

---

# Collaborative Learning in the Jungle (Decentralized, Byzantine, Heterogeneous, Asynchronous and Nonconvex Learning)

---

**El-Mahdi El-Mhamdi \***

École Polytechnique  
Palaiseau, France  
el-mahdi.el-mhamdi@polytechnique.edu

**Sadegh Farhadkhani \***

IC School, EPFL  
Lausanne, Switzerland  
sadegh.farhadkhani@epfl.ch

**Rachid Guerraoui \***

IC School, EPFL  
Lausanne, Switzerland  
rachid.guerraoui@epfl.ch

**Arsany Guirguis \***

IC School, EPFL  
Lausanne, Switzerland  
arsany.guirguis@epfl.ch

**Lê-Nguyên Hoang \***

IC School, EPFL  
Lausanne, Switzerland  
le.hoang@epfl.ch

**Sébastien Rouault \***

IC School, EPFL  
Lausanne, Switzerland  
sebastien.rouault@epfl.ch

## Abstract

We study *Byzantine collaborative learning*, where  $n$  nodes seek to collectively learn from each others' local data. The data distribution may vary from one node to another. No node is trusted, and  $f < n$  nodes can behave arbitrarily. We prove that collaborative learning is equivalent to a new form of agreement, which we call *averaging agreement*. In this problem, nodes start each with an initial vector and seek to approximately agree on a common vector, which is close to the average of honest nodes' initial vectors. We present two asynchronous solutions to averaging agreement, each we prove optimal according to some dimension. The first, based on the minimum-diameter averaging, requires  $n \geq 6f + 1$ , but achieves asymptotically the best-possible averaging constant up to a multiplicative constant. The second, based on reliable broadcast and coordinate-wise trimmed mean, achieves optimal Byzantine resilience, i.e.,  $n \geq 3f + 1$ . Each of these algorithms induces an optimal Byzantine collaborative learning protocol. In particular, our equivalence yields new impossibility theorems on what any collaborative learning algorithm can achieve in adversarial and heterogeneous environments.

## 1 Introduction

The distributed nature of data, the prohibitive cost of data transfers and the privacy concerns all call for collaborative machine learning. The idea consists for each machine to keep its data locally and to “simply” exchange with other machines what it learned so far. If all machines correctly communicate and execute the algorithms assigned to them, collaborative learning is rather easy. It can be achieved through the standard workhorse optimization algorithm: stochastic gradient descent (SGD) [37], which can be effectively distributed through averaging [26].

---

\* Authors are listed alphabetically.

But in a practical distributed setting, hardware components may crash, software can be buggy, communications can be slowed down, data can be corrupted and machines can be hacked. Besides, large-scale machine learning systems are trained on user-generated data, which may be crafted maliciously. For example, recommendation algorithms have such a large influence on social medias that there are huge incentives from industries and governments to fabricate data that bias the learning algorithms and increase the visibility of some contents over others [6, 33]. In the parlance of distributed computing, “nodes” can be *Byzantine* [28], i.e., they can behave arbitrarily maliciously, to confuse the system. Given that machine learning (ML) is now used in many critical applications (e.g., driving, medication, content moderation), its ability to tolerate Byzantine behavior is of paramount importance.

In this paper, we precisely define and address, for the first time, the problem of *collaborative learning* in a *fully decentralized, Byzantine, heterogeneous* and *asynchronous* environment with *non-convex* loss functions. We consider  $n$  nodes, which may be machines or different accounts on a social media. Each node has its own local data, drawn from data distributions that may *greatly vary* across nodes. The nodes seek to collectively learn from each other, without however exchanging their data. *None of the nodes is trusted*, and any  $f < n$  nodes can be Byzantine.

**Contributions.** We first precisely formulate the collaborative learning problem. Then, we give our main contribution: an equivalence between collaborative learning and a new more abstract problem we call *averaging agreement*. More precisely, we provide two reductions: from collaborative learning to averaging agreement and from averaging agreement to collaborative learning. We prove that both reductions essentially preserve the correctness guarantees on the output. The former reduction is the most challenging one to design and to prove correct. First, to update nodes’ models, we use averaging agreement to aggregate nodes’ stochastic gradients. Then, to avoid model drift, we regularly “contract” the nodes’ models using averaging agreement. To prove correctness, we bound the diameter of honest nodes’ models, and we analyze the *effective gradient* [12]. We then carefully select a halting iteration, for which correctness can be guaranteed.

Our tight reduction allows to derive both impossibility results and optimal algorithms for collaborative learning by studying the “simpler” averaging agreement problem. We prove lower bounds on the correctness and Byzantine resilience that any averaging agreement algorithm can achieve, which implies the same lower bounds for collaborative learning. We then propose two optimal algorithms for averaging agreement. Our first algorithm is asymptotically optimal with respect to correctness, up to a multiplicative constant, when nearly all nodes are honest. Our second algorithm achieves optimal Byzantine resilience. Each of these algorithms induces an optimal collaborative learning protocol.

While our algorithms apply in a very general setting, they can easily be tweaked for more specific settings with additional assumptions, such as the presence of a trusted parameter server [30], the assumption of homogeneous (i.e. i.i.d.) local data or synchrony (Section 3.3 and Section 6).

We implemented and evaluated our algorithms in a distributed environment with 3 ResNet models [24]. More specifically, we present their throughput overhead when compared to a non-robust collaborative learning approach with both i.i.d. and non-i.i.d. data (i.e. we highlight the cost of heterogeneity). Essentially, we show that our first algorithm is more lightweight with a slowdown of at most 1.7X in the i.i.d. case and almost the triple in the non-i.i.d. case. Our second algorithm adds slightly more than an order of magnitude overhead: here the non-i.i.d. slowdown is twice the i.i.d. one.

**Related work: Byzantine learning.** Several techniques have recently been proposed for Byzantine distributed learning, where different workers collaborate through a central *parameter server* [30] to minimize the average of their loss functions [26]. In each round, the server sends its model parameters to the workers which use their local data to compute gradients. Krum and Multi-Krum [4] use a distance-based scheme to eliminate Byzantine inputs and average the remaining ones. Median-based aggregation alternatives were also considered [38]. Bulyan [13] uses a meta-algorithm against a strong adversary that can fool the aforementioned aggregation rules in high-dimensional spaces. Coding schemes were used in Draco [8] and Detox [35]. In [2], quorums of workers enable to reach an information theoretical learning optimum, assuming however a strong convex loss function. Kardam [9] uses filters to tolerate Byzantine workers in an asynchronous setting. All these approaches assume a central *trusted* (parameter server) machine.

The few decentralized approaches that removed this single point of failure, restricted however the problem to (a) homogeneous data distribution, (b) convex functions, and/or (c) a weak (non-

Byzantine) adversary. MOZI [22] combines a distance-based aggregation rule with a performance-based filtering technique, assuming that adversaries send models with high loss values, restricting thereby the arbitrary nature of a Byzantine agent that can craft poisoned models whose losses are small only with respect to the honest nodes' incomplete loss functions. The technique is also inapplicable to heterogeneous learning, where nodes can have a biased loss function compared to the average of all loss functions<sup>2</sup>. BRIDGE [39] and ByRDIE [40] consider *gradient* descent (GD) and *coordinate* descent (CD) optimizations, respectively. Both rely on trimmed-mean to achieve Byzantine resilience assuming a synchronous environment and strongly convex loss functions<sup>3</sup> with homogeneous data distribution (*i.i.d.*). In addition, none of their optimization methods is stochastic: at each step, each node is supposed to compute the gradient on its entire local data set. ByzSGD [12] starts from the classical model of several workers and one server, which is then replicated for Byzantine resilience. It is assumed that up to 1/3 of the server replicas and up to 1/3 of the workers can be Byzantine, which is stronger than what we assume in the present paper where nodes play both roles and tolerate any subset of 1/3 Byzantine nodes. More importantly, ByzSGD assumes that all communication patterns between honest servers eventually hold with probability 1; we make no such assumption here. Additionally, our present paper is more general, considering heterogeneous data distributions, as opposed to [12]. Furthermore, heterogeneity naturally calls for *personalized* collaborative learning [14, 23, 10, 15], where nodes aim to learn local models, but still leverage collaborations to improve their local models. Interestingly, our general scheme encompasses personalized collaborative learning.

Maybe more importantly, our reduction to averaging agreement yields new more precise bounds that improve upon all the results listed above. These are we believe of interest, even in more centralized, homogeneous, synchronous and convex settings. In particular, our reduction can easily be adapted to settings where parameter servers and workers with local data are different entities, as in [12].

**Related work: Agreement.** A major challenge in collaborative learning is to guarantee "some" agreement between nodes about the appropriate parameters to consider. Especially in non-convex settings, this is critical as, otherwise, the gradients computed by a node may be completely irrelevant for another node. The agreement could be achieved using the traditional *consensus* abstraction [28]. Yet, consensus is impossible in asynchronous environments [18] and when it is possible (with partial synchrony), its usage is expensive and would be prohibitive in the context of modern ML models, with a dimension  $d$  in the order of billions. In fact, and as we show in this paper, consensus is unnecessary.

An alternative candidate abstraction is *approximate agreement*. This is a weak form of consensus introduced in [11] where honest nodes converge to values that are close to each other, while *remaining in the convex hull* of the values proposed by honest nodes. In the one-dimensional case, optimal convergence rate has been achieved in both synchronous [16] and asynchronous environments [17], while optimal asynchronous Byzantine tolerance was attained by [1]. The multi-dimensional version was addressed by [31], requiring however  $n^d$  local computations in each round, and assuming  $n > f(d + 2)$ . This is clearly impractical in the context of modern ML.

By leveraging some distributed computing techniques [36, 1], we prove that collaborative learning can be reduced to *averaging agreement*, which is even weaker than approximate agreement. This enables us to bring down the requirement on the number of honest nodes from  $n > f(d + 2)$  to  $n > 3f$ , and only require linear computation time in  $d$ .

**Structure.** The rest of the paper is organized as follows. In Section 2, we precisely define the problems we aim to solve. Section 3 states our main result, namely, the equivalence between collaborative learning and averaging agreement. Section 4 describes our two solutions to averaging agreement, and proves their optimality. Section 5 reports on our empirical evaluation and highlight important takeaways. Finally, Section 6 concludes. The full proofs are provided in the supplementary material, as well as the optimized algorithm for homogeneous local data.

---

<sup>2</sup>Besides, MOZI, focusing on convex optimization, assumes that eventually, models on honest nodes do not drift among each others, which may not hold for Byzantine nodes could influence the honest models to drift away from each other [3].

<sup>3</sup>Convexity greatly helps, as the average of good models will necessarily be a good model. This is no longer the case in non-convex optimization, which includes the widely used neural network framework.

## 2 Model and Problem Definitions

### 2.1 Distributed computing assumptions

We consider a standard distributed computing model with a set  $[n] = \{1, \dots, n\}$  of nodes, out of which  $h$  are honest and  $f = n - h$  are Byzantine. For presentation simplicity, we assume that the first  $h$  nodes are honest. But crucially, no honest node knows which  $h - 1$  other nodes are honest. The  $f$  Byzantine nodes know each other, can collude, and subsequently know who the  $h$  remaining honest nodes are. Essentially, we assume a single adversary that controls all the Byzantine nodes. These nodes can send arbitrary messages, and they can send different messages to different nodes. In the terminology of distributed computing, the adversary is omniscient but not omnipotent. Such an adversary has access to all learning and deployment information, including the learning objective, the employed algorithm, as well as the dataset. We consider a general asynchronous setting [5]: the adversary can delay messages to honest nodes: no bound on communication delays or relative speeds is assumed. We denote BYZ the algorithm adopted by the adversary. Yet, we assume that the adversary is not able to delay all messages indefinitely [7]. Besides, the adversary is not able to alter the messages from the honest nodes, which can authenticate the source of a message to prevent spoofing and Sybil attacks.

Also for presentation simplicity, we assume that processes communicate in a round-based manner [5]. In each round, every honest node broadcasts a message (labelled with the round number) and waits until it successfully gathers messages from at most  $q \leq h$  other nodes (labelled with the correct round number), before performing some local computation and moving to the next round. Even though the network is asynchronous, each round is guaranteed to eventually terminate for all honest nodes, as the  $h$  honest nodes' messages will all be eventually delivered. Evidently, however, some of them may be delivered after the node receives  $q$  messages (including Byzantine nodes'). Such messages will fail to be taken into account. Our learning algorithm will then rely on main rounds (denoted  $t$  in Section 3), each of which is decomposed into sub-rounds that run averaging agreements.

### 2.2 Machine learning assumptions

We assume each honest node  $j \in [h]$  has a local data distribution  $\mathcal{D}_j$ . The node's local loss function is derived from the parameters  $\theta \in \mathbb{R}^d$ , the model and the local data distribution, typically through  $\mathcal{L}^{(j)}(\theta) = \mathbb{E}_{x \sim \mathcal{D}_j}[\ell(\theta, x)]$ , where  $\ell(\theta, x)$  is the loss for data point  $x$ , which may or may not include some regularization of the parameter  $\theta$ . Our model is agnostic to whether the local data distribution is a uniform distribution over collected data (i.e., empirical risk), or whether it is a theoretical unknown distribution the node can sample from (i.e., statistical risk). We make the following assumptions about this loss function.

**Assumption 1.** *The loss functions are non-negative, i.e.,  $\mathcal{L}^{(j)} \geq 0$  for all honest nodes  $j \in [h]$ .*

**Assumption 2.** *The loss functions are  $L$ -smooth, i.e., there exists a constant  $L$  such that*

$$\forall \theta, \theta' \in \mathbb{R}^d, \forall j \in [h], \left\| \nabla \mathcal{L}^{(j)}(\theta) - \nabla \mathcal{L}^{(j)}(\theta') \right\|_2 \leq L \|\theta - \theta'\|_2. \quad (1)$$

**Assumption 3.** *The variance of the noise in the gradient estimations is uniformly bounded, i.e.,*

$$\forall j \in [h], \forall \theta \in \mathbb{R}^d, \mathbb{E}_{x \sim \mathcal{D}_j} \left\| \nabla_{\theta} \ell(\theta, x) - \nabla \mathcal{L}^{(j)}(\theta) \right\|_2^2 \leq \sigma^2. \quad (2)$$

*Moreover, the data samplings done by two different nodes are independent.*

**Assumption 4.** *There is a computable bound  $\mathcal{L}_{max}$  such that, at initial point  $\theta_1 \in \mathbb{R}^d$ , for any honest node  $j \in [h]$ , we have  $\mathcal{L}^{(j)}(\theta_1) \leq \mathcal{L}_{max}$ .*

While the first three assumptions are standard, the fourth assumption deserves further explanation. Notice first that  $\theta_1$  is a given parameter of our algorithms, which we could, just for the sake of the argument, set to 0. The assumption would thus be about the value of the local losses at 0, which will typically depend on the nodes' local data distribution. But losses usually depend on the data only as an average of the loss per data point. Moreover, the loss at 0 for any data point is usually bounded. In image classification tasks for example, each color intensity of each pixel of an input image has a bounded value. This usually suffices to upper-bound the loss at 0 for any data point, which then yields Assumption 4.

In iteration  $t$ , we require each node to average the stochastic gradient estimates over a batch of  $b_t$  i.i.d. samples. As a result, denoting  $\theta_t^{(j)}$  and  $g_t^{(j)} \triangleq \frac{1}{b_t} \sum_{i \in [b_t]} \nabla_{\theta} \ell(\theta_t^{(j)}, x_{t,i}^{(j)})$  node  $j$ 's parameters and computed stochastic gradient in iteration  $t$ , we have

$$\mathbb{E}_{x \sim \mathcal{D}_j} \left\| g_t^{(j)} - \nabla \mathcal{L}^{(j)}(\theta_t^{(j)}) \right\|_2^2 \leq \frac{\sigma^2}{b_t} \triangleq \sigma_t^2. \quad (3)$$

As  $t$  grows, we increase batch size  $b_t$  up to  $\Theta(1/\delta^2)$  (where  $\delta$  is a parameter of the collaborative learning problem, see Section 2.3), so that  $\sigma_t = \mathcal{O}(\delta)$  for  $t$  large enough (see Remark 2). This allows to dynamically mitigate the decrease of the norm of the true gradient. Namely, early on, while we are far from convergence, this norm is usually large. It is then desirable to have very noisy estimates, as these can be obtained more efficiently, and as the aggregation of these poor estimates will nevertheless allow progress. However, as we get closer to convergence, the norm of the true gradient becomes smaller, making the learning more vulnerable to Byzantine attacks [3]. Increasing the batch size then becomes useful. Our proofs essentially formalize this intuition. Note that, in the homogeneous setting, [25] theoretically proves that leveraging Momentum reduces variance and prevents Byzantine attacks without increasing the batch size, while [32] observes on several hundred experimental settings that Momentum indeed often offsets, or even cancels out, the effects of Byzantine attacks.

### 2.3 Collaborative learning

Given the  $\ell_2$  diameter  $\Delta_2(\vec{\theta}) = \max_{j,k \in [h]} \|\theta^{(j)} - \theta^{(k)}\|_2$ , collaborative learning consists in minimizing the average  $\bar{\mathcal{L}}(\vec{\theta}) \triangleq \frac{1}{h} \sum_{j \in [h]} \mathcal{L}^{(j)}(\vec{\theta})$  of local losses at the average  $\vec{\theta} \triangleq \frac{1}{h} \sum_{j \in [h]} \theta^{(j)}$ , while guaranteeing that the honest nodes' parameters have a small diameter.

This general model encompasses to the *personalized* federated learning problem introduced by [14, 23, 10, 15]. For instance, in [23], each node  $j$  aims to learn a local model  $x_j$  that minimizes  $f_j$ , with a penalty  $\frac{\lambda}{2} \|x_j - \bar{x}\|_2^2$  on their distance to the average  $\bar{x}$  of all models. This framework can be restated by considering that nodes must agree on a common parameter  $\theta = \bar{x}$ , but have local losses defined by  $\mathcal{L}^{(j)}(\theta) \triangleq \min_{x_j} f_j(x_j) + \frac{\lambda}{2} \|x_j - \theta\|_2^2$ . The problem of [23] then boils down to minimizing the average of local losses.

**Definition 1.** *An algorithm LEARN solves the Byzantine  $C$ -collaborative learning problem if, given any local losses  $\mathcal{L}^{(j)}$  for  $j \in [h]$  satisfying assumptions (1,2,3,4) and any  $\delta > 0$ , no matter what Byzantine attack BYZ is adopted by Byzantines, LEARN outputs a vector family  $\vec{\theta}$  of honest nodes such that*

$$\mathbb{E} \Delta_2(\vec{\theta})^2 \leq \delta^2 \quad \text{and} \quad \mathbb{E} \|\nabla \bar{\mathcal{L}}(\vec{\theta})\|_2^2 \leq (1 + \delta)^2 C^2 K^2, \quad (4)$$

where  $K \triangleq \sup_{j,k \in [h]}, \sup_{\theta \in \mathbb{R}^d} \|\nabla \mathcal{L}^{(j)}(\theta) - \nabla \mathcal{L}^{(k)}(\theta)\|_2$  is the largest difference between the true local gradients at the same parameter  $\theta$ , and where the randomness comes from the algorithm (typically the random sampling for gradient estimates).

In our definition above, the constant  $K$  measures the heterogeneity of local data distributions. Intuitively, this also captures the hardness of the problem. Indeed, the more heterogeneous the local data distributions, the more options Byzantine nodes have to bias the learning, the harder it is to learn in a Byzantine-resilient manner. Interestingly, for convex quadratic losses, our guarantee implies straightforwardly an upper-bound on the distance to the unique optimum of the problem, which is proportional to the hardness of the problem measured by  $K$ . In fact, our equivalence result conveys the tightness of this guarantee. In particular, the combination of our equivalence and of Theorem 5 implies that, for any  $\varepsilon > 0$ , asynchronous  $(2f/h - \varepsilon)$ -collaborative learning is impossible.

### 2.4 Averaging agreement

We address collaborative learning by reducing it to a new abstract distributed computing problem, which we call *averaging agreement*.

**Definition 2.** *A distributed algorithm AVG achieves Byzantine  $C$ -averaging agreement if, for any input  $N \in \mathbb{N}$ , any vector family  $\vec{x} \in \mathbb{R}^{d \cdot h}$  and any Byzantine attack BYZ, denoting  $\vec{y} \triangleq \text{AVG}_N(\vec{x}, \text{BYZ})$  the output of AVG given such inputs, we guarantee*

$$\mathbb{E} \Delta_2(\vec{y})^2 \leq \frac{\Delta_2(\vec{x})^2}{4^N} \quad \text{and} \quad \mathbb{E} \|\vec{y} - \bar{x}\|_2^2 \leq C^2 \Delta_2(\vec{x})^2, \quad (5)$$

where  $\bar{y}_N = \frac{1}{h} \sum_{j \in [h]} y_N^{(j)}$  is the average of honest nodes' vectors, and where the randomness comes from the algorithm. We simply say that an algorithm solves averaging agreement if there exists a constant  $C$  for which it solves  $C$ -averaging agreement.

In particular, for deterministic algorithms,  $C$ -averaging on input  $N$  ensures the following guarantee

$$\Delta_2(\bar{y}) \leq \frac{\Delta_2(\bar{x})}{2^N} \quad \text{and} \quad \|\bar{y} - \bar{x}\|_2 \leq C \Delta_2(\bar{x}). \quad (6)$$

In Section 4, we will present two solutions to the averaging agreement problem. These solutions typically involve several rounds. At each round, each node sends their current vector to all other nodes. Then, once a node has received sufficiently many vectors, it will execute a robust mean estimator to these vectors, the output of which will be their starting vector for the next round. The nodes then halt after a number of rounds dependent on the parameter  $N$ .

### 3 The Equivalence

The main result of this paper is that, for  $K > 0$ ,  $C$ -collaborative learning is equivalent to  $C$ -averaging agreement. We present two reductions, first from collaborative learning to averaging agreement, and then from averaging agreement to collaborative learning.

#### 3.1 From collaborative learning to averaging agreement

Given an algorithm AVG that solves Byzantine  $C$ -averaging agreement, we design a Byzantine collaborative learning algorithm LEARN. Recall that LEARN must take a constant  $\delta > 0$  as input, which determines the degree of agreement (i.e., learning quality) that LEARN must achieve.

All honest parameter vectors are initialized with the same random values (i.e.,  $\forall j \in [h], \theta_1^{(j)} = \theta_1$ ) using a pre-defined seed. At iteration  $t$ , each honest node  $j \in [h]$  first computes a local gradient estimate  $g_t^{(j)}$  given its local loss function  $\mathcal{L}^{(j)}(\cdot)$  and its local parameters  $\theta_t^{(j)}$ , with a batch size  $b_t$ . But, instead of performing a learning step with this gradient estimate, LEARN uses an aggregate of all local gradients, which we compute using the averaging agreement algorithm AVG.

Recall from Definition 2 that AVG depends on a parameter which defines the degree of agreement. We set this parameter at  $N(t) \triangleq \lceil \log_2 t \rceil$  at iteration  $t$ , so that  $1/4^{N(t)} \leq 1/t^2$ . Denoting  $\vec{\gamma}_t$  the output of  $\overrightarrow{\text{AVG}}_{N(t)}$  applied to vectors  $\vec{g}_t$ , we then have the following guarantee:

$$\mathbb{E} \Delta_2(\vec{\gamma}_t)^2 \leq \frac{\Delta_2(\vec{g}_t)^2}{t^2} \quad \text{and} \quad \mathbb{E} \|\vec{\gamma}_t - \bar{g}_t\|_2^2 \leq C^2 \Delta_2(\vec{g}_t)^2, \quad (7)$$

where the expectations are conditioned on  $\vec{g}_t$ . We then update node  $j$ 's parameters by  $\theta_{t+1/2}^{(j)} = \theta_t^{(j)} - \eta \gamma_t^{(j)}$ , for a fixed learning rate  $\eta \triangleq \delta/12L$ . But before moving on to the next iteration, we run once again AVG, with its parameter set to 1. Moreover, this time, AVG is run on local nodes' parameters. Denoting  $\vec{\theta}_{t+1}$  the output of  $\overrightarrow{\text{AVG}}_1$  executed with vectors  $\vec{\theta}_{t+1/2}$ , we then have

$$\mathbb{E} \Delta_2(\vec{\theta}_{t+1})^2 \leq \frac{\Delta_2(\vec{\theta}_{t+1/2})^2}{4} \quad \text{and} \quad \mathbb{E} \left\| \vec{\theta}_{t+1} - \bar{\theta}_{t+1/2} \right\|_2^2 \leq C^2 \Delta_2(\vec{\theta}_{t+1/2})^2, \quad (8)$$

where the expectations are conditioned on  $\vec{\theta}_{t+1/2}$ . On input  $\delta$ , LEARN then runs  $T \triangleq T_{\text{LEARN}}(\delta)$  learning iterations. The function  $T_{\text{LEARN}}(\delta)$  will be given explicitly in the proof of Theorem 1, where we will stress the fact that it can be computed from the inputs of the problem<sup>4</sup> ( $L, K, C, n, f, \sigma_t, \mathcal{L}_{\max}$  and  $\delta$ ). Finally, instead of returning  $\vec{\theta}_{T_{\text{LEARN}}(\delta)}$ , LEARN chooses uniformly randomly an iteration  $* \in [T_{\text{LEARN}}(\delta)]$ , using a predefined common seed, and returns the vector family  $\vec{\theta}_*$ . We recapitulate the local execution of LEARN (at each node) in Algorithm 1. We stress that on steps 6 and 8, when the averaging agreement algorithm AVG is called, the Byzantines can adopt any procedure BYZ, which

<sup>4</sup>Note that upper-bounds (even conservative) on such values suffice. Our guarantees still hold, though  $T_{\text{LEARN}}(\delta)$  would take larger values, which makes LEARN slower to converge to  $\vec{\theta}_*$ .

consists in sending any message to any node at any point based on any information in the system, and in delaying for any amount of time any message sent by any honest node. Note that, apart from this, all other steps of Algorithm 1 is a purely local operation.

**Data:** Local loss gradient oracle, parameter  $\delta > 0$   
**Result:** Model parameters  $\theta$

```

1 Initialize local parameters  $\theta_1$  using a fixed common seed;
2 Fix learning rate  $\eta \triangleq \delta/12L$ ;
3 Fix number of iterations  $T \triangleq T_{\text{LEARN}}(\delta)$ ;
4 for  $t \leftarrow 1, \dots, T$  do
5    $g_t \leftarrow \text{GradientOracle}(\theta_t, b_t)$ ;
6    $\gamma_t \leftarrow \text{AVG}_{N(t)}(\vec{g}_t, \text{BYZ})$  // Vulnerable to Byzantine attacks
7    $\theta_{t+1/2} \leftarrow \theta_t - \eta\gamma_t$ ;
8    $\theta_{t+1} \leftarrow \text{AVG}_1(\vec{\theta}_{t+1/2}, \text{BYZ})$  // Vulnerable to Byzantine attacks
9 end
10 Draw  $* \sim \mathcal{U}([T])$  using the fixed common seed;
11 Return  $\theta_*$ ;

```

**Algorithm 1:** LEARN execution on a honest node.

**Remark 1.** In practice, it may be more efficient to return the last computed vector family, though our proof applies to a randomly selected iteration.

**Theorem 1.** Under assumptions (1, 2, 3, 4) and  $K > 0$ , given a  $C$ -averaging agreement oracle  $\text{AVG}$ , on any input  $0 < \delta < 3$ , LEARN solves Byzantine  $C$ -collaborative learning.

The proof is quite technical, and is provided in the supplementary material. Essentially, it focuses on the average of all honest nodes' local parameters, and on the *effective gradients* that they undergo, given the local gradient updates and the applications of the averaging agreement oracle  $\text{AVG}$ .

**Remark 2.** Our proof requires  $T_{\text{LEARN}}(\delta) = \Theta\left(\delta^{-1} \max\{\delta^{-2}, (t \mapsto \sup_{s \geq t} \sigma_s)^{-1}(\delta)\}\right)$ . To prevent the noise from being the bottleneck for the convergence rate, we then need  $\sigma_t = \Theta(\delta)$ , so that  $T_{\text{LEARN}} = \Theta(\delta^{-3})$ . Interestingly, this can be obtained by, for example, setting  $b_t \triangleq t$ , for  $t \leq T_1 = \Theta(\delta^{-2})$ , and  $b_t \triangleq T_1$  for  $t > T_1$ , where  $T_1$  is precisely defined by the proof provided in the supplementary material. In particular, we do not need to assume that  $b_t \rightarrow \infty$ .

### 3.2 Converse reduction

We also prove the converse reduction, from averaging agreement to collaborative learning.

**Theorem 2.** Given a Byzantine  $C$ -collaborative learning oracle, then, for any  $\delta > 0$ , there is a solution to Byzantine  $(1 + \delta)C$ -averaging agreement.

The proof is given in the supplementary material. It is obtained by applying a collaborative learning algorithm LEARN to the local loss functions  $\mathcal{L}^{(j)}(\theta) \triangleq \|\theta - x^{(j)}\|_2^2$ .

This converse reduction proves the tightness of our former reduction, and allows to straightforwardly derive impossibility theorems about collaborative learning from impossibility theorems about averaging agreement. In particular, in the sequel, we prove that no asynchronous algorithm can achieve better than Byzantine  $\frac{h+2f-q}{h}$ -averaging agreement. It follows that no algorithm can achieve better than Byzantine  $\frac{h+2f-q}{h}$ -collaborative learning. Similarly, Theorem 6 implies the impossibility of Byzantine asynchronous collaborative learning for  $n \leq 3f$ .

### 3.3 Particular cases

**Trusted server.** It is straightforward to adapt our techniques and prove the equivalence between averaging agreement and collaborative learning in a context with a trusted server. Our lower bounds for asynchronous collaborative learning still apply to  $C$ -averaging agreement, and thus to  $C$ -collaborative learning. Note, however, that the trusted server may allow to improve the speed of collaborative learning, as it no longer requires contracting the parameters of local nodes' models.

**Homogeneous learning.** In the supplementary material, we propose a faster algorithm for i.i.d. data, called HOM-LEARN, which skips the averaging agreement of nodes' gradients. Despite requiring fewer communications, HOM-LEARN remains correct, as the following theorem shows.

**Theorem 3.** *Under assumptions (1, 2, 3, 4), for i.i.d. local data and given a C-averaging agreement oracle AVG, on input  $\delta > 0$ , HOM-LEARN guarantees  $\mathbb{E} \Delta_2(\vec{\theta}_*)^2 \leq \delta^2$  and  $\mathbb{E} \|\nabla \mathcal{L}(\vec{\theta}_*)\|_2^2 \leq \delta^2$ .*

## 4 Solutions to Averaging Agreement

We now present two solutions to the averaging agreement problem, called MINIMUM-DIAMETER AVERAGING<sup>5</sup> (MDA) and RELIABLE BROADCAST - TRIMMED MEAN (RB-TM), each thus inducing a solution to collaborative learning. We prove each optimal according to some dimension.

### 4.1 Optimal averaging

Given a family  $\vec{z} \in \mathbb{R}^{d \cdot q}$  of vectors, MDA first identifies a subfamily  $S_{\text{MDA}}(\vec{z})$  of  $q - f$  vectors of minimal  $\ell_2$  diameter, i.e.,

$$S_{\text{MDA}}(\vec{z}) \in \arg \min_{\substack{S \subset [q] \\ |S|=q-f}} \Delta_2(\vec{z}^{(S)}) = \arg \min_{\substack{S \subset [q] \\ |S|=q-f}} \max_{j,k \in S} \|z^{(j)} - z^{(k)}\|_2. \quad (9)$$

We denote  $\vec{z}^{(\text{MDA})} \in \mathbb{R}^{d \cdot (q-f)}$  the subfamily thereby selected. MDA then outputs the average of this subfamily, i.e.,

$$\text{MDA}(\vec{z}) \triangleq \frac{1}{q-f} \sum_{j \in S_{\text{MDA}}(\vec{z})} z^{(j)}. \quad (10)$$

On input  $N \in \mathbb{N}$ ,  $\text{MDA}_N$  iterates MDA  $T_{\text{MDA}}(N) = \lceil N \ln 2 / \tilde{\varepsilon} \rceil$  times on vectors received from other nodes at each communication round such that the output of round  $t$  will be the input of round  $t + 1$ . The correctness of MDA is then ensured under the following assumption.

**Assumption 5** (Assumption for analysis of MDA). *There is  $0 < \varepsilon < 1$  such that  $n \geq \frac{6+2\varepsilon}{1-\varepsilon} f$ . This then allows to set  $q \geq \frac{1+\varepsilon}{2} h + \frac{5+3\varepsilon}{2} f$ . In this case, we define  $\tilde{\varepsilon} \triangleq \frac{2\varepsilon}{1+\varepsilon}$ .*

**Theorem 4.** *Under Assumption 5, MDA achieves Byzantine  $\frac{(2f+h-q)q+(q-2f)f}{h(q-f)\tilde{\varepsilon}}$ -averaging agreement.*

The proof is given in the supplementary material. It relies on the observation that, because of the filter, no Byzantine vector can significantly harm the estimation of the average.

**Remark 3.** *Although MDA runs in linear time in  $d$ , it runs in exponential time in  $q$ . Interestingly, assuming that each honest node fully trusts its computations and its data (which may not hold if parameter-servers do not compute gradients as in [30]), each honest node can use its own vector to filter out the  $f$  most dissimilar gradients in linear time in  $q$ , and can average out all remaining vectors. Using a similar proof as for MDA, the algorithm thereby defined can be shown to achieve asymptotically the same averaging constant as MDA, in the limit  $q \gg f$ ; but it now runs in  $\mathcal{O}(dq)$ , (requiring however  $n \geq 7f + 1$ ).*

**Theorem 5.** *No asynchronous algorithm can achieve better than Byzantine  $\frac{h+2f-q}{h}$ -averaging agreement.*

The proof is given in the supplementary material. It relies on the quasi-unanimity lemma, which shows that if a node receives at least  $q - f$  identical vectors  $x$ , then it must output  $x$ . We then construct an instance where, because of this and of Byzantines, honest nodes cannot agree. Note that, as a corollary, in the regime  $q = h \rightarrow \infty$  and  $f = o(h)$ , MDA achieves asymptotically the best-possible averaging constant, up to a multiplicative constant equal to  $3/2$ .

<sup>5</sup>Introduced by [13] in the context of robust machine learning, it uses the same principle as the minimal volume ellipsoid, that was introduced by [36] in the context of robust statistics.



## 4.2 Optimal Byzantine resilience

Our second algorithm makes use of reliable broadcast<sup>6</sup>: each Byzantine node broadcasts only a single vector (the uniqueness property of reliable broadcast in [1]). We denote  $\vec{w} \in \mathbb{R}^{d \cdot n}$  the family of vectors proposed by all nodes. For each  $j \in [h]$ ,  $w^{(j)} = x^{(j)}$  is the vector of an honest node, while a Byzantine node proposes  $w^{(j)}$  for each  $j \in [h + 1, n]$ . Moreover, [1] showed the existence of a multi-round algorithm which, by using reliable broadcast and a witness mechanism, guarantees that any two honest nodes  $j$  and  $k$  will collect at least  $q$  similar inputs. Formally, denoting  $Q^{(j)} \subseteq [n]$  the set of nodes whose messages were successfully delivered to node  $j$  (including through relays), the algorithm by [1] guarantees that  $|Q^{(j)} \cap Q^{(k)}| \geq q$  for any two honest nodes  $j, k \in [h]$ . At each iteration of our RB-TM algorithm, each node  $j$  exploits the same reliable broadcast and witness mechanism techniques to collect other nodes' vectors. Now, given its set  $Q^{(j)}$  of collected messages, each node  $j$  applies coordinate-wise trimmed mean, denoted TM, as follows. For each coordinate  $i$ , it discards the  $f$  smallest  $i$ -th coordinates it collected, as well as the  $f$  largest. We denote  $\vec{z}^{(j)} = \vec{w}^{(Q^{(j)})}$  the subfamily received by node  $j$ , and  $S(\vec{z}^{(j)}[i]) \subset [n]$  the subset of nodes whose  $i$ -th coordinates remain after trimming. Node  $j$  then computes the average  $y^{(j)}$  of the  $i$ -th coordinates of this subset, i.e.

$$y^{(j)}[i] \triangleq \frac{1}{|Q^{(j)}| - 2f} \sum_{k \in S(\vec{z}^{(j)}[i])} w^{(k)}[i]. \quad (11)$$

RB-TM consists of iterating TM, on vectors received from other nodes at each communication round. Namely, given input  $N \in \mathbb{N}$ , RB-TM iterates TM  $T_{\text{RB-TM}}(N)$  times, where

$$T_{\text{RB-TM}}(N) \triangleq \left\lceil \frac{(N + 1) \ln 2 + \ln \sqrt{h}}{\tilde{\varepsilon}} \right\rceil. \quad (12)$$

The correctness of RB-TM can then be guaranteed under the following assumption.

**Assumption 6** (Assumption for analysis of RB-TM). *There is  $\varepsilon > 0$  such that  $n \geq (3 + \varepsilon)f$ . We then set  $q = n - f$ , and define  $\tilde{\varepsilon} \triangleq \frac{\varepsilon}{1 + \varepsilon}$ .*

**Theorem 6.** *Under Assumption 6, RB-TM guarantees Byzantine  $\frac{4f}{\sqrt{h}}$ -averaging agreement. This is optimal in terms of Byzantine resilience. Indeed, for  $n \leq 3f$ , no algorithm can achieve Byzantine averaging agreement.*

The proof is provided in the supplementary material. The correctness relies on a coordinate-wise analysis, and on the study of a so-called *coordinate-wise diameter*, and its relation with the  $\ell_2$  diameter. The lower bound exploits the quasi-unanimity lemma. Note that while RB-TM tolerates more Byzantine nodes, its averaging constant is larger than that of MDA by a factor of  $\mathcal{O}(\sqrt{h})$ .

## 5 Empirical Evaluation

We implemented our collaborative learning algorithms using *Garfield* library [20] and PyTorch [34]. Each agreement algorithm comes in two variants: one assuming i.i.d. data (See supplementary material) and one tolerating non-i.i.d. data (Algorithm 1). In each case, the first variants require fewer communications. We report below on the empirical evaluation of the overhead of our four variants when compared to a non-robust collaborative learning approach. Our baseline is indeed a vanilla fully decentralized implementation in which all nodes share their updates with each other and then aggregate these updates by *averaging* (a deployment that cannot tolerate even one Byzantine node).

We focus on throughput, measuring the number of updates the system performs per second. As we consider an asynchronous network, we report on the fastest node in each experiment. We consider image classification tasks, using MNIST [29] and CIFAR-10 [27] datasets. MNIST is a dataset of handwritten digits with 70,000  $28 \times 28$  images in 10 classes. CIFAR-10 consists of 60,000  $32 \times 32$  colour images in 10 classes. We use batches of size 100, and we experimented with 5 models with different sizes ranging from simple models like small convolutional neural network (*MNIST\_CNN* and *CifarNet*), training a few thousands of parameters, to big models like ResNet-50 with around 23M

<sup>6</sup>Note that MDA can also be straightforwardly upgraded into RB-MDA to gain Byzantine resilience.




Figure 1: Convergence of our algorithms and the vanilla baseline.




Figure 2: Slowdown of our algorithms normalized to the vanilla baseline throughput.

parameters. Our experimental platform is Grid5000 [19]. We always employ nodes from the same cluster, each having 2 CPUs (Intel Xeon E5-2630 v4) with 14 cores, 768 GiB RAM,  $2 \times 10$  Gbps Ethernet, and 2 Nvidia Tesla P100 GPUs. We set  $f = 1$ , except when deploying our vanilla baseline.

Figure 1 compares the convergence of our algorithms to the vanilla baseline w.r.t. the training epochs. We use 7 nodes in this experiment, and we train Resnet-18 with CIFAR10. We verify from this figure that our algorithms can follow the same convergence trajectory as the vanilla baseline. It is clear from the figure that the *i.i.d.* versions outperform the *non-i.i.d.* ones.

Figure 2 depicts the throughput overhead of our algorithms (with both *i.i.d.* and *non-i.i.d.* data) compared to our vanilla baseline, with 10 nodes from the same cluster. Three observations from this figure are in order. First, the MDA-based algorithm performs better than the RB-TM one. The reason is that the latter incurs much more communication messages than the former as the latter uses reliable broadcast and a witness mechanism. Second, tolerating Byzantine nodes with *i.i.d.* data is much cheaper than the *non-i.i.d.* case. The reason is that it is harder to detect Byzantine behavior when data is not identically distributed on the nodes, which translates into more communication steps. Third, the slowdown is much higher with small models (i.e., *MNIST\_CNN* and *CifarNet*). This is because the network bandwidth is not saturated by the small models in the vanilla case, where it gets congested with the many communication rounds required by our algorithms. On the other hand, with the larger models, the vanilla deployment saturates the network bandwidth, making the extra communication messages account only for linear overhead.

Finally, it is important to notice that our evaluation is by no means exhaustive and our implementation has not been optimized. Our goal was to give an overview of the relative overheads. With proper optimizations, we believe the actual throughput could be increased for all implementations.

## 6 Conclusion

We defined and solved collaborative learning in a fully decentralized, Byzantine, heterogeneous, asynchronous and non-convex setting. We proved that the problem is equivalent to a new abstract form of agreement, which we call averaging agreement. We then described two solutions to averaging agreement, inducing two original solutions to collaborative learning. Each solution is optimal along some dimension. In particular, our lower bounds for the averaging agreement problem provide lower bounds on what any collaborative learning algorithm can achieve. Such impossibility results would have been challenging to obtain without our reduction. Our algorithms and our impossibility theorems are very general but can also be adapted for specific settings, such as the presence of a trusted parameter server, the assumption of *i.i.d.* data or a synchronous context<sup>7</sup>. In the latter case for instance, our two algorithms would only require  $n \geq 4f + 1$  and  $n \geq 2f + 1$ , respectively.

**Limitations and potential negative social impacts.** Like all Byzantine learning algorithms, we essentially filter out outliers. In practice, this may discard minorities with vastly diverging views. Future research should aim to address this fundamental trade-off between inclusivity and robustness. We also note that the computation time of MDA grows exponentially with  $q$ , when  $f$  is a constant fraction of  $q$ .

<sup>7</sup>In the synchronous case of MDA, with  $q = n$ ,  $n \geq \frac{4+2\epsilon}{1-\epsilon} f$  is sufficient to guarantee  $q \geq \frac{1+\epsilon}{2} h + \frac{5+3\epsilon}{2} f$  in Assumption 5. Also, note that no synchronous algorithm can achieve better than  $\frac{f}{h}$ -averaging agreement.

## Acknowledgments and Disclosure of Funding

We thank Rafaël Pinot and Nirupam Gupta for their useful comments. This work has been supported in part by the Swiss National Science Foundation projects: 200021\_182542, Machine learning and 200021\_200477, Controlling the spread of Epidemics. Most experiments presented in this paper were carried out using the Grid’5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations (see <https://www.grid5000.fr>).

## References

- [1] Ittai Abraham, Yonatan Amit, and Danny Dolev. Optimal resilience asynchronous approximate agreement. In *International Conference on Principles of Distributed Systems*, pages 229–239. Springer, 2004.
- [2] Dan Alistarh, Zeyuan Allen-Zhu, and Jerry Li. Byzantine stochastic gradient descent. In *Neural Information Processing Systems*, pages 4613–4623, 2018.
- [3] Gilad Baruch, Moran Baruch, and Yoav Goldberg. A little is enough: Circumventing defenses for distributed learning. In *Neural Information Processing Systems*, pages 8635–8645, 2019.
- [4] Peva Blanchard, El-Mahdi El-Mhamdi, Rachid Guerraoui, and Julien Stainer. Machine learning with adversaries: Byzantine tolerant gradient descent. In *Neural Information Processing Systems*, pages 118–128, 2017.
- [5] Gabriel Bracha. Asynchronous byzantine agreement protocols. *Information and Computation*, 75(2):130–143, 1987.
- [6] Samantha Bradshaw and Philip N Howard. *The global disinformation order: 2019 global inventory of organised social media manipulation*. Project on Computational Propaganda, 2019.
- [7] Miguel Castro, Barbara Liskov, et al. Practical Byzantine fault tolerance. In *Operating Systems Design and Implementation*, volume 99, pages 173–186, 1999.
- [8] Lingjiao Chen, Hongyi Wang, Zachary Charles, and Dimitris Papailiopoulos. Draco: Byzantine-resilient distributed training via redundant gradients. In *International Conference on Machine Learning*, pages 902–911, 2018.
- [9] Georgios Damaskinos, El-Mahdi El-Mhamdi, Rachid Guerraoui, Rhicheck Patra, and Mahsa Taziki. Asynchronous Byzantine machine learning (the case of SGD). In *Neural Information Processing Systems*, volume 80, pages 1145–1154, 2018.
- [10] Canh T. Dinh, Nguyen H. Tran, and Tuan Dung Nguyen. Personalized federated learning with moreau envelopes. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [11] Danny Dolev, Nancy A Lynch, Shlomit S Pinter, Eugene W Stark, and William E Weihl. Reaching approximate agreement in the presence of faults. *Journal of the ACM*, 33(3):499–516, 1986.
- [12] El-Mahdi El-Mhamdi, Rachid Guerraoui, Arsany Guirguis, Lê Nguyễn Hoàng, and Sébastien Rouault. Genuinely distributed Byzantine machine learning. In *Principles of Distributed Computing*, pages 355–364, 2020.
- [13] El-Mahdi El-Mhamdi, Rachid Guerraoui, and Sébastien Rouault. The hidden vulnerability of distributed learning in Byzantium. In *International Conference on Machine Learning*, volume 80, pages 3521–3530, 2018.

- [14] Alireza Fallah, Aryan Mokhtari, and Asuman E. Ozdaglar. Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [15] Sadegh Farhadkhani, Rachid Guerraoui, and Lê-Nguyên Hoang. Strategyproof learning: Building trustworthy user-generated datasets. *ArXiv*, 2021.
- [16] AD Fekete. Asymptotically optimal algorithms for approximate agreement. In *Principles of Distributed Computing*, pages 73–87, 1986.
- [17] Alan David Fekete. Asynchronous approximate agreement. In *Principles of Distributed Computing*, pages 64–76, 1987.
- [18] Michael J Fischer, Nancy A Lynch, and Michael S Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374–382, 1985.
- [19] Grid5000. Grid5000. <https://www.grid5000.fr/>, 2019.
- [20] Rachid Guerraoui, Arsany Guirguis, Jérémy Plassmann, Anton Ragot, and Sébastien Rouault. Garfield: System support for byzantine machine learning (regular paper). In *2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pages 39–51. IEEE, 2021.
- [21] Arsany Guirguis. LEARN source code. <https://github.com/LPD-EPFL/garfield/tree/decentralized>.
- [22] Shangwei Guo, Tianwei Zhang, Xiaofei Xie, Lei Ma, Tao Xiang, and Yang Liu. Towards byzantine-resilient learning in decentralized systems. *arXiv preprint arXiv:2002.08569*, 2020.
- [23] Filip Hanzely, Slavomír Hanzely, Samuel Horváth, and Peter Richtárik. Lower bounds and optimal algorithms for personalized federated learning. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [25] Sai Praneeth Karimireddy, Lie He, and Martin Jaggi. Learning from history for byzantine robust optimization. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 5311–5319. PMLR, 2021.
- [26] Jakub Konečný, H Brendan McMahan, Daniel Ramage, and Peter Richtárik. Federated optimization: distributed machine learning for on-device intelligence. *arXiv preprint arXiv:1610.02527*, 2016.
- [27] Alex Krizhevsky. Cifar dataset. <https://www.cs.toronto.edu/~kriz/cifar.html>, 2009.
- [28] Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine generals problem. *Transactions on Programming Languages and Systems*, 4(3):382–401, 1982.
- [29] Yann Lecunn. Mnist dataset. <http://yann.lecun.com/exdb/mnist/>, 1998.
- [30] Mu Li, David G Andersen, Jun Woo Park, Alexander J Smola, Amr Ahmed, Vanja Josifovski, James Long, Eugene J Shekita, and Bor-Yiing Su. Scaling distributed machine learning with the parameter server. In *Operating Systems Design and Implementation*, volume 1, page 3, 2014.
- [31] Hammurabi Mendes and Maurice Herlihy. Multidimensional approximate agreement in byzantine asynchronous systems. In *Symposium on Theory of computing*, pages 391–400, 2013.

- [32] El Mahdi El Mhamdi, Rachid Guerraoui, and Sébastien Rouault. Distributed momentum for byzantine-resilient stochastic gradient descent. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [33] Lisa-Maria Neudert, Philip Howard, and Bence Kollanyi. Sourcing and automation of political news and information during three european elections. *Social Media+ Society*, 5(3):2056305119863147, 2019.
- [34] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703*, 2019.
- [35] Shashank Rajput, Hongyi Wang, Zachary Charles, and Dimitris Papailiopoulos. Detox: A redundancy-based framework for faster and more robust gradient aggregation. In *Neural Information Processing Systems*, pages 10320–10330, 2019.
- [36] Peter J Rousseeuw. Multivariate estimation with high breakdown point. *Mathematical Statistics and Applications*, 8:283–297, 1985.
- [37] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986.
- [38] Cong Xie, Oluwasanmi Koyejo, and Indranil Gupta. Generalized Byzantine-tolerant SGD. *arXiv preprint arXiv:1802.10116*, 2018.
- [39] Zhixiong Yang and Waheed U Bajwa. Bridge: Byzantine-resilient decentralized gradient descent. *arXiv preprint arXiv:1908.08098*, 2019.
- [40] Zhixiong Yang and Waheed U Bajwa. Byrdie: Byzantine-resilient distributed coordinate descent for decentralized learning. *Transactions on Signal and Information Processing over Networks*, 5(4):611–627, 2019.

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? **[Yes]**
  - (b) Did you describe the limitations of your work? **[Yes]** The assumptions are fully explicit.
  - (c) Did you discuss any potential negative societal impacts of your work? **[Yes]** The conclusion notes the tension between security and inclusivity. We call for further research in addressing this tradeoff.
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]**
2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? **[Yes]** Our assumptions are clearly highlighted in Section 2.2
  - (b) Did you include complete proofs of all theoretical results? **[Yes]** They are provided in the supplementary material
3. If you ran experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[Yes]** The code is open-sourced and can be found here [21].
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[No]** Choosing the best hyperparameters is out of the scope of this paper. In our experimental evaluation, we focus only on the systems metrics, e.g., throughput.
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[No]** Error bars in our experiments are very small; we omitted them for better readability.
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[Yes]** See Section 5.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? **[Yes]** See Section 5.
  - (b) Did you mention the license of the assets? **[No]** The license can be found using the provided URLs.
  - (c) Did you include any new assets either in the supplemental material or as a URL? **[No]**
  - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? **[No]** The datasets and models we use are already open-sourced.
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[N/A]**
5. If you used crowdsourcing or conducted research with human subjects...
  - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[N/A]**
  - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? **[N/A]**
  - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **[N/A]**

# Supplementary Material

In the proofs, we call the first and second properties of averaging agreement “asymptotic agreement” and “C-averaging” guarantee respectively (see Equation (5) in the paper).

Moreover, we denote by  $\overrightarrow{\text{BYZ}}^{(j)}(\vec{x}) = \left( \text{BYZ}^{(j,1)}(\vec{x}), \dots, \text{BYZ}^{(j,q)}(\vec{x}) \right)$  the family of inputs (of size  $q$ ) collected by node  $j$ . There thus exists a bijection  $\tau : Q_t^{(j)} \rightarrow [q]$ , where  $Q_t^{(j)}$  is the set of nodes that successfully delivered messages to  $j$  at round  $t$ , such that  $\text{BYZ}^{(j,\tau(k))} = x^{(k)}$  for all honest nodes  $k \in [h]$ .

## A The Equivalence

### Preliminary lemmas

**Lemma 1.** For any  $\alpha > 0$  and any two vectors  $u$  and  $v$ , we have

$$\|u + v\|_2^2 \leq (1 + \alpha^{-1}) \|u\|_2^2 + (1 + \alpha) \|v\|_2^2. \quad (13)$$

As an immediate corollary, for any two families  $\vec{u}$  and  $\vec{v}$  of vectors, we have

$$\Delta_2(\vec{u} + \vec{v})^2 \leq (1 + \alpha^{-1}) \Delta_2(\vec{u})^2 + (1 + \alpha) \Delta_2(\vec{v})^2. \quad (14)$$

*Proof.* We have the following inequalities:

$$(1 + \alpha^{-1}) \|u\|_2^2 + (1 + \alpha) \|v\|_2^2 - \|u + v\|_2^2 = \alpha^{-1} \|u\|_2^2 + \alpha \|v\|_2^2 - 2u \cdot v \quad (15)$$

$$= \left\| \alpha^{-1/2} u - \alpha^{1/2} v \right\|_2^2 \geq 0. \quad (16)$$

Rearranging the terms yields the lemma.  $\square$

**Lemma 2.** For any vector family  $u_1, \dots, u_N$ , we have

$$\left\| \sum_{j \in [N]} u_j \right\|_2^2 \leq N \sum_{j \in [N]} \|u_j\|_2^2. \quad (17)$$

As an immediate corollary, for any family of vector families  $\vec{u}_1, \dots, \vec{u}_N$ , we have

$$\Delta_2 \left( \sum_{j \in [N]} \vec{u}_j \right)^2 \leq N \sum_{j \in [N]} \Delta_2(\vec{u}_j)^2. \quad (18)$$

*Proof.* Notice that  $u \mapsto \|u\|_2^2$  is a convex function. As a result,

$$\left\| \frac{1}{N} \sum_{j \in [N]} u_j \right\|_2^2 \leq \frac{1}{N} \sum_{j \in [N]} \|u_j\|_2^2. \quad (19)$$

Multiplying both sides by  $N^2$  allows to conclude.  $\square$

**Lemma 3.** For any vector family  $\vec{u} \in \mathbb{R}^{d \cdot h}$ , we have

$$\Delta_2(\vec{u}) \leq 2 \max_{j \in [h]} \left\| u^{(j)} \right\|_2. \quad (20)$$

*Proof.* We have the inequalities

$$\Delta_2(\vec{u}) = \max_{j, k \in [h]} \left\| u^{(j)} - u^{(k)} \right\|_2 \leq \max_{j, k \in [h]} \left\| u^{(j)} \right\|_2 + \left\| u^{(k)} \right\|_2 \quad (21)$$

$$= \max_{j \in [h]} \left\| u^{(j)} \right\|_2 + \max_{k \in [h]} \left\| u^{(k)} \right\|_2 = 2 \max_{j \in [h]} \left\| u^{(j)} \right\|_2, \quad (22)$$

which is the lemma.  $\square$

We now prove that LEARN solves collaborative learning. Note that all the proofs depend on some quantity  $\alpha_t$ , which will eventually be defined as  $\alpha_t \triangleq \max\{1/\sqrt{t}, \sigma_t\}$ . Note also that we then have  $\alpha_t \leq \bar{\alpha} \triangleq \max\{1, \sigma\}$  since  $b_t \geq 1$ . We also define

$$\xi_t^{(j)} \triangleq g_t^{(j)} - \nabla \mathcal{L}^{(j)}(\theta_t^{(j)}), \quad (23)$$

the gradient estimation error (noise) of node  $j$  at round  $t$ , whose norm is bounded by  $\sigma_t$  (see Section 2.2).

**Lemma 4.** *Under assumptions (2, 3), for any  $0 < \alpha_t \leq \bar{\alpha}$ , we have the following bound on the expected  $\ell_2$  diameter of gradients:*

$$\mathbb{E}_{\xi_t | \bar{\theta}_t} \Delta_2(\vec{g}_t)^2 \leq (1 + \alpha_t)K^2 + 16\bar{\alpha}\alpha_t^{-1} \left( L^2 \Delta_2(\vec{\theta}_t)^2 + h\sigma_t^2 \right). \quad (24)$$

*Proof.* Note that we have

$$g_t^{(j)} = \nabla \mathcal{L}^{(j)}(\theta_t^{(j)}) + \xi_t^{(j)} \quad (25)$$

$$= \nabla \mathcal{L}^{(j)}(\bar{\theta}_t) + \left( \nabla \mathcal{L}^{(j)}(\theta_t^{(j)}) - \nabla \mathcal{L}^{(j)}(\bar{\theta}_t) \right) + \xi_t^{(j)}. \quad (26)$$

Applying Lemma 1 with  $\alpha = \alpha_t^{-1}$ , and then Lemma 2 to the last terms then yields

$$\Delta_2(\vec{g}_t)^2 \leq (1 + \alpha_t) \Delta_2(\vec{\nabla} \bar{\mathcal{L}}(\bar{\theta}_t))^2 \quad (27)$$

$$+ (1 + \alpha_t^{-1}) \left( 2\Delta_2(\vec{\nabla} \bar{\mathcal{L}}(\vec{\theta}_t) - \vec{\nabla} \bar{\mathcal{L}}(\bar{\theta}_t))^2 + 2\Delta_2(\vec{\xi}_t)^2 \right). \quad (28)$$

Note that

$$\Delta_2(\vec{\nabla} \bar{\mathcal{L}}(\bar{\theta}_t))^2 = \max_{j, k \in [h]} \left\| \nabla \mathcal{L}^{(j)}(\bar{\theta}_t) - \nabla \mathcal{L}^{(k)}(\bar{\theta}_t) \right\|_2^2 \leq K^2. \quad (29)$$

The second term can be controlled using Lemma 3, which yields

$$\Delta_2(\vec{\nabla} \bar{\mathcal{L}}(\vec{\theta}_t) - \vec{\nabla} \bar{\mathcal{L}}(\bar{\theta}_t)) \leq 2 \max_{j \in [h]} \left\| \nabla \mathcal{L}^{(j)}(\theta_t^{(j)}) - \nabla \mathcal{L}^{(j)}(\bar{\theta}_t) \right\|_2 \quad (30)$$

$$\leq 2 \max_{j \in [h]} L \left\| \theta_t^{(j)} - \bar{\theta}_t \right\|_2 \leq 2L \Delta_2(\vec{\theta}_t). \quad (31)$$

To bound the third term, first note that Lemma 3 implies that  $\Delta_2(\vec{\xi}_t) \leq 2 \max_{j \in [h]} \left\| \xi_t^{(j)} \right\|_2$ . Thus,

$$\mathbb{E}_{\xi_t | \bar{\theta}_t} \Delta_2(\vec{\xi}_t)^2 \leq 4 \mathbb{E}_{\xi_t | \bar{\theta}_t} \max_{j \in [h]} \left\| \xi_t^{(j)} \right\|_2^2 \leq 4 \mathbb{E}_{\xi_t | \bar{\theta}_t} \sum_{j \in [h]} \left\| \xi_t^{(j)} \right\|_2^2 \quad (32)$$

$$= 4h \mathbb{E}_{\xi_t | \bar{\theta}_t} \left\| \xi_t^{(j)} \right\|_2^2 \leq 4h\sigma_t^2. \quad (33)$$

Combining it all, and using  $1 + \alpha_t^{-1} \leq (\bar{\alpha} + 1)\alpha_t^{-1} \leq 2\bar{\alpha}\alpha_t^{-1}$  for  $\alpha_t \leq \bar{\alpha}$  and  $\bar{\alpha} = \max\{1, \sigma\} \geq 1$ , yields the result.  $\square$

**Definition 3.** *The (stochastic) effective gradient  $G_t^{(j)}$  of node  $j$  is defined by*

$$G_t^{(j)} = \frac{\theta_t^{(j)} - \theta_{t+1}^{(j)}}{\eta}. \quad (34)$$

In particular, we shall focus on the effective gradient of the average parameter, which turns out to also be the average of the effective gradients, that is,

$$\bar{G}_t \triangleq \frac{1}{h} \sum_{j \in [h]} G_t^{(j)} = \frac{\bar{\theta}_t - \bar{\theta}_{t+1}}{\eta}. \quad (35)$$



**Lemma 5.** Under assumptions (2, 3), for any  $0 < \alpha_t \leq \bar{\alpha}$ , the expected discrepancy between the average effective gradient and the true gradient at the average parameter is bounded as follows:

$$\begin{aligned} \mathbb{E}_{\vec{\xi}_t, \text{AVG}|\bar{\theta}_t} \left\| \bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t) \right\|_2^2 &\leq (1 + \alpha_t) C^2 \mathbb{E}_{\vec{\xi}_t|\bar{\theta}_t} \Delta_2(\vec{g}_t)^2 \\ &+ 6\bar{\alpha}\alpha_t^{-1} \left[ \frac{\sigma_t^2}{h} + \left( L^2 + \frac{2C^2}{\eta^2} \right) \Delta_2(\bar{\theta}_t)^2 + \frac{2C^2}{t^2} \mathbb{E}_{\vec{\xi}_t|\bar{\theta}_t} \Delta_2(\vec{g}_t)^2 \right]. \end{aligned} \quad (36)$$

*Proof.* Note that

$$\bar{\theta}_{t+1} - \bar{\theta}_t = (\bar{\theta}_{t+1} - \bar{\theta}_{t+1/2}) + (\bar{\theta}_{t+1/2} - \bar{\theta}_t) \quad (37)$$

$$= (\bar{\theta}_{t+1} - \bar{\theta}_{t+1/2}) - \eta\bar{\gamma}_t. \quad (38)$$

As a result  $\bar{G}_t = \bar{g}_t + (\bar{\gamma}_t - \bar{g}_t) - \frac{1}{\eta}(\bar{\theta}_{t+1} - \bar{\theta}_{t+1/2})$ . Moreover, we have

$$\bar{g}_t = \frac{1}{h} \sum_{j \in [h]} \nabla \mathcal{L}^{(j)}(\theta_t^{(j)}) + \frac{1}{h} \sum_{j \in [h]} \xi_t^{(j)} \quad (39)$$

$$= \nabla \bar{\mathcal{L}}(\bar{\theta}_t) + \frac{1}{h} \sum_{j \in [h]} \left( \nabla \mathcal{L}^{(j)}(\theta_t^{(j)}) - \nabla \mathcal{L}^{(j)}(\bar{\theta}_t) \right) + \frac{1}{h} \sum_{j \in [h]} \xi_t^{(j)}, \quad (40)$$

where  $\nabla \bar{\mathcal{L}}(\bar{\theta}_t) = \frac{1}{h} \sum_{j \in [h]} \nabla \mathcal{L}^{(j)}(\bar{\theta}_t)$  is the average gradient at the average parameter. This then yields:

$$\begin{aligned} \bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t) &= \frac{1}{h} \sum_{j \in [h]} \left( \nabla \mathcal{L}^{(j)}(\theta_t^{(j)}) - \nabla \mathcal{L}^{(j)}(\bar{\theta}_t) \right) + \frac{1}{h} \sum_{j \in [h]} \xi_t^{(j)} \\ &\quad + \frac{1}{\eta} \left( \bar{\theta}_{t+1/2} - \bar{\theta}_{t+1} \right) + (\bar{\gamma}_t - \bar{g}_t). \end{aligned} \quad (41)$$

Applying Lemma 1 for  $\alpha = \alpha_t$  (by isolating the first three terms), and then Lemma 2 to the first three terms then yields

$$\begin{aligned} \left\| \bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t) \right\|_2^2 &\leq 3(1 + \alpha_t^{-1}) \left\| \frac{1}{h} \sum_{j \in [h]} \left( \nabla \mathcal{L}^{(j)}(\theta_t^{(j)}) - \nabla \mathcal{L}^{(j)}(\bar{\theta}_t) \right) \right\|_2^2 \\ &\quad + 3(1 + \alpha_t^{-1}) \left\| \frac{1}{h} \sum_{j \in [h]} \xi_t^{(j)} \right\|_2^2 + \frac{3(1 + \alpha_t^{-1})}{\eta^2} \left\| \bar{\theta}_{t+1/2} - \bar{\theta}_{t+1} \right\|_2^2 \\ &\quad + (1 + \alpha_t) \left\| \bar{\gamma}_t - \bar{g}_t \right\|_2^2. \end{aligned} \quad (42)$$

We now note that the expectation of each term can be bounded. Indeed,

$$\left\| \frac{1}{h} \sum_{j \in [h]} \left( \nabla \mathcal{L}^{(j)}(\theta_t^{(j)}) - \nabla \mathcal{L}^{(j)}(\bar{\theta}_t) \right) \right\|_2 \leq \frac{1}{h} \sum_{j \in [h]} \left\| \nabla \mathcal{L}^{(j)}(\theta_t^{(j)}) - \nabla \mathcal{L}^{(j)}(\bar{\theta}_t) \right\|_2 \quad (43)$$

$$\leq \frac{1}{h} \sum_{j \in [h]} L \left\| \theta_t^{(j)} - \bar{\theta}_t \right\|_2 \leq \frac{1}{h} \sum_{j \in [h]} L \Delta_2(\bar{\theta}_t) = L \Delta_2(\bar{\theta}_t). \quad (44)$$

Moreover, using the conditional non-correlation of  $\xi_t^{(j)}$ , we have

$$\mathbb{E}_{\xi_t|\bar{\theta}_t} \left\| \frac{1}{h} \sum_{j \in [h]} \xi_t^{(j)} \right\|_2^2 = \mathbb{E}_{\xi_t|\bar{\theta}_t} \frac{1}{h^2} \sum_{j,k \in [h]} \xi_t^{(j)} \cdot \xi_t^{(k)} = \frac{1}{h^2} \sum_{j,k \in [h]} \mathbb{E}_{\xi_t|\bar{\theta}_t} \xi_t^{(j)} \cdot \xi_t^{(k)} \quad (45)$$

$$= \frac{1}{h^2} \sum_{j \in [h]} \mathbb{E}_{\xi_t|\bar{\theta}_t} \left\| \xi_t^{(j)} \right\|_2^2 \leq \frac{1}{h^2} \sum_{j \in [h]} \sigma_t^2 = \frac{\sigma_t^2}{h}. \quad (46)$$

For the third term, we use the  $C$ -averaging guarantee of AVG to obtain

$$\mathbb{E}_{\text{AVG}|\vec{\theta}_{t+1/2}} \left\| \bar{\theta}_{t+1} - \bar{\theta}_{t+1/2} \right\|_2^2 \leq C^2 \Delta_2(\vec{\theta}_{t+1/2})^2. \quad (47)$$

Since  $\vec{\theta}_{t+1/2} = \vec{\theta}_t - \eta \vec{\gamma}_t$ , using Lemma 2 and taking the expectation over  $\vec{\xi}_t$ , we then have

$$\mathbb{E}_{\vec{\xi}_t, \text{AVG}|\vec{\theta}_t} \left\| \bar{\theta}_{t+1} - \bar{\theta}_{t+1/2} \right\|_2^2 \leq 2C^2 \Delta_2(\vec{\theta}_t)^2 + 2C^2 \eta^2 \mathbb{E}_{\vec{\xi}_t, \text{AVG}|\vec{\theta}_t} \Delta_2(\vec{\gamma}_t)^2 \quad (48)$$

$$\leq 2C^2 \Delta_2(\vec{\theta}_t)^2 + 2C^2 \eta^2 \frac{\mathbb{E}_{\vec{\xi}_t|\vec{\theta}_t} \Delta_2(\vec{g}_t)^2}{t^2} \quad (49)$$

Finally, for the last term, we again use the  $C$ -averaging guarantee of the aggregation AVG and take the expectation over  $\vec{\xi}_t$ :

$$\mathbb{E}_{\vec{\xi}_t, \text{AVG}|\vec{\theta}_t} \|\bar{\gamma}_t - \bar{g}_t\|_2^2 \leq C^2 \mathbb{E}_{\vec{\xi}_t|\vec{\theta}_t} \Delta_2(\vec{g}_t)^2. \quad (50)$$

Combining it all, and using  $1 + \alpha_t^{-1} \leq 2\bar{\alpha}\alpha_t^{-1}$  finally yields the lemma.  $\square$

**Lemma 6.** *Under assumptions (2, 3), for any  $0 < \alpha_t \leq \bar{\alpha}$  and  $\alpha_t \geq 1/\sqrt{t}$ , there exist constants  $A$  and  $B$  which can be computed explicitly given  $\bar{\alpha}$ ,  $C$ ,  $L$ ,  $h$  and  $\eta$ , such that*

$$\mathbb{E}_{\vec{\xi}_t, \text{AVG}|\vec{\theta}_t} \left\| \bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t) \right\|_2^2 \leq (1 + \alpha_t)^2 (1 + \kappa_t) C^2 K^2 + \alpha_t^{-1} \left( A \Delta_2(\vec{\theta}_t)^2 + B \sigma_t^2 \right), \quad (51)$$

where  $\kappa_t \leq \frac{12\bar{\alpha}}{t^{(3/2)}}$ .

*Proof.* Combining the two previous lemmas, this bound can be guaranteed by setting

$$\kappa_t = \frac{12\bar{\alpha}}{(1 + \alpha_t)\alpha_t t^2} \quad (52)$$

$$A_t = 16(1 + \alpha_t)\bar{\alpha}C^2L^2 + 6\bar{\alpha}L^2 + \frac{12\bar{\alpha}C^2}{\eta^2} + \frac{192\bar{\alpha}^2C^2L^2}{\alpha_t t^2} \quad (53)$$

$$B_t = 16(1 + \alpha_t)\bar{\alpha}C^2h + \frac{6\bar{\alpha}}{h} + \frac{192\bar{\alpha}^2C^2h}{\alpha_t t^2}. \quad (54)$$

Assumptions  $0 < \alpha_t \leq \bar{\alpha}$  (which implies  $1 + \alpha_t \leq 2\bar{\alpha}$ ) and  $\alpha_t \sqrt{t} \geq 1$  allow to conclude, with

$$\kappa_t \leq \frac{12\bar{\alpha}}{t^{(3/2)}} \quad (55)$$

$$A = 32\bar{\alpha}^2C^2L^2 + 6\bar{\alpha}L^2 + \frac{12\bar{\alpha}C^2}{\eta^2} + 192\bar{\alpha}^2C^2L^2 \quad (56)$$

$$B = 32\bar{\alpha}^2C^2h + \frac{6\bar{\alpha}}{h} + 192\bar{\alpha}^2C^2h. \quad (57)$$

This shows in particular that  $A$  and  $B$  can indeed be computed from the different constants of the problem.  $\square$

**Lemma 7.** *We have the following bound on parameter drift:*

$$\mathbb{E}_{\text{AVG}|\vec{\theta}_t, \vec{g}_t} \Delta_2(\vec{\theta}_{t+1})^2 \leq \frac{1}{2} \Delta_2(\vec{\theta}_t)^2 + \frac{\eta^2}{2t^2} \Delta_2(\vec{g}_t)^2. \quad (58)$$

*Proof.* Recall that  $\vec{\theta}_{t+1} = \overrightarrow{\text{AVG}}_1 \circ \overrightarrow{\text{BYZ}}_{t, \theta}(\vec{\theta}_{t+1/2})$ . Thus, by the asymptotic agreement property of  $\text{AVG}_1$ , we know that

$$\mathbb{E}_{\text{AVG}|\vec{\theta}_t, \vec{g}_t} \Delta_2(\vec{\theta}_{t+1})^2 \leq \frac{1}{4} \Delta_2(\vec{\theta}_{t+1/2})^2. \quad (59)$$

Now recall that  $\vec{\theta}_{t+1/2} = \vec{\theta}_t - \eta \vec{\gamma}_t$ . Applying Lemma 1 for  $\alpha = 1$  thus yields

$$\Delta_2(\vec{\theta}_{t+1/2})^2 \leq 2\Delta_2(\vec{\theta}_t)^2 + 2\eta^2 \Delta_2(\vec{\gamma}_t)^2. \quad (60)$$

We now use the asymptotic agreement property of  $\text{AVG}_{N(t)}$ , which yields

$$\mathbb{E}_{\text{AVG}|\vec{\theta}_t, \vec{g}_t} \Delta_2(\vec{\gamma}_t)^2 \leq \frac{\Delta_2(\vec{g}_t)^2}{t^2}. \quad (61)$$

Combining it all then yields the result.  $\square$

**Lemma 8.** *Under assumptions (2, 3),  $0 < \alpha_t \leq \bar{\alpha}$  and  $\alpha_t = \max\{1/\sqrt{t}, \sigma_t\}$ , there exists a constant  $D$  such that*

$$\mathbb{E}_{\vec{\xi}_{1:t}} \Delta_2(\vec{\theta}_t)^2 \leq D/t^2. \quad (62)$$

Note that the constant  $D$  can be computed from the constants  $L, \eta, K, h$  and the functions  $\sigma_t$  and  $\alpha_t$ .

*Proof.* Denote  $u_t = \mathbb{E}_{\vec{\xi}_{1:t}} \Delta_2(\vec{\theta}_t)^2$ . Combining Lemmas 4 and 7 yields

$$u_{t+1} \leq \rho_t u_t + \delta_t, \quad (63)$$

where  $\rho_t$  and  $\delta_t$  are given by

$$\rho_t \triangleq \frac{1}{2} + \frac{8\bar{\alpha}L^2\eta^2}{\alpha_t t^2} \quad (64)$$

$$\delta_t \triangleq \frac{\eta^2}{2t^2} \left( (1 + \alpha_t)K^2 + 16\bar{\alpha}\alpha_t^{-1}h\sigma_t^2 \right). \quad (65)$$

Given that  $\alpha_t \geq 1/\sqrt{t}$ , we know that  $\rho_t \leq \frac{1}{2} + 8\bar{\alpha}L^2\eta^2 t^{-3/2}$ . Thus, for  $t \geq t_0 \triangleq (32\bar{\alpha}L^2\eta^2)^{2/3}$ , we know that  $\rho_t \leq \rho \triangleq 3/4$ . Moreover, using now  $\sigma_t \leq \alpha_t \leq \bar{\alpha}$ , we know that

$$\delta_t \leq \frac{\eta^2}{2t^2} \left( (1 + \bar{\alpha})K^2 + 16\bar{\alpha}h\sigma_t \right) \leq \frac{\eta^2}{2t^2} \left( (1 + \bar{\alpha})K^2 + 16h\bar{\alpha}^2 \right) \triangleq \delta_t^+. \quad (66)$$

Now note that  $\delta_t^+$  is decreasing. In particular, for  $t \geq t_0$  we now have

$$u_{t+1} \leq \rho u_t + \delta_t^+. \quad (67)$$

By induction we see that, for  $t \geq 1$ , we have

$$u_{t+t_0} \leq \rho^t u_{t_0} + \sum_{s=0}^{t-1} \rho^s \delta_{t+t_0-s-1}^+. \quad (68)$$

We now separate the sum into two parts. Calling  $t_1$  the separation point for  $t_1 \geq 1$ , and using the fact that  $\delta_t^+$  is decreasing yields

$$u_{t+t_0} \leq \rho^t u_{t_0} + \sum_{s=0}^{t_1-1} \rho^s \delta_{t+t_0-s-1}^+ + \sum_{s=t_1}^{t-1} \rho^s \delta_{t+t_0-s-1}^+ \quad (69)$$

$$\leq \rho^t u_{t_0} + \delta_{t+t_0-t_1}^+ \sum_{s=0}^{t_1-1} \rho^s + \delta_{t_0}^+ \sum_{s=t_1}^{t-1} \rho^s \quad (70)$$

$$\leq \rho^t u_{t_0} + \delta_{t+t_0-t_1}^+ \sum_{s=0}^{\infty} \rho^s + \delta_{t_0}^+ \sum_{s=t_1}^{\infty} \rho^s \quad (71)$$

$$\leq \rho^t u_{t_0} + \frac{\delta_{t+t_0-t_1}^+}{1-\rho} + \frac{\rho^{t_1} \delta_{t_0}^+}{1-\rho} \quad (72)$$

$$= \rho^t u_{t_0} + 4\delta_{t+t_0-t_1}^+ + 4\rho^{t_1} \delta_{t_0}^+. \quad (73)$$

We now take  $t_1 = \lfloor \frac{t+t_0}{2} \rfloor$ . As a result,

$$\delta_{t+t_0-t_1}^+ = \delta_{\lfloor \frac{t+t_0}{2} \rfloor}^+ \leq \frac{2\eta^2}{(t+t_0)^2} \left( (1 + \bar{\alpha})K^2 + 16h\bar{\alpha}^2 \right). \quad (74)$$

Now define  $v_t$  by  $v_0 = 0$  and  $v_{t+1} = \rho_t v_t + \delta_t$ . Note that  $u_{t_0}$  can be upper-bounded given  $L, \eta, \alpha_{1:t_0}, K, h$  and  $\sigma_{1:t_0}$ , by computing  $v_{t_0}$ . Indeed, by induction we then clearly have  $u_t \leq v_{t_0}$ , and thus the bound

$$u_{t+t_0} \leq \frac{8\eta^2 ((1 + \bar{\alpha})K^2 + 16h\bar{\alpha}^2)}{(t + t_0)^2} + \rho^t v_{t_0} + 4\rho^{t_1} \delta_{t_0}^+, \quad (75)$$

where the right-hand side is perfectly computable given the constants of the problem. Given that  $\rho^{t_1} = \mathcal{O}(1/t^2)$  and  $\rho^t = \mathcal{O}(1/t^2)$ , we can also compute a constant  $D$  from these constants, such that for all iterations  $t$ , we have  $u_t \leq D/t^2$ .  $\square$

### A.1 Reduction from collaborative learning to averaging agreement

Now we proceed with the proof of our theorem.

*Proof of Theorem 1.* At any iteration  $t$ , Taylor's theorem implies the existence of  $\lambda \in [0, 1]$  such that

$$\bar{\mathcal{L}}(\bar{\theta}_{t+1}) = \bar{\mathcal{L}}(\bar{\theta}_t - \eta \bar{G}_t) \quad (76)$$

$$= \bar{\mathcal{L}}(\bar{\theta}_t) - \eta \bar{G}_t \cdot \nabla \bar{\mathcal{L}}(\bar{\theta}_t) + \frac{1}{2} (\eta \bar{G}_t)^T \nabla^2 \bar{\mathcal{L}}(\bar{\theta}_t - \lambda \eta \bar{G}_t) (\eta \bar{G}_t). \quad (77)$$

Lipschitz continuity of the gradient implies that  $\nabla^2 \bar{\mathcal{L}}(\bar{\theta}_t - \lambda \eta \bar{G}_t) \preceq LI$ , which thus implies

$$\bar{\mathcal{L}}(\bar{\theta}_{t+1}) \leq \bar{\mathcal{L}}(\bar{\theta}_t) - \eta \bar{G}_t \cdot \nabla \bar{\mathcal{L}}(\bar{\theta}_t) + \frac{L\eta^2}{2} \|\bar{G}_t\|_2^2. \quad (78)$$

For the second term, using the inequality  $2u \cdot v \geq -\|u\|_2^2 - \|v\|_2^2$ , note that

$$\bar{G}_t \cdot \nabla \bar{\mathcal{L}}(\bar{\theta}_t) = (\bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t) + \nabla \bar{\mathcal{L}}(\bar{\theta}_t)) \cdot \nabla \bar{\mathcal{L}}(\bar{\theta}_t) \quad (79)$$

$$= (\bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t)) \cdot \nabla \bar{\mathcal{L}}(\bar{\theta}_t) + \|\nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2 \quad (80)$$

$$\geq -\frac{1}{2} \|\bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2 - \frac{1}{2} \|\nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2 + \|\nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2 \quad (81)$$

$$= -\frac{1}{2} \|\bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2 + \frac{1}{2} \|\nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2. \quad (82)$$

For the last term, we use  $\|a + b\|_2^2 \leq 2\|a\|_2^2 + 2\|b\|_2^2$  to derive

$$\|\bar{G}_t\|_2^2 = \|\bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t) + \nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2 \quad (83)$$

$$\leq 2\|\bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2 + 2\|\nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2. \quad (84)$$

Combining the two above bounds into Equation (78) yields

$$\bar{\mathcal{L}}(\bar{\theta}_{t+1}) \leq \bar{\mathcal{L}}(\bar{\theta}_t) - \left(\frac{\eta}{2} - L\eta^2\right) \|\nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2 + \left(\frac{\eta}{2} + L\eta^2\right) \|\bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2. \quad (85)$$

Rearranging the terms then yields

$$\left(\frac{\eta}{2} - L\eta^2\right) \|\nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2 \leq \bar{\mathcal{L}}(\bar{\theta}_t) - \bar{\mathcal{L}}(\bar{\theta}_{t+1}) + \left(\frac{\eta}{2} + L\eta^2\right) \|\bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2. \quad (86)$$

We now use the fact that  $\eta \leq \delta/12L$ . Denoting  $\nu \triangleq \delta/6$ , this implies that  $\frac{\eta}{2} - L\eta^2 \geq (1 - \nu)\eta/2$  and  $\frac{\eta}{2} + L\eta^2 \leq (1 + \nu)\eta/2$ . As a result,

$$\|\nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2 \leq \frac{4}{\eta} (\bar{\mathcal{L}}(\bar{\theta}_t) - \bar{\mathcal{L}}(\bar{\theta}_{t+1})) + \frac{1 + \nu}{1 - \nu} \|\bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2. \quad (87)$$

Taking the expectation and the average over  $t \in [T]$  yields

$$\mathbb{E}_{\xi_{1:T}} \frac{1}{T} \sum_{t \in [T]} \|\nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2 \leq \frac{4(\bar{\mathcal{L}}(\bar{\theta}_1) - \bar{\mathcal{L}}(\bar{\theta}_{T+1}))}{\eta T} + \frac{1 + \nu}{1 - \nu} \frac{1}{T} \sum_{t \in [T]} \mathbb{E}_{\xi_{1:T}} \|\bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2. \quad (88)$$

Note that  $\mathbb{E}_{\vec{\xi}_{1:T}} \frac{1}{T} \sum_{t \in [T]} \|\nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2 = \mathbb{E} \|\nabla \bar{\mathcal{L}}(\bar{\theta}_*)\|_2^2$ , since the second term is obtained by taking uniformly randomly one of the values averaged in the first term. Using also the fact that  $\bar{\mathcal{L}}(\bar{\theta}_{T+1}) \geq \inf_{\theta} \bar{\mathcal{L}}(\theta) \geq 0$  (Assumption 1) and  $\mathcal{L}^{(j)}(\theta_1) \leq \mathcal{L}_{max}$ , we then obtain

$$\mathbb{E} \|\nabla \bar{\mathcal{L}}(\bar{\theta}_*)\|_2^2 \leq \frac{4\mathcal{L}_{max}}{\eta T} + \frac{1+\nu}{1-\nu} \frac{1}{T} \sum_{t \in [T]} \mathbb{E}_{\vec{\xi}_{1:T}} \|\bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2. \quad (89)$$

Now recall that all nodes started with the same value  $\theta_1$ , and thus know  $\bar{\theta}_1$ . As a result, each node  $j$  can compute  $\mathcal{L}^{(j)}(\bar{\theta}_1)$ , but it cannot compute  $\bar{\mathcal{L}}(\bar{\theta}_1)$ .

Let us focus on the last term. Taking Lemma 6 and averaging over all noises  $\vec{\xi}_{1:t}$  yields

$$\mathbb{E}_{\vec{\xi}_{1:t}} \|\bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2 \leq (1 + \alpha_t)^2 (1 + \kappa_t)^2 C^2 K^2 + \alpha_t^{-1} \left( A \mathbb{E}_{\vec{\xi}_{1:t}} \Delta_2(\bar{\theta}_t)^2 + B\sigma_t^2 \right). \quad (90)$$

Recall that, by Lemma 8,  $\mathbb{E}_{\vec{\xi}_{1:t}} \Delta_2(\bar{\theta}_t)^2 \leq D/t^2$ . Recall also that  $\alpha_t \triangleq \max\{1/\sqrt{t}, \sigma_t\}$  and  $\kappa_t \leq 12\bar{\alpha}/t^{(3/2)}$ . Then we obtain

$$\mathbb{E}_{\vec{\xi}_{1:t}} \|\bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2 \leq (1 + \alpha_t)^2 \left( 1 + \frac{12\bar{\alpha}}{t^{3/2}} \right) C^2 K^2 + \frac{AD}{t^{3/2}} + B\sigma_t. \quad (91)$$

Now recall that we regularly increase the batch size (see Section 2.2). Thus, there exists some iteration  $T_1$ , such that  $\sigma_{T_1} \leq \min\{\nu, \nu C^2 K^2 / 8B\} = \mathcal{O}(\delta)$ . Defining  $T_2 \triangleq 1/\nu^2$ ,  $T_3 \triangleq (12\bar{\alpha}/\nu)^{2/3}$  and  $T_4 \triangleq (8AD/\nu C^2 K^2)^{2/3}$ , for  $t \geq T_5 \triangleq \max\{T_1, T_2, T_3, T_4\}$ , we have

$$\mathbb{E}_{\vec{\xi}_{1:t}} \|\bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2 \leq (1 + \nu)^3 C^2 K^2 + \nu C^2 K^2 / 8 + \nu C^2 K^2 / 8 \quad (92)$$

$$\leq (1 + 5\nu) C^2 K^2, \quad (93)$$

using the inequality  $\nu \leq 1/2$  to show that  $(1 + \nu)^3 \leq 1 + 3\nu + 3\nu^2 + \nu^3 \leq 1 + 3\nu + 3\nu/2 + \nu/4$ . If we now average this quantity over  $t$  from 1 to  $T$ , assuming  $T \geq T_5$ , we can separate the sum from 1 to  $T_5$ , and the sum from  $T_5 + 1$  to  $T$ . This yields

$$\frac{1}{T} \sum_{t \in T} \mathbb{E}_{\vec{\xi}_{1:t}} \|\bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2 \leq \frac{T_5 E}{T} + \frac{T - T_5}{T} (1 + 5\nu) C^2 K^2, \quad (94)$$

where  $E = (1 + \bar{\alpha})^2 (1 + 12\bar{\alpha}) C^2 K^2 + AD + B\sigma$ . Now consider  $T_6 \triangleq T_5 E / \nu C^2 K^2$ . For  $T \geq T_6$ , we then have

$$\frac{1}{T} \sum_{t \in T} \mathbb{E}_{\vec{\xi}_{1:t}} \|\bar{G}_t - \nabla \bar{\mathcal{L}}(\bar{\theta}_t)\|_2^2 \leq (1 + 6\nu) C^2 K^2. \quad (95)$$

Plugging this into Equation (89), and using  $1/(1 - \nu) \leq 1 + 2\nu$  for  $0 < \nu \leq 1/2$  then yields, for  $T \geq T_6$ ,

$$\mathbb{E} \|\nabla \bar{\mathcal{L}}(\bar{\theta}_*)\|_2^2 \leq \frac{4\mathcal{L}_{max}}{\eta T} + (1 + \nu)(1 + 2\nu)(1 + 6\nu) C^2 K^2. \quad (96)$$

Now note that  $(1 + \nu)(1 + 2\nu) \leq 1 + 4\nu$  for  $\nu \leq 1/2$ . Now consider  $T_7 \triangleq \mathcal{L}_{max} / (\nu^2 \eta C^2 K^2)$ . Then for  $T \geq T_8 \triangleq \max\{T_6, T_7\}$ , we have the guarantee

$$\mathbb{E} \|\nabla \bar{\mathcal{L}}(\bar{\theta}_*)\|_2^2 \leq ((1 + 4\nu)(1 + 6\nu) + \nu^2) C^2 K^2 \leq (1 + 6\nu)^2 C^2 K^2. \quad (97)$$

Now recall that  $\nu = \delta/6$ , and consider  $T_9 \triangleq D\tau^2 / 24\delta^2$  and<sup>8</sup>  $T = T_{\text{LEARN}}(\delta) \triangleq \max\{T_8, T_9\}$ . Note that we have

$$\mathbb{E} \Delta_2(\bar{\theta}_*)^2 = \frac{1}{T} \sum_{t \in [T]} \mathbb{E} \Delta_2(\bar{\theta}_t)^2 \leq \frac{1}{T} \sum_{t \in [T]} \frac{D}{t^2} \quad (98)$$

$$\leq \frac{D}{T} \sum_{t=1}^{\infty} \frac{1}{t^2} = \frac{D}{T} \frac{\tau^2}{24} \leq \delta^2. \quad (99)$$

Combining it all yields

$$\mathbb{E} \Delta_2(\bar{\theta}_*)^2 \leq \delta^2 \quad \text{and} \quad \mathbb{E} \|\nabla \bar{\mathcal{L}}(\bar{\theta}_*)\|_2^2 \leq (1 + \delta)^2 C^2 K^2, \quad (100)$$

which corresponds to saying that LEARN solves collaborative learning.  $\square$

<sup>8</sup>Here,  $\tau \approx 6.2832$  is the ratio of the circumference of the circle by its radius.

## A.2 Reduction from averaging agreement to collaborative learning

*Proof of Theorem 2.* Without loss of generality, assume  $0 < \delta \leq 1$ . Let  $\vec{x} \in \mathbb{R}^{d \cdot h}$  be a family of vectors. For any honest node  $j \in [h]$ , consider the losses defined by  $\mathcal{L}^{(j)}(\theta) \triangleq \frac{1}{2} \|\theta - x^{(j)}\|_2^2$ . Note that we thus have  $\nabla \mathcal{L}^{(j)}(\theta) = \theta - x^{(j)}$ . As a result,

$$\left\| \nabla \mathcal{L}^{(j)}(\theta) - \nabla \mathcal{L}^{(k)}(\theta) \right\|_2 = \left\| x^{(j)} - x^{(k)} \right\|_2 \leq \Delta_2(\vec{x}), \quad (101)$$

which corresponds to saying that local losses satisfy the definition of collaborative learning (Definition 1) with  $K = \Delta_2(\vec{x})$ . Now consider a Byzantine-resilient  $C$ -collaborative learning algorithm `LEARN`. For any  $N \in \mathbb{N}$ , we run `LEARN` with parameter  $\delta \triangleq \min\{1, \Delta_2(\vec{x})/2^N\}$ , which outputs  $\vec{x}_N \triangleq \vec{\theta}_*$ .

We then have the guarantee  $\mathbb{E} \Delta_2(\vec{x}_N)^2 = \mathbb{E} \Delta_2(\vec{\theta}_*)^2 \leq \delta^2 \leq \Delta_2(\vec{x})^2/4^N$ , which corresponds to asymptotic agreement. Moreover, we notice that

$$\nabla \bar{\mathcal{L}}(\vec{\theta}_*) = \frac{1}{h} \sum_{j \in [h]} (\vec{\theta}_* - x^{(j)}) = \vec{\theta}_* - \bar{x} = \vec{x}_N - \bar{x}. \quad (102)$$

The second collaborative learning guarantee of algorithm `LEARN` (Equation (4) in the paper) then yields

$$\mathbb{E} \|\vec{x}_N - \bar{x}\|_2^2 = \mathbb{E} \|\nabla \bar{\mathcal{L}}(\vec{\theta}_*)\|_2^2 \leq (1 + \delta)^2 C^2 K^2 = ((1 + \delta)C)^2 \Delta_2(\vec{x})^2. \quad (103)$$

This shows  $(1 + \delta)C$ -averaging, and concludes the proof.  $\square$

## B Efficient i.i.d. algorithm

From a practical viewpoint, a disadvantage of Algorithm 1 is that it requires a large number of communication rounds due to the use of two instances of the averaging algorithm `AVG`, one of which requires a stronger agreement (which induces more communication rounds) as the iteration number  $t$  grows. This is because when the data distributions of nodes vary from each other, data heterogeneity enables the Byzantines to bias the models in their favor and induce model drift more easily.

In this section, we show that we can perform better in the homogeneous (i.i.d.) setting, where the ability of Byzantines is more limited. More specifically, we present a simpler algorithm (Algorithm 2) that uses the averaging algorithm `AVG` only once at each iteration with a fixed parameter ( $N = 1$ ). This results in much lower communication/computation overhead compared to Algorithm 1. We evaluate the throughput overhead of our two algorithms compared to the non-robust vanilla implantation in Section 5, which, not surprisingly, shows the superiority of the i.i.d. algorithm.

Note that in the homogeneous setting, all local losses are equal, i.e.,  $\mathcal{L}^{(j)}(\cdot) = \mathcal{L}(\cdot)$  for all honest nodes  $j \in [h]$ .

<p><b>Data:</b> Global loss gradient oracle  <b>Result:</b> Model parameters <math>\theta_t</math></p> <pre style="margin: 0;"> 1 Initialize local parameters <math>\theta_1</math> using a fixed seed <math>s</math>; 2 Fix learning rate <math>\eta \leq \frac{1}{2L}</math>; 3 Fix number of rounds <math>T \triangleq T_{\text{HOM-LEARN}}(\delta)</math>; 4 <b>for</b> <math>t \leftarrow 1, \dots, T</math> <b>do</b> 5   <math>g_t \leftarrow \text{GradientOracle}(\theta_t)</math>; 6   <math>\theta_{t+1/2} \leftarrow \theta_t - \eta g_t</math>; 7   <math>\theta_{t+1} \leftarrow \text{AVG}_1 \circ \text{BYZ}_{t,\theta}(\vec{\theta}_{t+1/2})</math> // Vulnerable to Byzantine attacks 8 <b>end</b> 9 Draw <math>*</math> <math>\sim \mathcal{U}([T])</math>; 10 Return <math>\theta_*</math>;</pre>
---

**Algorithm 2:** HOM-LEARN execution on an honest node.

### B.1 Proof of Theorem 3

Before proving our theorem, we prove some preliminary lemmas.

**Lemma 9.** *Under Assumption 3, we have*

$$\mathbb{E}_{\vec{\xi}_t | \vec{\theta}_t} \Delta_2(\vec{\xi}_t)^2 \leq 4\sigma_t^2 h, \quad (104)$$

where  $\xi_t^{(j)}$  is the gradient estimation error of node  $j$  at round  $t$  defined in (23).

*Proof.*

$$\mathbb{E}_{\vec{\xi}_t | \vec{\theta}_t} \Delta_2(\vec{\xi}_t)^2 = \mathbb{E}_{\vec{\xi}_t | \vec{\theta}_t} \max_{j, k \in [h]} \left\| \xi_t^{(j)} - \xi_t^{(k)} \right\|_2^2 \quad (105)$$

$$\leq 4 \mathbb{E}_{\vec{\xi}_t | \vec{\theta}_t} \max_{j \in [h]} \left\| \xi_t^{(j)} \right\|_2^2, \quad (106)$$

where we used Lemma 3. We now use the fact that the maximum over nodes  $j \in [h]$  is smaller than the sum over nodes  $j \in [h]$ , yielding

$$\mathbb{E}_{\vec{\xi}_t | \vec{\theta}_t} \Delta_2(\vec{\xi}_t)^2 \leq 4 \mathbb{E}_{\vec{\xi}_t | \vec{\theta}_t} \sum_{j \in [h]} \left\| \xi_t^{(j)} \right\|_2^2 = 4 \sum_{j \in [h]} \mathbb{E}_{\vec{\xi}_t | \vec{\theta}_t} \left\| \xi_t^{(j)} \right\|_2^2 \quad (107)$$

$$\leq 4 \sum_{j \in [h]} \sigma_t^2 = 4\sigma_t^2 h, \quad (108)$$

where the last inequality uses Assumption 3.  $\square$

**Lemma 10.** *Under assumptions (2,3), the expected  $\ell_2$  diameter between honest gradient estimations is upper-bounded as follows*

$$\mathbb{E}_{\vec{\xi}_t | \vec{\theta}_t} \Delta_2(\vec{g}_t)^2 \leq 2L^2 \Delta_2(\vec{\theta}_t)^2 + 8\sigma_t^2 h. \quad (109)$$

*Proof.* By Lemma 2, we know that

$$\Delta_2(\vec{g}_t)^2 \leq 2\Delta_2(\nabla \mathcal{L}(\vec{\theta}_t))^2 + 2\Delta_2(\vec{\xi}_t)^2. \quad (110)$$

Assumption 2 then guarantees that

$$\Delta_2(\nabla \mathcal{L}(\vec{\theta}_t))^2 = \max_{j, k \in [h]} \left\| \nabla \mathcal{L}(\theta_t^{(j)}) - \nabla \mathcal{L}(\theta_t^{(k)}) \right\|_2^2 \quad (111)$$

$$\leq \max_{j, k \in [h]} L^2 \left\| \theta_t^{(j)} - \theta_t^{(k)} \right\|_2^2 \quad (112)$$

$$= L^2 \max_{j, k \in [h]} \left\| \theta_t^{(j)} - \theta_t^{(k)} \right\|_2^2 = L^2 \Delta_2(\vec{\theta}_t)^2. \quad (113)$$

Combining this with the previous lemmas completes the proof.  $\square$

**Lemma 11.** *Under assumptions (2, 3), we have*

$$\mathbb{E}_{\vec{\xi}_t | \vec{\theta}_t} \Delta_2(\vec{\theta}_{t+1})^2 \leq \left( \frac{1}{2} + L^2 \eta^2 \right) \Delta_2(\vec{\theta}_t)^2 + 4\sigma_t^2 \eta^2 h. \quad (114)$$

*Proof.* We first bound the diameter of  $\theta_{t+1/2}$ , using Lemma 2 and the bound of Lemma 10. This yields

$$\mathbb{E}_{\vec{\xi}_t | \vec{\theta}_t} \Delta_2(\theta_{t+1/2})^2 = \mathbb{E}_{\vec{\xi}_t | \vec{\theta}_t} \Delta_2(\vec{\theta}_t - \eta \vec{g}_t)^2 \leq 2\Delta_2(\vec{\theta}_t)^2 + 2\Delta_2(\eta \vec{g}_t)^2 \quad (115)$$

$$= 2\Delta_2(\vec{\theta}_t)^2 + 2\eta^2 \Delta_2(\vec{g}_t)^2 \leq (2 + 4L^2 \eta^2) \Delta_2(\vec{\theta}_t)^2 + 16\eta^2 \sigma_t^2 h. \quad (116)$$

We now apply the asymptotic agreement guarantee of AVG, which yields

$$\mathbb{E}_{\xi_t | \bar{\theta}_t} \Delta_2 \left( \bar{\theta}_{t+1} \right)^2 \leq \frac{1}{4} \mathbb{E}_{\xi_t | \bar{\theta}_t} \Delta_2 \left( \bar{\theta}_{t+1/2} \right)^2 \quad (117)$$

$$\leq \frac{1}{4} \left( (2 + 4L^2\eta^2) \Delta_2 \left( \bar{\theta}_t \right)^2 + 16\eta^2\sigma_t^2 h \right) \quad (118)$$

$$\leq \left( \frac{1}{2} + L^2\eta^2 \right) \Delta_2 \left( \bar{\theta}_t \right)^2 + 4\eta^2\sigma_t^2 h, \quad (119)$$

which is the lemma.  $\square$

**Lemma 12.** *Under assumptions (2, 3), the diameter of the parameters is upper-bounded, i.e., there exists a constant  $D$  such that for any  $t \geq 1$  we have*

$$\mathbb{E}_{\xi_{1:t}} \Delta_2 \left( \bar{\theta}_t \right)^2 \leq D. \quad (120)$$

Moreover, for any  $\varepsilon > 0$ , there exists an iteration  $T^*(\varepsilon)$ , such that for all  $t \geq T^*(\varepsilon)$ , we have

$$\mathbb{E}_{\xi_{1:t}} \Delta_2 \left( \bar{\theta}_t \right)^2 \leq \varepsilon. \quad (121)$$

*Proof.* We know that  $\eta \leq 1/2L$ . Denoting  $u_t \triangleq \mathbb{E} \Delta_2 \left( \bar{\theta}_t \right)^2$ , by Lemma 11, we have

$$u_{t+1} \leq \frac{3}{4} u_t + 4\sigma_t^2 \eta^2 h. \quad (122)$$

By induction, we observe that, for all  $t \geq 1$ ,

$$u_{t+1} \leq 4\eta^2 h \sum_{\tau=0}^{t-1} \left( \frac{3}{4} \right)^\tau \sigma_{t-\tau}^2. \quad (123)$$

Now recall that  $\sigma_t$  is decreasing (see Section 2.2), thus, for all  $t \geq 1$ , we know  $\sigma_t \leq \sigma_1$ . Therefore,

$$u_{t+1} \leq 4\eta^2 h \sigma_1^2 \sum_{\tau=0}^{t-1} \left( \frac{3}{4} \right)^\tau \leq 4\eta^2 h \sigma_1^2 \sum_{\tau=0}^{\infty} \left( \frac{3}{4} \right)^\tau = 16\eta^2 h \sigma_1^2 \triangleq D. \quad (124)$$

For the second part of the lemma, recall also from Section 2.2 that we regularly increase the batch size. Thus, there is an iteration  $T_1(\varepsilon)$  such that for all  $t \geq T_1(\varepsilon)$  we have  $\sigma_t^2 \leq \frac{\varepsilon}{32\eta^2 h}$ . By (123), we then have

$$u_{t+1} \leq 4\eta^2 h \sum_{\tau=0}^{t-T_1(\varepsilon)} \left( \frac{3}{4} \right)^\tau \sigma_{t-\tau}^2 + 4\eta^2 h \sum_{\tau=t-T_1(\varepsilon)+1}^{t-1} \left( \frac{3}{4} \right)^\tau \sigma_{t-\tau}^2 \quad (125)$$

$$\leq \frac{4\eta^2 h \varepsilon}{32\eta^2 h} \sum_{\tau=0}^{t-T_1(\varepsilon)} \left( \frac{3}{4} \right)^\tau + 4\eta^2 h \sigma_1^2 \sum_{\tau=t-T_1(\varepsilon)+1}^{t-1} \left( \frac{3}{4} \right)^\tau \quad (126)$$

$$\leq \frac{\varepsilon}{8} \sum_{\tau=0}^{\infty} \left( \frac{3}{4} \right)^\tau + 4\eta^2 h \sigma_1^2 \left( \frac{3}{4} \right)^{t-T_1(\varepsilon)+1} \sum_{\tau=0}^{\infty} \left( \frac{3}{4} \right)^\tau \quad (127)$$

$$\leq \frac{\varepsilon}{2} + 16\eta^2 h \sigma_1^2 \left( \frac{3}{4} \right)^{t-T_1(\varepsilon)+1}. \quad (128)$$

Defining  $T^*(\varepsilon) \triangleq T_1(\varepsilon) - 1 + \frac{\ln(32\eta^2 h \sigma_1^2 / \varepsilon)}{\ln(4/3)}$ , we then have  $u_{t+1} \leq \varepsilon$  for  $t \geq T^*(\varepsilon)$ , which is what we wanted.  $\square$



**Lemma 13.** *Under assumptions (2, 3), there exist constants  $A$  and  $B$ , such that for all  $t \geq 1$ , we have*

$$\mathbb{E}_{\bar{\xi}_t | \bar{\theta}_t} \|\bar{G}_t - \nabla \mathcal{L}(\bar{\theta}_t)\|_2^2 \leq A\sigma_t^2 + B\Delta_2(\bar{\theta}_t)^2, \quad (129)$$

where  $\bar{G}_t$  is the average of the effective gradients of the nodes defined in (35).

*Proof.* Note that

$$\bar{\theta}_{t+1} - \bar{\theta}_t = (\bar{\theta}_{t+1} - \bar{\theta}_{t+1/2}) + (\bar{\theta}_{t+1/2} - \bar{\theta}_t) \quad (130)$$

$$= (\bar{\theta}_{t+1} - \bar{\theta}_{t+1/2}) - \eta \bar{g}_t. \quad (131)$$

As a result,  $\bar{G}_t = \bar{g}_t - \frac{1}{\eta}(\bar{\theta}_{t+1} - \bar{\theta}_{t+1/2})$ . Moreover, we have

$$\bar{g}_t = \frac{1}{h} \sum_{j \in [h]} \nabla \mathcal{L}(\theta_t^{(j)}) + \frac{1}{h} \sum_{j \in [h]} \xi_t^{(j)} \quad (132)$$

$$= \nabla \mathcal{L}(\bar{\theta}_t) + \frac{1}{h} \sum_{j \in [h]} (\nabla \mathcal{L}(\theta_t^{(j)}) - \nabla \mathcal{L}(\bar{\theta}_t)) + \frac{1}{h} \sum_{j \in [h]} \xi_t^{(j)}. \quad (133)$$

This then yields:

$$\bar{G}_t - \nabla \mathcal{L}(\bar{\theta}_t) = \frac{1}{h} \sum_{j \in [h]} (\nabla \mathcal{L}(\theta_t^{(j)}) - \nabla \mathcal{L}(\bar{\theta}_t)) + \frac{1}{h} \sum_{j \in [h]} \xi_t^{(j)} + \frac{\bar{\theta}_{t+1/2} - \bar{\theta}_{t+1}}{\eta}. \quad (134)$$

Taking the  $\ell_2$  norm on both sides and invoking Lemma 2 then implies that

$$\begin{aligned} \|\bar{G}_t - \nabla \mathcal{L}(\bar{\theta}_t)\|_2^2 &\leq 3 \left\| \frac{1}{h} \sum_{j \in [h]} (\nabla \mathcal{L}(\theta_t^{(j)}) - \nabla \mathcal{L}(\bar{\theta}_t)) \right\|_2^2 \\ &\quad + 3 \left\| \frac{1}{h} \sum_{j \in [h]} \xi_t^{(j)} \right\|_2^2 + 3 \left\| \frac{\bar{\theta}_{t+1/2} - \bar{\theta}_{t+1}}{\eta} \right\|_2^2. \end{aligned} \quad (135)$$

We now note that the expectation of each term can be bounded. Indeed,

$$\left\| \frac{1}{h} \sum_{j \in [h]} (\nabla \mathcal{L}(\theta_t^{(j)}) - \nabla \mathcal{L}(\bar{\theta}_t)) \right\|_2 \leq \frac{1}{h} \sum_{j \in [h]} \|\nabla \mathcal{L}(\theta_t^{(j)}) - \nabla \mathcal{L}(\bar{\theta}_t)\|_2 \quad (136)$$

$$\leq \frac{1}{h} \sum_{j \in [h]} L \|\theta_t^{(j)} - \bar{\theta}_t\|_2 \leq \frac{1}{h} \sum_{j \in [h]} L\Delta_2(\bar{\theta}_t) = L\Delta_2(\bar{\theta}_t). \quad (137)$$

Moreover,

$$\mathbb{E} \left\| \frac{1}{h} \sum_{j \in [h]} \xi_t^{(j)} \right\|_2^2 = \mathbb{E} \frac{1}{h^2} \sum_{j, k \in [h]} \xi_t^{(j)} \cdot \xi_t^{(k)} = \frac{1}{h^2} \sum_{j, k \in [h]} \mathbb{E} \xi_t^{(j)} \cdot \xi_t^{(k)} \quad (138)$$

$$= \frac{1}{h^2} \sum_{j \in [h]} \mathbb{E} \|\xi_t^{(j)}\|_2^2 \leq \frac{1}{h^2} \sum_{j \in [h]} \sigma_t^2 = \frac{\sigma_t^2}{h}, \quad (139)$$

using the fact that the noises are independent to move from one line to the other. For the last term, we use the  $C$ -averaging guarantee of AVG, yielding

$$\left\| \bar{\theta}_{t+1} - \bar{\theta}_{t+1/2} \right\|_2^2 \leq C^2 \Delta_2(\bar{\theta}_{t+1/2})^2. \quad (140)$$

To bound the right-hand side, note that we have  $\bar{\theta}_{t+1/2} = \bar{\theta}_t - \eta \bar{g}_t = \bar{\theta}_t - \eta \bar{\nabla} \bar{\mathcal{L}}(\bar{\theta}_t) - \eta \bar{\xi}_t$ , where  $\bar{\nabla} \bar{\mathcal{L}}(\bar{\theta}_t) = \left( \nabla \mathcal{L}(\theta_t^{(1)}), \dots, \nabla \mathcal{L}(\theta_t^{(h)}) \right)$ . Lemma 2 then implies

$$\left\| \frac{\bar{\theta}_{t+1} - \bar{\theta}_{t+1/2}}{\eta} \right\|_2^2 \leq \frac{3C^2}{\eta^2} \Delta_2(\bar{\theta}_t)^2 + 3C^2 L^2 \Delta_2(\bar{\theta}_t)^2 + 3C^2 \Delta_2(\bar{\xi}_t)^2. \quad (141)$$

Combining it all, applying Lemma 9 and defining  $B \triangleq \frac{9C^2}{\eta^2} + 3L^2 + 9C^2 L^2$ , then yields

$$\mathbb{E}_{\bar{\xi}_t | \bar{\theta}_t} \left\| \bar{G}_t - \nabla \mathcal{L}(\bar{\theta}_t) \right\|_2^2 \leq \frac{3\sigma_t^2}{h} + \left( \frac{9C^2}{\eta^2} + 3L^2 + 9C^2 L^2 \right) \Delta_2(\bar{\theta}_t)^2 + 9C^2 \mathbb{E}_{\bar{\xi}_t | \bar{\theta}_t} \Delta_2(\bar{\xi}_t)^2 \quad (142)$$

$$\leq \left( \frac{3}{h} + 36C^2 h \right) \frac{\sigma_t^2}{h} + B \Delta_2(\bar{\theta}_t)^2. \quad (143)$$

Defining  $A \triangleq \frac{3}{h} + 36C^2 h$  then yields the desired result.  $\square$

We now proceed with the proof of our theorem.

*Proof of Theorem 3.* A direct consequence of Lipschitz continuity of the gradient of the loss function (Assumption 2) is that for all  $\phi, \psi \in \mathbb{R}^d$ , we have

$$\mathcal{L}(\psi) \leq \mathcal{L}(\phi) + (\psi - \phi) \cdot \nabla \mathcal{L}(\phi) + \frac{L}{2} \|\psi - \phi\|_2^2. \quad (144)$$

Therefore, by the definition of the effective gradient, we have

$$\mathcal{L}(\bar{\theta}_{t+1}) \leq \mathcal{L}(\bar{\theta}_t) - \eta \bar{G}_t \cdot \nabla \mathcal{L}(\bar{\theta}_t) + \frac{L\eta^2}{2} \|\bar{G}_t\|_2^2. \quad (145)$$

The fact that  $\eta \leq 1/L$  then implies

$$\mathcal{L}(\bar{\theta}_{t+1}) \leq \mathcal{L}(\bar{\theta}_t) - \eta \bar{G}_t \cdot \nabla \mathcal{L}(\bar{\theta}_t) + \frac{\eta}{2} \|\bar{G}_t\|_2^2 \quad (146)$$

$$= \mathcal{L}(\bar{\theta}_t) - \frac{\eta}{2} \|\nabla \mathcal{L}(\bar{\theta}_t)\|_2^2 + \frac{\eta}{2} \left( \|\bar{G}_t\|_2^2 - 2\bar{G}_t \cdot \nabla \mathcal{L}(\bar{\theta}_t) + \|\nabla \mathcal{L}(\bar{\theta}_t)\|_2^2 \right) \quad (147)$$

$$= \mathcal{L}(\bar{\theta}_t) - \frac{\eta}{2} \|\nabla \mathcal{L}(\bar{\theta}_t)\|_2^2 + \frac{\eta}{2} \|\bar{G}_t - \nabla \mathcal{L}(\bar{\theta}_t)\|_2^2. \quad (148)$$

By rearranging the terms, we then have

$$\|\nabla \mathcal{L}(\bar{\theta}_t)\|_2^2 \leq \frac{2}{\eta} (\mathcal{L}(\bar{\theta}_t) - \mathcal{L}(\bar{\theta}_{t+1})) + \|\bar{G}_t - \nabla \mathcal{L}(\bar{\theta}_t)\|_2^2. \quad (149)$$

Now taking the expectation over all of the stochastic noises and averaging over  $t$ , yields

$$\frac{1}{T} \sum_{t \in [T]} \mathbb{E}_{\bar{\xi}_{1:T}} \|\nabla \mathcal{L}(\bar{\theta}_t)\|_2^2 \leq \frac{2}{\eta T} (\mathcal{L}(\bar{\theta}_1) - \mathcal{L}(\bar{\theta}_{T+1})) + \frac{1}{T} \sum_{t \in [T]} \mathbb{E}_{\bar{\xi}_{1:T}} \|\bar{G}_t - \nabla \mathcal{L}(\bar{\theta}_t)\|_2^2 \quad (150)$$

$$\leq \frac{2\mathcal{L}_{max}}{\eta T} + \frac{A}{T} \sum_{t \in [T]} \sigma_t^2 + \frac{B}{T} \sum_{t \in [T]} \mathbb{E}_{\bar{\xi}_{1:T}} \Delta_2(\bar{\theta}_t)^2, \quad (151)$$

where in the last inequality we used Lemma 13. The regular increase of the batch size (Section 2.2) then implies the existence of an iteration  $T_1$ , after which we have  $\sigma_t^2 \leq \frac{\delta^2}{8A}$ . Therefore, for  $T \geq T_2 \triangleq \max \left\{ \frac{8\mathcal{L}_{max}}{\eta\delta^2}, T_1 \right\}$ , we have

$$\frac{1}{T} \sum_{t \in [T]} \mathbb{E}_{\bar{\xi}_{1:T}} \|\nabla \mathcal{L}(\bar{\theta}_t)\|_2^2 \leq \frac{\delta^2}{4} + \frac{A}{T} \left( T_2 \sigma_1^2 + (T - T_2) \frac{\delta^2}{8A} \right) + \frac{B}{T} \sum_{t \in [T]} \mathbb{E}_{\bar{\xi}_{1:T}} \Delta_2(\bar{\theta}_t)^2 \quad (152)$$

$$\leq \frac{3\delta^2}{8} + \frac{T_2}{T} A \sigma_1^2 + \frac{B}{T} \sum_{t \in [T]} \mathbb{E}_{\bar{\xi}_{1:T}} \Delta_2(\bar{\theta}_t)^2. \quad (153)$$

For  $T \geq T_3 \triangleq \frac{8T_2 A \sigma_1^2}{\delta^2}$ , we then have

$$\frac{1}{T} \sum_{t \in [T]} \mathbb{E}_{\vec{\xi}_{1:T}} \|\nabla \mathcal{L}(\vec{\theta}_t)\|_2^2 \leq \frac{\delta^2}{2} + \frac{B}{T} \sum_{t \in [T]} \mathbb{E}_{\vec{\xi}_{1:T}} \Delta_2(\vec{\theta}_t)^2. \quad (154)$$

We now invoke Lemma 12 with  $\varepsilon = \frac{1}{2} \min\{\delta^2, \delta^2/2B\}$ . For  $T \geq T^*(\varepsilon)$  we then have

$$\frac{1}{T} \sum_{t \in [T]} \mathbb{E}_{\vec{\xi}_{1:T}} \Delta_2(\vec{\theta}_t)^2 = \frac{1}{T} \sum_{t=1}^{T^*(\varepsilon)} \mathbb{E}_{\vec{\xi}_{1:T}} \Delta_2(\vec{\theta}_t)^2 + \frac{1}{T} \sum_{t=T^*(\varepsilon)+1}^T \mathbb{E}_{\vec{\xi}_{1:T}} \Delta_2(\vec{\theta}_t)^2 \quad (155)$$

$$\leq \frac{T^*(\varepsilon)}{T} D + \frac{T - T^*(\varepsilon)}{T} \varepsilon \leq \frac{T^*(\varepsilon)}{T} D + \varepsilon. \quad (156)$$

Thus, for  $T \geq T_4 \triangleq \frac{DT^*(\varepsilon)}{\varepsilon}$  we have

$$\mathbb{E} \Delta_2(\vec{\theta}_*)^2 = \frac{1}{T} \sum_{t \in [T]} \mathbb{E}_{\vec{\xi}_{1:T}} \Delta_2(\vec{\theta}_t)^2 \leq 2\varepsilon = \min\{\delta^2, \delta^2/2B\}. \quad (157)$$

Combining this with (154), for  $T = T_{\text{HOM-LEARN}}(\delta) \triangleq \max\{T_3, T_4\}$  we then have

$$\mathbb{E} \Delta_2(\vec{\theta}_*)^2 \leq \delta^2 \quad \text{and} \quad \mathbb{E} \|\nabla \mathcal{L}(\vec{\theta}_*)\|_2^2 \leq \delta^2, \quad (158)$$

which is the desired result.  $\square$

## C MDA Algorithm

### C.1 Correctness proof of MDA

We first note a few important properties of MDA.

**Lemma 14.** *The  $\ell_2$  diameter of the MDA subfamily is upper-bounded by that of the honest vectors. In other words, for any Byzantine attack  $\overrightarrow{\text{BYZ}}$ , denoting  $\vec{z} \triangleq \overrightarrow{\text{BYZ}}(\vec{x})$ , we have*

$$\Delta_2(\vec{z}^{(\text{MDA})}) \leq \Delta_2(\vec{x}) \quad (159)$$

*Proof.* Since  $\overrightarrow{\text{BYZ}}$  selects  $q$  vectors, out of which at most  $f$  are Byzantine vectors, we know that there exists a subset  $H \subset [q]$  of cardinal at least  $q - f$  that only contains honest vectors. But then, we have

$$\Delta_2(\vec{z}^{(\text{MDA})}) = \min_{\substack{S \subset [q] \\ |S|=q-f}} \Delta_2(\vec{z}^{(S)}) \leq \Delta_2(\vec{z}^{(H)}) \leq \Delta_2(\vec{x}), \quad (160)$$

which is the lemma.  $\square$

**Lemma 15.** *Under Assumption 5, MDA guarantees Byzantine asymptotic agreement. In other words, for any input  $N \in \mathbb{N}$  and any family  $\vec{x} \in \mathbb{R}^{d \cdot h}$ , denoting  $\vec{x}_N \triangleq \overrightarrow{\text{MDA}}_N \circ \overrightarrow{\text{BYZ}}_N(\vec{x})$ ,*

$$\Delta_2(\vec{x}_N) \leq \frac{\Delta_2(\vec{x})}{2^N}. \quad (161)$$

*Proof.* Denote  $\vec{z}^{(1)} \triangleq \overrightarrow{\text{BYZ}}^{(1)}(\vec{x})$  and  $\vec{z}^{(2)} \triangleq \overrightarrow{\text{BYZ}}^{(2)}(\vec{x})$  the results of the two Byzantine attacks,  $S_1 = S_{\text{MDA}}(\vec{z}^{(1)})$  and  $S_2 = S_{\text{MDA}}(\vec{z}^{(2)})$  the subsets selected by MDA in the two cases.

Moreover, we write  $S_1 = H_1 \cup F_1$  and  $S_2 = H_2 \cup F_2$ , where  $H_1$  and  $H_2$  are subsets of honest vectors within  $S_1$  and  $S_2$ . Without loss of generality, we assume both  $H_1$  and  $H_2$  to be of cardinal  $q - 2f$ . As a result, we know that there exist injective functions  $\sigma_1 : H_1 \rightarrow [h]$  and  $\sigma_2 : H_2 \rightarrow [h]$  such that  $z^{(1,j)} = x^{(\sigma_1(j))}$  and  $z^{(2,k)} = x^{(\sigma_2(k))}$ , for all  $j \in H_1$  and  $k \in H_2$ .

Finally, we denote  $y^{(1)} \triangleq \text{MDA}(\bar{z}^{(1)})$  and  $y^{(2)} \triangleq \text{MDA}(\bar{z}^{(2)})$ . We then have

$$(q-f) \left\| y^{(1)} - y^{(2)} \right\|_2 = \left\| \sum_{j \in S_1} z^{(1,j)} - \sum_{k \in S_2} z^{(2,k)} \right\|_2 \quad (162)$$

$$= \left\| \sum_{j \in F_1} z^{(1,j)} - \sum_{k \in F_2} z^{(2,k)} + \sum_{j \in \sigma_1(H_1)} x^{(j)} - \sum_{k \in \sigma_2(H_2)} x^{(k)} \right\|_2 \quad (163)$$

$$\leq \left\| \sum_{j \in F_1} z^{(1,j)} - \sum_{k \in F_2} z^{(2,k)} \right\|_2 + \left\| \sum_{j \in \sigma_1(H_1) - \sigma_2(H_2)} x^{(j)} - \sum_{k \in \sigma_2(H_2) - \sigma_1(H_1)} x^{(k)} \right\|_2. \quad (164)$$

Note that  $|F_1| = |S_1 - H_1| = f = |S_2 - H_2| = |F_2|$ . Moreover,

$$|\sigma_1(H_1) - \sigma_2(H_2)| = |\sigma_1(H_1) \cup \sigma_2(H_2) - \sigma_2(H_2)| \quad (165)$$

$$= |\sigma_1(H_1) \cup \sigma_2(H_2)| - |\sigma_2(H_2)| \leq |[h]| - |H_2| = 2f + h - q, \quad (166)$$

and similarly for  $\sigma_2(H_2) - \sigma_1(H_1)$ . Note that  $|\sigma_1(H_1)|, |\sigma_2(H_2)| \geq q - 2f$ . Therefore, Assumption 5 implies that

$$|\sigma_1(H_1)| + |\sigma_2(H_2)| \geq 2 \left( \frac{1+\varepsilon}{2} h + \frac{5+3\varepsilon}{2} f - 2f \right) > h, \quad (167)$$

which yields  $\sigma_1(H_1) \cap \sigma_2(H_2) \neq \emptyset$ . Now let  $\gamma$  be an element of the intersection of  $\sigma_1(H_1)$  and  $\sigma_2(H_2)$ , and consider any bijections  $\tau_F : F_1 \rightarrow F_2$  and  $\tau_H : \sigma_1(H_1) - \sigma_2(H_2) \rightarrow \sigma_2(H_2) - \sigma_1(H_1)$ . Using triangle inequality and Lemma 14, for any  $j \in F_1$ , we then have

$$\left\| z^{(1,j)} - z^{(2,\tau_F(j))} \right\|_2 \leq \left\| z^{(1,j)} - x^{(\gamma)} \right\|_2 + \left\| x^{(\gamma)} - z^{(2,\tau_F(j))} \right\|_2 \leq 2\Delta_2(\bar{x}). \quad (168)$$

Combining it all, we obtain

$$(q-f) \left\| y^{(1)} - y^{(2)} \right\|_2 \leq \sum_{j \in F_1} \left\| z^{(1,j)} - z^{(2,\tau_F(j))} \right\|_2 + \sum_{j \in \sigma_1(H_1) - \sigma_2(H_2)} \left\| x^{(j)} - x^{(\tau_H(j))} \right\|_2 \quad (169)$$

$$\leq 2f\Delta_2(\bar{x}) + (2f + h - q)\Delta_2(\bar{x}), \quad (170)$$

which implies

$$\left\| y^{(1)} - y^{(2)} \right\|_2 \leq \frac{4f + h - q}{q - f} \Delta_2(\bar{x}). \quad (171)$$

We then apply Assumption 5, which implies that

$$\frac{4f + h - q}{q - f} \leq \frac{4f + h - \frac{1+\varepsilon}{2}h - \frac{5+3\varepsilon}{2}f}{\frac{1+\varepsilon}{2}h + \frac{5+3\varepsilon}{2}f - f} = \frac{(1-\varepsilon)h + 3(1-\varepsilon)f}{(1+\varepsilon)h + 3(1+\varepsilon)f} \quad (172)$$

$$= \frac{1-\varepsilon}{1+\varepsilon} = \frac{1+\varepsilon-2\varepsilon}{1+\varepsilon} = 1 - \frac{2\varepsilon}{1+\varepsilon} = 1 - \tilde{\varepsilon}. \quad (173)$$

This shows that  $\Delta_2(\bar{y}) \leq (1 - \tilde{\varepsilon})\Delta_2(\bar{x})$ . In other words, one iteration of MDA is guaranteed to multiply the  $\ell_2$  diameter of honest nodes by at most  $(1 - \tilde{\varepsilon})$ . It follows that  $T_{\text{MDA}}(N) = \lceil N \ln 2 / \tilde{\varepsilon} \rceil$  iterations will multiply this diameter by at most  $(1 - \tilde{\varepsilon})^{T_{\text{MDA}}(N)} \leq \exp\left(\frac{\ln(1-\tilde{\varepsilon}) \ln 2}{\tilde{\varepsilon}}\right)^N \leq 2^{-N}$ , using the inequality  $\ln(1 - \tilde{\varepsilon}) \leq -\tilde{\varepsilon}$ .  $\square$

**Remark 4.** Note that we can set  $\varepsilon = \frac{n-6f}{n+2f}$ . Thus, in the regime  $f \ll n$ , we have  $\varepsilon \rightarrow 1$ , which implies  $\tilde{\varepsilon} \rightarrow 1$ . Thus, for any fixed value of  $N$ , given that our proof showed that  $\Delta_2(\bar{y}) \leq (1 - \tilde{\varepsilon})\Delta_2(\bar{x})$ , when  $f$  is sufficiently small compared to  $n$ , MDA actually achieves asymptotic agreement in only one communication round.

**Lemma 16.** One iteration of MDA returns a vector close to the average of the honest vectors. Denoting  $\bar{y} \triangleq \text{MDA} \circ \text{BYZ}(\bar{x})$ , we have

$$\|\bar{y} - \bar{x}\|_2 \leq \frac{(2f + h - q)q + (q - 2f)f}{h(q - f)} \Delta_2(\bar{x}). \quad (174)$$

In the synchronous case where  $q = n = f + h$ , the right-hand side becomes  $\frac{2f}{h} \Delta_2(\bar{x})$ .

*Proof.* Let us write  $S_{\text{MDA}}(\vec{z}) = H \cup F$ , where  $H$  are honest vectors and  $F$  are Byzantine vectors. We know that  $|H| \geq q - 2f$  and  $|H| + |F| = q - f$ . In fact, without loss of generality, we can assume  $|H| = q - 2f$  (since this is equivalent to labeling honest vectors not in  $H$  as Byzantine vectors).

Let us also denote  $\sigma : H \rightarrow [h]$  the injective function that maps honest vectors to the index of their node, and  $\bar{H} = [h] - \sigma(H)$  the unqueried nodes. We have

$$\|y - \bar{x}\|_2 = \left\| \frac{|H| \bar{z}^{(H)} + |F| \bar{z}^{(F)}}{|H| + |F|} - \frac{|\sigma(H)| \bar{x}^{(\sigma(H))} + |\bar{H}| \bar{x}^{(\bar{H})}}{|\sigma(H)| + |\bar{H}|} \right\|_2 \quad (175)$$

$$= \left\| \frac{|H| \bar{z}^{(H)} + |F| \bar{z}^{(F)}}{|H| + |F|} - \frac{|H| \bar{z}^{(H)} + |\bar{H}| \bar{x}^{(\bar{H})}}{|H| + |\bar{H}|} \right\|_2 \quad (176)$$

$$= \frac{\left\| |H| (|\bar{H}| - |F|) \bar{z}^{(H)} + |F| (|H| + |\bar{H}|) \bar{z}^{(F)} - |\bar{H}| (|H| + |F|) \bar{x}^{(\bar{H})} \right\|_2}{(|H| + |F|) (|H| + |\bar{H}|)} \quad (177)$$

$$= \frac{\left\| |H| |\bar{H}| (\bar{z}^{(H)} - \bar{x}^{(\bar{H})}) + |F| |H| (\bar{z}^{(H)} - \bar{z}^{(F)}) + |\bar{H}| |F| (\bar{z}^{(F)} - \bar{x}^{(\bar{H})}) \right\|_2}{(|H| + |F|) (|H| + |\bar{H}|)} \quad (178)$$

$$\leq \frac{|H| |\bar{H}| \left\| \bar{z}^{(H)} - \bar{x}^{(\bar{H})} \right\|_2 + |F| |H| \left\| \bar{z}^{(H)} - \bar{z}^{(F)} \right\|_2 + |\bar{H}| |F| \left\| \bar{z}^{(F)} - \bar{x}^{(\bar{H})} \right\|_2}{(|H| + |F|) (|H| + |\bar{H}|)}. \quad (179)$$

Now note that

$$\left\| \bar{z}^{(H)} - \bar{x}^{(\bar{H})} \right\|_2 = \left\| \bar{x}^{(\sigma(H))} - \bar{x}^{(\bar{H})} \right\|_2 \quad (180)$$

$$= \frac{1}{|\sigma(H)| |\bar{H}|} \left\| |\bar{H}| \sum_{j \in \sigma(H)} x^{(j)} - |H| \sum_{k \in \bar{H}} x^{(k)} \right\|_2 \quad (181)$$

$$= \frac{1}{|\sigma(H)| |\bar{H}|} \left\| \sum_{k \in \bar{H}} \sum_{j \in \sigma(H)} x^{(j)} - \sum_{j \in \sigma(H)} \sum_{k \in \bar{H}} x^{(k)} \right\|_2 \quad (182)$$

$$= \frac{1}{|\sigma(H)| |\bar{H}|} \left\| \sum_{j \in \sigma(H)} \sum_{k \in \bar{H}} (x^{(j)} - x^{(k)}) \right\|_2 \quad (183)$$

$$\leq \frac{1}{|\sigma(H)| |\bar{H}|} \sum_{j \in \sigma(H)} \sum_{k \in \bar{H}} \left\| x^{(j)} - x^{(k)} \right\|_2 \quad (184)$$

$$\leq \frac{1}{|\sigma(H)| |\bar{H}|} \sum_{j \in \sigma(H)} \sum_{k \in \bar{H}} \Delta_2(\vec{x}) = \Delta_2(\vec{x}). \quad (185)$$

Similarly, we show that  $\left\| \bar{z}^{(H)} - \bar{z}^{(F)} \right\|_2 \leq \Delta_2(\vec{z}^{\text{MDA}}) \leq \Delta_2(\vec{x})$ . Finally, we use the triangle inequality to show that

$$\left\| \bar{z}^{(F)} - \bar{x}^{(\bar{H})} \right\|_2 \leq \left\| \bar{z}^{(F)} - \bar{z}^{(H)} \right\|_2 + \left\| \bar{z}^{(H)} - \bar{x}^{(\bar{H})} \right\|_2 \leq 2\Delta_2(\vec{x}). \quad (186)$$

Therefore, we now have

$$\|y - \bar{x}\|_2 \leq \frac{|H| |\bar{H}| + |F| |H| + 2 |\bar{H}| |F|}{(|H| + |F|) (|H| + |\bar{H}|)} \Delta_2(\vec{x}) \quad (187)$$

$$= \frac{(q - 2f)(2f + h - q) + f(q - 2f) + 2(2f + h - q)f}{h(q - f)} \Delta_2(\vec{x}) \quad (188)$$

$$= \frac{(2f + h - q)q + (q - 2f)f}{h(q - f)} \Delta_2(\vec{x}), \quad (189)$$

which is the lemma.  $\square$

Now we can prove our theorem.

*Proof of Theorem 4.* Lemma 15 already proved asymptotic agreement. Moreover, using Lemma 16, and denoting  $\alpha \triangleq \frac{(2f+h-q)q+(q-2f)f}{h(q-f)}$  and  $\bar{x}_t$  the vector family obtained after  $t$  iterations of MDA, we also know that

$$\|\bar{x}_{t+1} - \bar{x}_t\|_2 \leq \alpha \Delta_2(\bar{x}_t) \leq \alpha(1 - \tilde{\varepsilon})^t \Delta_2(\bar{x}_0). \quad (190)$$

Using triangle inequality then yields, for any number  $T_{\text{MDA}}(N)$  of iterations of MDA,

$$\|\bar{x}_{T_{\text{MDA}}(N)} - \bar{x}_0\|_2 \leq \sum_{t=0}^{T_{\text{MDA}}(N)-1} \|\bar{x}_{t+1} - \bar{x}_t\|_2 \leq \sum_{t=0}^{T_{\text{MDA}}(N)-1} \alpha(1 - \tilde{\varepsilon})^t \Delta_2(\bar{x}_0) \quad (191)$$

$$\leq \alpha \Delta_2(\bar{x}_0) \sum_{t=0}^{\infty} (1 - \tilde{\varepsilon})^t = \frac{\alpha \Delta_2(\bar{x}_0)}{\tilde{\varepsilon}}, \quad (192)$$

which is the guarantee of the theorem.  $\square$

## C.2 Lower bound on the averaging constant

We prove here a lower bound on the averaging constant that any algorithm can achieve. Our proof requires the construction of hard instances. We use the following notation.

**Definition 4** ( $\star$  notation). We denote by  $x \star h \triangleq \underbrace{(x, \dots, x)}_{h \text{ times}}$  the repetition of a value  $x$   $h$  times.

**Lemma 17** (Quasi-unanimity). For any averaging agreement algorithm AVG, whenever a node  $j$  only hears from  $q$  nodes, assuming  $q - f$  of these nodes act like honest nodes with the same initial value  $x$ , then AVG must make node  $j$  output  $x$ .

*Proof.* For any agreeing initial family  $\vec{x} = x \star h$ , we have  $\Delta_2(\vec{x}) = 0$  and  $\bar{x} = x$ . Then averaging agreement implies that, for any  $N \in \mathbb{N}$ , we output  $\bar{x}_N$  such that

$$\Delta_2(\bar{x}_N) \leq \frac{\Delta_2(\vec{x})}{2^N} = 0 \quad \text{and} \quad \|\bar{x}_N - x\|_2 \leq C \Delta_2(\vec{x}) = 0. \quad (193)$$

In other words, we must have  $\bar{x}_N = \bar{x}$ .

But then, if node  $j$  only hears from  $q$  nodes, and if it receives  $q - f$  nodes agreeing on a value  $x$ , then it cannot exclude the possibility that the remaining  $f$  nodes come from Byzantine nodes. As a result, node  $j$  cannot exclude that the initial family was  $\vec{x}$ . To satisfy averaging agreement, node  $j$  must then output  $x$ .  $\square$

*Proof of Theorem 5.* Consider the vector family defined by

$$\vec{x} \triangleq (0 \star (q - 2f), 1 \star (h + 2f - q)). \quad (194)$$

For any algorithm AVG used by honest nodes, Byzantine nodes can slow down all messages from nodes in  $[q - f + 1, h]$  to nodes in  $[q - f]$ . Thus, the first  $q - f$  honest nodes would be making decisions without receiving any input from nodes in  $[q - f + 1, h]$ . Assume now that the Byzantine nodes all act exactly like the first  $q - 2f$  nodes. Then, all first  $q - f$  nodes would see  $q - f$  nodes acting like honest nodes with initial vector 0, and  $f$  nodes acting like honest nodes with initial vector 1. By quasi-unanimity (Lemma 17), the  $q - 2f$  first nodes must output 0.

But now, by asymptotic agreement, this implies that any other honest node must output a vector at distance at most  $\Delta_2(\vec{x})/2^N = 1/2^N$  of 0. As a result, as  $N \rightarrow \infty$ , denoting  $\bar{x}_N$  the output of AVG for input  $N \in \mathbb{N}$ , we must have  $\bar{x}_N \rightarrow 0$ .

Since  $\Delta_2(\vec{x}) = 1$  and  $\bar{x} = (h + 2f - q)/h$ , we then have

$$\lim_{N \rightarrow \infty} |\bar{x}_N - \bar{x}| = |0 - \bar{x}| = \frac{h + 2f - q}{h} \geq \frac{h + 2f - q}{h} \Delta_2(\vec{x}). \quad (195)$$

This shows that AVG cannot achieve better than  $\frac{h+2f-q}{h}$ -averaging agreement which is equal to  $\frac{2f}{h}$ , for  $q = h$ .

Now we show MDA indeed achieves this bound up to a multiplicative constant. From Assumption 5, we can set  $\varepsilon = \frac{n-6f}{n+2f} = \frac{h-5f}{h+3f}$ . In the regime  $q = h$  and  $f \ll h$ , we thus have  $\varepsilon \rightarrow 1$ , which then implies  $\tilde{\varepsilon} \rightarrow 1$ . Now notice that

$$\frac{(2f + h - q)q + (q - 2f)f}{h(q - f)\tilde{\varepsilon}} = \frac{2f/h}{(1 - \frac{f}{q})\tilde{\varepsilon}} + \frac{f}{h} \frac{1 - \frac{2f}{q}}{(1 - \frac{f}{q})\tilde{\varepsilon}} = \frac{3f}{h} + o(1). \quad (196)$$

Yet, for  $q = h$ , we showed that the lower bound on the averaging constant is  $2f/h$ . MDA is thus asymptotically optimal, up to the multiplicative constant  $3/2$ .  $\square$

### C.3 Note on Byzantine tolerance

Our MDA algorithm tolerates a small fraction of Byzantine nodes, namely  $n > 6f$ .

**Proposition 1.** *Assume  $n \leq 6f$ . For any parameter  $N$ , no matter how the number  $T_{\text{MDA}}(N)$  of iterations is chosen, Byzantines can make MDA fail to achieve asymptotic agreement. As a result, MDA cannot guarantee averaging agreement for  $n \leq 6f$ .*

*Proof.* Define  $\delta \triangleq \min \{1, 4 - 4 \cdot 2^{-(N-1)/T_{\text{MDA}}(N)}\}$ , and consider the honest vector family

$$\vec{x} \triangleq (-1 \star 2f, 0 \star f, 1 \star 2f). \quad (197)$$

For the first  $2f$  nodes, Byzantine nodes can block  $f$  of the messages of the last  $2f$  nodes, and add  $f$  values equal to  $-2 + \delta$ . The first  $2f$  honest nodes then observe

$$\vec{z}_1 \triangleq ((-2 + \delta) \star f, -1 \star 2f, 0 \star f, 1 \star f). \quad (198)$$

MDA would then remove the largest  $f$  inputs, as it then achieves a diameter equal to  $2 - \delta < 2$ . The first  $2f$  nodes would then output

$$\frac{(-2 + \delta)f - 2f}{4f} = -\frac{4f - \delta f}{4f} = -1 + \frac{\delta}{4}. \quad (199)$$

For the middle  $f$  nodes, Byzantine nodes can simply allow perfect communication and not intervene. By symmetry, the middle  $f$  nodes would output 0 (note that to be rigorous, we could define MDA as taking the average of the outputs over all subsets of inputs of minimal diameter).

For the last  $2f$  nodes, Byzantine nodes can block  $f$  messages of the first  $2f$  nodes, and add  $f$  values equal to  $2 - \delta$ . The situation is then symmetric to the case for the first  $2f$  nodes, and make the last  $2f$  nodes output  $1 - \frac{\delta}{4}$ .

As a result, one iteration of MDA multiplies the diameter of the honest nodes by  $1 - \frac{\delta}{4} \geq 2^{-(N-1)/T_{\text{MDA}}(N)}$ . Byzantine nodes can use the same strategy in all other iterations. Thus, denoting  $\vec{y} \triangleq \overrightarrow{\text{MDA}} \circ \overrightarrow{\text{BYZ}}(\vec{x})$ , we have

$$\Delta_2(\vec{y}) = \left(1 - \frac{\delta}{4}\right)^T \Delta_2(\vec{x}) \geq \left(2^{-(N-1)/T_{\text{MDA}}(N)}\right)^{T_{\text{MDA}}(N)} \Delta_2(\vec{x}) = 2 \frac{\Delta_2(\vec{x})}{2^N} > \frac{\Delta_2(\vec{x})}{2^N}, \quad (200)$$

which shows that the final vectors obtained by MDA violate asymptotic agreement.  $\square$

## D RB-TM Algorithm

In this section, we prove that RB-TM solves averaging agreement under Assumption 6.

### D.1 Diameters

Before doing so, we introduce the notion of  $\ell_r$ -diameters and we prove a few useful lemmas.

**Definition 5.** *For any  $r \in [1, \infty]$ , we define the diameter along coordinate  $i$  by*

$$\Delta^{cw}(\vec{x})[i] = \max_{j,k \in [h]} |x^{(j)}[i] - x^{(k)}[i]|, \quad (201)$$

and the coordinate-wise  $\ell_r$ -diameters by  $\Delta_r^{cw}(\vec{x}) = \|\Delta^{cw}(\vec{x})\|_r$ .

Interestingly, we have the following bounds between diameters.

**Lemma 18.** *The  $\ell_r$ -diameters are upper-bounded by coordinate-wise  $\ell_r$ -diameters, i.e.,*

$$\forall r, \Delta_r \leq \Delta_r^{cw} \leq \min \left\{ d^{1/r}, 2h^{1/r} \right\} \Delta_r. \quad (202)$$

*Note that in ML applications, we usually expect  $d \gg h$ , in which case the more relevant right-hand side inequality is  $\Delta_r^{cw} \leq 2h^{1/r} \Delta_r$ .*

*Proof.* Consider  $j^*, k^* \in [h]$  such that  $\Delta_r(\vec{x}) = \|x^{(j^*)} - x^{(k^*)}\|_r$ . But then, we note that on each coordinate  $i \in [d]$ ,

$$\left| x^{(j^*)}[i] - x^{(k^*)}[i] \right| \leq \max_{j,k \in [h]} \left| x^{(j)}[i] - x^{(k)}[i] \right| = \Delta^{cw}(x)[i]. \quad (203)$$

As a result,  $\Delta_r(\vec{x}) = \|x^{(j^*)} - x^{(k^*)}\|_r \leq \|\Delta^{cw}(\vec{x})\|_r = \Delta_r^{cw}(\vec{x})$ . For the right-hand side, first note that a coordinate-wise diameter is smaller than the  $\ell_r$  diameter, which yields

$$\Delta_r^{cw}(\vec{x})^r = \sum_{i \in [d]} \max_{j,k \in [h]} \left| x^{(j)}[i] - x^{(k)}[i] \right|^r \leq \sum_{i \in [d]} \max_{j,k \in [h]} \left\| x^{(j)} - x^{(k)} \right\|_r^r \quad (204)$$

$$= \sum_{i \in [d]} \Delta_r(\vec{x})^r = d \Delta_r(\vec{x})^r. \quad (205)$$

Taking the  $r$ -th root shows that  $\Delta_r^{cw} \leq d^{1/r} \Delta_r$ . What is left to prove is that  $\Delta_r^{cw} \leq 2h^{1/r} \Delta_r$ . To prove this, note that for any  $i \in [d]$ , we have

$$\max_{j,k \in [h]} \left| x^{(j)}[i] - x^{(k)}[i] \right| \leq \max_{j \in [h]} \left( 2 \left| x^{(j)}[i] - x^{(1)}[i] \right| \right). \quad (206)$$

Indeed, assuming the former maximum is reached for  $j^*$  and  $k^*$ , the latter maximum will be reached for  $j^*$  or  $k^*$ , depending on whether  $x^{(1)}[i]$  is closer to  $x^{(j^*)}[i]$  or  $x^{(k^*)}[i]$ . In either case, the above inequality holds. As a result,

$$\Delta_r^{cw}(\vec{x})^r = \sum_{i \in [d]} \max_{j,k \in [h]} \left| x^{(j)}[i] - x^{(k)}[i] \right|^r \leq \sum_{i \in [d]} \max_{j \in [h]} \left( 2 \left| x^{(j)}[i] - x^{(1)}[i] \right| \right)^r \quad (207)$$

$$= 2^r \sum_{i \in [d]} \max_{j \in [h]} \left| x^{(j)}[i] - x^{(1)}[i] \right|^r \leq 2^r \sum_{i \in [d]} \sum_{j \in [h]} \left| x^{(j)}[i] - x^{(1)}[i] \right|^r \quad (208)$$

$$= 2^r \sum_{j \in [h]} \sum_{i \in [d]} \left| x^{(j)}[i] - x^{(1)}[i] \right|^r = 2^r \sum_{j \in [h]} \left\| x^{(j)} - x^{(1)} \right\|_r^r \quad (209)$$

$$\leq 2^r \sum_{j \in [h]} \Delta_r(\vec{x})^r = 2^r h \Delta_r(\vec{x})^r. \quad (210)$$

Taking the  $r$ -th root yields  $\Delta_r^{cw} \leq 2h^{1/r} \Delta_r$ , which concludes the proof.  $\square$

As an immediate corollary, asymptotic agreement is equivalent to showing that *any* of the diameters we introduce in this section goes to zero.

Interestingly, our diameters satisfy the triangle inequality, as shown by the following lemma.

**Lemma 19.** *The diameters and coordinate-wise diameters satisfy the triangle inequality. Namely, for any two families of vectors  $\vec{x}$  and  $\vec{y}$ , we have the following inequality*

$$\Delta^{cw}(\vec{x} + \vec{y}) \leq \Delta^{cw}(\vec{x}) + \Delta^{cw}(\vec{y}). \quad (211)$$

*As an immediate corollary, by triangle inequality of norms, for any  $r \in [1, \infty]$ , we also have  $\Delta_r^{cw}(\vec{x} + \vec{y}) \leq \Delta_r^{cw}(\vec{x}) + \Delta_r^{cw}(\vec{y})$ . We also have  $\Delta_r(\vec{x} + \vec{y}) \leq \Delta_r(\vec{x}) + \Delta_r(\vec{y})$ .*



*Proof.* For any coordinate  $i \in [d]$ , the following holds:

$$\Delta^{cw}(\vec{x} + \vec{y})[i] = \max_{j,k \in [h]} \left| x^{(j)}[i] + y^{(j)}[i] - x^{(k)}[i] - y^{(k)}[i] \right| \quad (212)$$

$$\leq \max_{j,k \in [h]} \left\{ \left| x^{(j)}[i] - x^{(k)}[i] \right| + \left| y^{(j)}[i] - y^{(k)}[i] \right| \right\} \quad (213)$$

$$\leq \max_{j,k \in [h]} \left| x^{(j)}[i] - x^{(k)}[i] \right| + \max_{j',k' \in [h]} \left| y^{(j')}[i] - y^{(k')}[i] \right| \quad (214)$$

$$= \Delta^{cw}(\vec{x})[i] + \Delta^{cw}(\vec{y})[i], \quad (215)$$

which concludes the proof for coordinate-wise diameters. The proof for  $\ell_r$  diameters is similar.  $\square$

## D.2 Correctness proof of RB-TM and Lower bound on Byzantine tolerance

We now move on to the proof of correctness of RB-TM for  $n \geq 3f + 1$ . First, denoting  $S^{(j)}[i] = S(\vec{z}^{(j)}[i])$ , note that we have the following lemma.

**Lemma 20.** *For any two honest nodes  $j$  and  $k$  and any coordinate  $i$ , we have*

$$\left| S^{(j)}[i] - S^{(k)}[i] \right| \leq f. \quad (216)$$

*Proof.* Note that node  $j$  receives messages from at most  $n$  nodes, and in the trimming step,  $2f$  nodes are discarded, which yields  $|S^{(j)}[i]| \leq n - 2f$ . Moreover, among the nodes in  $Q^{(j)} \cap Q^{(k)}$  at most,  $f$  nodes with the smallest and  $f$  nodes with the largest  $i$ -th coordinates will be trimmed. Now recall that  $|Q^{(j)} \cap Q^{(k)}| \geq q$ , thus,  $|S^{(j)}[i] \cap S^{(k)}[i]| \geq q - 2f = n - 3f$ . We then obtain

$$\left| S^{(j)}[i] - S^{(k)}[i] \right| = \left| S^{(j)}[i] \right| - \left| S^{(j)}[i] \cap S^{(k)}[i] \right| \leq f, \quad (217)$$

which is the lemma.  $\square$

We denote  $x^{(min)}[i] = \min_{j \in [h]} x^{(j)}[i]$  and  $x^{(max)}[i] = \max_{j \in [h]} x^{(j)}[i]$  the minimal and maximal  $i$ -th coordinate among the parameters of the honest nodes.

**Lemma 21.** *All the values that are not discarded in the trimming step are within the range of the values proposed by the honest nodes, i.e.,*

$$\forall i \in [d], \forall j \in [h], \forall k \in S^{(j)}[i], x^{(min)}[i] \leq w^{(k)}[i] \leq x^{(max)}[i] \quad (218)$$

*Proof.* Note that there exist at most  $f$  Byzantine nodes that might broadcast vectors with the  $i$ -th coordinate larger than  $x^{(max)}[i]$  or smaller than  $x^{(min)}[i]$ , and all of these nodes will be removed by trimming.  $\square$

**Lemma 22** (Contraction by TM). *Under Assumption 6, TM guarantees the contraction of the coordinate-wise diameters, that is,*

$$\forall \vec{x}, \forall \overrightarrow{\text{BYZ}}, \Delta^{cw} \left( \overrightarrow{\text{TM}} \circ \overrightarrow{\text{BYZ}}(\vec{x}) \right) \leq (1 - \varepsilon) \Delta^{cw}(\vec{x}). \quad (219)$$

*As an immediate corollary, this inequality holds by taking the  $\ell_r$ -norm on both sides, which means that the coordinate-wise  $\ell_r$ -diameter is also contracted by the same factor.*

*Proof.* Let us bound the distance between  $y^{(j)}[i]$  and  $y^{(k)}[i]$ , the  $i$ -th coordinate of the outputs of two arbitrary nodes  $j$  and  $k$  after a Byzantine attack. Denote

$$m = \frac{1}{|S^{(j)}[i] \cap S^{(k)}[i]|} \sum_{l \in S^{(j)}[i] \cap S^{(k)}[i]} w^{(l)}[i], \quad (220)$$

the average of the  $i$ -th coordinates of the nodes in  $S^{(j)}[i] \cap S^{(k)}[i]$ . Without loss of generality, assume  $y^{(j)}[i] \geq y^{(k)}[i]$ . We then obtain

$$\left| y^{(j)}[i] - y^{(k)}[i] \right| = \frac{1}{|S^{(j)}[i]|} \sum_{l \in S^{(j)}[i]} w^{(l)}[i] - \frac{1}{|S^{(k)}[i]|} \sum_{l \in S^{(k)}[i]} w^{(l)}[i] \quad (221)$$

$$= \left( m + \frac{1}{|S^{(j)}[i]|} \sum_{l \in S^{(j)}[i] - (S^{(j)}[i] \cap S^{(k)}[i])} (w^{(l)}[i] - m) \right) \quad (222)$$

$$- \left( m + \frac{1}{|S^{(k)}[i]|} \sum_{l \in S^{(k)}[i] - (S^{(j)}[i] \cap S^{(k)}[i])} (w^{(l)}[i] - m) \right) \quad (223)$$

$$= \frac{1}{|S^{(j)}[i]|} \sum_{l \in S^{(j)}[i] - S^{(k)}[i]} (w^{(l)}[i] - m) - \frac{1}{|S^{(k)}[i]|} \sum_{l \in S^{(k)}[i] - S^{(j)}[i]} (w^{(l)}[i] - m) \quad (224)$$

$$\leq \frac{1}{|S^{(j)}[i]|} \sum_{l \in S^{(j)}[i] - S^{(k)}[i]} (x^{(max)}[i] - m) - \frac{1}{|S^{(k)}[i]|} \sum_{l \in S^{(k)}[i] - S^{(j)}[i]} (x^{(min)}[i] - m) \quad (225)$$

$$= \frac{|S^{(j)}[i] - S^{(k)}[i]|}{|S^{(j)}[i]|} (x^{(max)}[i] - m) + \frac{|S^{(k)}[i] - S^{(j)}[i]|}{|S^{(k)}[i]|} (m - x^{(min)}[i]), \quad (226)$$

where the inequality uses Lemma 21. Note that Lemma 21 also implies that  $x^{(min)}[i] \leq m \leq x^{(max)}[i]$  since  $m$  is the average of some real numbers, all of which are within the range of the values proposed by the honest nodes. Now notice that using Lemma 20 we have

$$\frac{|S^{(j)}[i] - S^{(k)}[i]|}{|S^{(j)}[i]|} = \frac{|S^{(j)}[i] - S^{(k)}[i]|}{|S^{(j)}[i] - S^{(k)}[i]| + |S^{(j)}[i] \cap S^{(k)}[i]|} \quad (227)$$

$$\leq \frac{f}{f + |S^{(j)}[i] \cap S^{(k)}[i]|} \quad (228)$$

$$\leq \frac{f}{f + \varepsilon f} = 1 - \tilde{\varepsilon}, \quad (229)$$

where we used the fact that  $|S^{(j)}[i] \cap S^{(k)}[i]| \geq q - 2f = n - 3f \geq \varepsilon f$ , and similarly,

$$\frac{|S^{(k)}[i] - S^{(j)}[i]|}{|S^{(k)}[i]|} \leq 1 - \tilde{\varepsilon}. \quad (230)$$

Combining these inequalities with Equation (226), we then obtain

$$\left| y^{(j)}[i] - y^{(k)}[i] \right| \leq (1 - \tilde{\varepsilon}) (x^{(max)}[i] - x^{(min)}[i]) = (1 - \tilde{\varepsilon}) \Delta^{cw}(\vec{x})[i]. \quad (231)$$

Therefore, we have

$$\Delta^{cw}(\vec{y})[i] \leq (1 - \tilde{\varepsilon}) \Delta^{cw}(\vec{x})[i], \quad (232)$$

which is what we wanted.  $\square$

**Remark 5.** Note that, in the regime  $f \ll n$ , we can set  $\varepsilon \rightarrow \infty$ , in which case we have  $\tilde{\varepsilon} \rightarrow 1$ . Thus, for any fixed value of  $N$ , when  $f$  is sufficiently small compared to  $n$ , RB-TM actually achieves asymptotic agreement in only one communication round.

We have the following corollary regarding RB-TM.

**Corollary 1.** Under Assumption 6, for any input  $N \in \mathbb{N}$ , the algorithm RB-TM guarantees Byzantine asymptotic agreement, i.e.,

$$\Delta_2 \left( \overrightarrow{\text{RB-TM}}_N \circ \overrightarrow{\text{BYZ}}_N(\vec{x}) \right) \leq \frac{\Delta_2(\vec{x})}{2^N}. \quad (233)$$

*Proof.* First note that since RB-TM iterates TM, there is actually a sequence of attacks  $\overrightarrow{\text{BYZ}}_t$  at each iteration  $t \in [T_{\text{RB-TM}}(N)]$ . In fact, we have a sequence of families  $\vec{y}_t$  defined by  $\vec{y}_0 \triangleq \vec{x}$  and  $\vec{y}_{t+1} \triangleq \overrightarrow{\text{TM}} \circ \overrightarrow{\text{BYZ}}_t(\vec{y}_t)$  for  $t \in [T_{\text{RB-TM}}(N) - 1]$ . We eventually have  $\overrightarrow{\text{RB-TM}} \circ \overrightarrow{\text{BYZ}}(\vec{x}) = \vec{y}_{T_{\text{RB-TM}}(N)}$ .

Note that the previous lemma implies that

$$\Delta^{cw}(\vec{y}_{t+1}) \leq (1 - \tilde{\varepsilon}) \Delta^{cw}(\vec{y}_t). \quad (234)$$

Taking the  $\ell_2$  norm on both sides then implies that

$$\Delta_2^{cw}(\vec{y}_{t+1}) = \|\Delta^{cw}(\vec{y}_{t+1})\|_2 \leq (1 - \tilde{\varepsilon}) \|\Delta^{cw}(\vec{y}_t)\|_2 = (1 - \tilde{\varepsilon}) \Delta_2^{cw}(\vec{y}_t). \quad (235)$$

It follows straightforwardly that

$$\Delta_2^{cw}(\vec{y}_{T_{\text{RB-TM}}(N)}) \leq (1 - \tilde{\varepsilon})^{T_{\text{RB-TM}}(N)} \Delta_2^{cw}(\vec{x}) \quad (236)$$

$$\leq (1 - \tilde{\varepsilon})^{\frac{(N+1) \ln 2 + \ln \sqrt{h}}{\tilde{\varepsilon}}} \Delta_2^{cw}(\vec{x}) \quad (237)$$

$$= \exp\left(\frac{\ln(1 - \tilde{\varepsilon})}{\tilde{\varepsilon}} ((N+1) \ln 2 + \ln \sqrt{h})\right) \Delta_2^{cw}(\vec{x}) \quad (238)$$

$$\leq \exp(-(N+1) \ln 2) \exp(-\ln \sqrt{h}) \Delta_2^{cw}(\vec{x}) \quad (239)$$

$$= \frac{1}{2^{1+N} \sqrt{h}} \Delta_2^{cw}(\vec{x}), \quad (240)$$

where, in Equation (239), we used  $\ln(1 + u) \leq u$  for  $u \in (-1, 0]$ . We now conclude by invoking Lemma 18, which implies

$$\Delta_2(\vec{y}_{T_{\text{RB-TM}}(N)}) \leq \Delta_2^{cw}(\vec{y}_{T_{\text{RB-TM}}(N)}) \leq 2^{-N} \frac{\Delta_2^{cw}(\vec{x})}{2\sqrt{h}} \leq 2^{-N} \Delta_2(\vec{x}), \quad (241)$$

which proves that RB-TM achieves asymptotic agreement.  $\square$

We now prove our theorem.

*Proof of theorem 6.* Consider a family  $\vec{x}_0 \in \mathbb{R}^{d \cdot h}$ . We first focus on coordinate  $i \in [d]$  only. We sort the family using a permutation  $\sigma$  of  $[h]$ , so that

$$x_0^{(\sigma(1))}[i] \leq x_0^{(\sigma(2))}[i] \leq \dots \leq x_0^{(\sigma(h-1))}[i] \leq x_0^{(\sigma(h))}[i]. \quad (242)$$

Now denote  $q_j = |Q^{(j)}|$ , and  $\vec{z} = \overrightarrow{\text{BYZ}}_0^{(j)}(\vec{x}_0) = \vec{w}_0^{(Q^{(j)})} \in \mathbb{R}^{d \cdot q_j}$  the result of a Byzantine attack. Again, we sort the vectors of this family, using a permutation  $\tau$  of  $[q_j]$ , so that

$$z^{(\tau(1))}[i] \leq z^{(\tau(2))}[i] \leq \dots \leq z^{(\tau(q_j-1))}[i] \leq z^{(\tau(q_j))}[i]. \quad (243)$$

Now, denoting  $y \triangleq \text{TM}(\vec{z})$ , we note that

$$y[i] = \frac{1}{q_j - 2f} \sum_{k=1}^{q_j - 2f} z^{(\tau(f+k))}[i]. \quad (244)$$

Moreover, note that there are  $f + k - 1$  values of  $\vec{z}$  that are smaller than  $z^{(\tau(f+k))}[i]$ . These can include  $f$  Byzantine vectors. But the remaining  $k - 1$  values must then come from the family of honest vectors. Yet, the  $k - 1$  smallest vectors of this family are  $x_0^{(\sigma(1))}[i], \dots, x_0^{(\sigma(k-1))}[i]$ . But then,  $z^{(\tau(f+k))}[i]$  will have to take a value on the right of  $x_0^{(\sigma(k-1))}[i]$  in the list of honest vectors, which corresponds to saying that

$$\forall k \in [q_j - f], z^{(\tau(f+k))}[i] \geq x_0^{(\sigma(k))}[i]. \quad (245)$$

But then, we know that

$$y[i] \geq \frac{1}{q_j - 2f} \sum_{k=1}^{q_j - 2f} x_0^{(\sigma(k))}[i]. \quad (246)$$

As an immediate corollary, we see that  $y[i] \geq x_0^{(\sigma(1))}[i]$ , which also implies that

$$x_0^{(\sigma(k))}[i] \leq x_0^{(\sigma(1))}[i] + \max_{l \in [h]} \left( x_0^{(\sigma(l))}[i] - x_0^{(\sigma(1))}[i] \right) \leq y[i] + \Delta^{cw}(\vec{x}_0)[i]. \quad (247)$$

But now notice that

$$\bar{x}_0[i] = \frac{1}{h} \sum_{k=1}^h x_0^{(\sigma(k))} = \frac{1}{h} \sum_{k=1}^{q_j-2f} x_0^{(\sigma(k))} + \frac{1}{h} \sum_{k=q_j-2f+1}^h x_0^{(\sigma(k))} \quad (248)$$

$$\leq \frac{1}{h} ((q_j - 2f)y[i]) + \frac{1}{h} \sum_{k=q_j-2f+1}^h (y[i] + \Delta^{cw}(\bar{x}_0)[i]) \quad (249)$$

$$= y[i] + \frac{h - q_j + 2f}{h} \Delta^{cw}(\bar{x}_0)[i]. \quad (250)$$

Similarly, we can also prove that  $\bar{x}_0[i] \geq y[i] - \frac{h - q_j + 2f}{h} \Delta^{cw}(\bar{x}_0)[i]$ , which implies that

$$|y[i] - \bar{x}_0[i]| \leq \frac{h - q_j + 2f}{h} \Delta^{cw}(\bar{x}_0)[i] \leq \frac{2f}{h} \Delta^{cw}(\bar{x}_0)[i], \quad (251)$$

where we used the fact that  $q_j \geq q = h$ . Thus,  $\|y - \bar{x}_0\|_2 \leq \frac{2f}{h} \|\Delta^{cw}(\bar{x}_0)\|_2 = \frac{2f}{h} \Delta_2^{cw}(\bar{x}_0)$ . In fact, more generally, we showed that, for any Byzantine attack  $\overrightarrow{\text{BYZ}}_0^{(j)}$ , we have

$$\text{TM} \circ \overrightarrow{\text{BYZ}}_0^{(j)}(\bar{x}_0) \in Y_0 = \bar{x}_0 + \frac{2f}{h} \prod_{i \in [d]} [-\Delta^{cw}(\bar{x}_0)[i], +\Delta^{cw}(\bar{x}_0)[i]]. \quad (252)$$

Yet Lemma 21 shows that any such parallelepiped was stable under application of TM despite Byzantine attacks. Thus, for any iteration  $t \geq 1$ , we still have  $x_t^{(j)} \in Y_0$ , which then guarantees that

$$\|\bar{x}_t - \bar{x}_0\|_2 = \left\| \frac{1}{h} \sum_{j \in [h]} (x_t^{(j)} - \bar{x}_0) \right\|_2 \leq \frac{1}{h} \sum_{j \in [h]} \|x_t^{(j)} - \bar{x}_0\|_2 \quad (253)$$

$$\leq \frac{1}{h} \sum_{j \in [h]} \frac{2f}{h} \|\Delta^{cw}(\bar{x}_0)\|_2 = \frac{2f}{h} \Delta_2^{cw}(\bar{x}_0). \quad (254)$$

Lemma 18 then guarantees  $\Delta_2^{cw}(\bar{x}_0) \leq 2\sqrt{h}\Delta_2(\bar{x}_0)$ . By noting that RB-TM corresponds to iterating TM, we conclude that RB-TM achieves  $\frac{4f}{\sqrt{h}}$ -averaging agreement.

Now we show that for  $n \leq 3f$ , no algorithm can achieve Byzantine averaging agreement. If  $n \leq 3f$ , then  $h = n - f \leq 2f$ . Thus honest nodes can be partitioned into two subsets of cardinals at most  $f$ . In particular, for any subset, Byzantine nodes can block all messages coming from the other subset. Any subset would thus only hear from nodes of the subset and from the Byzantine nodes.

Assume now by contradiction that AVG achieves Byzantine-resilience averaging agreement for  $n \leq 3f$ . Note that we then have  $q \leq 2f$ . As a result  $q - f \leq f$ . But as a result, if Byzantines send  $\vec{z} = z \star f$  to all honest nodes, quasi-unanimity (Lemma 17) applies, which means that all honest nodes must output  $z$ .

But this hold for any value  $z$  chosen by the Byzantine nodes. Clearly, this prevents averaging. Thus AVG fails to achieve averaging agreement.  $\square$