimal point x_t^* in χ_t , and establish a relation similar to (3). Once we have (3), a simple application of the triangle inequality and the Lipschitz condition, implies the result.

Our contributions.

- We show that under the defined information structure, the proposed algorithm simultaneously achieves $R_d(\mathbf{x}^*) \leq O(V_{\mathbf{x}^*})$ and $P'_g(\mathbf{x}^*) \leq P_g(\mathbf{x}^*) \leq O(V_{\mathbf{x}^*})$ for any $\text{dist} > 0$. Importantly, no information about $x_t^*, V_{\mathbf{x}^*}, T$ or D is needed.
- As a function of information variable $\text{dist} > 0$, both $R_d(\mathbf{x}^*)$ and $P_g(\mathbf{x}^*)$ scale inverse polynomially, which is natural to expect since for any algorithm as information availability is decreased, (smaller value of dist), the regret should worsen. We do not know at this point if the algorithm achieves the best scaling in terms of dist .
- In Remark 3, we also argue that our result is the best one can hope for, given the minimal information structure.

Notation: For the rest of the paper, we follow the notation described as follows. For a set $I \in \mathbb{R}^n$, its interior is defined as $\text{int}(I)$, while its boundary as $\text{boundary}(I)$. $\mathcal{B}(x, r)$ is the ball of radius r centered at x . For a discrete set of points S , $\text{convex hull}(x \in S)$ represents the convex hull of points $x \in S$. $\text{Proj}(x, S)$ is the projection of point x on set S , i.e. $\text{Proj}(x, S) = \arg \min_{y \in S} \|x - y\|$.

II. SYSTEM MODEL

Time is slotted with total time horizon T , and time slots are indexed as $t = 1, \dots, T$. Let $\chi \subset \mathbb{R}^n$ be a compact and convex set. For each t , two functions f_t and g_t are of interest, that are defined over χ . The feasible set at time t is defined as $\chi_t = \{x \in \chi : g_t(x) \leq 0\}$. Let the optimizer for f_t over the constraint set $g_t(x) \leq 0$ be x_t^* , i.e., $x_t^* = \arg \min_{\{x \in \chi_t\}} f_t(x)$.

We make the following standard assumptions about f_t and g_t . Functions f_t and g_t are assumed to be Lipschitz with Lipschitz constants \mathcal{L}_f and \mathcal{L}_g , respectively. Moreover, functions f_t and g_t are assumed to be smooth, i.e., the

gradients ∇f_t and ∇g_t are assumed to be Lipschitz with Lipschitz constants L_f and L_g , respectively. Moreover, for all $1 \leq t \leq T$, $\sup_{x \in \mathcal{X}} \|\nabla f_t(x)\| \leq G$ and $\sup_{x \in \mathcal{X}} \|\nabla g_t(x)\| \leq G$.¹ Compared to prior work [18], [19] that assume that f_t and g_t are convex, we assume that f_t and g_t are *strongly* convex with strong convexity parameters ν_f, ν_g , respectively.²

At each time t , an action a_t is chosen by an algorithm, for which the cost is $f_t(a_t)$. The goal of the algorithm to choose a_t such that the cost $f_t(a_t)$ is as small as possible while making sure that $a_t \in \chi_t$. However, the information available with the algorithm to choose a_t is limited and described as follows.

Information structure: Similar to [17]–[19], once the action $a_t \in \mathbb{R}^n$ is chosen at time t , $g_t(a_t)$ is revealed. Moreover, the algorithm can also access $\nabla f_t(x), \nabla g_t(x)$ for at most one point x of its choice. As described in the Introduction, additionally, in this paper, we assume that, set $\chi_t(a_t) = \chi_t \cap \mathcal{B}(a_t, \text{dist})$ is also revealed at time t for a fixed constant $\text{dist} > 0$, after a_t has been chosen. Note that dist can be arbitrarily small but is a constant that is fixed throughout the time horizon and does not depend on t or T . Compared to prior work, [17]–[19], acquiring this information is less imposing and does not involve finding any x_t^* . The set $\chi_t(a_t)$ maps the local behaviour of g_t in a very small neighborhood of a_t . Note that convexity implies that $\chi_t(a_t)$ is convex for any a_t .

Remark 1. *For the considered problem to be meaningful, once a_t is chosen, $g_t(a_t)$ has to be revealed, as already assumed in prior work [17]–[19]. In this work, in addition, we are assuming that $\chi_t(a_t)$ is also known which in turn requires that $g_t(x)$ for $x \in \mathcal{B}(a_t, \text{dist})$ is also known. When $\text{dist} \rightarrow 0$, this new information is equivalent to just acquiring $g_t(a_t)$. Since dist is allowed to be any arbitrarily small constant, the extra information assumed is very minimal and can be obtained similar to obtaining $g_t(a_t)$ (necessary), and can be done efficiently by exploiting the convexity of g_t .*

The performance metric for an online algorithm that chooses actions $a_t, t = 1, \dots, T$ is defined as the dynamic regret

$$R_d(\mathbf{x}^*) = \max_{f_t, g_t} \sum_{t=1}^T \|f_t(x_t^*) - f_t(a_t)\|,$$

and penalty for constraint violation as

$$P_g(\mathbf{x}^*) = \max_{f_t, g_t} \sum_{t=1}^T \|g_t(x_t^*) - g_t(a_t)\|,$$

where a_t 's are the causal actions of the algorithm that can depend on the information acquired till time slot $t - 1$. Moreover, f_t, g_t can be chosen by an adversary (can be adaptive, i.e., depend on previous actions $a_\tau, \tau \leq t - 1$) and are not required to follow any structure, other than what has been described earlier.

¹For notational simplicity we are assuming the same constant G , which can be generalized without any change in following analysis.

²We are assuming that all f_t 's and g_t 's have the same smoothness parameter L_f and L_g only for notational simplicity. All results will go through with different parameters as well.

Note that P_g is stronger than the penalty considered in earlier work [19] that is defined as $P'_g(\mathbf{x}^*) = \max_{f_t, g_t} \sum_{t=1}^T g_t(a_t)$, in two aspects. $P'_g(\mathbf{x}^*)$ can be negative, while $P_g(\mathbf{x}^*)$ is always positive, and $P_g(\mathbf{x}^*) \geq P'_g(\mathbf{x}^*)$ since $g_t(x_t^*)$ can be negative.

III. ALGORITHM

We present the proposed algorithm as a pseudo code in Algorithm 1, and describe it as follows. Let at time t ,

$$x_t^* = \arg \min_{x \in \chi_t} f_t(x).$$

Let the action chosen at time t be a_t . We want to choose a_{t+1} in such a way that

$$\|x_t^* - a_{t+1}\|^2 < c \|x_t^* - a_t\|^2, \quad (4)$$

for some constant $0 < c < 1$ that does not depend on t . Recall that while choosing a_{t+1} , no information about f_{t+1}, g_{t+1} is available. Thus, relation (4) is useful in the sense that it ensures that a_{t+1} is closer to x_t^* compared to a_t , in hope that if x_t^* and x_{t+1}^* are close, then a_{t+1} will be close to x_{t+1}^* as well.

The main idea of the algorithm is to accomplish this goal (showing that (4) holds) depending on three possible cases, namely : i) $g_t(a_t) < 0$, i.e. a_t is strictly feasible for g_t , ii) $g_t(a_t) = 0$, i.e. a_t is on the boundary of the feasible region for g_t , and finally, iii) $g_t(a_t) > 0$, i.e. a_t is strictly infeasible for g_t . Just to be clear, all the described actions in the following are taken after a_t is chosen and the information has been revealed about $\nabla f_t(x), g_t(a_t), \nabla g_t(x)$ and $\chi_t(a_t)$ for some one point x .

In case i) $g_t(a_t) < 0$, and we know that a_t is strictly feasible and potentially there is room to move to a point closer to x_t^* , the optimizer of f_t . Using the Lipschitz property of g_t 's, this implies that each point in the ball $\mathcal{B}(a_t, \|\frac{g_t(a_t)}{2L_g}\|)$ is also feasible. Thus, Algorithm 1 chooses the set $\mathcal{B}(a_t, \|\frac{g_t(a_t)}{2L_g}\|)$ as the feasible region to execute the PGD.

In case, the radius $\|\frac{g_t(a_t)}{2L_g}\|$ of the identified feasible region is smaller than the fixed constant dist , then using the extra local information $\chi_t(a_t)$ as described earlier, the feasible region is chosen as $\chi_t(a_t)$. A local gradient descent algorithm over the chosen feasible region using subroutine OPTIMIZE Algorithm 2 (online gradient descent) is used to find the next action a_{t+1} .

In case ii) $g_t(a_t) = 0$, a_t is on the boundary of the feasible region. In this case, we use the local information about g_t around a_t and choose $\chi_t(a_t)$ as the feasible region. Next, a local gradient descent is executed using subroutine OPTIMIZE Algorithm 2 in the identified feasible region to find the next action.

Finally in case iii) $g_t(a_t) > 0$ a_t is strictly infeasible. Since the current choice of a_t is infeasible for g_t , and $g_t(x_t^*) \leq 0$, it is sufficient to move towards the region for which $g_t(x) \leq 0$ to ensure (4) while staying infeasible. In fact, if we 'blindly' move into the feasible region, we cannot guarantee that a_{t+1} is closer to x_t^* than a_t , for example if $g_t(x_t^*) = 0$. However, using the Lipschitz condition, we know that each point in the ball

$\mathcal{B}(a_t, \|\frac{g_t(a_t)}{2\mathcal{L}_g}\|)$ is infeasible given that a_t is infeasible. Thus, in this case as long as $\|\frac{g_t(a_t)}{2\mathcal{L}_g}\| \geq \|\nabla g_t(a_t)\frac{1}{L_g}\|$ we move a distance of $\|\nabla g_t(a_t)\frac{1}{L_g}\|$ from a_t in the direction of negative gradient of g_t at a_t . Thus the new point a_{t+1} is still infeasible, but as we show in Lemma 7, a_{t+1} is closer to x_t^* than a_t . In case $\|\frac{g_t(a_t)}{\mathcal{L}_g}\| < \|\nabla g_t(a_t)\frac{1}{L_g}\|$, the algorithm finds a feasible region similar to case i) using the local information $\chi_t(a_t)$ and follow a local gradient descent in this feasible region using subroutine OPTIMIZE Algorithm 2 to find the next action. In case, $\chi_t(a_t)$ turns out to be an empty set, we proceed similar to the case when $\|\frac{g_t(a_t)}{\mathcal{L}_g}\| \geq \|\nabla g_t(a_t)\frac{1}{L_g}\|$, since the whole of $\mathcal{B}(a_t, \text{dist})$ is infeasible.

Algorithm 1 Algorithm

```

1: Input  $L_f, L_g, \mathcal{L}_g, \text{dist} > 0$ , feasible set  $\chi \subset \mathbb{R}^n$ 
2: Initialize  $t = 0$ 
3: Choose action  $a_1$  arbitrarily belonging to  $\chi$ 
4: while  $t \leq T - 1$  do
5:    $t = t + 1$ 
6:   if  $g_t(a_t) < 0$  then %previous action  $a_t$  was strictly
feasible for  $g_t$ 
7:      $\delta_t = \|\frac{g_t(a_t)}{2\mathcal{L}_g}\|$ 
8:     if  $\delta_t \geq \text{dist}$  then
9:        $I_t = \mathcal{B}(a_t, \delta_t)$ 
10:       $a_{t+1} = \text{OPTIMIZE}(f_t, I_t, \frac{1}{2L_f}, a_t)$ 
11:     else
12:       Find the feasible region  $\chi_t(a_t) = \chi_t \cap$ 
 $\mathcal{B}(a_t, \text{dist})$ 
13:        $I_t = \chi_t(a_t)$ 
14:        $a_{t+1} = \text{OPTIMIZE}(f_t, I_t, \frac{1}{2L_f}, a_t)$ 
15:     end if
16:     else if  $g_t(a_t) = 0$  then % $a_t$  is on the boundary of the
feasible region
17:       Find the feasible region  $\chi_t(a_t) = \chi_t \cap \mathcal{B}(a_t, \text{dist})$ 
18:        $I_t = \chi_t(a_t)$ 
19:        $a_{t+1} = \text{OPTIMIZE}(f_t, I_t, \frac{1}{2L_f}, a_t)$ 
20:     else if  $g_t(a_t) > 0$  then % $a_t$  is infeasible
21:       if  $\delta_t \geq \|\nabla g_t(a_t)\frac{1}{L_g}\|$  then
22:          $a_{t+1} = a_t + \alpha(\hat{a}_t - a_t)$ , where -
23:          $\hat{a}_t = a_t - \nabla g_t(a_t)\frac{1}{L_g}$ ,
24:       else
25:         Find the feasible region  $\chi_t(a_t) = \chi_t \cap$ 
 $\mathcal{B}(a_t, \text{dist})$ 
26:         if  $\chi_t(a_t) \neq \emptyset$  then
27:            $I_t = \chi_t(a_t)$ 
28:            $a'_t = \text{Proj}(a_t, I_t)$ 
29:            $a_{t+1} = \text{OPTIMIZE}(f_t, I_t, \frac{1}{2L_f}, a'_t)$ 
30:         else
31:            $a_{t+1} = a_t + \alpha(\hat{a}_t - a_t)$ , where -
32:            $\hat{a}_t = a_t - \nabla g_t(a_t)\frac{\text{dist}}{\|\nabla g_t(a_t)\|}$ 
33:         end if
34:       end if
35:     end if
36: end while

```

Theorem 1. When both f_t, g_t are strongly convex, Lipschitz, and smooth for all $t \leq T$ and $1 \leq t \leq T$, $\sup_{x \in \chi} \nabla f_t(x) \leq G$ and $\sup_{x \in \chi} \nabla g_t(x) \leq G$, with information structure as defined, for Algorithm 1

$$\|x_t^* - a_{t+1}\| < c\|x_t^* - a_t\|,$$

for some constant $0 < c < 1$ that does not depend on t . In particular,

$$c = \max\{c_2, c_3, c_4, c_5\} < 1,$$

where $c_2 = \left(1 - \frac{\alpha\nu_f}{2L_f}\right)^{1/2}$, $c_3 = \frac{D+\alpha\text{dist}}{D+\text{dist}}$ and $c_4 = (1 - \alpha\nu_g/L_g)^{1/2}$, $c_5 = \left(1 - \alpha\frac{\nu_g}{\max\{G/\text{dist}, L_g\}}\right)^{1/2}$, and $0 < \alpha < 1$ is a constant to be chosen by subroutine OPTIMIZE, D is the diameter of the feasible region and $0 < \text{dist}$ is a constant to be chosen by Algorithm 1. Note that $\nu_g \leq L_g$ and $\nu_f \leq L_f$ always, thus $0 < c < 1$.

Algorithm 2 OPTIMIZE

```

1: Input( $h, I, \mu, x_t$ )
2: Constant  $0 < \alpha < 1$ ,  $x_{t+1} = x_t + \alpha(\hat{x}_t - x_t)$ , where

$$\hat{x}_t = \text{Proj}(x_t - \frac{1}{\mu}\nabla h(x_t), I).$$


```

Using Theorem 1, we get the main result of the paper as follows.

Theorem 2. When both f_t, g_t are strongly convex, Lipschitz, and smooth ($\nabla f_t, \nabla g_t$ are Lipschitz) for all $t \leq T$ and $1 \leq t \leq T$, $\sup_{x \in \chi} \nabla f_t(x) \leq G$ and $\sup_{x \in \chi} \nabla g_t(x) \leq G$, with information structure as defined, with Algorithm 1, simultaneously,

$$R_d(\mathbf{x}^*) = O(V_{\mathbf{x}^*}), \text{ and, } P_g(\mathbf{x}^*) = O(V_{\mathbf{x}^*}).$$

Remark 2. Both the regret and constraint violation penalty bounds derived in Theorem 2 are inverse polynomially proportional to the chosen constant dist . In particular, they grow as $\frac{1}{1-c_5}$ where $c_5 = \left(1 - \alpha\frac{\nu_g}{\max\{G/\text{dist}, L_g\}}\right)^{1/2}$. It is natural to expect that regret grows with decreasing dist since for any algorithm as information availability is decreased, (in this case smaller value of dist), the regret should worsen. However, dist can be any constant and not necessarily has to be $\ll 1$, and there is a tradeoff between regret and the amount of available local feasibility information $\mathcal{B}(a_t, \text{dist})$.

Remark 3. For the unconstrained OCO, when at each step gradient information is available only at a single point, the best known algorithm when each f_t is smooth, and strongly convex, has $R_d(\mathbf{x}^*) \leq O(V_{\mathbf{x}^*})$ [12]. Note that $V_{\mathbf{x}^*}$ in the constrained and the unconstrained OCO problem are different, therefore directly we cannot compare our result with that of [12]. However, since functions g_t and f_t are allowed to be arbitrary with the constrained OCO, $g_t = f_t$ for each t is a valid choice for g_t and f_t . With $g_t = f_t$, the constrained OCO collapses to the unconstrained OCO, for which the best

known result on regret is $O(V_{\mathbf{x}^*})$, making the derived result (which also needs gradient availability at only one point) the best possible.

Proof of Theorem 2. Using the triangle inequality, we get that $\sum_{t=1}^T \|x_t^* - a_t\|$

$$\begin{aligned} &\leq \|x_1^* - a_1\| + \sum_{t=2}^T \|x_{t-1}^* - a_t\| + \sum_{t=2}^T \|x_t^* - x_{t-1}^*\|, \\ &\stackrel{(a)}{\leq} \|x_1^* - a_1\| + c \sum_{t=2}^T \|x_{t-1}^* - a_{t-1}\| + \sum_{t=2}^T \|x_t^* - x_{t-1}^*\|, \\ &\stackrel{(b)}{\leq} \|x_1^* - a_1\| - c \|x_T^* - a_T\| + c \sum_{t=1}^T \|x_t^* - a_t\| \\ &\quad + \sum_{t=2}^T \|x_t^* - x_{t-1}^*\|, \end{aligned} \quad (5)$$

where (a) is obtained by using Theorem 1, while to obtain (b) we added and subtracted $c\|a_T - x_T^*\|$ and rearranged terms.

Regrouping terms in (5), we get $\sum_{t=1}^T \|x_t^* - a_t\|$

$$\leq \frac{\|x_1^* - a_1\| - c\|x_T^* - a_T\|}{1-c} + \frac{1}{1-c} \sum_{t=2}^T \|x_t^* - x_{t-1}^*\|. \quad (6)$$

Thus, using the Lipschitz property of f_t and g_t , (6) implies that $R_d(\mathbf{x}^*) = \sum_{t=1}^T \|f_t(x_t^*) - f_t(a_t)\|$

$$\leq \frac{L_f}{1-c} \sum_{t=2}^T \|x_t^* - x_{t-1}^*\| + \frac{D}{1-c} = O(V_{\mathbf{x}^*})$$

and $P_g(\mathbf{x}^*) = \sum_{t=1}^T \|g_t(x_t^*) - g_t(a_t)\|$

$$\leq \frac{L_g}{1-c} \sum_{t=2}^T \|x_t^* - x_{t-1}^*\| + \frac{D}{1-c} = O(V_{\mathbf{x}^*}),$$

where $V_{\mathbf{x}^*} = \sum_{t=2}^T \|x_t^* - x_{t-1}^*\|$, the accumulated variation of the per-step minimizers. \square

Next, we first briefly discuss the basic difference between the proposed algorithm and the relevant prior work. In [18], a primal dual algorithm has been proposed using the Lagrangian

$$L(x, \lambda) = f_t(x) + \lambda^T g_t(x) + \frac{\eta}{\|\lambda\|^2},$$

where a_t is updated using the gradient descent over the Lagrangian to move towards the optimizer of f_t with penalty function λ as

$$a_{t+1} = a_t - \eta \nabla_x L(x, \lambda),$$

while gradient ascent is used to increase the penalty in case of constraint violation as $\lambda_{t+1} = \lambda_t - \eta \nabla_\lambda L(x, \lambda)$.

Similarly, in [19], a primal dual algorithm is proposed where the increase in λ is derived by minimizing the expected ‘drift’ of the constraint violation. In particular, it is given by

$$\lambda_{t+1} = \max\{\lambda_t + \eta_t g_t(a_t), -\eta_t g_t(a_t)\},$$

while

$$a_{t+1} = \nabla f_t^T(a_t)(x - a_t) + \mu_t \|x - a_t\|^2 + [\lambda_t + \eta_t g_t(a_t)]^T \eta_t g_t(a_t).$$

Both these algorithms [18], [19] are long-term in the sense that they want to remain close to x_t^* while minimizing the constraint violation penalty $P'_g(\mathbf{x}^*)$ in the long term, i.e., they nudge the updates ‘slowly’ in the direction of constraint satisfaction to avoid large accumulated constraint violation penalty. In contrast, the proposed algorithm in this paper is local, and is trying to go close to the optimal point in every single step as shown in Lemma 5, 6 and 7. Thus, conceptually our algorithm is entirely different than [18], [19].

In terms of restrictions, over and above [18], [19], we assume that f_t and g_t are strongly convex, however in terms of information, we require far less. In particular, at time t , after a_t has been chosen, Algorithm 1 requires only $g_t(a_t)$, $\chi_t(a_t)$ and $\nabla f_t(x)$, $\nabla g_t(x)$, at $x = a_t$ or some $x \in \chi_t(a_t)$. In contrast, [19] assumes that once a_t is chosen, full f_t, g_t are revealed, making x_t^* known. Moreover, it requires the knowledge of the diameter D . In [18], knowledge of $V_{\mathbf{x}^*}, D, T$ is needed over and above $\nabla f_t(a_t), \nabla g_t(a_t), g_t(a_t)$.

In the rest of the paper, we prove Theorem 1, for which we need the following Lemma regarding the subroutine OPTIMIZE.

Lemma 3. [12] *If function h is ν_h -strongly convex, and ∇h is Lipschitz with parameter L_h , and $x^{I^*} = \arg \min_{x \in I} h(x)$, then if parameter $\mu \geq L_h$, the output x_{t+1} from subroutine OPTIMIZE satisfies*

$$\|x^{I^*} - x_{t+1}\| \leq c \|x^{I^*} - x_t\|, \quad (7)$$

for $c = \left(1 - \alpha \frac{\nu_h}{\mu}\right)^{1/2} < 1$.

Corollary 4. *For subroutine OPTIMIZE, let $\tilde{x} \in I$ be such that $h(\tilde{x}) < h(\hat{x}_t)$, then with parameter $\mu \geq L_h$, the output x_{t+1} from subroutine OPTIMIZE satisfies*

$$\|\tilde{x} - x_{t+1}\| \leq c \|\tilde{x} - x_t\|, \quad (8)$$

for $c = \left(1 - \alpha \frac{\nu_h}{\mu}\right)^{1/2} < 1$ as long as function h is ν_h -strongly convex, and ∇h is Lipschitz with parameter L_h .

Proof. The only place where optimality of x^{I^*} is used in the proof of Lemma 3 in [12] is to show that $h(\hat{x}_t) > h(x^{I^*})$. Thus, the proof goes through as it is, even with this weaker condition that $h(\tilde{x}) < h(\hat{x}_t)$. For completeness, the full proof is given in Section V. Another way to see the result is that by pruning I to get I' such that $\tilde{x} = \arg \min_{I'} h(x)$, while keeping h a strongly convex function over I' . Thus, the result follows directly from Lemma 3. \square

For ease of exposition, we break the proof of Theorem 1 into three parts corresponding to $g_t(a_t) < 0, g_t(a_t) = 0$, and $g_t(a_t) > 0$ in the next three lemmas.

Lemma 5. *When both f_t, g_t are strongly convex and smooth for all $t \leq T$ with information structure as defined, with Algorithm 1, for the case when $g_t(a_t) < 0$*

$$\|x_t^* - a_{t+1}\| < c_1 \|x_t^* - a_t\|,$$

where $0 < c_1 = \max\{c_2, c_3\} < 1$ for $c_2 = \left(1 - \frac{\alpha\nu_f}{2L_f}\right)^{1/2}$, $c_3 = \frac{D+\alpha\text{dist}}{D+\text{dist}} < 1$ that does not depend on t . Since $L_f \geq \nu_f$ (always), $c_2 < 1$.

Proof. Recall that $\delta_t = \left\| \frac{g_t(a_t)}{2L_g} \right\|$.

Case a) $\delta_t \geq \text{dist}$. In this case, $I_t = \mathcal{B}(a_t, \delta_t)$ and $I_t \subseteq \chi_t$ using the Lipschitz condition on g_t .

Subroutine OPTIMIZE is executed with set I_t and starting point a_t . The output of Subroutine OPTIMIZE is

$$\begin{aligned} a_{t+1} &= a_t + \alpha(\hat{a}_t - a_t), \\ \hat{a}_t &= \text{Proj}\left(a_t - \nabla f_t(a_t) \frac{1}{2L_f}, I_t\right). \end{aligned} \quad (9)$$

Subcase a-i) $\hat{a}_t \in \text{convexhull}(a_t, x_t^*)$ (just the line segment connecting a_t and x_t^*). If $x_t^* \in I_t$, then directly from Lemma 3, we get

$$\|a_{t+1} - x_t^*\| \leq c_2 \|a_t - x_t^*\|, \quad (10)$$

where $c_2 = \left(1 - \alpha \frac{\nu_f}{2L_f}\right)^{1/2}$ as we have chosen $\mu = 2L_f$.

Otherwise, if $x_t^* \notin I_t$, then $\hat{a}_t \in \text{boundary}(I_t)$ since f_t is strongly convex, $\hat{a}_t \in \text{convexhull}(a_t, x_t^*)$ and $x_t^* \notin I_t$. Thus, the distance between \hat{a}_t and a_t is at least dist since $\delta_t \geq \text{dist}$, and the distance between a_{t+1} and a_t is at least αdist , while the distance between a_{t+1} and x_t^* is at most D (the diameter). Thus, we get that

$$\|a_{t+1} - x_t^*\| \leq c_3 \|a_t - x_t^*\|, \quad (11)$$

where $c_3 = \frac{D+\alpha\text{dist}}{D+\text{dist}}$.

Subcase a-ii) $\hat{a}_t \notin \text{convexhull}(a_t, x_t^*)$

If $x_t^* \in I_t$, then directly from Lemma 3, we get that

$$\|a_{t+1} - x_t^*\| \leq c_2 \|a_t - x_t^*\|, \quad (12)$$

as we have chosen $\mu = 2L_f$.

Thus, consider the case when $x_t^* \notin I_t$. Let $I'_t = \text{convex hull}(a_t, \hat{a}_t, x_t^*)$ where $I'_t \subseteq \chi_t$, i.e. full set I'_t is feasible, since g_t is convex.

Now, consider that if Subroutine OPTIMIZE is executed with set I'_t and the same starting point a_t , the output of Subroutine OPTIMIZE will be the same as (9), since

$$\text{Proj}\left(a_t - \nabla f_t(a_t) \frac{1}{2L_f}, I_t\right) = \text{Proj}\left(a_t - \nabla f_t(a_t) \frac{1}{2L_f}, I'_t\right),$$

irrespective of whether $\tilde{a}_t = a_t - \nabla f_t(a_t) \frac{1}{2L_f}$ belongs to I_t or not. However, since $x_t^* \in I'_t$, we get from Lemma 3 that

$$\|a_{t+1} - x_t^*\| \leq c_2 \|a_t - x_t^*\|. \quad (13)$$

An illustration of the basic idea of the proof when $\tilde{a}_t \notin I_t$ is presented in Fig. 1.

Case b) $\delta_t < \text{dist}$ Except for the choice of set I_t which is now $\chi_t(a_t)$ everything else is same as in case a). Moreover, since $\chi_t(a_t) \in \chi_t$ is feasible by definition, the same arguments as detailed in case a) apply, and we either get (11) or (13).

The two distinct choices of I_t are essentially made to speed up the algorithm. Always choosing $I_t = \chi_t(a_t)$ is sufficient for analysis.




Fig. 1: Illustration for the proof of Lemma 5 case a-ii), where the blue dashed triangle is $I'_t = \text{convex hull}(a_t, \hat{a}_t, x_t^*) \subseteq \chi_t$.

□

At this point it is difficult to appreciate the power of Lemma 5. What Lemma 5 saying is that irrespective of the size (how small) of set I_t chosen by the algorithm, as well as independent of the distance of x_t^* (however far) from I_t , we get a relation (13), that states that the distance between the optimal point and the updated point a_{t+1} contracts by a fixed amount compared to the original point a_t . The main tool that we are exploiting to prove Lemma 5 is both the strong convexity as well as the smoothness (gradient being Lipschitz) of the function f_t , and in some measure of g_t . To gather more intuition we consider a one-dimensional case in Figs. 2 and 3 to show how strong convexity together with smoothness indicates that contraction of distance from the optimal holds independent of the distance between the present point a_t , the updated point a_{t+1} , and the optimal point x_t^* .

In Fig. 2, for function $f(x)$ which is assumed to be strongly convex and smooth, we consider that the feasible set is $\chi_1 = (-\infty, x_2)$ and $x^* = x_1$, while in Fig. 3 it is $\chi_2 = (-\infty, x_2)$ and $x^* = x_2$. Clearly, by construction, a_{t+1} remains the same when OPTIMIZE is executed with starting point a_t , input function $h = f$ with $I = \chi_1$ or χ_2 and an identical choice of μ . Thus, from Lemma 3, we get that

$$\|a_{t+1} - x_1\| \leq c \|a_t - x_1\|. \quad (14)$$

as well as

$$\|a_{t+1} - x_2\| \leq c \|a_t - x_2\|. \quad (15)$$

for the same $c < 1$. Clearly, as x_2 is moved sufficiently far away to the right, one does not expect (15) to hold together with (14). However, since f is both strongly convex and smooth, there is a limit on how far x_2 can be compared to x_1 , before f starts to increase. This is the key reason behind both (14) and (15) to be true. Essentially, when f is both strongly convex and smooth, it is ‘trapped’ between a lower and an upper envelope.

In general, coming back to Lemma 5, because of the strong convexity and smoothness, the function f_t cannot continue to decrease beyond a point, and the estimate one gets for the contraction in (13) is an underestimate when x_t^* is close to a_t and a_{t+1} , while becomes tighter as x_t^* is drawn away from a_t and a_{t+1} .




Fig. 2: $\chi_1 = (-\infty, x_1)$




Fig. 3: $\chi_2 = (-\infty, x_2)$

Lemma 6. When both f_t, g_t are strongly convex and smooth for all $t \leq T$ with information structure as defined, with Algorithm 1, for the case when $g_t(a_t) = 0$

$$\|x_t^* - a_{t+1}\| < c_2 \|x_t^* - a_t\|.$$

Proof. When $g_t(a_t) = 0$, a_t is on the boundary, and the chosen set is $I_t = \chi_t(a_t)$ which by definition is feasible. Thus, the analysis is identical to that of Lemma 5, and we get the same relation as in Lemma 5 as required. \square

Next, we consider the final case when $g(a_t) > 0$, which is the most involved of the lot.

Lemma 7. When both f_t, g_t are strongly convex and smooth for all $t \leq T$, $1 \leq t \leq T$, $\sup_{x \in \chi} \nabla f_t(x) \leq G$ and $\sup_{x \in \chi} \nabla g_t(x) \leq G$, with information structure as defined, with Algorithm 1, for the case when $g_t(a_t) > 0$

$$\|x_t^* - a_{t+1}\| < c_6 \|x_t^* - a_t\|,$$

where $c_6 = \max\{c_2, c_3, c_4, c_5\} < 1$ for $c_4 = (1 - \alpha \nu_g / L_g)^{1/2}$, $c_5 = \left(1 - \alpha \frac{\nu_g}{\max\{G/\text{dist}, L_g\}}\right)^{1/2}$.

For proving Lemma 7, we will use the strong convexity of g_t as well as f_t .

Proof. Case a) $\delta_t = \|\frac{g_t(a_t)}{2L_g}\| \geq \|\nabla g_t(a_t)\frac{1}{L_g}\|$, in which case the update is

$$a_{t+1} = a_t + \alpha(\hat{a}_t - a_t), \quad (16)$$

where

$$\hat{a}_t = a_t - \nabla g_t(a_t) \frac{1}{L_g}.$$

Since $\delta_t = \|\frac{g_t(a_t)}{2L_g}\| \geq \|\nabla g_t(a_t)\frac{1}{L_g}\|$, the Lipschitz condition on g_t implies that $\mathcal{B}(a_t, \delta_t) \cap \chi_t = \emptyset$. Thus, $g_t(\hat{a}_t) > 0$, i.e. \hat{a}_t is still infeasible for g_t , and we want to show that

$$\|x_t^* - a_{t+1}\|^2 \leq c \|x_t^* - a_t\|^2, \quad (17)$$

for some fixed constant $c < 1$ that does not depend on t . Towards this end, we will exploit the strong convexity of g_t .

We will connect the update (16) with an update Subroutine OPTIMIZE will make on a suitable initial point x_t , function h , step size μ , and a feasible set I . Recall that $x_t^* = \arg \min_{x \in \chi_t} f_t(x)$. Consider a new set $I'_t = \text{convex hull}(a_t, \hat{a}_t, x_t^*)$, where a_t and \hat{a}_t are as defined in (16). As discussed above, both $g_t(a_t) > 0$ and $g_t(\hat{a}_t) > 0$. Important to note that x_t^* is not necessarily equal to $\arg \min_{x \in I'_t} g_t(x)$. However, $g_t(x_t^*) < g_t(\hat{a}_t) < g_t(a_t)$ since $g_t(x_t^*) \leq 0$, while $g_t(a_t) > 0$ and $g_t(\hat{a}_t) > 0$, and $g_t(\hat{a}_t) < g_t(a_t)$ since $\nabla g_t(a_t)$ is a descent direction for g_t .

Consider the update x_{t+1} which Subroutine OPTIMIZE will make if the initial/starting point $x_t = a_t$, the set $I = I'_t$ with step size $\mu = L_g$ and $h = g_t$. Since $\hat{a}_t \in I'_t$, $\hat{a}_t = \text{Proj}(\hat{a}_t, I'_t)$. Hence from Subroutine OPTIMIZE we get that

$$\hat{x}_t = \text{Proj}(a_t - \nabla g_t(a_t) \frac{1}{L_g}, I'_t) = a_t - \nabla g_t(a_t) \frac{1}{L_g} = \hat{a}_t, \quad (18)$$

$$x_{t+1} = a_t + \alpha(\hat{x}_t - a_t), \quad (19)$$

coinciding with (16). Thus, the update of the algorithm (16) is equivalent to executing Subroutine OPTIMIZE with starting point $x_t = a_t$, set $I = I'_t$ with step size $\mu = L_g$, for function $h = g_t$. So we would like to use Lemma 3. However, since x_t^* need not be $\arg \min_{x \in I'_t} g_t(x)$, we cannot use Lemma 3 directly. Instead we exploit the fact that $g_t(x_t^*) < g_t(\hat{a}_t) < g_t(a_t)$. Hence Corollary 4 becomes applicable, and we get that

$$\|a_{t+1} - x_t^*\| \leq c_4 \|a_t - x_t^*\|, \quad (20)$$

with $c_4 = (1 - \alpha \nu_g / L_g)^{1/2}$ since we have chosen $\mu = L_g$, inverse of the step size in Subroutine OPTIMIZE.

Case b) $\delta_t = \|\frac{g_t(a_t)}{2L_g}\| < \|\nabla g_t(a_t)\frac{1}{L_g}\|$

In this case, we have no sufficiently sized estimate of the infeasible region around a_t . Thus, we will exploit the strong convexity and smoothness of f_t , as follows.

b-i) Let $\chi_t(a_t) \neq \emptyset$. In this case, $I_t = \chi_t(a_t)$, $a'_t = \text{Proj}(a_t, I_t)$ and $a_{t+1} = \text{OPTIMIZE}(f_t, I_t, \frac{1}{2L_f}, a'_t)$.

Subcase b-i-i) Let $I_t = \chi_t(a_t) \subseteq \text{int}(\mathcal{B}(a_t, \text{dist}))$ which implies that $x_t^* \in I_t$. Recall that by definition, $a'_t = \text{Proj}(a_t, I_t)$. Thus, with $x_t^* \in I_t$, we get directly from Lemma 3 that

$$\|a_{t+1} - x_t^*\| \leq c_2 \|a'_t - x_t^*\| \leq c_2 \|a_t - x_t^*\|,$$

where the last inequality follows since $a_t \notin I_t$.

Subcase b-i-ii) Let $I_t = \chi_t(a_t) \not\subseteq \mathcal{B}(a_t, \text{dist})$. In this subcase, we get that $g_t(a'_t) = 0$ and is identical to the case considered in Lemma 6, except the starting point is a'_t instead of a_t . Thus, similar to (13), we get the first inequality

$$\|x_t^* - a_{t+1}\| \leq \max\{c_2, c_3\} \|x_t^* - a'_t\| \leq \max\{c_2, c_3\} \|x_t^* - a_t\|, \quad (21)$$

where the second inequality follows since $a_t \notin I_t$.

Case b-ii) Let $\chi_t(a_t) = \emptyset$. In this case, the update is

$$a_{t+1} = a_t + \alpha(\hat{a}_t - a_t), \quad (22)$$

where $\hat{a}_t = a_t - \nabla g_t(a_t) \frac{\text{dist}}{\|\nabla g_t(a_t)\|}$. Since $\chi_t(a_t) = \mathcal{B}(a_t, \text{dist}) \cap \chi_t$ is empty, $g_t(\hat{a}_t) > 0$. Thus, we can exploit the strong convexity and smoothness of g_t as in case a).

Given the assumption that $\nabla g_t(a_t) \leq G$, we get that (22) is equivalent to executing OPTIMIZE with set $I_t = \text{convex hull}(a_t, \hat{a}_t, x_t^*)$, $h = g_t$, and $\mu = \max\{G/\text{dist}, L_g\}$. Thus, similar to (20), since $g_t(x_t^*) < g_t(\hat{a}_t) < g_t(a_t)$, we get

$$\|a_{t+1} - x_t^*\| \leq c_5 \|a_t - x_t^*\|, \quad (23)$$

where $c_5 = \left(1 - \alpha \frac{\nu_g}{\max\{G/\text{dist}, L_g\}}\right)^{1/2}$.

IV. CONCLUSIONS

In this paper, we considered a constrained OCO problem, and provided the best (simultaneously) possible bounds for the regret and the constraint violation penalty, when both the loss function and the function defining the constraint are strongly convex and smooth. Compared to prior work, we proposed an algorithm that has better regret and penalty bounds while using significantly less information requirement about the loss function and the function defining the constraints. Extending these results when the respective functions are just convex and not strongly convex, remains open. \square

REFERENCES

- [1] N. Srebro, K. Sridharan, and A. Tewari, "Smoothness, low noise and fast rates," *Advances in neural information processing systems*, vol. 23, 2010.
- [2] P. Bartlett, E. Hazan, and A. Rakhlin, "Adaptive online gradient descent," in *Advances in Neural Information Processing Systems 20: Proceedings of the 2007 Conference*. Neural Information Processing Systems (NIPS) Foundation, 2009, pp. 65–72.
- [3] D. Garber, G. Korcia, and K. Levy, "Online convex optimization in the random order model," in *International Conference on Machine Learning*. PMLR, 2020, pp. 3387–3396.
- [4] E. Hazan, "Introduction to online convex optimization," *arXiv preprint arXiv:1909.05207*, 2019.
- [5] S. Shalev-Shwartz *et al.*, "Online learning and online convex optimization," *Foundations and trends in Machine Learning*, vol. 4, no. 2, pp. 107–194, 2011.
- [6] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *Proceedings of the 20th international conference on machine learning (icml-03)*, 2003, pp. 928–936.
- [7] E. Hall and R. Willett, "Dynamical models and tracking regret in online convex programming," in *International Conference on Machine Learning*. PMLR, 2013, pp. 579–587.
- [8] L. Zhang, S. Lu, and Z.-H. Zhou, "Adaptive online learning in dynamic environments," in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 2018, pp. 1330–1340.
- [9] O. Besbes, Y. Gur, and A. Zeevi, "Non-stationary stochastic optimization," *Operations research*, vol. 63, no. 5, pp. 1227–1244, 2015.
- [10] A. Jadbabaie, A. Rakhlin, S. Shahrampour, and K. Sridharan, "Online optimization: Competing with dynamic comparators," in *Artificial Intelligence and Statistics*. PMLR, 2015, pp. 398–406.
- [11] T. Yang, L. Zhang, R. Jin, and J. Yi, "Tracking slowly moving clairvoyant: Optimal dynamic regret of online learning with true and noisy gradient," in *International Conference on Machine Learning*. PMLR, 2016, pp. 449–457.
- [12] A. Mokhtari, S. Shahrampour, A. Jadbabaie, and A. Ribeiro, "Online optimization in dynamic environments: Improved regret rates for strongly convex problems," in *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 7195–7201.
- [13] L. Zhang, T. Yang, J. Yi, R. Jin, and Z.-H. Zhou, "Improved dynamic regret for non-degenerate functions," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017, pp. 732–741.
- [14] P. Zhao, G. Wang, L. Zhang, and Z.-H. Zhou, "Bandit convex optimization in non-stationary environments," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 1508–1518.
- [15] L. Zhang, S. Lu, and T. Yang, "Minimizing dynamic regret and adaptive regret simultaneously," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 309–319.
- [16] P. Zhao and L. Zhang, "Improved analysis for dynamic regret of strongly convex and smooth functions," in *Learning for Dynamics and Control*. PMLR, 2021, pp. 48–59.
- [17] T. Chen and G. B. Giannakis, "Bandit convex optimization for scalable and dynamic iot management," *IEEE Internet of Things Journal*, vol. 6, no. 1, pp. 1276–1286, 2018.
- [18] X. Cao and K. R. Liu, "Online convex optimization with time-varying constraints and bandit feedback," *IEEE Transactions on automatic control*, vol. 64, no. 7, pp. 2665–2680, 2018.
- [19] Q. Liu, W. Wu, L. Huang, and Z. Fang, "Simultaneously achieving sublinear regret and constraint violations for online convex optimization with time-varying constraints," *Perform. Evaluation*, vol. 152, p. 102240, 2021. [Online]. Available: <https://doi.org/10.1016/j.peva.2021.102240>

V. PROOF OF COROLLARY 4

Using the ν_h -strong convexity, and L smoothness of function h , from (31) [12], we have that for any $x \in I$ and $x_t \in I$ as the starting point and \hat{x}_t as defined in OPTIMIZE, for $\mu \geq L$, $h(x) - \frac{\nu_h}{2} \|x - x_t\|^2$

$$\geq h(\hat{x}_t) + \frac{\mu}{2} \|\hat{x}_t - x_t\|^2 + \mu(x_t - \hat{x}_t)^T(x - x_t).$$

Choosing $x = x' \in I$ such that $h(x') < h(\hat{x}_t)$, and rearranging terms, we get that $h(x') - h(\hat{x}_t)$

$$\geq \frac{\mu}{2} \|\hat{x}_t - x_t\|^2 + \frac{\nu_h}{2} \|x' - x_t\|^2 + \mu(x_t - \hat{x}_t)^T(x' - x_t).$$

Using the fact that $h(x) < h(\hat{x}_t) < h(x_t)$, the LHS is negative, and dividing both sides by μ , and rearranging terms we get

$$(x_t - \hat{x}_t)^T(x_t - x') \geq \frac{1}{2} \|\hat{x}_t - x_t\|^2 + \frac{\nu_h}{2\mu} \|x' - x_t\|^2. \quad (24)$$

Recall that $x_{t+1} = (1-\alpha)x_t + \alpha\hat{x}_t$. Using this, $\|x_{t+1} - x'\|^2 = \|x_t - x'\|^2 + \alpha^2 \|x_t - \hat{x}_t\|^2 - 2\alpha(x_t - x')^T(x_t - \hat{x}_t)$.

Using the bound on $(x_t - x')^T(x_t - \hat{x}_t)$ obtained in (24), we get

$$\|x_{t+1} - x'\|^2 \leq \left(1 - \frac{\alpha\nu_h}{\mu}\right) \|x_t - x'\|^2 + \alpha(\alpha-1) \|x_t - \hat{x}_t\|^2.$$

Since $\alpha \in (0, 1]$, the second term in RHS is non-positive. Thus, we get

$$\|x_{t+1} - x'\|^2 \leq \left(1 - \frac{\alpha\nu_h}{\mu}\right) \|x_t - x'\|^2, \quad (25)$$

as required.