

ECO-PANDA: A COMPUTATIONALLY ECONOMIC, GEOMETRICALLY CONVERGING DUAL OPTIMIZATION METHOD ON TIME-VARYING UNDIRECTED GRAPHS

Marie Maros and Joakim Jaldén

Department of Information Science and Engineering, EECS
Royal Institute of Technology (KTH)
SE-100 44 Stockholm, Sweden

ABSTRACT

In this paper we consider distributed convex optimization over time-varying undirected graphs. We propose a linearized version of primarily averaged network dual ascent (PANDA) that keeps the advantages of PANDA while requiring less computational costs. The proposed method, economic primarily averaged network dual ascent (Eco-PANDA), provably converges at R-linear rate to the optimal point given that the agents' objective functions are strongly convex and have Lipschitz continuous gradients. Therefore, the method is competitive, in terms of type of rate, with both DIGing and PANDA. The proposed method halves the communication costs of methods like DIGing while still converging R-linearly and having the same per iterate complexity.

Index Terms— Distributed Optimization, Time-varying networks, Dual methods

1. INTRODUCTION

In this paper we solve convex optimization problems of the form

$$\min_{\bar{\mathbf{x}} \in \mathbb{R}^p} \sum_{i=1}^n f_i(\bar{\mathbf{x}}) \quad (1)$$

in a decentralized manner within a network of n agents. Agent i possesses exclusive knowledge of its private objective function f_i . Problems of the form of (1) naturally arise in applications such as distributed estimation and control [7, 8, 9, 10], decentralized source localization [11], power grids [12, 13] and distributed learning [14, 15, 16]. Methods to solve problems of the form (1) in a decentralized manner over static networks have received a lot of attention in the last few years. If the individual objective functions, f_i , are strongly convex and have Lipschitz continuous gradients, linearly converging methods such as the decentralized exact first-order algorithm (EXTRA) [17] have previously been proposed. In some applications, the agents may be communicating via wireless links. Due to the random nature of the wireless channel the communication links can be unreliable. This justifies the need for procedures according to which the agents cooperate to solve the optimization problem in (1) over time varying networks. DIGing [2] is the first known method to have established linear convergence over time-varying networks. However, DIGing requires the exchange of both the primal variables and the gradients at each iteration.

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant Agreement no. 742648)

Dual methods have been shown to reach optimal convergence rates in distributed optimization [18]. However, with the exception of PANDA [4, 3], no dual method has been proposed to date that can provably converge linearly on time-varying networks. In [3] we establish PANDA's R-linear converge and demonstrate that it requires the exchange of only the primal variables as opposed to DIGing. However, at each iteration the nodes must solve a convex optimization problem. This is not problematic whenever the objective function is dual friendly [19], but may be computationally expensive otherwise. To alleviate this we propose herein Eco-PANDA, a linearized version of PANDA, capable of converging linearly on time varying graphs while still requiring the exchange of only the primal variables at each iteration.

The remainder of this paper is organized as follows. In Section 2 we formalize the problem further, introduce Eco-PANDA and provide the formal statement regarding Eco-PANDA's convergence. In Section 3 we provide a skeleton of the proof the statement made in Section 2. In Section 4 we empirically illustrate Eco-PANDA's convergence. Finally we end the paper with some concluding remarks.

2. PROBLEM FORMULATION AND ALGORITHM

Let $\mathcal{G}(k) = (\mathcal{V}, \mathcal{E}(k))$ denote the undirected graph representing the network that connects the n agents at time k . The agents are represented by the vertices \mathcal{V} and the connections between the agents by the edges $\mathcal{E}(k)$. The agents i and j can communicate if the edge $(i, j) \in \mathcal{E}(k)$. Further, we assume the graph is undirected, and therefore, if $(i, j) \in \mathcal{E}(k)$ then $(j, i) \in \mathcal{E}(k)$. Further, the set of agents with which agent i can communicate will be denoted $\mathcal{N}_i(k)$. In order to solve the optimization problem (1) in a distributed manner over the graph $\mathcal{G}(k)$ we introduce for each of the agents a copy \mathbf{x}_i of the optimization variable $\bar{\mathbf{x}}$ and define the optimization problem

$$\min_{\mathbf{x} \in \mathbb{R}^{np}} f(\mathbf{x}) = \sum_{i=1}^n f_i(\mathbf{x}_i) \quad (2a)$$

$$\text{s.t. } (\mathbf{I}_{np} - \mathbf{W}(k) \otimes \mathbf{I}_p) \mathbf{x} = 0, \quad (2b)$$

where $\mathbf{x} \triangleq (\mathbf{x}_1^T, \dots, \mathbf{x}_n^T)^T$ and $\mathbf{W}(k)$ fulfils the requirements provided in Assumption 1.

Assumption 1 (Mixing Matrix Sequence $\{\mathbf{W}(k)\}$). *For any $k = 0, \dots, 1$, the mixing matrix $\mathbf{W}(k) \in \mathbb{R}^{n \times n}$ satisfies the following relations:*

(P₁) *Decentralized property: if $i \neq j$ and $(i, j) \notin \mathcal{E}(k)$ $w_{ij}(k) = 0$, i.e., $\mathbf{W}(k)$ is defined on the edges of the graph $\mathcal{G}(k)$.*

(P₂) *Doubly stochastic: $\mathbf{W}(k)\mathbf{1}_n = \mathbf{1}_n$, $\mathbf{1}_n^T \mathbf{W}(k) = \mathbf{1}_n^T$.*

(P₃) *Joint spectrum property*: Let σ_{\max} denote the largest singular value of a matrix and let

$$\mathbf{W}_b(k) \triangleq \mathbf{W}(k)\mathbf{W}(k-1)\dots\mathbf{W}(k-b+1), \quad (3)$$

for $k \geq 0$ and $b \geq k-1$, with $\mathbf{W}_b(k) = \mathbf{I}_n$ for $k < 0$ and $\mathbf{W}_0(k) = \mathbf{I}_n$. Then, there exists a positive B such that

$$\sup_{k \geq B-1} \delta(k) = \delta < 1, \quad (4)$$

where

$$\delta(k) = \sigma_{\max} \left\{ \mathbf{W}_B(k) - \frac{1}{n} \mathbf{1}_n^T \mathbf{1}_n \right\}, \quad \forall k = 0, 1, \dots \quad (5)$$

Properties (P₁) and (P₂) in Assumption 1 are common in the consensus literature while (P₃) is a requirement on the time-varying nature of the graph $\mathcal{G}(k)$. Examples and discussion on condition (P₃) can be found in [2].

Eco-PANDA follows the same intuition as PANDA. For an intuition on PANDA please refer to [3]. PANDA solves the optimization problem

$$\max_{\mathbf{y}: (\mathbf{I}_n - \frac{1}{n} \mathbf{1}_n \mathbf{1}_n^T) \otimes \mathbf{I}_p \mathbf{y} = 0} f^*(\mathbf{y}), \quad (6)$$

where $\mathbf{1}_n \in \mathbb{R}^n$ denotes the vector of all ones and $f^*(\mathbf{y}) \triangleq \min_{\mathbf{x} \in \mathbb{R}^{np}} \mathbf{y}^T \mathbf{x} - f(\mathbf{x})$ denotes the Fenchel conjugate of f by applying approximated projected gradient descent on (6), where the projection is approximated via a gradient tracking scheme much like the schemes in [2] and [5]. While PANDA achieves R-linear convergence rates by communicating half as many variables as DIGing per iteration, the communication cost is decreased at the expense of increased computational costs. In particular, PANDA requires obtaining the minimizer

$$\mathbf{x}_p(k+1) := \nabla f^*(\mathbf{y}(k)) = \min_{\mathbf{x} \in \mathbb{R}^{np}} f(\mathbf{x}) - \mathbf{y}(k)^T \mathbf{x}, \quad (7)$$

for a given value of $\mathbf{y}(k)$ at each iteration, while DIGing requires only a simple update in the direction of the approximated average gradient. Eco-PANDA addresses this by using a quadratic upper bound on the objective functions that leads to simple iterates. Before introducing Eco-PANDA we will introduce the following assumption on the objective functions.

Assumption 2 (Strong convexity and Lipschitz continuity of gradients). *The objective function f is μ -strongly convex and its gradient is L -Lipschitz continuous.*

In order to make PANDA competitive to DIGing in terms of iterate complexity we propose to upper bound the objective function in (7) by

$$f(\mathbf{x}(k)) + (\nabla f(\mathbf{x}(k)) - \mathbf{y}(k))^T (\mathbf{x} - \mathbf{x}(k)) + \frac{\eta}{2} \|\mathbf{x} - \mathbf{x}(k)\|^2. \quad (8)$$

Since the function f is convex and is L -Lipschitz continuous, as long as $\eta > L$ (8) is an upper bound to the objective function in (7). Further, the upper bound becomes tight as $\mathbf{y}(k) \rightarrow \mathbf{y}^*$ and $\mathbf{x}(k) \rightarrow \mathbf{x}^*$. With this we are ready to introduce Eco-PANDA which is formally written in Algorithm 1.

We are now ready to provide the paper's main statement.

Algorithm 1 Eco-PANDA

- 1: Choose step size $c > 0$, $\eta > L$ and pick $\mathbf{z}(0) = \mathbf{x}(0)$ and $\mathbf{y}(0)$ such that $(\mathbf{\Pi}_{\mathbf{1}_n} \otimes \mathbf{I}_p) \mathbf{y}(0) = \mathbf{0}$.
- 2: **for** $k = 0, 1, \dots$ **do** each agent $i = 1, \dots, n$:
- 3: **computes**

$$\mathbf{x}_i(k+1) := \mathbf{x}_i(k) - \frac{1}{\eta} (\nabla f_i(\mathbf{x}_i) - \mathbf{y}_i(k))$$

- 4: **exchanges** $\mathbf{z}_i(k)$ with the agents in $\mathcal{N}_i(k)$.
- 5: **computes**

$$\mathbf{z}_i(k+1) := \sum_{j \in \mathcal{N}_i(k) \cup \{i\}} w_{ij}(k) \mathbf{z}_j(k) + \mathbf{x}_i(k+1) - \mathbf{x}_i(k)$$

- 6: **and computes**

$$\mathbf{y}_i(k+1) := \mathbf{y}_i(k) - c(\mathbf{x}_i(k+1) - \mathbf{z}_i(k+1))$$

- 7: **end for**
-

Theorem 1 (Eco-PANDA converges R-linearly). *Let Assumptions 1 and 2 hold. Also, let $\kappa \triangleq \frac{L}{\mu}$ denote the condition number of f and $q \triangleq 1 - \frac{\mu}{\eta}$. Define the quantities*

$$C \triangleq \frac{q(1 - \sqrt{q})}{(\max\{\frac{1}{\eta}, \frac{q}{\mu}\} + B(1 - \delta))(3 + q)}, \quad \alpha \triangleq \frac{C}{\mu} < 1 \quad (9)$$

$$\lambda_{\bar{c}} \triangleq \frac{\alpha\delta + \sqrt{4\kappa^{3/2}(4\kappa^{3/2} - \alpha\delta^2)}}{\alpha + \kappa^{3/2}}, \quad \bar{c} \triangleq \frac{\alpha\mu(\lambda_{\bar{c}}^B - \delta)^2}{2\sqrt{\kappa}} \quad (10)$$

Then for any step size

$$c \in \left(0, \frac{\alpha\mu(1 - \delta)^2}{2} \right) \quad (11)$$

the sequence $\{\mathbf{y}(k)\}_{k \geq 0}$ converges to \mathbf{y}^* , the unique solution of (6), and the sequences $\{\mathbf{x}_i(k)\}_{k \geq 0}$ converge to $\bar{\mathbf{x}}^*$ the unique solution of (1), at a global R-linear rate of at least $\mathcal{O}(\lambda^k)$, where $\lambda < 1$ is given by

$$\lambda = \begin{cases} \sqrt[2B]{1 - \frac{c}{2L}} & c \in (0, \bar{c}] \\ \sqrt[2B]{\delta + \sqrt{\frac{2c\sqrt{\kappa}}{\alpha\mu}}} & c \in (\bar{c}, \frac{\alpha\mu(1-\delta)^2}{2}) \end{cases} \quad (12)$$

When comparing the statement in Theorem 1 to the main statement in [3] one can see that the upper bound on the convergence rate is affected by the parameter α . In particular, for PANDA $\alpha = 1$, while for Eco-PANDA $\alpha < 1$.

Due to space restrictions, the convergence proof is not provided fully in the following section. A full version of the paper with the convergence proofs can be found in [1]. However, a skeleton of the proof of convergence is provided in the following section in the sense that the main intermediate results required for the overall proof are stated. Further, an intuition and connections to PANDA are established at every point.

3. SKELETON OF PROOF OF CONVERGENCE

This section is devoted to providing the guidelines to proof Theorem 1. We will start the section by introducing some notation and the small gain theorem, which is the main tool used to establish the

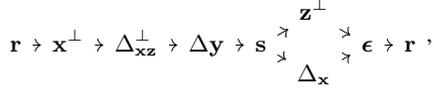


Fig. 1. Cycle of arrows

convergence of Eco-PANDA. Let $\mathbf{s}^i \triangleq \{\mathbf{s}^i(0), \mathbf{s}^i(1), \dots\}$ denote an infinite sequence of vectors $\mathbf{s}^i(k) \in \mathbb{R}^{np}$, for $i = 1, \dots, m$. Further, let

$$\|\mathbf{s}^i\|^{\lambda, K} \triangleq \sup_{k=0, \dots, K} \frac{1}{\lambda^k} \|\mathbf{s}^i(k)\|, \quad \|\mathbf{s}^i\|^\lambda \triangleq \sup_{k \geq 0} \frac{1}{\lambda^k} \|\mathbf{s}^i(k)\|. \quad (13)$$

Theorem 2 (Small gain theorem [2]). *Suppose $\mathbf{s}^1, \dots, \mathbf{s}^m$ are vector sequences such that for all positive k and for each $i = 1, \dots, m$, we have an arrow $\mathbf{s}^i \rightarrow \mathbf{s}^{(i \bmod m)+1}$, i.e.,*

$$\|\mathbf{s}^{(i \bmod m)+1}\|^{\lambda, K} \leq \gamma_i \|\mathbf{s}^i\|^{\lambda, K} + \omega_i, \quad (14)$$

where the constants $\gamma_1, \dots, \gamma_m$ and $\omega_1, \dots, \omega_m$ are independent of K . Further, suppose that the constants $\gamma_1, \dots, \gamma_m$ are non-negative and satisfy $\gamma_1 \dots \gamma_m < 1$. Then we have that

$$\|\mathbf{s}^1\|^\lambda \leq \frac{1}{1 - \gamma_1 \gamma_2 \dots \gamma_m} (\omega_1 \gamma_1 \gamma_3 \dots \gamma_m + \omega_2 \gamma_3 \gamma_4 \dots \gamma_m + \dots + \omega_{m-1} \gamma_m + \omega_m). \quad (15)$$

Further, if $\|\mathbf{s}^1\|^\lambda < C$ where $C < \infty$, $\|\mathbf{s}^1(k)\|$ converges to zero exponentially fast and at rate λ . Proof of this statement can be found in [2]. Further, due to the cyclic nature of the small gain theorem all sequences will converge to zero exponentially fast at rate λ .

In order to establish the convergence of Eco-PANDA we will first establish the cycle of arrows in Fig. 1 where, given that $\Pi_{\mathbf{I}_n}^\perp \triangleq (\mathbf{I}_n - \frac{1}{n} \mathbf{1}_n \mathbf{1}_n^T) \otimes \mathbf{I}_p$, the sequences are defined as $\mathbf{r}(k) \triangleq \mathbf{y}(k) - \mathbf{y}^*$, $\mathbf{x}^\perp \triangleq \Pi_{\mathbf{I}_n}^\perp \mathbf{x}(k+1)$, $\Delta_{\mathbf{xz}}^\perp(k) \triangleq \Pi_{\mathbf{I}_n}^\perp (\mathbf{x}(k) - \mathbf{z}(k))$, $\Delta \mathbf{y}(k) \triangleq \mathbf{y}(k) - \mathbf{y}(k-1)$, $\mathbf{s}(k) \triangleq (\|\mathbf{x}(k) - \mathbf{x}(k-1)\|, \|\mathbf{x}(k) - \mathbf{x}_p(k)\|)^T$, (c.f. (7) for the definition of $\mathbf{x}_p(k)$), $\mathbf{z}^\perp(k) \triangleq \Pi_{\mathbf{I}_n}^\perp \mathbf{z}(k)$, $\Delta_{\mathbf{x}}(k) \triangleq \mathbf{x}(k) - \mathbf{x}_p(k)$, and $\boldsymbol{\epsilon}(k) \triangleq \Pi_{\mathbf{I}_n}^\perp (\mathbf{x}(k) - \mathbf{x}_p(k) - \mathbf{z}(k))$. For consistency in the proofs we will define $\mathbf{y}(-1) = \mathbf{y}(0)$ and $\mathbf{x}(-1) = \mathbf{x}(0) = \mathbf{z}(0)$.

Before starting to prove that each of the arrows holds true we require some intermediate results. Proof of the two lemmas introduced below can be found in [3].

Lemma 1 (Well conditioned Fenchel Conjugate [3]). *The Fenchel conjugate f^* is $\frac{1}{L}$ -strongly convex and has $\frac{1}{\mu}$ -Lipschitz continuous gradients if and only if the objective function f is μ -strongly convex and has L -Lipschitz continuous gradients.*

Lemma 2 (Equivalent iterates [3]). *The Eco-PANDA iterates can be equivalently written as*

$$\mathbf{x}(k+1) := \mathbf{x}(k) - \frac{1}{\eta} (\nabla f(\mathbf{x}) - \mathbf{y}(k)) \quad (16a)$$

$$\mathbf{z}(k+1) := (\mathbf{W}(k) \otimes \mathbf{I}_p) \mathbf{z}(k) + \mathbf{x}(k+1) - \mathbf{x}(k) \quad (16b)$$

$$\mathbf{y}(k+1) := \mathbf{y}(k) - c(\Pi_{\mathbf{I}_n}^\perp \otimes \mathbf{I}_p)(\mathbf{x}(k+1) - \mathbf{z}(k+1)) \quad (16c)$$

Where the iterates in \mathbf{x} can be seen as a step in the direction of the gradient of (7) with time varying \mathbf{y} , the iterate in \mathbf{z} tracks the moving average of \mathbf{x} and the iterate in \mathbf{y} corresponds to an inexact projected gradient step. We now formally express each of the arrows in the form of (14).

Lemma 3 ($\mathbf{r} \rightarrow \mathbf{x}^\perp$). *If f is μ -strongly convex (c.f. Assumption 2), it holds that*

$$\|\mathbf{x}^\perp\|^{\lambda, K} \leq \frac{2\lambda + (1+q)}{\lambda^2 \mu (1 - \frac{q}{\lambda})} \|\mathbf{r}\|^{\lambda, K}, \quad (17)$$

for all $\lambda \in (q, 1)$ and $K \geq 0$, where $q \triangleq 1 - \frac{\mu}{\eta}$ for $\eta > L$.

The proof of Lemma 3 uses the fact that the iterate (16a) of Eco-PANDA can be seen as running gradient descent on the time varying objective function $f(\mathbf{x}) - \mathbf{y}(k)^T \mathbf{x}$. In other words, every time the objective function changes, via the change of $\mathbf{y}(k)$ a single gradient step is taken using the previous iteration as a warm start. This strategy has been used to track the time varying optimizer of a sequence of time varying functions in [6]. In particular, we use the results in [6] to establish a connection between $\mathbf{x}(k+1)$ and $\mathbf{x}_p(k+1)$ (c.f. (7)). Then, we use the results in [3] to finalize the proof.

Lemma 4 ($\mathbf{x}^\perp \rightarrow \Delta_{\mathbf{xz}}^\perp$). *Under Assumption 1 it holds that*

$$\|\Delta_{\mathbf{xz}}^\perp\| \leq \frac{2(1 - \lambda^B)}{(1 - \lambda)(\lambda^B - \delta)} \|\mathbf{x}^\perp\|^{\lambda, K} + \frac{\lambda^B}{\lambda^B - \delta} \sum_{t=1}^B \lambda^{1-t} \|\Delta_{\mathbf{xz}}^\perp(t-1)\|, \quad (18)$$

for all $\lambda \in (\delta^{\frac{1}{B}}, 1)$ and for all $K \geq 0$.

Lemma 5 ($\Delta_{\mathbf{xz}}^\perp \rightarrow \Delta \mathbf{y}$). *It holds that*

$$\|\Delta_{\mathbf{xz}}\|^{\lambda, K} \leq c \|\Delta \mathbf{y}\|^{\lambda, K}, \quad (19)$$

for any $c > 0$, $K > 0$ and $\lambda \in (0, 1)$.

Lemma 6 ($\Delta \mathbf{y} \rightarrow \mathbf{s}$). *Given that the objective function f is μ -strongly convex (c.f. Assumption 2) it holds that*

$$\|\mathbf{s}\|^{\lambda, K} \leq \frac{\lambda \max\{\frac{1}{\eta}, \frac{q}{\mu}\}}{\lambda - \rho} \|\Delta \mathbf{y}\|^{\lambda, K}, \quad \text{for } \lambda \in (\rho, 1). \quad (20)$$

In order to establish how well the average tracking sequence $\mathbf{z}(k+1)$ is performing, we require to understand how the quantity $\mathbf{x}(k+1) - \mathbf{x}(k)$ is behaving. In the case of PANDA, this is fully characterized by the quantity $\mathbf{y}(k) - \mathbf{y}(k-1)$ through Lipschitz continuity of the gradients of the Fenchel conjugate. However, the variation $\mathbf{x}(k+1) - \mathbf{x}(k)$ in the case of Eco-PANDA does not exclusively depend on the difference $\mathbf{y}(k) - \mathbf{y}(k-1)$ but also on how close to $\mathbf{x}_p(k)$ $\mathbf{x}(k)$ was. This is captured in the statement in Lemma 6.

The following lemmas are introduced with no proof. For a proof of the lemmas see full version of the paper.

Lemma 7 ($\mathbf{s} \rightarrow \Delta_{\mathbf{x}}$). *It holds that $\|\Delta_{\mathbf{x}}\|^{\lambda, K} \leq \|\mathbf{s}\|^{\lambda, K}$.*

Lemma 8 ($\mathbf{s} \rightarrow \mathbf{z}^\perp$). *Under assumption 1 it holds that*

$$\|\mathbf{z}^\perp\|^{\lambda, K} \leq \frac{1 - \lambda^B}{(1 - \lambda)(\lambda^B - \delta)} \|\mathbf{s}\|^{\lambda, K} + \frac{\lambda^B}{\lambda^B - \delta} \sum_{t=1}^B \lambda^{1-t} \|\mathbf{z}(t-1)\|, \quad (21)$$

for $\lambda^B \in (\delta, 1)$.

Lemma 9 ($\Delta_{\mathbf{x}}, \mathbf{z}^\perp \rightarrow \epsilon$). It holds that $\|\epsilon\|^{\lambda, K} \leq \|\mathbf{s}\|^{\lambda, K} + \|\mathbf{z}^\perp\|^{\lambda, K}$.

Lemma 10 ($\epsilon \rightarrow \mathbf{r}$). Under Assumptions 1 and 2, $c \in (0, \frac{\mu}{2}]$ and $\lambda \in [\sqrt{1 - \frac{c}{2L}}, 1)$ it holds that

$$\|\mathbf{r}\|^{\lambda, K} \leq \sqrt{L\mu}\|\epsilon\|^{\lambda, K} + 2\|\mathbf{r}(0)\|. \quad (22)$$

We now have established all the arrows in Fig. 1 and are ready to invoke the small gain theorem. In order to be able to claim that Eco-PANDA converges R-linearly we have to establish that there exists a tuple (λ, c, η) such that the conditions

$$\frac{2(2\lambda + 1 + q)(1 - \lambda^B)A(1 - \lambda^B + (1 - \lambda)(\lambda^B - \delta))}{(\lambda - q)(\lambda - \rho)((1 - \lambda)(\lambda^B - \delta))^2} \times c < \frac{1}{\sqrt{\kappa}}, \text{ with } A \triangleq \max\left\{\frac{1}{\eta}, \frac{q}{\mu}\right\} \quad (23)$$

$$\lambda \in \left(\max\left\{q, \rho, \delta^{1/B}, \sqrt{1 - \frac{c}{2L}}\right\}, 1\right), \quad c \in \left(0, \frac{\mu}{2}\right] \quad (24)$$

for $q = 1 - \frac{\mu}{\eta}$, $\eta > 0$ are fulfilled.

By using that for $0.5 \leq \lambda \leq 1$ it holds that $(1 - \lambda^B)/(1 - \lambda) \leq B$ [2], we have that if c fulfils

$$c \leq \frac{(\lambda - q)(\lambda - \rho)(\lambda^B - \delta)^2}{\sqrt{\kappa}(A + B(\lambda^B - \delta))(2(2\lambda + 1 + q))} \quad (25)$$

it will also fulfil (23). Recall now that $q = 1 - \frac{\mu}{\eta}$ and $\rho = \max\{q, \frac{L}{\eta}\}$. Further, by inspecting (24) we may conclude that $\lambda \geq \sqrt{1 - \frac{\mu}{4L}} \triangleq \lambda_{\min}$. In particular, by choosing $\eta = 4L$ we have that $\sqrt{q} = \lambda_{\min}$ and also $\rho = q$. Therefore, we conclude that if we chose $\eta = 4L$ we have that any c fulfilling $c \leq \frac{C(\lambda^B - \delta)^2}{2\sqrt{\kappa}}$, where fulfils (23). This implies that we are to establish that there exists a pair (λ, c) fulfilling

$$c \leq \frac{C(\lambda^B - \delta)^2}{2\sqrt{\kappa}}, \quad c \geq 2L(1 - \lambda^2) \quad (26)$$

$$\lambda \in \left(\delta^{1/B}, 1\right), \quad c \in \left(0, \frac{\mu}{2}\right] \quad (27)$$

given that $\eta = 4L$. Note that evaluating the value of q and A in C for a more detailed discussion refer to [full version]. For convenience let us write $C = \alpha\mu$ for $\alpha \in (0, 1)$. Then, the procedure to establish the result is identical to that in Section 4.C in [3].

4. NUMERICAL EXPERIMENTS

In order to illustrate the performance of Eco-PANDA, consider the decentralized estimation problem

$$\min_{\mathbf{x} \in \mathbb{R}^p} \frac{1}{2nd} \sum_{i=1}^n \|\mathbf{H}_i \mathbf{x} - \mathbf{b}_i\|^2 + \frac{r}{2} \|\mathbf{x}\|^2. \quad (28)$$

The elements of a vector $\hat{\mathbf{x}} \in \mathbb{R}^d$ are generated independently and according to the distribution $\mathcal{N}(0, 10)$. The elements of the matrices \mathbf{H}_i are generated independently and according to the distribution $\mathcal{N}(0, 0.1)$. Then, the vectors \mathbf{b}_i are generated according to the equation $\mathbf{b}_i = \mathbf{H}_i \hat{\mathbf{x}} + \boldsymbol{\xi}_i$ where the elements of $\boldsymbol{\xi}_i$ are generated independently and according to the distribution $\mathcal{N}(0, 0.1)$. The sequence of graphs $\mathcal{G}(\mathcal{V}, \mathcal{E}(k))$ on which the problem (28) is solved

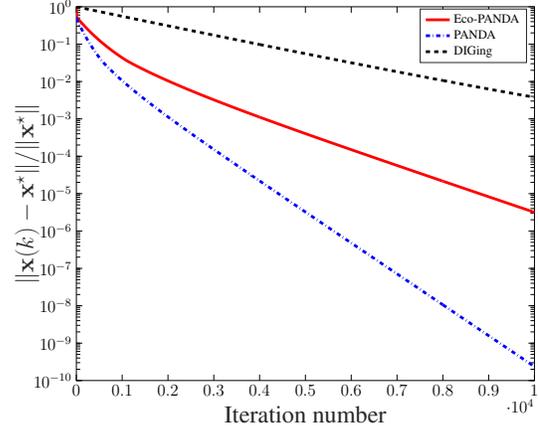


Fig. 2. Performance of Eco-PANDA, PANDA and DIGing.

is randomly generated. In particular, each bi-directional arc will belong to the set of edges with probability π . The Metropolis Hastings matrix is chosen for both methods, i.e. the mixing matrix is given by

$$w_{i,j}(k) = \begin{cases} \frac{1}{\max\{d_i(k), d_j(k)\}} & (i, j) \in \mathcal{E}(k) \\ 1 - \sum_{j \in \mathcal{N}_i(k)} w_{i,j}(k) & i = j \\ 0 & \text{otherwise.} \end{cases} \quad (29)$$

Fig. 2 is generated by setting $n = 10$, $p = 3$, $d = 5$, $r = 1$ and $\pi = 0.1$. Picking $r = 1$ allows for the optimization problems to be well conditioned allowing for PANDA and Eco-PANDA to perform well. On ill conditioned problems, DIGing is expected to perform better [3]. This observation is consistent with the empirical evidence in [3] and the analysis in both [3] and [2], where in [3] we require the strong convexity of each of the function components and in [2] the strong convexity of only one of them is required. The step-sizes for the methods have been hand-optimized and are set to $c = 0.0005$, $\eta = 0.5$ for Eco-PANDA, $c_{\text{PANDA}} = 0.001$ for PANDA and $\alpha_{\text{DIGing}} = 0.003$ for DIGing respectively.

From Fig. 2 one can see that Eco-PANDA is outperformed by PANDA in terms of the number of iterations required to achieve certain accuracy. This is reasonable as the iterates of Eco-PANDA are an approximation of those of PANDA. Note however that Eco-PANDA could still be preferable in terms of computation cost due to the simpler iterates. Further, as could be concluded from Theorem 1 Eco-PANDA tolerates a smaller step size before starting to diverge than PANDA. However, the iterates of PANDA require a matrix inversion at each iterate to solve the optimization problem (6) while the iterates of Eco-PANDA require just the computation of the gradient.

5. CONCLUSIONS

In this paper we proposed Eco-PANDA a linearized version of PANDA that we have proven converges linearly on time-varying graphs. By doing this we have circumvented the main disadvantage PANDA has when compared to DIGing, that is, that the iterates are computationally speaking more expensive. Eco-PANDA alleviates the nodes' computational burden while keeping the low communication cost of PANDA. We analytically and numerically observe the performance loss due to the having simpler iterates. The loss in performance is directly related to the shrinking of the largest allowed step-size which provides for the best convergence bound.

6. REFERENCES

- [1] M. Maros, J. Jaldén, “Eco-PANDA: A Computationally Economic, Geometrically Converging Dual Optimization Method on Time-Varying Undirected graphs,” Full version, arXiv:1810.12240, October 2018.
- [2] 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.
- [3] M. Maros, J. Jaldén, “A geometrically converging dual method for distributed optimization over time-varying graphs,” October 2018.
- [4] M. Maros, J. Jaldén, “PANDA: A Dual Linearly Converging Method for Distributed Optimization over Time-Varying Undirected Graphs,” arXiv:105.01660, March 2018.
- [5] M. Zhu, S. Martínez, “Discrete-time dynamic average consensus,” *Automatica*, vol. 46, no. 2, pp. 322-329, February 2010.
- [6] A.Yu. Popkov, “Gradient Methods for Nonstationary Unconstrained Optimization Problems,” *Automation and Remote Control*, vol.66, no. 6, pp.883-891, 2005.
- [7] J. Xiao, A. Ribeiro, Z.Q. Luo, and G.B. Giannakis, “Distributed compression-estimation using wireless sensor networks,” *IEEE Signal Processing Magazine*, vol.23, no.4, pp.27-41, July 2006.
- [8] I. Neocara, V. Nedelcu, and I. Dumitrache, “Parallel and distributed optimization methods for estimation and control in networks,” *Journal of Process Control*, vol.21, no. 5, pp. 756-766, 2011.
- [9] R.L. Raffard, C.J. Tomlin, and S.P. Boyd, “Distributed optimization for cooperative agents: application to formation flight,” in *2004 43rd IEEE Conference on Decision and Control (CDC)*, Dec. 2004, vol.3, pp. 2453-2459.
- [10] I.D. Schizas, A. Ribeiro, and G.B. Giannakis, “Consensus in ad hoc WSNs with noisy links; Part I: Distributed estimation of deterministic signals,” *IEEE Transactions on Signal Processing*, vol. 56, no.1, pp. 350-364, Jan 2008.
- [11] M.G. Rabbat, and R.D. Nowak, “Decentralized source localization and tracking,” in *2004 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, May 2004.
- [12] G.B. Giannakis, V. Kekatos, N. Gatsis, S.J. Kim, H. Zhu, and B.F. Wollenberg, “Monitoring and optimization for power grids: A signal processing perspective,” *IEEE Signal Processing Magazine*, vol. 30, no. 5, pp. 107-128, Sept. 2013.
- [13] T. Yang, J. Lu, D. Wu, J. Wu, et. al., “A distributed algorithm for economic dispatch over time-varying directed networks with delays,” *IEEE Transactions on Industrial Electronics*, vol. 64, no.6, pp. 5095-5106, June 2017.
- [14] A. Nedić, A. Olshevsky, and C.A. Uribe, “Fast convergence rates for distributed non-bayesian learning,” *IEEE Transactions on Automatic Control*, vol. 62, no. 11, pp. 5538-5553, Nov. 2017.
- [15] G. Giannakis, Q. Ling, G. Mateos, et. al. “Decentralized learning for wireless communications and networking,” in *Splitting methods in Communications, Imaging, Science and Engineering*, Eds. R. Glowinski, S. Osher, and W. Yin, pp. 461-467, Springer, 2016.
- [16] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, “Non-bayesian social learning,” *Games and Economic Behavior*, vol. 76, no. 1, pp. 48-61, 2012.
- [17] W. Shi, Q. Ling, G. Wu, and W. Yin, “EXTRA: An Exact first-order algorithm for decentralized consensus optimization,” *SIAM Journal on Optimization*, vol. 25, no. 2, pp. 944-966, 2015.
- [18] K. Scaman, F. Back, S. Bubeck, et. al. “Optimal algorithms for smooth and strongly convex distributed optimization in networks,” arXiv: 1702.08704, 2017.
- [19] C.A. Uribe, S. Lee, A. Gasnikov, and A. Nedić, “Optimal Algorithms for Distributed Optimization,” arXiv:1712.00232, Dec. 2017.