

A Multi-Agent Primal-Dual Strategy for Composite Optimization over Distributed Features

Sulaiman A. Alghunaim

Department of Electrical Engineering
Kuwait University
Kuwait
salghunaim@ucla.edu

Ming Yan

Departments of CMSE & Math
Michigan State University
East Lansing, MI, USA
myan@msu.edu

Ali H. Sayed

School of Engineering
Ecole Polytechnique Federale de Lausanne
Switzerland
ali.sayed@epfl.ch

Abstract—This work studies multi-agent sharing optimization problems with the objective function being the sum of smooth local functions plus a convex (possibly non-smooth) function coupling all agents. This scenario arises in many machine learning and engineering applications, such as regression over distributed features and resource allocation. We reformulate this problem into an equivalent saddle-point problem, which is amenable to decentralized solutions. We then propose a proximal primal-dual algorithm and establish its linear convergence to the optimal solution when the local functions are strongly-convex. To our knowledge, this is the first linearly convergent decentralized algorithm for multi-agent sharing problems with a general convex (possibly non-smooth) coupling function.

Index Terms—Decentralized composite optimization, primal-dual methods, linear convergence, distributed learning.

I. INTRODUCTION

We consider K agents connected through a graph. Each agent can only send and receive information from its immediate neighbors. Its goal is to find its corresponding solution, denoted by $w_k^* \in \mathbb{R}^{Q_k}$, of the following coupled multi-agent optimization problem:

$$\min_{w_1, \dots, w_K} \sum_{k=1}^K J_k(w_k) + g\left(\sum_{k=1}^K B_k w_k\right), \quad (1)$$

where the smooth function $J_k : \mathbb{R}^{Q_k} \rightarrow \mathbb{R}$ and the matrix $B_k \in \mathbb{R}^{E \times Q_k}$ are known by agent k only, and $g : \mathbb{R}^E \rightarrow \mathbb{R} \cup \{+\infty\}$ is a convex possibly non-smooth function known by all agents. Problem (1) is the sharing formulation, where the agents own different variables but are coupled through the function g . Problems of the form (1) appear in many machine learning applications, such as regression over distributed features [1], [2], dictionary learning over distributed models [3], and clustering in graphs [4]. They also appear in engineering applications, including smart grid control [5] and network utility maximization [6]. For a general convex function g , centralized algorithms for (1) have been shown to achieve global linear convergence if the matrix $[B_1, \dots, B_K]$ has full row rank and $\sum_{k=1}^K J_k(w_k)$ is strongly-convex [7]–[9]. On the other hand, *decentralized* algorithms have only been shown to converge linearly under stricter conditions than centralized ones. In this work, we aim toward closing the gap in linear

convergence between centralized and decentralized algorithms for problem (1).

Literature Review. Sharing problems have been studied in different fields and date back to studies in economics [10]—see the discussion in [1]. The earliest center-free algorithm to solve such problems dates back to [11]. Decentralized solutions for the sharing formulation (1) have only been shown to achieve global linear convergence for special cases and under stricter conditions compared to the centralized ones as we now explain. The works [12]–[15] establish linear convergence for resource allocation formulations, where g is an indicator function of zero (i.e., $g(x) = 0$ if $x = 0$ and ∞ otherwise) and $B_k = I$, for smooth and strongly-convex local costs. The work [16] also establishes linear convergence for resource allocation problems in the presence of simple local constraints (i.e., $\underline{w} \leq w_k \leq \bar{w}$), but under stronger assumptions on the costs such as twice differentiability of J_k and knowledge of the conjugate function of J_k . The works [17], [18] establish linear convergence for special cases with g being an indicator function of zero; moreover, each B_k is required to have full row rank in [17] or satisfy a certain rank condition in [18]. The works [19], [20] establish linear convergence for strongly convex objectives and a smooth coupling function g .

Note that problem (1) recovers the consensus problem [21]–[23] if we choose g such that it returns zero when $w_1 = \dots = w_K$ and ∞ otherwise. In this case, the matrix $B = [B_1, \dots, B_K]$ is sparse and encodes the communication graph between agents. The works [24]–[26] studied linear convergence of consensus problems in the presence of a common non-smooth term, which is not applicable for the sharing problem. Different from the consensus problem, the matrix B in the sharing problem is not necessarily sparse, and B_k is a private matrix known by agent k only. Thus, solution methods for these two problems are different [1]. To the best of our knowledge, establishing linearly convergent algorithms for the *sharing* problem (1) with a general convex g are still missing.

Contribution. As mentioned before, there is a theoretical gap between centralized and decentralized algorithms for the sharing problem (1). In this work, we propose a decentralized algorithm for (1) and establish its linear convergence to the global solution. The derivation of our algorithm is based on

reformulating (1) into an equivalent problem that is amenable to decentralized solutions. This technique motivates the derivation of many other decentralized algorithms.

Notation. We let $\|x\|_D^2 = x^\top D x$ for a square matrix D . The symbol I_S denotes the identity matrix of size S , and S is removed when there is no confusion. The symbol $\mathbb{1}_N$ denotes the $N \times 1$ vector with all entries being one. The Kronecker product of two matrices is denoted by $A \otimes B$. We use $\text{col}\{x_k\}_{k=1}^K$ to denote a column vector formed by stacking x_1, \dots, x_K on top of each other and $\text{blkdiag}\{X_k\}_{k=1}^K$ to denote a block diagonal matrix consisting of diagonal blocks $\{X_k\}$. The subdifferential $\partial f(x)$ of a function f at x is the set of all subgradients. The proximal operator of a function $f(x)$ with step-size μ is $\text{prox}_{\mu f}(x) = \arg \min_u f(u) + \frac{1}{2\mu} \|x - u\|^2$. The conjugate of a function f is defined as $f^*(v) = \sup_x v^\top x - f(x)$. A differentiable function f is δ -smooth and ν -strongly-convex if $\|\nabla f(x) - \nabla f(y)\| \leq \delta \|x - y\|$ and $(x - y)^\top (\nabla f(x) - \nabla f(y)) \geq \nu \|x - y\|^2$, respectively, for any x and y .

II. SADDLE-POINT REFORMULATION

In this section, we provide the main assumption on the objective and explain how (1) is reformulated into an equivalent saddle-point problem. We introduce the quantities:

$$w \triangleq \text{col}\{w_1, \dots, w_K\} \in \mathbb{R}^Q, \quad Q \triangleq \sum_{k=1}^K Q_k, \quad (2a)$$

$$\mathcal{J}(w) \triangleq \sum_{k=1}^K J_k(w_k), \quad (2b)$$

$$B \triangleq [B_1 \ \dots \ B_K] \in \mathbb{R}^{E \times Q}. \quad (2c)$$

Then, problem (1) can be rewritten as:

$$\min_w \mathcal{J}(w) + g(Bw). \quad (3)$$

Throughout this work, the following assumption holds.

Assumption 1. (Objective Functions) *Problem (3) has a solution w^* , and the function $\mathcal{J} : \mathbb{R}^Q \rightarrow \mathbb{R}$ is δ -smooth and ν -strongly-convex with $0 < \nu \leq \delta$. Moreover, the function $g : \mathbb{R}^E \rightarrow \mathbb{R} \cup \{+\infty\}$ is proper lower semi-continuous and convex, and there exists $w \in \mathbb{R}^Q$ such that Bw belongs to the relative interior domain of g .*

Under Assumption 1, strong duality holds [27, Corollary 31.2.1], and problem (3) is equivalent to the following saddle-point problem [28, Proposition 19.18]:

$$\boxed{\min_w \max_y \mathcal{J}(w) + y^\top Bw - g^*(y)} \quad (4)$$

where y is the dual variable and g^* is the conjugate function of g . A primal-dual pair (w^*, y^*) is optimal if, and only if, it satisfies the optimality conditions [28, Proposition 19.18]:

$$-B^\top y^* = \nabla \mathcal{J}(w^*), \quad (5a)$$

$$Bw^* \in \partial g^*(y^*). \quad (5b)$$

Directly solving (4) does not result in a decentralized algorithm. This is because the dual variable y couples all agents, and these algorithms would require a centralized unit to compute

the dual update. Therefore, further reformulations are needed to derive a decentralized solution.

III. DECENTRALIZED REFORMULATION

In this section, we reformulate (4) into another equivalent saddle-point problem that can be solved in a decentralized manner. Since the dual variable y couples all agents, we introduce local copies of y at all agents. Let y_k denote a local copy of y at agent k . In addition, we introduce the following network quantities:

$$y \triangleq \text{col}\{y_k\}_{k=1}^K \in \mathbb{R}^{EK}, \quad \mathcal{G}^*(y) \triangleq \frac{1}{K} \sum_{k=1}^K g^*(y_k),$$

$$\mathcal{B}_d \triangleq \text{blkdiag}\{B_k\}_{k=1}^K \in \mathbb{R}^{EK \times Q},$$

and the symmetric matrix $\mathcal{L} \in \mathbb{R}^{EK \times EK}$ such that:

$$\mathcal{L}y = 0 \iff y_1 = \dots = y_K. \quad (7)$$

Consider the saddle-point problem:

$$\boxed{\min_{w, \mathcal{X}} \max_y \mathcal{J}(w) + y^\top \mathcal{B}_d w + y^\top \mathcal{L} \mathcal{X} - \mathcal{G}^*(y)} \quad (8)$$

with an optimal solution $(w^*, \mathcal{X}^*, y^*)$ satisfying [28]:

$$-\mathcal{B}_d^\top y^* = \nabla \mathcal{J}(w^*), \quad (9a)$$

$$\mathcal{L}y^* = 0, \quad (9b)$$

$$B_d w^* + \mathcal{L} \mathcal{X}^* \in \partial \mathcal{G}^*(y^*). \quad (9c)$$

Problem (8) can be solved in a decentralized manner because the matrix \mathcal{L} encodes the network sparsity structure and the matrix \mathcal{B}_d is block diagonal. We now show that problems (8) and (4) are equivalent.

Lemma 1. (Saddle-Point) *If $(w^*, \mathcal{X}^*, y^*)$ satisfies the optimality condition (9), then it holds that $y^* = \mathbb{1}_K \otimes y^*$ with (w^*, y^*) satisfying the optimality condition (5).*

Proof. From equations (7) and (9b), we have that $y^* = \mathbb{1}_K \otimes y^*$ for some y^* . Thus, equation (9a) can be written as:

$$-\mathcal{B}_d^\top y^* = -B^\top y^* = \nabla \mathcal{J}(w^*). \quad (10)$$

Multiplying (9c) by $\mathbb{1}_K^\top \otimes I_E$ and using $y^* = \mathbb{1}_K \otimes y^*$ gives:

$$\begin{aligned} (\mathbb{1}_K^\top \otimes I_E) \mathcal{B}_d w^* + (\mathbb{1}_K^\top \otimes I_E) \mathcal{L} \mathcal{X}^* &\in (\mathbb{1}_K^\top \otimes I_E) \partial \mathcal{G}^*(y^*) \\ \implies Bw^* &\in \partial g^*(y^*), \end{aligned} \quad (11)$$

where we used the fact that $(\mathbb{1}_K^\top \otimes I_E) \mathcal{L} = 0$, which holds from (7). From equations (10) and (11), we see that (w^*, y^*) satisfies the optimality condition (5). \square

IV. DECENTRALIZED STRATEGY

In this section, we propose an algorithm to solve (8) and show how to implement it in a decentralized manner.

A. General Algorithm

Let w_{-1}, y_{-1} be any values and $x_{-1} = 0$. The iteration is:

$$w_i = w_{i-1} - \mu_w \nabla \mathcal{J}(w_{i-1}) - \mu_w \mathcal{B}_d^\top y_{i-1}, \quad (12a)$$

$$z_i = y_{i-1} + \mu_y \mathcal{B}_d w_i + \mathcal{L} x_{i-1}, \quad (12b)$$

$$x_i = x_{i-1} - \mathcal{L} z_i, \quad (12c)$$

$$y_i = \text{prox}_{\mu_y \mathcal{G}^*}(\bar{\mathcal{A}} z_i), \quad (12d)$$

where $\bar{\mathcal{A}} = \bar{A} \otimes I_E$ with $\bar{A} \in \mathbb{R}^{K \times K}$ satisfying Assumption 2.

Assumption 2. (Combination Matrices) We assume that \bar{A} is a primitive symmetric doubly-stochastic matrix. Moreover, we assume that the matrix \mathcal{L} satisfies condition (7) and

$$0 < I - \mathcal{L}^2, \quad \bar{A}^2 \leq I - \mathcal{L}^2. \quad (13)$$

The condition on \bar{A} can be easily satisfied for any undirected connected network – see [23]. Note that the eigenvalues of \bar{A} belong to $(-1, 1]$. Given \bar{A} , there are many choices for \mathcal{L} . For instance, we can let $\mathcal{L}^2 = I - \bar{A}^2$ and check whether $I - \mathcal{L}^2$ is positive definite or not. If it is positive definite, then the assumption is satisfied. If not, we can let $\mathcal{L}^2 = c(I - \bar{A}^2)$ for any $c \in (0, 1)$. Although many choices for \mathcal{L} and \bar{A} are possible, due to space considerations, we only focus on one choice in this work.

We first construct a primitive symmetric doubly-stochastic matrix $A = [a_{sk}] \in \mathbb{R}^{K \times K}$ such that $a_{sk} = 0$ if two agents s and k are not connected through an edge. Then, we let $\bar{A} = 0.5(I + A)$, which is also a primitive symmetric doubly-stochastic matrix. In this case, the eigenvalues of \bar{A} belong to $(0, 1]$, and we can let $\mathcal{L}^2 = I - \bar{A}$, which satisfies Assumption 2. We now show how to implement (12) using these choices.

B. Proximal Exact Dual Diffusion (PED²)

From (12b)–(12c), it holds that, for all $i \geq 1$,

$$z_i = (I - \mathcal{L}^2) z_{i-1} + y_{i-1} - y_{i-2} + \mu_y \mathcal{B}_d (w_i - w_{i-1}). \quad (14)$$

It eliminates x_i , and we can rewrite (12) as

$$\begin{aligned} w_i &= w_{i-1} - \mu_w \nabla \mathcal{J}(w_{i-1}) - \mu_w \mathcal{B}_d^\top y_{i-1}, \\ z_i &= (I - \mathcal{L}^2) z_{i-1} + y_{i-1} - y_{i-2} + \mu_y \mathcal{B}_d (w_i - w_{i-1}), \\ \phi_i &= \bar{\mathcal{A}} z_i, \\ y_i &= \text{prox}_{\mu_y \mathcal{G}^*}(\phi_i). \end{aligned}$$

With our choice of $\bar{\mathcal{A}}$ and \mathcal{L} , the k -th block in the vectors can be updated by agent k . Let $w_{k,-1}, y_{k,-1}$ be any values and $\phi_{k,-1} = \psi_{k,-1}$. For each agent k , repeat for $i \geq 0$:

$$w_{k,i} = w_{k,i-1} - \mu_w \nabla J_k(w_{k,i-1}) - \mu_w B_k^\top y_{k,i-1}, \quad (15a)$$

$$\psi_{k,i} = y_{k,i-1} + \mu_y B_k w_{k,i}, \quad (15b)$$

$$z_{k,i} = \phi_{k,i-1} + \psi_{k,i} - \psi_{k,i-1}, \quad (15c)$$

$$\phi_{k,i} = \sum_{s \in \mathcal{N}_k} \bar{a}_{sk} z_{s,i}, \quad (15d)$$

$$y_{k,i} = \text{prox}_{\mu_y / K \mathcal{G}^*}(\phi_{k,i}). \quad (15e)$$

V. LINEAR CONVERGENCE RESULT

In this section, we establish the linear convergence of (12). We first show that the fixed-point of (12) is optimal.

Lemma 2. (Fixed Point) A fixed-point (w^o, x^o, y^o, z^o) of (12) exists, i.e.,

$$0 = \nabla \mathcal{J}(w^o) + \mathcal{B}_d^\top y^o, \quad (16a)$$

$$z^o = y^o + \mu_y \mathcal{B}_d w^o + \mathcal{L} x^o, \quad (16b)$$

$$0 = \mathcal{L} z^o, \quad (16c)$$

$$y^o = \text{prox}_{\mu_y \mathcal{G}^*}(\bar{\mathcal{A}} z^o). \quad (16d)$$

Also, for any fixed-point (w^o, x^o, y^o, z^o) , it holds that $y^o = \mathbb{1}_K \otimes y^o$ with (w^o, y^o) satisfying the optimality condition (5).

Proof. Given an optimal solution (w^*, y^*) of (4) that satisfies (5), we let $w^o \triangleq w^*$ and $y^o \triangleq y^* = \mathbb{1}_K \otimes y^*$. Then, (16a) holds because of (9a). We define

$$z^o \triangleq \mathbb{1}_K \otimes (y^* + \frac{\mu_y}{K} B w^*) \triangleq \mathbb{1}_K \otimes z^o, \quad (17)$$

which satisfies condition (16c). Because of the construction of $\bar{\mathcal{A}}$ and (16c), we have $\bar{\mathcal{A}} z^o = z^o$, and equation (16d) is equivalent to $z^o - y^o \in \mu_y \partial \mathcal{G}^*(y^o)$. Thus, using the definition of \mathcal{G}^* , equation (16d) holds from (5b) and (17). Finally, we construct x^o such that (16b) holds. To see this, note that

$$\begin{aligned} (\mathbb{1}_K^\top \otimes I_E)(z^o - y^o - \mu_y \mathcal{B}_d w^o) \\ = K(z^o - y^o) - \mu_y B w^o \stackrel{(17)}{=} 0 \end{aligned} \quad (18)$$

The above equation implies that $z^o - y^o - \mu_y \mathcal{B}_d w^o$ belongs to the range space of \mathcal{L} (null space of $(\mathbb{1}_K^\top \otimes I_E)$). Thus, there exists x^o such that (16b) holds.

Now, if (w^o, x^o, y^o, z^o) is a fixed-point of (12), then we have $z^o = \mathbb{1}_K \otimes z^o$ because of (16c). Therefore, $\bar{\mathcal{A}} z^o = z^o$, and (16d) shows that $y^o = \mathbb{1}_K \otimes y^o$. Combining (16b) and (16d), we have $\mathcal{B}_d w^o + \frac{1}{\mu_y} \mathcal{L} x^o \in \partial \mathcal{G}^*(y^o)$. Therefore, (w^o, y^o) satisfies (5), which follows from Lemma 1. \square

Note from (12c) that if $x_{-1} = 0$, then $x_1 = -\mathcal{L} z_1$ belongs to the range space of \mathcal{L} . Consequently, $\{x_i\}_{i \geq 0}$ will always remain in the range space of \mathcal{L} . By following similar arguments to [29, Lemma 2], we can always assume that (w^o, x^o, y^o, z^o) is a fixed-point with x^o in the range space of \mathcal{L} because adding a vector in the null space of \mathcal{L} to x^o does not change the optimality condition. To analyze the algorithm (12), we consider the error quantities:

$$\begin{aligned} \tilde{w}_i &= w_i - w^o, & \tilde{y}_i &= y_i - y^o, \\ \tilde{z}_i &= z_i - z^o, & \tilde{x}_i &= x_i - x^o. \end{aligned}$$

From (12) and (16), the error quantities evolve as:

$$\tilde{w}_i = \tilde{w}_{i-1} - \mu_w (\nabla \mathcal{J}(w_{i-1}) - \nabla \mathcal{J}(w^o)) - \mu_w \mathcal{B}_d^\top \tilde{y}_{i-1}, \quad (20a)$$

$$\tilde{z}_i = \tilde{z}_{i-1} + \mu_y \mathcal{B}_d \tilde{w}_i + \mathcal{L} \tilde{x}_{i-1}, \quad (20b)$$

$$\tilde{x}_i = \tilde{x}_{i-1} - \mathcal{L} \tilde{z}_i, \quad (20c)$$

$$\tilde{y}_i = \text{prox}_{\mu_y \mathcal{G}^*}(\bar{\mathcal{A}} z_i) - \text{prox}_{\mu_y \mathcal{G}^*}(\bar{\mathcal{A}} z^o). \quad (20d)$$

To state our main result, we note that condition (13) implies that $0 \leq \underline{\mathcal{L}}^2 < I$. Therefore, $0 < \underline{\sigma}(\mathcal{L}) < 1$, where $\underline{\sigma}(\mathcal{L})$ denotes the smallest *non-zero* singular value of \mathcal{L} . Let $\sigma_{\max}(\mathcal{L})$ denote the largest singular value of \mathcal{L} . The following result establishes the linear convergence of the algorithm (12).

Theorem 1. (Linear Convergence) *Let Assumptions 1 and 2 hold. Assume that each B_k has full row rank. If the step-sizes μ_w and μ_y are strictly positive and satisfy*

$$\mu_w \leq \frac{2}{\delta + \nu}, \quad \mu_y < \frac{2\delta\nu}{(\delta + \nu)\sigma_{\max}^2(\mathcal{B}_d)}, \quad (21)$$

then it holds that $\|\tilde{w}_i\|^2 \leq \gamma^i C_o$ for all $i \geq 0$ and some $C_o \geq 0$ where

$$\gamma = \max \left\{ \frac{1 - \frac{2\mu_w\delta\nu}{\delta + \nu}}{1 - \mu_y\mu_w\sigma_{\max}^2(\mathcal{B}_d)}, 1 - \mu_w\mu_y\lambda_{\min}(\mathcal{B}_d\mathcal{B}_d^\top), 1 - \underline{\sigma}^2(\mathcal{L}) \right\} < 1, \quad (22)$$

with $\lambda_{\min}(\mathcal{B}_d\mathcal{B}_d^\top)$ being the smallest eigenvalue of $\mathcal{B}_d\mathcal{B}_d^\top$.

Proof. See Appendix A. \square

This theorem shows that the proposed algorithm has linear convergence for non-smooth g if B_k has full row rank. Centralized algorithms can achieve linear convergence when $B = [B_1 \ \cdots \ B_K]$ has full row rank. We leave it to future work to verify whether decentralized algorithms can also achieve linear convergence under the same condition.

Remark 1 (SEMI-STRONGLY-CONVEX). Since (4) has the same form as (8), it can be solved with existing algorithms. Thus, one can utilize existing algorithms to derive other decentralized solutions for problem (8). However, existing linear convergence results [7]–[9] require the saddle-point to be strongly-convex with respect to (w.r.t.) the primal-variable. Problem (8) is only strongly-convex w.r.t. the primal block vector w but not strongly-convex w.r.t. to the whole variable $\text{col}\{w, x\}$. Therefore, the linear convergence results from [7]–[9] are not applicable in our setup.

VI. NUMERICAL SIMULATION

In this section, we apply algorithm (15) to solve the following problem:

$$\min_{w_1, \dots, w_K} \frac{1}{2} \sum_{k=1}^K w_k^\top R_k w_k + r_k^\top w_k, \quad \text{s.t.} \quad \sum_{k=1}^K w_k \leq b. \quad (23)$$

This problem fits into (1) with $B_k = I$ and $g(x) = 0$ if $x \leq b$ and ∞ otherwise. We randomly generate positive-definite matrices $R_k \in \mathbb{R}^{10 \times 10}$ and vectors $r_k \in \mathbb{R}^{10}$. The entries of the vector b are uniformly chosen between $(0, 1)$. The combination matrix A is generated using the Metropolis rule [23]. We consider a randomly generated network with $K = 20$ agents shown on the left side of Fig. 1. The simulation result is shown

on the right side of Fig. 1, where the linearized prox-ascent algorithm is:

$$w_i = w_{i-1} - \mu_w \nabla \mathcal{J}(w_{i-1}) - \mu_w B^\top \lambda_{i-1}, \quad (24a)$$

$$\lambda_i = \text{prox}_{\mu_y g^*}(\lambda_{i-1} + \mu_y B w_i). \quad (24b)$$

The plot clearly shows PED² (15) converges linearly in this setup, and it is slightly slower than the centralized algorithm (24).




Fig. 1: The network topology used in the simulation and squared error $\|w_i - w^*\|^2$ evolution of the proposed PED² (15) and the centralized algorithm (24) with $\mu_w = 0.03$ and $\mu_y = 2$.

VII. CONCLUSION

We studied the sharing problem (1), where agents are coupled through a convex possibly non-smooth composite function. To solve (1) in a decentralized manner, we reformulated it into an equivalent saddle-point problem. We then proposed a proximal decentralized algorithm and established its linear convergence. To our knowledge, this is the first *decentralized* linear convergence result for the multi-agent sharing problem (1) with a general non-smooth coupling function.

REFERENCES

- [1] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, “Distributed optimization and statistical learning via alternating direction method of multipliers,” *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2011.
- [2] S. Sundhar Ram, A. Nedić, and V. V. Veeravalli, “A new class of distributed optimization algorithms: Application to regression of distributed data,” *Optimization Methods and Software*, vol. 27, no. 1, pp. 71–88, 2012.
- [3] J. Chen, Z. J. Towfic, and A. H. Sayed, “Dictionary learning over distributed models,” *IEEE Trans. Signal Process.*, vol. 63, no. 4, pp. 1001–1016, 2015.
- [4] D. Hallac, J. Leskovec, and S. Boyd, “Network LASSO: Clustering and optimization in large graphs,” in *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2015, pp. 387–396.
- [5] T.-H. Chang, A. Nedić, and A. Scaglione, “Distributed constrained optimization by consensus-based primal-dual perturbation method,” *IEEE Trans. Autom. Control*, vol. 59, no. 6, pp. 1524–1538, 2014.
- [6] D. P. Palomar and M. Chiang, “A tutorial on decomposition methods for network utility maximization,” *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 8, pp. 1439–1451, 2006.
- [7] P. Chen, J. Huang, and X. Zhang, “A primal-dual fixed point algorithm for convex separable minimization with applications to image restoration,” *Inverse Problems*, vol. 29, no. 2, p. 025011, Jan. 2013.
- [8] N. Dhillon, S. Khong, and M. Jovanovic, “The proximal augmented Lagrangian method for nonsmooth composite optimization,” *IEEE Trans. Autom. Control*, vol. 64, no. 7, pp. 2861–2868, 2019.
- [9] W. Deng and W. Yin, “On the global and linear convergence of the generalized alternating direction method of multipliers,” *Journal of Scientific Computing*, vol. 66, no. 3, pp. 889–916, 2016.
- [10] L. Walras, *Elements Deconomie Politique Pure, ou, Theorie de la Richesse Sociale*. F. Rouge, 1896.

- [11] Y. Ho, L. Servi, and R. Suri, "A class of center-free resource allocation algorithms," in *IFAC Proceedings*, vol. 13, Toulouse, France, 1980, pp. 475–482.
- [12] I. Necoara, "Random coordinate descent algorithms for multi-agent convex optimization over networks," *IEEE Trans. Autom. Control*, vol. 58, no. 8, pp. 2001–2012, 2013.
- [13] T. T. Doan and A. Olshevsky, "Distributed resource allocation on dynamic networks in quadratic time," *Systems & Control Letters*, vol. 99, pp. 57–63, 2017.
- [14] A. Nedić, A. Olshevsky, and W. Shi, "Improved convergence rates for distributed resource allocation," in *IEEE Conference on Decision and Control (CDC)*, Miami Beach, FL, USA, 2018, pp. 172–177.
- [15] J. Xu, S. Zhu, Y. C. Soh, and L. Xie, "A dual splitting approach for distributed resource allocation with regularization," *IEEE Transactions on Control of Network Systems*, vol. 6, no. 1, pp. 403–414, 2018.
- [16] T. Yang, D. Wu, H. Fang, W. Ren, H. Wang, Y. Hong, and K. H. Johansson, "Distributed energy resources coordination over time-varying directed communication networks," *IEEE Transactions on Control of Network Systems*, vol. 6, no. 3, pp. 1124–1134, 2019.
- [17] T.-H. Chang, M. Hong, and X. Wang, "Multi-agent distributed optimization via inexact consensus ADMM," *IEEE Trans. Signal Process.*, vol. 63, no. 2, pp. 482–497, Jan. 2015.
- [18] S. A. Alghunaim, K. Yuan, and A. H. Sayed, "A proximal diffusion strategy for multi-agent optimization with sparse affine constraints," *IEEE Trans. Autom. Control*, 2020, early access (to appear).
- [19] B. Ying, K. Yuan, and A. H. Sayed, "Supervised learning under distributed features," *IEEE Trans. Signal Process.*, vol. 67, no. 4, pp. 977–992, 2019.
- [20] L. He, A. Bian, and M. Jaggi, "COLA: Decentralized linear learning," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2018, pp. 4536–4546.
- [21] J. Tsitsiklis, D. Bertsekas, and M. Athans, "Distributed asynchronous deterministic and stochastic gradient optimization algorithms," *IEEE Trans. Autom. Control*, vol. 31, no. 9, pp. 803–812, 1986.
- [22] A. Nedic and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Trans. Autom. Control*, vol. 54, pp. 48–61, 2009.
- [23] A. H. Sayed, "Adaptation, learning, and optimization over networks," *Foundations and Trends in Machine Learning*, vol. 7, pp. 311–801, 2014.
- [24] Y. Sun, A. Daneshmand, and G. Scutari, "Convergence rate of distributed optimization algorithms based on gradient tracking," *arXiv preprint:1905.02637*, May 2019.
- [25] S. A. Alghunaim, K. Yuan, and A. H. Sayed, "A linearly convergent proximal gradient algorithm for decentralized optimization," in *Advances in Neural Information Processing Systems*, 2019, pp. 2844–2854.
- [26] S. A. Alghunaim, E. K. Ryu, K. Yuan, and A. H. Sayed, "Decentralized proximal gradient algorithms with linear convergence rates," *submitted for publication*, Sept. 2019, available on arXiv:1909.06479.
- [27] R. T. Rockafellar, *Convex Analysis*. Citeseer, 1970.
- [28] H. H. Bauschke and P. L. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, 2011, vol. 408.
- [29] S. A. Alghunaim and A. H. Sayed, "Linear convergence of primal-dual gradient methods and their performance in distributed optimization," *Automatica*, vol. 117, pp. 1–8, July 2020.
- [30] Y. Nesterov, *Introductory Lectures on Convex Optimization: A Basic Course*. Springer, 2013, vol. 87.

APPENDIX A PROOF OF THEOREM 1

The following lemma establishes a useful equality, whose proof is omitted due to space limitations.

Lemma 3. (Equality) *Assume that the step-sizes μ_w and μ_y are strictly positive. The iterates of (12) satisfy:*

$$\begin{aligned} & \|\tilde{w}_i\|_{I-\mu_y\mu_w\mathcal{B}_d^T\mathcal{B}_d}^2 + c_y\|\tilde{z}_i\|_{I-\mathcal{L}^2}^2 + c_y\|\tilde{x}_i\|^2 \\ &= \|\tilde{w}_{i-1} - \mu_w(\nabla\mathcal{J}(w_{i-1}) - \nabla\mathcal{J}(w^o))\|^2 \\ & \quad + c_y\|\tilde{y}_{i-1}\|_{I-\mu_y\mu_w\mathcal{B}_d^T\mathcal{B}_d}^2 + c_y\|\tilde{x}_{i-1}\|_{I-\mathcal{L}^2}^2, \end{aligned} \quad (25)$$

where $c_y = \mu_w/\mu_y$. \square

It can be verified that

$$\begin{aligned} & \|\tilde{w}_{i-1} - \mu_w(\nabla\mathcal{J}(w_{i-1}) - \nabla\mathcal{J}(w^o))\|^2 \\ & \leq (1 - \frac{2\mu_w\delta\nu}{\delta+\nu})\|\tilde{w}_{i-1}\|^2 \\ & \leq \left(\frac{1 - \frac{2\mu_w\delta\nu}{\delta+\nu}}{1 - \mu_y\mu_w\sigma_{\max}^2(\mathcal{B}_d)}\right)\|\tilde{w}_{i-1}\|_{I-\mu_y\mu_w\mathcal{B}_d^T\mathcal{B}_d}^2. \end{aligned} \quad (26)$$

where the first inequality holds under Assumption 1 for $\mu_w \leq \frac{2}{\delta+\nu}$ – see [30]. The second inequality holds if $I - \mu_y\mu_w\mathcal{B}_d^T\mathcal{B}_d$ is positive definite. Let

$$\gamma_1 \triangleq \frac{1 - \frac{2\mu_w\delta\nu}{\delta+\nu}}{1 - \mu_y\mu_w\sigma_{\max}^2(\mathcal{B}_d)}. \quad (27)$$

Then, we have $\gamma_1 < 1$ if $\mu_y < \frac{2\delta\nu}{(\delta+\nu)\sigma_{\max}^2(\mathcal{B}_d)}$. Since the proximal mapping is nonexpansive, it holds from (20d) that:

$$\begin{aligned} \|\tilde{y}_i\|^2 &= \|\text{prox}_{\mu_y\mathcal{G}^*}(\tilde{\mathcal{A}}z_i) - \text{prox}_{\mu_y\mathcal{G}^*}(\tilde{\mathcal{A}}z^o)\|^2 \\ &\leq \|\tilde{\mathcal{A}}\tilde{z}_i\|^2 = \|\tilde{z}_i\|_{\tilde{\mathcal{A}}^2}^2 \leq \|\tilde{z}_i\|_{I-\mathcal{L}^2}^2, \end{aligned} \quad (28)$$

where the last inequality holds due to Assumption 2. Since each B_k has full row rank, it holds that

$$0 < \lambda_{\min}(\mathcal{B}_d\mathcal{B}_d^T)I \leq \mathcal{B}_d\mathcal{B}_d^T \quad (29)$$

where $\lambda_{\min}(\cdot)$ denotes the smallest eigenvalue of its argument. Therefore,

$$\|\tilde{y}_{i-1}\|_{I-\mu_y\mu_w\mathcal{B}_d^T\mathcal{B}_d}^2 \leq (1 - \mu_w\mu_y\lambda_{\min}(\mathcal{B}_d\mathcal{B}_d^T))\|\tilde{y}_{i-1}\|^2. \quad (30)$$

Finally, since $x_{-1} = 0$ and x^o are in the range space of \mathcal{L} , the error quantity \tilde{x}_{i-1} always belongs to the range space of \mathcal{L} . This implies that $\|\tilde{x}_{i-1}\|_{\mathcal{L}^2}^2 \geq \underline{\sigma}^2(\mathcal{L})\|\tilde{x}_{i-1}\|^2$ where $\underline{\sigma}^2(\mathcal{L})$ denotes the minimum non-zero singular value of \mathcal{L} – see [29, Lemma 1]. Therefore,

$$\|\tilde{x}_{i-1}\|_{I-\mathcal{L}^2}^2 \leq (1 - \underline{\sigma}^2(\mathcal{L}))\|\tilde{x}_{i-1}\|^2 \quad (31)$$

Substituting the bounds (26), (28), (30), and (31) into (25) gives:

$$\begin{aligned} & \|\tilde{w}_i\|_{I-\mu_y\mu_w\mathcal{B}_d^T\mathcal{B}_d}^2 + c_y\|\tilde{y}_i\|^2 + c_y\|\tilde{x}_i\|^2 \\ & \leq \gamma_1\|\tilde{w}_{i-1}\|_{I-\mu_y\mu_w\mathcal{B}_d^T\mathcal{B}_d}^2 + \gamma_2c_y\|\tilde{y}_{i-1}\|^2 + \gamma_3c_y\|\tilde{x}_{i-1}\|^2, \end{aligned} \quad (32)$$

where $\gamma_2 \triangleq 1 - \mu_w\mu_y\lambda_{\min}(\mathcal{B}_d\mathcal{B}_d^T)$ and $\gamma_3 \triangleq 1 - \underline{\sigma}^2(\mathcal{L})$. Under the step-size conditions given in (21), it holds that $I - \mu_y\mu_w\mathcal{B}_d^T\mathcal{B}_d > 0$ and $\gamma = \max\{\gamma_1, \gamma_2, \gamma_3\} < 1$. Moreover, we have that $(1 - \mu_y\mu_w\sigma_{\max}^2(\mathcal{B}_d))\|\tilde{w}_i\|^2 \leq \|\tilde{w}_i\|_{I-\mu_y\mu_w\mathcal{B}_d^T\mathcal{B}_d}^2$. Let

$$C_o = \frac{\|\tilde{w}_{-1}\|_{I-\mu_y\mu_w\mathcal{B}_d^T\mathcal{B}_d}^2 + c_y\|\tilde{y}_{-1}\|^2 + c_y\|\tilde{x}_{-1}\|^2}{1 - \mu_y\mu_w\sigma_{\max}^2(\mathcal{B}_d)}. \quad (33)$$

Iterating inequality (32) yields the result.