Optimal distributed convex optimization on slowly time-varying graphs

Alexander Rogozin* César A. Uribe* Alexander Gasnikov Nikolay Malkovsky Angelia Nedić

Abstract—We study optimal distributed first-order optimization algorithms when the network (i.e., communication constraints between the agents) changes with time. This problem is motivated by scenarios where agents experience network malfunctions. Under specific constraints on the dual function, we provide a sufficient condition that guarantees a convergence rate with optimal (up to logarithmic terms) dependencies on the network and function parameters if the network changes are constrained to a small percentage α of the total number of iterations. We call such networks *slowly* time-varying networks. Moreover, we show that Nesterov's method has an iteration complexity of $\Omega\left(\left(\sqrt{\kappa_\Phi \cdot \bar{\chi}} + \alpha \log(\kappa_\Phi \cdot \bar{\chi})\right) \log(1/\varepsilon)\right)$ for decentralized algorithms, where κ_Φ is condition number of the objective function, and $\bar{\chi}$ is a worst case bound on the condition number of the sequence of communication graphs. Additionally, we provide an explicit upper bound on α in terms of the condition number of the objective function and network topologies.

Index Terms—distributed optimization, time-varying graph, accelerated method.

I. INTRODUCTION

INCREASING amounts of data and privacy constraints in distributed storage systems, as well as the distributed nature of data sources, has driven the development of distributed optimization algorithms that can be executed over networks. For example, consider the machine learning problem with a vector of parameters $y \in \mathbb{R}^d$ and a loss function $L(\mathbf{A}, y)$, where $\mathbf{A}$ is a training set of l samples, and each sample is a vector of $\mathbb{R}^m$. Moreover, assume the dataset $\mathbf{A}$ is not available in the memory of a single computer due to its size and communication costs, but is divided into n parts $\{\mathbf{A_i}\}_{i=1}^n$ and stored on n different machines. Therefore, one seeks to take into account the information constraints induced by the distributed nature of the data.

The distributed data generation and storage requires the study of the fundamental performance limits of distributed optimization algorithm that can be executed over a network [1]–[6]. One primary objective is to understand whether one can achieve the same convergence rate of centralized algorithms by using distributed methods. In [7], the authors first showed that distributed algorithms could achieve linear convergence rates when optimizing sums of strongly convex and smooth functions, in comparison with previous algorithms such as the distributed sub-gradient [8]. In [9], [10], the authors show that one can accelerate distributed algorithms and achieve convergence rate close to centralized methods. However, dependencies on the function parameter and network topology were not optimal. In [1], the authors proposed a dual-based approach [11], [12] and provided the first result on complexity lower bounds for distributed optimization over networks for sums of strongly convex and smooth functions. Later in [2], the authors extended these results to non-smooth problems or non-strongly convex problems. It was shown that distributed optimization algorithms could achieve the same convergence rates as their centralized counterparts with an additional multiplicative cost related to the communication network. Recent work has shown complexity lower bounds for distributed non-convex optimization problems as well [5].

Existing approaches show optimal convergence rates where graphs are assumed fixed [1], [13], [14]. We focus on the case where the network is allowed to *change with time*. These changes occur, for example, due to technical malfunctions and loss of connectivity between nodes [13]. As a result, the change in topology induces a change in the distributed problem formulation and leads to time-varying optimization. These changes can come both from changing cost function and changing constraints [15], [16]. Time-varying problems with continuous time [17], [18] and discrete time [19], [20] have been studied before. Also, there are distributed algorithms that can be executed over time-varying networks and achieve linear convergence rates, such as DIGing [7], Push-Pull Gradient Method [21], PANDA [22] and Acc-DNGD [9]. Nevertheless, optimal convergence rate dependencies with respect to the network and function parameters are not yet fully understood. We provide a comparative performance analysis with such methods.

In this paper, we study the convex optimization problem

$$\varphi(y) = \sum_{i=1}^n \varphi_i(y) \longrightarrow \min_{y \in \mathbb{R}^d}, \qquad (1)$$

where $\varphi_i : \mathbb{R}^d \to \mathbb{R}$ is a convex function for each $i = 1, \cdots, n$. We focus on the distributed problem where each of the functions φ_i is privately held by a computational entity in a network. That is, each node or agent i on a network has

access to φ_i only, and yet, the group of agents seek to solve the optimization problem in (1) by repeated interactions with other agents following the communication constraints imposed by the network. The interactions between the agents are driven by a sequence of graphs $\{\mathcal{G}_k\}_{k=1}^{\infty}$, where $\mathcal{G}_k = (V, E_k)$ is a connected undirected graph with $V = \{1, \ldots, n\}$ and E_k is a set of edges such that $(j, i) \in E_k$ if a pair of nodes $i, j \in V$ can communicate at time instant k, see Figure 1.
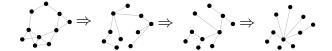

Fig. 1. A sequence of graphs with 10 nodes each, but with different edge set at each time instant.

We analyze how first-order methods behave when the objective function changes from time to time (under some restrictions to be specified later), with a special interest in distributed optimal accelerated methods [13]. Our main contribution is a sufficient condition that certifies that an optimal rate can be achieved, with an additional cost proportional to the number of changes in the network, expressed as a fraction of the number of iterations. Whether this condition is also necessary is left for future work.

This paper is organized as follows. Problem statement and dual formulation are described in Section II. Section III presents the general analysis for first-order methods with changing functions; we include the convergence rates of the gradient descent and Nesterov's fast gradient method. In Section IV we build upon the results of Section III to propose an algorithm for distributed optimization over time-varying graphs and we provide its convergence rate. We compare our results with other distributed algorithms in Section V. In Section VI, we provide numerical experiments to illustrate and numerically evaluate our theoretical results. Finally, some conclusive notes and future work are described in Section VII.

II. PROBLEM STATEMENT

In this section, initially we recall some basic definitions, then we present a formulation of the distributed optimization problem that incorporates the communication constraints induced by the network. The use of the network constraints allows for the formulation of a dual problem with a suitable structure for distributed computation. Finally, we formally pose the problem of distributed optimization over time-varying networks.

A. Preliminaries

This paper is focused on μ-strongly convex L-smooth functions.

Definition II.1. *Let f be a differentiable function on $\mathbb{X}$, where $\mathbb{X} = \mathbb{R}^n$ or $\mathbb{X} = \mathbb{R}^{d \times n}$. We say that f μ-strongly convex ($\mu > 0$) w.r.t. $\|\cdot\|$ if*

$$\forall x, y \in \mathbb{X}: \ f(y) \geqslant f(x) + \langle \nabla f(x), y - x \rangle + \frac{\mu}{2}\|y - x\|^2.$$

Moreover, we say that f L-smooth w.r.t. $\|\cdot\|$ if $\nabla f(x)$ is L-Lipschitz continuous w.r.t. to the dual norm $\|\cdot\|_*$, i.e.,

$$\forall x, y \in \mathbb{X}: \ f(y) \leqslant f(x) + \langle \nabla f(x), y - x \rangle + L/2\|y - x\|^2.$$

For simplicity of exposition we will present our results mainly on the 2-norm $\|\cdot\|_2$ in $\mathbb{R}^n$ and Frobenius norm $\|\cdot\|_F$ in $\mathbb{R}^{d \times n}$. Note that $\|\cdot\|_{2*} = \|\cdot\|_2$ and $\|\cdot\|_{F*} = \|\cdot\|_F$. We also denote $\|\cdot\|_{op}$ the operator norm in $\mathbb{R}^{n \times n}$ generated by $\|\cdot\|_2$, which is define as $\|A\|_{op} = \sup_{x \in \mathbb{R}^n} \|Ax\|_2/\|x\|_2$.

Definition II.2. *The scalar product for $X, Y \in \mathbb{R}^{d \times n}$ is given by $\langle X, Y \rangle = \sum_{i=1}^{d} \sum_{j=1}^{n} X_{ij} Y_{ij}$. The Frobenius norm is given by $\|X\|_F = \sqrt{\langle X, X \rangle}$.*

Definition II.3. *Let $\mathbb{X}$ be a Euclidean space with a scalar product $\langle \cdot, \cdot \rangle$ and $\phi : \mathbb{X} \to \mathbb{R}$. Then the conjugate function to ϕ, denoted by ϕ^*, is given by $\phi^*(Y) = \sup_{X \in \mathbb{X}} \left(\langle X, Y \rangle - \phi(X) \right)$, and the dual norm $\|Y\|_*$ is defined as $\|Y\|_* = \sup_{X \in \mathbb{X}} \{ \langle X, Y \rangle : \|X\| \leqslant 1 \}$.*

B. Dual problem formulation for static graphs

Problem (1) can be equivalently rewritten as

$$\sum_{i=1}^{n} \varphi_i(y_i) \longrightarrow \min \ s.t. \ y_1 = \cdots = y_n, \ y_i \in \mathbb{R}^d \ \forall i \in V. \quad (2)$$

Additionally, the consensus constrains in (2) can be equivalently represented by the communication constraints imposed by the network topology. Particularly, we define the Laplacian of the graph $\mathcal{G}$ as: $[W]_{ij} = -1$ if $(i, j) \in E$, $[W]_{ij} = \deg(i)$ if $i = j$, and $[W]_{ij} = 0$ otherwise, where $\deg(i)$ is the degree of the node i, i.e., the number of neighbors of the node.

If a graph $\mathcal{G}$ is undirected and connected. Then, the Laplacian matrix W is symmetric and positive semidefinite. Moreover, by Perron-Frobenius theorem [23], it holds that $YW = 0$ if and only if $Y\sqrt{W} = 0$ if and only if $y_1 = \cdots = y_n$. Therefore, for static graphs (2) is equivalent to

$$\Phi(Y) = \sum_{i=1}^{n} \varphi_i(y_i) \longrightarrow \min_{Y\sqrt{W}=0}, \quad (3)$$

where $Y = [y_1, \ldots, y_n] \in \mathbb{R}^{d \times n}$. Note that Y is a matrix in $\mathbb{R}^{d \times n}$ consisting of local copies y_i of the decision vector y in the original problem (1). Using this equivalent constraint leads us to the dual function

$$f(X) = \max_{Y \in \mathbb{R}^{d \times n}} \left[-\langle X, Y\sqrt{W} \rangle - \Phi(Y) \right]. \quad (4)$$

Remark II.4. *It is not necessary to restrict our attention to the Laplacian of the graph as communication matrix. We may use arbitrary positive weights and weighted degrees that follow the same sparsity structure. This is equivalent to using constraints of form $\omega_{ij} y_i = \omega_{ji} y_j$ with $\omega_{ij} > 0$ instead of $y_i = y_j$. All the required properties are induced, and the rest of the analysis stays the same. On the other hand, it gives more flexibility for practical purposes, e.g., proper weight choice can induce better conditioning [24]. However, choosing specific weights is a separate question and therefore stays beyond our scope.*

Now that we have related the optimization problem with the network structure and the communication constraints it imposes, in this subsection, we show the connection between

properties of function $\Phi(Y)$ in (3) and the dual function $f(X)$ given in (4).

Lemma II.5 (Lemma 3.1 in [25], Proposition 12.60 in [26], Theorem 1 in [27], Theorem 6 in [28]). *Let ϕ be a closed convex function. Then ϕ is μ-strongly convex w.r.t. $\|\cdot\|$ if and only if f is $1/\mu$-smooth w.r.t. $\|\cdot\|_*$.*

Lemma II.5 allows us to establish the relationship between strong convexity and smoothness of functions $\Phi(Y)$ in (3) and $f(X)$ in (4). This relationship is formally stated in the next theorem, which extends the results of [28] on matrices.

Theorem II.6. *Let $\sigma_{\max}(W)$ be the largest eigenvalue and $\tilde{\sigma}_{\min}(W)$ be the least nonzero eigenvalue of $W^T W = W^2$, where W is the Laplacian of the communication graph $\mathcal{G} = (V, E)$. Let $\Phi(Y)$ be L_Φ-smooth and μ_Φ-strongly convex w.r.t. $\|\cdot\|_F$. Then, $f(X)$ in (4) is strongly convex with constant $\mu_f = \sqrt{\tilde{\sigma}_{\min}(W)}/L_\Phi$ on the subspace $(\mathrm{Ker}\, W)^\perp$ and smooth with constant $L_f = \sqrt{\sigma_{\max}(W)}/\mu_\Phi$ on $\mathbb{R}^{d \times n}$.*

The proof of Theorem II.6 is presented in Appendix A.

C. Dual problem over time-varying networks

We are now ready to discuss the distributed optimization problem when the communication network changes with time. Particularly, we explicitly define this *time-varying* setting as the case where the edge set changes. Thus, we consider a sequence of graphs $\{\mathcal{G}_k\}_{k=1}^\infty$, such that $\mathcal{G} = (V, E_k)$, i.e., the set of nodes remain the same but the edges might change with time. Therefore, the Laplacian matrix of the graph changes as well, which defines a sequence of graph Laplacians $\{W_k\}_{k=1}^\infty$. As a result, contrary to the fixed network setup, we work with a sequence of dual functions $f_k(X)$, such that

$$f_k(X) = \max_{Y \in \mathbb{R}^{d \times n}} \left(-\langle X, Y\sqrt{W_k}\rangle - \Phi(Y) \right). \tag{5}$$

We assume that the network is connected for all $k \geq 0$. Then, all W_k have the same nullspace: $\mathrm{Ker}(W_k) = \{y_1 = \ldots = y_n\} = \mathrm{Ker}(\sqrt{W_k})$. Consequently, the description of the constraint set changes from time to time, while the constraint set itself remains the same.

Remark II.7. *We impose a rather strong assumption of graphs being connected at every iteration. Such strong assumption is driven by the dual nature of the proposed algorithm, and the focus on guaranteeing optimal convergence rate dependencies on the function and network parameters. Primal-based methods allow for weaker connectivity assumptions, such as uniform connectivity, but to the best of the authors knowledge, optimal dependencies are not guaranteed [7], [9], [21], [22].*

Moreover, we define

$$\theta_{\max} = \sup_{k \geq 0} \{\sigma_{\max}(W_k)\}, \text{ and } \theta_{\min} = \inf_{k \geq 0} \{\tilde{\sigma}_{\min}(W_k)\}. \tag{6}$$

Then, by Theorem II.6, every $f_k(X)$ is μ-strongly convex on $\left(\mathrm{Ker}\, W\right)^\perp$ and L-smooth on $\mathbb{R}^n$, where $\mu = \sqrt{\theta_{\min}}/L_\Phi$, $L = \sqrt{\theta_{\max}}/\mu_\Phi$. We remark that what is changing with time is not the objective function Φ but the constraints representation, and thus the dual function f_k. Consequently all $f_k(X)$ have a common point of minimum and the same value of minimum. For clarity of exposition, we group some

properties of the dual function in the following assumption, which holds for (5) by definition.

Assumption II.8. *The sequence of convex functions $\{f_k(x)\}_{k=0}^\infty$ has the following properties:*
- *There is a point x^* which is a common minimum for all the functions f_k.*
- *Every function $f_k(x)$ is μ-strongly convex and L-smooth.*

Remark II.9. *In Assumption II.8, we require all functions f_k to have a common point of minimum. This statement holds, for example, if for each dual function f_k defined in (5) its minimizer $X_k^* \in \mathrm{Ker}\, W_k$, i.e. $[X_k^*]_1 = \ldots = [X_k^*]_n$, where $[X_k^*]_i$ denotes the i-th column of X_k^*.*

As an example, consider $\varphi(y) = \frac{1}{2}\sum_{i=1}^n \|y - a_i\|^2$, where $\sum_{i=1}^n a_i = 0$. Then

$$\Phi(Y) = \frac{1}{2}\|y_i - a_i\|^2 = \frac{1}{2}\|Y - A\|^2$$
$$F_k(X) = \arg\max_Y (-\Phi(Y) - \langle Y, XW_k\rangle)$$
$$= \frac{1}{2}\|XW_k - A\|^2 - \frac{1}{2}\|A\|^2$$
$$X_k^* = 0$$

Remark II.10. *Note that we require f_k to be strongly convex not only on $(\mathrm{Ker}\, W)^\perp$, but on the whole $\mathbb{R}^d$. This assumption is discussed in Section IV. Assumption II.8 does not mean that strong convexity constant of every $f_k(x)$ strictly equals to μ. Instead, $\mu = \min_k \mu(f_k)$, where $\mu(f_k)$ denotes the strong convexity parameter of f_k. Analogously, $L = \max_k L(f_k)$.*

III. ANALYSIS OF FIRST-ORDER METHODS ON TIME-VARYING FUNCTIONS

In this section, we start by studying the convergence of the gradient descent and Nesterov's fast gradient method for the general case where the objective function changes with time but remains L-smooth and μ-strongly convex on $\mathbb{R}^n$. This is precisely the case of the dual function (5). Later in Section IV, we will show that the trajectories of both methods are situated in $x_0 + (\mathrm{Ker}\, W)^\perp$, where x_0 is the initial point, and thus even if the functions are μ-strongly convex only on $(\mathrm{Ker}\, W)^\perp$ and not on $\mathbb{R}^n$ (which is the case for the dual of the distributed optimization problem) the studied methods still maintain the same convergence rates. Until now, we have been working with matrix argument $X \in \mathbb{R}^{d \times n}$. For simplicity of exposition and without loss of generality, the following results are derived for the vector argument $x \in \mathbb{R}^n$.

A. Gradient descent

In this subsection, we show that convergence of gradient descent on time-varying functions is the same that on static functions. The proofs are omitted, because non-accelerated gradient descent is not the main focus of this paper. One can carry out the proof using a classical bound

$$f(x_{k+1}) \leqslant f(x_k) - \frac{1}{2L}\|\nabla f(x_k)\|_2^2 \tag{7}$$

and a result given in Theorem 2.5.11 in [29], which states that

$$\langle \nabla f(x) - \nabla f(y), x - y \rangle \geqslant \frac{\mu L}{\mu + L} \|x - y\|_2^2$$
$$+ \frac{1}{\mu + L} \|\nabla f(x) - \nabla f(y)\|_2^2$$

for μ-strongly convex L-smooth functions.

Theorem III.1. *Let $\{f_k(x_k)\}_{k=0}^{\infty}$ be a sequence of functions for which Assumption II.8 hold. Then, the sequence $\{x_k\}_{k=0}^{\infty}$ generated by the gradient descent method, i.e.,*

$$x_{k+1} = x_k - \frac{1}{L} \nabla f_k(x_k), \tag{8}$$

has the following property:

$$\|x_k - x^*\|_2 \leqslant \left(\frac{L - \mu}{L + \mu} \right)^k \|x_0 - x^*\|_2 \quad \text{for all } k \geqslant 0.$$

Next, we provide a Corollary that relates the convergence rate estimate in Theorem III.1 and the minimum number of iterations required to obtain an arbitrarily close approximation of the optimal solution of the optimization problem.

Corollary III.1.1. *Let $\{f_k(x_k)\}_{k=0}^{\infty}$ be a sequence of functions for which Assumption II.8 hold. Then, for any $\varepsilon > 0$, the sequence generated by the iterations in (8) has the following property: for any $k \geq N + 1$ it holds that $\|x_k - x^*\| \leqslant \varepsilon$, where $N \geq \left\lceil (\log((L + \mu)/(L - \mu))^{-1} \log(\|x_0 - x^*\|/\varepsilon) \right\rceil$.*

B. Nesterov fast gradient method

In this subsection, we provide a potential-based proof for the convergence of the Nesterov's fast gradient method [29] for time-varying functions under Assumption II.8, i.e.,

$$y_{k+1} = x_k - \frac{1}{L} \nabla f_k(x_k), \tag{9a}$$

$$x_{k+1} = \left(1 + \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right) y_{k+1} - \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} y_k, \tag{9b}$$

with initial points $y_0 = x_0$ and $\kappa = L/\mu$.

We will follow the potential function proof methods presented in [30]. The general idea of such proof is the use of auxiliary potential function of the following form:

$$\Psi_k = a_k \cdot (f_k(x_k) - f(x^*)) + b_k \cdot \|x_k - x^*\|_2^2,$$

with $a_k, b_k \geqslant 0$. If we denote

$$\Delta \Psi_k = \Psi_{k+1} - \Psi_k. \tag{10}$$

Then $\Psi_N = \Psi_0 + \sum_{k=1}^{N-1} \Delta \Psi_k$ and

$$f_N(y_N) - f(x^*) \leqslant \left(\Psi_0 + \sum_{k=1}^{N-1} \Delta \Psi_k \right) / a_N. \tag{11}$$

If an upper bound on $\Delta \Psi_k$ is obtained, then (11) shows the convergence rate for the method.

Nesterov's method is not a strict descent method. This becomes an obstacle in the time-varying case, because sudden changes of the function may happen too often so that the Nesterov method is run for too short periods of time and thus does not manage to make enough progress. However, this method's convergence can be proved if the number of function changes is finite. Next, we formally introduce the definition of a change in a sequence of functions.

Definition III.2. *Consider a sequence $\{f_k(x)\}_{k=0}^{\infty}$ of functions and let $f_n \not\equiv f_{n+1}$. Then we say that the sequence $\{f_k(x)\}_{k=0}^{\infty}$ of functions has a change at the moment n.*

We are now ready to state the main auxiliary result of this paper, that relates the convergence rate of an accelerated method on time-varying functions under Assumption II.8.

Theorem III.3. *Let $N > 0$ be a time horizon, and let Assumption II.8 hold for a sequence of functions $\{f_k(x)\}_{k=0}^{N}$ with $0 \leq m \leq N$ changes. Then, the sequence generated by (9) has the following property:*

$$f_N(y_N) - f^* \leqslant \frac{L + \mu}{2} R^2 \kappa^m \left(1 - \frac{1}{\sqrt{\kappa}} \right)^N,$$

where $\kappa = L/\mu$ and $\|x_0 - x^*\|_2 \leq R$.

Before proceeding to the proof of Theorem III.3, we provide a sequence of technical lemmas that will facilitate the analysis.

Following the technique for strongly convex functions described in [30], we introduce the following potential:

$$\Psi_k = (1 + \gamma)^k \cdot \left(f_k(y_k) - f^* + \frac{\mu}{2} \|z_k - x^*\|_2^2 \right), \tag{12}$$

where $\gamma = 1/\sqrt{\kappa} - 1$ and z_k will be defined shortly.

The next lemma provides an intermediate result regarding an auxiliary sequence $\{z_k\}$ that will come handy later in the proofs.

Lemma III.4. *Consider updates in (9) and define*

$$\tau = \frac{1}{\sqrt{\kappa} + 1}, \text{ and } z_{k+1} = \frac{1}{\tau} x_{k+1} - \frac{1 - \tau}{\tau} y_{k+1}.$$

Then, $z_{k+1} = \frac{1}{1+\gamma} z_k + \frac{\gamma}{1+\gamma} x_k - \frac{\gamma}{\mu(1+\gamma)} \nabla f_k(x_k)$, where $\gamma = \frac{1}{\sqrt{\kappa}-1}$.

Proof. By the update rule for x_{k+1} given in (9) and the definition of τ, we have that

$$x_{k+1} = \left(1 + \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right) y_{k+1} - \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} y_k$$
$$= (2 - 2\tau) y_{k+1} - (1 - 2\tau) y_k.$$

Moreover, by the definition of z_{k+1}, it follows that

$$z_{k+1} = \frac{1}{\tau} x_{k+1} - \frac{1 - \tau}{\tau} y_{k+1}$$
$$= \frac{1}{\tau} ((2 - 2\tau) y_{k+1} - (1 - 2\tau) y_k) - \frac{1 - \tau}{\tau} y_{k+1}$$
$$= \frac{1}{\tau} ((1 - \tau) y_{k+1} - (1 - 2\tau) y_k).$$

Now we use the update rule for y_{k+1} given in (9) and also note that $x_k = (1 - \tau) y_k + \tau z_k$:

$$z_{k+1} = \frac{1}{\tau} \left[(1 - \tau)(x_k - \frac{1}{L} \nabla f_k(x_k)) - \frac{1 - 2\tau}{1 - \tau}(x_k - \tau z_k) \right]$$
$$= \frac{1 - 2\tau}{1 - \tau} z_k + \frac{\tau}{1 - \tau} x_k - \frac{1 - \tau}{L\tau} \nabla f_k(x_k)$$
$$\overset{①}{=} \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa}} z_k + \frac{1}{\sqrt{\kappa}} x_k - \frac{1}{\mu \sqrt{\kappa}}$$
$$\overset{②}{=} \frac{1}{1 + \gamma} z_k + \frac{\gamma}{1 + \gamma} x_k - \frac{\gamma}{\mu(1 + \gamma)} \nabla f_k(x_k),$$

where ① is obtained by using the definitions of τ and κ, and ② is obtained by using the definition of γ. $\qquad \square$

The next lemma will help us towards quantification of the maximum function value change in the sequence of time-varying functions.

Lemma III.5. *Define $\delta_k(x) = f_{k+1}(x) - f_k(x)$, and let both $f_k(x)$ and $f_{k+1}(x)$ be μ-strongly convex and L-smooth, and have the same minimizer x^*. Then*

$$\delta_k(x) \leqslant \frac{L-\mu}{\mu}(f_k(x) - f^*),$$

where f^* is the common value of minimum for $\{f_k\}_{k=1}^{\infty}$.

Proof. By strong convexity and smoothness obtain

$$\frac{\mu}{2}\|x - x^*\|_2^2 \leqslant f_k(x) - f^* \leqslant \frac{L}{2}\|x - x^*\|_2^2.$$

The same holds for f_{k+1}.

$$f_{k+1}(x) - f_k(x) \leqslant \frac{L-\mu}{2}\|x - x^*\|_2^2 \leqslant \frac{L-\mu}{\mu}(f_k(x) - f^*).$$
$\square$

Finally, the next lemma relates the upper bounds on the function values of the sequence of functions with the changes of a specific potential function.

Lemma III.6. *Let $\{f_k(x)\}_{k=0}^{\infty}$ be a sequence of functions for which Assumption II.8 hold, and let Ψ_k be the potential function given in* (12). *Then, it holds that*

$$\Delta\Psi_k \leqslant (1+\gamma)^{k+1}\delta_k(y_{k+1}). \tag{13}$$

Proof. The proof is analogous to the proof in Section 5.4 in [30]. We use the definitions of τ, z_k given in Lemma III.4. We have

$$\Delta\Psi_k \cdot (1+\gamma)^{-k} = (1+\gamma)\Big(f(y_{k+1}) - f^* + \frac{\mu}{2}\|z_{k+1} - x^*\|_2^2\Big)$$
$$- \Big(f(y_k) - f^* + \frac{\mu}{2}\|z_k - x^*\|_2^2\Big)$$
$$= (1+\gamma)\big(f_{k+1}(y_{k+1}) - f_{k+1}(x^*)\big) - \big(f_k(y_k) - f_k(x^*)\big)$$
$$+ \frac{\mu}{2}\Big[(1+\gamma)\|z_{k+1} - x^*\|_2^2 - \|z_k - x^*\|_2^2\Big]. \tag{14}$$

Note that from basic gradient step inequality given in (7), it follows that

$$f_k(y_{k+1}) \leqslant f_k(x_k) - \frac{1}{2L}\|\nabla f_k(x_k)\|_2^2,$$

and by using the definition of δ_k in Lemma III.5, we have

$$f_{k+1}(y_{k+1}) \leqslant f_k(x_k) - \frac{1}{2L}\|\nabla f_k(x_k)\|_2^2 + \delta_k(y_{k+1}).$$

Therefore the first term in (14) can be bounded as follows:

$$(1+\gamma)\big(f_{k+1}(y_{k+1}) - f_{k+1}(x^*)\big) - \big(f_k(y_k) - f_k(x^*)\big)$$
$$\leqslant (1+\gamma)\Big(f_k(x_k) - \frac{1}{2L}\|\nabla f_k(x_k)\|_2^2 + \delta_k(y_{k+1}) - f^*\Big)$$
$$- \big(f_k(y_k) - f^*\big)$$
$$= f_k(x_k) - f_k(y_k) + \gamma(f_k(x_k) - f^*)$$
$$- (1+\gamma)\frac{\|\nabla f_k(x_k)\|_2^2}{2L} + (1+\gamma)\delta_k(y_{k+1})$$
$$\leqslant \langle\nabla f_k(x_k), x_k - y_k\rangle$$
$$+ \gamma\big(\langle\nabla f_k(x_k), x_k - x^*\rangle - \frac{\mu}{2}\|x_k - x^*\|_2^2\big)$$
$$- \frac{1+\gamma}{2L}\|\nabla f_k(x_k)\|_2^2 + (1+\gamma)\delta_k(y_{k+1}). \tag{15}$$

It is convenient to rewrite the above expression without references to y_k, by using Lemma III.4. Thus, by Lemma III.4, by using the definitions of z_k, γ, τ and κ, we deduce

$$z_k = \Big(\frac{1}{\tau} - 1\Big)(x_k - y_k) + x_k = \sqrt{\kappa}(x_k - y_k) + x_k$$
$$\gamma(z_k - x^*) = \sqrt{\kappa}\gamma(x_k - y_k) + \gamma(x_k - x^*).$$

Keeping in mind that $\sqrt{\kappa}\gamma = 1 + \gamma$, we obtain

$$(x_k - y_k) + \gamma(x_k - x^*) =$$
$$\frac{1}{1+\gamma} \cdot \Big[\gamma(z_k - x^*) + \gamma^2(x_k - x^*)\Big].$$

The expression on the right hand side of (15) can be written as follows:

$$\frac{1}{1+\gamma}\langle\nabla f_k(x_k), \gamma(z_k - x^*) + \gamma^2(x_k - x^*)\rangle -$$
$$\frac{\mu\gamma}{2}\|x_k - x^*\|_2^2 - \frac{1+\gamma}{2L}\|\nabla f_k(x_k)\|_2^2 + (1+\gamma)\delta_k(y_{k+1}). \tag{16}$$

The obtained bound (15) is almost the same as (5.69) in [30]. The only difference is the additional term $(1+\gamma)\delta_k(y_{k+1})$.

The second term in (14) is bounded in the same way as in [30]. By Lemma III.4:

$$\frac{\mu}{2}\Big[(1+\gamma)\|z_{k+1} - x^*\|_2^2 - \|z_k - x^*\|_2^2\Big]$$
$$= \frac{\mu}{2}(1+\gamma)\Big\|\frac{1}{1+\gamma}(z_k - x^*) + \frac{\gamma}{1+\gamma}(x_k - x^*)$$
$$- \frac{\gamma}{\mu(1+\gamma)}\nabla f_k(x_k)\Big\|_2^2 - \frac{\mu}{2}\|z_k - x^*\|_2^2$$
$$= \frac{\mu}{2}\frac{1}{1+\gamma}\Big[\|z_k - x^*\|_2^2 + \gamma^2\|x_k - x^*\|_2^2 + \frac{\gamma^2}{\mu^2}\|\nabla f_k(x_k)\|_2^2$$
$$+ 2\gamma\langle z_k - x^*, x_k - x^*\rangle - \frac{2\gamma}{\mu}\langle z_k - x^*, \nabla f_k(x_k)\rangle$$
$$- \frac{2\gamma^2}{\mu}\langle x_k - x^*, \nabla f_k(x_k)\rangle\Big] - \frac{\mu}{2}\|z_k - x^*\|_2^2. \tag{17}$$

Now by adding (16) and (17), we obtain a final bound on $\Delta\Psi_k$. Moreover, note that terms involving $\langle\nabla f_k(x_k), x_k - x^*\rangle$ and $\langle\nabla f_k(x_k), z_k - x^*\rangle$ cancel out.

$$\Delta\Psi_k(1+\gamma)^{-k} \leqslant \Big(-\frac{1+\gamma}{2L} + \frac{\gamma^2}{2\mu(1+\gamma)}\Big)\|\nabla f_k(x_k)\|_2^2$$
$$+ \frac{\mu\gamma}{2}\Big(\frac{\gamma}{1+\gamma} - 1\Big)\|x_k - x^*\|_2^2 + \frac{\mu}{2}\Big(\frac{1}{1+\gamma} - 1\Big)\|z_k - x^*\|_2^2$$
$$+ \frac{\mu\gamma}{1+\gamma}\langle z_k - x^*, x_k - x^*\rangle + (1+\gamma)\delta_k(y_{k+1})$$
$$\leqslant -\frac{\mu\gamma}{2(1+\gamma)}\big(\|x_k - x^*\|_2^2 + \|z_k - x^*\|_2^2$$
$$- 2\langle z_k - x^*, x_k - x^*\rangle\big) + (1+\gamma)\delta_k(y_{k+1})$$
$$= -\frac{\mu\gamma}{2(1+\gamma)}\|(x_k - x^*) - (z_k - x^*)\|_2^2 + (1+\gamma)\delta_k(y_{k+1})$$
$$\leqslant (1+\gamma)\delta_k(y_{k+1}),$$

and the proof is complete. $\square$

Now all the auxiliary lemmas are proved, and we move to the proof of Theorem III.3.

Proof of Theorem III.3. Lemmas III.5 and III.6 establish the connection between a potential change and the function residual, which enables to perform the proof by induction on the number of changes m.

Induction base for $m = 0$ holds due to smoothness of $\{f_k(x)\}_{k=0}^{\infty}$ and to the fact $x_0 = y_0 = z_0$:

$$\Psi_0 = f_0(y_0) - f^* + \frac{\mu}{2}\|z_0 - x^*\|_2^2$$

$$\leqslant \frac{L}{2}\|y_0 - x^*\|_2^2 + \frac{\mu}{2}\|z_0 - x^*\|_2^2 = \frac{L+\mu}{2}R^2$$

$$f_N(y_N) - f^* \leqslant \frac{\Psi_0}{a^N} \leqslant \frac{L+\mu}{2}\frac{R^2}{(1+\gamma)^N}.$$

Let the induction hypothesis hold for $0, 1, ..., m$. By Lemma III.5 and using the fact that (13) implies $\Delta\Psi_k \leq 0$ unless $k = n_i$ for some i we get

$$f_N(y_N) - f^* \leqslant \left(\Psi_0 + \sum_{k=1}^{m}\Delta\Psi_{n_k}\right)/a_N$$

$$\leqslant \left(\Psi_0 + \sum_{k=1}^{m}(1+\gamma)^{n_k+1}\delta_{n_k}(y_{n_k+1})\right)/(1+\gamma)^N$$

$$\leqslant \Psi_0 + \sum_{k=1}^{m}(1+\gamma)^{n_k+1} \cdot \frac{L-\mu}{\mu}(f_{n_k}(y_{n_k+1}) - f^*)/(1+\gamma)^N.$$

Since the function changes take place at the moments $n_1, ..., n_m$, the bound is true for $f_{n_1}, ..., f_{n_m}$.

$$(1+\gamma)^N(f_N(y_k) - f^*)$$

$$\leqslant \Psi_0 + \sum_{k=1}^{m}(1+\gamma)^{n_k+1} \cdot \frac{L-\mu}{\mu}\left(\frac{L}{\mu}\right)^{k-1}\frac{L+\mu}{2}\frac{R^2}{(1+\gamma)^{n_k+1}}$$

$$\leqslant \frac{L+\mu}{2}R^2\left(1 + \sum_{k=1}^{m}\frac{L-\mu}{\mu}\left(\frac{L}{\mu}\right)^{k-1}\right)$$

$$\leqslant \frac{L+\mu}{2}R^2\left(1 + \frac{L-\mu}{\mu} \cdot \frac{\left(\frac{L}{\mu}\right)^m - 1}{\frac{L}{\mu} - 1}\right)$$

$$= \left(\frac{L}{\mu}\right)^m\frac{L+\mu}{2}R^2. \tag{18}$$

Dividing (18) by $(1+\gamma)^N$ finishes the proof. $\qquad\square$

In the next section, we will use the result in Theorem III.3 for the convergence rate analysis of accelerated methods in distributed optimization over time-varying graphs. Note that the result in Theorem III.3 we set a fixed time horizon N and a fixed number of changes m. Our main result is going to be stated for the general case where the number of changes is a fixed percentage of the number of iterations N.

IV. AN ACCELERATED METHOD FOR DISTRIBUTED OPTIMIZATION OVER TIME-VARYING FUNCTIONS

In this section, we present the main result regarding the convergence rate of the distributed Nesterov fast gradient method for time-varying functions. It states that this method is linearly convergent on a slowly time-varying network. More specifically, we show that optimal rates are guaranteed if the number of changes in the network, expressed as a fraction of the number of iterations, is bounded. We provide this explicit bound and its depedency witht the worst case condition number in the sequene of graphs.

Theorems III.1 and III.3 hold for time-varying functions which are L-smooth and μ-strongly convex on $\mathbb{R}^n$. However, our initial aim was to find a common minimum of the sequence of functions defined in (5). In (5), every function $f_k(x)$ is μ-strongly convex only on the subspace $\left(\text{Ker } W_k\right)^{\perp}$ and L-smooth on $\mathbb{R}^n$. Therefore, we need to show that the Theorems III.1 and III.3 can be generalized on strong convexity on a subspace. To do so, we show that the iterates generated by the studied algorithms are always in the space where the functions are strongly convex.

In the next lemma, we show that the gradients of the dual function are always in $\left(\text{Ker } W\right)^{\perp}$.

Lemma IV.1. *Consider the function* $f(x) = \max_{y \in \mathbb{R}^n}(-\langle x, \sqrt{W}y\rangle - \varphi(y))$. *Then, it holds that* $\nabla f(x) \in \left(\text{Ker } W\right)^{\perp}$.

Proof. Initially, denote the optimal point of the inner maximization problem of the dual function as

$$y(x) = \arg\max_{y \in \mathbb{R}^n}\left(-\langle x, \sqrt{W}y\rangle - \varphi(y)\right).$$

Thus, by the Demianov-Danskin formula [31]–[33] it follows that $\nabla f(x) = -\sqrt{W}y(x)$. Therefore, it is sufficient to show $\langle\nabla f(x), z\rangle = 0 \; \forall z \in \text{Ker } W$, which follows from

$$\langle\nabla f(x), z\rangle = \langle-\sqrt{W}y(x), z\rangle = \langle-\sqrt{W}z, y(x)\rangle = 0.$$

$\qquad\square$

In the next lemma, we show that the iterates generated by the gradient descent method and the fast gradient method are always in the space where the strong convexity of the dual function holds. This will allow us to use the results in Section III for the specific problem of distributed optimization over time-varying graphs.

Lemma IV.2. *The algorithm in* (9)*, with initial point* $y_0 = x_0$*, generates sequences that are always in* $x_0 + \left(\text{Ker } W\right)^{\perp}$*.*

Proof. The proof follows by induction. Let $x_k - x_0 \in \left(\text{Ker } W\right)^{\perp}$ and $y_k - y_0 = y_k - x_0 \in \left(\text{Ker } W\right)^{\perp}$ (note that it holds for $k = 0$). Then, from (9) it holds that

$$y_{k+1} - x_0 = (x_k - x_0) - \frac{1}{\mu}\nabla f_k(x_k) \in (\text{Ker } W)^{\perp}$$

$$x_{k+1} - x_0 = \left(1 + \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}\right)(y_{k+1} - y_0) - \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}(y_k - y_0)$$

$$x_{k+1} - x_0 \in \left(\text{Ker } W\right)^{\perp}.$$

$\qquad\square$

A. Algorithm and main result

Now that we have shown that the general analysis in Section III hold for the iterates generated by the studied methods in (8) and (9), we proceed to explicitly write the proposed accelerated distributed optimization algorithm for each of the agents in the network. Moreover, we analyze its convergence rate.

Remark IV.3.

- *We are working with the Lagrange dual and running Nesterov method on it. Line 3 of Algorithm 1 is a result of introducing $Z = X\sqrt{W}$.*
- *The algorithm uses $\arg\max$ computation. We assume that functions $\varphi_i(\cdot)$ are dual-friendly and this operation is cheap. Relaxations of this assumption follows from [4].*

Algorithm 1 Distributed Nesterov Method
___
Require: Each agent $i \in V$ locally holds φ_i and some iteration number N.
1: Choose $\tilde{z}_0^i = z_0^i$ for all $i \in V$
2: **for** $k = 0, 1, 2, \cdots, N-1$ **do**
3:     $\tilde{y}_i(z_i^k) = \arg\max\limits_{y \in \mathbb{R}^d}\left[\langle z_i^k, y \rangle - \varphi_i(y_i)\right]$
4:     Send $\tilde{y}_i(z_i^k)$ to every neighbor and receive $\tilde{y}_j(z_j^k)$ from every neighbor.
5:     $\tilde{z}_i^{k+1} = z_i^k - \frac{1}{L}\sum_{j=1}^{n}[W_k]_{ij}\tilde{y}_j(z_j^k)$
6:     $z_i^{k+1} = \left(1 + \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}\right)\tilde{z}_i^{k+1} - \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}\tilde{z}_i^k$
7: **end for**
___

We are now ready to state the main result of this paper, that provides the convergence rate of the distributed Nesterov fast gradient method over slowly time-varying networks.

Theorem IV.4. *Let Φ be a μ_Φ-strongly convex L_Φ-smooth function. Also, for $N \geq 0$ let $\{\mathcal{G}_k\}_{k=1}^N$ be a sequence of undirected connected graphs with at most αN changes, where $\alpha \in (0, 1/(\sqrt{\kappa}\log(\kappa)))$. For any $\varepsilon > 0$, the output $\tilde{Z}_N$ of Algorithm 1 has the following property:*

$$\tilde{f}_N(\tilde{Z}_N) - f^* \leqslant \varepsilon$$

for $N \geq (\sqrt{\kappa} + \alpha\log(\kappa))\log\left((L+\mu)R^2/(2\varepsilon)\right)$, where $\tilde{f}(Z) = \Phi^*(-Z) = \max\{-\langle Z, Y\rangle - \Phi(Y)\}$, $\tilde{Z}_N = [\tilde{z}_1^N, \cdots, \tilde{z}_n^N]$, $R = \|X_0 - X^*\|_2$, $\kappa = L_\Phi/\mu_\Phi \cdot \sqrt{\theta_{\min}/\theta_{\max}}$, where $\theta_{\max}, \theta_{\min}$ are defined in (6).

Proof. Given that Φ is μ_Φ-strongly convex L_Φ-smooth it follows from Lemma II.5, that f (defined in (5)) is μ-strongly convex L-smooth, with $L = \sqrt{\theta_{\max}}/\mu_\Phi$, and $\mu = \sqrt{\theta_{\min}}/L_\Phi$. Now, from Theorem III.3 we have that

$$\tilde{f}_N(\tilde{Z}_N) - f^* \leq \frac{L+\mu}{2}R^2\exp(-N(1/\sqrt{\kappa} - \alpha\log\kappa)).$$

Initially, note that in order to guarantee convergence we require $\alpha \in (0, 1/(\sqrt{\kappa}\log(\kappa)))$. Moreover, We need to find a bound on N such that

$$\frac{L+\mu}{2}R^2\left(\kappa^\alpha\left(1 - \frac{1}{\sqrt{\kappa}}\right)\right)^N \leq \varepsilon.$$

Thus, $N \geq (\sqrt{\kappa} + \alpha\log(\kappa))\log\left((L+\mu)R^2/(2\varepsilon)\right)$ □

In the next corollary, we provide convergence results for the primal problem.

Corollary IV.4.1. *Let $\tilde{f}(Z) - f^* \leqslant \varepsilon$. Then*

$$\Phi(\tilde{Y}(Z)) - \Phi^* \leqslant 2\kappa\varepsilon + L\|X^*\|\sqrt{\frac{2\varepsilon}{\mu}},$$

where $\tilde{Y}(Z) = \arg\max\limits_{Y\in\mathbb{R}^{d\times n}}\{-\langle Z, Y\rangle - \Phi(Y)\} = [\tilde{y}_1(z_i^N), \ldots, \tilde{y}_n(z_n^N)]$.

Proof. By definition of $\tilde{f}$, $\tilde{f}(Z) = f(X\sqrt{W})$ for some X. Strong convexity on $(\text{Ker }W)^\perp$ and smoothness of f yield

$$\frac{\mu}{2}\|X - X^*\|^2 \leqslant \varepsilon,$$

$$\|\nabla f(X) - \nabla f(X^*)\| \leqslant L_f\|X - X^*\| \leqslant L_f\sqrt{\frac{2\varepsilon}{\mu}}.$$

Now, note that $\nabla f(X) = -Y(X)W$, and trivially $f(X) \geqslant f(X^*)$, thus

$$-\left\langle X, Y(X)\sqrt{W}\right\rangle - \Phi(Y(X)) \geqslant$$
$$-\left\langle X^*, Y(X^*)\sqrt{W}\right\rangle - \Phi(Y(X^*)),$$
$$\Phi(Y(X)) - \Phi^* \leqslant \langle X, \nabla f(X)\rangle$$
$$\leqslant (\|X - X^*\|_2 + \|X^*\|) \cdot L\sqrt{\frac{2\varepsilon}{\mu}}$$
$$\leqslant \left(\sqrt{\frac{2\varepsilon}{\mu}} + \|X^*\|\right) \cdot L\sqrt{\frac{2\varepsilon}{\mu}}.$$

This completes the proof. □

Remark IV.5. *The convergence result of Nesterov method in Theorem IV.4 depends on $\kappa_\Phi = L_\Phi/\min\limits_i(\mu_{\varphi_i})$. This term can be reduced by changing the functions φ_i. Let strong convexity parameter of φ_i equal to μ_i and denote $\overline{\mu} = (1/n)\sum_{i=1}^n \mu_i$. Introduce $\widehat{\varphi}_i(y) = \varphi_i(y_i) + (\overline{\mu} - \mu_i)\|y_i\|^2/2$.*

$$\Phi(Y) = \sum_{i=1}^n \varphi_i(y_i) = \sum_{i=1}^n\left(\varphi_i(y_i) + (\overline{\mu} - \mu_i)\frac{\|y_i\|^2}{2}\right)$$
$$= \sum_{i=1}^n \widehat{\varphi}_i(y_i)$$

Now, if we work with $\widehat{\varphi}_i$ instead of φ_i, it will result in a better condition number $\widehat{\kappa}_\Phi = L_\Phi/\overline{\mu}_\Phi < L_\Phi/\min\limits_i\{\mu_{\varphi_i}\} = \kappa_\Phi$.

Note that the number of steps in Theorem IV.4 reaches the lower bound for decentralized methods in [1], which means that the Algorithm 1 is optimal for time-varying graphs with a finite number of changes. Moreover, since $\kappa = \chi(W) \cdot \kappa_\Phi$, it follows that the factor κ is proportional to $\chi(W)$, which is the communication graph condition number. The lower the graph condition number is, the better convergence rate we obtain. Note that if we used restriction $YW = 0$ instead of $Y\sqrt{W} = 0$, then it would be $\kappa \sim \chi(W)^2$, which would result in slower convergence.

When the number of changes is not finite, but rather a percentage α of the total number of iterations, then, α needs to be upper bounded by $1/(\sqrt{\kappa}\log\kappa)$. This shows that optimal rates can only be achieved if the graph changes slowly. This provides only a sufficient condition, and it remains an open question whether this bound on α is also necessary.

V. Discussion and comparison to other methods

In this section, we compare the performance of the accelerated gradient method to several distributed algorithms presented in other works. Particularly, we consider PANDA [22], DIGing [7] and Nesterov method for static networks Acc-DNGD [9]. These algorithms are designed to solve problem (1) and are based on a mixing matrix sequence $\{V(k)\}_{k=1}^\infty$, which has the following properties:

Assumption V.1. *1) (Decentralized property) If $i \neq j$ and edge $(i, j) \notin E_k$, then $V(k)_{ij} = 0$;*
2) (Double stochasticity) $V(k)1_n = 1_n$, $1_n^T V(k) = 1_n^T$;
3) (Joint spectrum property) There exists $B \in \mathbb{Z}$, $B > 0$, such that

$$\delta = \sup_{k \geqslant B-1} \sigma_{\max} \left\{ V_B(k) - \frac{1}{n} 1_n 1_n^T \right\} < 1. \qquad (19)$$

Here $1_n = [1 \ 1... \ 1]^T \in \mathbb{R}^n$ and $V_B(k) = V(k)V(k - 1)...V(k - B + 1)$.

Following the arguments in [7], one can establish that matrices $(I_n - n^{-1}W(k))$ meets all the requirements in Assumption V.1 with $B = 1$.

A. Relation to DIGing

Let us give a lower bound on the theoretical convergence rate of the DIGing algorithm, which is linearly convergent and originally presented in [7], and compare it with the rate of accelerated gradient method obtained in Theorem III.3.

Assumption V.2. *Every φ_i in problem (1) is μ_i-strongly convex and L_i-smooth w.r.t. $\| \cdot \|_2$.*

Proposition V.3. *Under Assumptions V.1 and V.2 the theoretical result for DIGing algorithm given in [7] does not guarantee a convergence rate faster than $O(\lambda_0^N)$, where λ_0 is defined as*

$$\lambda_0 = 1 - 1/(12\overline{\kappa}^{3/2}\sqrt{n}).$$

Here $\overline{\kappa} = 1/n \sum_{i=1}^{n} L_i/\mu_i$ and n is the number of vertices in the network graph.

The proof of Proposition V.3 is presented in Appendix B.
The convergence rate of Nesterov gradient method obtained in Theorem III.3 is $O(\lambda_1^N)$ where

$$\lambda_1 = 1 - 1/\kappa^{1/2} (\theta_{\max}/\theta_{\min})^{1/4}. \qquad (20)$$

Note that κ is the condition number of f in (1), while $\overline{\kappa}$ is an average condition number of summands φ_i.
Accelerated gradient method has several advantages as well as disadvantages in comparison with the DIGing algorithm.

- Typically, the objective function condition number κ is rather large, and the graph condition number $(\theta_{\max}/\theta_{\min})^{1/2}$ corresponds to the diameter of network graph [34] and therefore is not larger than n. Moreover, if we are working with a machine learning problem and the dataset is uniformly distributed between the computers in the network, then the summands φ_i in (1) have approximately the same condition number, i.e. $\overline{\kappa} \approx \kappa$. In this case, Nesterov accelerated method outperforms DIGing, since $\kappa^{1/2} \ll \overline{\kappa}^{3/2}$ and $(\theta_{\max}/\theta_{\min})^{1/4} \leqslant \sqrt{n}$.
- The case where the graph remains connected at every time step corresponds to $B = 1$ in DIGing. The DIGing algorithm is capable of working with an arbitrary number of changes and with graphs which do not stay connected all the time.
- Nesterov accelerated method's number of iterations grows linearly with the number of changes in the network, while the number of iterations of the DIGing algorithm does not depend on the number of changes.

Finally, we emphasise that to get a full comparison between the proposed algorithm and DIGing [7] one needs to revisit the analysis done in [7] for this particular sequence of graphs, which is beyond the scope of their work.

B. Relation to PANDA

PANDA is a linearly-convergent dual-based algorithm presented in [22].

Assumption V.4. *Let φ in (1) be L-smooth and μ-strongly convex w.r.t. $\| \cdot \|_2$.*

Proposition V.5. *Let Assumptions V.1 and V.4 hold. Then the theoretical result for PANDA in [22] does not guarantee a convergence rate better then $O(\lambda_0^N)$ where λ_0 is given by*

$$\lambda_0 = 1 - \frac{9}{64} \frac{1}{\kappa^{3/2}},$$

if PANDA step-size $c \in (0, \alpha]$, where α is defined as

$$\alpha = 2\sqrt{\kappa}\mu \left(\frac{\sqrt{(1 - \delta^2)\kappa^{-2/3} + 8} - 8\delta}{\kappa^{-3/2} + 8} \right)^2.$$

The proof of Proposition V.5 is provided in Appendix C.
One can make sure that the PANDA algorithm can work with step size $c > \alpha$. Although it is interesting to compare Nesterov accelerated method and PANDA with a bigger step size, the analysis, in this case, seems to be complicated and therefore is left for future work.
Analogously to Section V-A, let us discuss advantages and disadvantages of the results of this paper in comparison with PANDA (i.e. compare λ_0 in (V.5) to λ_1 in (20)).

- If the objective function is badly conditioned, i.e. $\kappa \gg 1$, and the communication graph is well-conditioned, then Nesterov method outperforms PANDA. On the other hand, if $\kappa \ll (\theta_{\max}/\theta_{\min})^{1/4}$, PANDA converges faster.
- Analogously to DIGing, PANDA works under weaker assumptions and does not depend on the number of changes in the network.

C. Relation to Nesterov method on static network

In this section we provide a theoretical comparison between our method and Nesterov method on a static network presented in [9].

Proposition V.6. *The theoretical bound for Algorithm 1 is better than the bound for distributed Nesterov method in [9] if and only if*

$$(\lambda_2(1 - \lambda_2))^{3/2}/250 \cdot \sqrt{\lambda_{\max}/\lambda_{\min}} < (L_\Phi/\mu_\Phi)^{3/14},$$

where λ_2 is the second largest eigenvalue of W.

The proof of Proposition V.6 is provided in Appendix D.

VI. NUMERICAL EXPERIMENTS

In this section, we present simulation results for the Algorithm 1 for the *rigde regression* (strongly convex and smooth) problem. Moreover, we compare its performance with the centralized fast gradient method [29], DIGing [7], Acc-DNGD [9], and PANDA [22].
The synthetic *rigde regression* problem is defined as

$$\min_{z \in \mathbb{R}^m} \frac{1}{2nl} \|b - Hz\|_2^2 + \frac{1}{2}c\|z\|_2^2. \qquad (21)$$
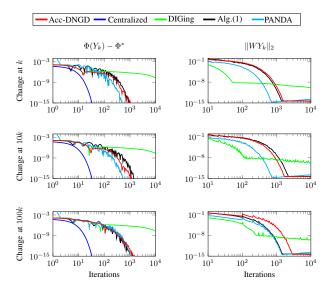
Fig. 2. Distance to optimality and distance to consensus for a sequence of Erdős-Rényi random graphs with 100 agents. Top row shows the results for graphs that change at every step, middle row shows the results for changes every 10 steps, and bottom row shows the results for changes every 100 steps.
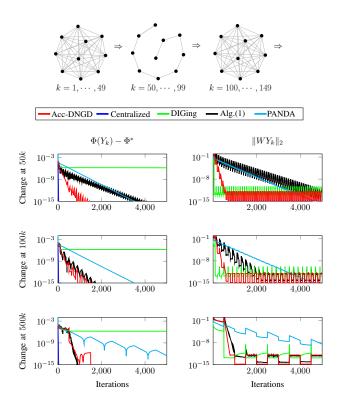

Fig. 3. Distance to optimality and distance to consensus for a network of 100 Agents on a sequence of graphs that shuffles between a complete graph and a path graph every 50 iterations, 100 iterations, and 500 iterations.

Moreover, we seek to solve (21) distributedly over a network. Each entry of the data matrix $H \in \mathbb{R}^{nl \times m}$ is generated as an independent identically distributed random variable $H_{ij} \sim \mathcal{N}(0,1)$, the vector of associated values $b \in \mathbb{R}^{nl}$ is generated as a vector of random variables where $b = Hx^* + \epsilon$ for some predefined $x^* \in \mathbb{R}^m$ and $\epsilon \sim \mathcal{N}(0,0.1)$. The columns of the data matrix H and the output vector b are evenly distributed among the agents with a total of l data points per agent. The regularization constant is set to $c = 0.1$. Thus, each agent has access to a subset of points such that

$$b^T = [\underbrace{b_1^T}_{\text{Agent }1} | \cdots | \underbrace{b_n^T}_{\text{Agent }n}] \quad \text{and} \quad H^T = [\underbrace{H_1^T}_{\text{Agent }1} | \cdots | \underbrace{H_n^T}_{\text{Agent }n}],$$

where $b_i \in \mathbb{R}^l$ and $H_i \in \mathbb{R}^{l \times m}$ for each agent $i \in V$. Therefore, in this setup each agent $i \in V$ has a private local function

$$\varphi_i(x_i) \triangleq \frac{1}{2nl}\|b_i - H_i x_i\|_2^2 + \frac{1}{2}\frac{c}{n}\|x_i\|_2^2.$$

Moreover, the optimization problem 21 is equivalent to

$$\min_{\sqrt{W}x=0} \sum_{i=1}^n \left(\frac{1}{2}\frac{1}{nl}\|b_i - H_i x_i\|_2^2 + \frac{1}{2}\frac{c}{n}\|x_i\|_2^2 \right),$$

where $W = \bar{W} \otimes I_m$.

Figure 2 shows the numerical results when the network is a sequence Erdős-Rényi random graphs with 100 agents and the graph changes at: every step, every 10 steps, and every 100 steps. Given that Erdős-Rényi random graphs condition number scales logarithmically with the number of agents, the changes do not affect the rate of convergence. In all three cases, the performance of Acc-DNGD is comparable with the proposed method. In the next examples we will see how abrupt changes can lead to the instability of the algorithm.

Figure 3 shows the numerical results for a sequence of graphs that changes between a complete graph and a path graph every 50, 100 or 500 iterations. Even if the graph changes every 50 iterations, the convergence is maintained, due to the connectivity of the complete graph. The DIGing

algorithm reaches consensus faster, but not on the optimal point, which is slower than other methods. Every time there is a change in the topology there is an increase in the distance to consensus due the changes in the neighbor sets. When the graph changes every 500 steps, one can see that after the initial steps as a line graph, the algorithm converges fast once the graph switches to the complete graph. For fast changing graphs PANDA has comparable performance as Alg. 1, however, performance improves as the changes happen less frequently. For this rapid abrupt change in the network topology, Acc-DNGD convergence faster for rapid changes. However, as the changes frequency decreases, the performance of Algorithm 1 is comparable with Acc-DNGD.

Figure 4 shows the numerical results for a sequence of graphs that changes between a cycle and a star graph every 50, 100 and 500 iterations. If the graph changes quickly, every 100 for this case, Alg. 1 diverges, i.e., the proposed accelerated method is not able to keep up with the network changes. This is evident in this case since we are switching between two graphs with relatively large condition number. It is only when the graph changes happen rarely enough, i.e., every 500 steps, than the proposed method converges. In this third example, we observe that Acc-DNGD outperforms Algorithm 1, at it maintains convergence even if the network changes rapidly. However, as the changes become infrequent Algorithm 1 shows faster convergence.

Appendix E provides additional simulations of the proposed algorithm for the problem of regularized logistic regression for
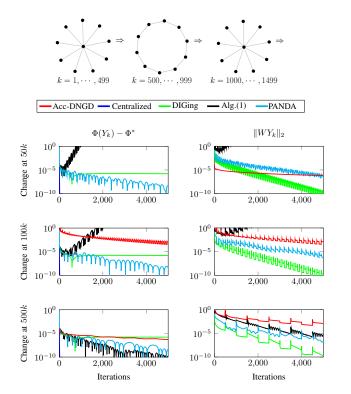
Fig. 4. Distance to optimality and distance to consensus for a network of 100 Agents on a sequence of graphs that shuffles between a star graph and a cycle graph every 50 iterations, 100 iterations, and 500 iterations.

training linear classifiers. , i.e.,

$$\min_{x \in \mathbb{R}^m} \frac{1}{2nl} \sum_{i=1}^{nl} \log\left(1 + \exp\left(-y_i \cdot A_i^T x\right)\right) + \frac{1}{2}c\|x\|_2^2, \quad (22)$$

where $A_i \in \mathbb{R}^d$ is a data point with $y_i \in \{-1, 1\}$ as its corresponding class assignment. We applied the proposed method to datasets from the library LIBSVM [35]. We seek to distributedly solve the logistic regression problem over the following datasets: A9A, MUSHROOMS, IJCNN1 and COVTYPE.

VII. CONCLUSIONS AND FUTURE WORK

We study the convergence of gradient descent, and Nesterov accelerated method on time-varying networks. We theoretically prove and empirically illustrate that these methods are linearly convergent under strong convexity and smoothness of the objective function and specific assumptions on the network structure. Nesterov accelerated method performs better in terms of the objective condition number than other methods in the literature. However, the number of required iterations grows linearly with the number of changes in the network, while other algorithms' performance does not depend on how often the network changes.

The rate $1/(1+\gamma)^N$ can be improved to $\prod_{i=1}^T (1+\gamma_i)^{-1}$, where $\gamma_i = 1/(\sqrt{\kappa_i} - 1)$. This would be a better bound, because not all of the functions $f_k(x)$ may be badly conditioned with $\kappa = L/\mu$. Also, the convergence rate of the accelerated gradient method may be improved by using restarts, which is left for future work. Other accelerating schemes such as the use

of Chebyshev accelerations require further work [1]. Finally, note that we provide only a sufficient condition on the number of changes in the network under which we can guarantee an optimal convergence rate (up to logarithmic terms). Whether this is also a necessary condition remains an open question that we will address in future work.

APPENDIX A
PROOF OF THEOREM II.6

Proof. The proof bases on the connection of strong convexity and smoothness of a function and its conjugate (Lemma II.5).

1) First, we show that the dual norm to $\|\cdot\|_F$ is $\|\cdot\|_F$ itself.

$$\|Y\|_{F*} = \sup_{X \in \mathbb{R}^{d \times n}} \left(\langle X, Y \rangle : \|X\|_F \leqslant 1\right)$$

Note that $\langle X, Y \rangle \leqslant \|X\|_F \|Y\|_F \leqslant 1 \cdot \|Y\|_F$ by Cauchy-Schwarz and $\langle \frac{Y}{\|Y\|_F}, Y \rangle = \|Y\|_F$. Thus, $\|Y\|_{F*} = \|Y\|_F$.

2) Second, let $X \in \mathbb{R}^{d \times n}, A \in \mathbb{R}^{n \times k}, B \in \mathbb{R}^{m \times d}$ and denote $\|X\|_{op}$ the operator norm of X generated by $\|\cdot\|_2$ in $\mathbb{R}^n$, i.e. $\|X\|_{op} = \sup_{y \in \mathbb{R}^n \setminus \{0\}} \frac{\|Xy\|_2}{\|y\|_2}$, where $\|\cdot\|_2$ is the euclidean norm in $\mathbb{R}^d$. Then $\|XA\|_F \leqslant \|X\|_{op} \cdot \|A\|_F$, and $\|BX\|_F \leqslant \|X\|_{op} \cdot \|B\|_F$.

Denote $a_1, ..., a_k$ the columns of A.

$$\|XA\|_F^2 = \|X(a_1...a_k)\|_F^2 = \|(Xa_1 \; Xa_2...Xa_k)\|_F^2$$
$$= \|Xa_1\|_F^2 + ... + \|Xa_k\|_F^2$$
$$\leqslant \|X\|_{op}^2 \cdot \|a_1\|_2^2 + ... + \|X\|_{op}^2 \cdot \|a_k\|_2^2$$
$$= \|X_{op}\|^2 \cdot \|A\|_F^2$$

The inequality $\|BX\|_F \leqslant \|X\|_{op} \cdot \|B\|_F$ is proved analogically.

3) Third, we show the smoothness of $f(X)$.

$$f(X) = \max_{Y \in \mathbb{R}^{d \times n}} \left[-\Psi(Y) - \langle Y, X\sqrt{W} \rangle \right]$$
$$= \max_{Y \in \mathbb{R}^{d \times n}} \left[\langle Y, -X\sqrt{W} \rangle - \Psi(Y) \right] = \Psi^*(-X\sqrt{W})$$

The function $\Psi(Y)$ is μ_Ψ-strongly convex w.r.t. $\|\cdot\|_F$, and thus $\Psi^*(Z)$ is $(1/\mu_\Psi)$-smooth w.r.t. $\|\cdot\|_F$ by Lemma II.5. It remains to show that $f(X) = \Psi^*(-X\sqrt{W})$ is $L_f = -\sqrt{\sigma_{\max}(W)}/(\mu_\Psi)$-smooth, thus

$$df(X) = \langle \nabla\Psi^*(-X\sqrt{W}), -dX\sqrt{W} \rangle$$
$$= \langle -\nabla\Psi^*(-X\sqrt{W})\sqrt{W}, dX \rangle$$
$$\nabla f(X) = -\nabla\Psi^*(-X\sqrt{W})\sqrt{W}. \quad (23)$$

Now when the gradient is computed, we explicitly show that it is Lipschitz with constant L_f.

$$\|\nabla f(X_2) - \nabla f(X_1)\|_F$$
$$\leqslant \|\sqrt{W}\|_{op} \cdot \|\nabla\Psi^*(-X_1\sqrt{W}) - \nabla\Psi^*(-X_2\sqrt{W})\|_F$$
$$\leqslant \|\sqrt{W}\|_{op} \cdot \frac{1}{\mu_\Psi} \|(X_1 - X_2)\sqrt{W}\|_F$$
$$\leqslant \frac{\|\sqrt{W}\|_{op}^2}{\mu_\Psi} \cdot \|X_1 - X_2\|_F = \frac{\sqrt{\sigma_{\max}(W)}}{\mu_\Psi} \|X_1 - X_2\|_F.$$

4) Finally, we prove the strong convexity of $f(X)$. It is sufficient to show

$$f(X + dX) - f(X) \geqslant \langle \nabla f(X), dX \rangle + \frac{\mu_f}{2} \|dX\|_F^2.$$

Keeping in mind that $f(X) = \Psi^*(-X\sqrt{W})$ and $\nabla f(X) = -\nabla\Psi^*(-X\sqrt{W})\sqrt{W}$, we obtain

$$f(X + dX) - f(X)$$
$$= \Psi^*(-(X + dX)\sqrt{W}) - \Psi^*(-X\sqrt{W})$$
$$\geqslant \langle \nabla\Psi^*(-X\sqrt{W}), -dX\sqrt{W}\rangle + 1/(2L_\Psi)\|dX\sqrt{W}\|_F^2$$
$$= \langle -\nabla\Psi^*(-X\sqrt{W})\sqrt{W}, dX\rangle + 1/(2L_\Psi)\|dX\sqrt{W}\|_F^2$$
$$= \langle \nabla f(X), dX\rangle + 1/(2L_\Psi)\|dX\sqrt{W}\|_F^2.$$

Noting that $\|dX\sqrt{W}\|_F^2 \geqslant \|dX\|_F^2 \cdot \sqrt{\tilde{\sigma}_{\min}(W)}$ concludes the proof.

$\square$

APPENDIX B
PROOF OF PROPOSITION V.3

We will need the original result for DIGing obtained in [7]:

Proposition B.1. *Let assumptions V.1 and V.2 hold. Denote $\overline{\mu} = 1/n\sum_{i=1}^n \mu_i$, $\overline{\kappa} = 1/n\sum_{i=1}^n L_i/\mu_i$ and $J = 3\sqrt{\overline{\kappa}}B^2(1 + 4\sqrt{n}\sqrt{\overline{\kappa}})$. Then the DIGing algorithm [7] generates a sequence $\{x_k\}$ such that $\|x_N - x^*\| = O(\lambda^N)$, where λ is defined as*

$$\lambda = \begin{cases} \sqrt[2B]{1 - \frac{\alpha\overline{\mu}}{1.5}}, & \text{if } \alpha \in (0, \alpha_0] \\ \sqrt[B]{\sqrt{\frac{\alpha\overline{\mu}J}{1.5}}, +\delta} & \text{if } \alpha \in \left(\alpha_0, \frac{1.5(1-\delta)^2}{\overline{\mu}J}\right] \end{cases} \quad (24)$$

where $\alpha_0 = 1.5\left(\sqrt{J^2 + (1-\delta^2)J} - \delta J\right)^2/(\overline{\mu}J(J+1)^2)$.

Proof of Proposition V.3. Let us consider the two cases: $\alpha \in (0, \alpha_0]$ and $\alpha \in (\alpha_0, 1.5(1-\delta)^2/(\overline{\mu}J)]$.

1) $\alpha \in (0, \alpha_0]$.

$$\alpha \leqslant \frac{1.5(\sqrt{J^2+J})^2}{\overline{\mu}J(J+1)^2} = \frac{1.5}{\overline{\mu}(J+1)}$$

$$\lambda = \sqrt{1 - \frac{\alpha\overline{\mu}}{1.5}} \leqslant \sqrt{1 - \frac{\overline{\mu}}{1.5}\frac{1.5}{\overline{\mu}(J+1)}}$$

$$\geqslant \sqrt{1 - \frac{1}{3\overline{\kappa}(1 + 4\sqrt{n}\sqrt{\overline{\kappa}}) + 1}} \geqslant \sqrt{1 - \frac{1}{12\overline{\kappa}^{3/2}\sqrt{n}}}$$

$$\geqslant 1 - 1/(12\overline{\kappa}^{3/2}\sqrt{n}).$$

2) $\alpha \in (\alpha_0, \frac{1.5(1-\delta)^2}{\overline{\mu}J}]$.

$$\lambda \geqslant \delta + \sqrt{\frac{\overline{\mu}J}{1.5} \cdot 1.5\frac{\left(\sqrt{J^2 + (1-\delta^2)J} - \delta J\right)^2}{\overline{\mu}J(J+1)^2}}$$

$$= \delta + \frac{\sqrt{J^2 + (1-\delta^2)J} - \delta J}{J+1} = \frac{\sqrt{J^2 + (1-\delta^2)J} + \delta}{J+1}$$

$$\overset{①}{\geqslant} \frac{\sqrt{J^2}}{J+1} = 1 - \frac{1}{J+1} \geqslant 1 - \frac{1}{J} \geqslant 1 - \frac{1}{12\overline{\kappa}^{3/2}\sqrt{n}},$$

where ① is because $0 \leqslant \delta \leqslant 1$. In both cases, $\lambda \geqslant \lambda_0$. $\square$

APPENDIX C
PROOF OF PROPOSITION V.5

The original result for PANDA in [22] states that

Proposition C.1. *Under assumptions V.1 and V.4 the convergence rate of PANDA with step size c is $O(\lambda^k)$, where*

$$\lambda = \sqrt[2B]{1 - \frac{c}{2L}}, \; c \in (0, \alpha], \; and$$

$$\alpha = 2\sqrt{\kappa^{-1}}\mu\left(\frac{\sqrt{(1-\delta^2)\kappa^{-2/3} + 8} - 8\delta}{\kappa^{-3/2} + 8}\right)^2$$

Proof of Proposition V.5. It suffices to show that $\lambda \geqslant \lambda_0$.

$$\alpha \overset{①}{\leqslant} 2\sqrt{\kappa^{-1}}\mu\left(\frac{\sqrt{\kappa^{-2/3} + 8}}{\kappa^{-3/2} + 8}\right)^2 \leqslant 2\sqrt{\kappa^{-1}}\mu \cdot \left(\frac{\sqrt{8+1}}{8+0}\right)^2$$

$$= \frac{9}{32}\sqrt{\kappa^{-1}}\mu$$

$$\lambda \geqslant \sqrt{1 - \frac{\alpha}{2L}} \geqslant \sqrt{1 - \frac{9}{32}\sqrt{\kappa^{-1}}\mu \cdot \frac{1}{2L}} = \sqrt{1 - \frac{9}{64}\kappa^{-3/2}}$$

$$\overset{②}{\geqslant} 1 - \frac{9}{64}\frac{1}{\kappa^{3/2}} = \lambda_0.$$

Here ① is because $\delta \geqslant 0$ due to its definition in Assumption V.1 and ② is since $\sqrt{z} \geqslant z$ for all $z \in [0, 1]$. $\square$

APPENDIX D
PROOF OF PROPOSITION V.6

Proposition D.1. *[9, Theorem 3] Nesterov method on static networks has a convergence rate of $O((1 - C(\mu_\Phi/L_\Phi)^{5/7})^N)$, where $C = (\lambda_2(1 - \lambda_2))^{3/2}/250$ and λ_2 is the second largest eigenvalue of W.*

Proof of Proposition V.6. Recall the notations $\kappa_\Phi = L_\Phi/\mu_\Phi$ and $\chi(W) = \lambda_{\max}/\lambda_{\min}$.

$$\frac{(\lambda_2(1 - \lambda_2))^{3/2}}{250} \cdot \sqrt{\frac{\lambda_{\max}}{\lambda_{\min}}} < \left(\frac{L_\Phi}{\mu_\Phi}\right)^{3/14}$$

$$\frac{(\lambda_2(1 - \lambda_2))^{3/2}}{250} < \kappa_\Phi^{3/14} \cdot (\chi(W))^{-1/2}$$

$$\frac{(\lambda_2(1 - \lambda_2))^{3/2}}{250}\kappa_\Phi^{-5/7} < (\kappa_\Phi\chi(W))^{-1/2}$$

$$\left(1 - \frac{(\lambda_2(1 - \lambda_2))^{3/2}}{250}\kappa_\Phi^{-5/7}\right)^N > \left(1 - \frac{1}{\sqrt{\kappa_\Phi\chi(W)}}\right)^N$$

$\square$

APPENDIX E
LOGISTIC REGRESSION SIMULATIONS

We assume there is a total of nl data points distributed evenly among n agents, where each agent holds l data points. Note that each of the agents in the network will have a local function

$$\varphi_i(x) = \frac{1}{2nl}\sum_{j=1}^l \log\left(1 + \exp\left(-[y^i]_j \cdot [A^i]_j^T x\right)\right) + \frac{1}{2n}c\|x\|_2^2,$$

where $A^j \in \mathbb{R}^{l \times m}$ and $y^j \in \{-1, 1\}^l$ are the data points held by agent j and their corresponding class assignments.

Figure 5 shows a performance comparison between the methods discussed in this paper, for problem (22). We assume a network of 100 agents, that changes every 10 iterations, where each instance is a random geometric graph simulating a group of sensors uniformly distributed over an area of unit length, and a radius that guarantees connectivity of the network. Each agent is assigned a random sample of 100 data
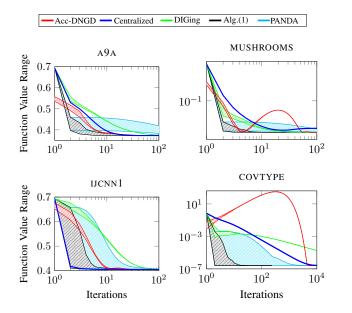
Fig. 5. Function value range achieved by the studied distributed methods for the distributed regularized logistic regression problem. The width of of each band represents the area between the maximum and minimum value achieved by the iterates held by the agents.

points from each of the datasets. For each of the methods, and each of the datasets, we show the range of the function value among all agents. That is, the width of each band corresponds to the values achieved by the current iterates held by the agents. A narrow band represents a relatively high consensus.

REFERENCES

[1] K. Scaman, F. Bach, S. Bubeck, Y. T. Lee, and L. Massoulié, "Optimal algorithms for smooth and strongly convex distributed optimization in networks," in *International Conference on Machine Learning*, 2017, pp. 3027–3036.

[2] K. Scaman, F. Bach, S. Bubeck, L. Massoulié, and Y. T. Lee, "Optimal algorithms for non-smooth distributed optimization in networks," in *Advances in Neural Information Processing Systems 31*, 2018, pp. 2740–2749.

[3] L. He, A. Bian, and M. Jaggi, "Cola: Decentralized linear learning," in *Advances in Neural Information Processing Systems 31*, 2018, pp. 4536–4546.

[4] C. A. Uribe, S. Lee, A. Gasnikov, and A. Nedić, "A dual approach for optimal algorithms in distributed optimization over networks," *arXiv:1809.00710*, 2018.

[5] H. Sun and M. Hong, "Distributed non-convex first-order optimization and information processing: Lower complexity bounds and rate optimal algorithms," in *2018 52nd Asilomar Conference on Signals, Systems, and Computers*, 2018, pp. 38–42.

[6] H. Li, C. Fang, W. Yin, and Z. Lin, "A sharp convergence rate analysis for distributed accelerated gradient methods," *arXiv:1810.01053*, 2018.

[7] A. Nedić, A. Olshevsky, and W. Shi, "Achieving geometric convergence for distributed optimization over time-varying graphs," *SIAM Journal on Optimization*, vol. 27, no. 4, pp. 2597–2633, 2017.

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

[9] G. Qu and N. Li, "Accelerated distributed nesterov gradient descent," *2016 54th Annual Allerton Conference on Communication, Control, and Computing*, 2016.

[10] D. Jakovetic, "A unification and generalization of exact distributed first order methods," *IEEE Transactions on Signal and Information Processing over Networks*, pp. 31–46, 2019.

[11] X. Wu and J. Lu, "Fenchel dual gradient methods for distributed convex optimization over time-varying networks," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, Dec 2017, pp. 2894–2899.

[12] G. Zhang and R. Heusdens, "Distributed optimization using the primal-dual method of multipliers," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 4, no. 1, pp. 173–187, 2018.

[13] A. Gasnikov, "Universal gradient descent," *arXiv:1711.00394*, 2018, [in Russian].

[14] G. Lan, S. Lee, and Y. Zhou, "Communication-efficient algorithms for decentralized and stochastic optimization," *Mathematical Programming*, 2018. [Online]. Available: https://doi.org/10.1007/s10107-018-1355-4

[15] A. Simonetto and G. Leus, "Distributed asynchronous time-varying constrained optimization," in *2014 48th Asilomar Conference on Signals, Systems and Computers*, 2014, pp. 2142–2146.

[16] T. Chen and G. B. Giannakis, "Bandit convex optimization for scalable and dynamic iot management," *IEEE Internet of Things Journal*, vol. 6, no. 1, pp. 1276–1286, 2019.

[17] S. Rahili and W. Ren, "Distributed continuous-time convex optimization with time-varying cost functions," *IEEE Transactions on Automatic Control*, vol. 62, no. 4, pp. 1590–1605, 2017.

[18] M. Fazlyab, C. Nowzari, G. J. Pappas, A. Ribeiro, and V. M. Preciado, "Self-triggered time-varying convex optimization," in *2016 IEEE 55th Conference on Decision and Control (CDC)*, 2016, pp. 3090–3097.

[19] A. Bernstein, E. Dall'Anese, and A. Simonetto, "Online primal-dual methods with measurement feedback for time-varying convex optimization," *IEEE Transactions on Signal Processing*, vol. 67, no. 8, pp. 1978–1991, 2019.

[20] A. B. Y. Tang, E. Dall'Anese and S. Low, "Running primal-dual gradient method for time-varying nonconvex problems," *arXiv:1812.00613*, 2018.

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

[22] M. Maros and J. Jaldén, "Panda: A dual linearly converging method for distributed optimization over time-varying undirected graphs," *2018 IEEE Conference on Decision and Control (CDC)*, pp. 6520–6525, 2018.

[23] H. Nikaido, *Convex Structures and Economic Theory*. Academic Press, Cambridge, 1968.

[24] S. Boyd, P. Diaconis, and L. Xiao, "Fastest mixing markov chain on a graph," *SIAM review*, vol. 46, no. 4, pp. 667–689, 2004.

[25] A. Beck and M. Teboulle, "A fast dual proximal gradient algorithm for convex minimization and applications," *Operations Research Letters*, vol. 42, no. 1, pp. 1–6, 2014.

[26] R. Rockafellar and R. Wets, *Variational analysis*. Springer, 2011, vol. 317.

[27] Y. Nesterov, "Smooth minimization of non-smooth functions," *Mathematical Programming*, vol. 103, no. 1, pp. 127–152, 2005.

[28] S. M. Kakade, S. Shalev-Shwartz, and A. Tewari, "On the duality of strong convexity and strong smoothness: learning applications and matrix regularization," Toyota Technological Institute, Tech. Rep., 2009.

[29] Y. Nesterov, *Introductory Lectures on Convex Optimization. A Basic Course*. Springer Science & Business Media, 2013.

[30] N. Bansal and G. Anupam, "Potential-function proofs for first-order methods," *arXiv:1712.04581*, 2017.

[31] D. P. Bertsekas, *Convex Optimization Theory*. Athena Scientific, 2009.

[32] J. M. Danskin, *The Theory of Max-Min and its Applications to Weapons Allocation Problems*. Springer-Verlag, 1967.

[33] V. F. Demianov and V. N. Malozemov, *Introduction into Min-Max*. Nauka, 1972, [in Russian].

[34] A. Olshevsky, "Linear time average consensus and distributed optimization on fixed graphs," *SIAM Journal on Control and Optimization*, vol. 55, no. 6, pp. 3990–4014, 2017.

[35] C.-C. Chang and C.-J. Lin, "Libsvm: a library for support vector machines," *ACM transactions on intelligent systems and technology (TIST)*, vol. 2, no. 3, p. 27, 2011.