

Distributed Strongly Convex Optimization

Konstantinos I. Tsianos

Department of Electrical and Computer Engineering

McGill University

Montreal, Quebec H3A 0E9

Email: konstantinos.tsianos@gmail.com

and

Michael G. Rabbat

Department of Electrical and Computer Engineering

McGill University

Montreal, Quebec H3A 0E9

Email: michael.rabbat@mcgill.ca

Abstract

A lot of effort has been invested into characterizing the convergence rates of gradient based algorithms for non-linear convex optimization. Recently, motivated by large datasets and problems in machine learning, the interest has shifted towards distributed optimization. In this work we present a distributed algorithm for strongly convex constrained optimization. Each node in a network of n computers converges to the optimum of a strongly convex, L -Lipchitz continuous, separable objective at a rate $O\left(\frac{\log(\sqrt{nT})}{T}\right)$ where T is the number of iterations. This rate is achieved in the online setting where the data is revealed one at a time to the nodes, and in the batch setting where each node has access to its full local dataset from the start. The same convergence rate is achieved in expectation when the subgradients used at each node are corrupted with additive zero-mean noise.

I. INTRODUCTION

In this work we focus on solving optimization problems of the form

$$\underset{w \in \mathcal{W}}{\text{minimize}} F(w) = \frac{1}{T} \sum_{t=1}^T f^t(w) \quad (1)$$

where each function $f^1(w), f^2(w), \dots$, is convex over a convex set $\mathcal{W} \subseteq \mathbb{R}^d$. This formulation applies widely in machine learning scenarios, where $f^t(w)$ measures the loss of model w with respect to data point t , and $F(w)$ is the average loss over T data points. In particular, we are interested in the behavior of online distributed optimization algorithms for this sort of problem as the number of data points T tends to infinity. We describe a distributed algorithm which, for strongly convex functions f^t , converges at a rate $O\left(\frac{\log(\sqrt{nT})}{T}\right)$. To the best of our knowledge this is the first distributed algorithm to achieve this converge rate for constrained optimization without relying on

smoothness assumptions on the objective or non-trivial communication mechanisms between the nodes. The result is true both in the online and the batch optimization setting.

When faced with a non-linear convex optimization problem, gradient-based methods can be applied to find the solution. The behavior of these algorithms is well-understood in the single-processor (centralized) setting. Under the assumption that the objective is L -Lipschitz continuous, projected gradient descent-type algorithms converge at a rate $O(\frac{1}{\sqrt{T}})$ [1], [2]. This rate is achieved both in an online setting where the f^t 's are revealed to the algorithm sequentially and in the batch setting where all f^t are known in advance. If the cost functions are also strongly convex then gradient algorithms can achieve linear rates, $O(\frac{1}{T})$, in the batch setting [3] and nearly-linear rates, $O(\frac{\log(T)}{T})$, in the online setting [4]. Under additional smoothness assumptions, such as Lipschitz continuous gradients, the same rate of convergence can also be achieved by second order methods in the online setting [5], [6], while accelerated methods can achieve a quadratic rate in the batch setting; see [7] and references therein.

The aim of this work is to extend the aforementioned results to the distributed setting where a network of processors jointly optimize a similar objective. Assuming the network is arranged as an expander graph with constant spectral gap, for general convex cost functions that are only L -Lipschitz continuous, the rate at which existing algorithms on a network of n processors will all reach the optimum value is $O(\frac{\log(T\sqrt{n})}{\sqrt{T}})$, i.e., similar to the optimal single processor algorithms up to a logarithmic factor [8], [9]. This is true both in a batch setting and in an online setting, even when the gradients are corrupted by noise. The technique proposed in [10] makes use of mini-batches to obtain asymptotic rates $O(\frac{\sqrt{\log(n)}}{\sqrt{nT}})$ for online optimization of smooth cost functions that have Lipschitz continuous gradients corrupted by bounded variance noise, and $O(\frac{1}{nT})$ for smooth strongly convex functions. However, this technique requires that each node exchange messages with every other node at the end of each iteration. Finally, if the objective function is strongly convex and three times differentiable, a distributed version of Nesterov's accelerated method [11] achieves a rate of $O(\frac{\log(T)}{T})$ for unconstrained problems in the batch setting, but the dependence on n is not characterized.

The algorithm presented in this paper achieves a rate $O(\frac{\log(\sqrt{nT})}{T})$ for strongly convex functions. Our formulation allows for convex constraints in the problem and assumes the objective function is Lipschitz continuous and strongly convex; no higher-order smoothness assumptions are made. Our algorithm works in both the online and batch setting and it scales nearly-linearly in number of iterations for network topologies with fast information diffusion. In addition, at each iteration nodes are only required to exchange messages with a subset of other nodes in the network (their neighbors).

The rest of the paper is organized as follows. Section II introduces notation and formalizes the problem. Section III describes the proposed algorithm and states our main results. These results are proven in Section IV, and Section V extends the analysis to the case where gradients are noisy. Section VI presents the results of numerical experiments illustrating the performance of the algorithm, and the paper concludes in Section VII.

II. ONLINE CONVEX OPTIMIZATION

Consider the problem of minimizing a convex function $F(w)$ over a convex set $\mathcal{W} \subseteq \mathbb{R}^d$. Of particular interest is the setting where the algorithm sequentially receives noisy samples of the (sub)gradients of $F(w)$. This setting arises in online loss minimization for machine learning when the data arrives as a stream and the (sub)gradient is evaluated using an individual data point at each step [1]. Suppose the t th data point $x(t) \in \mathcal{X} \subseteq \mathbb{R}^d$ is drawn i.i.d. from an unknown distribution \mathcal{D} , and let $f^t(w) = f(w, x(t))$ denote the loss of this data point with respect to a particular model w . In this setting one would like to find the model w that minimizes the expected loss $\mathbb{E}_{\mathcal{D}}[f(w, x)]$, possibly with the constraint that w be restricted to a model space \mathcal{W} . Clearly, as $T \rightarrow \infty$, the objective $F(w) = \frac{1}{T} \sum_{t=1}^T f^t(w) \rightarrow \mathbb{E}_{\mathcal{D}}[f(w, x)]$, and so if the data stream is finite this motivates minimizing the empirical loss $F(w)$.

An online convex optimization algorithm observes a data stream $x(1), x(2), \dots$, and sequentially chooses a sequence of models $w(1), w(2), \dots$, after each observation. Upon choosing $w(t)$, the algorithm receives a subgradient $g(t) \in \partial f^t(w(t))$. The goal is for the sequence $w(1), w(2), \dots$ to converge to a minimizer w^* of $F(w)$.

The performance of an online optimization algorithm is measured in terms of the *regret*:

$$R(T) = \sum_{t=1}^T f^t(w(t)) - \min_{w \in \mathcal{W}} \sum_{t=1}^T f^t(w). \quad (2)$$

The regret measures the gap between the cost accumulated by the online optimization algorithm over T steps and that of a model chosen to simultaneously minimize the total regret over all T cost terms. If the costs f^t are allowed to be arbitrary convex functions then it can be shown that the best achievable rate for any online optimization algorithm is $\frac{R(T)}{T} = \Omega(\frac{1}{\sqrt{T}})$, and this bound is also achievable [1]. The rate can be significantly improved if the cost functions has more favourable properties.

A. Assumptions

Assumption 1: We assume for the rest of the paper that each cost function $f^t(w) = f(w, x(t))$ is σ -strongly convex for all $x(t) \in \mathcal{X}$; i.e., there is a $\sigma > 0$ such that for all $\theta \in [0, 1]$ and all $u, w \in \mathcal{W}$

$$f^t(\theta u + (1 - \theta)w) \leq \theta f^t(u) + (1 - \theta)f^t(w) - \frac{\sigma}{2}\theta(1 - \theta) \|u - w\|^2. \quad (3)$$

If each $f^t(w)$ is σ -strongly convex, it follows that $F(w)$ is also σ -strongly convex. Moreover, if $F(w)$ is strongly convex then it is also strictly convex, and so $F(w)$ has a unique minimizer which we denote by w^* .

Assumption 2: We also assume that the subgradients $g(t)$ of each cost function f^t are bounded by a known constant $L > 0$; i.e., $\|g(t)\| \leq L$ where $\|\cdot\|$ is the (ℓ_2) Euclidean norm.

B. Example: Training a Classifier

For a specific example of this setup, consider the problem of training an SVM classifier using a hinge-loss with ℓ_2 regularization [4]. In this case, the data stream consists of pairs $\{x(t), y(t)\}$ such that $x(t) \in \mathcal{X}$ and

$y(t) \in \{-1, +1\}$. The goal is to minimize the misclassification error as measured by the ℓ_2 -regularized hinge loss. Formally, we wish to find the $w^* \in \mathcal{W} \subseteq \mathbb{R}^d$ that solves

$$\text{minimize}_{w \in \mathcal{W}} \frac{\sigma}{2} \|w\|^2 + \frac{1}{m} \sum_{t=1}^m \max\{0, 1 - y(t)\langle w, x(t)\rangle\} \quad (4)$$

which is σ -strongly convex¹. For these types of problems, using a single-processor stochastic gradient descent algorithm, one can achieve $\frac{R(T)}{T} = O(\frac{\log T}{T})$ [4] or $\frac{R(T)}{T} = O(\frac{1}{T})$ [12] by using different update schemes.

C. Distributed Online Convex Optimization

In this paper, we are interested in solving online convex optimization problems with a network of computers. The computers are organized as a network $G = (V, E)$ with $|V| = n$ nodes, and messages are only exchanged between nodes connected with an edge in E .

Assumption 3: In this work we assume that G is connected and undirected.

Each node i receives a stream of data $x_i(1), x_i(2), \dots$, similar to the serial case, and the nodes must collaborate to minimize the network-wide objective

$$F(w) = \frac{1}{nT} \sum_{t=1}^T \sum_{i=1}^n f_i^t(w), \quad (5)$$

where $f_i^t(w) = f(w, x_i(t))$ is the cost incurred at processor i at time t . In the distributed setting, the definition of regret is naturally extended to

$$R(T) = \sum_{t=1}^T \sum_{i=1}^n f(w_i(t), x_i(t)) - \min_{w \in \mathcal{W}} \sum_{t=1}^T \sum_{i=1}^n f(w, x_i(t)). \quad (6)$$

For general convex cost functions, the distributed algorithm proposed in [8] has been proven to have an average regret that decreases at a rate \sqrt{T} , similar to the serial case, and this result holds even when the algorithm receives noisy, unbiased, observations of the true subgradients at each step. In the next section, we present a distributed algorithm that achieves a nearly-linear rate of decrease of the average regret (up to a logarithmic factor) when the cost functions are strongly convex.

III. ALGORITHM

Nodes must collaborate to solve the distributed online convex optimization problem described in the previous section. To that end, the network is endowed with a $n \times n$ consensus matrix P which respects the structure of G , in the sense that $[P]_{ji} = 0$ if $(i, j) \notin E$. We assume that P is doubly stochastic, although generalizations to the case where P is row stochastic or column stochastic (but not both) are also possible [13], [14].

A detailed description of the proposed algorithm, *distributed online gradient descent* (DOGD), is given in Algorithm 1. In the algorithm, each node performs a total of T updates. One update involves processing a single data point $x_i(t)$ at each processor. The updates are performed over k rounds, and T_s updates are performed in round

¹Although the hinge loss itself is not strongly convex, adding a strongly convex regularizer makes the overall cost function strongly convex.

Algorithm 1 DOGD

```

1: Initialize:  $T_1 = \lceil \frac{2}{\sigma} \rceil$ ,  $a_1 = 1$ ,  $k = 1$ ,  $z_i^1(1) = w_i^1(1) = 0$ 
2:
3: while  $\sum_{s=1}^k T_s \leq T$  do ▷ Each node  $i$  repeats
4:   for  $t = 1$  to  $T_k$  do
5:     Send/receive  $z_i^k(t)$  and  $z_j^k(t)$  to/from neighbors
6:     Obtain next subgradient  $g_i(t) \in \partial_w f_i^t(w_i^k(t))$ 
7:      $z_i^k(t+1) = \sum_{j=1}^n p_{ij} z_j^k(t) - a_k g_i(t)$ 
8:      $w_i^k(t+1) = \Pi_{\mathcal{W}} [z_i^k(t+1)]$ 
9:   end for
10:   $w_i^{k+1}(1) = w_i^k(T_k)$ 
11:   $z_i^{k+1}(1) = w_i^{k+1}(1)$ 
12:   $\hat{w}_i^{k+1} = \frac{1}{T_k} \sum_{t=1}^{T_k} w_i^k(t)$ 
13:
14:   $T_{k+1} \leftarrow 2T_k$ 
15:   $a_{k+1} \leftarrow \frac{a_k}{2}$ 
16:   $k = k + 1$ 
17: end while

```

$s \leq k$. The main steps within each round (lines 9–11) involve updating an accumulated gradient variable, $z_i^k(t)$, by simultaneously incorporating the information received from neighboring nodes and taking a local gradient-descent like step. The accumulated gradient is projected onto the constraint set to obtain $w_i^k(t)$, where

$$\Pi_{\mathcal{W}} [z] = \operatorname{argmin}_{w \in \mathcal{W}} \|w - z\| \quad (7)$$

denotes the Euclidean projection of z onto \mathcal{W} , and then this projected value is merged into a running average $\hat{w}_i(r)$. The step size parameter a_k remains constant within each round, and the step size is reduced by half at the end of each round. The number of updates per round doubles from one round to the next.

Note that the algorithm proposed here differs from the distributed dual averaging algorithm described in [8], where a proximal projection is used rather than the Euclidean projection. Also, in contrast to the distributed subgradient algorithms described in [15], DOGD maintains an accumulated gradient variable in $z_i^k(t+1)$ which is updated using $\{z_j^k(t)\}$ as opposed to the primal feasible variables $\{w_j^k(t)\}$. Finally, key to achieving fast convergence is the exponential decrease of the learning rate after performing an exponentially increasing number of gradient steps together with a proper initialization of the learning rate.

The next section provides theoretical guarantees on the performance of DOGD.

IV. CONVERGENCE ANALYSIS

Our main convergence result, stated below, guarantees that the average regret decreases at a rate which is nearly linear.

Theorem 1: Let Assumptions 1–3 hold and suppose that the consensus matrix P is doubly stochastic with constant λ_2 . Let w^* be the minimizer of $F(w)$. Then the sequence $\{\hat{w}_i^k\}$ produced by nodes running DOGD to minimize $F(w)$ obeys

$$F(\hat{w}_i^{k+1}) - F(w^*) = O\left(\frac{\log(\sqrt{nT})}{T}\right), \quad (8)$$

where $k = \lfloor \log_2(T/2 + 1) \rfloor$ is the number of rounds executed during a total of T gradient steps per node, and \hat{w}_i^k is the running average maintained locally at each node.

Remark 1: We state the result for the case where λ_2 is constant. This is the case when G is, e.g., a complete graph or an expander graph [16]. For other graph topologies where λ_2 shrinks with n and consensus does not converge fast, the convergence rate dependence on n is going to be worse due to a factor $1 - \sqrt{\lambda_2}$ in the denominator; see the proof of Theorem 1 below for the precise dependence on the spectral gap $1 - \sqrt{\lambda_2}$.

Remark 2: The theorem characterizes performance of the online algorithm DOGD, where the data and cost functions f_i^t are processed sequentially at each node in order to minimize an objective of the form

$$F(w) = \frac{1}{n} \sum_{i=1}^n \frac{1}{T} \sum_{t=1}^T f_i^t(w). \quad (9)$$

However, as pointed out in [4], if the entire dataset is available in advance, we can use the same scheme to do batch minimization by effectively setting $f_i^t(w) = f_i^1(w)$, where $f_i^1(w)$ is the objective function accounting for the entire dataset available to node i . Thus, the same result holds immediately for a batch version of DOGD.

The remainder of this section is devoted to the proof of Theorem 1. Our analysis follows arguments that can be found in [1], [8], [12] and references therein. We first state and prove some intermediate results.

A. Properties of Strongly Convex Functions

Recall the definition of σ -strong convexity given in Assumption 1. A direct consequence of this definition is that if $F(w)$ is σ -strongly convex then

$$F(w) - F(w^*) \geq \frac{\sigma}{2} \|w - w^*\|^2. \quad (10)$$

Strong convexity can be combined with the assumptions above to upper bound the difference $F(w) - F(w^*)$ for an arbitrary point $w \in \mathcal{W}$.

Lemma 1: Let w^* be the minimizer of $F(w)$. For all $w \in \mathcal{W}$, we have $F(w) - F(w^*) \leq \frac{2L^2}{\sigma}$.

Proof: For any subgradient g of F at w , by convexity we know that $F(w) - F(w^*) \leq \langle g, w - w^* \rangle$. It follows from Assumption 2 that $F(w) - F(w^*) \leq L \|w - w^*\|$. Furthermore, from Assumption 1, we obtain that $\frac{\sigma}{2} \|w - w^*\|^2 \leq L \|w - w^*\|$ or $\|w - w^*\| \leq \frac{2L}{\sigma}$. As a result, $F(w) - F(w^*) \leq \frac{2L^2}{\sigma}$. ■

B. The Lazy Projection Algorithm

The analysis of DOGD below involves showing that the average state, $\frac{1}{n} \sum_{i=1}^n w_i^k(t)$, evolves according to the so-called (single processor) *lazy projection* algorithm [1], which we discuss next. The lazy projection algorithm is an online convex optimization scheme for the serial problem discussed at the beginning of Section II. A single processor sequentially chooses a new variable $w(t)$ and receives a subgradient $g(t)$ of $f(w(t), x(t))$. The algorithm chooses $w(t+1)$ by repeating the steps

$$z(t+1) = z(t) - ag(t) \quad (11)$$

$$w(t+1) = \Pi_{\mathcal{W}} [z(t+1)]. \quad (12)$$

By unwrapping the recursive form of (11), we get

$$z(t+1) = -a \sum_{s=1}^t g(s) + z(1). \quad (13)$$

The following is a typical result for subgradient descent-style algorithms, and is useful towards eventually characterizing how the regret accumulates. Its proof can be found in the appendix of the extended version of [1].

Theorem 2 (Zinkevich [1]): Let $w(1) \in \mathcal{W}$, let $a > 0$, and set $z(1) = w(1)$. After T rounds of the serial lazy projection algorithm (11)–(12), we have

$$\sum_{t=1}^T \langle g(t), w(t) - w^* \rangle \leq \frac{\|w(1) - w^*\|^2}{2a} + \frac{TaL^2}{2}. \quad (14)$$

Theorem 2 immediately yields the same bound for the regret of lazy projection [1].

C. Evolution of Network-Average Quantities in DOGD

We turn our attention to Algorithm 1. A standard approach to studying convergence of distributed optimization algorithms, such as DOGD, is to keep track of the discrepancy between every node's state and an average state sequence defined as

$$\bar{z}^k(t) = \frac{1}{n} \sum_{i=1}^n z_i^k(t) \quad \text{and} \quad \bar{w}^k(t) = \Pi_{\mathcal{W}} [\bar{z}^k(t)]. \quad (15)$$

Observe that $\bar{z}^k(t)$ evolves in a simple recursive manner,

$$\bar{z}^k(t+1) = \frac{1}{n} \sum_{i=1}^n z_i^k(t+1) \quad (16)$$

$$= \frac{1}{n} \sum_{i=1}^n \left[\sum_{j=1}^n p_{ij} z_j^k(t) - a_k g_i(t) \right] \quad (17)$$

$$= \frac{1}{n} \sum_{j=1}^n z_j^k(t) \sum_{i=1}^n p_{ij} - \frac{a_k}{n} \sum_{i=1}^n g_i(t) \quad (18)$$

$$= \bar{z}^k(t) - \frac{a_k}{n} \sum_{i=1}^n g_i(t) \quad (19)$$

$$= -a_k \sum_{s=1}^t \frac{1}{n} \sum_{i=1}^n g_i(s) + \frac{1}{n} \sum_{i=1}^n z_i^k(1) \quad (20)$$

where equation (19) holds since P is doubly stochastic. Notice (cf. eqn. (13)) that the states $\{\bar{z}^k(t), \bar{w}^k(t)\}$ evolve according to the lazy projection algorithm with gradients $\bar{g}(t) = \frac{1}{n} \sum_{i=1}^n g_i(t)$ and learning rate a_k . In the sequel, we will also use an analytic expression for $z_i^k(t)$ derived by back substituting in its recursive update equation. After some algebraic manipulation, we obtain

$$\begin{aligned} z_i^k(t) = & -a_k \sum_{s=1}^{t-1} \sum_{j=1}^n [P^{t-s+1}]_{ij} g_j(s-1) - a_k g_i(t-1) \\ & + \sum_{j=1}^n [P^t]_{ij} z_j^k(1), \end{aligned} \quad (21)$$

and since the projection is non-expansive and $z_i^1(1) = 0, \forall i$,

$$\|z_i^{k+1}(1)\| = \|w_i^{k+1}(1)\| = \|w_i^k(T_k)\| = \|\Pi_{\mathcal{W}} [z_i^k(T_k)]\| \quad (22)$$

$$\leq \|z_i^k(T_k)\| \quad (23)$$

$$\begin{aligned} \leq & \left\| -a_k \sum_{t=1}^{T_k-1} \sum_{i=1}^n [P^{T_k-s+1}]_{ij} g_i(s-1) \right\| \\ & + \|-a_k g_i(T_k-1)\| + \sum_{j=1}^n [P^{T_k}]_{ij} \|z_j^k(1)\| \end{aligned} \quad (24)$$

$$\leq a_k T_k L + \sum_{j=1}^n [P^{T_k}]_{ij} \|z_j^k(1)\| \quad (25)$$

$$\leq \dots \quad (26)$$

$$\leq L \sum_{s=1}^k a_s T_s. \quad (27)$$

D. Analysis of One Round of DOGD

Next, we focus on bounding the amount of regret accumulated during the k th round of DOGD (lines 5–12 of Algorithm 1) during which the learning rate remains fixed at a_k . Using Assumptions 1, 2, and the triangle inequality

we have that

$$\begin{aligned} & \sum_{t=1}^{T_k} [F(w_i^k(t)) - F(w^*)] \\ &= \sum_{t=1}^{T_k} [F(\bar{w}^k(t)) - F(w^*) + F(w_i^k(t)) - F(\bar{w}^k(t))] \end{aligned} \quad (28)$$

$$\leq \sum_{t=1}^{T_k} [F(\bar{w}^k(t)) - F(w^*) + L \|w_i^k(t) - \bar{w}^k(t)\|] \quad (29)$$

$$\begin{aligned} & \leq \sum_{t=1}^{T_k} \frac{1}{n} \sum_{i=1}^n [f_i(w_i^k(t)) - f_i(w^*)] \\ & \quad + \sum_{t=1}^{T_k} \frac{1}{n} \sum_{i=1}^n [f_i(\bar{w}^k(t)) - f_i(w_i^k(t))] \\ & \quad + \sum_{t=1}^{T_k} L \|w_i^k(t) - \bar{w}^k(t)\| \end{aligned} \quad (30)$$

$$\begin{aligned} & \leq \underbrace{\sum_{t=1}^{T_k} \frac{1}{n} \sum_{i=1}^n \langle g_i(t), w_i^k(t) - w^* \rangle}_{A_1} \\ & \quad + \sum_{t=1}^{T_k} \frac{1}{n} \sum_{i=1}^n L \|\bar{w}^k(t) - w_i^k(t)\| \\ & \quad + \sum_{t=1}^{T_k} L \|w_i^k(t) - \bar{w}^k(t)\|. \end{aligned} \quad (31)$$

For the first summand we have

$$A_1 = \sum_{t=1}^{T_k} \frac{1}{n} \sum_{i=1}^n \langle g_i(t), w_i^k(t) - w^* \rangle \quad (32)$$

$$\begin{aligned} & \leq \sum_{t=1}^{T_k} \frac{1}{n} \sum_{i=1}^n \langle g_i(t), \bar{w}^k(t) - w^* \rangle \\ & \quad + \sum_{t=1}^{T_k} \frac{1}{n} \sum_{i=1}^n \langle g_i(t), w_i^k(t) - \bar{w}^k(t) \rangle \end{aligned} \quad (33)$$

$$\begin{aligned} & \leq \underbrace{\sum_{t=1}^{T_k} \frac{1}{n} \sum_{i=1}^n \langle g_i(t), \bar{w}^k(t) - w^* \rangle}_{A_2} \\ & \quad + \sum_{t=1}^{T_k} \frac{1}{n} \sum_{i=1}^n L \|w_i^k(t) - \bar{w}^k(t)\|. \end{aligned} \quad (34)$$

To bound term A_2 we invoke Theorem 2 for the average sequences $\{\bar{w}^k(t)\}$ and $\{\bar{z}^k(t)\}$.

$$A_2 = \sum_{t=1}^{T_k} \frac{1}{n} \sum_{i=1}^n \langle g_i(t), \bar{w}^k(t) - w^* \rangle \quad (35)$$

$$= \sum_{t=1}^{T_k} \left\langle \frac{1}{n} \sum_{i=1}^n g_i(t), \Pi_{\mathcal{W}} [\bar{z}^k(t)] - w^* \right\rangle \quad (36)$$

$$= \sum_{t=1}^{T_k} \left\langle \bar{g}(t), \Pi_{\mathcal{W}} [\bar{z}^k(t)] - w^* \right\rangle \quad (37)$$

$$\leq \frac{\|\bar{w}^k(1) - w^*\|^2}{2a_k} + \frac{T_k a_k \left\| \frac{1}{n} \sum_{i=1}^n g_i(t) \right\|^2}{2} \quad (38)$$

$$= \frac{\|\bar{w}^k(1) - w^*\|^2}{2a_k} + \frac{T_k a_k L^2}{2}. \quad (39)$$

Collecting now all the partial results and bounds, so far we have shown that

$$\begin{aligned} \sum_{t=1}^{T_k} [F(w_i^k(t)) - F(w^*)] &\leq \frac{\|\bar{w}^k(1) - w^*\|^2}{2a_k} + \frac{T_k a_k L^2}{2} \\ &\quad + \sum_{t=1}^{T_k} \frac{2}{n} \sum_{i=1}^n L \|w_i^k(t) - \bar{w}^k(t)\| \\ &\quad + \sum_{t=1}^{T_k} L \|w_i^k(t) - \bar{w}^k(t)\|. \end{aligned} \quad (40)$$

and since the projection operator is non-expansive, we have

$$\begin{aligned} \sum_{t=1}^{T_k} [F(w_i^k(t)) - F(w^*)] &\leq \frac{\|\bar{w}^k(1) - w^*\|^2}{2a_k} + \frac{T_k a_k L^2}{2} \\ &\quad + \sum_{t=1}^{T_k} \frac{2}{n} \sum_{i=1}^n L \|z_i^k(t) - \bar{z}^k(t)\| \\ &\quad + \sum_{t=1}^{T_k} L \|z_i^k(t) - \bar{z}^k(t)\|. \end{aligned} \quad (41)$$

The first two terms are standard for subgradient algorithms using a constant step size. The last two terms depend on the error between each node's iterate $z_i^k(t)$ and the network-wide average $\bar{z}^k(t)$, which we bound next.

E. Bounding the Network Error

What remains is to bound the term $\|z_i^k(t) - \bar{z}^k(t)\|$ which describes an error induced by the network since the different nodes do not agree on the direction towards the optimum. By recalling that P is doubly stochastic and manipulating the recursive expressions (21) and (20) for $z_i(t)$ and $\bar{z}^k(t)$ using arguments similar to those in [8],

[14], we obtain the bound,

$$\begin{aligned} \|z_i^k(t) - \bar{z}^k(t)\| &\leq a_k L \sum_{s=1}^{t-1} \sum_{j=1}^n \left| \frac{1}{n} \mathbf{1}^T - [P^{t-s-1}]_{ij} \right|_1 + 2a_k L \\ &\quad + \sum_{j=1}^n \left| \frac{1}{n} - [P^t]_{ij} \right| \|z_j^k(1)\| \end{aligned} \quad (42)$$

$$\begin{aligned} &= a_k L \sum_{s=1}^{t-1} \left\| \frac{1}{n} \mathbf{1}^T - [P^{t-s-1}]_{i,:} \right\|_1 + 2a_k L \\ &\quad + \sum_{j=1}^n \left| \frac{1}{n} - [P^t]_{ij} \right| \|z_j^k(1)\|. \end{aligned} \quad (43)$$

The ℓ_1 norm can be bounded using Lemma 2, which is stated and proven in the Appendix, and using (27) we arrive at

$$\|z_i^k(t) - \bar{z}^k(t)\| \leq 2a_k L \frac{\log(T_k \sqrt{n})}{1 - \sqrt{\lambda_2}} + 3a_k L + \frac{L \sum_{s=1}^{k-1} a_s T_s}{T_k} \quad (44)$$

where λ_2 is the second largest eigenvalue of P . Using this bound in equation (41), along with the fact that $F(w)$ is convex, we conclude that

$$F(\hat{w}_i^{k+1}) - F(w^*) = F\left(\frac{1}{T_k} \sum_{t=1}^{T_k} w_i^k(t)\right) - F(w^*) \quad (45)$$

$$\leq \frac{1}{T_k} \sum_{t=1}^{T_k} [F(w_i^k(t)) - F(w^*)] \quad (46)$$

$$\begin{aligned} &\leq \frac{\|\bar{w}^k(1) - w^*\|^2}{2a_k T_k} + \frac{a_k L^2}{2} \\ &\quad + L^2 a_k \left[6 \frac{\log(T_k \sqrt{n})}{1 - \sqrt{\lambda_2}} + 9 \right] \\ &\quad + \frac{3L^2 \sum_{s=1}^{k-1} a_s T_s}{T_k}, \end{aligned} \quad (47)$$

where $\bar{w}^k(1) = \Pi_{\mathcal{W}} \left[\frac{1}{n} \sum_{i=1}^n z_i^k(1) \right]$.

F. Analysis of DOGD over Multiple Rounds

As our last intermediate step, we must control the learning rate and update of T_k from round-to-round to ensure linear convergence of the error. From strong convexity of F we have

$$\|\bar{w}^k(1) - w^*\|^2 \leq 2 \frac{F(\bar{w}^k(1)) - F(w^*)}{\sigma} \quad (48)$$

and thus

$$\begin{aligned} F(\hat{w}_i^{k+1}) - F(w^*) &\leq \frac{F(\bar{w}^k(1)) - F(w^*)}{\sigma a_k T_k} \\ &\quad + \frac{L^2 a_k}{2} \left[12 \frac{\log(T_k \sqrt{n})}{1 - \sqrt{\lambda_2}} + 19 \right] \\ &\quad + \frac{3L^2 \sum_{s=1}^{k-1} a_s T_s}{T_k}. \end{aligned} \quad (49)$$

Now, from Theorem 3 in [1] which is a direct consequence of Theorem 2 for the average sequence \bar{w} viewed as a single processor lazy projection algorithm, we have that after executing T_{k-1} gradient steps in round $k-1$,

$$F(\bar{w}^k(1)) - F(w^*) \leq \frac{\|\bar{w}^{k-1}(1) - w^*\|^2}{2a_{k-1}T_{k-1}} + \frac{a_{k-1}L^2}{2} \quad (50)$$

and by repeatedly using strong convexity and Theorem 2 we see that

$$F(\bar{w}^k(1)) - F(w^*) \leq \frac{F(\bar{w}^{k-1}(1)) - F(w^*)}{\sigma a_{k-1}T_{k-1}} + \frac{a_{k-1}L^2}{2} \quad (51)$$

$$\leq \dots \quad (52)$$

$$\begin{aligned} &\leq \frac{F(\bar{w}^1(1)) - F(w^*)}{\prod_{j=0}^{k-1} (\sigma a_{k-j} T_{k-j})} \\ &\quad + \sum_{j=1}^{k-1} \frac{a_{k-j} L^2}{2 \prod_{s=1}^{j-1} (\sigma a_{k-s} T_{k-s})}. \end{aligned} \quad (53)$$

Now, let us fix positive integers b and c , and suppose we use the following rules to determine the step size and number of updates performed within each round:

$$a_k = \frac{a_{k-1}}{b} = \dots = \frac{a_1}{b^{k-1}} \quad (54)$$

$$T_k = cT_{k-1} = \dots = c^{k-1}T_1. \quad (55)$$

Combining (53) with (49) and invoking Lemma 1, we have

$$\begin{aligned} F(\hat{w}_i^{k+1}) - F(w^*) &\leq \frac{2L^2}{\sigma \prod_{j=0}^{k-1} (\sigma a_1 T_1 (\frac{c}{b})^{k-j-1})} \\ &\quad + \sum_{j=1}^{k-1} \frac{a_1 L^2}{2b^{k-j-1} \prod_{s=0}^{j-1} (\sigma a_1 T_1 (\frac{c}{b})^{k-s-1})} \\ &\quad + \frac{L^2 a_1}{2b^{k-1}} \left[12 \frac{\log(T_1 c^{k-1} \sqrt{n})}{1 - \sqrt{\lambda_2}} + 19 \right] \\ &\quad + \frac{3L^2 \sum_{s=1}^{k-1} a_1 T_1 (\frac{c}{b})^{s-1}}{T_1 c^{k-1}}. \end{aligned} \quad (56)$$

To ensure convergence to zero, we need $c \geq b$ and $\sigma a_1 T_1 > 1$ or $a_1 > \frac{1}{T_1 \sigma}$. Given these restrictions, let us make the choices

$$a_1 = 1, \quad T_1 = \left\lceil \frac{2}{\sigma} \right\rceil, \quad c = b = 2. \quad (57)$$

To simplify the exposition, let us assume that $T_1 = \frac{2}{\sigma}$ is an integer. Using the selected values, we obtain

$$\begin{aligned}
F(\hat{w}_i^{k+1}) - F(w^*) &\leq \frac{2L^2}{\sigma \prod_{j=0}^{k-1} \left(2 \left(\frac{2}{2}\right)^{k-j-1}\right)} \\
&\quad + \sum_{j=1}^{k-1} \frac{L^2}{2 \cdot 2^{k-j-1} \prod_{s=0}^{j-1} \left(2 \left(\frac{2}{2}\right)^{k-s-1}\right)} \\
&\quad + \frac{L^2}{2 \cdot 2^{k-1}} \left[12 \frac{\log\left(\frac{2}{\sigma} \cdot 2^{k-1} \sqrt{n}\right)}{1 - \sqrt{\lambda_2}} + 19 \right] \\
&\quad + \frac{3L^2 \sum_{s=1}^{k-1} \left(\frac{2}{2}\right)^{s-1}}{2^{k-1}} \tag{58}
\end{aligned}$$

$$\begin{aligned}
&\leq \frac{2L^2}{\sigma 2^k} + \sum_{j=1}^{k-1} \frac{L^2}{2^{k-j} 2^j} \\
&\quad + \frac{L^2}{2^k} \left[12 \frac{\log\left(\frac{2^k \sqrt{n}}{\sigma}\right)}{1 - \sqrt{\lambda_2}} + 19 \right] + \frac{3L^2(k-1)}{2^{k-1}} \tag{59}
\end{aligned}$$

$$\begin{aligned}
&\leq \frac{2L^2}{\sigma 2^k} + \frac{L^2(k-1)}{2^k} \\
&\quad + \frac{L^2}{2^k} \left[12 \frac{\log\left(\frac{2^k \sqrt{n}}{\sigma}\right)}{1 - \sqrt{\lambda_2}} + 19 \right] + \frac{6L^2(k-1)}{2^k} \tag{60}
\end{aligned}$$

$$\begin{aligned}
&\leq \frac{2L^2}{\sigma 2^k} + \frac{L^2(k-1)}{2^k} \\
&\quad + \frac{L^2}{2^k} \left[12 \frac{\log\left(\frac{2^k \sqrt{n}}{\sigma}\right)}{1 - \sqrt{\lambda_2}} + 19 \right] + \frac{6L^2(k-1)}{2^k}. \tag{61}
\end{aligned}$$

Finally, we have all we need to complete the analysis of Algorithm 1.

G. Proof of Theorem 1

Suppose we run Algorithm 1 for T total steps at each node. This allows for \tilde{k} rounds, where \tilde{k} is determined by solving

$$\sum_{i=1}^{\tilde{k}} T_i \leq T \iff \sum_{i=1}^{\tilde{k}} 2 \cdot 2^i \leq T \iff \tilde{k} \leq \log_2 \left(\frac{T}{2} + 1 \right). \tag{62}$$

Using this value for k we see that

$$F(\hat{w}_i^{\tilde{k}+1}) - F(w^*) \leq \frac{L^2}{\sigma} 2^{\tilde{k}} + \frac{L^2(\tilde{k}-1)}{2^{\tilde{k}}} + \frac{L^2}{2^{\tilde{k}}} \left[12 \frac{\log\left(\frac{2^{\tilde{k}}\sqrt{n}}{\sigma}\right)}{1-\sqrt{\lambda_2}} + 19 \right] + \frac{6L^2(\tilde{k}-1)}{2^{\tilde{k}}} \quad (63)$$

$$\begin{aligned} &\leq \frac{L^2}{\sigma\left(\frac{T}{2}+1\right)} + \frac{L^2(\log_2\left(\frac{T}{2}+1\right)-1)}{\left(\frac{T}{2}+1\right)} \\ &\quad + \frac{L^2}{\left(\frac{T}{2}+1\right)} \left[12 \frac{\log\left(\frac{\left(\frac{T}{2}+1\right)\sqrt{n}}{\sigma}\right)}{1-\sqrt{\lambda_2}} + 19 \right] \\ &\quad + \frac{6L^2\left(\left(\frac{T}{2}+1\right)-1\right)}{\left(\frac{T}{2}+1\right)} \\ &= O\left(\frac{\log(\sqrt{n}T)}{T(1-\sqrt{\lambda_2})}\right) = O\left(\frac{\log(\sqrt{n}T)}{T}\right), \end{aligned} \quad (64)$$

when λ_2 is constant and does not scale with n , and this concludes the proof of Theorem 1.

V. EXTENSION TO STOCHASTIC OPTIMIZATION

The proof presented in the previous section can easily be extended to the case where each node receives a random estimate $\hat{g}(t)$ of the gradient, satisfying $\mathbb{E}[\hat{g}(t)] = g(t)$, instead of receiving $g(t)$ directly. We assume that noisy gradients still have bounded variance i.e., $\mathbb{E}[\|\hat{g}_i(t)\|^2] \leq L^2$. In this setting, instead of equation (35), we have

$$A_2 = \sum_{t=1}^{T_k} \frac{1}{n} \sum_{i=1}^n \langle g_i(t), \bar{w}^k(t) - w^* \rangle \quad (65)$$

$$\begin{aligned} &= \sum_{t=1}^{T_k} \left\langle \frac{1}{n} \sum_{i=1}^n \hat{g}_i(t), \bar{w}^k(t) - w^* \right\rangle \\ &\quad + \sum_{t=1}^{T_k} \frac{1}{n} \sum_{i=1}^n \langle g_i(t) - \hat{g}_i(t), \bar{w}^k(t) - w^* \rangle. \end{aligned} \quad (66)$$

However, the proof of Theorem 2 does not depend on the gradients being correct; rather, it holds for noisy gradients $\hat{g}(t)$ as well. Moreover, we have $\mathbb{E}[\|\hat{g}_i(t)\|] \leq L$, and by Hölder's inequality $\mathbb{E}[\|\hat{g}_i(t)\| \|\hat{g}_j(t)\|] \leq L^2$. Thus,

$$\mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n \hat{g}_i(t) \right\|^2 \right] \leq \frac{1}{n^2} \sum_{i,j=1}^n \mathbb{E}[\|\hat{g}_i(t)\| \|\hat{g}_j(t)\|] \leq L^2. \quad (67)$$

Thus, invoking Theorem 2, if the new data and thus the subgradients are independent of the past, and since $\mathbb{E}[\hat{g}_i(t)] = g_i(t)$, we have

$$\begin{aligned} \mathbb{E}[A_2] &\leq \frac{\|\bar{w}^k(1) - w^*\|^2}{2a_k} + \frac{T_k a_k L^2}{2} \\ &\quad + \mathbb{E}\left[\sum_{t=1}^{T_k} \frac{1}{n} \sum_{i=1}^n \langle g_i(t) - \hat{g}_i(t), \bar{w}^k(t) - w^* \rangle\right] \end{aligned} \quad (68)$$

$$\begin{aligned} &= \frac{\|\bar{w}^k(1) - w^*\|^2}{2a_k} + \frac{T_k a_k L^2}{2} \\ &\quad + \sum_{t=1}^{T_k} \frac{1}{n} \sum_{i=1}^n \langle \mathbb{E}[g_i(t) - \hat{g}_i(t)], \bar{w}^k(t) - w^* \rangle \end{aligned} \quad (69)$$

$$= \frac{\|\bar{w}^k(1) - w^*\|^2}{2a_k} + \frac{T_k a_k L^2}{2}. \quad (70)$$

Furthermore, the network error bound holds in expectation as well, i.e.,

$$\begin{aligned} \mathbb{E}[\|\bar{w}^k(t) - w_i^k(t)\|] &\leq \mathbb{E}[\|\bar{z}^k(t) - z_i^k(t)\|] \\ &\leq 2a_k L \frac{\log(T_k \sqrt{n})}{1 - \sqrt{\lambda_2}} + 3a_k L + \frac{L \sum_{s=1}^{k-1} a_s T_s}{T_k} \end{aligned} \quad (71)$$

Collecting all these observations we have shown that, in expectation,

$$\begin{aligned} \mathbb{E}[F(\hat{w}_i^{k+1}) - F(w^*)] &\leq \frac{\|\bar{w}^k(1) - w^*\|^2}{2a_k T_k} + \frac{a_k L^2}{2} \\ &\quad + L^2 a_k \left[6 \frac{\log(T_k \sqrt{n})}{1 - \sqrt{\lambda_2}} + 9 \right] \\ &\quad + \frac{3L^2 \sum_{s=1}^{k-1} a_s T_s}{T_k} \end{aligned} \quad (72)$$

which, after using the update rules for a_k and T_k , is exactly the same rate as before. We note however that there may still be room for improvement in the distributed stochastic optimization setting since [12] describes a single-processor algorithm that converges at a rate $O(\frac{1}{T})$.

VI. SIMULATION

To illustrate the performance of DOGD we simulate online training of a classifier by solving the problem (4) using a network of 10 nodes arranged as a random geometric graph. Each node is given $T = 600$ data points, and the input dimension is $d = 100$. We set $\sigma = 0.1$ and generate the data from a standard normal distribution and classify them as -1 or 1 depending on their relative position to a randomly drawn hyperplane in \mathbb{R}^d . As we see in Figure 1, DODG minimizes the objective much faster than Distributed Dual Averaging (DDA) [8] which has a convergence rate of $O\left(\frac{\log(T\sqrt{n})}{\sqrt{T}}\right)$. DDA is simulated using the learning rate that is suggested in [8]. We have observed that boosting this learning rate may yield faster convergence, but still not as fast as DOGD. Figure 1 also shows the performance of a version of Fast Distributed Gradient Descent (FDGD) [11]. As we can see, FDGD fails to converge in an online or stochastic setting and ends up oscillating.




Fig. 1. Optimization of a $d = 100$ dimensional problem of the form (4) with a random network of 10 nodes. Our proposed algorithm DOGD(red) converges faster than DDA(green) as expected from the T instead of \sqrt{T} in the denominator of the convergence rate bound. FDGD(black), is unable to converge in the online problem.

VII. FUTURE WORK

In this paper we have proposed and analyzed a novel distributed optimization algorithm which we call Distributed Online Gradient Descent (DOGD). Our analysis shows that DOGD converges at a rate $O(\frac{\log(\sqrt{nT})}{T})$ when solving online, stochastic or batch constrained convex optimization problems if the objective function is strongly convex. This rate is optimal in the number of iterations for the online and batch setting and slower than a serial algorithm only by a logarithmic factor in the stochastic optimization setting.

In its current form, DOGD requires the nodes in the network to exchange gradient information at every iteration. Our preliminary investigation suggests that gradually performing more and more updates between each communication can speed up distributed optimization algorithms in the batch setting when one explicitly accounts for the time required to communicate data. Our future work will carry out a similar analysis for online and stochastic optimization algorithms.

APPENDIX

Lemma 2: If P is a doubly stochastic matrix defined over a strongly connected graph $G = (V, E)$ with $|V| = n$ nodes so that $p_{ji} = 0$ if $(i, j) \notin E$, then for any $t \leq T$,

$$\sum_{s=1}^{t-1} \left\| \frac{1}{n} \mathbf{1}^T - [P^{t-s+1}]_{i,:} \right\|_1 \leq 1 + \frac{\log(T\sqrt{n})}{1 - \sqrt{\lambda_2}} \quad (73)$$

where λ_2 is the second largest eigenvalue of P .

Proof: If the consensus matrix P is doubly stochastic it is straightforward to show that $P^t \rightarrow \frac{1}{n} \mathbf{1} \mathbf{1}^T$ as $t \rightarrow \infty$. Moreover, from standard Perron-Frobenius it is easy to show (see e.g., [17])

$$\left\| \frac{1}{n} \mathbf{1}^T - [P^t]_{i,:} \right\|_1 = 2 \left\| \frac{1}{n} \mathbf{1}^T - [P^t]_{i,:} \right\|_{TV} \leq \sqrt{n} \left(\sqrt{\lambda_2} \right)^t \quad (74)$$

so in our case $\left\| \frac{1}{n} \mathbf{1}^T - [P^{t-s+1}]_{i,:} \right\|_1 \leq \sqrt{n} (\sqrt{\lambda_2})^{t-s+1}$. Next, demand that the right hand side bound is less than $\sqrt{n}\delta$ with δ to be determined:

$$\sqrt{n} (\sqrt{\lambda_2})^{t-s+1} \leq \sqrt{n}\delta \Rightarrow t-s+1 \geq \frac{\log(\delta^{-1})}{\log(\sqrt{\lambda_2}^{-1})}. \quad (75)$$

So with the choice $\delta^{-1} = \sqrt{nT}$,

$$\left\| \frac{1}{n} \mathbf{1}^T - [P^{t-s+1}]_{i,:} \right\|_1 \leq \sqrt{n} \frac{1}{\sqrt{nT}} = \frac{1}{T} \quad (76)$$

if $t-s+1 \geq \frac{\log(\delta^{-1})}{\log(\sqrt{\lambda_2}^{-1})} = \hat{t}$. When s is large and $t-s+1 < \hat{t}$ we take $\left\| \frac{1}{n} \mathbf{1}^T - [P^{t-s+1}]_{i,:} \right\|_1 \leq 2$. The desired bound is not obtained as follows

$$\sum_{s=1}^{t-1} \left\| \frac{1}{n} \mathbf{1}^T - [P^{t-s+1}]_{i,:} \right\|_1 = \sum_{s=1}^{t-\hat{t}-1} \left\| \frac{1}{n} \mathbf{1}^T - [P^{t-s+1}]_{i,:} \right\|_1 \quad (77)$$

$$+ \sum_{s=t-\hat{t}}^{t-1} \left\| \frac{1}{n} \mathbf{1}^T - [P^{t-s+1}]_{i,:} \right\|_1$$

$$\leq \sum_{s=1}^{t-\hat{t}-1} \frac{1}{T} + \sum_{s=t-\hat{t}}^{t-1} 2 \quad (78)$$

$$\leq \frac{t-\hat{t}}{T} + 2\hat{t} \leq 1 + 2\hat{t} \quad (79)$$

Since $t \leq T$ we know that $t - \hat{t} < T$. Moreover, $\log(\sqrt{\lambda_2})^{-1} \geq 1 - \sqrt{\lambda_2}$. Using these two facts we arrive at the result. The same bound is true for any individual entry of P^t approaching $\frac{1}{n}$. ■

REFERENCES

- [1] M. A. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *20th International Conference on Machine Learning (ICML)*, 2003.
- [2] Y. Nesterov, "Primal-dual subgradient methods for convex problems," *Mathematical Programming Series B*, vol. 120, pp. 221–259, 2009.
- [3] L. Bottou, "Large-scale machine learning with stochastic gradient descent," in *Proceedings of the 19th International Conference on Computational Statistics*, Y. Lechevallier and G. Saporta, Eds., Paris, France, August 2010, pp. 177–187.
- [4] S. Shalev-Shwartz and Y. Singer, "Logarithmic regret algorithms for strongly convex repeated games," in *The Hebrew University*, 2007.
- [5] E. Hazan, A. Kalai, S. Kale, and A. Agarwal, "Logarithmic regret algorithms for online convex optimization," in *19th COLT*, 2006, pp. 499–513.
- [6] P. L. Bartlett, E. Hazan, and A. Rakhlin, "Adaptive online gradient descent," in *Advances in Neural Information Processing Systems 20*, J. C. Platt, D. Koller, Y. Singer, and S. Roweis, Eds. MIT Press, 2007.
- [7] P. Tseng, "On accelerated proximal gradient methods for convex-concave optimization," *SIAM Journal on Optimization*, vol. (Submitted), 2008.
- [8] J. Duchi, A. Agarwal, and M. Wainwright, "Dual averaging for distributed optimization: Convergence analysis and network scaling," *IEEE Transactions on Automatic Control*, vol. 57, no. 3, pp. 592–606, 2011.
- [9] A. Nedic, A. Ozdaglar, and P. A. Parrilo, "Constrained consensus and optimization in multi-agent networks," *IEEE Transactions on Automatic Control*, vol. 55, no. 4, pp. 922–938, 2010.
- [10] O. Dekel, R. Gilad-Bachrach, O. Shamir, and L. Xiao, "Optimal distributed online prediction using mini-batches," *Journal of Machine Learning Research*, vol. 13, pp. 165–202, 2012.
- [11] D. Jakovetic, J. Xavier, and J. M. Moura, "Fast distributed gradient methods," *arXiv:1112.2972v1*, 2011.

- [12] E. Hazan and S. Kale, “Beyond the regret minimization barrier: an optimal algorithm for stochastic strongly-convex optimization,” in *24th Annual Conference on Learning Theory (COLT)*, 2011.
- [13] K. I. Tsianos, S. Lawlor, and M. G. Rabbat, “Push-sum distributed dual averaging for convex optimization,” in *51st IEEE Conference on Decision and Control*, 2012.
- [14] K. I. Tsianos and M. G. Rabbat, “Distributed dual averaging for convex optimization under communication delays,” in *American Control Conference (ACC)*, 2012.
- [15] S. S. Ram, A. Nedic, and V. V. Veeravalli, “Distributed stochastic subgradient projection algorithms for convex optimization,” *Journal of Optimization Theory and Applications*, vol. 147, no. 3, pp. 516–545, 2011.
- [16] O. Reingold, S. Vadhan, and A. Wigderson, “Entropy waves, the zig-zag graph product, and new constant-degree expanders,” *Annals of Mathematics*, vol. 155, no. 2, pp. 157–187, 2002.
- [17] P. Diaconis and D. Stroock, “Geometric bounds for eigenvalues of markov chains,” *The Annals of Applied Probability*, vol. 1, no. 1, pp. 36–61, 1991.