

# ByRDIE: Byzantine-resilient Distributed Coordinate Descent for Decentralized Learning

Zhixiong Yang, *Graduate Student Member, IEEE*, and Waheed U. Bajwa, *Senior Member, IEEE*

**Abstract**—Distributed machine learning algorithms enable learning of models from datasets that are distributed over a network without gathering the data at a centralized location. While efficient distributed algorithms have been developed under the assumption of faultless networks, failures that can render these algorithms nonfunctional occur frequently in the real world. This paper focuses on the problem of Byzantine failures, which are the hardest to safeguard against in distributed algorithms. While Byzantine fault tolerance has a rich history, existing work does not translate into efficient and practical algorithms for high-dimensional learning in *fully* distributed (also known as decentralized) settings. In this paper, an algorithm termed *Byzantine-resilient distributed coordinate descent* (ByRDIE) is developed and analyzed that enables distributed learning in the presence of Byzantine failures. Theoretical analysis (convex settings) and numerical experiments (convex and nonconvex settings) highlight its usefulness for high-dimensional distributed learning in the presence of Byzantine failures.

**Index Terms**—Byzantine failure, consensus, coordinate descent, decentralized learning, distributed optimization, empirical risk minimization, machine learning

## I. INTRODUCTION

One of the fundamental goals in machine learning is to learn a model that minimizes the statistical risk. This is typically accomplished through stochastic optimization techniques, with the underlying principle referred to as *empirical risk minimization* (ERM) [2]–[4]. The ERM principle involves the use of a training dataset and tools from optimization theory. Traditionally, training data have been assumed available at a centralized location. Many recent applications of machine learning, however, involve the use of a dataset that is either distributed across different locations (e.g., the *Internet of Things*) or that cannot be processed at a single machine due to its size (e.g., social network data). The ERM framework in this setting of distributed training data is often referred to as *decentralized* or *distributed* learning [5], [6].

While there exist excellent works that solve the problem of distributed learning, all these works make a simplified assumption that all nodes in the network function as expected. Unfortunately, this assumption does not always hold true in practice; examples include cyber attacks, malfunctioning equipments and undetected failures [7], [8]. When a node

arbitrarily deviates from its intended behavior, it is termed to have undergone Byzantine failure [9]. While Byzantine failures are hard to detect in general, they can easily jeopardize the operation of the whole network [10]–[12].

In particular, with just a simple strategy, one can show that a single Byzantine node in the network can lead to failures of most state-of-the-art distributed learning algorithms [13]–[15]. The main contribution of this paper is to introduce and analyze an algorithm that completes the distributed learning task in the presence of Byzantine failures in the network.

### A. Related Works

To achieve the goal of distributed learning, one usually sets up a distributed optimization problem by defining and minimizing a (regularized) loss function on the training data of each node. The resulting problem can then be solved by distributed optimization algorithms. Several types of distributed algorithms have been introduced in the past [16]–[24]. The most common of these are gradient-based methods [16], [17], [21], which usually have low local computational complexity. Another type includes augmented Lagrangian-based methods [18]–[20], which iteratively update the primal and dual variables. These methods often require the ability to locally solve an optimization subproblem at each node. A third type includes second-order distributed methods [22], [24], which tend to have high computational and/or communications costs. While any of these distributed optimization algorithms can be used for distributed learning, they all make the assumption that there are no failures in the network.

Byzantine-resilient algorithms for a variety of problems have been studied extensively over the years [9], [25]–[28]. Byzantine-resilient algorithms for scalar- and vector-averaging distributed consensus were studied in [29]–[33]. The algorithms proposed in [13], [14] extend some of these works from scalar consensus to scalar-valued distributed optimization, but they cannot be used in a straightforward manner for vector-valued distributed optimization problems. In the context of distributed learning, [15] introduces a method to implement distributed support vector machines (SVMs) under Byzantine failures, but it neither provides theoretical guarantees nor generalizes to other learning problems. A number of recent works have also investigated Byzantine-resilient distributed learning in networks that are equipped with a central processor (often referred to as a *parameter server*) [34]–[44]. Theoretical guarantees developed in such works make extensive use of the fact that the parameter server is connected to all network nodes; as such, these guarantees cannot be generalized to Byzantine-resilient learning in *fully* distributed networks.

Some of the results reported in this paper were presented at the 2018 IEEE Data Science Workshop (DSW'18) [1]. This work is supported in part by the NSF under award CCF-1453073, by the ARO under award W911NF-17-1-0546, and by the DARPA Lagrange Program under ONR/SPAWAR contract N660011824020.

The authors are with the Department of Electrical and Computer Engineering, Rutgers University–New Brunswick, 94 Brett Rd, Piscataway, NJ 08854, USA (Emails: {zhixiong.yang, waheed.bajwa}@rutgers.edu).

## B. Our Contributions

We have already noted several limitations of works like [13]–[15], [29]–[44] within the context of Byzantine-resilient learning in fully distributed settings. Additionally, one of the limitations of existing Byzantine-resilient algorithms such as [33] is that, when required to work with vector-valued data, they make a strong assumption on the network topology. Specifically, the smallest size of neighborhood of each node in the vector setting depends linearly on the dimensionality of the problem. This is impractical for most learning problems since the dimensionality of the training samples is usually much larger than the size of the neighborhood of each node. Within the specific context of Byzantine-resilient distributed machine learning, the main limitation of fully distributed algorithms proposed in [13], [14] is that they only yield the minimizer of a convex combination of local empirical risk functions for scalar-valued problems. Since this (scalar) minimizer is usually different from the minimizer of the exact average of local loss functions, these works cannot guarantee by themselves alone that the outputs of any vector-valued algorithms that leverage similar ideas will converge to either the minimum empirical risk or the minimum statistical risk.

In contrast to prior works, our work has two main contributions. First, we propose a Byzantine-resilient algorithm that scales well with the dimensionality of distributed learning problems. The proposed algorithm is an *inexact* distributed variant of coordinate descent [45] that first breaks vector-valued distributed learning problems into a (possibly infinite) sequence of scalar-valued distributed subproblems and then solves these subproblems in a Byzantine-resilient manner by leveraging ideas from works such as [13], [14]. The inexactness here stems from the fact that—even when using an exact line search procedure—the subproblems cannot be solved exactly in the presence of Byzantine failures (cf. Sec. III). This inexactness in the solution of each subproblem, whether or not an exact line search procedure is used, is one of the reasons the analytical techniques of prior works such as [13], [14] do not lead to theoretical guarantees for the proposed coordinate descent algorithm. In contrast, under the assumption of independent and identically distributed training samples across the network, we provide theoretical guarantees that the output of the proposed algorithm will lead to the minimum statistical risk with high probability. Our theoretical analysis, which forms the second main contribution of this work, also highlights the fact that the output of our algorithm can statistically converge to the minimizer of the statistical risk faster than using only local information by a factor that will be shown in the sequel.

## C. Notation and Organization

Given a vector  $v$  and a matrix  $A$ ,  $[v]_k$  and  $[A]_{ij}$  denote their  $k$ -th and  $(i, j)$ -th element, respectively. We use  $v|_{[v]_k=a'}$  to denote the vector formed by replacing the  $k$ -th element of  $v$  with  $a' \in \mathbb{R}$ , while  $(\cdot)^T$  denotes the transpose operation. We use  $\|v\|$  and  $\|v\|_\infty$  to denote the  $\ell_2$ - and  $\ell_\infty$ -norms of  $v$ , respectively,  $\mathbf{1}$  to denote the vector of all ones, and  $I$  to denote the identity matrix. Given a vector  $v$  and a constant

$\gamma > 0$ , we use  $B(v, \gamma) := \{v' : \|v - v'\| \leq \gamma\}$  to denote an  $\ell_2$ -ball of radius  $\gamma$  centered around  $v$ . Given a set,  $|\cdot|$  denotes its cardinality. We use the scaling relation  $f(n) = \mathcal{O}(g(n))$  if  $\exists c_o > 0, n_o : \forall n \geq n_o, f(n) \leq c_o g(n)$ . Finally, unless stated otherwise, we use  $\nabla f(w, (x, y))$  to denote the gradient of a function  $f(w, (x, y))$  with respect to  $w$ .

The rest of this paper is organized as follows. Section II gives the problem formulation. Section III discusses the proposed algorithm along with theoretical guarantees for consensus as well as statistical and algorithmic convergence. Numerical results corresponding to distributed convex and nonconvex learning on two different datasets are provided in Section IV, while Section V concludes the paper.

## II. PROBLEM FORMULATION

Given a connected network in which nodes have access to local training data, our goal is to learn a machine learning model from the distributed data, even in the presence of Byzantine failures. We begin with a model of our basic problem, which will be followed with a model for Byzantine failures and the final problem statement.

Consider a connected network of  $M$  nodes, expressed as a directed, static graph  $\mathcal{G}(J, \mathcal{E})$ . Here, the set  $J := \{1, \dots, M\}$  represents nodes in the network, while the set of edges  $\mathcal{E}$  represents communication links between different nodes. Specifically,  $(j, i) \in \mathcal{E}$  if and only if node  $i$  can receive information from node  $j$ . Each node  $j$  has access only to a local training set  $S_j = \{(x_{jn}, y_{jn})\}_{n=1}^{|S_j|}$ . Let  $x \in \mathbb{R}^P$  represent the training features satisfying  $\|x\| \leq B$  for some constant  $B$  and  $y$  be the corresponding label. For classification,  $y \in \{-1, 1\}$ , while  $y \in \mathbb{R}$  for regression. We assume that the training samples are independent and identically distributed (i.i.d.) and drawn from an unknown distribution  $D$ , i.e.,  $(x_{jn}, y_{jn}) \sim D$ . For simplicity, we assume that the cardinalities of the local training sets are the same, i.e.,  $|S_j| = N$ . The generalization to the case when  $S_j$ 's are not equal is trivial.

*Remark 1.* Note that while the main problem is being formulated here under the supervised setting, our proposed framework and the final results are equally applicable under both unsupervised and semi-supervised settings.

In machine learning, one would ideally like to collect all the data into one set  $S = \{(x_n, y_n)\}_{n=1}^{|S|}$  and perform centralized training on  $S$ . The goal in that case is to learn a function that reliably maps  $x$  to  $y$ . One popular mapping is  $y = w^T x$  (sometimes this is defined as  $y = w^T x + b$ , which can be transformed into  $y = w^T x$  by adding one more dimension to  $x$ ). To find a “good”  $w$ , one first defines a (non-negative) loss function  $\ell(w, (x, y))$ , where the value of loss function increases when the difference between the mapping of  $x$  and  $y$  increases. To avoid overfitting, a (non-negative) regularizer  $R(w)$  is often added to the loss function. Then one can solve for  $w$  by statistically minimizing a regularized loss function

$$f(w, (x, y)) := \ell(w, (x, y)) + R(w). \quad (1)$$

The regularized  $f(w, (x, y))$  is often referred to as *risk function*. In this paper, we focus on the class of convex differentiable loss functions and strictly convex and smooth regular-

izers.<sup>1</sup> Examples include square loss  $(1 - y \cdot w^T x)^2$ , square hinge loss  $\max(0, 1 - y \cdot w^T x)^2$ , logistic loss  $\ln(1 + e^{-y \cdot w^T x})$  and  $R(w) = \frac{\lambda}{2} \|w\|^2$ . We also assume the gradient of the loss function  $\ell(w, (x, y))$  is  $L$ -Lipschitz continuous. Since  $R(w)$  is smooth, we formally state the Lipschitz assumption as follows.

**Assumption 1.** The risk function  $f = \ell(w, (x, y)) + R(w)$  satisfies

$$\forall w_1, w_2, \|\nabla f(w_1, (x, y)) - \nabla f(w_2, (x, y))\| \leq L \|w_1 - w_2\|.$$

Note that Assumption 1 implies the risk function itself is also Lipschitz continuous [46]:

$$\forall w_1, w_2, \|f(w_1, (x, y)) - f(w_2, (x, y))\| \leq L' \|w_1 - w_2\|.$$

Centralized machine learning focuses on learning the “best” mapping from  $x$  to  $y$  in terms of the following stochastic optimization problem (referred to as *statistical risk minimization*):

$$w^* := \arg \min_{w \in \mathbb{R}^p} \mathbb{E}[f(w, (x, y))], \quad (2)$$

where the expectation is with respect to the unknown distribution  $D$ . In this work, we assume without loss of generality that  $\|w^*\|_\infty \leq \gamma^*$  for some constant  $\gamma^* \in \mathbb{R}_+$ . Owing to the equivalence of norms in finite-dimensional spaces, extensions of our analysis to bounds on  $w^*$  in norms other than the  $\ell_\infty$ -norm are straightforward, but would not be pursued in this work. In the following, we will use  $W$  to denote the smallest  $\ell_\infty$ -ball centered around the origin that contains  $w^*$  and the *sublevel set* of the *statistical risk*  $\mathbb{E}[f(w, (x, y))]$  at level  $c_0 := \mathbb{E}[f(0, (x, y))]$ .<sup>2</sup> Specifically, defining  $W_\gamma := \{w \in \mathbb{R}^p : \|w\|_\infty \leq \gamma\}$  and

$$\Gamma := \arg \inf_{\gamma} \gamma \text{ s.t. } W_\gamma \supset \{w : \mathbb{E}[f(w, (x, y))] \leq c_0\},$$

we have that  $W := \{w \in \mathbb{R}^p : \|w\|_\infty \leq \Gamma\}$ . We now make another assumption.

**Assumption 2.** For any  $w \in W$ , the loss function  $\ell(w, (x, y))$  is bounded almost surely over all training samples, i.e.,

$$\forall (x, y) \in \bigcup_{j \in J} S_j, \quad 0 \leq \ell(w, (x, y)) \leq C < \infty.$$

Note that Assumption 2 would be satisfied for datasets with finite-valued training samples because of the Lipschitz continuity of  $\ell(w, (x, y))$  and the compactness of  $W$ .

*Remark 2.* While the set  $W$  has no algorithmic significance, it is later shown that the iterates of the proposed algorithm stay within it. Because of this reason,  $W$  makes an appearance in our theoretical guarantees concerning statistical convergence.

Since  $D$  is unknown in (2), one cannot solve for  $w^*$  directly. A broadly adopted alternative then is to utilize the training

<sup>1</sup>A function  $R(w)$  is strictly convex if it satisfies  $R(zw_1 + (1-z)w_2) < zR(w_1) + (1-z)R(w_2)$  for any  $z \in (0, 1)$ . Further,  $R(w)$  is smooth if it is differentiable for all orders.

<sup>2</sup>Our choice of the level  $c_0$  is tied to our initialization of the proposed algorithm, which is currently taken to be  $w = 0$ . This initialization, however, is without loss of generality and can be changed according to one’s preference, resulting in a minor change in the definition of the set  $W$ .

data and minimize the *empirical risk*  $\hat{f}(w, S)$  in lieu of the statistical risk  $\mathbb{E}[f(w, (x, y))]$ , where

$$\begin{aligned} \hat{f}(w, S) &:= \frac{1}{|S|} \sum_{n=1}^{|S|} f(w, S) \\ &:= \frac{1}{|S|} \sum_{n=1}^{|S|} \ell(w, (x_n, y_n)) + R(w). \end{aligned} \quad (3)$$

In particular, the minimizer of the empirical risk  $\hat{f}(w, S)$  can be shown to converge to  $w^*$  with high probability [3], [47]. In the case of i.i.d. training samples, and for a fixed probability of failure and ignoring some log terms, the rate of this *statistical convergence* is known to be  $\mathcal{O}\left(\frac{1}{\sqrt{|S|}}\right) \equiv \mathcal{O}\left(\frac{1}{\sqrt{MN}}\right)$  under mild assumptions on the problem [47].

In many distributed learning problems, training data cannot be made available at a single location. This then requires learning  $w^*$  in a distributed fashion, which can be done by employing distributed optimization techniques. The main idea of distributed optimization-based learning is to minimize the average of local empirical risks, i.e.,

$$\frac{1}{M} \sum_{j=1}^M \hat{f}(w, S_j) = \frac{1}{MN} \sum_{j=1}^M \sum_{n=1}^N \ell(w, (x_{jn}, y_{jn})) + R(w).$$

To achieve this goal, we need nodes to cooperate with each other by communicating over network edges. Specifically, define the neighborhood of node  $i$  as  $\mathcal{N}_i := \{j \in J : (j, i) \in \mathcal{E}\}$ . We say that node  $j$  is a neighbor of node  $i$  if  $j \in \mathcal{N}_i$ . Distributed learning algorithms proceed iteratively. In each iteration  $(r + 1)$  of the algorithm, node  $j$  is expected to accomplish two tasks:

- 1) Update a local variable  $w_j^r$  according to some (deterministic or stochastic) rule  $g_j(\cdot)$ , and
- 2) Broadcast the updated local variable to other nodes, where node  $i$  can receive the broadcasted information from node  $j$  only if  $j \in \mathcal{N}_i$ .

#### A. Byzantine Failure Model

While distributed learning via message passing is well understood [18], [48], existing protocols require all nodes to operate as intended. In contrast, the main assumption in this paper is that some of the nodes in the network can undergo Byzantine failures, formally defined as follows.

**Definition 1.** A node  $j \in J$  is said to be Byzantine if during any iteration, it either updates its local variable using an update function  $g_j'(\cdot) \neq g_j(\cdot)$  or it broadcasts some value other than the intended update to its neighbors.

In this paper, we assume there are  $b$  Byzantine nodes in the network. Knowing the exact value of  $b$ , however, is not necessary. One can, for example, set  $b$  to be an upper bound on the number of Byzantine nodes. Let  $J' \subset J$  denote the set of nonfaulty nodes. Without loss of generality, we assume the nonfaulty nodes are labeled from 1 to  $|J'|$ . We now provide some definitions and an assumption that are common in the literature; see, e.g., [13], [14].

**Definition 2.** A subgraph  $\mathcal{G}_r$  of  $\mathcal{G}$  is called a reduced graph if it is generated by (i) removing all Byzantine nodes along with their incoming and outgoing edges, and (ii) removing additionally up to  $b$  incoming edges from each nonfaulty node.

**Definition 3.** A ‘‘source component’’ of a graph is a collection of nodes such that each node in the source component has a directed path to every other node in the graph.

**Assumption 3.** All reduced graphs  $\mathcal{G}_r$  generated from  $\mathcal{G}(J, \mathcal{E})$  contain a source component of cardinality at least  $(b + 1)$ .

The purpose of Assumption 3 is to ensure there is enough redundancy in the graph to tolerate Byzantine failures. Note that the total number of different reduced graphs one can generate from  $\mathcal{G}$  is finite as long as  $M$  is finite. So, in theory, Assumption 3 can be certified for any graph. But efficient certification of this assumption remains an open problem. In the case of Erdős–Rényi graphs used in our experiments, however, we have observed that Assumption 3 is typically satisfied whenever the ratio of the average incoming degree of the graph and the number of Byzantine nodes is high enough.

*Remark 3.* Assumption 3 has been stated to guarantee resilience against the *worst-case* attack scenario in which all Byzantine nodes concentrate in the neighborhood of any one of the network nodes. While such worst-case analysis is customary in the literature on Byzantine fault tolerance [14], [31], [32], it does impose stringent constraints on the network topology. Such topology constraints, however, seem to be unavoidable for the types of ‘‘local screening’’ algorithms being considered in the fully distributed setting of this paper.

### B. Problem Statement

Our goal is to develop a Byzantine fault-tolerant algorithm for distributed learning under Assumptions 1–3. Specifically, under the assumption of at most  $b$  Byzantine nodes in the network, we need to accomplish the following:

- 1) Achieve consensus among nonfaulty nodes in the network, i.e.,  $w_j^r = w_i^r \forall i, j \in J'$  as the number of algorithmic iterations  $r \rightarrow \infty$ ; and
- 2) Guarantee  $w_j^r \rightarrow w^* \forall j \in J'$  as the number of training samples at nonfaulty nodes  $N \rightarrow \infty$ .

In particular, the latter objective requires understanding both the statistical convergence and the algorithmic convergence of the distributed empirical risk minimization problem in the presence of Byzantine failures.

## III. BYZANTINE-RESILIENT DISTRIBUTED COORDINATE DESCENT FOR DECENTRALIZED LEARNING

In distributed learning, one would ideally like to solve the empirical risk minimization (ERM) problem

$$w_{opt} = \arg \min_{w \in \mathbb{R}^P} \frac{1}{MN} \sum_{j=1}^M \sum_{n=1}^N \ell(w, (x_{jn}, y_{jn})) + R(w) \quad (4)$$

at each node  $j \in J'$  and show that  $w_{opt} \rightarrow w^*$  as  $N \rightarrow \infty$ . However, we know from prior work [49] that this is infeasible when  $b \geq 1$ . Nonetheless, we establish in the following

### Algorithm 1 Byzantine-resilient distributed coordinate descent

---

**Input:**  $S_1, S_2, \dots, S_M, \{\rho(\tau)\}_{\tau=1}^\infty, b \in \mathbb{N}, \bar{r} \in \mathbb{N}, T \in \mathbb{N}$

- 1: **Initialize:**  $r \leftarrow 1, t \leftarrow 1$ , and  $\forall j \in J', w_j^r(t) \leftarrow 0$
- 2: **for**  $r = 1, 2, \dots, \bar{r}$  **do**
- 3:     **for**  $k = 1, 2, \dots, P$  **do**
- 4:         **for**  $t = 1, 2, \dots, T$  **do**
- 5:             **for**  $j = 1, 2, \dots, |J'|$  **do (in parallel)**
- 6:                 Receive  $[w_i^r(t)]_k$  from all  $i \in \mathcal{N}_j$
- 7:                 Find  $\mathcal{N}_j^s(r, k, t), \mathcal{N}_j^l(r, k, t), \mathcal{N}_j^*(r, k, t)$  according to (5), (6), and (7)
- 8:                 Update  $[w_j^r(t+1)]_k$  as in (8)
- 9:             **end for**
- 10:         **end for**
- 11:     **end for**
- 12:      $w_j^{r,T} \leftarrow w_j^r(T+1), \forall j \in J'$
- 13:      $w_j^{r+1}(1) \leftarrow w_j^{r,T}, \forall j \in J'$
- 14: **end for**

**Output:**  $\{w_j^{\bar{r}, T}\}_{j \in J'}$

---

that distributed strategies based on coordinate descent algorithms [45] can still be used to solve a variant of (4) at nonfaulty nodes and guarantee that the solutions converge to the minimizer  $w^*$  of the statistical risk for the case of i.i.d. training data. We refer to our general approach as **Byzantine-Resilient Distributed coordinate dEscent (ByRDIE)**, which is based on key insights gleaned from two separate lines of prior work. First, it is known that certain types of scalar-valued distributed optimization problems can be inexactly solved in the presence of Byzantine failures [13], [14]. Second, coordinate descent algorithms break down vector-valued optimization problems into a sequence of scalar-valued problems [45]. The algorithmic aspects of ByRDIE leverage these insights, while the theoretical aspects leverage tools from Byzantine-resilient distributed consensus [30], optimization theory [45], stochastic convex optimization [47], and statistical learning theory [3].

### A. Algorithmic Details

ByRDIE involves splitting the ERM problem (4) into  $P$  one-dimensional subproblems using coordinate descent and then solving each scalar-valued subproblem using the Byzantine-resilient approach described in [13]. The exact implementation is detailed in Algorithm 1. The algorithm can be broken into an outer loop (Step 2) and an inner loop (Step 4). The outer loop is the coordinate descent loop, which breaks the vector-valued optimization problem in each iteration  $r$  into  $P$  scalar-valued subproblems. The inner loop solves a scalar-valued optimization problem in each iteration  $t$  and ensures resilience to Byzantine failures. We assume the total number of iterations  $\bar{r}$  for coordinate descent are specified during initialization. We use  $[w_j^r(t)]_k$  to denote the  $k$ -th element of  $w_j$  at the  $r$ -th iteration of the coordinate descent loop and the  $t$ -th iteration of the  $k$ -th subproblem (coordinate). Without loss of generality, we initialize  $[w_j^1(1)]_k = 0, \forall k = 1, \dots, P$ .

We now fix some  $r$  and  $k$ , and focus on the implementation of the inner loop (Step 4). Every node has some  $[w_j^r(1)]_k$  at the

start of the inner loop ( $t = 1$ ). During each iteration  $t$  of this loop, all (nonfaulty) nodes engage in the following: *broadcast*, *screening*, and *update*. In the broadcast step (Step 6), all nodes  $i \in J$  broadcast  $[w_i^r(t)]_k$ 's and each node  $j \in J$  receives  $[w_i^r(t)]_k, \forall i \in \mathcal{N}_j$ . During this step, a node can receive values from both nonfaulty and Byzantine neighbors. The main idea of the screening step (Step 7) is to reject values at node  $j$  that are either "too large" or "too small" so that the values being used for update by node  $j$  in each iteration will be upper and lower bounded by a set of values generated by nonfaulty nodes. To this end, we partition  $\mathcal{N}_j$  into 3 subsets  $\mathcal{N}_j^*(r, k, t)$ ,  $\mathcal{N}_j^s(r, k, t)$  and  $\mathcal{N}_j^l(r, k, t)$ , which are defined as following:

$$\mathcal{N}_j^s(r, k, t) = \arg \min_{X: X \subset \mathcal{N}_j, |X|=b} \sum_{i \in X} [w_i^r(t)]_k, \quad (5)$$

$$\mathcal{N}_j^l(r, k, t) = \arg \max_{X: X \subset \mathcal{N}_j, |X|=b} \sum_{i \in X} [w_i^r(t)]_k, \quad (6)$$

$$\mathcal{N}_j^*(r, k, t) = \mathcal{N}_j \setminus \mathcal{N}_j^s(r, k, t) \setminus \mathcal{N}_j^l(r, k, t). \quad (7)$$

The step is called screening because node  $j$  only uses  $[w_i^r(t)]_k$ 's from  $\mathcal{N}_j^*(r, k, t)$  to update its local variable. Note that there might still be  $[w_i^r(t)]_k$ 's received from Byzantine nodes in  $\mathcal{N}_j^*(r, k, t)$ . We will see later, however, that this does not effect the workings of the overall algorithm.

The final step of the inner loop in ByRDIE is the update step (Step 8). Using  $[\nabla \hat{f}(w_j^r(t), S_j)]_k$  to denote the  $k$ -th element of  $\nabla \hat{f}(w_j^r(t), S_j)$ , we can write this update step as follows:

$$[w_j^r(t+1)]_k = \frac{1}{|\mathcal{N}_j^*(r, k, t)| - 2b + 1} \sum_{i \in \mathcal{N}_j^*(r, k, t) \cup \{j\}} [w_i^r(t)]_k - \rho(r+t-1) [\nabla \hat{f}(w_j^r(t), S_j)]_k, \quad (8)$$

where  $\{\rho(\tau)\}_{\tau=1}^\infty$  are square-summable (but not summable), diminishing stepizes:  $0 < \rho(\tau+1) \leq \rho(\tau)$ ,  $\sum_\tau \rho(\tau) = \infty$ , and  $\sum_\tau \rho^2(\tau) < \infty$ . Notice that  $[w_j^r(T+1)]_k$  is updated after the  $k$ -th subproblem of coordinate descent in iteration  $r$  finishes and it stays fixed until the start of the  $k$ -th subproblem in the  $(r+1)$ -th iteration of coordinate descent. An iteration  $r$  of the coordinate descent loop is considered complete once all  $P$  subproblems within the loop are solved. The local variable at each node  $j$  at the end of this iteration is then denoted by  $w_j^{r,T}$  (Step 12). We also express the output of the whole algorithm as  $\{w_j^{r,T}\}_{j \in J'}$ . Finally, note that while Algorithm 1 cycles through the  $P$  coordinates of the optimization variables in each iteration  $r$  in the natural order, one can use any permutation of  $\{1, \dots, P\}$  in place of this order.

We conclude this discussion by noting that the parameter  $T$  in ByRDIE, which can take any value between 1 and  $\infty$ , trades off consensus among the nonfaulty nodes and the convergence rate; see Sec. III-E for further discussion on this tradeoff and Sec. IV for numerical experiments that highlight this tradeoff. Our theoretical analysis of ByRDIE focuses on the two extreme cases of  $T \rightarrow \infty$  and  $T = 1$ . In both cases, we establish in the following that the output of ByRDIE at nonfaulty nodes converges in probability to the minimum of the statistical risk. Convergence guarantees for a finite-valued  $T > 1$  can be obtained from straightforward modifications of the analytical techniques used in the following.

## B. Theoretical Guarantees: Consensus

We now turn our attention to theoretical guarantees for ByRDIE. We first show that it leads to consensus among nonfaulty nodes in the network, i.e., all nonfaulty nodes agree on the same variable, in the limit of large  $\bar{r}$  and/or  $T$ . Then we show in the next sections that the output of ByRDIE converges to the statistical optimum in the limit of large  $\bar{r}$  for the two extreme choices of  $T$ :  $T \rightarrow \infty$  and  $T = 1$ .

Let us begin by establishing the claim of consensus. To this end, focusing exclusively on dimension  $k$  in ByRDIE, we see that the  $k$ -th dimension of each local variable is updated  $\bar{r}T$  times. Fixing any  $k$ , we can use an index  $m := (r-1)T + t$  to denote the sequence generated for the  $k$ -th dimension up to iteration  $(r, t)$  in ByRDIE. We define a vector  $\Omega(m) \in \mathbb{R}^{|\mathcal{J}'|}$  such that  $[\Omega(m)]_j := [w_j^r(t)]_k$  and let  $\omega_j(m)$  denote  $[G(m)]_j$ . Similarly, we define a vector  $G(m) \in \mathbb{R}^{|\mathcal{J}'|}$  such that  $[G(m)]_j := [\nabla \hat{f}(w_j^r(t), S_j)]_k$ . Next, let  $\bar{\rho}(m) := \rho(r+t-1)$  and note that  $\bar{\rho}(m)$  also satisfies  $\bar{\rho}(m) \rightarrow 0$ ,  $\sum_{m=1}^\infty \bar{\rho}(m) = \infty$  and  $\sum_{m=1}^\infty \bar{\rho}^2(m) < \infty$ . We can now express the update of the sequence corresponding to the  $k$ -th dimension at nonfaulty nodes in a matrix form as follows:

$$\Omega(m+1) = Y(m)\Omega(m) - \bar{\rho}(m)G(m), \quad (9)$$

where  $Y(m)$  is a matrix that is fully specified in the following.

Let  $\mathcal{N}_j^s$  and  $\mathcal{N}_j^b$  denote the nonfaulty nodes and the Byzantine nodes, respectively, in the neighborhood of  $j \in J'$ , i.e.,  $\mathcal{N}_j^s = J' \cap \mathcal{N}_j$  and  $\mathcal{N}_j^b = \mathcal{N}_j \setminus \mathcal{N}_j^s$ . Notice that one of two cases can happen during each iteration at node  $j \in J'$ :

$$\mathcal{N}_j^{s*}(m) \cap \mathcal{N}_j^b \neq \emptyset, \quad \text{or} \quad (10(a))$$

$$\mathcal{N}_j^{s*}(m) \cap \mathcal{N}_j^b = \emptyset. \quad (10(b))$$

For case (10(a)), since  $|\mathcal{N}_j^b| \leq b$  and  $|\mathcal{N}_j^{s*}(m)| = |\mathcal{N}_j^s(m)| = b$ , we must have  $\mathcal{N}_j^{s*}(m) \cap \mathcal{N}_j^b \neq \emptyset$  and  $\mathcal{N}_j^l(m) \cap \mathcal{N}_j^s \neq \emptyset$ . Then for each  $i \in \mathcal{N}_j^{s*}(m) \cap \mathcal{N}_j^b$ ,  $\exists s_j^i \in \mathcal{N}_j^s(m) \cap \mathcal{N}_j^l$  and  $l_j^i \in \mathcal{N}_j^l(m) \cap \mathcal{N}_j^s$  satisfying  $\omega_{s_j^i}(m) \leq \omega_i(m) \leq \omega_{l_j^i}(m)$ . Therefore, we have for each  $i \in \mathcal{N}_j^{s*}(m) \cap \mathcal{N}_j^b$  that  $\exists \theta_i^j(m) \in [0, 1]$  such that the following holds:

$$\omega_i(m) = \theta_i^j(m)\omega_{s_j^i}(m) + (1 - \theta_i^j(m))\omega_{l_j^i}(m). \quad (11)$$

We can now rewrite the update at node  $j \in J'$  as follows:

$$\begin{aligned} \omega_j(m+1) &= \frac{1}{|\mathcal{N}_j^*(m)| - 2b + 1} \left( \omega_j(m) + \sum_{i \in \mathcal{N}^s \cap \mathcal{N}_j^{s*}(m)} \omega_i(m) \right. \\ &+ \left. \sum_{i \in \mathcal{N}^b \cap \mathcal{N}_j^{s*}(m)} (\theta_i^j(m)\omega_{s_j^i}(m) + (1 - \theta_i^j(m))\omega_{l_j^i}(m)) \right) \\ &+ \bar{\rho}(m)[G(m)]_j. \end{aligned} \quad (12)$$

It can be seen from (12) that the screening rule in ByRDIE effectively enables nonfaulty nodes to replace data received from Byzantine nodes with convex combinations of data received from nonfaulty nodes in their neighborhoods. This

enables us to express the updates at nonfaulty nodes in the form (9), with the entries of  $Y(m)$  given by

$$[Y(m)]_{ji} = \begin{cases} \frac{1}{|\mathcal{N}_j|-2b+1}, & i = j, \\ \frac{1}{|\mathcal{N}_j|-2b+1}, & i \in \mathcal{N}' \cap \mathcal{N}_j^*(m), \\ \sum_{n \in \mathcal{N}^b \cap \mathcal{N}_j^*(m)} \frac{\theta_n^j(m)}{|\mathcal{N}_j|-2b+1}, & i \in \mathcal{N}' \cap \mathcal{N}_j^s(m), \\ \sum_{n \in \mathcal{N}^b \cap \mathcal{N}_j^*(m)} \frac{1-\theta_n^j(m)}{|\mathcal{N}_j|-2b+1}, & i \in \mathcal{N}' \cap \mathcal{N}_j^l(m), \\ 0, & \text{otherwise.} \end{cases} \quad (13)$$

Notice further that, since  $\mathcal{N}_j^*(m) \cap \mathcal{N}_j^b = \emptyset$  (cf. (12)), case (10(b)) corresponds to a special case of case (10(a)) in which we keep only the first, second, and last rows of (13). It is worth noting here that our forthcoming proof does not require knowledge of  $Y(m)$ ; in particular, since the choices of  $s_j^i$  and  $l_j^i$  are generally not unique,  $Y(m)$  itself is also generally not unique. The main thing that matters here is that  $Y(m)$  will always be a row stochastic matrix; we refer the reader to [30] for further properties of  $Y(m)$ .

In order to complete our claim that nonfaulty nodes achieve consensus under ByRDIE, even in the presence of Byzantine failures in the network, fix an arbitrary  $m_0 \geq 0$  and consider  $m > m_0$ . It then follows from (9) that

$$\begin{aligned} \Omega(m+1) &= Y(m)\Omega(m) - \bar{\rho}(m)G(m) \\ &= Y(m)Y(m-1) \cdots Y(m_0)\Omega(m_0) - \bar{\rho}(m)G(m) \\ &\quad - \sum_{\tau=m_0}^{m-1} Y(m)Y(m-1) \cdots Y(\tau+1)\bar{\rho}(\tau)G(\tau). \end{aligned} \quad (14)$$

We now define matrices

$$\Phi(m, m_0) := Y(m)Y(m-1) \cdots Y(m_0),$$

$\Phi(m, m) := Y(m)$ , and  $\Phi(m, m+1) := I$ . Notice that  $\Phi(m, m_0)$  is also row stochastic since it is a product of row-stochastic matrices. We can then express (14) as

$$\Omega(m+1) = \Phi(m, m_0)\Omega(m_0) - \sum_{\tau=m_0}^m \Phi(m, \tau+1)\bar{\rho}(\tau)G(\tau). \quad (15)$$

Next, we need two key properties of  $\Phi(m, m_0)$  from [49].

**Lemma 1** ([49]). *Suppose Assumption 3 holds. Then for any  $m_0 \geq 0$ , there exists a stochastic vector  $\pi(m_0)$  such that*

$$\lim_{m \rightarrow \infty} \Phi(m, m_0) = \mathbf{1}\pi^T(m_0). \quad (16)$$

In words, Lemma 1 states that the product of the row-stochastic matrices  $Y(m)$  converges to a steady-state matrix whose rows are identical and stochastic. We can also characterize the rate of this convergence; to this end, let  $\psi$  denote the total number of reduced graphs that can be generated from  $\mathcal{G}$ , and define  $\nu := \psi|J'|$ ,  $\mathcal{N}_{max} := \max_{j \in J'} |\mathcal{N}_j|$ , and

$$\mu := 1 - \frac{1}{(2\mathcal{N}_{max} - 2b + 1)^\nu}.$$

**Lemma 2** ([49]). *Suppose Assumption 3 holds. We then have that  $\forall m_0 \geq 0$ ,*

$$\left| [\Phi(m, m_0)]_{ji} - [\pi(m_0)]_i \right| \leq \mu^{\left(\frac{m-m_0+1}{\nu}\right)}. \quad (17)$$

Lemma 2 describes the rate at which the rows of  $\Phi(m, m_0)$  converge to  $\pi(m_0)$ . We now leverage this result and show that the nonfaulty nodes achieve consensus under ByRDIE in the limit of large  $m$ , which translates into  $\bar{r} \rightarrow \infty$  and/or  $T \rightarrow \infty$ . To this end, under the assumption of  $m_0 = 1$ , we have from (15) the following expression:

$$\Omega(m+1) = \Phi(m, 1)\Omega(1) - \sum_{\tau=1}^m \Phi(m, \tau+1)\bar{\rho}(\tau)G(\tau). \quad (18)$$

Next, suppose the nonfaulty nodes stop computing local gradients at time step  $m$  and use  $G(m+m') = 0$  for  $m' \geq 0$ . Then, defining  $V(m) := \lim_{m' \rightarrow \infty} \Omega(m+m'+1)$ , we obtain:

$$\begin{aligned} V(m) &= \lim_{m' \rightarrow \infty} \Phi(m+m', 1)\Omega(1) \\ &\quad - \lim_{m' \rightarrow \infty} \sum_{\tau=1}^{m+m'} \Phi(m+m', \tau)\bar{\rho}(\tau)G(\tau) \\ &= \mathbf{1}\pi^T(1)\Omega(1) - \sum_{\tau=1}^{m-1} \mathbf{1}\pi^T(\tau)\bar{\rho}(\tau)G(\tau). \end{aligned} \quad (19)$$

Notice from (19) that all elements in the vector  $V(m) \in \mathbb{R}^{|J'|}$  are identical. Recall that  $V(m)$  is obtained by looking at only one dimension  $k$  of the optimization variable in ByRDIE; in the following, we use  $v^k(m)$  to denote the identical elements of  $V(m)$  corresponding to dimension  $k$ . We then have the following result concerning nonfaulty nodes.

**Theorem 1** (Consensus Behavior of ByRDIE). *Let Assumptions 1 and 3 hold and fix  $\bar{m} = \bar{r}T + 1$ . Then,*

$$\forall j \in J', \forall k \in \{1, \dots, P\}, \left[ w_j^{\bar{r}, T} \right]_k \rightarrow v^k(\bar{m}) \quad (20)$$

as  $\bar{r} \rightarrow \infty$  and/or  $T \rightarrow \infty$ .

Theorem 1, whose proof is given in Appendix A, establishes consensus at the nonfaulty nodes under ByRDIE when  $\bar{r} \rightarrow \infty$  and/or  $T \rightarrow \infty$ . We conclude this discussion by also stating the rate at which the iterates of ByRDIE achieve consensus. To this end, we define the consensus vector  $\bar{V}(r) \in \mathbb{R}^P$  in iteration  $r$  as  $[\bar{V}(r)]_k := v^k(r)$ . To keep the notation simple, we limit ourselves to  $T = 1$  and use  $w_j^r$  to denote  $w_j^{r,1}$ . Nonetheless, a similar result holds for other values of  $T$ .

**Theorem 2** (Consensus Rate for ByRDIE). *Let Assumptions 1 and 3 hold. Then, fixing  $T = 1$ , the iterates of ByRDIE satisfy:*

$$\forall j \in J', \|w_j^r - \bar{V}(r)\| = \mathcal{O}\left(\sqrt{P}\rho(r)\right), \quad (21)$$

where  $\bar{V}(r)$  denotes the consensus vector in iteration  $r$ .

Theorem 2, which is a straightforward consequence of the proof of Theorem 1 (cf. (40)), guarantees a sublinear rate for consensus; indeed, choosing the stepsize evolution to be  $\rho(r) = \mathcal{O}(1/r)$  gives us  $\|w_j^r - \bar{V}(r)\| = \mathcal{O}\left(\sqrt{P}/r\right)$ .

### C. Theoretical Guarantees: Convergence for $T \rightarrow \infty$

We now move to the second (and perhaps the most important) claim of this paper. This involves showing that the output of ByRDIE converges in probability to the minimizer (and minimum) of the statistical risk (cf. (2)) for two extreme cases: *Case I*:  $T \rightarrow \infty$  and *Case II*:  $T = 1$ . We start our discussion with the case of  $T \rightarrow \infty$ , in which case an auxiliary lemma simply follows from [13, Theorem 2] (also, see [30]).

**Lemma 3.** *Let Assumptions 1 and 3 hold, and let the  $k$ -th subproblem of the coordinate descent loop in iteration  $r$  of ByRDIE be initialized with some  $\{w_j\}_{j \in J'}$ . Then, as  $T \rightarrow \infty$ ,*

$$\forall j \in J', \left[ w_j^{r,T} \right]_k \rightarrow \arg \min_{w' \in \mathbb{R}} \sum_{j \in J'} \alpha_j(r, k) \hat{f}(w_j |_{[w_j]_k = w'}, S_j)$$

for some  $\alpha_j(r, k) \geq 0$  such that  $\sum_{j \in J'} \alpha_j(r, k) = 1$ .

Lemma 3 shows that each subproblem of the coordinate descent loop in ByRDIE under Case I converges to the minimizer of some convex combination of local empirical risk functions of the nonfaulty nodes with respect to each coordinate. In addition, Lemma 3 guarantees that consensus is achieved among the nonfaulty nodes at the end of each coordinate descent loop under Case I. Note that while this fact is already known from Theorems 1 and 2, Lemma 3 helps characterize the consensus point. In summary, when nonfaulty nodes begin a coordinate descent subproblem with identical local estimates and  $T \rightarrow \infty$ , they are guaranteed to begin the next subproblem with identical local estimates.

We now fix  $(r, k)$  and use  $\tilde{w}_k^r$  to denote the identical initial local estimates at nonfaulty nodes at the beginning of  $k$ -th subproblem of the coordinate descent loop in the  $r$ -th iteration of ByRDIE under Case I. Next, we define  $h_k^r(w')$  and  $H_k^r(w')$  for  $w' \in \mathbb{R}$  as

$$h_k^r(w') := \mathbb{E}[f(\tilde{w}_k^r |_{[\tilde{w}_k^r]_k = w'}, (x, y))], \quad \text{and} \quad (22)$$

$$H_k^r(w') := \sum_{j \in J'} \alpha_j(r, k) \hat{f}(\tilde{w}_k^r |_{[\tilde{w}_k^r]_k = w'}, S_j) \quad (23)$$

for some  $\alpha_j(r, k) \geq 0$  such that  $\sum_{j \in J'} \alpha_j(r, k) = 1$ . Note that  $h_k^r(\cdot)$  is strictly convex and Lipschitz continuous. Now for fixed  $r$  and  $k$ , define

$$w^* := \arg \min_{w' \in \mathbb{R}} h_k^r(w'), \quad \text{and} \quad (24)$$

$$\hat{w} := \arg \min_{w' \in \mathbb{R}} H_k^r(w'). \quad (25)$$

*Remark 4.* It should be evident to the reader from (22) and (24) that the univariate stochastic function  $h_k^r(w')$  depends on  $\tilde{w}_k^r$  and its (scalar-valued) minimizer  $w^*$ , which should not be confused with the vector-valued statistical minimizer  $w^*$  in (2), is a function of  $r$  and  $k$ . Similarly, it should be obvious from (23) and (25) that  $H_k^r(w')$  depends on  $\tilde{w}_k^r$  and  $\{\alpha_j(r, k)\}_{j \in J'}$ , while its minimizer  $\hat{w}$  is also a function of  $r$  and  $k$ . We are dropping these explicit dependencies here for ease of notation.

In words, if one were to solve the statistical risk minimization problem (2) using (centralized) coordinate descent then  $w^*$  will be the  $k$ -th component of the output of coordinate descent after update of each coordinate  $k$  in every iteration

$r$ . In contrast,  $\hat{w}$  is the  $k$ -th component of the outputs of ByRDIE after update of each coordinate  $k$  in every iteration  $r$  (cf. Lemma 3). While there exist works that relate the empirical risk minimizers to the statistical risk minimizers (see, e.g., [47]), such works are not directly applicable here because of the fact that  $H_k^r(w')$  in this paper changes from one pair  $(r, k)$  of indices to the next. Nonetheless, we can provide the following uniform statistical convergence result for ByRDIE under Case I that relates the empirical minimizers  $\{\hat{w}\}$  to the statistical minimizers  $\{w^*\}$ .

**Theorem 3** (Statistical Convergence Rate for ByRDIE). *Let  $P$  and  $|J'|$  be fixed,  $\bar{r}$  be any (arbitrarily large) positive integer, and  $\{\alpha_j(r, k) \geq 0, j \in J'\}_{r, k=1}^{\bar{r}, P}$  be any arbitrary collection satisfying  $\sum_{j \in J'} \alpha_j(r, k) = 1$ . Let  $|\hat{w}| \leq \Gamma$ ,  $|w^*| \leq \Gamma$ , and  $\{\tilde{w}_k^r\}_{r, k} \subset W$ , and define  $\bar{a} := \max_{(r, k)} \sqrt{\sum_{j \in J'} \alpha_j^2(r, k)}$ . Then, as long as Assumptions 1 and 2 hold, we have  $\forall \epsilon > 0$*

$$\sup_{r, k} [h_k^r(\hat{w}) - h_k^r(w^*)] < \epsilon \quad (26)$$

with probability exceeding

$$1 - 2 \exp \left( -\frac{4|J'|N\epsilon^2}{c_1^2|J'|\bar{a}^2 + \epsilon^2} + |J'| \log \left( \frac{c_2}{\epsilon} \right) + P \log \left( \frac{c_3}{\epsilon} \right) \right), \quad (27)$$

where  $c_1 := 8C$ ,  $c_2 := 24C|J'|$ , and  $c_3 := 24LTP$ .

The proof of this theorem is provided in Appendix B. In words, ignoring minor technicalities that are resolved in Theorem 4 in the following, Theorem 3 states that the coordinate-wise outputs  $\{\hat{w}\}$  of ByRDIE for *all*  $(r, k)$  under Case I achieve, with high probability, almost the same statistical risk as that obtained using the corresponding coordinate-wise statistical risk minimizers  $\{w^*\}$ . We now leverage this result to prove that the iterates of ByRDIE at individual nodes achieve statistical risk that converges to the minimum statistical risk achieved by the statistical risk minimizer (vector)  $w^*$  (cf. (2)).

**Theorem 4** (Convergence Behavior of ByRDIE). *Let Assumptions 1–3 hold. Then,  $\forall j \in J', \forall \epsilon > 0$ , and  $T \rightarrow \infty$ , we have*

$$\lim_{\bar{r} \rightarrow \infty} \left[ \mathbb{E}[f(w_j^{\bar{r}, T}, (x, y))] - \mathbb{E}[f(w^*, (x, y))] \right] < \epsilon \quad (28)$$

with probability exceeding

$$1 - 2 \exp \left( -\frac{4|J'|N\epsilon^2}{c_1'^2|J'|\bar{a}^2 + \epsilon^2} + |J'| \log \left( \frac{c_2'}{\epsilon} \right) + P \log \left( \frac{c_3'}{\epsilon} \right) \right), \quad (29)$$

where  $c_1' := c_1 c_4$ ,  $c_2' := c_2 c_4$ , and  $c_3' := c_3 c_4$  for  $c_4 := 2PL\Gamma$ , and  $(\bar{a}, c_1, c_2, c_3)$  are as defined in Theorem 3.

The proof of this theorem is given in Appendix C. We now make a couple of remarks concerning Theorem 4. First, note that the uniqueness of the minimum of strictly convex functions coupled with the statement of Theorem 4 guarantee that  $\forall j \in J', w_j^{\bar{r}, T} \rightarrow w^*$  with high probability.

Second, Theorem 4 helps crystallize the advantages of distributed learning over *local learning*, in which nodes individually solve the empirical risk minimization problem using only their local data samples. Prior work on stochastic convex

optimization (see, e.g., [47, Theorem 5 and (11)]) tells us that, with high probability and ignoring the log terms, the gap between the statistical risk achieved by the empirical risk minimizer and the statistical risk minimizer scales as  $\mathcal{O}(1/\sqrt{\# \text{ of samples}})$  in the centralized setting. This *learning rate* translates into  $\mathcal{O}(1/\sqrt{N})$  for local learning and  $\mathcal{O}(1/\sqrt{MN})$  for the idealized centralized learning. In contrast, Theorem 4 can be interpreted as resulting in the following *distributed learning rate* (with high probability):<sup>3</sup>

$$\mathbb{E}[f(w_j^{\bar{r},T}, (x, y))] - \mathbb{E}[f(w^*, (x, y))] = \mathcal{O}\left(\frac{1}{\sqrt{N_{\text{eff}}}}\right), \quad (30)$$

where  $N_{\text{eff}} := N/\bar{a}^2$  denotes the *effective* number of training samples available during distributed learning.

In order to understand the significance of (30), notice that

$$\frac{1}{M} \leq \frac{1}{|J'|} \leq \bar{a}^2 \leq 1 \Rightarrow N \leq N_{\text{eff}} \leq |J'|N \leq NM. \quad (31)$$

In particular,  $\bar{a}^2 = 1/M$  if and only if there are no Byzantine failures in the network, resulting in the coordinate descent-based distributed learning of  $\mathcal{O}(1/\sqrt{NM})$ , which matches the centralized learning rate. (This, to the best of our knowledge, is the first result on the explicit learning rate of coordinate descent-based distributed learning.) In the presence of Byzantine nodes, however, the *maximum* number of *trustworthy* samples in the network is  $|J'|N$ , and (30) and (31) tell us that the learning rate of ByRDIE in this scenario will be somewhere between the idealized learning rate of  $\mathcal{O}(1/\sqrt{|J'|N})$  and the local learning rate of  $\mathcal{O}(1/\sqrt{N})$ .

Our discussion so far has focused on the rate of statistical convergence (i.e., learning rate) of ByRDIE. The proof of Theorem 4, however, also contains within itself the algorithmic rate of convergence for ByRDIE. We state this convergence rate in terms of the following theorem, which uses the notation  $\bar{f}^0$  to denote the starting statistical risk  $\mathbb{E}[f(0, (x, y))]$ ,  $\bar{f}^*$  to denote the minimum statistical risk  $\mathbb{E}[f(w^*, (x, y))]$  and  $1 - \delta(\epsilon, N, \bar{a})$  to express the probability expression in (29).

**Theorem 5** (Algorithmic Convergence Rate for ByRDIE). *Let Assumptions 1–3 hold. Then,  $\forall j \in J', \forall r \in \mathbb{N}, \forall \epsilon > 0$ , and  $T \rightarrow \infty$ , we get with probability exceeding  $1 - \delta(\epsilon, N, \bar{a})$  that*

$$\mathbb{E}\left[f(w_j^{r,T}, (x, y))\right] - \bar{f}^* < \max\left\{\left(\bar{f}^0 - \bar{f}^*\right)\left(1 - \frac{r\epsilon}{c_5}\right), \epsilon\right\}, \quad (32)$$

where the parameter  $\delta(\epsilon, N, \bar{a})$  is given by

$$2 \exp\left(-\frac{4|J'|N\epsilon^2}{c_1^2|J'|\bar{a}^2 + \epsilon^2} + |J'| \log\left(\frac{c_2'}{\epsilon}\right) + P \log\left(\frac{c_3'}{\epsilon}\right)\right), \quad (33)$$

while  $c_5 := c_4 L' \sqrt{P} \gamma^*$ , and the constants  $c_1', c_2', c_3'$ , and  $c_4$  are as defined in Theorem 4.

The proof of this theorem is given in Appendix D. It can be seen from Theorem 5 that, with high probability, ByRDIE

<sup>3</sup>We are once again ignoring the log terms in our discussion; it can be checked, however, that the log terms resulting from Theorem 4 match the ones in prior works such as [47] on centralized learning.

requires  $r = \mathcal{O}(1/\epsilon)$  iterations to bring the *excess risk*  $\mathbb{E}[f(w_j^{r,T}, (x, y))] - \bar{f}^*$  down to  $\epsilon$ . In terms of minimization of the statistical risk, therefore, ByRDIE achieves a sublinear rate of algorithmic convergence with high probability, even in the presence of Byzantine failures in the network.

#### D. Theoretical Guarantees: Convergence for $T = 1$

Case I for ByRDIE, in which  $T \rightarrow \infty$ , is akin to doing exact line search during minimization of each coordinate, which is one of the classic ways of implementing coordinate descent. Another well-adopted way of performing coordinate descent is to take only one step in the direction of descent in a dimension and then switch to another dimension [45]. Within the context of ByRDIE, this is equivalent to setting  $T = 1$  (Case II); our goal here is to provide convergence guarantees for ByRDIE in this case. Our analysis in this section uses the compact notation  $w_j^r := w_j^{r,1} \equiv w_j^{r+1}(1)$ . In the interest of space, and since the probabilistic analysis of this section is similar to that of the previous section, our probability results are stated asymptotically here, rather than in terms of precise bounds.

The starting point of our discussion is Theorem 1. Recall from Sec. III-B the definition of the index  $m := (r-1)T + t$ . When  $T = 1$ , we have  $r = m$  and therefore  $[w_j^r]_k \rightarrow v^k(r)$  as  $r \rightarrow \infty$  according to Theorem 1. In order to provide convergence guarantee, we only need to show that  $v^k(r) \xrightarrow{r, N} [w^*]_k$  in probability. To this end, we define a sequence  $Q(q) \in \mathbb{R}^P$  as follows:  $\forall k, [Q(1)]_k := v^k(1)$ , while for the parameterized index  $q = (r-1)P + k \neq 1$ ,  $Q(q)$  is obtained by replacing  $[Q(q-1)]_k$  with  $v^k(r)$  and keeping the other dimensions fixed. Similarly, we define a sequence  $\eta(q)$  satisfying  $\eta(q) = \rho(r)$  for  $Pr \leq q < (P+1)r$ . Notice that  $0 < \eta(q+1) \leq \eta(q)$ ,  $\sum_q \eta(q) = \infty$  and  $\sum_q \eta^2(q) < \infty$ . Since we have from (19) that  $v^k(r+1) = v^k(r) - \rho(r) \sum_{i=1}^{|J'|} [\pi(r+1)]_i [\nabla \hat{f}(w_i^r, S_i)]_k$ , we can write the following iteration:

$$Q(q+1) = Q(q) - \eta(q) \sum_{i=1}^{|J'|} [\pi(r+1)]_i [\nabla \hat{f}(w_i^r, S_i)]_k e_k, \quad (34)$$

where  $e_k$  denotes the standard basis vector (i.e., it is zero in every dimension except  $k$  and  $[e_k]_k = 1$ ) and the relationship between  $r, k$ , and  $q$  is as defined earlier. The sequence  $Q(q)$  effectively helps capture update of the optimization variable after each coordinate-wise update of ByRDIE. In particular, we have the following result concerning the sequence  $Q(q)$ .

**Lemma 4.** *Let Assumptions 1, 2, and 3 hold for ByRDIE and choose  $T = 1$ . We then have that  $Q(q) \xrightarrow{q, N} w^*$  in probability.*

The proof of this lemma is provided in Appendix E. We are now ready to state the convergence result for ByRDIE under Case II (i.e,  $T = 1$ ).

**Theorem 6** (Asymptotic Convergence of ByRDIE). *Let Assumptions 1, 2, and 3 hold for ByRDIE and choose  $T = 1$ . Then,  $\forall j \in J', w_j^{\bar{r}, N} \xrightarrow{\bar{r}, N} w^*$  in probability.*

*Proof.* We have from Theorem 1 that  $\forall j \in J', [w_j^{\bar{r}}]_k \xrightarrow{\bar{r}} v^k(\bar{r})$  for all  $k \in \{1, 2, \dots, P\}$ . The definition of  $Q(q)$  along with



Lemma 4 also implies that  $v^k(\bar{r}) \xrightarrow{\bar{r}, N} [w^*]_k$  in probability. This completes the proof of the theorem. ■

### E. How to Choose the Parameter $T$ in ByRDIE?

The parameter  $T$  in ByRDIE trades off consensus among the nonfaulty nodes and the convergence rate as a function of the number of communication iterations  $t_c(r, k, t)$ , defined as

$$t_c := (r - 1)TP + [(k - 1)T + t]. \quad (35)$$

In particular, given a fixed  $T$ , each iteration  $r$  of ByRDIE involves  $TP$  (scalar-valued) communication exchanges among the neighboring nodes. In the previous section, we have provided theoretical guarantees for the two extreme cases of  $T \rightarrow \infty$  and  $T = 1$ . In the limit of large  $\bar{r}$ , our results establish that both extremes result in consensus and convergence to the statistical risk minimizer. In practice, however, different choices of  $T$  result in different behaviors as a function of  $t_c$  ( $\equiv t_c(r, k, t)$ ), as discussed in the following and as illustrated in our numerical experiments in the next section.

When  $T$  is large, the two time-scale nature of ByRDIE ensures the disagreement between nonfaulty nodes does not become too large in the initial stages of the algorithm; in particular, the larger the number of iterations  $T$  in the inner loop, the smaller the disagreement among the nonfaulty nodes at the beginning of the algorithm. Nonetheless, this comes at the expense of slower convergence to the desired minimizer as a function of the number of communication iterations  $t_c$ .

On the other hand, while choosing  $T = 1$  also guarantees consensus among nonfaulty nodes, it only does so asymptotically (cf. Theorem 1). Stated differently, ByRDIE cannot guarantee in this case that the disagreement between nonfaulty nodes will be small in the initial stages of the algorithm (cf. Theorem 2). This tradeoff between small consensus error and slower convergence (as a function of communication iterations  $t_c$ ) should be considered by a practitioner when deciding the value of  $T$ . We conclude by noting that the different nature of the two extreme cases requires markedly different proof techniques, which should be of independent interest to researchers.

*Remark 5.* Our discussion so far has focused on the use of a static parameter  $T$  within ByRDIE. Nonetheless, it is plausible that one could achieve somewhat better tradeoffs between consensus and convergence through the use of an adaptive parameter  $T_r$  in lieu of  $T$  that starts with a large value and gradually decreases as  $r$  increases. Careful analysis and investigation of such an adaptive two-time scale variant of ByRDIE, however, is beyond the scope of this paper.

## IV. NUMERICAL RESULTS

In this section, we validate our theoretical results and make various observations about the performance of ByRDIE using two sets of numerical experiments. The first set of experiments involves learning of a binary classification model from the infinite MNIST dataset<sup>4</sup> that is distributed across a network of




Fig. 1: Average classification accuracy of ByRDIE, local CD, and DGD (on test data) for different values of  $N$  as a function of the number of algorithmic iterations for distributed training of a linear SVM on (infinite) MNIST dataset.

nodes. This set of experiments fully satisfies all the assumptions in the theoretical analysis of ByRDIE. The second set of experiments involves training of a small-scale neural network for classification of the *Iris* dataset [50] distributed across a network. The learning problem in this case corresponds to a nonconvex one, which means this set of experiments does not satisfy the main assumptions of our theorems. Nonetheless, we show in the following that ByRDIE continues to perform well in such distributed nonconvex learning problems.

### A. Distributed SVM Using Infinite MNIST Dataset

We consider a distributed linear binary classification problem involving MNIST handwritten digits dataset. The (infinite) MNIST dataset comprises images of handwritten digits from ‘0’ to ‘9’. Since our goal is demonstration of the usefulness of ByRDIE in the presence of Byzantine failures, we focus only on distributed training of a linear support vector machine (SVM) for classification between digits ‘5’ and ‘8’, which tend to be the two most inseparable digits. In addition to highlighting the robustness of ByRDIE against Byzantine failures in this problem setting, we evaluate its performance for different choices of the parameters  $T$ ,  $N$ , and  $b$ .

In terms of the experimental setup, we generate Erdős–Rényi networks ( $p = 0.5$ ) of  $M$  nodes,  $b$  of which are randomly chosen to be Byzantine nodes. All nonfaulty nodes are allocated  $N$  samples—equally divided between the two classes—from the dataset, while a Byzantine node broadcasts random data uniformly distributed between 0 and 1 to its neighbors in each iteration. When running ByRDIE algorithm, each node updates one dimension  $T$  times before proceeding to the next dimension. All tests are performed on the same test set with 1000 samples of digits ‘5’ and ‘8’ each.

We first report results that confirm the idea that ByRDIE can take advantage of cooperation among different nodes to achieve better performance even when there are Byzantine failures in the network. This involves varying the local sample size  $N$  and comparing the classification accuracy on the test data. We generate a network of  $M = 50$  nodes, randomly

<sup>4</sup><https://leon.bottou.org/projects/infimnist>



(a) Classification accuracy for different values of  $T$ (b) Consensus behavior for different values of  $T$ 

Fig. 2: Convergence and consensus behavior of ByRDIE as a function of the number of communication iterations  $t_c$  for different values of the parameter  $T$ . All plots correspond to distributed training of a linear SVM using the MNIST dataset.

pick  $b = 10$  nodes within the network to be Byzantine nodes (20% failures), vary  $N$  from 10 to 30, and average the final set of results over 10 independent (over network, Byzantine nodes, and data allocation) Monte Carlo trials of ByRDIE. The performance of ByRDIE is compared with two approaches: (i) coordinate descent-based training using only local data (local CD); and (ii) distributed gradient descent-based [16] training involving network data (DGD). To achieve the best convergence rate for ByRDIE,  $T$  is chosen to be 1 in these experiments. The final set of results are shown in Fig 1, in which the average classification accuracy is plotted against the number of (scalar-valued) communication iterations for ByRDIE, the number of per-dimension updates for local CD, and the number of (vector-valued) communication iterations for DGD. It can be seen that the performance of local CD is not good enough due to the small local sample size. On the other hand, when trying to improve performance by cooperating among different nodes, DGD fails for lack of robustness against Byzantine failures. In contrast, the higher accuracy of ByRDIE shows that ByRDIE can take advantage of the larger



(a) Classification accuracy for different values of  $b$ (b) Consensus behavior for different values of  $b$ 

Fig. 3: Convergence and consensus behavior of ByRDIE within a small network ( $M = 20$ ) as a function of the number of communication iterations  $t_c$  for different number  $b$  of Byzantine nodes in the network. All plots correspond to distributed training of a linear SVM using the MNIST dataset.

distributed dataset while being Byzantine resilient.

Next, we investigate the impact of different values of  $T$  in ByRDIE on the tradeoff between consensus and convergence rate. This involves generating a network of  $M = 50$  nodes that includes randomly placed  $b = 5$  Byzantine nodes within the network (10% failures), randomly allocating  $N = 60$  training samples to each nonfaulty node, and averaging the final set of results over 10 independent trials. The corresponding results are reported in Fig. 2 for four different values of  $T$  as a function of the number of communication iterations  $t_c$  in terms of (i) average classification accuracy (Fig. 2a) and (ii) average pairwise distances between local classifiers (Fig. 2b). It can be seen from these figures that  $T = 1$  leads to the fastest convergence in the initial stages of the algorithm, but this fast convergence comes at the expense of the largest differences among local classifiers, especially in the beginning of the algorithm. In contrast, while  $T = 4$  results in the slowest convergence, it ensures closeness of the local classifiers at all stages of the algorithm.

Finally, we investigate the impact of different values of  $b$

(the actual number of Byzantine nodes) on the performance of ByRDIE. In order to amplify this impact, we focus on a smaller network of  $M = 20$  nodes with a small number of  $N = 10$  samples per node and  $T = 3$ . Assuming resilience against the worst-case Byzantine scenario (cf. Remark 3), ByRDIE requires the neighborhood of each node to have at least  $(2b+1)$  nodes. Under this constraint, we find that  $b \geq 5$  in this setting, which translates into 25% or more of the nodes as being Byzantine, often renders ByRDIE unusable.<sup>5</sup> We therefore report our results in terms of both consensus and convergence behavior by varying  $b$  from 1 to 4. The final results, averaged over 10 independent trials, are shown in Fig. 3. It can be seen from these figures that both the classification accuracy (Fig. 3a) and the consensus performance (Fig. 3b) of ByRDIE suffer as  $b$  increases from 1 to 4. Such behavior, however, is in line with our theoretical developments. First, as  $b$  increases, the *post-screening* graph becomes sparser, which slows down information diffusion and consensus. Second, as  $b$  increases, fewer samples are incorporated into distributed learning, which limits the final classification accuracy of distributed SVM.

### B. Distributed Neural Network Using Iris Dataset

While the theoretical guarantees for ByRDIE have been developed for convex learning problems, we now demonstrate the usefulness of ByRDIE in distributed nonconvex learning problems. Our setup in this regard involves distributed training of a small-scale neural network using the three-class Iris dataset. This dataset consists of 50 samples each from three species of irises, with each sample being a four-dimensional feature vector. While one of the species in this data is known to be linearly separable from the other two, the remaining two species cannot be linearly separated. We employ a fully connected three-layer neural network for classification of this dataset, with a four-neuron input layer, a three-neuron hidden layer utilizing the *rectified linear unit* (ReLU) activation function and a three-neuron output layer using the softmax function for classification. The distributed setup corresponds to a random network of  $M = 10$  nodes with one Byzantine node ( $b = 1$ ), in which each nonfaulty node is allocated  $N = 15$  samples that are equally divided between the three classes.

Algorithm	Iterations to achieve 95% accuracy	Consensus
DGD	$\infty$	No
ByRDIE	19	Yes
Centralized CD	17	N/A

TABLE I: Distributed training of a neural network in the presence of Byzantine failures using ByRDIE and DGD.

Since distributed training of this neural network requires solving a distributed nonconvex problem, our theoretical guarantees for ByRDIE do not hold in this setting. Still, we use ByRDIE with  $T = 1$  to train the neural network for a total of 200 independent trials with independent random initializations. Simultaneously, we use DGD for distributed training and also train the neural network in a centralized setting using CD for

<sup>5</sup>As noted in Remark 3, one could overcome this limit by opting instead for resilience against the average-case scenario and replacing  $b$  with  $b_{\text{ave}}$ .

comparison purposes. The final results are reported in Table I, which lists the average number of *outer* iterations needed by each training algorithm to achieve 95% classification accuracy. It can be seen from these results that while DGD completely breaks down in the presence of a *single* Byzantine node, ByRDIE continues to be resilient to Byzantine failures in distributed nonconvex learning problems and comes close to matching the performance of centralized CD in this case.

## V. CONCLUSION

In this paper, we have introduced a coordinate descent-based distributed algorithm, termed ByRDIE, that can carry out distributed learning tasks in the presence of Byzantine failures in the network. The proposed algorithm pursues the minimizer of the statistical risk in a fully distributed environment and theoretical results guarantee that ByRDIE achieves this objective under mild assumptions on the risk function and the network topology. In addition, numerical results presented in the paper validate resilience of ByRDIE against Byzantine failures in the network for distributed convex and nonconvex learning tasks. Future works in this direction include strengthening the results in terms of almost sure convergence, deriving explicit rates of convergence, and extensions of theoretical analysis to constant step size and nonconvex risk functions.

## APPENDIX A PROOF OF THEOREM 1

Fix any dimension  $k \in \{1, \dots, P\}$  and recall from (19) that

$$\begin{aligned} V(m+1) &= \mathbf{1}\pi^T(1)\Omega(1) \\ &\quad - \sum_{\tau=1}^{m-1} \mathbf{1}\pi^T(\tau+1)\bar{\rho}(\tau)G(\tau) - \mathbf{1}\pi^T(m+1)\bar{\rho}(m)G(m). \end{aligned} \quad (36)$$

Since  $m := (r-1)T + t$  and  $[\Omega(m)]_j := [w_j^r(t)]_k$ , we get

$$\begin{aligned} v^k(m+1) &= \sum_{i=1}^{|J'|} [\pi(1)]_i [w_i^1(1)]_k - \sum_{\tau=1}^{m-1} \bar{\rho}(\tau) \sum_{i=1}^{|J'|} [\pi(\tau+1)]_i [G(\tau)]_i \\ &\quad - \bar{\rho}(m) \sum_{i=1}^{|J'|} [\pi(r+1)]_i [G(m)]_i. \end{aligned} \quad (37)$$

We also get from (18) that

$$\begin{aligned} [w_j^r(t+1)]_k &= \sum_{i=1}^{|J'|} [\Phi(m, 1)]_{ji} [w_i^1(1)]_k \\ &\quad - \sum_{\tau=1}^{m-1} \bar{\rho}(\tau) \sum_{i=1}^{|J'|} [\Phi(m, \tau+1)]_{ji} [G(\tau)]_i - \bar{\rho}(m) [G(m)]_j, \end{aligned} \quad (38)$$

where we have used the fact that  $\Phi(m, m+1) = I$ . Thus,

$$\begin{aligned} & \left| [w_j^r(t+1)]_k - v^k(m+1) \right| = \\ & \left| \sum_{i=1}^{|J'|} ([\Phi(m, 1)]_{ji} - [\pi(1)]_i) [w_i^1(1)]_k \right. \\ & \quad - \sum_{\tau=1}^{m-1} \bar{\rho}(\tau) \sum_{i=1}^{|J'|} ([\Phi(m, \tau+1)]_{ji} - [\pi(\tau+1)]_i) [G(\tau)]_i \\ & \quad \left. - \bar{\rho}(m) \sum_{i=1}^{|J'|} [\pi(m+1)]_i ([G(m)]_j - [G(m)]_i) \right|, \quad (39) \end{aligned}$$

where the last summation follows from the observation that  $\sum_{i=1}^{|J'|} [\pi(m+1)]_i = 1$ . Further, Assumption 1 implies there exists a coordinate-wise bound  $L_\nabla \leq L'$  such that  $\forall i, m, |[G(m)]_i| \leq L_\nabla$ . Using this and Lemma 2, we obtain

$$\begin{aligned} & \left| [w_j^r(t+1)]_k - v^k(m+1) \right| \\ & \leq \sum_{i=1}^{|J'|} \left| [\Phi(m, 1)]_{ji} - [\pi(1)]_i \right| \left| [w_i^1(1)]_k \right| \\ & \quad + \sum_{\tau=1}^{m-1} \bar{\rho}(\tau) \sum_{i=1}^{|J'|} \left| [\Phi(m, \tau+1)]_{ji} - [\pi(\tau+1)]_i \right| \left| [G(\tau)]_i \right| \\ & \quad + \bar{\rho}(m) \sum_{i=1}^{|J'|} [\pi(m+1)]_i \left| [G(m)]_j - [G(m)]_i \right| \\ & \leq |J'| L_\nabla \sum_{\tau=1}^{m-1} \bar{\rho}(\tau) \mu^{\frac{m-\tau}{\nu}} + 2\bar{\rho}(m) L_\nabla. \quad (40) \end{aligned}$$

Note that the last inequality in the above expression exploits the fact that  $[w_i^1(1)]_k \equiv 0$ . In the case of an arbitrary initialization of ByRDIE, however, we could have bounded the first term in the second inequality in (40) as  $|J'| \Gamma \mu^{\frac{m+1}{\nu}}$ , which still goes to zero as  $m \rightarrow \infty$ .

We now expand the term  $\sum_{\tau=1}^{m-1} \bar{\rho}(\tau) \mu^{\frac{m-\tau}{\nu}}$  in (40) as<sup>6</sup>

$$\begin{aligned} \sum_{\tau=1}^{m-1} \bar{\rho}(\tau) \mu^{\frac{m-\tau}{\nu}} &= \left( \sum_{\tau=1}^{\frac{m}{2}-1} \bar{\rho}(\tau) \mu^{\frac{m-\tau}{\nu}} + \sum_{\tau=\frac{m}{2}}^{m-1} \bar{\rho}(\tau) \mu^{\frac{m-\tau}{\nu}} \right) \\ &\leq \bar{\rho}(1) \sum_{\tau=1}^{\frac{m}{2}-1} \mu^{\frac{m-\tau}{\nu}} + \bar{\rho}\left(\frac{r}{2}\right) \sum_{\tau=\frac{m}{2}}^{m-1} \mu^{\frac{m-\tau}{\nu}} \\ &= \frac{\bar{\rho}(1) \mu^{\frac{m}{2}+1} (1 - \mu^{\frac{m}{2\nu}})}{\mu^{1-\frac{1}{\nu}}} + \frac{\bar{\rho}\left(\frac{m}{2}\right) \mu^{\frac{1}{\nu}} (1 - \mu^{\frac{m}{2\nu}})}{\mu^{1-\frac{1}{\nu}}}. \end{aligned}$$

It can be seen from the above expression that  $\lim_{m \rightarrow \infty} \sum_{\tau=1}^{m-1} \bar{\rho}(\tau) \mu^{\frac{m-\tau}{\nu}} \leq 0$ . Since  $\sum_{\tau=1}^{m-1} \bar{\rho}(\tau) \mu^{\frac{m-\tau}{\nu}} \geq 0$ , it follows that  $\lim_{m \rightarrow \infty} \sum_{\tau=1}^{m-1} \bar{\rho}(\tau) \mu^{\frac{m-\tau}{\nu}} = 0$ . This fact in concert with (40) and the diminishing nature of  $\bar{\rho}(m)$  give

$$\lim_{m \rightarrow \infty} \left| [w_j^r(t+1)]_k - v^k(m+1) \right| = 0. \quad (41)$$

Next, take  $r = \bar{r}$ ,  $t = T$ , and note that  $w_j^r(t+1) = w_j^{\bar{r}, T}$  and  $\bar{m} = \bar{r}T + 1 = [(r-1)T + T] + 1 = m + 1$  in this case.

<sup>6</sup>We focus here only on the case of an even  $m$  for the sake of brevity; the expansion for an odd  $m$  follows in a similar fashion.

Further,  $\bar{m} \rightarrow \infty$  when  $\bar{r} \rightarrow \infty$  and/or  $T \rightarrow \infty$ . It therefore follows from (41) that  $[w_j^{\bar{r}, T}]_k \rightarrow v^k(\bar{m})$  as  $\bar{r} \rightarrow \infty$  and/or  $T \rightarrow \infty$ . Finally, since  $k$  in our analysis was arbitrary, the same holds for all  $k = 1, 2, \dots, P$ . ■

## APPENDIX B

### PROOF OF THEOREM 3

We begin by fixing  $w' \in \mathbb{R}$  such that  $|w'| \leq \Gamma$  and defining  $\alpha_k^r \in \mathbb{R}^{|J'|}$  as  $\alpha_k^r := [\alpha_j(r, k) : j \in J']$  and  $\tilde{F}_k^r(w') \in \mathbb{R}^{|J'|}$  as  $\tilde{F}_k^r(w') := [\tilde{f}(\tilde{w}_k^r |_{[\tilde{w}_k^r]_k = w'}, S_j) : j \in J']$ . We can then write

$$H_k^r(w') = \alpha_k^{rT} \tilde{F}_k^r(w'). \quad (42)$$

Next, notice from (3), (22), and (23) that

$$\forall (r, k), \mathbb{E}[H_k^r(w')] = h_k^r(w'). \quad (43)$$

We now fix indices  $(r, k)$  and note that the random variables  $\alpha_j(r, k) \ell(\tilde{w}_k^r |_{[\tilde{w}_k^r]_k = w'}, (x_{jn}, y_{jn}))$  involved in the definition of the univariate function  $H_k^r(w')$  are (i) independent due to the independence of the training samples, and (ii) bounded as  $0 \leq \alpha_j \ell(w, (x_{jn}, y_{jn})) \leq \alpha_j C$  due to Assumption 2. Therefore, the following holds  $\forall \epsilon' > 0$  due to Hoeffding's inequality [51]:

$$\begin{aligned} & \mathbb{P} \left( \left| \alpha_k^{rT} \tilde{F}_k^r(w') - h_k^r(w') \right| \geq \epsilon' \right) = \\ & \mathbb{P} \left( |H_k^r(w') - h_k^r(w')| \geq \epsilon' \right) \leq 2 \exp \left( -\frac{2N\epsilon'^2}{C^2 \|\alpha_k^r\|^2} \right). \quad (44) \end{aligned}$$

Further, since the  $|J'|$ -dimensional vector  $\alpha_k^r$  is an arbitrary element of the standard simplex, defined as

$$\Delta := \{v \in \mathbb{R}^{|J'|} : \sum_{j=1}^{|J'|} [v]_j = 1 \text{ and } \forall j, [v]_j \geq 0\}, \quad (45)$$

the probability bound in (44) also holds for any  $v \in \Delta$ , i.e.,

$$\mathbb{P} \left( \left| v^T \tilde{F}_k^r(w') - h_k^r(w') \right| \geq \epsilon' \right) \leq 2 \exp \left( -\frac{2N\epsilon'^2}{C^2 \|v\|^2} \right). \quad (46)$$

We now define the set  $\mathcal{S}_\alpha := \{\alpha_k^r\}_{r, k=1}^{\bar{r}, P}$ . Our next goal is to leverage (46) and derive a probability bound similar to (44) that *uniformly* holds for *all*  $v \in \mathcal{S}_\alpha$ . To this end, let

$$\mathbb{C}_\xi := \{c_1, \dots, c_{d_\xi}\} \subset \Delta \quad \text{s.t.} \quad \mathcal{S}_\alpha \subseteq \bigcup_{q=1}^{d_\xi} B(c_q, \xi) \quad (47)$$

denote an  $\xi$ -covering of  $\mathcal{S}_\alpha$  in terms of the  $\ell_2$  norm and define  $\bar{c} := \arg \max_{c \in \mathbb{C}_\xi} \|c\|$ . It then follows from (46) and the union bound that

$$\begin{aligned} & \mathbb{P} \left( \sup_{c \in \mathbb{C}_\xi} \left| c^T \tilde{F}_k^r(w') - h_k^r(w') \right| \geq \epsilon' \right) \\ & \leq 2d_\xi \exp \left( -\frac{2N\epsilon'^2}{C^2 \|\bar{c}\|^2} \right). \quad (48) \end{aligned}$$

In addition, we have

$$\begin{aligned} \sup_{v \in \mathcal{S}_\alpha} \left| v^T \tilde{F}_k^r(w') - h_k^r(w') \right| &\stackrel{(a)}{\leq} \sup_{c \in \mathbb{C}_\xi} \left| c^T \tilde{F}_k^r(w') - h_k^r(w') \right| + \\ & \sup_{v \in \mathcal{S}_\alpha, c \in \mathbb{C}_\xi} \|v - c\| \|\tilde{F}_k^r(w')\|, \quad (49) \end{aligned}$$

where (a) is due to triangle and Cauchy–Schwarz inequalities. Trivially,  $\sup_{v \in \mathcal{S}_\alpha, c \in \mathbb{C}_\xi} \|v - c\| \leq \xi$  from the definition of  $\mathbb{C}_\xi$ , while  $\|\tilde{F}_k^r(w')\| \leq \sqrt{|J'|}C$  from the definition of  $\tilde{F}_k^r(w')$  and Assumption 2. Combining (48) and (49), we get

$$\begin{aligned} \mathbb{P} \left( \sup_{v \in \mathcal{S}_\alpha} \left| v^T \tilde{F}_k^r(w') - h_k^r(w') \right| \geq \epsilon' + \sqrt{|J'|}\xi C \right) \\ \leq 2d_\xi \exp \left( -\frac{2N\epsilon'^2}{C^2\|\bar{c}\|^2} \right). \end{aligned} \quad (50)$$

We now define  $\bar{\alpha} := \arg \max_{v \in \mathcal{S}_\alpha} \|v\|$ . It can then be shown from the definitions of  $\mathbb{C}_\xi$  and  $\bar{c}$  that

$$\|\bar{c}\|^2 \leq 2(\|\bar{\alpha}\|^2 + \xi^2). \quad (51)$$

Therefore, picking any  $\epsilon'' > 0$ , and defining  $\epsilon' := \epsilon''/2$  and  $\xi := \epsilon''/(2C\sqrt{|J'|})$ , we have from (50) and (51) that

$$\begin{aligned} \mathbb{P} \left( \sup_{v \in \mathcal{S}_\alpha} \left| v^T \tilde{F}_k^r(w') - h_k^r(w') \right| \geq \epsilon'' \right) \\ \leq 2d_\xi \exp \left( -\frac{4|J'|N\epsilon''^2}{4C^2|J'|\|\bar{\alpha}\|^2 + \epsilon''^2} \right). \end{aligned} \quad (52)$$

In order to obtain the desired uniform bound, we next need to remove the dependence on  $\tilde{w}_k^r$  and  $w'$  in (52). To this end, we write  $\tilde{F}_k^r(w')$  and  $h_k^r(w')$  as  $\tilde{F}_k^r(\tilde{w}_k^r, w')$  and  $h_k^r(\tilde{w}_k^r, w')$ , respectively, to highlight their dependence on  $\tilde{w}_k^r$  and  $w'$ . Next, we define

$$\mathbb{U}_\zeta := \{u_1, \dots, u_{m_\zeta}\} \subset W \quad \text{s.t.} \quad W \subseteq \bigcup_{q=1}^{m_\zeta} B(u_q, \zeta) \quad (53)$$

to be a  $\zeta$ -covering of  $W$  in terms of the  $\ell_2$  norm. It then follows from (52) that

$$\begin{aligned} \mathbb{P} \left( \sup_{v \in \mathcal{S}_\alpha, u \in \mathbb{U}_\zeta} \left| v^T \tilde{F}_k^r(u, [u]_k) - h_k^r(u, [u]_k) \right| \geq \epsilon'' \right) \\ \leq 2d_\xi m_\zeta \exp \left( -\frac{4|J'|N\epsilon''^2}{4C^2|J'|\|\bar{\alpha}\|^2 + \epsilon''^2} \right). \end{aligned} \quad (54)$$

Similar to (49), and using notation  $\mathcal{L} := \{w' \in \mathbb{R} : |w'| \leq \Gamma\}$ , we can also write

$$\begin{aligned} \sup_{v \in \mathcal{S}_\alpha, w \in W, w' \in \mathcal{L}} \left| v^T \tilde{F}_k^r(w, w') - h_k^r(w, w') \right| \\ \leq \sup_{u \in \mathbb{U}_\zeta, v \in \mathcal{S}_\alpha} \left| v^T \tilde{F}_k^r(u, [u]_k) - h_k^r(u, [u]_k) \right| + \\ \sup_{u \in \mathbb{U}_\zeta, v \in \mathcal{S}_\alpha, w \in W, w' \in \mathcal{L}} \left[ \left| v^T \tilde{F}_k^r(w, w') - v^T \tilde{F}_k^r(u, [u]_k) \right| + \right. \\ \left. |h_k^r(u, [u]_k) - h_k^r(w, w')| \right]. \end{aligned} \quad (55)$$

Further, since  $w|_{[u]_k=w'} \in W$  for any  $(w, w') \in W \times \mathcal{L}$ , we have from Assumption 1 and definition of the set  $\mathbb{U}_\zeta$  that

$$\sup_{u, v, w, w'} \left| v^T \tilde{F}_k^r(w, w') - v^T \tilde{F}_k^r(u, [u]_k) \right| \leq L'\zeta, \quad \text{and} \quad (56)$$

$$\sup_{u, v, w, w'} |h_k^r(u, [u]_k) - h_k^r(w, w')| \leq L'\zeta. \quad (57)$$

We now fix  $\epsilon''' > 0$ , and define  $\epsilon'' := \epsilon'''/2$  and  $\zeta := \epsilon'''/4L'$ . We then obtain the following from (52)–(57):

$$\begin{aligned} \mathbb{P} \left( \sup_{v \in \mathcal{S}_\alpha, \{\tilde{w}_k^r\}, w'} |H_k^r(w') - h_k^r(w')| \geq \epsilon''' \right) \\ \leq 2d_\xi m_\zeta \exp \left( -\frac{4|J'|N\epsilon'''^2}{16C^2|J'|\|\bar{\alpha}\|^2 + \epsilon'''^2} \right). \end{aligned} \quad (58)$$

To conclude, let us define the event

$$\mathcal{A}_\epsilon := \left\{ \sup_{v \in \mathcal{S}_\alpha, \{\tilde{w}_k^r\}, w'} |H_k^r(w') - h_k^r(w')| < \frac{\epsilon}{2} \right\} \quad (59)$$

for any  $\epsilon > 0$ . Conditioned on this event, we have

$$\begin{aligned} \forall (r, k), \quad h_k^r(\hat{w}) - h_k^r(w^*) = \underbrace{h_k^r(\hat{w}) - H_k^r(\hat{w})}_{< \epsilon/2} \\ + \underbrace{H_k^r(w^*) - h_k^r(w^*)}_{< \epsilon/2} + \underbrace{H_k^r(\hat{w}) - H_k^r(w^*)}_{\leq 0} < \epsilon. \end{aligned} \quad (60)$$

Therefore, given any  $\epsilon > 0$ , we have from (58)–(60) that

$$\begin{aligned} \mathbb{P} \left( \sup_{r, k} [h_k^r(\hat{w}) - h_k^r(w^*)] \geq \epsilon \right) \\ \leq 2d_\xi m_\zeta \exp \left( -\frac{4|J'|N\epsilon^2}{64C^2|J'|\|\bar{\alpha}\|^2 + \epsilon^2} \right), \end{aligned} \quad (61)$$

where  $\xi = \frac{\epsilon}{8C|J'|}$  and  $\zeta = \frac{\epsilon}{8L'}$ . The proof now follows from (61) and the following facts about the covering numbers of the sets  $\mathcal{S}_\alpha$  and  $W$ : (1) Since  $\mathcal{S}_\alpha$  is a subset of  $\Delta$ , which can be circumscribed by a sphere in  $\mathbb{R}^{|J'|-1}$  of radius  $\sqrt{\frac{|J'|-1}{|J'|}} < 1$ , we can upper bound  $d_\xi$  by  $\left(\frac{24C|J'|}{\epsilon}\right)^{|J'|}$  [52]; and (2) Since  $W \subset \mathbb{R}^P$  can be circumscribed by a sphere in  $\mathbb{R}^P$  of radius  $\Gamma\sqrt{P}$ , we can upper bound  $m_\zeta$  by  $\left(\frac{24L'\Gamma\sqrt{P}}{\epsilon}\right)^P$ . ■

## APPENDIX C PROOF OF THEOREM 4

Let us begin with any coordinate descent iteration  $r$  and dimension  $k$ , which starts with some  $\tilde{w}_k^r \in \mathbb{R}^P$  (recall that  $\tilde{w}_1^1 \equiv 0$ ). In order to facilitate the proof, we explicitly write  $w^* \in \mathbb{R}$ , defined in (24), and  $\hat{w} \in \mathbb{R}$ , defined in (25), as  $w_k^{*r}$  and  $\hat{w}_k^r$ , respectively, to bring out their dependence on the indices  $(r, k)$ . We now define  $\hat{w}^r \in \mathbb{R}^P$  as  $[\hat{w}^r]_k := \hat{w}_k^r$  and recall from Lemma 3 and the subsequent definitions that  $\forall r \in \mathbb{N}, \forall j \in J', \lim_{T \rightarrow \infty} w_j^{r,T} = \hat{w}^r$ . It therefore suffices to show that, in the limit of large  $\bar{r}$ , the statistical risk of  $\hat{w}^{\bar{r}}$  approaches the minimum statistical risk.

We now fix an arbitrary  $\epsilon' \in (0, 1)$  and note that

$$\sup_{r, k} [h_k^r(\hat{w}_k^r) - h_k^r(w_k^{*r})] < \epsilon' \quad (62)$$

with high probability due to Theorem 3. Note that invoking Theorem 3 requires the conditions  $|\hat{w}_k^r| \leq \Gamma$ ,  $|w_k^{*r}| \leq \Gamma$ , and  $\{\tilde{w}_k^r\}_{r, k} \subset W$ . We will return to these conditions in the latter

part of the proof. Going forward, we condition on the event described by (62) and notice that  $\forall w' \in \mathbb{R}$ , we have

$$h_k^r(w^{*r}_k) \stackrel{(a)}{\leq} h_k^r(w') \stackrel{(b)}{\leq} h_k^r(\hat{w}_k^{r-1}) + \frac{L}{2}|w' - \hat{w}_k^{r-1}|^2 + [\nabla h_k^r(\hat{w}_k^{r-1})]_k(w' - \hat{w}_k^{r-1}), \quad (63)$$

where (a) follows from the definition of  $w^{*r}_k$  and (b) follows from Assumption 1. Plugging  $w' = \hat{w}_k^{r-1} - \frac{1}{L}[\nabla h_k^r(\hat{w}_k^{r-1})]_k$  in (63), we obtain

$$h_k^r(\hat{w}_k^{r-1}) - h_k^r(w^{*r}_k) \geq \frac{1}{2L}[\nabla h_k^r(\hat{w}_k^{r-1})]_k^2 \quad (64)$$

$$\Leftrightarrow h_k^r(\hat{w}_k^{r-1}) - h_k^r(\hat{w}_k^r) > \frac{1}{2L}[\nabla h_k^r(\hat{w}_k^{r-1})]_k^2 - \epsilon' \quad (65)$$

where (64) follows from (62). We have from (65) that  $h_k^r(\hat{w}_k^r)$  is a strict monotonically decreasing function of  $r$  for all  $k$  as long as  $[\nabla h_k^r(\hat{w}_k^{r-1})]_k^2 \geq 2L\epsilon'$ . It therefore follows that there exists some  $r_0 \in \mathbb{N}$  such that

$$\forall k, \forall r \geq r_0, [\nabla h_k^r(\hat{w}_k^r)]_k^2 < 4L\epsilon' \quad (66)$$

$$\Rightarrow \forall r \geq r_0, \|\nabla \mathbb{E}[f(\hat{w}^r, (x, y))]\| < 2\sqrt{P}L\epsilon'. \quad (67)$$

In addition, convexity of  $\mathbb{E}[f(\cdot, (x, y))]$  dictates

$$\mathbb{E}[f(w^*, (x, y))] \geq \mathbb{E}[f(\hat{w}^r, (x, y))] + \nabla \mathbb{E}[f(\hat{w}^r, (x, y))]^T(w^* - \hat{w}^r). \quad (68)$$

Using (67), (68), and the Cauchy–Schwarz inequality yields

$$\forall r \geq r_0, \mathbb{E}[f(\hat{w}^r, (x, y))] - \mathbb{E}[f(w^*, (x, y))] \leq \|\nabla \mathbb{E}[f(\hat{w}^r, (x, y))]\| \|\hat{w}^r - w^*\| < 2PL\Gamma\epsilon'. \quad (69)$$

Setting  $\epsilon' = \epsilon/c_4$  in (68) and removing conditioning on (62) using Theorem 3 give us the desired bound.

We conclude by commenting on the validity of  $|\hat{w}_k^r| \leq \Gamma$ ,  $|w^{*r}_k| \leq \Gamma$ , and  $\{\hat{w}_k^r\}_{r,k} \subset W$  needed for Theorem 3. The condition holds for  $\hat{w}_1^1$  and  $w^{*1}_1$  from the definition of the set  $W$ . The proof of Theorem 3 and a union bound argument also tells us that we can augment the bound in Theorem 3 with  $|h_1^1(\hat{w}_1^1) - h_1^1(w^{*1}_1)| < \epsilon'$  without either requiring  $|\hat{w}_1^1| \leq \Gamma$  or exploding the probability of failure. This however leads to the condition  $h_1^1(\hat{w}_1^1) < h_1^1(0)$  due to (65), which implies  $|\hat{w}_1^1| \leq \Gamma$ ,  $\hat{w}_2^1 \in W$ , and  $|w^{*2}_1| \leq \Gamma$ . We can therefore revert to the original probability bound of Theorem 3 and start over the same argument with the knowledge that  $\hat{w}_2^1 \in W$ ,  $|\hat{w}_1^1| \leq \Gamma$ , and  $|w^{*2}_1| \leq \Gamma$ . The rest of the claim follows by induction. ■

#### APPENDIX D PROOF OF THEOREM 5

We begin by defining the notation  $\bar{f}(w) := \mathbb{E}[f(w, (x, y))]$  for  $w \in \mathbb{R}^P$ . Next, we borrow the notation of  $\hat{w}^r \in \mathbb{R}^P$  and some facts from the proof of Theorem 4 in Appendix C. These facts include the following. First,  $\hat{w}^r$  is the output of ByRDIE after each iteration  $r$  at all nonfaulty nodes, i.e.,  $\forall r \in \mathbb{N}, \forall j \in J', \lim_{T \rightarrow \infty} w_j^{r,T} = \hat{w}^r$ . Second, defining  $\epsilon' := \epsilon/c_4$ , we have

$$\forall(r, k), h_k^r(\hat{w}_k^r) - h_k^r(\hat{w}_k^{r-1}) < -\epsilon' \quad (70)$$

with probability  $\geq 1 - \delta(\epsilon, N, \bar{a})$  as long as  $\bar{f}(\hat{w}^r) - \bar{f}^* > \epsilon$  (see, e.g., (65) and the discussion around it). Conditioning on

the probability event described by (70), the definition of  $h_k^r(\cdot)$  and (70) then give us the following recursion in  $r$ :

$$\forall r, \bar{f}(\hat{w}^r) - \bar{f}(\hat{w}^{r-1}) < -\epsilon' \quad (71)$$

as long as  $\bar{f}(\hat{w}^r) - \bar{f}^* > \epsilon$ . Note here that  $\hat{w}^0 = \hat{w}_1^1 \equiv 0$ . Therefore, (71) and the telescoping sum argument give us

$$\bar{f}(\hat{w}^r) < \bar{f}^0 - r\epsilon' \Leftrightarrow \bar{f}(\hat{w}^r) - \bar{f}^* < \bar{f}^0 - \bar{f}^* - r\epsilon' \quad (72)$$

$$\Leftrightarrow \bar{f}(\hat{w}^r) - \bar{f}^* < (\bar{f}^0 - \bar{f}^*) \left(1 - \frac{r\epsilon'}{\bar{f}^0 - \bar{f}^*}\right). \quad (73)$$

Plugging the inequality  $\bar{f}^0 - \bar{f}^* \leq L'\|w^*\| \leq L'\sqrt{P}\gamma^*$  in (72) completes the proof. ■

#### APPENDIX E PROOF OF LEMMA 4

Similar to the proof of Theorem 5, we once again use the notation  $\bar{f}(Q(q)) := \mathbb{E}[f(Q(q), (x, y))]$  to denote the statistical risk incurred by  $Q(q)$  and show  $\bar{f}(Q(q)) \xrightarrow{q, N} \bar{f}(w^*)$  in probability. It then follows from [53, Theorem 4.4] and our assumptions that  $w^*$  is a *strong minimizer* of  $\bar{f}(\cdot)$  and, therefore,  $Q(q) \xrightarrow{q, N} w^*$  in probability.

In order to prove the aforementioned claim, we fix any  $\epsilon > 0$  and show that  $\bar{f}(Q(q)) - \bar{f}(w^*) \leq \epsilon$  for all large enough  $q$  with probability that approaches 1 as  $N \rightarrow \infty$ . To this end, we claim that  $\bar{f}(Q(q))$  for all  $q$  greater than some  $q_0 \in \mathbb{N}$  is a strictly monotonically decreasing function with probability 1, which is also lower bounded by  $\bar{f}(w^*)$ . By the monotone convergence theorem, therefore,  $\bar{f}(Q(q))$  converges. We claim this convergence only takes place when  $|\nabla \bar{f}(Q(q))|_k \leq \epsilon_\nabla$ . Relegating the validity of this claim to the latter part of this proof, this means that  $|\nabla \bar{f}(Q(q))|_k$  eventually becomes smaller than  $\epsilon_\nabla$  with probability 1 for large enough  $q$ , which implies

$$\begin{aligned} \bar{f}(Q(q)) - \bar{f}(w^*) &\leq -\nabla \bar{f}(Q(q))^T(w^* - Q(q)) \\ &\leq \|\nabla \bar{f}(Q(q))\|_2(\|w^*\|_2 + \|Q(q)\|_2) \\ &\leq (\sqrt{P}\epsilon_\nabla) \cdot (2\sqrt{P}\Gamma) < \epsilon \end{aligned} \quad (74)$$

because of convexity of  $\bar{f}(\cdot)$ , the Cauchy–Schwarz inequality, and our assumptions. Since this is the desired result, we need now focus on the claim of strict monotonicity of  $\bar{f}(Q(q))$  for this lemma. To prove this claim, note from Assumption 1 that

$$\begin{aligned} \bar{f}(Q(q+1)) &\leq \bar{f}(Q(q)) + \nabla \bar{f}(Q(q))^T(Q(q+1) - Q(q)) \\ &\quad + \frac{L}{2}\|Q(q+1) - Q(q)\|_2^2 \\ &\stackrel{(a)}{=} \bar{f}(Q(q)) + [\nabla \bar{f}(Q(q))]_k[(Q(q+1) - Q(q))]_k \\ &\quad + \frac{L}{2}|[Q(q+1) - Q(q)]_k|^2, \end{aligned} \quad (75)$$

where (a) follows from the fact that  $Q(q+1) - Q(q)$  is only nonzero in dimension  $k$ .

Next, we rewrite (34) as follows:

$$\begin{aligned}
Q(q+1) &= Q(q) - \eta(q) \sum_{i=1}^{|J'|} [\pi(r+1)]_i \left( [\nabla \hat{f}(Q(q), S_i)]_k - \right. \\
&\quad \left. [\nabla \hat{f}(w_i^r, S_i)]_k + [\nabla \hat{f}(w_i^r, S_i)]_k \right) e_k \\
&= Q(q) - \eta(q) \sum_{i=1}^{|J'|} [\pi(r+1)]_i [\nabla \hat{f}(Q(q), S_i)]_k e_k + E(q),
\end{aligned} \tag{76}$$

where  $E(q) := \eta(q) \sum_{i=1}^{|J'|} [\pi(r+1)]_i ([\nabla \hat{f}(Q(q), S_i)]_k - [\nabla \hat{f}(w_i^r, S_i)]_k) e_k$ . Plugging this into (75) results in

$$\begin{aligned}
\bar{f}(Q(q)) - \bar{f}(Q(q+1)) &\geq -[E(q)]_k [\nabla \bar{f}(Q(q))]_k \\
&+ \eta(q) \sum_{i=1}^{|J'|} [\pi(r+1)]_i [\nabla \hat{f}(Q(q), S_i)]_k [\nabla \bar{f}(Q(q))]_k \\
&- \frac{L}{2} \left| [E(q)]_k - \eta(q) \sum_{i=1}^{|J'|} [\pi(r+1)]_i [\nabla \hat{f}(Q(q), S_i)]_k \right|^2. \tag{77}
\end{aligned}$$

The right-hand side of (77) strictly lower bounded by 0 implies strict monotonicity of  $\bar{f}(Q(q))$ . Simple algebraic manipulations show that this is equivalent to the condition

$$\begin{aligned}
&\frac{L\eta(q)^2}{2} \left( \sum_{i=1}^{|J'|} [\pi(r+1)]_i [\nabla \hat{f}(Q(q), S_i)]_k \right)^2 \\
&< \eta(q) \sum_{i=1}^{|J'|} [\pi(r+1)]_i [\nabla \hat{f}(Q(q), S_i)]_k [\nabla \bar{f}(Q(q))]_k \\
&\quad + L\eta(q) [E(q)]_k \sum_{i=1}^{|J'|} [\pi(r+1)]_i [\nabla \hat{f}(Q(q), S_i)]_k \\
&\quad - [E(q)]_k [\nabla \bar{f}(Q(q))]_k - \frac{L}{2} [E(q)]_k^2. \tag{78}
\end{aligned}$$

Next, notice  $\mathbb{E}[\sum_{i=1}^{|J'|} [\pi(r+1)]_i [\nabla \hat{f}(Q(q), S_i)]_k] = \mathbb{E}[\nabla f(Q(q), (x, y))]_k$ . We now make an assumption whose validity is also discussed at the end of the proof. We assume  $\exists q'_0 \in \mathbb{N} : \forall q \geq q'_0, Q(q) \in W$ , in which case we can show using arguments similar to the ones in the proof of Theorem 3 (cf. Appendix B) that

$$\mathbb{P} \left( \left| \sum_{i \in J'} [\pi(r+1)]_i [\nabla \hat{f}(Q(q), S_i)]_k - [\nabla \bar{f}(Q(q))]_k \right| \leq \epsilon' \right) \tag{79}$$

converges to 1 uniformly for all  $(r, k)$  (equivalently, all  $q$ ) for any  $\epsilon' > 0$ . We therefore have with probability 1 (as  $N \rightarrow \infty$ )

$$\begin{aligned}
\left( \sum_{i=1}^{|J'|} [\pi(r+1)]_i [\nabla \hat{f}(Q(q), S_i)]_k \right)^2 &\leq [\nabla \bar{f}(Q(q))]_k^2 + \epsilon'^2 \\
&+ 2\epsilon' \left| [\nabla \bar{f}(Q(q))]_k \right|. \tag{80}
\end{aligned}$$

Next, we consider two cases: (i)  $[\nabla \bar{f}(Q(q))]_k > 0$ , and (ii)  $[\nabla \bar{f}(Q(q))]_k < 0$ . When  $[\nabla \bar{f}(Q(q))]_k > 0$ , we have in probability  $\sum_{i \in J'} [\pi(r+1)]_i [\nabla \hat{f}(Q(q), S_i)]_k \geq [\nabla \bar{f}(Q(q))]_k - \epsilon'$ . This fact along with (80), the realization that  $\frac{L\eta(q)^2}{2} - \eta(q) < 0$  for large enough  $q$  since  $\eta(q) \rightarrow 0$ , and some tedious but

straightforward algebraic manipulations show that the following condition is sufficient for (78) to hold in probability:

$$\begin{aligned}
&\left| [\nabla \bar{f}(Q(q))]_k \right| > \\
&\left( \frac{2\epsilon' \left| [\nabla \bar{f}(Q(q))]_k \right|}{2 - L\eta(q)} + \frac{2 \left| [E(q)]_k \right| \left| [\nabla \bar{f}(Q(q))]_k \right|}{2\eta(q) - L\eta(q)^2} \right. \\
&\quad + \frac{2L \left| [E(q)]_k \right| \left| \sum_{i \in J'} [\pi(r+1)]_i [\nabla \hat{f}(Q(q), S_i)]_k \right|}{2 - L\eta(q)} \\
&\quad \left. + \frac{L \left| [E(q)]_k \right|^2}{2\eta(q) - L\eta(q)^2} + \frac{2L\eta(q)\epsilon' \left| [\nabla \bar{f}(Q(q))]_k \right|}{2 - L\eta(q)} + \frac{L\eta(q)\epsilon'^2}{2 - L\eta(q)} \right)^{\frac{1}{2}}. \tag{81}
\end{aligned}$$

Using similar arguments, one can also show that the case  $[\nabla \bar{f}(Q(q))]_k < 0$  also results in (81) as a sufficient condition for (78) to hold in probability. We now note that  $\left| [\nabla \bar{f}(\cdot)]_k \right|$  and  $\left| [\nabla \hat{f}(\cdot, \cdot)]_k \right|$  in (81) can be upper bounded by some constants  $L_{\nabla}$  and  $L_{\nabla}$  by virtue of Assumption 1 and the definition of  $W$ . This results in the following sufficient condition for (78):

$$\begin{aligned}
&\left| [\nabla \bar{f}(Q(q))]_k \right| > \left( \frac{2L_{\nabla}\epsilon' + 2LL_{\nabla}\eta(q)\epsilon' + L\eta(q)\epsilon'^2}{2 - L\eta(q)} \right. \\
&\quad \left. + \frac{2LL_{\nabla} \left| [E(q)]_k \right|}{2 - L\eta(q)} + \frac{2L_{\nabla} \left| [E(q)]_k \right| + L \left| [E(q)]_k \right|^2}{2\eta(q) - L\eta(q)^2} \right)^{\frac{1}{2}}. \tag{82}
\end{aligned}$$

The right-hand side of (82) can be made arbitrarily small (and, in particular, equal to  $\epsilon_{\nabla}$ ) through appropriate choice of  $\epsilon'$  and large enough  $q$ ; indeed, we have from our assumptions, Theorem 1, and the definitions of  $\eta(q)$  and  $E(q)$  that both  $[E(q)]_k$  and  $[E(q)]_k/\eta(q)$  converge to 0 as  $q \rightarrow \infty$ .

This completes the proof, except that we need to validate one remaining claim and discuss one assumption. The claim is that  $\bar{f}(Q(q))$  cannot converge when  $\left| [\nabla f(Q(q))]_k \right| > \epsilon_{\nabla}$ . We prove this by contradiction. Suppose  $\exists k \in \{1, \dots, P\}$  and  $\epsilon_0 > 0$  such that  $\left| [\nabla \bar{f}(Q(q))]_k \right| - \epsilon_{\nabla} > \epsilon_0$  for all  $q$ . We know  $\exists q_0 \in \mathbb{N}$  such that the right hand side of (82) becomes smaller than  $\epsilon_{\nabla}$  for all  $q \geq q_0$ . Therefore, adding  $\epsilon_0$  to the right hand side of (82) and combining with (77) gives  $\forall q \geq q_0$ :

$$\bar{f}(Q(q)) - \bar{f}(Q(q+1)) \geq (2\eta(q) - L\eta(q)^2)(\epsilon_0^2 + 2\epsilon_{\nabla}\epsilon_0). \tag{83}$$

Taking summation on both sides of (83) from  $q = q_0$  to  $\infty$ , and noting that  $\sum_{q=q_0}^{\infty} \eta(q) = \infty$  and  $\sum_{q=q_0}^{\infty} \eta(q)^2 < \infty$ , gives us  $\bar{f}(Q(q_0)) - \lim_{q \rightarrow \infty} \bar{f}(Q(q)) = \infty$ . This contradicts the fact that  $\bar{f}(\cdot)$  is lower bounded, thereby validating our claim.

Finally, the assumption  $\exists q'_0 \in \mathbb{N} : \forall q \geq q'_0, Q(q) \in W$  is true with probability 1 (as  $N \rightarrow \infty$ ) by virtue of the facts that  $W$  is defined in terms of the sublevel set of  $\bar{f}(\cdot)$ , (79) holds  $\forall q < q'_0$  without requiring the assumption,  $\exists q_0 \in \mathbb{N}$  such that  $\bar{f}(Q(q))$  is monotonic in  $q$  for all  $q \geq q_0$  due to (79), and the probabilistic ‘‘onion peeling’’ induction argument at the end of the proof of Theorem 4 (cf. Appendix C) is applicable in this case also (except that one will have to start the argument from some index  $q = q'_0 \geq q_0$ ). ■

## REFERENCES

- [1] Z. Yang and W. U. Bajwa, ‘‘ByRDIE: A Byzantine-resilient distributed learning algorithm,’’ in *Proc. IEEE Data Science Workshop (DSW'18)*, Lausanne, Switzerland, Jun. 2018, pp. 21–25.



- [2] V. Vapnik, "Principles of risk minimization for learning theory," in *Proc. Advances in Neural Information Processing Systems (NIPS'92)*, 1992, pp. 831–838.
- [3] —, *The Nature of Statistical Learning Theory*, 2nd ed. New York, NY: Springer-Verlag, 1999.
- [4] M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foundations of Machine Learning*. Cambridge, MA: MIT Press, 2012.
- [5] J. B. Predd, S. B. Kulkarni, and H. V. Poor, "Distributed learning in wireless sensor networks," *IEEE Signal Process. Mag.*, vol. 23, no. 4, pp. 56–69, 2006.
- [6] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. and Trends Mach. Learning*, vol. 3, no. 1, pp. 1–122, 2011.
- [7] K. Driscoll, B. Hall, H. Sivencrona, and P. Zumsteq, "Byzantine fault tolerance, from theory to reality," in *Proc. Int. Conf. Computer Safety, Reliability, and Security (SAFECOMP'03)*, 2003, pp. 235–248.
- [8] K. Driscoll, B. Hall, M. Paulitsch, P. Zumsteq, and H. Sivencrona, "The real Byzantine generals," in *Proc. Digital Avionics Syst. Conf. (DASC'04)*, 2004, pp. 1–11.
- [9] L. Lamport, R. Shostak, and M. Pease, "The Byzantine generals problem," *ACM Trans. Programming Languages and Syst.*, vol. 4, no. 3, pp. 382–401, 1982.
- [10] M. J. Fischer, N. A. Lynch, and M. S. Paterson, "Impossibility of distributed consensus with one faulty process," *J. ACM*, vol. 32, no. 2, pp. 374–382, 1985.
- [11] P. Dutta, R. Guerraoui, and M. Vukolic, "Best-case complexity of asynchronous Byzantine consensus," EPFL/IC/200499, Tech. Rep., 2005.
- [12] J. Sousa and A. Bessani, "From Byzantine consensus to BFT state machine replication: A latency-optimal transformation," in *Proc. 9th Euro. Dependable Computing Conf. (EDCC'12)*, 2012, pp. 37–48.
- [13] L. Su and N. Vaidya, "Fault-tolerant distributed optimization (Part IV): Constrained optimization with arbitrary directed networks," *arXiv preprint arXiv:1511.01821*, 2015.
- [14] L. Su and N. H. Vaidya, "Fault-tolerant multi-agent optimization: Optimal iterative distributed algorithms," in *Proc. ACM Symp. Principles of Distributed Computing*, 2016, pp. 425–434.
- [15] Z. Yang and W. U. Bajwa, "RD-SVM: A resilient distributed support vector machine," in *Proc. IEEE Intl. Conf. Acoustics, Speech, and Signal Processing (ICASSP'16)*, Shanghai, China, Mar. 2016, pp. 2444–2448.
- [16] A. Nedić and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Trans. Autom. Control*, vol. 54, no. 1, pp. 48–61, 2009.
- [17] S. S. Ram, A. Nedić, and V. Veeravalli, "Distributed stochastic subgradient projection algorithms for convex optimization," *J. Optim. Theory and Appl.*, vol. 147, no. 3, pp. 516–545, 2010.
- [18] P. A. Forero, A. Cano, and G. B. Giannakis, "Consensus-based distributed support vector machines," *J. Mach. Learning Research*, vol. 11, pp. 1663–1707, 2010.
- [19] J. F. Mota, J. M. Xavier, P. M. Aquiar, and M. Puschel, "D-ADMM: A communication-efficient distributed algorithm for separable optimization," *IEEE Trans. Signal Process.*, vol. 61, no. 10, pp. 2718–2723, 2013.
- [20] W. Shi, Q. Ling, K. Yuan, G. Wu, and W. Yin, "On the linear convergence of the ADMM in decentralized consensus optimization," *IEEE Trans. Signal Process.*, vol. 62, no. 7, pp. 1750–1761, 2014.
- [21] A. Nedić and A. Olshevsky, "Distributed optimization over time-varying directed graphs," *IEEE Trans. Autom. Control*, vol. 60, no. 3, pp. 601–615, 2015.
- [22] A. Mokhtari, W. Shi, Q. Ling, and A. Ribeiro, "A decentralized second-order method with exact linear convergence rate for consensus optimization," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 2, no. 4, pp. 507–522, 2016.
- [23] H. Raja and W. U. Bajwa, "Cloud K-SVD: A collaborative dictionary learning algorithm for big, distributed data," *IEEE Trans. Signal Process.*, vol. 64, no. 1, pp. 173–188, Jan. 2016.
- [24] A. Mokhtari, Q. Ling, and A. Ribeiro, "Network Newton distributed optimization methods," *IEEE Trans. Signal Process.*, vol. 65, no. 1, pp. 146–161, 2017.
- [25] Y. M. Minsky and F. B. Schneider, "Tolerating malicious gossip," *Distributed Computing*, vol. 16, no. 1, pp. 49–68, 2003.
- [26] A. Rawat, P. Anand, H. Chen, and P. Varshney, "Collaborative spectrum sensing in the presence of Byzantine attacks in cognitive radio networks," *IEEE Trans. Signal Process.*, vol. 59, no. 2, pp. 774–786, Feb. 2011.
- [27] A. Vempaty, L. Tong, and P. Varshney, "Distributed inference with Byzantine data: State-of-the-art review on data falsification attacks," *IEEE Signal Process. Mag.*, vol. 30, no. 5, pp. 65–75, May 2013.
- [28] Y. Chen, S. Kar, and J. M. F. Moura, "Attack resilient distributed estimation: A consensus+innovations approach," in *Proc. Annu. American Control Conference (ACC'18)*, Jun. 2018, pp. 1015–1020.
- [29] N. H. Vaidya, L. Tseng, and G. Liang, "Iterative approximate Byzantine consensus in arbitrary directed graphs," in *Proc. ACM Symp. Principles of Distributed Computing*, 2012, pp. 365–374.
- [30] N. Vaidya, "Matrix representation of iterative approximate Byzantine consensus in directed graphs," *arXiv preprint arXiv:1203.1888*, 2012.
- [31] H. J. LeBlanc, H. Zhang, X. Koutsoukos, and S. Sundaram, "Resilient asymptotic consensus in robust networks," *IEEE J. Sel. Areas in Commun.*, vol. 31, no. 4, pp. 766–781, 2013.
- [32] N. H. Vaidya and V. K. Garg, "Byzantine vector consensus in complete graphs," in *Proc. 2016 ACM Symp. Principles of Distributed Computing*, 2013, pp. 65–73.
- [33] N. H. Vaidya, L. Tseng, and G. Liang, "Iterative Byzantine vector consensus in incomplete graphs," in *Proc. 15th Int. Conf. Distributed Computing and Networking*, 2014, pp. 14–28.
- [34] P. Blanchard, E. E. Mhamdi, R. Guerraoui, and J. Stainer, "Byzantine-tolerant machine learning," *CoRR*, vol. abs/1703.02757, 2017.
- [35] Y. Chen, L. Su, and J. Xu, "Distributed statistical machine learning in adversarial settings: Byzantine gradient descent," in *Proc. ACM Measurement and Analysis of Computing Systems*, vol. 1, no. 2, Dec. 2017, pp. 44:1–44:25.
- [36] P. Blanchard, R. Guerraoui, and J. Stainer, "Machine learning with adversaries: Byzantine tolerant gradient descent," in *Proc. Advances in Neural Inf. Process. Syst.*, 2017, pp. 118–128.
- [37] G. Damaskinos, E. E. Mhamdi, R. Guerraoui, R. Patra, and M. Taziki, "Asynchronous Byzantine machine learning (the case of SGD)," in *Proc. 35th Int. Conf. Machine Learning*, vol. 80. PMLR, 2018, pp. 1145–1154.
- [38] E. E. Mhamdi, R. Guerraoui, and S. Rouault, "The hidden vulnerability of distributed learning in Byzantium," in *Proc. 35th Int. Conf. Machine Learning*, vol. 80. PMLR, 2018, pp. 3521–3530.
- [39] C. Xie, O. Koyejo, and I. Gupta, "Generalized Byzantine-tolerant SGD," *arXiv preprint arXiv:1802.10116*, 2018.
- [40] L. Chen, H. Wang, Z. Charles, and D. Papailiopoulos, "DRACO: Byzantine-resilient distributed training via redundant gradients," in *Proc. 35th Intl. Conf. Machine Learning (ICML'18)*, Jul. 2018, pp. 903–912.
- [41] D. Yin, Y. Chen, K. Ramchandran, and P. Bartlett, "Byzantine-robust distributed learning: Towards optimal statistical rates," in *Proc. 35th Int. Conf. Machine Learning*, vol. 80, 2018, pp. 5650–5659.
- [42] D. Alistarh, Z. Allen-Zhu, and J. Li, "Byzantine stochastic gradient descent," in *Proc. Advances in Neural Information Processing Systems*, 2018, pp. 4618–4628.
- [43] L. Su and J. Xu, "Securing distributed machine learning in high dimensions," *arXiv preprint arXiv:1804.10140*, 2018.
- [44] D. Yin, Y. Chen, K. Ramchandran, and P. Bartlett, "Defending against saddle point attack in Byzantine-robust distributed learning," *arXiv preprint arXiv:1806.05358*, 2018.
- [45] S. J. Wright, "Coordinate descent algorithms," *Math. Programming*, vol. 151, no. 1, pp. 3–34, 2015.
- [46] H. H. Sotrab, *Basic Real Analysis*, 2nd ed. New York, NY: Springer, 2003.
- [47] S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan, "Stochastic convex optimization," in *Proc. Conf. Learning Theory (COLT'09)*, Jun. 2009.
- [48] J. C. Duchi, A. Agarwal, and M. J. Wainwright, "Dual averaging for distributed optimization: Convergence analysis and network scaling," *IEEE Trans. Autom. Control*, vol. 57, no. 3, pp. 592–606, 2012.
- [49] L. Su and N. Vaidya, "Byzantine multi-agent optimization: Part I," *arXiv preprint arXiv:1506.04681*, 2015.
- [50] D. Dua and E. K. Taniskidou, "UCI machine learning repository," 2017. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [51] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *J. American stat. assoc.*, vol. 58, no. 301, pp. 13–30, 1963.
- [52] J. Verger-Gaugry, "Covering a ball with smaller equal balls in  $R^n$ ," *Discrete & Computational Geometry*, vol. 33, no. 1, pp. 143–155, 2005.
- [53] C. Planiden and X. Wang, "Strongly convex functions, Moreau envelopes, and the generic nature of convex functions with strong minimizers," *SIAM J. Optim.*, vol. 26, no. 2, pp. 1341–1364, 2016.