

Improving Dynamic Regret in Distributed Online Mirror Descent Using Primal and Dual Information

Nima Eshraghi NESHRAGHI@ECE.UTORONTO.CA and **Ben Liang** LIANG@ECE.UTORONTO.CA
Department of Electrical and Computer Engineering, University of Toronto, Ontario, Canada

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

We consider the problem of distributed online optimization, with a group of learners connected via a dynamic communication graph. The goal of the learners is to track the global minimizer of a sum of time-varying loss functions in a distributed manner. We propose a novel algorithm, termed Distributed Online Mirror Descent with Multiple Averaging Decision and Gradient Consensus (DOMD-MADGC), which is based on mirror descent but incorporates multiple consensus averaging iterations over local gradients as well as local decisions. The key idea is to allow the local learners to collect a sufficient amount of global information, which enables them to more accurately approximate the time-varying global loss, so that they can closely track the dynamic global minimizer over time. We show that the dynamic regret of DOMD-MADGC is upper bounded by the path length, which is defined as the cumulative distance between successive minimizers. The resulting bound improves upon the bounds of existing distributed online algorithms and removes the explicit dependence on T .

Keywords: Online convex optimization, online learning, distributed optimization, regret analysis

1. Introduction

In recent years, applications have emerged that require extremely large data volumes over many networked machines. As a result, distributed collection and processing of these datasets are not only desirable but often necessary (Boyd et al., 2011). Therefore, distributed optimization has become popular in solving problems that arise in various areas of control and learning (Boyd et al., 2011; Duchi et al., 2012; Xi et al., 2014; Rabbat, 2015).

The aforementioned distributed optimization methods assume a static loss function. Nevertheless, in many practical applications, the system parameters and loss functions vary over time. For example, time-varying loss functions frequently appear in online machine learning, where data samples arrive dynamically, so that newly observed data samples result in new losses. Another example is object-tracking, where the goal is to track the time-varying states of a moving target. These problems can be solved using online optimization algorithms that update the decisions based on the dynamically arriving data.

The performance of online optimization algorithms is usually measured in terms of regret. Depending on the problem environment and settings, different notions of regret have been proposed in the literature. For instance, static regret measures the accumulated loss incurred by the online optimization algorithm against that of an offline optimal solution, which is made in hindsight given the knowledge of all loss functions. The static regret of online optimization algorithms has been

extensively studied in the literature (Cesa-Bianchi and Lugosi, 2006; Shalev-Shwartz, 2012; Hazan, 2019; Orabona, 2019).

Dynamic regret is a more appropriate performance metric when the underlying target of interest is time-varying. Dynamic regret measures the accumulated loss experienced by the online algorithm versus that of a sequence of optimizers, which separately minimize the loss functions at every round. In contrast to the static regret, it is impossible to guarantee a sublinear upper bound on the dynamic regret for every problem instances, due to the arbitrary and potentially severe fluctuations in the problem environment. Instead, the growth of the dynamic regret depends on some regularity measures, which reflect the speed of changes in a dynamic environment. For instance, the *path length* is a popular regularity measure that is related to the rate of changes in the sequence of minimizers over time, i.e.,

$$C_T = \sum_{t=2}^T \|x_t^* - x_{t-1}^*\|, \quad (1)$$

where $x_t^* = \operatorname{argmin}_{x \in \mathcal{X}} f_t(x)$ is the minimizer of the time-varying loss function. When loss functions are convex, the dynamic regret of online convex optimization algorithms can be bounded by $O(\sqrt{T}(C_T + 1))$ (Zinkevich, 2003), which can be improved to $O(\sqrt{T}(C_T + 1))$ when knowledge of the path length C_T and the number of rounds T is available in advance (Yang et al., 2016). Recent studies have shown that the dependency of dynamic regret on the number of online rounds can be removed when stronger assumptions on the function curvature is available (Mokhtari et al., 2016; Zhang et al., 2017). In particular, a dynamic regret bound of $O(1 + C_T)$ has been derived for strongly convex and smooth loss functions (Mokhtari et al., 2016; Zhang et al., 2017), which improves the prior bound of (Zinkevich, 2003) by removing a \sqrt{T} factor.

However, (Mokhtari et al., 2016; Zhang et al., 2017) focus on centralized problems, where there is only a single learner. Distributed optimization can be more challenging since the learners have only partial information about the global problem, which necessitates engagement in communication so that they can complement their insufficient knowledge. None of the previously proposed distributed online convex optimization algorithms (Shahrampour and Jadbabaie, 2018; Dixit et al., 2019; Zhang et al., 2020) achieve $O(1 + C_T)$ dynamic regret.

In this paper, we investigate whether it is possible to obtain the aforementioned dynamic regret bound of $O(1 + C_T)$ in distributed online convex optimization. To this end, we propose a new algorithm termed Distributed Online Mirror Descent with Multiple Averaging Decision and Gradient Consensus (DOMD-MADGC), which is a variant of DOMD with an improved dynamic regret. Previous works on distributed mirror descent (Shahrampour and Jadbabaie, 2018; Xi et al., 2014; Rabbat, 2015) perform a single consensus step only on the local decision vectors, and then each learner tries to minimize its local loss using its local gradient information. The resulting lack of knowledge about the global gradient prevents the learners from closely tracking the time-varying global minimizer. In contrast, DOMD-MADGC allows the communication of the local gradient vectors in addition to the local decisions, leading to reduced global loss. We further observe that simply averaging the local decisions and the local gradients in a single consensus step does not provide the dynamic regret bound that we seek. Therefore, we design DOMD-MADGC to use multiple consensus averaging iterations per online round to collect the local information update from distant learner nodes. Our analysis reveals that a logarithmically increasing number of consensus iterations will result in sufficiently fast decrease of local errors. With the proposed design, DOMD-MADGC achieves a dynamic regret bound of $O(1 + C_T)$, without any prior knowledge of the path length C_T or the number of rounds T .

2. Related Works

Early works on online convex optimization mostly concern static regret. In the centralized setting, it has been established that various optimization methods including online gradient descent (Zinkevich, 2003), online dual averaging (Xiao, 2010), online mirror descent (Duchi et al., 2010), and many others (Shalev-Shwartz, 2012; Hazan, 2019) achieve an upper bound of $O(\sqrt{T})$ and $O(\log T)$ on the static regret, for convex and strongly convex loss functions, respectively. Furthermore, to bound constraint violation and static regret an algorithm based on the Lyapanov technique is developed in (Wei et al., 2020) that uses a primal-dual version of the mirror descent method. The static regret of distributed online convex optimization algorithms has also been extensively studied in the literature (Hosseini et al., 2013; Mateos-Núñez and Cortés, 2014; Akbari et al., 2015; Tsianos and Rabbat, 2016; Lee et al., 2016; Yuan et al., 2020), where the same regret rates have been derived under similar convexity assumptions. However, it is not previously known whether similar results hold for the more useful dynamic regret. That is, whether there exists a distributed online algorithm that has dynamic regret bound comparable to the most competitive centralized online algorithms.

There has been a surge of recent interest in studying dynamic regret in the centralized setting (Hall and Willett, 2013, 2015; Jadbabaie et al., 2015; Yang et al., 2016; Mokhtari et al., 2016; Zhang et al., 2017; Chang and Shahrampour, 2020). Previous works on online gradient descent (Zinkevich, 2003), and online mirror descent (Hall and Willett, 2013, 2015) show that the dynamic regret of both methods is bounded by $O(\sqrt{T}(1 + C_T))$ when the loss functions are convex, which can be tightened to $O(\sqrt{T(1 + C_T)})$ when the knowledge of path length C_T is present a priori (Yang et al., 2016). Furthermore, when the loss functions are strongly convex and smooth, the upper bound on the dynamic regret can be improved to $O(1 + C_T)$ (Mokhtari et al., 2016; Zhang et al., 2017). In another direction, researchers have used various regularity measures to bound the dynamic regret of different online optimization algorithms (Jadbabaie et al., 2015; Besbes et al., 2015; Chang and Shahrampour, 2020; Campolongo and Orabona, 2020, 2021). The above works do not apply to the distributed setting.

Distributed online mirror descent is studied in (Shahrampour and Jadbabaie, 2018), where a dynamic regret bound of $O(\sqrt{T(1 + C_T)})$ is derived when the joint knowledge of the path length C_T and the number of rounds T is available in advance. However, it is generally impossible to know C_T beforehand in practice, so most subsequent works do not require this assumption. A distributed online gradient tracking algorithm is proposed in (Lu et al., 2019), which has a dynamic regret bound of $O(\sqrt{1 + C_T}T^{3/4}\sqrt{\ln T})$. The dynamic regret of distributed online proximal gradient descent with an $O(\log t)$ communication steps per round is bounded by $O(\log T(1 + C_T))$ (Dixit et al., 2019). Different from (Shahrampour and Jadbabaie, 2018), the learners do not receive a new cost function at every round in (Dixit et al., 2019) since every single consensus step is counted as an online round. Using the gradient variation as the regularity measure $V_T = \sum_{t=2}^T \sup_{x \in \mathcal{X}} \|\nabla f_t(x) - \nabla f_{t-1}(x)\|$, a dynamic regret bound of $O(\sqrt{T}(1 + C_T) + V_T + \sqrt{V_T C_T})$ is derived in (Li et al., 2020), which is improved to $O(1 + C_T + V_T')$ in (Zhang et al., 2020), where V_T' is a variant of V_T in which gradients are evaluated at the optimal points. However, the work of (Zhang et al., 2020) studies only unconstrained online optimization problem. Therefore, the results in (Dixit et al., 2019; Zhang et al., 2020) are not directly comparable to ours due to the differences in modeling and problem setup.

In this work, we show that the proposed DOMD-MADGC has a dynamic regret bound of $O(1 + C_T)$. In contrast to most prior works, where the dynamic regret bound depends on a combination

of a sublinear term in T (e.g., $\log T$ or \sqrt{T}) and some regularity measure, our bound depends only on the path length C_T . Furthermore, the resulting bound of $O(1 + C_T)$ improves upon the existing regret bounds in dynamic distributed learning. Unlike (Shahrampour and Jadbabaie, 2018), the dynamic regret bound of our proposed algorithm can be achieved by selecting a step size without the knowledge of the path length C_T or the number of rounds T . Unlike (Zhang et al., 2020), in this work, we allow the optimization variables to be constrained by an arbitrary compact and convex set.

3. Problem Formulation

3.1. Network Model

We consider a distributed learning network in a time-slotted setting. The learner nodes interact with each other over a time-varying communication topology, which is modeled by an undirected graph $\mathcal{G}_t = (\mathcal{V}, \mathcal{E}^t)$, where $\mathcal{V} = \{1, \dots, n\}$ and $\mathcal{E}^t \subset \mathcal{V} \times \mathcal{V}$ denote the set of nodes and edges present at round t , respectively.

Different from online convex optimization where only a single learner is present, in distributed online optimization, every node $i \in \mathcal{V}$ is a local learner that receives a private local loss in every round. Each learner node i is associated with a sequence of time-varying loss functions, i.e., $\{f_{i,1}(x), f_{i,2}(x), \dots, f_{i,T}(x)\}$, where $f_{i,t}$ denotes the scalar-valued loss function of node i at round t , and $x \in \mathcal{X}$ represents a decision variable taken from a compact and convex set \mathcal{X} . At each round t , after each learner i submits a decision $x_{i,t}$, the local loss function $f_{i,t}(x)$ is revealed, and the learner suffers a corresponding loss of $f_{i,t}(x_{i,t})$. Each local loss is only observed by a single learner. Therefore, the learners need to interact with each other, to supplement their incomplete knowledge of the global task.

3.2. Distributed Online Optimization

We are interested in an optimization problem with a global loss function, represented by $f_t(x)$ at round t . It is based on the local loss functions that are distributed over the entire network, i.e.,

$$f_t(x) = \sum_{i=1}^n f_{i,t}(x). \quad (2)$$

The goal of the online learners is to minimize the total loss over a finite number of rounds T . Prior studies on online learning often measure the quality of decisions in terms of *static* regret, defined as the difference between the total loss incurred by the online algorithm and that of an optimal fixed offline solution, which is made in hindsight with knowledge of $f_t(x)$ for all t . The benchmark variable in static regret is a fixed point in the feasible set \mathcal{X} . Such a static metric can accurately reflect the performance of an online algorithm as long as the static benchmark performs consistently well over all rounds. However, this does not always hold in a dynamic environment, where the sequence of loss functions are time-varying. Thus, *dynamic* regret has been proposed as a more stringent metric that incorporates a time-varying comparator sequence. Most commonly, the performance of online algorithms is measured relative to a sequence of minimizers (Shahrampour and Jadbabaie, 2018; Zhang et al., 2020), i.e.,

$$\text{Reg}_T^d = \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T f_t(x_{i,t}) - \sum_{t=1}^T f_t(x_t^*). \quad (3)$$

where $x_t^* = \operatorname{argmin}_{x \in \mathcal{X}} f_t(x)$ is a minimizer of the global loss function at round t .

It is well-known that online optimization problems may be intractable in a dynamic setting, due to arbitrary fluctuation in the loss functions. Therefore, obtaining a sublinear bound on the dynamic regret may be impossible. However, it is possible to bound the dynamic regret in terms of regularity measures that reflect how fast an environment changes as time progresses. A popular quantity to represent such regularity is the path length C_T , defined in (1), which collects the variation between successive minimizers.

Since the local loss function is available only to the local learner, due to partial access to the global information, it is unlikely for each learner to accurately compute the global minimizer x_t^* at each round t . To minimize the dynamic regret, a careful design of the online algorithm is required to allow distributed learners to collect sufficient global information to track the time-varying global minimizer. Thus, it is necessary to use the local information as well as those collected from the neighboring learners, so that each learner can find its local estimates sufficiently close to the global gradient $\nabla f_t(x_t^*)$ and minimizer x_t^* .

3.3. Preliminaries

To be self-contained, we collect here several standard definitions that are used in this paper. We use $\|\cdot\|$ to denote the ℓ -2 norm throughout the paper.

Definition 1: A function $f(x)$ is Lipschitz continuous with factor G if for all x and y in \mathcal{X} , the following holds:

$$|f(x) - f(y)| \leq G\|x - y\|, \forall x, y \in \mathcal{X}.$$

Definition 2: A subdifferentiable function $f(x)$ is β -smooth with respect to some norm $\|\cdot\|$, if there exists a positive constant β such that

$$f(y) \leq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\beta}{2}\|y - x\|^2, \forall x, y \in \mathcal{X}.$$

where $\nabla f(x)$ stands for the subgradient of function $f(x)$.

Definition 3: A subdifferentiable convex function $f(x)$ is λ -strongly convex with respect to some norm $\|\cdot\|$, if there exists a positive constant λ such that

$$f(y) + \langle \nabla f(y), x - y \rangle + \frac{\lambda}{2}\|x - y\|^2 \leq f(x), \forall x, y \in \mathcal{X}.$$

Definition 4: The Bregman divergence with respect to the function $r(x)$ is defined as

$$D_r(x, y) = r(x) - r(y) - \langle \nabla r(y), x - y \rangle.$$

The Bregman divergence is a general distance-measuring function, which contains the Euclidean norm and the Kullback-Leibler divergence as two special cases.

4. Distributed Online Mirror Descent with Multiple Averaging Decision and Gradient Consensus

We now present DOMD-MADGC, a distributed online optimization algorithm that incorporates multiple consensus averaging of local gradients as well local decisions, over time-varying dynamic learning networks. We will bound the performance of DOMD-MADGC in terms of dynamic regret, showing that it achieves the same regret rate as the centralized algorithms of (Mokhtari et al., 2016; Zhang et al., 2017).

4.1. Algorithm Description

DOMD-MADGC uses mirror descent as its optimization engine. The core of the algorithm is the use of consensus averaging over both the local decisions and the local gradients. We note that in prior studies on mirror descent, only the local decisions are used to reach a consensus, while DOMD-MADGC applies consensus averaging to the local gradients as well. We motivate this design choice by comparing it with standard online mirror descent, where each learner performs consensus averaging over the local decision vectors, and then uses its own local gradient to update its decision. Therefore, in standard online mirror descent, while each learner aims to stay close to the averaged decision, it takes a step in a direction to minimize its local loss, which may not align well with the direction toward the global minimizer. In other words, each local learner has access to only one component of the global loss, and the resulting lack of knowledge may prevent the learner from closely tracking the time-varying global target. Therefore, DOMD-MADGC uses both the local gradients and the local decisions in the consensus phase. As a result, the learners can use a local approximation of the global gradient in their decision updates, which leads them to find decisions that reduces the global loss function.

However, simply adding gradient consensus to standard online mirror descent still does not sufficiently reduce the dynamic regret. The success of the algorithm relies on the accuracy of the learners' approximation of the global solution and gradient. In particular, in the commonly used strategy of single consensus averaging, every learner collects the information from immediate neighboring learners only. In a large network with a diverse set of learners, the knowledge of a small group of learners in a local neighborhood may not be sufficient to build an accurate approximation of the global loss, leading to decision update directions that deviate from the global minimizer. This issue can be resolved by applying consensus averaging for multiple times. It is important to have a proper number of consensus iterations to ensure that the errors due to the lack of access to global information are sufficiently small. Yet, an increase in the consensus iterations leads to a larger communication overhead. Therefore, this requires a careful design to balance the need for information sharing and communication overhead. In DOMD-MADGC, we will show that a logarithmic number of consensus iterations in time suffices to achieve $O(1 + C_T)$ dynamic regret.

The detailed procedure of DOMD-MADGC is as follows. At each round t , every learner node i maintains a decision vector $x_{i,t}$, a local estimate of the global minimizer $y_{i,t}$, and a local approximate gradient of the global gradient $g_{i,t}$. Let W^t denote the weight matrix of the network used at round t . Every learner i updates its approximation of the global minimizer after $K_t = \lceil \frac{-2 \log t}{\log \sigma_2(W^t)} \rceil$ consensus iterations, where $\sigma_2(W^t)$ is the second largest singular value of the communication network graph at round t , i.e., W^t . At iteration K_t , after collecting messages from all its neighboring learners, learner i 's estimate of the global minimizer is

$$y_{i,t} = \sum_{j=1}^n ((W^t)^{K_t})_{ij} x_{j,t}. \quad (4)$$

Then, the loss functions are revealed to the learners individually. Every learner i computes the gradient of the local loss function $\nabla f_{t,i}(y_{t,i})$, which is evaluated at its estimate of the global minimizer $y_{i,t}$. Next, every learner i updates its estimate of the global gradient using K_t consensus

Algorithm 1 DOMD-MADGC

Input: Arbitrary initialization of $\{x_{i,1}\} \in \mathcal{X}$; step size η ; time horizon T .

Output: Sequence of decisions $\{x_{i,t}, y_{i,t} : 1 \leq t \leq T\}$.

1: **for** $t = 1, 2, \dots, T$ **do**
 2: set $K_t = \lceil \frac{-2 \log t}{\log \sigma_2(W^t)} \rceil$
 3: $y_{i,t} = \sum_{j=1}^n ((W^t)^{K_t})_{ij} x_{j,t}$
 4: $g_{i,t} = \sum_{j=1}^n ((W^t)^{K_t})_{ij} \nabla f_{j,t}(y_{j,t})$
 5: $x_{i,t+1} = \operatorname{argmin}_{x \in \mathcal{X}} \left\{ \langle x, g_{i,t} \rangle + \frac{1}{\eta} D_r(x, y_{i,t}) \right\}$
 6: **end for**

averaging iterations over the local gradients, which leads to

$$g_{i,t} = \sum_{j=1}^n ((W^t)^{K_t})_{ij} \nabla f_{j,t}(y_{j,t}). \quad (5)$$

The objective of this step is to enable the learners to access an estimate of the global gradient, which is collaboratively built based on the knowledge of the individual local learner as well as the neighboring learners.

After computing the local estimate vectors $y_{i,t}$ and $g_{i,t}$, learner node i updates its local decision variable by

$$x_{i,t+1} = \mathbf{MD}_\eta(g_{i,t}, y_{i,t}) \triangleq \operatorname{argmin}_{x \in \mathcal{X}} \left\{ \langle x, g_{i,t} \rangle + \frac{1}{\eta} D_r(x, y_{i,t}) \right\}, \quad (6)$$

where $\mathbf{MD}_\eta(g, y)$ represents the mirror descent update with a step size of η , and $D_r(x, y)$ is the Bregman divergence between x and y corresponding to the regularization function $r(x)$. Recall that $x_{i,t}$ and $y_{i,t}$ are respectively the local decision and the local estimate of the global minimizer at round t . Thus, (6) suggests that every learner i aims to stay close to the locally estimated global minimizer $y_{i,t}$ as measured by the Bregman divergence, while taking a step in a direction close to $g_{i,t}$ to reduce the local estimate of the global loss function at the current round. The above procedure of DOMD-MADGC is summarized in Algorithm 1.

4.2. Improved Dynamic Regret

In this section, we show how DOMD-MADGC improves the dynamic regret via multiple consensus averaging over both the local gradients and the local decisions. We first start by stating several standard assumptions commonly used in the literature of learning theory.

The following set of assumptions are commonly used in the literature after the group of studies began by (Hazan et al., 2006; Shalev-Shwartz and Singer, 2007), to provide stronger regret bounds by constraining the curvature of loss functions. For example, they are also used in (Chang and Shahrampour, 2020; Hendriks et al., 2020; Zhang et al., 2020).

Assumption 1. The loss functions $f_{i,t}(x)$ are λ -strongly convex and β -smooth. The loss functions have bounded gradients, i.e., $\|\nabla f_{i,t}(x)\| \leq G$. The regularization function $r(x)$ is μ -strongly convex and μ' -smooth.

The next set of assumptions pertain to the distributed nature of the learning network. Following previous studies on distributed optimization with time-varying network topology, we make several standard assumptions on topology graph \mathcal{G}_t .

Assumption 2. The weight matrix W^t is doubly stochastic, i.e.,

$$\sum_{j=1}^n W_{ij}^t = \sum_{i=1}^n W_{ij}^t = 1, \quad \forall t \geq 0. \quad (7)$$

Also, there exists a path from any learner i to any learner j , i.e., the network is connected in each online round.

Assumption 2 ensures that the entries of the matrix $(W^t)^K = \underbrace{W^t \dots W^t}_{K \text{ times}}$ are close to $1/n$,

which is expressed by

$$\sum_{j=1}^n \left| \left((W^t)^K \right)_{ij} - \frac{1}{n} \right| \leq \sqrt{n} \sigma_2^K(W^t), \quad (8)$$

where $\sigma_2(W^t)$ is the second largest singular value of the weight matrix W^t . This is a standard property of doubly stochastic matrices ([Horn and Johnson, 2012](#)).

Now, we are ready to analyze the performance of DOMD-MADGC. We begin by bounding the distance between the exact average of the local decisions and the current global minimizer in the following lemma.

Lemma 1 *Under Assumptions 1 and 2, the sequence of decisions generated by DOMD-MADGC with the step size $\eta < \frac{\mu'}{\lambda}$ satisfy the following bound:*

$$\|\bar{x}_{t+1} - x_t^*\| \leq \rho \|\bar{x}_t - x_t^*\| + \|\Delta_t\| + \|\delta_t\|, \quad (9)$$

where x_t^* is the minimizer of $f_t(x)$, $\bar{x}_t = \frac{1}{n} \sum_{i=1}^n x_{i,t}$ represents the exact average of local decisions, $\rho = \frac{\mu' - \eta\lambda}{\mu}$, and Δ_t and δ_t are given by

$$\Delta_t = \bar{x}_{t+1} - \mathbf{MD}_\eta(\bar{g}_t, \bar{x}_t), \quad (10)$$

$$\delta_t = \mathbf{MD}_\eta(\bar{g}_t, \bar{x}_t) - \mathbf{MD}_\eta(\nabla f_t(\bar{x}_t), \bar{x}_t). \quad (11)$$

The following is a proof sketch of Lemma 1. We first decompose the left-hand side of (9) into multiple terms, each representing a distinct source of error. In particular, the time-varying minimizers and limited local knowledge are two prominent source of errors. The error term due to the first source, called tracking error, measures the distance between the time-varying minimizer and the local decision of a learner, if it had access to exact average of the local decisions and global loss functions. Using the smooth duality property, from Lemma 2.19 in ([Shalev-Shwartz, 2012](#)), we bound the tracking error based on $\|\bar{x}_t - x_t^*\|$. The other error terms are denoted by Δ_t and δ_t . They are due to the second source of error, which reflects the inadequacy of the knowledge of local learners. Thus, they are called network errors. Since the learners do not have access to the gradient of the global loss and the exact average of decisions \bar{x}_t , their approximation of these quantities will always have some errors. The detailed proof of Lemma 1 is given in App. A of ([Eshraghi and Liang, 2021](#)).

The result in Lemma 1 implies that after one round of DOMD-MADGC, the distance between the exact average of the local decisions \bar{x}_{t+1} and the current global minimizer x_t^* is less than a ρ -fraction of the distance between \bar{x}_t and x_t^* , plus the network errors Δ_t and δ_t .

The following lemma provides an alternative form of update for mirror descent, which we require in our regret analysis.

Lemma 2 *Suppose $r(x)$ is strongly convex and y satisfies $\nabla r(y) = \nabla r(u) - \eta \nabla f(l)$ for some convex function $f(x)$, and vectors $l, u \in \mathcal{X}$ and step size η . We have*

$$\operatorname{argmin}_{x \in \mathcal{X}} \left\{ \langle \nabla f(l), x \rangle + \frac{1}{\eta} D_r(x, u) \right\} = \operatorname{argmin}_{x \in \mathcal{X}} D_r(x, y).$$

Lemma 2 is given as *Proposition 17* in (Chiang et al., 2012).

We are now ready to upper bound the dynamic regret in Theorem 3.

Theorem 3 *Under Assumptions 1 and 2, the dynamic regret of DOMD-MADGC, with $K_t = \lceil \frac{-2 \log t}{\log \sigma_2(W^t)} \rceil$ and a fixed step size $\frac{\mu' - \mu}{\lambda} < \eta < \frac{\mu'}{\lambda}$, satisfies*

$$\operatorname{Reg}_T^d \leq \frac{G \|\bar{x}_1 - x_1^*\|}{1 - \rho} + \frac{G}{1 - \rho} \sum_{t=2}^T \|x_t^* - x_{t-1}^*\| + \left(\frac{GR\mu' + \eta G^2 + \eta \lambda GR}{\mu} + GR \right) \frac{\sqrt{n} \pi^2}{6}$$

where $R = \max_{x \in \mathcal{X}} \|x\|$, and C_T is the path length defined in (1).

The proof of Theorem 3 is given in App. B of (Eshraghi and Liang, 2021).

Remark 1. Theorem 3 indicates that the dynamic regret of DOMD-MADGC is bounded by $O(1 + C_T)$, where C_T is the path length, defined in (1).

Remark 2. Recall that the dynamic regret of the distributed online mirror descent algorithm in (Shahrampour and Jadbabaie, 2018) is upper bounded by $O(\sqrt{T(1 + C_T)})$ for convex and Lipschitz continuous cost functions. Theorem 3 shows that the upper bound on the dynamic regret can be improved to $O(1 + C_T)$ using both primal and dual information boosted with multiple consensus steps for strongly convex and smooth cost functions. The $O(1 + C_T)$ dynamic regret bound of DOMD-MADGC is smaller than $O(\sqrt{T(1 + C_T)})$ as long as C_T is in the order of $o(T)$. Furthermore, in contrast to (Shahrampour and Jadbabaie, 2018), our regret bound can be achieved by setting a constant step size without any prior knowledge of C_T or the number of rounds T .

Remark 3. Recent studies on distributed online convex optimization have shown that the upper bound on the dynamic regret can be as tight as $O(1 + C_T + V_T)$ or $O(\log T(1 + C_T))$ when the loss functions are strongly convex and smooth (Zhang et al., 2020; Dixit et al., 2019). Theorem 3 shows that by employing multiple consensus averaging iterations over both local decisions and local gradients, DOMD-MADGC can improve the dynamic regret bound to $O(1 + C_T)$. In particular, gradient exchange allows the learners to approximate and track the global gradient, which is boosted by multiple consensus iterations to allow them to collect the information of a larger number of nodes. These two steps together enable the learners to more precisely track the time-varying global minimizer, which lead to the improved regret bound of $O(1 + C_T)$.

5. Discussion

We now turn to investigate the effect of network topology and weight matrix on the communication load. Recall that W^t represents the weight matrix of the network topology in online round t . The largest singular value $\sigma_1(W^t) = 1$ since W^t is doubly stochastic. The number of consensus steps per round $K_t = \lceil \frac{-2 \log t}{\log \sigma_2(W^t)} \rceil \leq \lceil \frac{2 \log t}{1 - \sigma_2(W^t)} \rceil$, which follows from the fact that $\log \sigma_2^{-1}(W^t) \geq 1 - \sigma_2(W^t)$. Therefore, the parameter K_t is controlled by the spectral gap of the weight matrix W^t , i.e., $\gamma(W^t) = 1 - \sigma_2(W^t)$.

Using the spectral gap, we can derive explicit bounds on the number of consensus steps for several interesting networks.

Regular Grids: The grid graph is formed by placing the nodes on a two dimensional plane. In d -connected regular grid, each node directly connects with its d nearest neighbors in axis-aligned directions. The grid networks have often been used to model cluster computing and sensor networks. For a \sqrt{n} -by- \sqrt{n} d -regular grid network the second-largest singular value is bounded by $\sigma_2(W^t) = 1 - \Theta(\frac{k^2}{n})$ (Duchi et al., 2012; Chung and Graham, 1997). In this case, the number of consensus steps per round is bounded by $K_t \in O(\frac{n \log t}{d^2})$.

Random Geometric Graphs: A random geometric graph can be constructed by placing the nodes uniformly on a two dimensional plane and connecting any pair of nodes whose distance is less than some pre-determined radius r . These graphs are commonly used to model the connectivity pattern of a set of wireless devices, such as in mobile communication networks. The properties of a random geometric graph with $r = ((\log^{1+\epsilon} n)/n)^{1/2}$ and $\epsilon > 0$ are analyzed in (Von Luxburg et al., 2014). In this case, the second-largest singular value of W^t is bounded by $\sigma_2(W^t) = 1 - \Omega(\frac{\log^{1+\epsilon} n}{n})$, so K_t is in the order of $O(\frac{n \log t}{\log n})$.

Communication cost vs. regret bound: The communication cost of DOMD-MADGC, which is defined as the overall number of consensus iterations, is $O(T \log T)$ after T rounds. In comparison, prior works based on a single consensus iteration per online round has a communication cost of $O(T)$ after T rounds. On the other hand, those prior works arrive at an upper bound of $O(\sqrt{T}(1 + C_T))$ on the dynamic regret when the prior knowledge of C_T is not available, which has an additional \sqrt{T} factor compared with our bound. Thus, the extra communication cost of $O(\log T)$ can be viewed as the price of removing the $O(\sqrt{T})$ factor in the dynamic regret bound.

6. Conclusion

In this paper, we have presented a novel algorithm for distributed online optimization, with an aim to improve the dynamic regret. The proposed DOMD-MADGC runs multiple consensus averaging iterations over both the local decisions and the local gradients, which allows the distributed learners to accurately estimate the gradients of the time-varying global loss functions. Furthermore, our algorithm does not require any prior knowledge of the regularity measures, such as C_T , or the number of rounds T . Our theoretical analysis shows that the dynamic regret of DOMD-MADGC is bounded by $O(1 + C_T)$, which is the best known result for distributed online optimization.

7. Acknowledgments

This work has been funded by the Natural Sciences and Engineering Research Council (NSERC) of Canada.

References

- Mohammad Akbari, Bahman Ghahesifard, and Tamás Linder. Distributed online convex optimization on time-varying directed graphs. *IEEE Transactions on Control of Network Systems*, 4(3): 417–428, 2015.
- Omar Besbes, Yonatan Gur, and Assaf Zeevi. Non-stationary stochastic optimization. *Operations Research*, 63(5):1227–1244, 2015.
- Stephen Boyd, Neal Parikh, and Eric Chu. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 2011.
- Nicolò Campolongo and Francesco Orabona. Temporal variability in implicit online learning. In *Advances in Neural Information Processing Systems*, 2020.
- Nicolò Campolongo and Francesco Orabona. A closer look at temporal variability in dynamic online learning. *arXiv preprint, arXiv:2102.07666*, 2021.
- Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge university press, 2006.
- Chih-Chung Chang and Chih-Jen Lin. LIBSVM: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2(3):1–27, 2011.
- Ting-Jui Chang and Shahin Shahrampour. Unconstrained online optimization: Dynamic regret analysis of strongly convex and smooth problems. *arXiv preprint, arXiv:2006.03912*, 2020.
- Chao-Kai Chiang, Tianbao Yang, Chia-Jung Lee, Mehrdad Mahdavi, Chi-Jen Lu, Rong Jin, and Shenghuo Zhu. Online optimization with gradual variations. In *Proceedings of Annual Conference on Learning Theory*, 2012.
- Fan RK Chung and Fan Chung Graham. *Spectral graph theory*. Number 92. American Mathematical Soc., 1997.
- Rishabh Dixit, Amrit Singh Bedi, Ketan Rajawat, and Alec Koppel. Distributed online learning over time-varying graphs via proximal gradient descent. In *Proceedings of IEEE International Conference on Decision and Control*, 2019.
- John C Duchi, Shai Shalev-Shwartz, Yoram Singer, and Ambuj Tewari. Composite objective mirror descent. In *Proceedings of the Conference on Learning Theory*, 2010.
- John C. Duchi, Alekh Agarwal, and Martin J. Wainwright. Dual averaging for distributed optimization: Convergence analysis and network scaling. *IEEE Transactions on Automatic Control*, 57(3), 2012.

- Nima Eshraghi and Ben Liang. Improving dynamic regret in distributed online mirror descent using primal and dual information. *arXiv preprint, arXiv:4062148*, 2021.
- Eric Hall and Rebecca Willett. Dynamical models and tracking regret in online convex programming. In *Proceedings of International Conference on Machine Learning*, 2013.
- Eric C Hall and Rebecca M Willett. Online convex optimization in dynamic environments. *IEEE Journal of Selected Topics in Signal Processing*, 9(4):647–662, 2015.
- Elad Hazan. Introduction to online convex optimization. *arXiv preprint, arXiv:1909.05207*, 2019.
- Elad Hazan, Adam Tauman Kalai, Satyen Kale, and Amit Agarwal. Logarithmic regret algorithms for online convex optimization. In *Proceedings of the Conference on Learning Theory*, 2006.
- Hadrien Hendrikx, Lin Xiao, Sebastien Bubeck, Francis Bach, and Laurent Massoulié. Statistically preconditioned accelerated gradient method for distributed optimization. 2020.
- Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge university press, 2012.
- Saghar Hosseini, Airlie Chapman, and Mehran Mesbahi. Online distributed optimization via dual averaging. In *Proceedings of IEEE International Conference on Decision and Control*, 2013.
- Ali Jadbabaie, Alexander Rakhlin, Shahin Shahrampour, and Karthik Sridharan. Online optimization: Competing with dynamic comparators. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 2015.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009. URL <http://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>.
- Soomin Lee, Angelia Nedić, and Maxim Raginsky. Coordinate dual averaging for decentralized online optimization with nonseparable global objectives. *IEEE Transactions on Control of Network Systems*, 5(1):34–44, 2016.
- Xiuxian Li, Xinlei Yi, and Lihua Xie. Distributed online convex optimization with an aggregative variable. *arXiv preprint, arXiv:2007.06844*, 2020.
- Kaihong Lu, Gangshan Jing, and Long Wang. Online distributed optimization with strongly pseudoconvex-sum cost functions. *IEEE Transactions on Automatic Control*, 65(1):426–433, 2019.
- David Mateos-Núñez and Jorge Cortés. Distributed online convex optimization over jointly connected digraphs. *IEEE Transactions on Network Science and Engineering*, 1(1):23–37, 2014.
- Aryan Mokhtari, Shahin Shahrampour, Ali Jadbabaie, and Alejandro Ribeiro. Online optimization in dynamic environments: Improved regret rates for strongly convex problems. In *Proceedings of IEEE International Conference on Decision and Control*, 2016.
- Francesco Orabona. A modern introduction to online learning. *arXiv preprint, arXiv:1912.13213*, 2019.

- Michael Rabbat. Multi-agent mirror descent for decentralized stochastic optimization. In *IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, 2015.
- Shahin Shahrampour and Ali Jadbabaie. Distributed online optimization in dynamic environments using mirror descent. *IEEE Transactions on Automatic Control*, 63(3):714–725, 2018.
- Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2012.
- Shai Shalev-Shwartz and Yoram Singer. *Logarithmic Regret Algorithms for Strongly Convex Repeated Games*. The Hebrew University, 2007.
- Konstantinos I Tsianos and Michael G Rabbat. Efficient distributed online prediction and stochastic optimization with approximate distributed averaging. *IEEE Transactions on Signal and Information Processing over Networks*, 2(4):489–506, 2016.
- Ulrike Von Luxburg, Agnes Radl, and Matthias Hein. Hitting and commute times in large random neighborhood graphs. *The Journal of Machine Learning Research*, 15(1):1751–1798, 2014.
- Xiaohan Wei, Hao Yu, and Michael J Neely. Online primal-dual mirror descent under stochastic constraints. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, pages 1–36, 2020.
- C Xi, Q Wu, and UA Khan. Distributed mirror descent over directed graphs. *arXiv preprint, arXiv:1412.5526*, 2014.
- Lin Xiao. Dual averaging methods for regularized stochastic learning and online optimization. *Journal of Machine Learning Research*, 11(88):2543–2596, 2010.
- Tianbao Yang, Lijun Zhang, Rong Jin, and Jinfeng Yi. Tracking slowly moving clairvoyant: Optimal dynamic regret of online learning with true and noisy gradient. In *Proceedings of International Conference on Machine Learning*, 2016.
- Deming Yuan, Yiguang Hong, Daniel WC Ho, and Shengyuan Xu. Distributed mirror descent for online composite optimization. *IEEE Transactions on Automatic Control*, 2020.
- Lijun Zhang, Tianbao Yang, Jinfeng Yi, Rong Jin, and Zhi-Hua Zhou. Improved dynamic regret for non-degenerate functions. In *Proceedings of the International Conference on Advances in Neural Information Processing Systems*, 2017.
- Yan Zhang, Robert J. Ravier, Vahid Tarokh, and Michael M. Zavlanos. Distributed online convex optimization with improved dynamic regret. *arXiv preprint, arXiv:1911.05127*, 2020.
- Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the International Conference on Machine Learning*, 2003.

Appendix A. Proof of Lemma 1

We begin by considering the difference between $\bar{x}_{t+1} = \frac{1}{n} \sum_{i=1}^n x_{i,t+1}$ and x_t^* . Using the update in (6), i.e., $x_{i,t+1} = \text{MD}_\eta(g_{i,t}, y_{i,t})$, and by adding and subtracting several terms we obtain

$$\begin{aligned}
 \|\bar{x}_{t+1} - x_t^*\| &= \left\| \frac{1}{n} \sum_{i=1}^n \text{MD}_\eta(g_{i,t}, y_{i,t}) - \text{MD}_\eta(\nabla f_t(x_t^*), x_t^*) \right\| \\
 &\leq \left\| \frac{1}{n} \sum_{i=1}^n \text{MD}_\eta(g_{i,t}, y_{i,t}) - \text{MD}_\eta(\bar{g}_t, \bar{x}_t) + \text{MD}_\eta(\bar{g}_t, \bar{x}_t) - \text{MD}_\eta(\nabla f_t(x_t^*), x_t^*) \right\| \\
 &\leq \left\| \text{MD}_\eta(\bar{g}_t, \bar{x}_t) - \text{MD}_\eta(\nabla f_t(\bar{x}_t), \bar{x}_t) + \text{MD}_\eta(\nabla f_t(\bar{x}_t), \bar{x}_t) - \text{MD}_\eta(\nabla f_t(x_t^*), x_t^*) \right\| \\
 &\quad + \left\| \frac{1}{n} \sum_{i=1}^n \text{MD}_\eta(g_{i,t}, y_{i,t}) - \text{MD}_\eta(\bar{g}_t, \bar{x}_t) \right\| \\
 &\leq \left\| \text{MD}_\eta(\nabla f_t(\bar{x}_t), \bar{x}_t) - \text{MD}_\eta(\nabla f_t(x_t^*), x_t^*) \right\| + \left\| \text{MD}_\eta(\bar{g}_t, \bar{x}_t) - \text{MD}_\eta(\nabla f_t(\bar{x}_t), \bar{x}_t) \right\| \\
 &\quad + \left\| \frac{1}{n} \sum_{i=1}^n \text{MD}_\eta(g_{i,t}, y_{i,t}) - \text{MD}_\eta(\bar{g}_t, \bar{x}_t) \right\|, \tag{12}
 \end{aligned}$$

where the first line follows from the fact that $x_t^* = \text{MD}_\eta(\nabla f_t(x_t^*), x_t^*)$. We bound each of the terms on the right hand-side of (12) separately. The first term of the above inequality can be expanded as

$$\begin{aligned}
 &\left\| \text{MD}_\eta(\nabla f_t(\bar{x}_t), \bar{x}_t) - \text{MD}_\eta(\nabla f_t(x_t^*), x_t^*) \right\|^2 \\
 &\leq \|\nabla r^*(\nabla r(\bar{x}_t) - \eta \nabla f_t(\bar{x}_t)) - \nabla r^*(\nabla r(x_t^*) - \eta \nabla f_t(x_t^*))\|^2 \\
 &\leq \frac{1}{\mu^2} \|\nabla r(\bar{x}_t) - \eta \nabla f_t(\bar{x}_t) - (\nabla r(x_t^*) - \eta \nabla f_t(x_t^*))\|^2, \tag{13}
 \end{aligned}$$

where the first line is obtained using the alternate form of mirror descent update stated in Lemma 2, and the fact that the inverse of $\nabla r(x)$ is $\nabla r^*(x)$ when $r(x)$ is strongly convex. Here $r^*(x)$ denotes the conjugate function (Shalev-Shwartz, 2012). We have also used the fact that the distance of the projection of two points into a convex set is smaller than the distance between these unprojected points. Note that since $r(x)$ is μ -strongly convex, its conjugate $r^*(x)$ is $1/\mu$ -smooth (Shalev-Shwartz, 2012).

We now proceed to bound the right hand-side of (13). The smoothness of $r(x)$ implies

$$r(y) \leq r(x) + \langle \nabla r(x), y - x \rangle + \frac{\mu'}{2} \|y - x\|^2, \forall x, y \in \mathcal{X}. \tag{14}$$

In addition, $f_t(x)$ is λ -strongly convex, i.e.,

$$f_t(y) \geq f_t(x) + \langle \nabla f_t(x), y - x \rangle + \frac{\lambda}{2} \|y - x\|^2, \forall x, y \in \mathcal{X}. \tag{15}$$

We multiply (15) by $-\eta$ and add it to (14) to obtain

$$r(y) - \eta f_t(y) \leq r(x) - \eta f_t(x) + \langle \nabla r(x) - \eta \nabla f_t(x), y - x \rangle + \frac{\mu' - \eta \lambda}{2} \|y - x\|^2, \forall x, y \in \mathcal{X}, \tag{16}$$

which shows that the function $r(x) - \eta f_t(x)$ is $(\mu' - \eta\lambda)$ -smooth when $\eta < \mu'/\lambda$. Therefore, we have

$$\|\nabla r(\bar{x}_t) - \eta \nabla f_t(\bar{x}_t) - (\nabla r(x_t^*) - \eta \nabla f_t(x_t^*))\|^2 \leq (\mu' - \eta\lambda)^2 \|\bar{x}_t - x_t^*\|^2. \quad (17)$$

By combining the above inequality and (13), we have

$$\left\| \mathbf{MD}_\eta(\nabla f_t(\bar{x}_t), \bar{x}_t) - \mathbf{MD}_\eta(\nabla f_t(x_t^*), x_t^*) \right\|^2 \leq \rho^2 \|\bar{x}_t - x_t^*\|^2, \quad (18)$$

where $\rho = \frac{\mu' - \eta\lambda}{\mu}$. Substituting (18) into (12), and using the definitions of Δ_t and δ_t in (10), we get

$$\|\bar{x}_{t+1} - x_t^*\| \leq \rho \|\bar{x}_t - x_t^*\| + \|\Delta_t\| + \|\delta_t\|. \quad (19)$$

Appendix B. Proof of Theorem 3

We provide a proof sketch for Theorem 3. We begin the proof by first bounding the network error term Δ_t . Using the definition of Δ_t and simple norm properties, we have

$$\begin{aligned} \|\Delta_t\| &= \left\| \frac{1}{n} \sum_{i=1}^n \mathbf{MD}_\eta(g_{i,t}, y_{i,t}) - \mathbf{MD}_\eta(\bar{g}_t, \bar{x}_t) \right\| \\ &\leq \frac{1}{n} \sum_{i=1}^n \left\| \mathbf{MD}_\eta(g_{i,t}, y_{i,t}) - \mathbf{MD}_\eta(\bar{g}_t, \bar{x}_t) \right\| \\ &\leq \frac{1}{n} \sum_{i=1}^n \left\| \nabla r^*(\nabla r(y_{i,t}) - \eta g_{i,t}) - \nabla r^*(\nabla r(\bar{x}_t) - \eta \bar{g}_t) \right\|, \end{aligned} \quad (20)$$

where the last line follows the result of Lemma 2, and the fact that the inverse of $\nabla r(x)$ is $\nabla r^*(x)$ when $r(x)$ is strongly convex (Shalev-Shwartz, 2012). We note that $r^*(x)$ denotes the conjugate of $r(x)$, which is defined by $r^*(x) = \max_u \{ \langle u, x \rangle - r(u) \}$. We have also used the fact that the distance between the projection of two points into a convex set is less than that of unprojected points. In addition, since $r(x)$ is μ -strongly convex, its conjugate $r^*(x)$ is $1/\mu$ -smooth (Shalev-Shwartz, 2012). Thus, we have

$$\begin{aligned} &\left\| \nabla r^*(\nabla r(y_{i,t}) - \eta g_{i,t}) - \nabla r^*(\nabla r(\bar{x}_t) - \eta \bar{g}_t) \right\| \\ &\leq \frac{1}{\mu} \left\| \nabla r(y_{i,t}) - \eta g_{i,t} - \nabla r(\bar{x}_t) + \eta \bar{g}_t \right\| \\ &\leq \frac{1}{\mu} \left[\left\| \nabla r(y_{i,t}) - \nabla r(\bar{x}_t) \right\| + \eta \left\| g_{i,t} - \bar{g}_t \right\| \right] \\ &\leq \frac{\mu'}{\mu} \|y_{i,t} - \bar{x}_t\| + \frac{\eta}{\mu} \|g_{i,t} - \bar{g}_t\|, \end{aligned} \quad (21)$$

where the last line is obtained due to the smoothness of $r(x)$. The above inequality implies that the norm of the network error term Δ_t , and ultimately the dynamic regret, are bounded by $\|y_{i,t} - \bar{x}_t\|$, which indicates the need for consensus over local decisions, and by $\|g_{i,t} - \bar{g}_t\|$, which indicates the necessity of consensus over local gradients. The resultant error bound based on these terms shows the importance of having accurate estimation of the exact average of decisions and gradients,

which motivates running a sufficient number of consensus iterations on both local decisions and local gradients to reduce the above terms. Thus, we combine (20) and (21), and noting the number of consensus iterations $K_t = \lceil \frac{-2 \log t}{\log \sigma_2(W^t)} \rceil$, to obtain

$$\begin{aligned}
 \|\Delta_t\| &\leq \frac{\mu'}{n\mu} \sum_{i=1}^n \left\| \sum_{j=1}^n \left(\left((W^t)^{K_t} \right)_{ij} - \frac{1}{n} \right) x_{j,t} \right\| \\
 &\quad + \frac{\eta}{n\mu} \sum_{i=1}^n \left\| \sum_{j=1}^n \left(\left((W^t)^{K_t} \right)_{ij} - \frac{1}{n} \right) \nabla f_{j,t}(y_{j,t}) \right\| \\
 &\leq \frac{\mu'}{n\mu} \sum_{i=1}^n \sum_{j=1}^n \left| \left((W^t)^{K_t} \right)_{ij} - \frac{1}{n} \right| \|x_{j,t}\| \\
 &\quad + \frac{\eta}{n\mu} \sum_{i=1}^n \sum_{j=1}^n \left| \left((W^t)^{K_t} \right)_{ij} - \frac{1}{n} \right| \|\nabla f_{j,t}(y_{j,t})\| \\
 &\leq \frac{\mu'}{n\mu} \sum_{i=1}^n \sqrt{n} \sigma_2^{K_t}(W^t) R + \frac{\eta}{n\mu} \sum_{i=1}^n \sqrt{n} \sigma_2^{K_t}(W^t) G \\
 &\leq \frac{\mu'}{\mu} \frac{R\sqrt{n}}{t^2} + \frac{\eta}{\mu} \frac{G\sqrt{n}}{t^2}, \tag{22}
 \end{aligned}$$

where the third inequality follows from (8). The above inequality implies that the network error Δ_t is bounded by $O(\frac{1}{t^2})$. This is an important result, since it implies that Δ_t is summable over time.

In the next step, we bound the other error term δ_t . Similar to the case with Δ_t , we use an alternative form of mirror descent update, and the properties of the conjugate function $r^*(x)$. We show that the norm of δ_t is bounded by $\|y_{i,t} - \bar{x}_t\|$. Our analysis further shows that a logarithmically increasing number of consensus averaging iterations $K_t = \lceil \frac{-2 \log t}{\log \sigma_2(W^t)} \rceil$ is sufficient to bound δ_t also by $O(1/t^2)$. Next, using the result of Lemma 1, and the upper bounds derived on Δ_t and δ_t , we bound the cumulative distance between the exact average of decisions and the time-varying minimizer $\sum_{t=1}^T \|\bar{x}_t - x_t^*\|$. We choose a step size $\frac{\mu' - \mu}{\lambda} < \eta < \frac{\mu'}{\lambda}$ to ensure that $\rho < 1$ in order to bound $\sum_{t=1}^T \|\bar{x}_t - x_t^*\|$ by $O(1 + C_T)$.

We note that Δ_t is bounded by $O(1/t^2)$, shown in (22). Now we proceed to bound the other term on network error δ_t as follows:

$$\begin{aligned}
 \|\delta_t\| &= \|\mathbf{MD}_\eta(\bar{g}_t, \bar{x}_t) - \mathbf{MD}_\eta(\nabla f_t(\bar{x}_t), \bar{x}_t)\| \\
 &\leq \|\nabla r^*(\nabla r(\bar{x}_t) - \eta \bar{g}_t) - \nabla r^*(\nabla r(\bar{x}_t) - \eta \nabla f_t(\bar{x}_t))\| \\
 &\leq \frac{1}{\mu} \|\nabla r(\bar{x}_t) - \eta \bar{g}_t - \nabla r(\bar{x}_t) + \eta \nabla f_t(\bar{x}_t)\| \\
 &\leq \frac{\eta}{\mu} \|\bar{g}_t - \nabla f_t(\bar{x}_t)\| \\
 &\leq \frac{\eta}{\mu} \left\| \frac{1}{n} \sum_{i=1}^n g_{i,t} - \frac{1}{n} \sum_{j=1}^n \nabla f_{j,t}(\bar{x}_t) \right\| \tag{23}
 \end{aligned}$$

where the first inequality follows from the alternative form of mirror descent update in Lemma 2. We also note that since $r(x)$ is μ -strongly convex, which implies its conjugate $r^*(x)$ is $1/\mu$ -smooth.

The last line shows that the network error δ_t relates to the distance between the average of the gradients of local learners and the gradient of the global loss function evaluated at \bar{x}_t . Therefore, we have

$$\begin{aligned}
 \|\delta_t\| &\leq \frac{\eta}{\mu} \left\| \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n \left((W^t)^{K_t} \right)_{ij} \nabla f_{j,t}(y_{j,t}) - \frac{1}{n} \sum_{j=1}^n \nabla f_{j,t}(\bar{x}_t) \right\| \\
 &\leq \frac{\eta}{\mu} \left\| \frac{1}{n} \sum_{j=1}^n \nabla f_{j,t}(y_{j,t}) - \frac{1}{n} \sum_{j=1}^n \nabla f_{j,t}(\bar{x}_t) \right\| \\
 &\leq \frac{\eta}{\mu} \frac{1}{n} \sum_{j=1}^n \left\| \nabla f_{j,t}(y_{j,t}) - \nabla f_{j,t}(\bar{x}_t) \right\| \\
 &\leq \frac{\eta}{\mu n} \sum_{j=1}^n \lambda \|y_{j,t} - \bar{x}_t\|, \tag{24}
 \end{aligned}$$

where the second line follows from the fact that the matrix $(W^t)^{K_t} = \underbrace{W^t \dots W^t}_{K_t \text{ times}}$ is doubly stochastic. We have also used the strong convexity of local loss functions to obtain the right-hand side of (24). The above inequality illustrates that the error δ_t is directly bounded by the distance between the local inexact average $y_{j,t}$ and the exact average of decisions \bar{x}_t . Thus, with K_t consensus iterations, we have

$$\begin{aligned}
 \|\delta_t\| &\leq \frac{\eta}{\mu n} \sum_{j=1}^n \lambda \sum_{i=1}^n \left| \left((W^t)^{K_t} \right)_{ij} - \frac{1}{n} \right| \|x_{i,t}\| \\
 &\leq \frac{\eta}{\mu} \sqrt{n} \lambda \sigma_2^{K_t}(W^t) R \\
 &\leq \frac{\eta \lambda \sqrt{n} R}{\mu t^2}, \tag{25}
 \end{aligned}$$

where the last line is obtained using $K_t = \lceil \frac{-2 \log t}{\log \sigma_2(W^t)} \rceil$. From (22) and (25), we observe that both error terms Δ_t and δ_t are bounded by $O(1/t^2)$.

Next, we upper bound the cumulative distance between \bar{x}_t and x_t^* . We first expand the summation and add and subtract several terms to obtain

$$\begin{aligned}
 \sum_{t=1}^T \|\bar{x}_t - x_t^*\| &= \|\bar{x}_1 - x_1^*\| + \sum_{t=2}^T \|\bar{x}_t - x_{t-1}^* + x_{t-1}^* - x_t^*\| \\
 &\leq \|\bar{x}_1 - x_1^*\| + \sum_{t=2}^T \|\bar{x}_t - x_{t-1}^*\| + \sum_{t=2}^T \|x_t^* - x_{t-1}^*\| \\
 &\leq \|\bar{x}_1 - x_1^*\| + \sum_{t=2}^T \left[\rho \|\bar{x}_{t-1} - x_{t-1}^*\| + \|\Delta_t\| + \|\delta_t\| \right] + \sum_{t=2}^T \|x_t^* - x_{t-1}^*\| \\
 &\leq \|\bar{x}_1 - x_1^*\| + \rho \sum_{t=1}^T \|\bar{x}_t - x_t^*\| + \sum_{t=1}^T \left(\frac{R\mu' + \eta G + \eta \lambda R}{\mu} \right) \frac{\sqrt{n}}{t^2} + \sum_{t=2}^T \|x_t^* - x_{t-1}^*\|, \tag{26}
 \end{aligned}$$

where we have applied Lemma 1 in the third line. We have further used the network error bounds in (22) and (25) to obtain the right hand-side of (26). The step size is chosen $\frac{\mu' - \mu}{\lambda} < \eta < \frac{\mu'}{\lambda}$ such that $\rho < 1$. Thus, we have

$$\sum_{t=1}^T \|\bar{x}_t - x_t^*\| \leq \frac{\|\bar{x}_1 - x_1^*\|}{1 - \rho} + \left(\frac{R\mu' + \eta G + \eta\lambda R}{\mu} \right) \frac{\sqrt{n}\pi^2}{6(1 - \rho)} + \frac{1}{1 - \rho} \sum_{t=2}^T \|x_t^* - x_{t-1}^*\|. \quad (27)$$

Now we are ready to establish an upper bound on the dynamic regret. Since each learner maintains two as follows:

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T f_t(y_{i,t}) - \sum_{t=1}^T f_t(x_t^*) &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \sum_{t=1}^T f_{j,t}(y_{i,t}) - \frac{1}{n} \sum_{t=1}^T \sum_{j=1}^n f_{j,t}(x_t^*) \\ &\leq \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \sum_{t=1}^T f_{j,t}(y_{i,t}) - f_{j,t}(x_t^*) \\ &\leq \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \sum_{t=1}^T G \|y_{i,t} - x_t^*\| \\ &\leq \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T G \|y_{i,t} - \bar{x}_t\| + \sum_{t=1}^T G \|\bar{x}_t - x_t^*\|, \end{aligned} \quad (28)$$

where we have used the Lipschitz continuity of local loss functions in the third line. It can be observed that the dynamic regret is a sum of two components. The first component reflects how close the local inexact average estimate is to the exact average of decisions. The second term measures the difference between the average of local decisions and the time-varying global minimizer, which is bounded in (27). Therefore, we now only need to bound the first term. Using (8), we have

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T G \|y_{i,t} - \bar{x}_t\| &\leq \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T G \sum_{j=1}^n \left| \left((W^t)^{K_t} \right)_{ij} - \frac{1}{n} \right| \|x_{j,t}\| \\ &\leq \sum_{t=1}^T GR \sqrt{n} \sigma_2 (W^t)^{K_t} \\ &\leq GR \sqrt{n} \sum_{t=1}^T \frac{1}{t^2} \\ &\leq GR \sqrt{n} \frac{\pi^2}{6}. \end{aligned} \quad (29)$$

Substituting (27) and (29) into (28) leads to the dynamic regret bound of

$$\text{Reg}_T^d \leq GR \sqrt{n} \frac{\pi^2}{6} + G \frac{\|\bar{x}_1 - x_1^*\|}{1 - \rho} + \left(\frac{GR\mu' + \eta G^2 + \eta\lambda GR}{\mu} \right) \frac{\sqrt{n}\pi^2}{6(1 - \rho)} + \frac{G}{1 - \rho} \sum_{t=2}^T \|x_t^* - x_{t-1}^*\|. \quad (30)$$

Appendix C. Experiments

We illustrate the performance of DOMD-MADGC with a series of numerical experiments on real datasets. In particular, we consider the problem of multi-class classification in the distributed online learning setting over notMNIST, CIFAR-10 (Krizhevsky, 2009), and the aloi dataset from the LIB-SVM repository (Chang and Lin, 2011). These three datasets respectively cover the cases of large number of data samples, large number of features, and large number of classes. They are summarized in Table 1.

Dataset	# Features	# Classes	# Examples
aloi	128	1000	108,000
notMNIST	784	10	500,000
CIFAR-10	3,072	10	60,000

Table 1: Summary of datasets

For all experiments presented in this section, the underlying network topology is time-varying. In particular, the communication graph switches sequentially in a round robin manner within a pool of randomly generated doubly stochastic matrices. In addition, each dataset is divided into equally-sized parts, and distributed among the learner nodes in the network.

In the first experiment, we consider multi-class classification with logistic regression over a large distributed network consisting of 1500 learner nodes. In this task, the learners observe a sequence of labeled examples (ω, z) taken from the aloi dataset, where $\omega \in \mathbb{R}^d$ denotes the feature vector, and $z \in \mathbb{R}$ represents the true label.

For logistic regression, the loss associated with each data point is given by $f(x, (\omega_t, z_t)) = \log(1 + \exp(-z_t x^T \omega_t))$. We use negative entropy as the regularization function $r(x)$. In addition, we set the batch size to 10 data examples per online round, and use the step size of $\eta = 0.01$.

In Fig. 1, we compare the performance of DOMD-MADGC with fixed and increasing number of consensus rounds versus DOMD (Shahrampour and Jadbabaie, 2018). Fig. 1 shows that DOMD-MADGC, even with fixed and small $K = 1$ and 5 outperforms the existing DOMD algorithm, which verifies the effectiveness of the proposed approach based on the communication of both local gradients and local decisions. We also find that the overall loss of DOMD-MADGC with logarithmically increasing number of consensus iterations decreases faster than all alternatives.

Next, we consider the problem of ridge regression on the well-known notMNIST dataset, and over a network with 200 distributed learners. For this task, the loss associated with each data point is given by $f(x, (\omega_t, z_t)) = (x^T \omega_t - z_t)^2$. In the experiment, we set the batch size to 40, set $\eta = 7 \times 10^{-4}$, and use the negative entropy as $r(x)$. Fig. 2 shows the overall loss versus the number of online rounds. We find that DOMD-MADGC reduces the overall loss faster than the alternatives.

We have also experimented on the CIFAR-10 dataset to solve the ridge regression problem in a distributed online manner. We consider a network of 500 learner nodes, and set the batch size to 10, and $\eta = 0.5$. Fig. 3 shows that DOMD-MADGC outperforms standard DOMD method in terms of the overall loss.

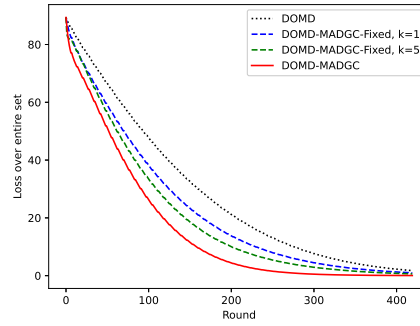


Figure 1: Logistic regression on aloi dataset.

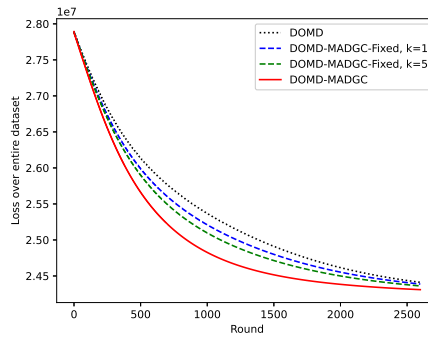


Figure 2: Ridge regression on notMNIST dataset.

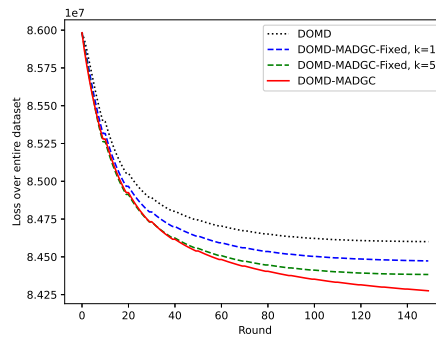


Figure 3: Ridge regression on CIFAR-10 dataset.