

Distributed gradient methods for convex machine learning problems in networks

Angelia Nedić

Abstract

In this paper we overview distributed gradient methods for solving convex machine learning problems of the form $\min_{x \in \mathbb{R}^n} \frac{1}{m} \sum_{i=1}^m f_i(x)$ in a system consisting of m agents that are embedded in a communication network. Each agent i has a collection of data captured by its privately known objective function $f_i(x)$. The distributed algorithms considered here obey two simple rules: privately known agent functions $f_i(x)$ cannot be disclosed to any other agent in the network, and every agent is aware of the local connectivity structure of the network, i.e., it knows its one-hop neighbors only. While obeying these two rules, the distributed algorithms that agents execute should find a solution to the overall system problem with the limited knowledge of the objective function and limited local communications. We overview such algorithms that typically involve two update steps: a gradient step based on the agent local objective function and a mixing step that essentially diffuses relevant information from one to all other agents in the network.

I. INTRODUCTION

Artificial intelligence is emerging as the driving technology that will enable future automated transportation, smart cities and smart power grids, as well as robots that will replace humans in hazardous workplaces and situations. At the core of this technology are the devices that can collect data (i.e., collect measurements from their surrounding), process data and communicate in order to share estimates for optimal and reliable system performance. Machine learning is at the core of the data processing, while the communication among the devices allows for collaborative learning from all the data within the system without the need to share the data locally stored at the devices. As such, these systems are inherently distributed, as the data collection is distributed, central for the development of these technologies are distributed computational models and coordination mechanisms relying on local communications.

The School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, AZ 85281 USA (e-mail: Angelia.Nedich@asu.edu). The work has been partially funded by the NSF grant CCF-1717391 and the Navy award N000141612245.

One possible architecture that can support operations in such distributed systems is the peer-to-peer architecture, which can also serve as a primitive for building more complex hierarchical architectures. In a peer-to-peer architecture, the consensus protocol has attracted a lot of interest and research for synchronization, formation control, learning and coordination of multi-robot tasks, to name a few. A common ground for addressing these problems is provided by distributed computational models (including machine learning in particular) that use a consensus mechanism as a virtual coordinator.

The goal of the paper is to overview the framework of distributed consensus-based gradient methods for convex machine learning problems over static graphs, both directed and undirected. These methods will be discussed in relation to the standard machine learning problem of minimizing the average of the sum of functions,

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

where each function $f_i(x)$ represents the loss associated with a given data point. Specifically, given a collection $\{(z_i, y_i), i = 1, \dots, p\}$ of data points, where $z_i \in \mathbb{R}^n$ is a feature vector and $y_i \in \{-1, +1\}$ is its corresponding label, a nonlinear classification problem consists of solving the following minimization problem:

$$\min_{x \in \mathbb{R}^n} \left(c\rho(x) + \frac{1}{p} \sum_{i=1}^p \ell(x; z_i, y_i) \right), \quad (2)$$

where $c > 0$ is a regularization parameter, $x = (x_1, \dots, x_n)$ is a vector, and the regularizing function $\rho(x)$ is a strongly convex function (such as the Euclidean norm), and $\ell(x; z_i, y_i)$ is a loss function associated with the data point $(z_i, y_i) \in \mathbb{R}^{n+1}$. The loss function $\ell(x; z, y)$ is typically convex but not necessarily differentiable. For linear classifiers, a common choice is the *logistic regression* loss function, given by:

$$\ell(x; z, y) = \log \left(1 + e^{-y \langle x, z \rangle} \right),$$

with $\langle x, z \rangle$ being the inner product of the vectors x and z , i.e., $\langle x, z \rangle = \sum_{j=1}^n x_j z_j$. The resulting problem, is known as the *logistic regression* (linear) classifier problem. Another common choice is the *hinge-loss function*

$$\ell(x; z, y) = \max\{0, 1 - y \langle x, z \rangle\},$$

giving rise to the maximum margin (linear) classifier problem.

As seen from the preceding discussion, the machine learning problem (2) falls under the general formulation in (1), where a function $f_i(x) = c\rho(x) + \ell(x; z_i, y_i)$ is associated with a data point (z_i, y_i) . When the number of data points p is not too large, the problem in (1) can be solved using the iterative

gradient method, as follows: We start with some initial guess x^0 . At iteration k , we have iterate x^k and we compute the gradient

$$\frac{1}{p} \sum_{i=1}^p \nabla f_i(x^k) \quad (3)$$

of the objective function at the current iterate. Then, we obtain the new iterate x^{k+1} by moving away from the point x^k along the opposite direction of the gradient with some positive stepsize α_k :

$$x^{k+1} = x^k - \frac{\alpha_k}{p} \sum_{i=1}^p \nabla f_i(x^k).$$

Convergence of this simple gradient method is known for the (positive) stepsize sequence $\{\alpha_k\}$ satisfying the following two conditions:

$$\sum_{k=0}^{\infty} \alpha_k = \infty \quad \text{and} \quad \sum_{k=0}^{\infty} \alpha_k^2 < \infty.$$

These conditions are satisfied for example with the stepsize $\alpha_k = \frac{1}{k}$, or more generally with the stepsize of the form $\alpha_k = \frac{a}{k^b}$, where $a > 0$ is arbitrary while $b \in (\frac{1}{2}, 1]$. When the objective function is convex, the gradient method with such a diminishing stepsize produces an iterate sequence $\{x^k\}$ converging to an optimal point x^* that solves the problem in (1), provided that an optimal point exists (e.g. see [1]). We note that, under some additional assumptions on the objective function, the gradient method with a suitably chosen constant stepsize $\alpha > 0$ (i.e., $\alpha_k = \alpha$ for all k) produces a sequence of iterates $\{x^k\}$ converging to the solution of the problem (see for example textbook [2]).

When the number p of the data points is large, computing the full gradient in (3) becomes expensive (time consuming) and alternative approaches have been developed, including the stochastic gradient method that uses $\nabla f_{i_k}(x^k)$ (with an index $i_k \in \{1, 2, \dots, p\}$ selected uniformly at random) instead of the full gradient. To improve the convergence properties of this basic stochastic gradient method, variance reduced methods have been proposed including SVRG [3] and SAGA [4].

The gradient method and its aforementioned stochastic variants are centralized in the sense that they are applicable to the situations where all the data is located at a single location or the data can be accessed from a central location. In what follows, we will consider the machine learning problem in the case when the data is spatially distributed among several locations and there is no central entity that can access all the data. Moreover, we assume that the data is private and cannot be shared among the data centers - a situation that arises in applications dealing with a medical data for example. Each data center is viewed as an agent capable of processing its data points locally and exchanging some estimates (extracted from its private data) with the other agents (data centers) over a communication graph. The agents communicate in order to coordinate their local computations with the other agents in the system, so as to jointly solve their aggregate problem (to be specified shortly).

The outline of the rest of the paper is as follows: Section II will provide a basic gradient algorithm with weighted-averaging consensus step for minimizing the average-sum objective function for the case when the computing agents communicate over an undirected communication graph. The basic idea of the algorithm will be explained as two-level approximation of a centralized gradient method, with insights on how and when these approximations are good enough for the method to work. The core references here are [5], [6], [7], [8], [9], [10]. Section III will be an analog of Section 2 for the directed graphs, where one resorts to a push-sum consensus instead of a weighted-averaging consensus. Key references include [11], [12], [13], [14], [15], [16].

II. DISTRIBUTED GRADIENT METHOD FOR UNDIRECTED GRAPHS

We consider the case when the data points pertinent to the same phenomena or a task are collected at m distinct centers that need not be colocated. We refer to these centers as agents, labeled by $1, 2, \dots, m$, and we let agent i have the set S_i of data points. The learning problem corresponding to the aggregate data of the system of all agents has the form as in (2), where p is the total number of data points at all agents, i.e., $p = |S_1| + \dots + |S_m|$, with $|S|$ denoting the cardinality of a finite set S . To account for the agent specific data, by aggregating the data per agent, the problem in (2) can be re-formulated, as follows:

$$\min_{x \in \mathbb{R}^n} \frac{1}{m} \sum_{i=1}^m f_i(x), \quad (4)$$

where

$$f_i(x) = \frac{m}{p} \sum_{s \in S_i} (c \rho(x) + \ell(x; z_s, y_s)). \quad (5)$$

Note that problem (4)–(5) is equivalent to the problem in (2). We will assume that the function $f_i : \mathbb{R}^n \rightarrow \mathbb{R}$ is the private loss function known only to agent i . Further, we will assume that the m agents are communicating over an undirected network represented by a graph $G = ([m], \mathcal{E})$, where $[m] = \{1, 2, \dots, m\}$ denotes the set of agents and \mathcal{E} denotes the set of undirected edges. We write $\{i, j\}$ to represent an edge connecting agents i and j .

The agents want to solve the problem in (4) collaboratively in the sense that they are willing to share some estimates with their immediate neighbors in the graph, but they are not willing (or not allowed) to share their data, which essentially means they do not reveal their loss functions f_i .

Given the communication graph $G = ([m], \mathcal{E})$, the problem can be reformulated as follows:

$$\begin{aligned} & \min_{x_i \in \mathbb{R}^n, i \in [m]} \frac{1}{m} \sum_{i=1}^m f_i(x_i) \\ & \text{subject to} \quad x_i = x_j \quad \text{for all } \{i, j\} \in \mathcal{E}. \end{aligned} \quad (6)$$

This reformulation is obtained from problem (4) by assigning a copy x_i of the decision variable x to each agent, and by imposing the requirement that these copies should be all the same, i.e., $x_i = x$ for all agents $i \in [m]$. This system of equations is then replaced with an equivalent system of equations requiring that $x_i = x_j$ for all $i, j \in [m]$. When the graph G is connected¹, the system of pairwise equations $x_i = x_j$ for all $i, j \in [m]$ is equivalent to the system of graph compatible equations $x_i = x_j$ for all $\{i, j\} \in \mathcal{E}$. Thus, when the graph G is connected, the problems in (4)–(5) and (6) are equivalent. Reformulation (6) serves as a departure point for the discussion of distributed first-order methods for machine learning problems which use a consensus-type update to enforce the equal-decision constraints in problem (6).

The objective function of problem (6) is decoupled, as each f_i depends on its own variable x_i , but these variables are coupled through edge-based constraints. The idea is to distribute the problem among the agents, by allowing each agent to know its neighbors in the graph, that is every agent i is aware of the agents j such that $\{i, j\} \in \mathcal{E}$, which constitute the set \mathcal{N}_i of the neighbors of i in the graph, i.e., $\mathcal{N}_i = \{j \mid \{i, j\} \in \mathcal{E}\}$. Using the local agent knowledge of the graph and the functions, every agent i can solve its own ‘local part’ of the overall problem. However, in order for the agents to collectively solve the overall problem, each agent needs to align its variables with the variables x_j corresponding to its neighbors $j \in \mathcal{N}_i$.

As a mechanism for alignment of the variables, a consensus algorithm is used. The consensus algorithm is a distributed method that the agents can use to asymptotically agree on a decision vector. Specifically, each agent starts with an arbitrary vector $x_i(0)$. At iteration k , every agent sends its current iterate x_i^k to its neighbors $j \in \mathcal{N}_i$ and receives x_j^k from its neighbors $j \in \mathcal{N}_i$. After this, every agent i , executes the consensus update step:

$$x_i^{k+1} = a_{ii}x_i^k + \sum_{j \in \mathcal{N}_i} a_{ij}x_j^k, \quad (7)$$

where $a_{ii} > 0$ and $a_{ij} > 0$ are such that $a_{ii} + \sum_{j \in \mathcal{N}_i} a_{ij} = 1$. The positive scalars $a_{ij}, j \in \mathcal{N}_i \cup \{i\}$ are referred to as convex weights, and the vector x_i^{k+1} is said to be a convex combination (or a weighted average) of the points $x_j, j \in \mathcal{N}_i \cup \{i\}$. Note that the positive weights $a_{ij}, j \in \mathcal{N}_i \cup \{i\}$, are selected by agent i .

To have a more compact representation of the consensus method, let us define an $m \times m$ weight matrix A with entries

$$a_{ii} > 0 \text{ and } a_{ij} > 0 \text{ with } a_{ii} + \sum_{j \in \mathcal{N}_i} a_{ij} = 1 \quad \text{and} \quad a_{ij} = 0 \text{ when } j \notin \mathcal{N}_i \cup \{i\}. \quad (8)$$

¹For any two distinct agents i and j , there is a path from i to j in the graph G .

In view of the fact that the nonnegative matrix A has a positive entry in the ij -th position only when $\{i, j\}$ is a link in the graph G , we will say that such a matrix A is *compatible with the structure of the graph G* . Using the weights a_{ij} defined in (8), the consensus algorithm in (7) assumes the following form: for every $i \in [m]$,

$$x_i^{k+1} = \sum_{j=1}^m a_{ij} x_j^k. \quad (9)$$

The sum of the entries in each row of the matrix A is equal to 1. Such a non-negative matrix is referred to as *row-stochastic*.

An illustration of a connected star-graph with four agents is given in Figure 1. As an example of

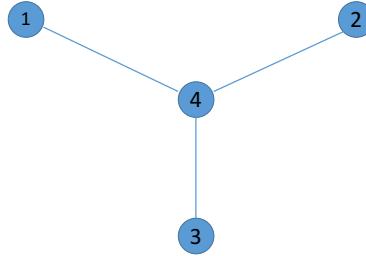


Fig. 1. A star-graph with 4 agents.

a row-stochastic matrix A that is compatible with the structure of the graph in Figure 1 consider the following:

$$A = \begin{bmatrix} \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \end{bmatrix}.$$

This matrix is constructed using an equal-weight rule, whereby each agent i gives the same weight to itself and all of its neighbors $j \in \mathcal{N}_i$.

When the graph $G = ([m], \mathcal{E})$ is connected, and the matrix A is row-stochastic and compatible with the graph structure (i.e., it satisfies the relations in (8)), the iterate sequences $\{x_i^k\}$, $i \in [m]$, generated by the consensus algorithm converge to a same limit point \tilde{x} . The limit point \tilde{x} , referred to as a *consensus or an agreement point*, is given as a convex combination of the initial values $\{x_i^0, i \in [m]\}$, i.e.,

$$\tilde{x} = \sum_{i=1}^m \pi_i x_i^0.$$

The weight vector $\pi = [\pi_1, \dots, \pi_m]$ is the unique (normalized) left-eigenvector of the matrix A corresponding to the eigenvalue $\lambda = 1$, i.e., $\pi A = \pi$. The convergence result is obtained by viewing the matrix

A as a one-step transition matrix of a Markov chain and employing the ergodicity theory for Markov chains. The detailed proof of the convergence of the consensus algorithm can be found in [5], where the consensus method was originally proposed. If the matrix A is doubly stochastic, then $\pi_i = \frac{1}{m}$ for all i and, consequently, the consensus point is the average of the initial values, i.e., $\tilde{x} = \frac{1}{m} \sum_{i=1}^m x_i^0$.

Using the consensus method as a mechanism for alignment of the agent variables, a distributed optimization method with a local information exchange can be constructed by executing two steps - a mixing (consensus) step (8) and a gradient-based step. At the beginning of a typical iteration k , every agent sends to its neighbors its current iterate x_i^k to its neighbors $j \in \mathcal{N}_i$ and receives x_j^k from its neighbors $j \in \mathcal{N}_i$. Then, every agent i , executes the following two update-steps:

$$\begin{aligned} v_i^k &= \sum_{j=1}^m a_{ij} x_j^k, \\ x_i^{k+1} &= v_i^k - \alpha_k \nabla f_i(v_i^k), \end{aligned} \quad (10)$$

where the weights a_{ij} satisfy the same relations as in the consensus method (i.e., the relations in (8)), and $\alpha_k > 0$ is a stepsize. We note that the vector v_i^k is a convex combination of the points $x_j, j \in \mathcal{N}_i \cup \{i\}$. After obtaining v_i^k through the mixing step, every agent i performs a gradient update step using v_i^k to obtain its new iterate x_i^{k+1} . The algorithm in (10) is *distributed* since every agent updates using a gradient of its own private function, and is local in the sense that it relies on local information exchange. The method is also often referred to as the *consensus-based* gradient method, due to its use of a mixing step that resembles the distributed consensus process.

Distributed algorithm (10) can be viewed as an extension of the gradient method, where the mixing step is introduced to align the agents' iterates. This mixing step plays a role of "virtual coordinator" of the agents iterates in a system that does not have a central coordinator. To shed more light into this, lets take a look at what is happening with the iterates of the distributed gradient method on the system level. Consider the iterate sequences $\{x_i^k\}, i \in [m]$, produced by the method at all agents in the system. By taking the average of these iterates (cf. (10)) across all the agents, at any given instance, we have

$$\frac{1}{m} \sum_{i=1}^m x_i^{k+1} = \frac{1}{m} \sum_{i=1}^m v_i^k - \frac{\alpha_k}{m} \sum_{i=1}^m \nabla f_i(v_i^k) = \frac{1}{m} \sum_{j=1}^m \left(\sum_{i=1}^m a_{ij} \right) x_j^k - \frac{\alpha_k}{m} \sum_{i=1}^m \nabla f_i(v_i^k),$$

where the second equality is obtained by using the definition of v_i^k in (10) and by exchanging the order of the resulting two sums. When the matrix A is doubly stochastic², we have $\sum_{i=1}^m a_{ij} = 1$ for all j , thus yielding

$$\frac{1}{m} \sum_{i=1}^m x_i^{k+1} = \frac{1}{m} \sum_{j=1}^m x_j^k - \frac{\alpha_k}{m} \sum_{i=1}^m \nabla f_i(v_i^k).$$

²The sum of its entries in each row and each column is equal to 1.

By letting \bar{x}^k denote the iterate average across the agents at time k , i.e.,

$$\bar{x}^k = \frac{1}{m} \sum_{i=1}^m x_i^k,$$

the preceding relation can be written as

$$\bar{x}^{k+1} = \bar{x}^k - \frac{\alpha_k}{m} \sum_{i=1}^m \nabla f_i(v_i^k). \quad (11)$$

Relation (11) is nearly the same as the (centralized) gradient method update for solving problem (6), with only difference being that the gradient of f_i is computed at the point v_i^k instead of the point \bar{x}^k . More concretely, by adding and subtracting the correct gradients, for the averaged iterates we have

$$\bar{x}^{k+1} = \bar{x}^k - \frac{\alpha_k}{m} \sum_{i=1}^m \nabla f_i(\bar{x}^k) + \epsilon^k, \quad \epsilon^k = \frac{\alpha_k}{m} \sum_{i=1}^m (\nabla f_i(\bar{x}^k) - \nabla f_i(v_i^k)).$$

The averaged iterates \bar{x}^{k+1} will follow an erroneous gradient method, where the error is proportional to the average gradient difference $\nabla f_i(\bar{x}^k) - \nabla f_i(v_i^k)$, as long as we can ensure that the error ϵ_k remains small. In fact, under some additional assumptions on the functions and the stepsize, the mixing step can ensure that the error ϵ_k does not blow up.

The role of the mixing matrix A is critical since it guarantees that the following properties hold under suitable conditions on the stepsize and the functions f_i (such as the convexity):

- (1) the averaged iterate sequence $\{\bar{x}^k\}$ converges to a solution x^* of problem (6) and
- (2) the disagreement sequence $\{\|x_i^k - \bar{x}^k\|\}$ converges to zero for every agent i .

Thus, the iterates $x_i^k, i \in [m]$, generated by the distributed method track the averaged process with increasing accuracy, while the average process converges to a solution of the problem. Consequently, the iterate sequences $\{x_i^k\}, i \in [m]$, all converge to the same point and this common point is a solution of the problem (6).

From the view point of the consensus process in (9), the distributed gradient method in (10) can be interpreted as a consensus-based process that is affected by two forces: the consensus force represented by the mixing step (influenced by the matrix A) and the agent-based gradient forces coming from the objective functions f_i . In this case, the gradients steer the limiting consensus point to a solution x^* of problem (6).

The consensus algorithm was originally proposed in [5] to model opinion dynamic leading to an agreement in a team of individuals. It was later used for estimation, control and optimization starting with works in [17], [18], [6]. In contrast to the machine learning problem considered in this paper, the optimization problem considered in [18], [6] involves a non-separable objective, i.e., the problem of the form

$$\min_{x=(x_1, \dots, x_m) \in \mathbb{R}^n} f(x_1, \dots, x_m),$$

where the decision vector $x \in \mathbb{R}^n$ is partitioned among m agents. The first variants of distributed gradient methods that employ a mixing (consensus) step to solve the problem with a separable structure (i.e., in the form of a machine learning problem (4)) was pioneered in [7], [8]. Subsequently, many distributed methods have been developed such as those with a faster convergence [19], [20] and with emphasis on the communication and the computation tradeoffs [21], [22]. A recent survey [10] gives a thorough review of the consensus problem, and variations and extensions of the distributed method that account for the implementations over time-varying graphs, over networks with delays and with noisy (or quantized) communication links, the method with noisy (stochastic) gradient evaluations, as well as some asynchronous implementations.

The mixing step of the distributed method in (10) is also referred to as a diffusion, as it allows for the local agent information to diffuse over the entire network. The mixing step and the gradient update step in (10) can be exchanged, leading to an alternative variant of the distributed method

$$\begin{aligned} v_i^k &= \frac{x_i^k}{m} - \alpha_k \nabla f_i(x_i^k), \\ x_i^{k+1} &= \sum_{j=1}^m a_{ij} v_j^k. \end{aligned} \quad (12)$$

Viewing the gradient update step as an *adaptation*, the distributed method in (10) is also referred to as *combine-then-adapt* (CTA) diffusion strategy, while the method in (12) is referred to as *adapt-then-combine* (ATC) diffusion strategy. One can find an elegant exposure of these strategies in [9], and follow the references therein for technical details of their analysis.

The convergence rate of the distributed method in (10) is of the order of $O(\frac{\log k}{\sqrt{k}})$ in the number k of iterations, which is dictated by the use of the diminishing stepsize satisfying the typical conditions in gradient methods, namely $\sum_{k=0}^{\infty} \alpha_k = \infty$ and $\sum_{k=0}^{\infty} \alpha_k^2 < \infty$. A centralized gradient method for a convex minimization problem that uses such a diminishing stepsize α_k converges in the order of $O(\frac{1}{\sqrt{k}})$. Thus, the distributed method is slower due to the $\log k$ term, which in turn is due to the extra mixing step in (10) that copes with the distributed knowledge of the overall objective function. Moreover, the hidden constant factor in $O(\frac{\log k}{\sqrt{k}})$ depends critically on the graph G connectivity structure and the spectral properties of the mixing matrix A . A more in depth discussion on how the graph topology affects the performance of the method can be found in a recent survey [23].

We conclude this section by summarizing the basic advantages and limitations of the distributed method (10). The method is considered to be slow since its convergence rate is in the order of $O(\frac{\log k}{\sqrt{k}})$ (resulting from the use of a diminishing stepsize). However, the method is suitable for the situations where noise or errors are present due to noisy links or failure-prone links, or due to the use of stochastic gradients. In such situations, the diminishing stepsize is beneficial as it attenuates the noise.

Two major limitations of the method (10) are:

- (1) It requires the use of a doubly stochastic matrix A in order to solve the problem of minimizing the average sum $\frac{1}{m} \sum_{i=1}^m f_i(x)$ of the agent objectives. The construction of such a matrix in a distributed fashion is a nontrivial problem when the underlying communication graph G is directed, as pointed out in [24]. To avoid this obstacle, one can consider a distributed gradient method that employs an alternative consensus process in the mixing step, which does not require a doubly stochastic matrix, as discussed in Section III.
- (2) In general, the method does not converge with a constant stepsize (i.e., $\alpha_k = \alpha$ for all k), which has been observed in [25]. Therefore, it can not match the performance of the centralized gradient methods that use such stepsizes and have a convergence rate faster than $O(\frac{1}{\sqrt{k}})$. A remedy for this drawback is provided by allowing the agents to exchange and mix both the decision vectors and some suitably defined estimates of the gradient directions, as seen in Section IV.

III. DISTRIBUTED GRADIENT METHOD FOR DIRECTED GRAPHS

When the underlying communication graph G is directed, the distributed method of Section II requiring a doubly stochastic matrix A cannot be easily implemented since the construction of the matrix A in a distributed way becomes computationally expensive. To deal with directed graphs, one can resort to a different consensus algorithm that makes use of a column stochastic weight matrix A , and combine it with a gradient method.

Specifically, the consensus algorithm is built using a column stochastic matrix A that is compatible with the given directed graph G , i.e., for all $j \in [m]$,

$$a_{jj} > 0 \text{ and } a_{ij} > 0 \text{ with } a_{jj} + \sum_{i \in \mathcal{N}_j^{\text{out}}} a_{ij} = 1 \quad \text{and} \quad a_{ij} = 0 \text{ when } i \notin \mathcal{N}_j^{\text{out}} \cup \{j\}, \quad (13)$$

where $\mathcal{N}_j^{\text{out}}$ is the set of all out-neighbors of agent j , i.e., the agents i that can be reached from j by a link $(j, i) \in \mathcal{E}$; formally $\mathcal{N}_j^{\text{out}} = \{i \mid (j, i) \in \mathcal{E}\}$. In this case, agent j chooses the weights in the j th column of A and ensures that its entries sum to 1. A common choice is to let all the values a_{ij} , $i \in \mathcal{N}_j^{\text{out}} \cup \{j\}$, be same and equal to the cardinality of the set $\mathcal{N}_j^{\text{out}} \cup \{j\}$.

Having such a matrix A , a consensus method is constructed that has two variables. In particular, each agent i maintains variables x_i^k and y_i^k at time k . At time $k+1$, each agent j sends $a_{ij}x_j^k$ and $a_{ij}y_j^k$ to all of its out-neighbors $i \in \mathcal{N}_j^{\text{out}}$. After that every agent i updates by simply summing the x - and y -variables it has received from its in-neighbors, i.e.,

$$x_i^{k+1} = \sum_{j \in \mathcal{N}_i^{\text{in}} \cup \{i\}} a_{ij}x_j^k,$$

$$y_i^{k+1} = \sum_{j \in \mathcal{N}_i^{\text{in}} \cup \{i\}} a_{ij} y_j^k, \quad (14)$$

where $\mathcal{N}_i^{\text{in}}$ is the set of in-neighbors of agent i , i.e.,

$$\mathcal{N}_i^{\text{in}} = \{j \mid (j, i) \in \mathcal{E}\}.$$

The algorithm in (14) is known as the *push-sum* method since every agent *pushes* some information along its outgoing links, while it *sums* the incoming information. The algorithm will lead to consensus, as long as the graph G is strongly connected³ and the y -variables are initiated with $y_i^0 = 1$ for all i . The x -variables can be vectors, while y -variables are always scalars. The variable that reaches a consensus is in fact a ratio of x - and y -variables at each agent, that is the variables

$$z_i^k = \frac{x_i^k}{y_i^k} \quad \text{for all } i \in [m],$$

converge to a common consensual point as the number k of the iterations increases to infinity. For this reason, the method is also referred to as a *ratio consensus*. The first push-sum algorithm has been proposed in [11] for a randomized gossip-type implementation. Subsequently, it has been further investigated in [26], [12], [27], [28] for networks with time-varying communication structure including the case of unreliable (packet-dropping) links.

One can use the push-sum consensus to construct a distributed gradient method for solving machine learning problem (6), where each function f_i is a private local objective function of agent i . Starting with arbitrary variables $x_i^0, i \in [m]$ and $y_i^0 = 1$ for all i , the algorithm proceeds as follows: each agent j sends $a_{ji}x_j^k$ and $a_{ji}y_j^k$ to its out-neighbors (i.e., the agents i that can be reached by a link $(j, i) \in \mathcal{E}$). After this communication step, all agents update their x - and y -variables as follows:

$$\begin{aligned} v_i^{k+1} &= \sum_{j \in \mathcal{N}_i^{\text{in}} \cup \{i\}} a_{ij} x_j^k, \\ y_i^{k+1} &= \sum_{j \in \mathcal{N}_i^{\text{in}} \cup \{i\}} a_{ij} y_j^k, \\ z_i^{k+1} &= \frac{v_i^{k+1}}{y_i^{k+1}}, \\ x_i^{k+1} &= v_i^{k+1} - \alpha_k \nabla f(z_i^{k+1}). \end{aligned} \quad (15)$$

³For every two distinct nodes i and j there exist directed paths connecting i to j and j to i .

The method will be referred to as *gradient-push method*. To gain the intuition behind this algorithm, consider dividing the update relation for x_i^{k+1} by y_i^{k+1} , which in view of the definition of z_i^{k+1} in (15) yields

$$\frac{x_i^{k+1}}{y_i^{k+1}} = z_i^{k+1} - \gamma_i^k \nabla f(z_i^{k+1}),$$

where $\gamma_i^k = \frac{\alpha_k}{y_i^{k+1}} > 0$. As the z -variables converge to a common point \tilde{x} and the stepsize α_k is diminishing, under suitable conditions on the gradients of the function f_i (such as boundedness), the ratio $\frac{x_i^{k+1}}{y_i^{k+1}}$ also converges to the point \tilde{x} . Additionally, the convexity of the functions f_i and the correct selection of the stepsize α_k are important to ensure that the consensual point \tilde{x} is in fact a solution to the aggregate-agent learning problem of minimizing the average sum $\frac{1}{m} \sum_{i=1}^m f_i(x)$ of their local objective functions (i.e., the problem in (4)). While the intuition behind the gradient-push method in (15) is simple, its formal analysis is quite involved.

The work in [13], [14] is the first to provide a distributed method that makes use of the push-sum consensus protocol to cope with directed graphs. There, the dual-averaging method of Nesterov has been used in the construction and its convergence properties have been established for a static directed graph. Subsequently, in [15] a simple (sub)gradient method using the push-sum consensus has been proposed for a general sequence of time-varying directed graphs, and later on extended to a stochastic method in [16]. The work in [15], [16] establishes the convergence rate of $O(\frac{\log k}{\sqrt{k}})$ and $O(\frac{\log k}{k})$ for the case when the functions f_i are convex and strongly convex, respectively. We note that, while the work in [15], [16] uses the special weight matrix A_k (equal-out-neighbor weights), it is straightforward to extend it to the case of an arbitrary sequence of column stochastic matrices that are compliant with the topology of the underlying graphs, and have nonzero entries that are uniformly lower-bounded away from zero. A more comprehensive review of the literature and the results for the push-pull consensus based methods is given in a recent survey paper [23].

The gradient-push method in (15) suffers from the same slow convergence as that of the distributed gradient method (10) due to their use of the diminishing stepsize. The advantages of the gradient-push method are the same as those of the distributed gradient method, mainly in noisy settings. If the stepsize is constant, both distributed gradient method (10) and gradient-push method (15) may fail to converge, while the constant step is desirable for a faster convergence. Thus, these algorithms are not readily rendering accelerated variants that can perform nearly as fast as the centralized gradient method.

Another explanation for the inability of these methods to be faster can be observed from their somewhat greedy (local) view of the objective function. In particular, in both of the distributed methods (cf. (10) and (15)) the agents use mixing (with their neighbors) for the decision variables, but the decision variables

are updated at every agent i using only a local objective function gradient ∇f_i . Thus, with respective to the global objective function, which is proportional to $\sum_{i=1}^m f_i(x)$, the agents perform a greedy update. To enhance the performance of these methods, the agents need to be less greedy, as we will see in the next section.

IV. GRADIENT-TRACKING ALGORITHMS

In this section, we assume that the agents are "aware" that they comprise of a system, while they need not be aware of the global connectivity structure nor the global objective function $\sum_{i=1}^m f_i(x)$. However, by being aware that there is a system objective, in addition to mixing their decision variables, they also use and mix directions that serve as estimates of the global objective gradient. This extra mixing step of directions needs to be properly constructed to track the gradients of the system objective.

Suppose that the graph G is undirected, and the mixing matrix is compatible with the graph (see (8)), then ideally, if the agents had access to all the objective functions, each agent would implement the distributed method in (10) by using the gradient of the system objective i.e., every agent i would execute the following two steps:

$$\begin{aligned} v_i^k &= \sum_{j=1}^m a_{ij} x_j^k, \\ x_i^{k+1} &= v_i^k - \frac{\alpha}{m} \sum_{i=1}^m \nabla f_i(v_i^k). \end{aligned}$$

The stepsize here is fixed since we are interested in a faster gradient methods that do not use diminishing step. In the absence of the knowledge of all functions f_i , every agent can track the aggregate function gradients, by employing a direction d_i^k , and tracking the average of these directions. In particular, every agent will update as follows: at time $k+1$, the agents have x_i^k and d_i^k , and exchange these vectors with their neighbors in the graph. Then, they will compute x_i^{k+1} and d_i^{k+1} , as follows:

$$\begin{aligned} x_i^{k+1} &= \sum_{j=1}^m a_{ij} x_j^k - \alpha d_i^k, \\ d_i^{k+1} &= \sum_{j=1}^m a_{ij} d_j^k + \nabla f_i(x_i^{k+1}) - \nabla f_i(x_i^k). \end{aligned} \tag{16}$$

We note that each agent still computes the gradients of its own local objective function f_i , but through the mixing of the directions in (16), it tracks the sum of the gradients of the overall system function objective. Subtracting the past gradient $\nabla f_i(x_i^k)$ in (16) is used to ensure that only new gradient information, captured by the difference $\nabla f_i(x_i^{k+1}) - \nabla f_i(x_i^k)$ influences the direction update.

To ensure the correct gradient-tracking, the method has to be initialized with

$$d_i^0 = \nabla f_i(x_i^0) \quad \text{for all agents } i \in [m],$$

while x_i^0 can be arbitrary. With such an initialization, when the matrix A is column stochastic, the directions d_i^{k+1} satisfy⁴

$$\sum_{i=1}^m d_i^{k+1} = \sum_{i=1}^m \nabla f_i(x_i^{k+1}).$$

Thus, the averaged iterates $\bar{x}^k = \frac{1}{m} \sum_{i=1}^m x_i^k$, satisfy

$$\bar{x}^{k+1} = \bar{x}^k - \frac{\alpha}{m} \sum_{i=1}^m \nabla f_i(x_i^k),$$

which can be interpreted as generated by an erroneous ‘‘centralized gradient’’ method

$$\bar{x}^{k+1} = \bar{x}^k - \frac{\alpha}{m} \sum_{i=1}^m \nabla f_i(\bar{x}^k) + \mathbf{e}_k, \quad \text{with error} \quad \mathbf{e}_k = \frac{\alpha}{m} \sum_{i=1}^m \left(\nabla f_i(\bar{x}^k) - \nabla f_i(x_i^k) \right).$$

These errors are well behaved even for the constant step due to an additional alignment of the directions in (16). In particular, work in [29] studies the method for smooth convex functions with Lipschitz continuous gradients⁵, and shows the convergence of the method in the order of $O(\frac{1}{k})$ when a constant stepsize α is suitably chosen.

Under additional requirement that each f_i is strongly convex⁶, in [30] it is shown that the method (with a carefully selected stepsize α) produces iterates x_i^k that converge to the point x^* minimizing the global objective function $\frac{1}{m} \sum_{i=1}^m f_i(x)$ with a geometric rate, i.e.,

$$\|x_i^k - x^*\| \leq q^k C \quad \text{for all agents } i \text{ and all } k \geq 0,$$

where $q \in (0, 1)$ and C are some constants that depend on the graph connectivity topology, the entries of the matrix A , the stepsize α , and the parameters characterizing the function f_i (such as the strong convexity and Lipschitz gradient constants). Thus, the distributed method with the gradient tracking mechanism matches the fastest convergence rate of the centralized gradient method (which is geometric).

When the graph G is directed, the gradient-push method can also be augmented to include a gradient tracking mechanism. In particular, the gradient-push method in (15) assumes the form:

$$\begin{aligned} v_i^{k+1} &= \sum_{j \in \mathcal{N}_i^{\text{in}} \cup \{i\}} (a_{ij} v_j^k - \alpha d_j^k), \\ y_i^{k+1} &= \sum_{j \in \mathcal{N}_i^{\text{in}} \cup \{i\}} a_{ij} y_j^k, \\ x_i^{k+1} &= \frac{v_i^{k+1}}{y_i^{k+1}}, \end{aligned}$$

⁴It can be shown by using the mathematical induction on k .

⁵A function f has Lipschitz continuous gradients with a constant L if $\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|$ for all $x, y \in \mathbb{R}^n$.

⁶A function f is strongly convex with a constant $\mu > 0$ if $\langle \nabla f(x) - \nabla f(y), x - y \rangle \geq \mu\|x - y\|^2$ for all $x, y \in \mathbb{R}^n$.

$$d_i^{k+1} = \sum_{j \in \mathcal{N}_i^{\text{in}} \cup \{i\}} a_{ij} d_j^k + \nabla f(x_i^{k+1}) - \nabla f(x_i^k). \quad (17)$$

For a column stochastic matrix A and under the conditions of strong convexity and Lipschitz continuous gradients of the functions f_i , similar to that of the distributed gradient-tracking method in (16), this algorithm also enjoys a geometric convergence rate. The convergence rate analysis of these distributed gradient-tracking methods have been shown in [30] for a more general case where the underlying communication graphs and the matrix sequence are time-varying.

Simultaneously and independently of [30], [29], the idea of tracking the gradients have been employed in [31], where a class of distributed algorithms have been developed for solving non-convex problems in networks. In summary, the distributed methods with gradient tracking in (16) and (17) can match the best known convergence rates of centralized gradient methods. The main distinction between the two methods is that the method in (16) works on undirected graphs, where the weight matrix A is required to be doubly stochastic, while the method in (17) works on directed graphs, where the weight matrix A is required to be column stochastic.

We note that each of the two tracking methods employs the same weight matrix A when mixing the x -variables and the direction related d -variables. We next consider an interesting gradient-tracking method that makes use of different matrices to mix the x - and d -variables. In particular, the method works for a directed strongly connected graph G . At iteration k , every agent i has two vectors x_i^k and d_i^k . At the time of iteration $k+1$, each i agent obtains (pulls) $x_j^k - \alpha_j d_j^k$ from its in-neighbors $j \in \mathcal{N}_i^{\text{in}}$, and sends (pushes) $b_{\ell i} d_i^k$ to each of its out-neighbors $\ell \in \mathcal{N}_i^{\text{out}}$. Upon this information exchange, the agents update as follows:

$$\begin{aligned} x_i^{k+1} &= \sum_{j \in \mathcal{N}_i^{\text{in}} \cup \{i\}} a_{ij} (x_j^k - \alpha_j d_j^k) \quad (\text{A-pull}), \\ d_i^{k+1} &= \sum_{\ell \in \mathcal{N}_i^{\text{out}} \cup \{i\}} b_{\ell i} d_i^k + \nabla f(x_i^{k+1}) - \nabla f(x_i^k) \quad (\text{B-push}). \end{aligned} \quad (18)$$

Note that in this method every agent uses its own stepsize α_i which need not be coordinated with the stepsizes of the other agents. For the algorithm to work properly, the initial directions are set to $d_i^0 = \nabla f_i(x_i^0)$ for all i , while $x_i^0 \in \mathbb{R}^n, i \in [m]$, are arbitrary initial decision vectors.

Agent i chooses the positive weights $a_{ij}, j \in \mathcal{N}_i^{\text{in}} \cup \{i\}$, for its in-neighbors involved in the pull-step, and also the positive weights $b_{\ell i}, \ell \in \mathcal{N}_i^{\text{out}} \cup \{i\}$, for its out-neighbors involved in the push-step. By defining these weights to be zero for all other pairs (i, j) , we obtain $m \times m$ matrices compatible with the structure of graph G ; specifically,

$$a_{ij} > 0 \text{ for all } j \in \mathcal{N}_i^{\text{in}} \cup \{i\}, \text{ and } a_{ij} = 0 \text{ otherwise,}$$

$$b_{\ell i} > 0 \text{ for all } \ell \in \mathcal{N}_i^{\text{out}} \cup \{i\}, \text{ and } b_{\ell i} = 0 \text{ otherwise.}$$

The interesting aspect of the *Apull-Bpush* method is that the matrix $A = [a_{ij}]$ is row stochastic, while the matrix $B = [b_{ij}]$ is column stochastic i.e.,

$$\sum_{j=1}^m a_{ij} = 1, \quad \sum_{\ell=1}^m b_{\ell i} = 1 \quad \text{for all } i \in [m].$$

Here, we have coined the name of the algorithm to be *Apull-Bpush* to fairly capture an independent development of two closely related methods, namely a *push-pull* method (cf. (18)) as given in [32] and a variant proposed and analyzed in [33], which performs slightly different updates. In particular, the algorithm in [33] is given by:

$$\begin{aligned} x_i^{k+1} &= \sum_{j \in \mathcal{N}_i^{\text{in}} \cup \{i\}} a_{ij} x_j^k - \eta d_i^k, \\ d_i^{k+1} &= \sum_{\ell \in \mathcal{N}_i^{\text{out}} \cup \{i\}} b_{\ell i} \left(d_j^k + \nabla f_j(x_j^{k+1}) - \nabla f_j(x_j^k) \right). \end{aligned} \quad (19)$$

Both algorithms use a row stochastic matrix A in the updates of the x -variables and a column stochastic matrix $B = [b_{ij}]$ in the updates of the directions. Also, with a suitably selected stepsizes both algorithms converge linearly to the solution of the problem

$$\min_{x \in \mathbb{R}^n} \frac{1}{m} \sum_{i=1}^m f_i(x)$$

when all the functions f_i are strongly convex and have Lipschitz continuous gradients (see [33], [32] for details).

In terms of communication requirements and the size of the variables exchanged among the agents, the distributed method for undirected graphs (Section II), requires $2|\mathcal{E}|$ communications per iteration since each link in the graph is used twice in every update time; the x -variables exchanged are of size n . The distributed method for directed graphs (Section III), requires $|\mathcal{E}|$ communications per iteration since each link in the graph is used once in every update time; the variables exchanged are of size $n+1$ (x -variable has size n and the auxiliary y -variable is scalar).

The distributed algorithm in (16) requires exchanging x - and d -variables with the total size of $2n$, and it requires $2|\mathcal{E}|$ communications among the agents per iteration (since the graph is undirected). The algorithm in (17) requires exchanging x - and d -variables, and an additional scalar y -variable giving a total size of $2n+1$. It requires $|\mathcal{E}|$ communications among the agents per iteration (since the graph is directed). Finally, the algorithm in (19) requires exchanging x - and d -variables with the total size of $2n$, and requires a subset of links forming a spanning tree -hence about $2m$ communications per update (see [32] for precise assumptions). Overall, the communication requirements for the algorithms

in this section are similar to those for the algorithms of Section II and Section III. However, in every communication (directed or not), the algorithms of this section require twice larger size of the variables exchanged than those of the algorithms of Section II and Section III.

V. SOME OPEN DIRECTIONS FOR FURTHER RESEARCH

We have discussed some of the basic distributed methods for solving machine learning problems. The algorithms are obeying two main principles: the local knowledge of the graph structure (reflected in communications with one-hop neighbors only) and (private) knowledge of its own objective function only. To obey with these principles, the distributed algorithms considered here must use some form of “mixing” or “diffusion” step that propagates the needed information from one to all the agents in the network. This mixing step combined with some basic gradient method is the basic idea in all the distributed methods discussed here. Other variants of distributed method can be constructed such as those based on the ADMM approach (see the recent survey paper [23] for relevant literature).

In practice, there are several issues that arise in the implementations of the distributed methods due to the use of a communication network, such as quantization and link delays. These two aspects are currently of a great interest as they are the bottleneck for the efficiency of the methods. The effects of quantization in distributed methods have been studied somewhat limited in the past decade [34], [35], [36], and very recently in [37]. The effects of communication delays for distributed computational systems have been considered in [38], while [39] considers the impact of delays in the weighted-average consensus method. Very recently, a distributed algorithm that factors in link delays in communications have been proposed and investigated in [40]. The method identifies and communicates more frequently over critical links, and saves communication time by using other links less frequently. Such communication efficient approaches have not yet been studied for the recently developed methods that employ a gradient-tracking mechanism (i.e., the methods discussed in Section IV).

Despite the prolific research done on distributed methods over the past decade, there are still many directions for further exploration. For example, the gradient-push method in (15) and its variant with gradient-tracking mechanism in (17) have been analyzed for solving an unconstrained machine learning problem (4) over time-varying undirected graphs [15], [30]. It is interesting to see how should these methods be modified and analyzed for a constrained learning problem, where the constraint set X may itself be distributed, that is, the set X has the form $X = \cap_{i=1}^m X_i$ and each X_i is known to one agent only, where X_i is either simple (easy to be projected onto) or it is given as an intersection of a collection of convex (algebraic) functional inequalities.

Another direction of research is posed by the *Apull-Bpush* method in (18) which have been analyzed for a static (directed) communication graph only [33], [32]. It would be desirable to see its performance over time-varying directed graphs, which would extend the method to a more realistic scenarios when the graph connectivity structure may change.

Yet another general direction of the research is the development and investigation of asynchronous implementations of the distributed methods, where the agents' information exchange and updates need not be synchronous. Some random gossip- and broadcast-gossip methods have been studied in the literature, but those studies are somewhat limited (see [23] for more detailed account of such randomized methods).

REFERENCES

- [1] D. Bertsekas, *Nonlinear Programming*, 2nd ed. Athena Scientific, Belmont, MA, 1999.
- [2] Y. Nesterov, *Introductory lectures on convex optimization: A basic course*. Springer Science & Business Media, 2013.
- [3] R. Johnson and T. Zhang, "Accelerating stochastic gradient descent using predictive variance reduction," in *NIPS*, C. J. C. Burges, L. Bottou, Z. Ghahramani, and K. Q. Weinberger, Eds., 2013, pp. 315–323.
- [4] A. Defazio, F. Bach, and S. Lacoste-Julien, "Saga: A fast incremental gradient method with support for non-strongly convex composite objectives," in *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, Eds. Curran Associates, Inc., 2014, pp. 1646–1654. [Online]. Available: <http://papers.nips.cc/paper/5258-saga-a-fast-incremental-gradient-method-with-support-for-non-strongly-convex-composite-objectives.pdf>
- [5] M. DeGroot, "Reaching a consensus," *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.
- [6] J. Tsitsiklis, D. 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. H. Sayed and C. G. Lopes, "Adaptive processing over distributed networks," *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, vol. E90-A, no. 8, pp. 1504–1510, 2007.
- [8] A. Nedić and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Transactions on Automatic Control*, vol. 54, no. 1, pp. 48–61, 2009.
- [9] A. Sayed, *Adaptation, Learning, and Optimization over Networks*. Foundations and Trends in Machine Learning, 2014, vol. 7.
- [10] A. Nedić and J. Liu, "Distributed optimization for control," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. I, pp. 77–103, 2018.
- [11] D. Kempe, A. Dobra, and J. Gehrke, "Gossip-based computation of aggregate information," in *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, 2003, pp. 482–491.
- [12] A. Dominguez-Garcia and C. Hadjicostis, "Distributed strategies for average consensus in directed graphs," in *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference*, Dec 2011.
- [13] K. Tsianos, S. Lawlor, and M. Rabbat, "Push-sum distributed dual averaging for convex optimization," in *Proceedings of the IEEE Conference on Decision and Control*, 2012.
- [14] K. Tsianos, "The role of the network in distributed optimization algorithms: Convergence rates, scalability, communication / computation tradeoffs and communication delays," Ph.D. dissertation, McGill University, Dept. of Electrical and Computer Engineering, 2013.

- [15] A. Nedić and A. Olshevsky, "Distributed optimization over time-varying directed graphs," *IEEE Transactions on Automatic Control*, vol. 60, no. 3, pp. 601–615, 2015.
- [16] —, "Stochastic gradient-push for strongly convex functions on time-varying directed graphs," *IEEE Transactions on Automatic Control*, vol. 61, no. 12, pp. 3936 – 3947, 2016.
- [17] V. Borkar and P. Varaiya, "Asymptotic agreement in distributed estimation," *IEEE Transactions on Automatic Control*, vol. 27, no. 3, pp. 650–655, 1982.
- [18] J. Tsitsiklis, "Problems in decentralized decision making and computation," Ph.D. dissertation, Dept. of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1984.
- [19] A. I. Chen and A. Ozdaglar, "A fast distributed proximal-gradient method," in *Proceedings of the IEEE 50th Annual Allerton Conference on Communication, Control, and Computing*, 2012, pp. 601–608.
- [20] D. Jakovetić, J. Xavier, and J. Moura, "Fast distributed gradient methods," *IEEE Transactions on Automatic Control*, vol. 59, no. 5, pp. 1131–1146, May 2014.
- [21] K. Tsianos, S. Lawlor, and M. Rabbat, "Communication/computation tradeoffs in consensus-based distributed optimization," in *Proc. Advances in Neural Information Processing Systems*, Lake Tahoe, USA, Dec. 2012, pp. 1943–1951.
- [22] A. Berahas, R. Bollapragada, N. Keskar, and E. Wei, "Balancing communication and computation in distributed optimization," *IEEE Transactions on Automatic Control*, pp. 3141–3155, 2018.
- [23] A. Nedić, A. Olshevsky, and M. Rabbat, "Network topology and communication - computation tradeoffs in decentralized optimization," *Proceedings of the IEEE*, vol. 106, no. 5, pp. 953–976, 2018.
- [24] B. Gharesifard and J. Cortés, "Distributed strategies for generating weight-balanced and doubly stochastic digraphs," *European Journal of Control*, vol. 18, no. 6, pp. 539–557, 2012.
- [25] 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.
- [26] F. Beneit, V. Blondel, P. Thiran, J. Tsitsiklis, and M. Vetterli, "Weighted gossip: distributed averaging using non-doubly stochastic matrices," in *Proceedings of the 2010 IEEE International Symposium on Information Theory*, Jun. 2010.
- [27] C. N. Hadjicostis and T. Charalambous, "Average consensus in the presence of delays in directed graph topologies," *IEEE Transactions on Automatic Control*, vol. 59, no. 3, pp. 763–768, 2014.
- [28] C. N. Hadjicostis, N. H. Vaidya, and A. D. Domínguez-García, "Robust distributed average consensus via exchange of running sums," *IEEE Transactions on Automatic Control*, vol. 61, no. 6, pp. 1492–1507, 2016.
- [29] G. Qu and N. Li, "Harnessing smoothness to accelerate distributed optimization," *IEEE Transactions on Control of Network Systems*, vol. 5, no. 3, pp. 1245–1260, 2018.
- [30] 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.
- [31] P. Di Lorenzo and G. Scutari, "Next: In-network nonconvex optimization," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 2, no. 2, pp. 120–136, 2016.
- [32] S. Pu, W. Shi, J. Xu, and A. Nedić, "Push-pull gradient methods for distributed optimization in networks," 2018, arXiv preprint at <https://arxiv.org/abs/1810.06653>.
- [33] R. Xin and U. A. Khan, "A linear algorithm for optimization over directed graphs with geometric convergence," *IEEE Control Systems Letters*, vol. 2, pp. 315–320, 2018.
- [34] M. G. Rabbat and R. D. Nowak, "Quantized incremental algorithms for distributed optimization," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 4, pp. 798–808, 2005.

- [35] A. Nedić, A. Olshevsky, A. Ozdaglar, and J. N. Tsitsiklis, “Distributed subgradient methods and quantization effects,” in *Proceedings of the 47th IEEE Conference on Decision and Control*, 2008, pp. 4177–4184.
- [36] A. Kashyap, T. Başar, and R. Srikant, “Quantized consensus,” *Automatica*, vol. 43, no. 7, pp. 1192–1203, 2007.
- [37] A. Koloskova, S. Stich, and M. Jaggi, “Decentralized stochastic optimization and gossip algorithms with compressed communication,” in *International Conference on Machine Learning*, 2019, pp. 3478–3487.
- [38] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and distributed computation: Numerical methods*, 1st ed. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1989.
- [39] A. Nedić and A. Ozdaglar, “Convergence rate for consensus with delays,” *Journal of Global Optimization*, vol. 47, pp. 437–456, 2010.
- [40] J. Wang, A. K. Sahu, Z. Yang, G. Joshi, and S. Kar, “MATCHA: Speeding up Decentralized SGD via Matching Decomposition Sampling,” 2019, available on arxiv at <https://arxiv.org/pdf/1905.09435.pdf>.

Angelia Nedić received her Ph.D. from Moscow State University, Moscow, Russia, in Computational Mathematics and Mathematical Physics (1994), and a Ph.D. from Massachusetts Institute of Technology, Cambridge, USA, in Electrical and Computer Science Engineering (2002). She has worked as a senior engineer in BAE Systems North America, Advanced Information Technology Division at Burlington, MA. She is a faculty member of the school of Electrical, Computer and Energy Engineering at Arizona State University, Tempe. Prior to joining Arizona State University, she has been a Willard Scholar faculty member at the University of Illinois at Urbana-Champaign. She is a recipient (jointly with co-authors) of the Best Paper Award at the Winter Simulation Conference 2013 and the Best Paper Award at the International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt) 2015. Her general research interest is in optimization, large scale complex systems dynamics, equilibrium and variational inequality problems. She serves as Associate Editor for SIAM Journal on Optimization, Operations Research, and IEEE Transactions on Signal and Information Processing over Networks.