

A Fast Distributed Proximal-Gradient Method

Annie I. Chen and Asuman Ozdaglar

Abstract—We present a distributed proximal-gradient method for optimizing the average of convex functions, each of which is the private local objective of an agent in a network with time-varying topology. The local objectives have distinct differentiable components, but they share a common nondifferentiable component, which has a favorable structure suitable for effective computation of the proximal operator. In our method, each agent iteratively updates its estimate of the global minimum by optimizing its local objective function, and exchanging estimates with others via communication in the network. Using Nesterov-type acceleration techniques and multiple communication steps per iteration, we show that this method converges at the rate $1/k$ (where k is the number of communication rounds between the agents), which is faster than the convergence rate of the existing distributed methods for solving this problem. The superior convergence rate of our method is also verified by numerical experiments.

I. INTRODUCTION

There has been a growing interest in developing distributed methods that enable the collection, storage, and processing of data using multiple agents connected through a network. Many of these problems can be formulated as

$$\min_{x \in \mathbb{R}^d} f(x) = \frac{1}{m} \sum_{i=1}^m f_i(x), \quad (1)$$

where m is the number of agents in the network, $f(x)$ is the global objective function, and for each $i = 1, \dots, m$, $f_i(x)$ is a local objective function determined by private information available to agent i . The goal is for agents to cooperatively solve problem (1). Most methods for solving this problem involve each agent maintaining an estimate of the global optimum of problem (1) and updating this estimate iteratively using his own private information and information exchanged with neighbors over the network. Examples include a team of sensors exploring an unknown terrain, where $f_i(x)$ may represent regularized least-squares fit to the measurement taken at agent i . As another example, in a distributed machine learning problem, $f_i(x)$ may represent a regularized loss function according to training samples accessible to agent i .

Most optimization algorithms developed for solving problem (1) and its variations are first-order methods (i.e., methods that use gradient or subgradient information of

the objective functions), which are computationally inexpensive and naturally lead to distributed implementations over networks. These methods typically *converge at rate* $1/\sqrt{n}$, where n is the number of communication steps in which agents exchange their estimates; in other words, the difference between the global objective function value at an agent estimate and the optimal value of problem (1) is inversely proportional to the square-root of the number of communication steps carried out (see [1] for a distributed subgradient method and [2] for a distributed dual averaging algorithm with this rate). An exception is the recent independent work [3], which developed a distributed gradient method with a diminishing step size rule, and showed that under certain conditions on the communication network and higher-order differentiability assumptions, the method converges at rate $\log(n)/n$.

In this paper, we focus on a structured version of problem (1) where the local objective function $f_i(x)$ takes the additive form $f_i(x) = g_i(x) + h(x)$ with g_i a differentiable function and h a common nondifferentiable function.¹ We develop a distributed proximal gradient method that solves this problem *at rate* $1/n$ *over a network with time-varying connectivity*. Our method involves each agent maintaining an estimate of the optimal solution to problem (1) and updating it through the following steps: at iteration k , each agent i takes a step along the negative gradient direction of g_i , the differentiable component of his local objective function, and then enters the *consensus stage* to exchange his estimate with his neighbors. The consensus stage consists of k *communication steps*. In each communication step, the agent updates his estimate to a linear combination of his current estimate and the estimates received from neighbors. After the consensus stage, each agent performs a proximal step with respect to h , the nondifferentiable part of his objective function, at his current estimate, followed by a Nesterov-type acceleration step.

This algorithm has two novel features: first, the multi-step consensus stage brings the estimates of the agents close together before performing the proximal step, hence allowing us to reformulate this method as an inexact centralized proximal gradient method with controlled error. Our analysis then uses the recent results on the convergence rate of an inexact (centralized) proximal-gradient method (see [4]) to

This work was supported by National Science Foundation under Career grant DMI-0545910 and AFOSR MURI FA9550-09-1-0538.

The authors are with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA 02139. {anniecia, asuman}@mit.edu

¹This models problems in which h represents a nondifferentiable regularization term. For example, a common choice for machine learning applications is $h(x) = \lambda \|x\|_1$, where $\|x\|_1$ is the sum of the absolute values of each element of the vector x .

establish the convergence rate of the distributed method. Second, exploiting the special structure in the objective functions allows for the use of a proximal-gradient method that can be accelerated using a Nesterov acceleration step, leading to the faster convergence rate of the algorithm.

Other than the papers cited above, our paper is related to the seminal works [5] and [6], which developed distributed methods for solving global optimization problems with a common objective (i.e., $f_i(x) = f(x)$ for all i in problem (1)) using parallel computations in multiple servers. It is also related to the literature on consensus problems and algorithms (see [7]–[10]) and a recent growing literature on multi-agent optimization where information is decentralized among multiple agents connected through a network (see [11]–[15] for subgradient algorithms and [16], [17] for algorithms based on the alternating direction method of multipliers). Finally, our paper builds on the seminal papers on (centralized) proximal-point and proximal-gradient methods (see [18]–[20]).

The paper is organized as follows: Section 2 describes preliminary results pertinent to our work. In Section 3, we introduce our fast distributed proximal-gradient method and establish its convergence rate. Section 4 presents numerical experiments that verify the effectiveness of our method. Finally, Section 5 concludes the paper with open questions for future work.

Notations and Definitions:

- For a vector or scalar that is local, we use subscript(s) to denote the agent(s) it belongs to, and superscripts with parentheses to denote the iteration number; for example, $x_i^{(k)}$ denotes the estimate of agent i at iteration k .
- For a vector or scalar that is common to every agent, or is part of a centralized formulation, the iteration number is also written in superscripts with parentheses; for example, $\bar{x}^{(k)}$ denotes the average estimate of all agents at iteration k . Similarly, $e^{(k)}$ and $\varepsilon^{(k)}$ denote the errors in the centralized formulation at iteration k .
- The standard inner product of two vectors $x, y \in \mathbb{R}^d$ is denoted $\langle x, y \rangle = x'y$. For $x \in \mathbb{R}^d$, its Euclidean norm is $\|x\| = \sqrt{\langle x, x \rangle}$, and its 1-norm is $\|x\|_1 = \sum_{l=1}^d |x(l)|$, where $x(l)$ is its l -th entry.
- For a matrix A , we denote its entry at the i -th row and j -th column as $[A]_{ij}$. We also write $[a_{ij}]$ to represent a matrix A with $[A]_{ij} = a_{ij}$. A matrix A is said to be *stochastic* if the entries in each row sum up to 1, and it is *doubly stochastic* if A and A' are both stochastic.
- We write $a(n) = O(b(n))$ if and only if there exists a positive real number M and a real number n_0 such that $|a(n)| \leq M|b(n)|$ for all $n \geq n_0$.
- For a function $F : \mathbb{R}^d \rightarrow (-\infty, \infty]$, we denote the domain of F by $\text{dom}(F)$, where

$$\text{dom}(F) = \{x \in \mathbb{R}^d \mid F(x) < \infty\}.$$

For a given vector $x \in \text{dom}(F)$, we say that $z_F(x) \in \mathbb{R}^d$ is a *subgradient* of the function F at x when the

following relation holds:

$$F(x) + \langle z_F(x), y - x \rangle \leq F(y) \quad \text{for all } x \in \text{dom}(F).$$

The set of all subgradients of F at x is denoted by $\partial F(x)$.

II. PRELIMINARIES

In this section, we introduce the main concepts and establish key results on which our subsequent analysis relies. Section 2.1 gives properties of the proximal operator; Section 2.2 summarizes convergence rate results for an inexact centralized proximal-gradient method characterized in terms of the errors introduced in the method.

A. Properties of the Proximal Operator

For a closed proper convex function $h : \mathbb{R}^d \rightarrow (-\infty, \infty]$ and a scalar $\alpha > 0$, we define the *proximal operator with respect to h* as

$$\text{prox}_h^\alpha \{x\} = \underset{z \in \mathbb{R}^d}{\text{argmin}} \left\{ h(z) + \frac{1}{2\alpha} \|z - x\|^2 \right\}.$$

It follows that the minimization in the preceding optimization problem is attained at a unique point $y = \text{prox}_h^\alpha \{x\}$, i.e., the proximal operator is a single-valued map [18]. Moreover, using the optimality condition for this problem

$$0 \in \partial h(y) + \frac{1}{\alpha}(y - x),$$

we can see that the proximal operator has the following properties [19]:

Proposition 1: (Basic properties of the proximal operator) Let $h : \mathbb{R}^d \rightarrow (-\infty, \infty]$ be a closed proper convex function. For a scalar $\alpha > 0$ and $x \in \mathbb{R}^d$, let $y = \text{prox}_h^\alpha \{x\}$.

- We have $\frac{1}{\alpha}(x - y) \in \partial h(y)$.
- The vector y can be written as $y = x - \alpha z$, where $z \in \partial h(y)$.
- We have $h(u) \geq h(y) + \frac{1}{\alpha} \langle x - y, u - y \rangle$ for all $u \in \mathbb{R}^d$.
- (Nonexpansiveness) For $x, \hat{x} \in \mathbb{R}^d$, we have

$$\|\text{prox}_h^\alpha \{x\} - \text{prox}_h^\alpha \{\hat{x}\}\| \leq \|x - \hat{x}\|.$$

B. Inexact Proximal-Gradient Method

Our approach for the analysis of the proposed distributed method is to view it as an inexact centralized proximal-gradient method, with the error controlled by multiple communication steps at each iteration. This allows us to use recent results on the convergence rate of an inexact centralized proximal-gradient method to establish the convergence rate of our distributed method. The following proposition from [4] characterizes the convergence rate of an inexact proximal-point method in terms of error sequences $\{e^{(k)}\}_{k=1}^\infty$ and $\{\varepsilon^{(k)}\}_{k=1}^\infty$.

Proposition 2: [4, Proposition 2] Let $g : \mathbb{R}^d \rightarrow \mathbb{R}$ be a convex function that has a Lipschitz continuous gradient with Lipschitz constant L , and let $h : \mathbb{R}^d \rightarrow (-\infty, \infty]$ be a lower semi-continuous proper convex function. Suppose the function $f = g + h$ attains its minimum at a certain $x^* \in \mathbb{R}^d$.

Given two sequences $\{e^{(k)}\}_{k=1}^{\infty}$ and $\{\varepsilon^{(k)}\}_{k=1}^{\infty}$, where $e^{(k)} \in \mathbb{R}^d$ and $\varepsilon \in \mathbb{R}$ for every k , consider the accelerated inexact proximal gradient method, which iterates the following recursion:

$$\begin{cases} x^{(k)} & \in \text{prox}_{h,\varepsilon^{(k)}}^{\alpha} \{y^{(k-1)} - \alpha (\nabla g(y^{(k-1)}) + e^{(k)})\} \\ y^{(k)} & = x^{(k)} + \frac{k-1}{k+2} (x^{(k)} - x^{(k-1)}) \end{cases} \quad (2)$$

where the step size is $\alpha = \frac{1}{L}$, and

$$\begin{aligned} \text{prox}_{h,\varepsilon}^{\alpha} \{y\} &= \{x \in \mathbb{R}^d \mid \\ & h(x) + \frac{1}{2\alpha} \|x - y\|^2 \leq \min_{z \in \mathbb{R}^d} \left(h(z) + \frac{1}{2\alpha} \|z - y\|^2 \right) + \varepsilon \} \end{aligned} \quad (3)$$

indicates the set of all ε -optimal solutions for the proximal operator.

Then, for all $n \geq 1$, we have

$$f(x(n)) - f(x^*) \leq \frac{2L \left(\|x^{(0)} - x^*\| + 2\tilde{A}_n + \sqrt{2\tilde{B}_n} \right)^2}{(n+1)^2}$$

where

$$\tilde{A}_n = \sum_{k=1}^n k \left(\frac{\|e^{(k)}\|}{L} + \sqrt{\frac{2\varepsilon^{(k)}}{L}} \right), \quad \tilde{B}_n = \sum_{k=1}^n \frac{k^2 \varepsilon^{(k)}}{L}.$$

Proposition 2 indicates that as long as the error sequences $\{\|e^{(k)}\|\}_{k=1}^{\infty}$ and $\{\varepsilon^{(k)}\}_{k=1}^{\infty}$ are such that the sequences $\{k\|e^{(k)}\|\}_{k=1}^{\infty}$ and $\{k\sqrt{\varepsilon^{(k)}}\}_{k=1}^{\infty}$ are both summable, then the accelerated inexact gradient method achieves the optimal convergence rate of $O(\frac{1}{n^2})$. It is straightforward to verify using the analysis in [4] that the result also holds for a constant step size $\alpha \leq \frac{1}{L}$.

We shall see that error sequences in our inexact formulation, introduced by the distributed nature of our problem and controlled by multi-step consensus, can be bounded by sequences of the form $\{p^{(k)}\gamma^k\}_{k=1}^{\infty}$ for some polynomial $p^{(k)}$ of k and some $\gamma \in (0, 1)$, which we shall henceforth refer to as *polynomial-geometric sequences*. The next proposition shows that such sequences are summable and allows us to use Proposition 2 in the convergence analysis of our method (the proof is omitted due to limited space).

Proposition 3: (Summability of polynomial-geometric sequences)

Let γ be a positive scalar such that $\gamma < 1$, and let

$P(k, N) = \{c_N k^N + \dots + c_1 k + c_0 \mid c_j \in \mathbb{R}, j = 0, \dots, N\}$ denote the set of all N -th order polynomials of k , where N is a nonnegative integer. Then for every polynomial $p^{(k)} \in P(k, N)$,

$$\sum_{k=0}^{\infty} p^{(k)} \gamma^k < \infty.$$

The result of this proposition for $p(k, N) = k^N$ will be particularly useful for our analysis in the upcoming sections. Therefore, we make the following definition:

$$S_N^{\gamma} := \sum_{k=0}^{\infty} k^N \gamma^k < \infty. \quad (4)$$

III. MODEL AND METHOD

We consider the optimization problem

$$\min_{x \in \mathbb{R}^d} f(x) := \frac{1}{m} \sum_{i=1}^m f_i(x), \quad (5)$$

where $f(x)$ is the global objective function, and $f_i(x) = g_i(x) + h(x)$, $i = 1, \dots, m$ are local objective functions that are private to each agent. For example, for regularized logistic regression, the local objective functions are given by $g_i(x) = \frac{1}{|N_i|} \sum_{j \in N_i} \log(1 + \exp(-b_j \langle a_j, x \rangle))$ and $h(x) = \lambda \|x\|_1$, where N_i is the training dataset of agent i , corresponding to $\{a_j \mid j \in N_i\}$, the set of feature vectors, and $\{b_j \mid j \in N_i\}$, the set of associated labels.

We adopt the following assumption on the functions $g_i(x)$ and $h(x)$.

Assumption 1: (a) For every i , $g_i : \mathbb{R}^d \rightarrow \mathbb{R}$ is convex, continuously differentiable, and has a Lipschitz-continuous gradient with Lipschitz constant $L > 0$, i.e.,

$$\|\nabla g_i(x) - \nabla g_i(y)\| \leq L \|x - y\| \quad \text{for all } x, y \in \mathbb{R}^d.$$

- (b) There exists a scalar G_g such that for every i and for every $x \in \mathbb{R}^d$, $\|\nabla g_i(x)\| < G_g$.
- (c) $h : \mathbb{R}^d \rightarrow \mathbb{R}$ is convex.
- (d) There exists a scalar G_h such that for every $x \in \mathbb{R}^d$, $\|z\| < G_h$ for each subgradient $z \in \partial h(x)$.
- (e) $f(x) = \frac{1}{m} \sum_{i=1}^m f_i(x) = \frac{1}{m} \sum_{i=1}^m g_i(x) + h(x)$ attains its minimum at a certain x^* .

These assumptions are standard in the analysis of distributed first-order methods (see [6], [21] and [1]).

We propose the following *distributed proximal-gradient method* for solving problem (5): Starting from initial estimates $\{y_i^{(0)}\}_{i=1, \dots, m}$ with $y_i^{(0)} \in \mathbb{R}^d$, each agent i updates his estimate $y_i^{(k-1)}$ at iteration k as follows:

$$q_i^{(k)} = y_i^{(k-1)} - \alpha \nabla g_i(y_i^{(k-1)}) \quad (6a)$$

$$\hat{q}_i^{(k)} = \sum_{j=1}^m \lambda_{ij}^{(k)} q_j^{(k)} \quad (6b)$$

$$x_i^{(k)} = \text{prox}_h^{\alpha} \{\hat{q}_i^{(k)}\} \quad (6c)$$

$$y_i^{(k)} = x_i^{(k)} + \frac{k-1}{k+2} (x_i^{(k)} - x_i^{(k-1)}) \quad (6d)$$

Here, $\alpha > 0$ is a constant stepsize which is also the constant scalar used in the proximal operator. The scalars $\lambda_{ij}^{(k)}$ are weights given by

$$\lambda_{ij}^{(k)} = [\Phi(t^{(k)} + k, t^{(k)})]_{ij},$$

for all $i, j = 1, \dots, m$ and all $k \geq s$, where $t^{(k)}$ is the total number of communication steps before iteration k , and Φ is a *transition matrix* representing the product of matrices $A(t)$, i.e., for $t \geq 0$,

$$\Phi(t, s) = A(t)A(t-1) \cdots A(s+1)A(s),$$

where $A(t) = [a_{ij}(t)]_{i,j=1, \dots, m}$ is a matrix of weights $a_{ij}(t) \geq 0$ for $i, j = 1, \dots, m$. Using a vector notation

$\hat{q}^{(k)} = [\hat{q}_i^{(k)}]_{i=1,\dots,m}$ and $q^{(k)} = [q_i^{(k)}]_{i=1,\dots,m}$, we can rewrite (6b) as

$$\begin{aligned}\hat{q}^{(k)} &= \Phi(t^{(k)} + k, t^{(k)})q^{(k)} \\ &= A(t^{(k)} + k)A(t^{(k)} + k - 1) \cdots A(t^{(k)})q^{(k)}.\end{aligned}$$

Hence, this step represents agents performing k communication steps at iteration k . At each communication step, agents exchange their values $q_i^{(k)}$ and update these values by linearly combining the received values using weights $A(t)$. We refer to (6b) as a *multi-step consensus stage*, since linear (in fact, convex, as we shall see) combinations of estimates will serve to bring the agent estimates close to each other.

Our method involves each agent updating his estimate along the negative gradient of the differentiable part of his local objective function (step (6a)), a multi-step consensus stage (step (6b)), and a proximal step with respect to the nondifferentiable part of his local objective function (step (6c)), which is then followed by a Nesterov-type acceleration step (step (6d)). Hence, it is a distributed proximal-gradient method with a multi-step consensus stage inserted before the proximal step. The multi-step consensus stage serves to bring the estimates close to each other before performing the proximal step with respect to the nondifferentiable function h . This enables us to control the error in the reformulation of the method as an inexact centralized proximal-gradient method.

We analyze the convergence behavior of this method under the information exchange model developed in [1], [5], which we summarize in this section. Let $A(t) = [a_{ij}(t)]$ be the weight matrix used in communication step t of the consensus stage. While the weight matrix $A(t)$ may be time-varying, we assume that it satisfies the following conditions for all t .

Assumption 2: (Weight Matrix and Network Conditions)

Consider the weight matrices $A(t) = [a_{ij}(t)]$, $t = 1, 2, \dots$

- (a) (Double stochasticity) For every t , $A(t)$ is doubly stochastic.
- (b) (Significant weights) There exists a scalar $\eta \in (0, 1)$ such that for all i , $a_{ii}(t) \geq \eta$, and for $j \neq i$, either $a_{ij}(t) \geq \eta$, in which case j is said to be a neighbor of i , and receives the estimate of i , at time t ; or $a_{ij}(t) = 0$, in which case j is not a neighbor of i at time t .
- (c) (Connectivity and bounded intercommunication intervals) Let

$$E_t = \{(j, i) \mid j \text{ receives the estimate of } i \text{ at time } t\},$$

$$E_\infty = \{(j, i) \mid j \text{ receives the estimate of } i \text{ for infinitely many } t\}.$$

Then E_∞ is connected. Moreover, there exists an integer $B \geq 1$ such that if $(j, i) \in E_\infty$, then $(j, i) \in E_t \cup E_{t+1} \cup \dots \cup E_{t+B-1}$.

In this assumption, part (a) ensures that each agent's estimate exerts an equal influence on the estimates of others in the network. Part (b) guarantees that in updating his estimate, each agent gives significant weight to his current

estimate and the estimates received from his neighbors. Part (c) states that the overall communication network is capable of exchanging information between any pair of agents in bounded time. An important implication of this assumption is that for $\bar{B} = (m-1)B$, then $E_t \cup E_{t+1} \cup \dots \cup E_{t+\bar{B}-1} = E_\infty$.

The following result from [1] on the limiting behavior of products of weight matrices will be key in establishing the convergence of our algorithm in the subsequent analysis.

Proposition 4: [1, Proposition 1(b)] Let Assumption 2 hold, and for $t \geq s$, let

$$\Phi(t, s) = A(t)A(t-1) \cdots A(s+1)A(s).$$

Then the entries $[\Phi(t, s)]_{ij}$ converges to $\frac{1}{m}$ as $t \rightarrow \infty$ with a geometric rate uniformly with respect to i, j . Specifically, for all $i, j \in \{1, \dots, m\}$ and all t, s with $t \geq s$,

$$\left| [\Phi(t, s)]_{ij} - \frac{1}{m} \right| \leq 2 \frac{1 + \eta^{-\bar{B}}}{1 - \eta^{\bar{B}}} \left(1 - \eta^{\bar{B}}\right)^{\frac{t-s}{\bar{B}}}.$$

For simplicity, we shall denote $\Gamma = 2 \frac{1 + \eta^{-\bar{B}}}{1 - \eta^{\bar{B}}}$, $\gamma = \left(1 - \eta^{\bar{B}}\right)^{\frac{1}{\bar{B}}}$, and restate this theorem as

$$\left| [\Phi(t, s)]_{ij} - \frac{1}{m} \right| \leq \Gamma \gamma^{t-s}. \quad (7)$$

This lemma will ensure that the distance between each agent's estimate and the average estimate decreases geometrically with respect to the number of communication steps taken in the consensus stage. In particular, it gives the following bound on the distance between iterates $\hat{q}_i^{(k)}$, the outcome outcomes of the multi-step consensus stage (6b), and their average, $\bar{q}^{(k)} = \frac{1}{m} \sum_{i=1}^m \hat{q}_i^{(k)}$:

$$\left\| \hat{q}_i^{(k)} - \bar{q}^{(k)} \right\| = \left\| \sum_{j=1}^m \lambda_{ij}^{(k)} q_j^{(k)} - \frac{1}{m} q_j^{(k)} \right\| \quad (8)$$

$$\begin{aligned}&\leq \sum_{j=1}^m \left| \lambda_{ij}^{(k)} - \frac{1}{m} \right| \left\| q_j^{(k)} \right\| \\ &\leq \Gamma \gamma^k \sum_{j=1}^m \left\| q_j^{(k)} \right\|. \quad (9)\end{aligned}$$

A. Formulation as an Inexact Method

We now show that our method can be formulated as an inexact centralized proximal gradient method in the framework of [4]:

Proposition 5: (Distributed proximal-gradient method as an inexact centralized proximal-gradient method) Let $x_i^{(k)}$ and $y_i^{(k)}$ be iterates generated by Algorithm (6). Let $\bar{x}^{(k)} = \frac{1}{m} \sum_{i=1}^m x_i^{(k)}$ and $\bar{y}^{(k)} = \frac{1}{m} \sum_{i=1}^m y_i^{(k)}$ be the average iterates at iteration k . Then Algorithm (6) can be written as

$$\begin{cases} \bar{x}^{(k)} &\in \text{prox}_{h, \varepsilon^{(k)}}^\alpha \left\{ \bar{y}^{(k-1)} - \alpha [\nabla g(\bar{y}^{(k-1)}) + e^{(k)}] \right\} \\ \bar{y}^{(k)} &= \bar{x}^{(k)} + \frac{k-1}{k+2} (\bar{x}^{(k)} - \bar{x}^{(k-1)}) \end{cases} \quad (10)$$

where the error sequences $\{e^{(k)}\}_{k=1}^\infty$ and $\{\varepsilon^{(k)}\}_{k=1}^\infty$ satisfy

$$\|e^{(k)}\| \leq \frac{L}{m} \sum_{i=1}^m \|y_i^{(k-1)} - \bar{y}^{(k-1)}\| \quad (11)$$

$$\begin{aligned} \varepsilon^{(k)} &\leq \frac{2G}{m} \sum_{i=1}^m \|\hat{q}_i^{(k)} - \bar{q}^{(k)}\| \\ &\quad + \frac{1}{2\alpha} \left(\frac{1}{m} \sum_{i=1}^m \|\hat{q}_i^{(k)} - \bar{q}^{(k)}\| \right)^2. \end{aligned} \quad (12)$$

Proof: By taking the average of (6a), we can see that

$$\bar{q}^{(k)} = \bar{y}^{(k-1)} - \alpha(\nabla g(\bar{y}^{(k-1)}) + e^{(k)}),$$

where

$$e^{(k)} = \frac{1}{m} \sum_{i=1}^m [\nabla g_i(y_i^{(k-1)}) - \nabla g_i(\bar{y}^{(k-1)})],$$

and therefore, due to the Lipschitz-continuity of the gradient of $g_i(x)$,

$$\|e^{(k)}\| \leq \frac{L}{m} \sum_{i=1}^m \|y_i^{(k-1)} - \bar{y}^{(k-1)}\|.$$

Let

$$z^{(k)} = \text{prox}_h^\alpha \{\bar{q}^{(k)}\} = \underset{x}{\text{argmin}} \left\{ h(x) + \frac{1}{2\alpha} \|x - \bar{q}^{(k)}\|^2 \right\}$$

denote the result of the exact centralized proximal step. Then $\bar{x}^{(k)} = \frac{1}{m} \sum_{i=1}^m x_i^{(k)} = \frac{1}{m} \sum_{i=1}^m \text{prox}_h^\alpha \{\hat{q}_i^{(k)}\}$, the result of the proximal step in the distributed method, can be seen as an approximation of $z^{(k)}$. We next relate $z^{(k)}$ and $\bar{x}^{(k)}$ by formulating the latter as an inexact proximal step with error $\varepsilon^{(k)}$. A simple algebraic expansion gives

$$\begin{aligned} &h(\bar{x}^{(k)}) + \frac{1}{2\alpha} \|\bar{x}^{(k)} - \bar{q}^{(k)}\|^2 \\ &\leq h(z^{(k)}) + G_h \|\bar{x}^{(k)} - z^{(k)}\| + \frac{1}{2\alpha} \left\{ \|z^{(k)} - \bar{q}^{(k)}\|^2 \right. \\ &\quad \left. + 2 \langle z^{(k)} - \bar{q}^{(k)}, \bar{x}^{(k)} - z^{(k)} \rangle + \|\bar{x}^{(k)} - z^{(k)}\|^2 \right\} \\ &= \min_{z \in \mathbb{R}^d} \left\{ h(z) + \frac{1}{2\alpha} \|z - \bar{q}^{(k)}\|^2 \right\} \\ &\quad + \|\bar{x}^{(k)} - z^{(k)}\| \left(G_h + \frac{1}{\alpha} \|z^{(k)} - \bar{q}^{(k)}\| \right) \\ &\quad + \frac{1}{2\alpha} \|\bar{x}^{(k)} - z^{(k)}\|^2, \end{aligned}$$

where in the inequality, we used the convexity of $h(x)$ and the bound on the subgradient $\partial h(\bar{x}^{(k)})$ to obtain $h(\bar{x}^{(k)}) \leq h(z^{(k)}) + G_h \|\bar{x}^{(k)} - z^{(k)}\|$; and in the equality, we used the fact that by definition, $z^{(k)}$ is the optimizer of $h(x) + \frac{1}{2\alpha} \|x - \bar{q}^{(k)}\|^2$.

With this expression, we can write

$$\bar{x}^{(k)} \in \text{prox}_{h, \varepsilon^{(k)}}^\alpha \{\bar{q}^{(k)}\},$$

where

$$\begin{aligned} \varepsilon^{(k)} &= \|\bar{x}^{(k)} - z^{(k)}\| \left(G_h + \frac{1}{\alpha} \|z^{(k)} - \bar{q}^{(k)}\| \right) \\ &\quad + \frac{1}{2\alpha} \|\bar{x}^{(k)} - z^{(k)}\|^2. \end{aligned}$$

By definition, $z^{(k)} = \text{prox}_h^\alpha \{\bar{q}^{(k)}\}$ also implies $\frac{1}{\alpha} (\bar{q}^{(k)} - z^{(k)}) \in \partial h(z^{(k)})$, and therefore its norm is bounded by G_h . As a result,

$$\varepsilon^{(k)} \leq 2G_h \|\bar{x}^{(k)} - z^{(k)}\| + \frac{1}{2\alpha} \|\bar{x}^{(k)} - z^{(k)}\|^2.$$

Combined with the nonexpansiveness of the proximal operator (Proposition 1(d)),

$$\begin{aligned} \|\bar{x}^{(k)} - z^{(k)}\| &\leq \frac{1}{m} \sum_{i=1}^m \left\| \text{prox}_h^\alpha \{\hat{q}_i^{(k)}\} - \text{prox}_h^\alpha \{\bar{q}^{(k)}\} \right\| \\ &\leq \frac{1}{m} \sum_{i=1}^m \|\hat{q}_i^{(k)} - \bar{q}^{(k)}\|, \end{aligned}$$

we arrive at the desired expression. \blacksquare

This proposition shows that the two error sequences $\|e^{(k)}\|$ and $\varepsilon^{(k)}$ have upper bounds in terms of $\frac{1}{m} \sum_{i=1}^m \|y_i^{(k-1)} - \bar{y}^{(k-1)}\|$ and $\frac{1}{m} \sum_{i=1}^m \|\hat{q}_i^{(k)} - \bar{q}^{(k)}\|$, respectively, which are in turn controlled by the multi-step consensus stage. According to [4, Proposition 2], if $\{k\|e^{(k)}\|\}$ and $\{k\sqrt{\varepsilon^{(k)}}\}$ are both summable, then the inexact proximal-gradient method exhibits the optimal exact convergence rate of $O(1/n^2)$. In the following sections, we shall see that this is indeed the case.

B. Convergence Rate Analysis

We next show that the sequences $\{\|e^{(k)}\|\}_{k=1}^\infty$ and $\{\varepsilon^{(k)}\}_{k=1}^\infty$ are bounded above by polynomial-geometric sequences. By Proposition 3, this establishes that the sequences $\{k\|e^{(k)}\|\}$ and $\{k\sqrt{\varepsilon^{(k)}}\}$ are summable. We first present some useful recursive expressions of the iterates.

Proposition 6: (Recursive expressions of iterates)

Let sequences $\{x_i^{(k)}\}_{k=1}^\infty$, $\{y_i^{(k)}\}_{k=1}^\infty$, $\{q_i^{(k)}\}_{k=1}^\infty$, $\{\hat{q}_i^{(k)}\}_{k=1}^\infty$, $i = 1, \dots, m$, be iterates generated by Algorithm (6). For every $k \geq 2$, we have

- $\sum_{i=1}^m \|q_i^{k+1}\| \leq \sum_{i=1}^m \|q_i^{(k)}\| + \alpha m(G_g + G_h) + \sum_{i=1}^m \|x_i^{(k)} - x_i^{(k-1)}\|$
- $\sum_{i=1}^m \|x_i^{(k)} - x_i^{(k-1)}\| \leq 2m\Gamma \sum_{l=1}^{k-1} \gamma^l \sum_{i=1}^m \|q_i^l\| + (k-1)\alpha m(G_g + G_h)$
- $\|y_i^{(k)} - \bar{y}^{(k)}\| \leq 4\Gamma\gamma^k \sum_{i=1}^m \|q_i^{(k)}\| + 2\Gamma\gamma^{k-1} \sum_{i=1}^m \|q_i^{(k-1)}\|$

The proof is given in the online appendix. These recursive expressions allows us to bound $\sum_{i=1}^m \|q_i^{(k)}\|$ with a second-order polynomial of k , as in the following lemma, whose proof is also given in the appendix.

Lemma 1: (Polynomial bound on $\sum_{j=1}^m \|q_i^{(k)}\|$)

Let sequences $\{x_i^{(k)}\}_{k=1}^\infty$, $\{y_i^{(k)}\}_{k=1}^\infty$, $\{q_i^{(k)}\}_{k=1}^\infty$, $\{\hat{q}_i^{(k)}\}_{k=1}^\infty$, $i = 1, \dots, m$, be generated by Algorithm (6). Then there exists scalars C_q, C'_q, C''_q such that for $k \geq 2$,

$$\sum_{i=1}^m \|q_i^{(k)}\| \leq C_q + C'_q k + C''_q k^2.$$

We now apply Lemma 1 on the error sequences in (10) to show that $\{k\|e^{(k)}\|\}_{k=1}^\infty$ and $\{k\sqrt{\varepsilon^{(k)}}\}_{k=1}^\infty$ are polynomial-geometric sequences, thus summable:

Lemma 2: (Summability of $\{k\|e^{(k)}\|\}_{k=1}^\infty$ and $\{k\sqrt{\varepsilon^{(k)}}\}_{k=1}^\infty$)

In the formulation (10), where

$$\begin{aligned} \|e^{(k)}\| &\leq \frac{L}{m} \sum_{i=1}^m \|y_i^{(k-1)} - \bar{y}^{(k-1)}\|, \\ \varepsilon^{(k)} &\leq \frac{2G_h}{m} \sum_{i=1}^m \|\hat{q}_i^{(k)} - \bar{q}^{(k)}\| \\ &\quad + \frac{1}{2\alpha} \left(\frac{1}{m} \sum_{i=1}^m \|\hat{q}_i^{(k)} - \bar{q}^{(k)}\| \right)^2, \end{aligned}$$

we have

- (a) $\sum_{k=1}^\infty k\|e^{(k)}\| < \infty$
- (b) $\sum_{k=1}^\infty k\sqrt{\varepsilon^{(k)}} < \infty$

Proof: In both cases, it suffices to show that the sequence is a polynomial-geometric sequence. The result then follows by Proposition 3.

- (a) By Proposition 6(c),

$$\begin{aligned} &\frac{1}{m} \sum_{i=1}^m \|y_i^{(k)} - \bar{y}^{(k)}\| \\ &\leq 4\Gamma\gamma^k \sum_{i=1}^m \|q_i^{(k)}\| + 2\Gamma\gamma^{k-1} \sum_{i=1}^m \|q_i^{(k-1)}\| \end{aligned}$$

and by Lemma 1(a), $\sum_{i=1}^m \|q_i^{(k)}\| \leq C_q + C'_q k + C''_q k^2$. Therefore,

$$\begin{aligned} k\|e^{(k)}\| &\leq 4L\Gamma\gamma^k k (C_q + C'_q k + C''_q k^2) \\ &\quad + 2L\Gamma\gamma^{k-1} k (C_q + C'_q(k-1) + C''_q(k-1)^2) \end{aligned}$$

which is a polynomial-geometric sequence.

- (b) Recall (8), $\|\hat{q}_i^{(k)} - \bar{q}^{(k)}\| \leq \Gamma\gamma^k \sum_{j=1}^m \|q_j^{(k)}\|$, and Lemma 1(a), $\sum_{j=1}^m \|q_j^{(k)}\| \leq C_q + C'_q k + C''_q k^2$. Therefore,

$$\begin{aligned} \varepsilon^{(k)} &\leq 2G_h\Gamma\gamma^k (C_q + C'_q k + C''_q k^2) \\ &\quad + \frac{1}{2\alpha} [\Gamma\gamma^k (C_q + C'_q k + C''_q k^2)]^2. \end{aligned}$$

Using the fact that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for all nonnegative real numbers a, b , we have

$$\begin{aligned} \sqrt{\varepsilon^{(k)}} &\leq \sqrt{2G_h\Gamma\gamma^k (C_q + C'_q k + C''_q k^2)} \\ &\quad + \frac{1}{\sqrt{2\alpha}} [\Gamma\gamma^k (C_q + C'_q k + C''_q k^2)] \\ &\leq \sqrt{2G_h\Gamma}\sqrt{\gamma^k} (\sqrt{C_q} + \sqrt{C'_q k} + \sqrt{C''_q k^2}) \\ &\quad + \frac{\Gamma}{\sqrt{2\alpha}} \gamma^k (C_q + C'_q k + C''_q k^2) \end{aligned}$$

where in the last line we used the fact that $\sqrt{k} \leq k$ for all $k \geq 1$. This is a polynomial-geometric sequence.

Therefore, $k\sqrt{\varepsilon^{(k)}}$ is also a polynomial-geometric sequence. ■

Using the lemma above, we can establish the convergence rate of our distributed proximal-gradient method:

Theorem 1: (Convergence rate of the distributed proximal-gradient method with multi-step consensus) Let $\{x_i^{(k)}\}_{k=1}^\infty$ be iterates generated by Algorithm (6), with a constant step size $\alpha \leq 1/L$ where L is the Lipschitz constant in Assumption 1. Let $\bar{x}^{(k)} = \frac{1}{m} \sum_{i=1}^m x_i^{(k)}$ be the average iterate at iteration k . Then, for all $t \geq 1$, where t is the total number of communication steps taken, we have

$$f(\bar{x}^{(t)}) - f(x^*) = O(1/t).$$

Proof: Since it takes k communication steps to complete iteration k , the total number of communication steps required to execute iterations $1, \dots, n$ is $\sum_{k=1}^n k = \frac{n(n+1)}{2}$.

In other words, after t communication steps, the number of iterations completed is n , where n is the greatest integer such that $\frac{n(n+1)}{2} = \frac{n^2+n}{2} \leq t$, or equivalently, $n = \left\lceil \frac{-1+\sqrt{1+8t}}{2} \right\rceil$. As a result,

$$(n+1)^2 \geq \left(\frac{-1+\sqrt{1+8t}}{2} \right)^2 = \frac{2+8t-2\sqrt{1+8t}}{4},$$

and thus,

$$f(\bar{x}^{(t)}) - f(x^*) \leq \frac{D}{(n+1)^2} \leq \frac{2D}{4t+1-\sqrt{1+8t}} = O(1/t).$$

Although this theorem is stated in terms of $\bar{x}^{(t)}$, it could also lead to a bound on $f(x_i^{(t)}) - f(x^*)$, using the gradient bound, nonexpansiveness of the proximal operator, and (9). Also, note that the results above hold for the fast distributed gradient method (where the objective functions are differentiable), which is clear by simply setting $h(x) = 0, G_h = 0$ and $\varepsilon^{(k)} = 0$. ■

C. Beyond $O(1/t)$

We have thus shown that taking k communication steps in the k -th iteration of Algorithm (6) results in the summability of error sequences $\{k\|e^{(k)}\|\}_{k=1}^\infty$ and $\{k\sqrt{\varepsilon^{(k)}}\}_{k=1}^\infty$. A natural question arises: can we do better? In particular, will the error sequences still converge if we took less than k communication steps in the k -th iteration? We address this question in this section.

Let s_k be the number of communication steps taken in the multi-step consensus stage at iteration k . In our method presented earlier, $s_k = k$. We wish to explore smaller choices of s_k that still preserves the guarantee for exact convergence.

With s_k , we see that (8) can be written as

$$\|\hat{q}_i^{(k)} - \bar{q}^{(k)}\| \leq \Gamma\gamma^{s_k} \sum_{j=1}^m \|q_j^{(k)}\|.$$

Therefore, Proposition 6(b) becomes

$$\begin{aligned} & \sum_{i=1}^m \left\| x_i^{(k)} - x_i^{(k-1)} \right\| \\ & \leq \sum_{l=1}^{k-1} 2m\Gamma\gamma^{s_l} \sum_{j=1}^m \left\| q_j^l \right\| + (k-1)\alpha m(G_g + G_h) \end{aligned}$$

As a result, if we have the equivalent of Proposition 3 for s_k , i.e., if $\sum_{k=0}^{\infty} k^N \gamma^{s_k} < \infty$ for any given $\gamma \in (0, 1)$ and nonnegative integer N , then Lemma 1 would hold, and so would Theorem 1.

Since $\sum_{k=0}^{\infty} k^a < \infty$ for $a < -1$, a sufficient condition for the above is $\gamma^{s_k} < k^{-N-1}$, or equivalently, $s_k > \frac{-N-1}{\log \gamma} \log k$. This is at the order of $O(\log k)$, which is smaller than our previous choice of $s_k = k = O(k)$. The hidden constant, $\frac{-N-1}{\log \gamma}$, depends on N and γ . In our case, we only require this condition to hold up to $N = 3$. Therefore, if γ is known, by choosing $s_k = \left\lceil \frac{4}{-\log \gamma} \log(k+1) \right\rceil$, the distributed proximal-gradient method is guaranteed to converge with rate $O(1/n^2)$, where n is the iteration number.

The time it takes to complete iterations $1, \dots, n$, which we denote by $T(n)$, is then

$$T(n) = \sum_{k=1}^n s_k = O(n \log n - n)$$

since $\int \log x dx = x(\log x - 1)$. Unfortunately, $n \log n = \log n \cdot e^{\log n}$ has no explicit inverse expression [22]. Therefore, we can only express the convergence rate as

$$f(\bar{x}^{(t)}) - f(x^*) = O(1/(T^{-1}(t))^2)$$

which we know to be better than $O(1/t)$, since $T(n)$ is bounded above by $O(n^2)$.

In closing, we remark that the improved choice of s_k above requires the knowledge of γ , which may not be readily available if detailed information or performance guarantees of the communication network is unknown. In such cases, the method could still be implemented with $s_k = k$.

IV. NUMERICAL EXPERIMENTS

Our theoretical findings are verified with numerical experiments on a machine learning task using 20 Newsgroups [23], [24], a benchmark dataset for text categorization. It consists of about 20,000 news articles, evenly chosen among 20 topics. The task is to perform L_1 -regularized logistic regression on the training data, so as to learn the classification model for a chosen topic label. Specifically, we wish to minimize

$$f(x) = \frac{1}{N} \sum_{j=1}^N \log(1 + \exp(-b_j \langle a_j, x \rangle)) + \lambda \|x\|_1,$$

where N is the total number of news articles, a_j is the 8615-dimensional feature vector of article j , and b_j is its the label for the chosen topic, which is equal to 1 if this article belongs to the topic, and -1 otherwise. x contains parameters of the classification model that we wish to learn, and $f(x)$ is its corresponding regularized loss function.

We distribute the training data across a network of $m = 10$ data centers, each with 1129 samples. Thus, each data center has the following private objective function:

$$f_i(x_i) = \frac{1}{|N_i|} \sum_{j \in N_i} \log(1 + \exp(-b_j \langle a_j, x_i \rangle)) + \lambda \|x_i\|_1$$

where N_i is the subset of data at center i , and $x_i, i = 1, \dots, m$, is its local estimate of the global classification model. In each communication step, a weight matrix is randomly chosen from a pool of 10 weight matrices generated from connected random graphs. All weight matrices satisfy Assumption 2.

To demonstrate the effect of using multiple communication steps after the gradient step in our method, we compare it with the following methods:

- The *basic* subgradient method with *single-step* consensus in [1]:

$$\begin{cases} x_i^{(k)} &= w_i^{(k-1)} - \alpha (\nabla g_i(w_i^{(k-1)}) + z_h(w_i^{(k-1)})) \\ w_i^{(k)} &= \sum_{j=1}^m a_{ij}^{(k)} x_j^{(k)} \end{cases} \quad (13)$$

where $z_h(w_i^{(k-1)}) \in \partial h(w_i^{(k-1)})$ and $[a_{ij}^{(k)}]$ is the randomly-chosen weight matrix.

- The *basic* proximal-gradient method with *single-step* consensus, similar to that of [1]:

$$\begin{cases} x_i^{(k)} &= \text{prox}_h^\alpha \{w_i^{(k-1)} - \alpha \nabla g_i(w_i^{(k-1)})\} \\ w_i^{(k)} &= \sum_{j=1}^m a_{ij}^{(k)} x_j^{(k)} \end{cases} \quad (14)$$

where $[a_{ij}^{(k)}]$ is the randomly-chosen weight matrix.

- The *accelerated* proximal-gradient method with *single-step* consensus:

$$\begin{cases} x_i^{(k)} &= \text{prox}_h^\alpha \{w_i^{(k-1)} - \alpha \nabla g_i(w_i^{(k-1)})\} \\ y_i^{(k)} &= x_i^{(k)} + \frac{k-1}{k+2} (x_i^{(k)} - x_i^{(k-1)}) \\ w_i^{(k)} &= \sum_{j=1}^m a_{ij}^{(k)} y_j^{(k)} \end{cases} \quad (15)$$

- The *accelerated* proximal-gradient method with *multi-step* consensus which is not inserted between the gradient and proximal steps, but instead performed only after the proximal step:

$$\begin{cases} x_i^{(k)} &= \text{prox}_h^\alpha \{w_i^{(k-1)} - \alpha \nabla g_i(w_i^{(k-1)})\} \\ y_i^{(k)} &= x_i^{(k)} + \frac{k-1}{k+2} (x_i^{(k)} - x_i^{(k-1)}) \\ w_i^{(k)} &= \sum_{j=1}^m \lambda_{ij}^{(k)} y_j^{(k)} \end{cases} \quad (16)$$

where $[\lambda_{ij}^{(k)}]$ is the product of k weight matrices randomly drawn from the pool of 10 weight matrices.

Figure 1 shows convergence rate results for each method. It is clear from the figure that Algorithm (13) converges to an error neighborhood at rate $O(1/t)$, as shown in [1]. Algorithms (14) and (15) also converge to an error neighborhood, but the latter exhibits more oscillation than the basic methods. Algorithm (16) converges with rate $O(1/t)$, but only to an error neighborhood instead of achieving exact convergence. This highlights the importance of having the consensus step before the proximal step instead of after it. Finally, our accelerated multi-step method attains exact convergence with rate $O(1/t)$, outperforming all others.




Fig. 1: Performance comparison for various distributed methods

V. CONCLUSION AND FUTURE WORK

We presented a distributed proximal-gradient method that solves for the optimum of the average of convex functions, each having a distinct differentiable component and a common nondifferentiable component. The method uses multiple communication steps and Nesterov’s acceleration technique. We established the convergence rate of this method as $O(1/t)$ (where t is the total number of communication steps), superior to most existing distributed methods.

Several questions remain open for future work. First, it would be useful to generalize the result for the case where the nondifferentiable functions $h_i(x)$ are distinct. Secondly, it is of interest to determine the condition under which the accelerated single-step proximal-gradient method (15) converges, and compare its performance with our multi-step consensus method. Last but not least, it would be useful to obtain a lower bound on the convergence rate of distributed first-order methods under our current framework.

REFERENCES

- [1] A. Nedic and A. Ozdaglar, “Distributed subgradient methods for multi-agent optimization,” in *LIDS report 2755, IEEE Transactions on Automatic Control*, vol. 54, no. 1, 2009, pp. 48–61.
- [2] J. Duchi, A. Agarwal, and M. Wainwright, “Dual averaging for distributed optimization: Convergence and network scaling,” *IEEE Transactions on Automatic Control*, 2012.
- [3] D. Jakovetic, J. Xavier, and J. M. F. Moura, “Fast distributed gradient methods,” *arXiv:1112.2972v1*, 2011.
- [4] M. Schmidt, N. L. Roux, and F. Bach, “Convergence rates of inexact proximal-gradient methods for convex optimization,” *CoRR*, vol. abs/1109.2415, 2011.
- [5] J. N. Tsitsiklis, “Problems in decentralized decision making and computation,” Ph.D. dissertation, Department of EECS, MIT, 1984.
- [6] J. N. Tsitsiklis, D. P. Bertsekas, and M. Athans, “Distributed asynchronous deterministic and stochastic gradient optimization algorithms,” *IEEE Transactions on Automatic Control*, vol. 31, no. 9, pp. 803–812, 1986.
- [7] A. Jadbabaie, J. Lin, and S. Morse, “Coordination of groups of mobile autonomous agents using nearest neighbor rules,” *IEEE Transactions on Automatic Control*, vol. 48, no. 6, pp. 988–1001, 2003.
- [8] R. Olfati-Saber and R. Murray, “Consensus problems in networks of agents with switching topology and time-delays,” *IEEE Transactions on Automatic Control*, vol. 49, no. 9, pp. 1520–1533, 2004.
- [9] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, “Gossip algorithms: Design, analysis, and applications,” *Proceedings of IEEE INFOCOM*, 2005.
- [10] A. Olshevsky and J. Tsitsiklis, “Convergence rates in distributed consensus and averaging,” *Proceedings of the 45th IEEE Conference on Decision and Control*, 2006.
- [11] B. Johansson, T. Keviczky, M. Johansson, and K. Johansson, “Subgradient methods and consensus algorithms for solving convex optimization problems,” *Proceedings of the 47th IEEE Conference on Decision and Control*, p. 41854190, 2008.
- [12] S. S. Ram, A. Nedic, and V. V. Veeravalli, “Asynchronous gossip algorithms for stochastic optimization,” *Proceedings of the 48th IEEE Conference on Decision and Control*, pp. 3581–3586, 2009.
- [13] A. Nedic, A. Ozdaglar, and P. A. Parrilo, “Constrained consensus and optimization in multi-agent networks,” *IEEE Transactions on Automatic Control*, vol. 55, no. 4, 2010.
- [14] I. Lobel and A. Ozdaglar, “Distributed subgradient methods for convex optimization over random networks,” *IEEE Transactions on Automatic Control*, vol. 56, no. 6, pp. 1291–1306, 2011.
- [15] M. Zhu and S. Martínez, “On distributed constrained formation control in operator-vehicle adversarial networks,” *Automatica*, submitted, 2012.
- [16] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, “Distributed optimization and statistical learning via the alternating direction method of multipliers,” *Foundations and Trends in Machine Learning*, vol. 3, no. 1, 2011.
- [17] E. Wei and A. Ozdaglar, “Distributed alternating direction method of multipliers.”
- [18] R. T. Rockafellar, “Monotone operators and the proximal point algorithm,” *SIAM Journal on Control and Optimization*, vol. 14, no. 5, pp. 877–898, 1976.
- [19] A. Beck and M. Teboulle, “A fast iterative shrinkage-thresholding algorithm for linear inverse problems,” *SIAM Journal on Imaging Sciences*, vol. 2, pp. 183–202, March 2009.
- [20] —, “Gradient-based algorithms with applications in signal recovery problems,” in *Convex Optimization in Signal Processing and Communications*, D. P. Palomar and Y. C. Eldar, Eds. Cambridge University Press, 2010.
- [21] V. D. Blondel, J. M. Hendrickx, A. Olshevsky, and J. N. Tsitsiklis, “Convergence in multiagent coordination, consensus, and flocking,” *Proceedings of the Joint 44th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC’05)*, 2005.
- [22] R. M. Corless, G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth, “On the lambert w function,” in *Advances in Computational Mathematics*, 1996, pp. 329–359.
- [23] K. Lang, “Newsweeder: Learning to filter netnews,” in *Proceedings of the Twelfth International Conference on Machine Learning*, 1995, pp. 331–339.
- [24] J. Rennie, “20 newsgroups,” <http://people.csail.mit.edu/jrennie/20Newsgroups/>.