

A Distributed Online Convex Optimization Algorithm with Improved Dynamic Regret

Yan Zhang, Robert J. Ravier, Michael M. Zavlanos, and Vahid Tarokh

Abstract—In this paper, we consider the problem of distributed online convex optimization, where a network of local agents aim to jointly optimize a convex function over a period of multiple time steps. The agents do not have any information about the future. Existing algorithms have established dynamic regret bounds that have explicit dependence on the number of time steps. In this work, we show that we can remove this dependence assuming that the local objective functions are strongly convex. More precisely, we propose a gradient tracking algorithm where agents jointly communicate and descend based on corrected gradient steps. We verify our theoretical results through numerical experiments.

I. INTRODUCTION

Distributed optimization has recently received considerable attention, particularly due to its wide applicability in the areas of control and learning [1–4]. The goal is to decompose large optimization problems into smaller, more manageable subproblems that are solved iteratively and in parallel by a group of communicating agents. As such, distributed algorithms avoid the cost and fragility associated with centralized coordination, and provide better privacy for the autonomous decision makers.

In this paper, we are concerned with distributed online convex optimization (OCO) problems, where the objective function can vary with time and the goal is to minimize a notion of *regret*. Most effort in OCO has been devoted to analyzing algorithms using *static* regret, which compares the computed candidate optima with the best fixed action in hindsight. It is well known that the static regret of distributed OCO algorithms grows sublinearly with the iterations; see [5–9] for examples. Our focus is on *dynamic* regret, which compares the output of the algorithm at each time with the current optimal value at that time. Dynamic regret is of more interest when the goal is to track a sequence of varying optimal solutions, as in distributed tracking of moving targets. It is well known that, in general, sublinear dynamic regret cannot be achieved even in the centralized case; bounds on dynamic regret are related to problem regularities, which are measures of how the individual optimization problems change over time [10–12]. The distributed problem has been analyzed in [13,14]. The work in [13] makes use

of an online version of mirror descent, while [14] uses a distributed variant of the celebrated ADAM algorithm [15]. In both cases, the bounds achieved depend on the path-length regularity as well as the total number of time steps.

The dependency of dynamic regret on the number of time steps was removed in [16] by assuming the objective function to be strongly convex, and then in [12] under weaker conditions than strong convexity, both in the centralized case. The key idea is that with sufficient assumptions, the standard offline algorithms can be shown to exhibit linear convergence provided a small enough step size is used. In this paper, we show that analogous results can hold in the distributed setting by employing the gradient tracking technique in [17]. Our main result shows that strong convexity allows us to eliminate the explicit dependence on time in the regret bound. The algorithm that we propose is an online modification of those featured in [17,18], which themselves are corrected versions of decentralized gradient descent [19]. In the process, we use another regularity in OCO for variations in the gradient.

The remainder of the work is organized as follows. In Section II, we formulate the problem and present assumptions necessary for the subsequent analysis. In Section III, we present our proposed algorithm and analyze its dynamic regret bounds. In Section IV we perform numerical experiments to evaluate the quality of our bounds and algorithm. We make concluding remarks in Section V.

II. REVIEW AND PRELIMINARIES

We briefly review necessary background for the distributed OCO problem at hand. At each time t , we are interested in optimizing a differentiable convex function $f_t(x) := \sum_{i=1}^n f_{i,t}(x)$, where n is the size of the network and $\{f_{i,t}\}$ are differentiable convex functions corresponding to each agent i in a connected network. Let each agent keep a local estimate $x_{i,t}$ of the global optimizer x_t^* at t . At time $t = 0$, the agents are initialized with starting points $x_{i,0}$ and the gradients $\nabla f_{i,0}(x_{i,0})$ are observed locally. Though many notions of regret exist, there are two notions that have historically been most popular in the OCO literature. The static regret is a measurement of performance with respect to the best fixed decision in hindsight and is given in the distributed setting by

$$R_T^s := \frac{1}{n} \sum_{i=1}^n \sum_{t=0}^T f_t(x_{i,t}) - \min_x \sum_{t=1}^T f_t(x). \quad (1)$$

The focus of our paper is on dynamic regret. In this measure, the performance of the algorithm is measured with respect

Yan Zhang and Michael M. Zavlanos are with the Department of Mechanical Engineering and Materials Science, Duke University, Durham, NC 27708, USA. {yan.zhang2, michael.zavlanos}@duke.edu Robert Ravier and Vahid Tarokh are with the Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708, USA. {robert.ravier, vahid.tarokh}@duke.edu This work is supported in part by AFOSR under award # FA9550-19-1-0169 and by DARPA under grant # FA8650-18-1-7837.

to the best decision at each time. Formally, it is defined as

$$R_T^d := \frac{1}{n} \sum_{i=1}^n \sum_{t=0}^T f_t(x_{i,t}) - \sum_{t=1}^T \min_{x_t} f_t(x_t). \quad (2)$$

Unlike the static regret R_T^s that can be shown to grow sublinearly with T , the dynamic regret R_T^d depends on problem regularities. The most common regularity is the optimizer path-length, which is the sum of the distance between the individual optimal points at each time, i.e.,

$$\mathcal{P}_T := \sum_{t=0}^{T-1} \|x_{t+1}^* - x_t^*\|, \quad (3)$$

where the norm is in the L_2 -sense. The distributed online algorithms in [13,14] both have regret that is asymptotically $O(\sqrt{T}\mathcal{P}_T)$. In this paper, we will also need what we call the gradient path-length. Let $g_t(x_t) = [\nabla f_{1,t}(x_{1,t}), \dots, \nabla f_{n,t}(x_{n,t})]^T$. Then, the gradient path-length is given by

$$\mathcal{V}_T := \sum_{t=0}^{T-1} \|g_{t+1} - g_t\|_\infty, \quad (4)$$

where $\|\cdot\|_\infty$ denotes the L_∞ norm. A similar squared version of \mathcal{V}_T known as the gradient variation has appeared in regret bounds before in [11,20]. We also make the following assumptions on the local objective function $f_{i,t}$ at each time.

Assumption II.1. For all i and t , the function $f_{i,t}$ is L_f -Lipschitz, that is, there exists a constant $L_f > 0$ such that

$$|f_{i,t}(x) - f_{i,t}(y)| \leq L_f \|x - y\|, \text{ for all } x, y.$$

Assumption II.2. For all i and t , the function $f_{i,t}$ is L_g -smooth, i.e. there exists a constant $L_g > 0$ such that

$$\|\nabla f_{i,t}(x) - \nabla f_{i,t}(y)\| \leq L_g \|x - y\|, \text{ for all } x, y.$$

As previously discussed, we will also make use of the additional assumption of strong convexity on $f_t(x)$.

Assumption II.3. For all t , the global objective function f_t is μ -strongly convex.

Strong convexity is an assumption used in offline convex optimization to prove linear convergence of gradient descent. This is expressed formally in the following result shown in [18], which we state without proof.

Lemma II.4. Let $F(x)$ be μ -strongly convex and its gradient be β -Lipschitz continuous. Moreover, let $x_0 \in \mathbb{R}^d$ and $0 < \eta < \frac{2}{\beta}$, and define $x_1 := x_0 - \eta \nabla F(x)$. Then, for $\lambda := \max(|1 - \eta\mu|, |1 - \eta\beta|) < 1$, we have $\|x_1 - x^*\| \leq \lambda \|x_0 - x^*\|$.

Furthermore, we introduce notations and an assumption on the network structure. Denote the graph as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the agent index set $\{1, \dots, n\}$ and \mathcal{E} is the edge set. $(i, j) \in \mathcal{E}$ if agent i can receive information from agent j . We define a weight matrix W that encodes the topology of graph \mathcal{G} . Specifically, if $W_{ij} \neq 0$, edge $(i, j) \in \mathcal{E}$ and

Algorithm 1: Distributed Online Optimization with Gradient Tracking

Require: The primal variable $x_{i,0}$, the local gradient $\nabla f_{i,0}(x_{i,0})$ and global gradient estimate $y_{i,0} = \nabla f_{i,0}(x_{i,0})$ for all i , and the doubly stochastic matrix $W := (w_{i,j})$.

1: **for** $1 \leq t \leq T$ **do**

2: For all $i = 1, \dots, n$, agent i computes

$$x_{i,t+1} = \sum_{j \in N_i} W_{ij}(x_{j,t} - \alpha_j y_{j,t}). \quad (5a)$$

3: For all $i = 1, \dots, n$, agent i computes

$$y_{i,t+1} = \sum_{j \in N_i} W_{ij} y_{j,t} + \nabla f_{i,t+1}(x_{i,t+1}) - \nabla f_{i,t}(x_{i,t}). \quad (5b)$$

4: **end for**

vice versa. We make the following assumption on graph \mathcal{G} and matrix W .

Assumption II.5. The graph \mathcal{G} is undirected and connected. Furthermore, W is doubly stochastic and symmetric. That is, $W\mathbf{1} = \mathbf{1}$ and $W = W^T$.

A direct consequence of Assumption II.5 is that the L_2 matrix norm $\|W - \frac{1}{n}\mathbf{1}\mathbf{1}^T\| = \sigma_W < 1$, where σ_W is also called the mixing rate. When σ_W is small, the network reaches consensus faster, see [21].

III. ALGORITHM AND PERFORMANCE ANALYSIS

Our proposed algorithm is detailed in Algorithm 1. Without loss of generality, we suppose that the iterates in Algorithm 1 satisfy $x_{i,t}, y_{i,t} \in \mathbb{R}$. Let $\alpha = \text{diag}(\alpha_1, \dots, \alpha_n)$. The update in (5) can be written in a compact form

$$x_{t+1} = W(x_t - \alpha y_t), \quad (6a)$$

$$y_{t+1} = W y_t + g_{t+1}(x_{t+1}) - g_t(x_t). \quad (6b)$$

The updates (6a) and (6b) are an online analogue to the algorithm proposed in [17]. We start by assuming that the agents have access to the local gradient at the present. The algorithm proceeds by alternating between descent and gradient computation steps, simultaneously communicating with other agents. Rather than performing standard distributed gradient descent, we employ a correction that corrects for the difference in the gradient between the current points and those in the immediate past. This technique is shown in [17] to improve the convergence rates for the offline problem. Denote by $\bar{x}_t = \frac{1}{n} \sum_i x_{i,t}$, $\bar{y}_t = \frac{1}{n} \sum_i y_{i,t}$ the average values of the primal variables and the gradients across the network. We have the following lemma.

Lemma III.1. Let Assumption II.1 hold. Then, the regret R_T^d is upper bounded by

$$R_T^d \leq \sqrt{n} L_f \sum_{t=0}^T \|x_t - \mathbf{1} \bar{x}_t\| + n L_f \sum_{t=0}^T \|\bar{x}_t - x_t^*\|.$$

Proof. Given Assumption II.1, for all i and t , we have that $f_t(x_{i,t}) - f_t(x_t^*) \leq nL_f \|x_{i,t} - x_t^*\|$. Summing both sides over i and t and dividing by n , we have that

$$R_T^d \leq L_f \sum_{t=0}^T \sum_{i=1}^n \|x_{i,t} - x_t^*\|. \quad (7)$$

Adding and subtracting \bar{x}_t inside the norm in (7) and applying the triangle inequality, we have that

$$\begin{aligned} R_T^d &\leq L_f \sum_{t=0}^T \sum_{i=1}^n \|x_{i,t} - \bar{x}_t\| + nL_f \sum_{t=0}^T \|\bar{x}_t - x_t^*\| \\ &\leq \sqrt{n}L_f \sum_{t=0}^T \|x_t - \mathbf{1}\bar{x}_t\| + nL_f \sum_{t=0}^T \|\bar{x}_t - x_t^*\|. \end{aligned}$$

The second inequality is due to the inequality of the root mean square and the arithmetic mean. The proof is complete. \square

From Lemma III.1, we observe that the regret is upper bounded by the network error $\sum_{t=0}^T \|x_t - \mathbf{1}\bar{x}_t\|$ and the tracking error $\sum_{t=0}^T \|\bar{x}_t - x_t^*\|$. Next, we shall show that under an appropriate step size, both errors are bounded by the regularity of the problem up to a scaling factor. Before presenting this result, we need the following lemma

Lemma III.2. Let Assumption II.5 hold. Then, for all t , we have $\mathbf{1}^T y_t = \mathbf{1}^T g_t(x_t)$.

Proof. According to the initialization of Algorithm 1, $y_{i,0} = \nabla f_{i,0}(x_{i,0})$, the above equality is satisfied when $t = 0$. Next, we use mathematical induction: Assuming that $\mathbf{1}^T y_t = \mathbf{1}^T g_t(x_t)$, we need to show $\mathbf{1}^T y_{t+1} = \mathbf{1}^T g_{t+1}(x_{t+1})$. Recalling the update in (6b), we have that $\mathbf{1}^T y_{t+1} = \mathbf{1}^T W y_t + \mathbf{1}^T g_{t+1}(x_{t+1}) - \mathbf{1}^T g_t(x_t) = \mathbf{1}^T y_t + \mathbf{1}^T g_{t+1}(x_{t+1}) - \mathbf{1}^T g_t(x_t) = \mathbf{1}^T g_{t+1}(x_{t+1})$. The proof is complete. \square

Lemma III.3. Let Assumptions II.2, II.3 and II.5 hold. Then, there exist small enough step sizes $[\alpha_1, \dots, \alpha_n]$ such that $\sum_{t=0}^T \|\bar{x}_t - x_t^*\|$, $\sum_{t=0}^T \|x_t - \mathbf{1}\bar{x}_t\|$ and $\sum_{t=0}^T \|y_t - \mathbf{1}\bar{y}_t\|$ are all of order $O(C_1 + C_2 + C_3 + \mathcal{P}_T + \mathcal{V}_T)$, where $C_1 = \|\bar{x}_0 - x_0^*\|$, $C_2 = \|x_0 - \mathbf{1}\bar{x}_0\|$ and $C_3 = \|y_0 - \mathbf{1}\bar{y}_0\|$. Furthermore, if $\alpha_1 = \dots = \alpha_n$, then the step sizes can be chosen as $\alpha_i < \frac{1 - \sigma_W^2}{1 - \sigma_W^2 + 2\sigma_W} \frac{1}{1 + n \frac{L_g}{\mu}} \frac{1}{L_g}$, for all i .

Proof. The first step is to show the contraction of $\|\bar{x}_t - x_t^*\|$, $\|x_t - \mathbf{1}\bar{x}_t\|$ and $\|y_t - \mathbf{1}\bar{y}_t\|$ over time. According to (6a), we have that

$$\begin{aligned} \|\bar{x}_{t+1} - x_{t+1}^*\| &= \left\| \frac{1}{n} \mathbf{1}^T x_{t+1} - x_{t+1}^* \right\| \\ &= \left\| \frac{1}{n} \mathbf{1}^T W(x_t - \alpha y_t) - x_t^* + x_t^* - x_{t+1}^* \right\| \quad (8) \\ &\leq \|\bar{x}_t - \frac{1}{n} \mathbf{1}^T \alpha y_t - x_t^*\| + \|x_t^* - x_{t+1}^*\|. \end{aligned}$$

Denote $\bar{y}_t = \frac{1}{n} \mathbf{1}^T y_t$, $\bar{g}_t = \frac{1}{n} \mathbf{1}^T g_t(\mathbf{1}\bar{x}_t)$ and $\alpha' = \frac{1}{n} \mathbf{1}^T \alpha \mathbf{1}$. Adding and subtracting $\alpha' \bar{g}_t$ and $\alpha' \bar{y}_t$ in $\|\bar{x}_t - \frac{1}{n} \mathbf{1}^T \alpha y_t - x_t^*\|$

in (8) and using the triangle inequality, we have that

$$\begin{aligned} \|\bar{x}_t - \frac{1}{n} \mathbf{1}^T \alpha y_t - x_t^*\| &\leq \|\bar{x}_t - \alpha' \bar{g}_t - x_t^*\| \\ &\quad + \alpha' \|\bar{g}_t - \bar{y}_t\| + \frac{\|\mathbf{1}^T \alpha\|}{n} \|y_t - \mathbf{1}\bar{y}_t\|. \end{aligned} \quad (9)$$

According to Lemma 10 in [18] and the μ -strong convexity of $f_t(x)$, we have that $\|\bar{x}_t - \alpha' \bar{g}_t - x_t^*\| \leq (1 - \frac{1}{n} \alpha' \mu) \|\bar{x}_t - x_t^*\|$ if the step size satisfies $\frac{1}{n} \alpha' \in (0, \frac{1}{n L_g})$. In addition, due to Lemma III.2 and Assumption II.2, we have that $\|\bar{g}_t - \bar{y}_t\| = \frac{1}{n} \|\mathbf{1}^T g_t(\mathbf{1}\bar{x}_t) - \mathbf{1}^T g_t(x_t)\| \leq \frac{L_g}{\sqrt{n}} \|x_t - \mathbf{1}\bar{x}_t\|$. Combining these inequalities with (8) and (9), we have that $\|\bar{x}_{t+1} - x_{t+1}^*\| \leq (1 - \frac{1}{n} \alpha' \mu) \|\bar{x}_t - x_t^*\| + \alpha' \frac{L_g}{\sqrt{n}} \|x_t - \mathbf{1}\bar{x}_t\| + \frac{\|\mathbf{1}^T \alpha\|}{n} \|y_t - \mathbf{1}\bar{y}_t\| + \|x_t^* - x_{t+1}^*\|$. Adding both sides of the above inequality from $t = 0$ to $t = T - 1$, adding $\|\bar{x}_0 - x_0^*\|$ to both sides and adding $(1 - \frac{1}{n} \alpha' \mu) \|\bar{x}_T - x_T^*\|$ on the right hand side, and rearranging the terms, we have that

$$\begin{aligned} \sum_{t=0}^T \|\bar{x}_t - x_t^*\| &\leq \frac{C_1}{\beta_1} + \frac{\sqrt{n}L_g}{\mu} \sum_{t=0}^T \|x_t - \mathbf{1}\bar{x}_t\| \\ &\quad + \frac{\|\mathbf{1}^T \alpha\|}{\alpha' \mu} \sum_{t=0}^T \|y_t - \mathbf{1}\bar{y}_t\| + \frac{n}{\alpha' \mu} \sum_{t=1}^T \|x_t^* - x_{t-1}^*\|, \end{aligned} \quad (10)$$

where $C_1 = \|\bar{x}_0 - x_0^*\|$, $\beta_1 = \frac{1}{n} \alpha' \mu$. Next, we bound the term $\sum_{t=0}^T \|x_t - \mathbf{1}\bar{x}_t\|$. According to (6a), we have that

$$\begin{aligned} \|x_{t+1} - \mathbf{1}\bar{x}_{t+1}\| &= \|W(x_t - \alpha y_t) - \frac{1}{n} \mathbf{1} \mathbf{1}^T W(x_t - \alpha y_t)\| \\ &\leq \|W - \frac{\mathbf{1} \mathbf{1}^T}{n}\| \|x_t - \mathbf{1}\bar{x}_t\| + \|W - \frac{\mathbf{1} \mathbf{1}^T}{n}\| \|\alpha\| \|y_t\|. \end{aligned} \quad (11)$$

Furthermore, we get that

$$\begin{aligned} \|y_t\| &\leq \|y_t - \mathbf{1}\bar{y}_t\| + \|\mathbf{1}\| \|\bar{y}_t\| \\ &\leq \|y_t - \mathbf{1}\bar{y}_t\| + \|\mathbf{1}\| \|\bar{y}_t - \bar{g}_t\| + \|\mathbf{1}\| \|\bar{g}_t - \frac{1}{n} \mathbf{1}^T g_t(\mathbf{1}x_t^*)\| \\ &\leq \|y_t - \mathbf{1}\bar{y}_t\| + L_g \|x_t - \mathbf{1}\bar{x}_t\| + \sqrt{n}L_g \|\bar{x}_t - x_t^*\|. \end{aligned} \quad (12)$$

The second inequality above is due to the fact that x_t^* is the optimizer at time t and $\mathbf{1}^T g_t(\mathbf{1}x_t^*) = 0$. The last inequality is due to the same reasoning following (9). Recalling that $\|W - \frac{\mathbf{1} \mathbf{1}^T}{n}\| = \sigma_W < 1$, and combining the inequalities (11) and (12), we have that

$$\begin{aligned} \|x_{t+1} - \mathbf{1}\bar{x}_{t+1}\| &\leq \sigma_W (1 + L_g \|\alpha\|) \|x_t - \mathbf{1}\bar{x}_t\| \\ &\quad + \sigma_W \|\alpha\| \|y_t - \mathbf{1}\bar{y}_t\| + \sigma_W \sqrt{n}L_g \|\alpha\| \|\bar{x}_t - x_t^*\|. \end{aligned} \quad (13)$$

Next, we bound the network error $\sum_{t=0}^T \|y_t - \mathbf{1}\bar{y}_t\|$. According to (6b), we have that $\|y_{t+1} - \mathbf{1}\bar{y}_{t+1}\| = \|(W - \frac{1}{n} \mathbf{1} \mathbf{1}^T)(y_t - \mathbf{1}\bar{y}_t) + (I - \frac{1}{n} \mathbf{1} \mathbf{1}^T)(g_{t+1}(x_{t+1}) - g_t(x_t))\| \leq \|W - \frac{1}{n} \mathbf{1} \mathbf{1}^T\| \|y_t - \mathbf{1}\bar{y}_t\| + \|I - \frac{1}{n} \mathbf{1} \mathbf{1}^T\| \|g_{t+1}(x_{t+1}) - g_t(x_t)\|$, where the last inequality is due to the triangle and the Cauchy-Schwartz inequalities. Adding and subtracting $g_{t+1}(x_t)$ in the norm $\|g_{t+1}(x_{t+1}) - g_t(x_t)\|$ and using the triangle inequality, we have that

$$\begin{aligned} \|y_{t+1} - \mathbf{1}\bar{y}_{t+1}\| &\leq \sigma_W \|y_t - \mathbf{1}\bar{y}_t\| + \|g_{t+1}(x_{t+1}) \\ &\quad - g_{t+1}(x_t)\| + \|g_{t+1}(x_t) - g_t(x_t)\| \leq \sigma_W \|y_t - \mathbf{1}\bar{y}_t\| \\ &\quad + L_g \|x_{t+1} - x_t\| + \|g_{t+1}(x_t) - g_t(x_t)\|. \end{aligned} \quad (14)$$

According to (6a), we have that $\|x_{t+1} - x_t\| = \|W(x_t - \alpha y_t) - x_t\| = \|(W - I)x_t - W\alpha y_t\| \leq \|W - I\| \|x_t - \mathbf{1}\bar{x}_t\| + \|W\| \|\alpha\| \|y_t\|$. Combining this inequality with (12) and (14), and recalling that $\|W\| = 1$, we have that

$$\begin{aligned} \|y_{t+1} - \mathbf{1}\bar{y}_{t+1}\| &\leq (\sigma_W + L_g \|\alpha\|) \|y_t - \mathbf{1}\bar{y}_t\| \\ &\quad + L_g (\|W - I\| + L_g \|\alpha\|) \|x_t - \mathbf{1}\bar{x}_t\| \\ &\quad + \sqrt{n} L_g^2 \|\alpha\| \|\bar{x}_t - x_t^*\| + \|g_{t+1}(x_t) - g_t(x_t)\|. \end{aligned} \quad (15)$$

Next, we telescope the inequalities (13) and (15) and rearrange terms as how we get (10). We can obtain

$$\begin{aligned} \sum_{t=0}^T \|x_t - \mathbf{1}\bar{x}_t\| &\leq \frac{C_2}{\beta_2} + \frac{\sigma_W \|\alpha\|}{\beta_2} \sum_{t=0}^T \|y_t - \mathbf{1}\bar{y}_t\| \\ &\quad + \frac{\sigma_W \sqrt{n} L_g \|\alpha\|}{\beta_2} \sum_{t=0}^T \|\bar{x}_t - x_t^*\|, \end{aligned} \quad (16)$$

where $C_2 = \|x_0 - \mathbf{1}\bar{x}_0\|$, $\beta_2 = 1 - \sigma_W(1 + L_g \|\alpha\|)$, and

$$\begin{aligned} \sum_{t=0}^T \|y_t - \mathbf{1}\bar{y}_t\| &\leq \frac{C_3}{\beta_3} + \frac{L_g (\|W - I\| + L_g \|\alpha\|)}{\beta_3} \\ &\quad \sum_{t=0}^T \|x_t - \mathbf{1}\bar{x}_t\| + \frac{\sqrt{n} L_g^2 \|\alpha\|}{\beta_3} \sum_{t=0}^T \|\bar{x}_t - x_t^*\| \\ &\quad + \frac{1}{\beta_3} \sum_{t=1}^T \|g_t(x_{t-1}) - g_{t-1}(x_{t-1})\|, \end{aligned} \quad (17)$$

where $C_3 = \|y_0 - \mathbf{1}\bar{y}_0\|$, $\beta_3 = 1 - \sigma_W - L_g \|\alpha\|$.

Next, we derive the upper bound on the network errors $\sum_{t=0}^T \|x_t - \mathbf{1}\bar{x}_t\|$ and $\sum_{t=0}^T \|y_t - \mathbf{1}\bar{y}_t\|$ using inequalities (10), (16) and (17). To achieve this, we scale (16) and (17) by positive factors M and N and add up the three inequalities. After rearranging terms, the desired upper bound in Lemma III.3 is derived if the following conditions are satisfied: the term $\sum_{t=0}^T \|\bar{x}_t - x_t^*\|$ on both sides is cancelled, the terms $\sum_{t=0}^T \|x_t - \mathbf{1}\bar{x}_t\|$ and $\sum_{t=0}^T \|y_t - \mathbf{1}\bar{y}_t\|$ appear on the left hand side and the regularity terms $\sum_{t=1}^T \|x_t^* - x_{t-1}^*\|$ and $\sum_{t=1}^T \|g_t(x_{t-1}) - g_{t-1}(x_{t-1})\|$ appear on the right hand side. In the following, we show that there always exists step size α so that positive scaling factors M and N can be found to satisfy these conditions. To do so, we let M and N satisfy the following conditions:

$$M \frac{\sigma_W \sqrt{n} L_g \|\alpha\|}{\beta_2} + N \frac{\sqrt{n} L_g^2 \|\alpha\|}{\beta_3} = 1, \quad (18a)$$

$$M > \frac{\sqrt{n} L_g}{\mu} + N \frac{L_g (\|W - I\| + L_g \|\alpha\|)}{\beta_3}, \quad (18b)$$

$$N > \frac{\|\mathbf{1}^T \alpha\|}{\alpha' \mu} + M \frac{\sigma_W \|\alpha\|}{\beta_2}. \quad (18c)$$

Moreover, we define $M = \frac{\beta_2}{\sigma_W \sqrt{n} L_g \|\alpha\|} b$ and $N = \frac{\beta_3}{\sqrt{n} L_g^2 \|\alpha\|} (1 - b)$, where $b \in (0, 1)$. Then, the condition (18a)

is automatically satisfied. Plugging the expressions of M and N into inequalities (18b) and (18c), we have that

$$b < \frac{\beta_3 - \sqrt{n} L_g^2 \frac{\|\mathbf{1}^T \alpha\|}{\alpha' \mu} \|\alpha\|}{\beta_3 + L_g \|\alpha\|} < 1, \quad (19a)$$

$$b > \frac{(n \frac{L_g^2}{\mu} + L_g) \|\alpha\| + \|W - I\|}{L_g \|\alpha\| + \|W - I\| + \frac{\beta_2}{\sigma_W}} > 0. \quad (19b)$$

As long as the interval of b defined in (19) is nonempty, we can use any b in this interval to find the corresponding M and N that satisfy the conditions in (18). When $\|\alpha\|$ goes to zero, the upper bound of b in (19a) goes to 1 and the lower bound of b in (19b) goes to $\frac{\|W - I\|}{\|W - I\| + \frac{1 - \sigma_W}{\sigma_W}} < 1$. Therefore, there always exists a step size α small enough so that b , M and N exist to derive the desired bound on $\sum_{t=0}^T \|x_t - \mathbf{1}\bar{x}_t\|$ and $\sum_{t=0}^T \|y_t - \mathbf{1}\bar{y}_t\|$ in Lemma III.3. According to the inequality (10), using the bound on $\sum_{t=0}^T \|x_t - \mathbf{1}\bar{x}_t\|$ and $\sum_{t=0}^T \|y_t - \mathbf{1}\bar{y}_t\|$, it is simple to see that $\sum_{t=0}^T \|\bar{x}_t - x_t^*\|$ is also bounded by the same order.

In the remainder of this proof, we show that if $\alpha_1 = \alpha_2 = \dots = \alpha_n$ and $\alpha_i < \frac{1 - \sigma_W^2}{1 - \sigma_W^2 + 2\sigma_W} \frac{1}{1 + n \frac{L_g}{\mu}} \frac{1}{L_g}$ for all i , then α satisfies the sufficient decrease requirement in Lemma II.4, the contraction requirement in (13) and (15), that is, $\beta_2, \beta_3 > 0$, and the requirement that the interval of b is nonempty in (19). To see this, let $\alpha_1 = \dots = \alpha_n = a$. First, we need to choose a such that $\beta_2, \beta_3 > 0$. Plugging a into the expression of β_2 and β_3 , we have that $a < \frac{1 - \sigma_W}{\sigma_W} \frac{1 - \sigma_W}{L_g}$ and $a < \frac{1 - \sigma_W}{L_g}$. Since $\sigma_W < 1$, it is sufficient to choose

$$a < \frac{1 - \sigma_W}{L_g} \quad (20a)$$

to obtain $\beta_2, \beta_3 > 0$. Moreover, we need to choose a so that the upper bound in (19a) is strictly greater than (19b). Because $\alpha_1 = \dots = \alpha_n = a$, we have that $\|\alpha\| = a$ and $\alpha' = \frac{1}{n} \mathbf{1}^T \alpha \mathbf{1} = a$. Plugging these into the bounds in (19), we obtain that

$$\frac{(n \frac{L_g^2}{\mu} + L_g) a + \|W - I\|}{L_g a + \frac{\beta_2}{\sigma_W} + \|W - I\|} < \frac{\beta_3 - n \frac{L_g^2}{\mu} a}{\beta_3 + L_g a}.$$

By Assumption II.5, we have that all eigenvalues of W satisfy $|\lambda(W)| \leq 1$ due to Perron-Frobenius theory. Therefore, $\|W - I\| \leq 2$. Therefore, to satisfy the above inequality, it is sufficient to let

$$\frac{(n \frac{L_g^2}{\mu} + L_g) a + 2}{L_g a + \frac{\beta_2}{\sigma_W} + 2} < \frac{\beta_3 - n \frac{L_g^2}{\mu} a}{\beta_3 + L_g a} \quad \text{and} \quad \beta_3 - n \frac{L_g^2}{\mu} a > 0,$$

due to the inequality $\frac{p+m}{q+m} > \frac{p}{q}$ for any $p > q > 0$ and $m > 0$. Simplifying the above two inequalities, we obtain the following conditions on a

$$a < \frac{1 - \sigma_W^2}{1 - \sigma_W^2 + 2\sigma_W} \frac{1}{1 + n \frac{L_g}{\mu}} \frac{1}{L_g}, \quad (20b)$$

$$a < (1 - \sigma_W) \frac{1}{1 + n \frac{L_g}{\mu}} \frac{1}{L_g}. \quad (20c)$$

In addition, recalling the analysis after (9), we have that

$$a < \frac{1}{L_g}. \quad (20d)$$

It is simple to see that all the step size conditions in (20) reduce to the condition (20b). The proof is complete. \square

Remark III.4. Compared to existing methods in distributed online optimization, e.g., [13], here we remove the dependency of the step size on the horizon T . According to the upper bound on the stepsize in Lemma III.3, when the network size n becomes larger, or the mixing rate σ_W is closer to 1, (that is, the mixing process is slower,) or $\frac{nL_g}{\mu}$ becomes larger, (that is, the central problem is more ill-conditioned), the step size needs to be more conservative.

Theorem III.5. Let Assumptions II.1-II.5 hold. Then, there exist small enough stepsizes $[\alpha_1, \dots, \alpha_n]$ such that the regret R_T^d is of the order $O(C_1 + C_2 + C_3 + \mathcal{P}_T + \mathcal{V}_T)$, where $C_1 = \|\bar{x}_0 - x_0^*\|$, $C_2 = \|x_0 - \mathbf{1}\bar{x}_0\|$ and $C_3 = \|y_0 - \mathbf{1}\bar{y}_0\|$. Furthermore, if $\alpha_1 = \dots = \alpha_n$, then the step sizes can be chosen as $\alpha_i < \frac{1 - \sigma_W^2}{1 - \sigma_W^2 + 2\sigma_W} \frac{1}{1 + n \frac{L_g}{\mu}} \frac{1}{L_g}$, for all i .

Proof. This theorem is a direct result of combining Lemma III.1 and III.3. \square

We conclude this section by remarking on our assumption that W is doubly stochastic. In [17], the matrix W in (5a) and (5b) is replaced by a row and column stochastic matrix respectively. It is easy to see that our analysis can be adapted to reach a similar conclusion, noting that additional terms should be involved in our bounds related to norm equivalence. We do not pursue this further.

IV. NUMERICAL SIMULATION

In this section, we validate our theoretical analysis using a distributed tracking example. Furthermore, we compare Algorithm 1 to the distributed online gradient descent algorithm in [13] which does not utilize the gradient tracking technique. Consider a sensor network of size $n = 10$. These sensors collaborate to track three moving targets. Every target i has a state variable of dimension 2, i.e. $x_i^*(t) = [x_{i,1}^*(t), x_{i,2}^*(t)]^T$. The position $x_{i,1}^*(t)$ is tracking a sinusoidal curve $x_{i,1}^*(t) = A_i \sin(\omega_i t + \phi_i)$ and the velocity satisfies $x_{i,2}^*(t) = \dot{x}_{i,1}^*(t)$. We stack all target states into one vector $x(t)^* \in \mathbb{R}^6$. Each sensor receives a local observation $y_i(t) \in \mathbb{R}^4$, with each entry of $y_i(t)$ being a linear combination of $x(t)^*$, $y_i(t) = C_i x(t)^*$, where $C_i \in \mathbb{R}^{4 \times 6}$. To estimate the state variable $x(t)^*$, the local optimization problem at each sensor is

$$\min_x f_{i,t}(x) = \frac{1}{2} \|C_i x(t) - y_i(t)\|^2,$$

which is under-determined, and the sensors need to communicate with each other to solve the global optimization problem at each time step

$$\min_x f_t(x) = \frac{1}{2} \|Cx(t) - y(t)\|^2,$$

where $C = [C_1^T | \dots | C_n^T]^T$ and $y(t) = [y_1(t)^T, \dots, y_n(t)^T]^T$. When there are enough measurements, the above

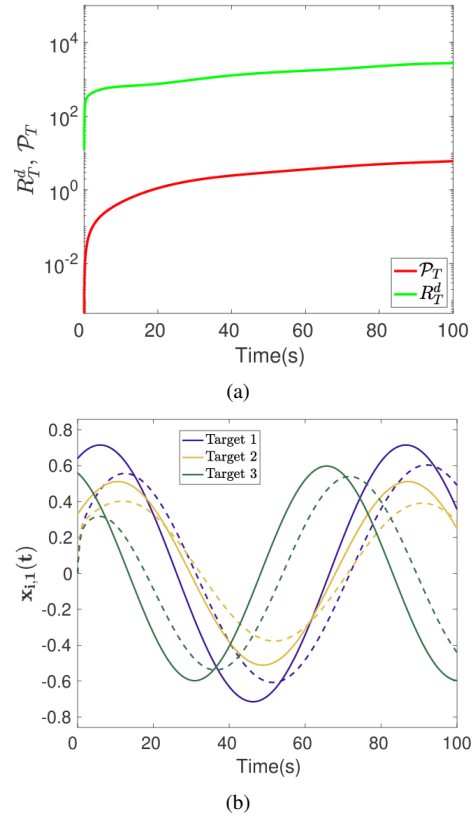


Fig. 1: Distributed tracking using sampling interval 0.01s and step size 0.003 for all agents. (a) shows the regret of the solution and the regularity of the optimizer path-length. (b) shows the tracking performance at agent 2 for three moving targets. The solid curves correspond to the true trajectory of target positions $\{x_{i,1}(t)\}$ over time. The dashed curves correspond to the estimated trajectory solved by Algorithm 1 at agent 2.

central problem is strongly convex. In this specific example, the gradient-path regularity \mathcal{V}_T in (4) is of the same order of \mathcal{P}_T in (3). To see this, we observe that the norm $\|\nabla f_{i,t+1} - \nabla f_{i,t}\|_\infty = \|C_i^T(y(t+1) - y(t))\| \leq \|C_i^T C_i\| \|x(t+1)^* - x(t)^*\|$. Therefore, in the following, we compare the regret directly with \mathcal{P}_T neglecting \mathcal{V}_T .

In our simulation, the target parameters $\{A_i, \omega_i, \phi_i\}$ and observation matrices $\{C_i\}$ are randomly generated. In addition, the doubly stochastic weight matrix W is randomly generated and its mixing rate is $\sigma_W = 0.59$. The stepsizes α_i are selected to be 0.003 according to the bound given in Theorem III.5. Applying Algorithm 1 with this step size and sampling from the target trajectory at frequency 100Hz, the regret of solving the online optimization problem and the tracking performance is presented in Figure 1. From Figure 1(a), we observe that the regret is of the same order as the regularity term \mathcal{P}_T , as we have shown in Theorem III.5, and the local sensors are tracking the moving targets within their neighborhood.

In what follows, we show that Algorithm 1 can admit more aggressive stepsizes than the algorithm without gradient tracking [13]. In this way, the achieved regret and tracking performance can be further improved. Specifically, we select the sampling frequency to be 10Hz, and present the compari-

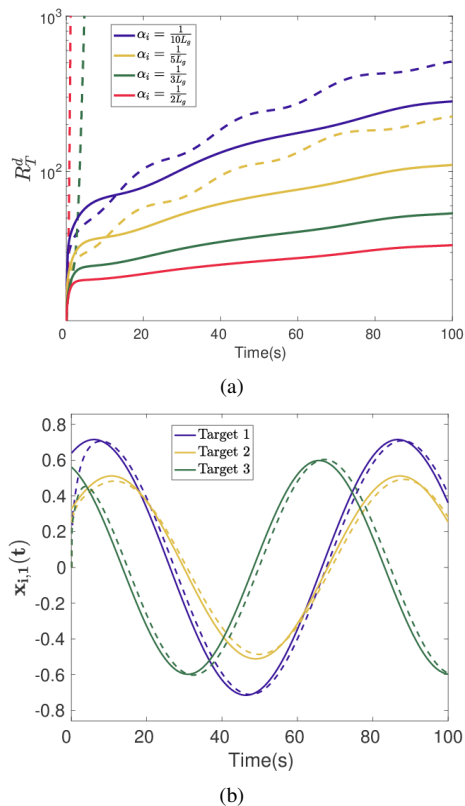


Fig. 2: Distributed tracking using sampling interval 0.1s. (a) shows the comparison between the distributed online optimization with (solid lines) and without gradient tracking (dashed lines) with different step sizes. (b) shows the tracking performance using sampling interval 0.1s and the stepsize $\frac{1}{2L_g}$.

son of the achieved regret between our algorithm and the one without gradient tracking in [13] in Figure 2(a). We observe that the regret of our algorithm always grows slower than the method in [13] with the same choice of stepsize. Moreover, the method in [13] diverges when the stepsizes are chosen to be $\frac{1}{3L_g}$ or $\frac{1}{2L_g}$ while our algorithm achieves lower regrets. In Figure 2(b), we demonstrate that such stepsizes can achieve better tracking performance than the one obtained using the proposed theoretical bound.

V. CONCLUSIONS

In this paper we proposed a new algorithm for distributed online convex optimization based on recent developments in the offline setting. We showed that under the additional assumption of strong convexity and fixed step size, our algorithm can achieve a bound on the dynamic regret without explicit dependence on the number of time steps, in contrast to existing relevant methods. We also illustrated the performance of our algorithm on a distributed tracking example to validate our theoretical results.

REFERENCES

[1] M. Rabbat and R. Nowak, "Distributed optimization in sensor networks," in *Proceedings of the 3rd international symposium on Information processing in sensor networks*. ACM, 2004, pp. 20–27.

[2] S. S. Ram, V. V. Veeravalli, and A. Nedic, "Distributed non-autonomous power control through distributed convex optimization," in *IEEE INFOCOM 2009*. IEEE, 2009, pp. 3001–3005.

[3] S. Shahrampour, A. Rakhlin, and A. Jadbabaie, "Distributed detection: Finite-time analysis and impact of network topology," *IEEE Transactions on Automatic Control*, vol. 61, no. 11, pp. 3256–3268, 2016.

[4] A. Nedić, A. Olshevsky, and C. A. Uribe, "Fast convergence rates for distributed non-bayesian learning," *IEEE Transactions on Automatic Control*, vol. 62, no. 11, pp. 5538–5553, 2017.

[5] 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.

[6] E. Hazan, A. Agarwal, and S. Kale, "Logarithmic regret algorithms for online convex optimization," *Machine Learning*, vol. 69, no. 2–3, pp. 169–192, 2007.

[7] S. Hosseini, A. Chapman, and M. Mesbahi, "Online distributed optimization via dual averaging," in *52nd IEEE Conference on Decision and Control*. IEEE, 2013, pp. 1484–1489.

[8] M. Akbari, B. Ghahsifard, and T. Linder, "Distributed online convex optimization on time-varying directed graphs," *IEEE Transactions on Control of Network Systems*, vol. 4, no. 3, pp. 417–428, 2015.

[9] S. Lee and M. M. Zavlanos, "On the sublinear regret of distributed primal-dual algorithms for online constrained optimization," *arXiv preprint arXiv:1705.11128*, 2017.

[10] E. C. Hall and R. M. Willett, "Online convex optimization in dynamic environments," *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 4, pp. 647–662, 2015.

[11] A. Jadbabaie, A. Rakhlin, S. Shahrampour, and K. Sridharan, "Online optimization: Competing with dynamic comparators," in *Artificial Intelligence and Statistics*, 2015, pp. 398–406.

[12] L. Zhang, T. Yang, J. Yi, J. Rong, and Z.-H. Zhou, "Improved dynamic regret for non-degenerate functions," in *Advances in Neural Information Processing Systems*, 2017, pp. 732–741.

[13] S. Shahrampour and A. Jadbabaie, "Distributed online optimization in dynamic environments using mirror descent," *IEEE Transactions on Automatic Control*, vol. 63, no. 3, pp. 714–725, 2018.

[14] P. Nazari, D. A. Tarzanagh, and G. Michailidis, "Dadam: A consensus-based distributed adaptive gradient method for online optimization," *arXiv preprint arXiv:1901.09109*, 2019.

[15] J. Duchi, E. Hazan, and Y. Singer, "Adaptive subgradient methods for online learning and stochastic optimization," *Journal of Machine Learning Research*, vol. 12, no. Jul, pp. 2121–2159, 2011.

[16] 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.

[17] S. Pu, W. Shi, J. Xu, and A. Nedić, "A push-pull gradient method for distributed optimization in networks," in *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, 2018, pp. 3385–3390.

[18] G. Qu and N. Li, "Harnessing smoothness to accelerate distributed optimization," *IEEE Transactions on Control of Network Systems*, vol. 5, no. 3, pp. 1245–1260, 2018.

[19] A. Nedic and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Transactions on Automatic Control*, vol. 1, no. 54, pp. 48–61, 2009.

[20] C.-K. Chiang, T. Yang, C.-J. Lee, M. Mahdavi, C.-J. Lu, R. Jin, and S. Zhu, "Online optimization with gradual variations," in *Conference on Learning Theory*, 2012, pp. 6–1.

[21] B. Johansson, T. Keviczky, M. Johansson, and K. H. Johansson, "Subgradient methods and consensus algorithms for solving convex optimization problems," in *2008 47th IEEE Conference on Decision and Control*. IEEE, 2008, pp. 4185–4190.