

This work is on a Creative Commons Attribution 4.0 International (CC BY 4.0) license, <https://creativecommons.org/licenses/by/4.0/>. Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

Tolerating Adversarial Attacks and Byzantine Faults in Distributed Machine Learning

Yusen Wu¹, Hao Chen¹, Xin Wang¹, Chao Liu¹, Phuong Nguyen^{1,2}, Yelena Yesha^{1,3}

¹University of Maryland, Baltimore County, Baltimore, MD, USA

²OpenKneck Inc, Halethorpe, MD, USA

³University of Miami, FL, USA

{ywu5, chen hao1, xinwang11, chaoliu717, phuong3, yeyesha}@umbc.edu

Abstract—Adversarial attacks attempt to disrupt the training, retraining, and utilizing of artificial intelligent and machine learning models in large-scale distributed machine learning systems. This causes security risks on its prediction outcome. For example, attackers attempt to poison the model by either presenting inaccurate misrepresentative data or altering the models’ parameters. In addition, Byzantine faults including software, hardware, network issues occur in distributed systems which also lead to a negative impact on the prediction outcome. In this paper, we propose a novel distributed training algorithm, partial synchronous stochastic gradient descent (ParSGD), which defends adversarial attacks and/or tolerates Byzantine faults. We demonstrate the effectiveness of our algorithm under three common adversarial attacks against the ML models and a Byzantine fault during the training phase. Our results show that using ParSGD, ML models can still produce accurate predictions as if it is not being attacked nor having failures at all when almost half of the nodes are being compromised or failed. We will report the experimental evaluations of ParSGD in comparison with other algorithms.

Index Terms—Data security, Byzantine-resilient SGD, Distributed ML

I. INTRODUCTION

Adversarial attacks against AI and ML attempt to disrupt the training and retraining process by either injecting inaccurate misrepresentative data samples or altering the models’ parameters. In distributed machine learning, each participant trains the global model using its own local data and shares the model’s updates only with central servers or one server. An adversary may compromise workers during the training phase with malicious data samples such as false labels or input. For example, attackers can compromise the mobile devices by sending the junk gradients and parameter updates to the server which severely impact the predicted outcome of the model.

Distributed Machine Learning (ML) or Federated Learning (FL) [18], [23] provides solutions for scaling ML models when training with real-world big data. Usually, with traditional single-data center or distributed data across a range of geographic places, multiple computation parties include the cluster of CPUs, multiple GPUs, or TPUs are aggregated for speed-up and collaboratively train a global model. For advanced distributed ML, the frameworks like Parameter Servers [22], MapReduce [12], Graph-parallel [19], and All-reduce [4] are widely used in popular open-source machine learning library

including Tensorflow [2], Horovod [31], Pytorch [28], Spark [37], etc.

In distributed training, synchronous and asynchronous modes are used in the above frameworks for data parallelism. Most approaches use stochastic gradient descent (SGD) based algorithms [7], which are the iterative algorithms for optimizing loss functions to train ML models. For the synchronous method, all workers train the same target model using the different portions of the dataset then take the average of gradients at each iteration (epoch) via All-reduce communication. In asynchronous mode, all workers train the model using their local batch of data independently, then send the updates to a parameter server asynchronously. Several security concerns arise in such a distributed environment since malicious attacks or failures are involved.

Previous works have been studied to mitigate the adversarial attacks in collaborative and distributed learning, [36], [6], [24], [11], [34], [35], [9], [26], and [15]. As we said, in distributed training settings, workers will send updates (gradients) to the server, the server will aggregate all the gradients, calculate the updates, and then return the updated value to all the workers for the next round of training. Workers can be compromised by malicious adversaries or are vulnerable to failures such as software, hardware, and network faults (Byzantine failures). What’s more, the averaging aggregation rule like Federated Averaging (FedAvg) [23], can not tolerate even one outlier because a Byzantine worker can simply tamper its gradient with a wrong direction, leading the averaging vector far away from the correct direction. Current approaches like Krum [6] are easy to tolerate outliers if the bad gradients are far away from the correct ones. But usually, the malicious workers may pretend they are correct nodes among the system. Some algorithms are not efficient to tolerate the bad gradients that are hidden among the correct ones. To sum up, building a privacy-preserving and Byzantine-resilient distributed machine learning system becomes extremely important and challenging.

In this paper, we analyze and address three main questions: (i), *are there any reasonable solutions to tolerate at most half Byzantine workers in a proper and explainable way?* (ii), *if some of the workers are to be crashed, how can we*

Byzantine-resilient GARs	Fault Tolerance	Mode	Assumption	Time Complexity
Median and Trimmed-mean [36]	$n \geq 2f + 2$	○	weak assumption	$\mathcal{O}(n \cdot d)$
Krum [6]	$n \geq 2f + 1$	○	strong assumption	$\mathcal{O}(n^2 \cdot d)$
multi-Krum [10]	$n \geq 2f + 3$	○	strong assumption	$\mathcal{O}(n^2 \cdot d)$
Bulyan [24]	$n \geq 4f + 3$	○	strong assumption	$\mathcal{O}(n^2 \cdot d)$
Kardam [11]	$n \geq 3f + 1$	●	strong assumption	$\mathcal{O}(d + n \cdot f)$
Zeno [34]	unbounded	○	weak assumption	$\mathcal{O}(n \cdot d)$
Zeno+ and Zeno++ [35]	unbounded	●	weak assumption	$\mathcal{O}(n \cdot d)$
Draco [9]	$n \geq 2f + 1$	○	strong assumption	$\mathcal{O}(n \cdot d)$, run twice
Hogwild! [26]	unbounded	●	strong assumption	$\mathcal{O}(n \cdot d)$
ByzSGD [15]	$n \geq 2f + 1$ $n \geq 3f + 2$	●	strong assumption	$\mathcal{O}(n^2 \cdot d)$
ParSGD	$n \geq 2f + 1$	◐	weak assumption	$\mathcal{O}(n \cdot d)$

TABLE I: Comparison of Byzantine-resilient GARs. ○ synchronous ● asynchronous ◐ partial synchronous.

detect and extract them from the total gradients so that the ML models can converge resiliently during the training? (iii), some Byzantine-resilient gradient aggregation rules (GARs) have been proposed and tested under the strong or weak byzantine assumption, but we found some of them may reach good accuracy in strong assumption and bad in weak assumption, and vice versa. Can we create a new aggregation rule which works in both strong and weak assumptions?

We propose a new GAR solution which is based on the mean of f nearest neighbors to the median among all gradients submitted by workers. We will show that this solution can tolerate at most $f \simeq \frac{n}{2} - 1$ Byzantine workers and degrade the time complexity from $\mathcal{O}(d \cdot n^2)$ to $\mathcal{O}(d \cdot n)$ in comparison with Krum, Multi-Krum and so on. Our GAR, ParSGD, can find g as a new median when $f \leq \frac{n-1}{2}$ workers are attacked. After finding this new median, we try to find f gradients closest to this new median and return the mean of $f + 1$ gradients (including the median) as a new parameter for the next epoch training. We define both strong Byzantine resilience and weak Byzantine resilience based on the ParSGD and the formal definitions are listed in Section IV. In short, *if f nodes are Byzantine workers and their gradients are far away from the median, we define it as strong Byzantine resilience because the aggregation rule will never choose the Byzantine gradients. However, if some Byzantine nodes are pretending to be correct workers and mix up with correct ones, we define it as weak Byzantine resilience.* Accuracy will be used as a performance metric for evaluating our ParSGD in comparison with other algorithms. We also propose to use an unknown bounded time δt for excluding the crash workers compare the results to other solutions.

Contribution. In summary, the three primary contributions of this work are as follows:

- We propose a new aggregation rule, named ParSGD, to tolerate Byzantine failures in distributed ML systems. Our experimental results show that the accuracy of ParSGD can reach near $f \simeq \frac{n}{2} - 1$ Byzantine workers among n workers with $\mathcal{O}(n \cdot d)$ time complexity. Compared with Mean, Median, and Krum, we get the best accuracy under three common attacks with two different datasets.

- Theoretically, we redefine strong and weak Byzantine resilience definitions based on ParSGD, and proved ParSGD can reach a relatively stable accuracy under both of these two assumptions.
- Our ParSGD is a new partial synchronous GAR based on an unknown bounded time δt which can efficiently exclude crash workers to make the training converge faster. Both theoretical and experimental analyses are also provided.

Paper Organization. Section II discusses related work. Section III introduces the motivations and some background definitions. Section IV introduces our new proposed aggregation rule ParSGD, definitions of strong and weak Byzantine resilience in partial synchronous SGD, time complexity analysis, and proof of its Byzantine resilience. Section V analyzes the convergence of a distributed SGD using ParSGD. Section VI presents our experimental evaluation of ParSGD and some discussion. We summarize our conclusion in Section VII.

II. RELATED WORK

A. Byzantine fault tolerance.

In previous studies on Byzantine fault tolerate (BFT), several consensus protocols and systems have been proposed to tolerate arbitrary faults in distributed system (e.g., PBFT [8], Raft [27], BFT-Smart [5], Honey Badger [25], Chios [13], etc). All of them need to make a consensus before delivery; however, it is costly in terms of communication overhead for employing consensus in distributed ML architecture. Because it may have thousands of workers processing very large datasets and all workers need to reach a consensus. On the other hand, the design of ParSGD is inspired by the concept of BFT to tolerate Byzantine workers which may submit malicious updates. However, we do not directly adopt BFT since we achieve a slightly different goal: BFT achieves the total order of events while we seek to tolerate malicious updates from the workers.

B. Byzantine-resilient SGD.

To tolerate the outliers, robust statistics have been proposed. We summarize them in Table I. Yin et al. [36] proposed

Median¹ and Trimmed-mean solutions, in which the server sorts all of the gradients and takes median as the global parameter for next round training, same as the Trimmed-mean which needs to remove a percentage of outliers after sorting the gradients. However, a recent paper proved that the Median aggregation rule is still under an order-optimal error rate [16]. Blanchard et al. [6] proposed Krum for selecting a valid vector update. This rule has local time complexity $\mathcal{O}(d \cdot n^2)$ which makes it relatively expensive to compute when the d and n are large. If Krum can fully tolerate Byzantine workers, time can be a good trade-off; however, the Krum assumes only one neighbor pretends as a valid vector among the correct vectors, and in each iteration, the algorithm chooses $n - f - 2$ neighbors which are not reasonable because Byzantine workers may corrode more workers to pretend they are correct. What's more, Chen et al. [9] proposed Draco which leverages a gradient-coding based algorithm for robust learning. EI Mahdi et al. [24] proposed a strong Byzantine algorithm, named Bulyan, which needs $n \geq 4f + 3$ for tolerating fewer Byzantine workers and convergence requires strong assumptions. To preserve training convergence, Sohn et al. [32] proposed a voting-based authentication to tolerate inaccurate training results. Alistarh et al. [3] utilize historical information to achieve the best sample complexity of training. Xie et al. [34] proposed Zeno to tolerate the cases where the majority of workers are not fully trusted during training. Xie et al. [33] also demonstrated that Median and Krum exhibit poor performance. Other asynchronous algorithms also have been proposed (e.g., Zeno++ [35], Kardam [11], Hogwild! [26], ByzSGD [15]), but inevitably the lack of information on the descent directions will cause a low training accuracy since asynchronous models do not depend on strict arrival times of messages for reliable operations.

III. PROBLEM SETTINGS AND BACKGROUND

We consider a large-scale distributed machine learning system with Byzantine-resilient GARs. At the core of our proposed gradient aggregation rule (GAR) ParSGD is to find a new vector \mathbf{g} as correct medians after $f \leq \frac{n-1}{2}$ workers are attacked in d -dimensional vectors. And calculate the mean of f gradients near the median. To better understand our algorithm, in this section, we first compare available approaches of Byzantine-resilient GARs. Then we describe problem settings, relevant attacks, and some background definitions of our ParSGD.

A. Byzantine-resilient GARs.

Several distributed ML aggregation rules $F(\cdot)$ have been proposed to tolerate Byzantine faults. We summarize some of them in Table I. Median [36] and Krum [6] have become the most efficient and effective ones for tolerating Byzantine faults. Authors showed that in theory, both of them can tolerate at most $\frac{1}{2}$ Byzantine workers, we say their *breakdown points* [17] can reach $\frac{1}{2}$. We will show that the sample median is

¹In the paper, we define the uppercase Median as a GAR solution and the lowercase median as the middle value.




Fig. 1: (a) the number of correct gradients larger than Byzantine gradients, (b) some Byzantine gradients pretend they are correct, B becomes the one closest to the *barycenter* among n gradients, (c) f gradients close to median.

an unbiased estimator which can also reach the *breakdown point* in Section 4. Krum proposed a solution using $n - f - 2$ squared-distance-based GAR to find a vector that is closest to the *barycenter*² among n vectors. The algorithm is resilient under random Gaussian attack when the f vectors are far from the correct vectors with a small variance. Fig1(a) shows f Byzantine workers are at a distance far away from the correct workers. As long as the number of correct workers is more than Byzantine workers, then the Krum can always find a correct gradient and sum up its nearby $n - f - 2$ neighbors to find a minimum vector among the correct workers. But it can hardly tolerate the Bit-flipping attack which may change all the gradients to opposite directions when there are 50% Byzantine workers. Fig1(b) shows under Bit-flipping attack where B can never find enough correct neighbors and would to count Byzantine gradients instead, and Krum will return B as the final vector for next round training. Our ParSGD will find the median first, then take the mean of f gradients near the median, as shown in Fig1(c). The Median solution has the same issue, if $n > 2f$, and the variance is large enough, the real median still can be changed to an imprecise gradient even there is only one Byzantine worker, although the new median is among the correct workers. In addition, both Krum and Median are just in search of one optimal vector among n vectors and this vector is vulnerable to be attacked. We will show that these two solutions are not available and have not enough training accuracy under around $\frac{1}{3}$ of Byzantine workers in the evaluation section, even both of their *breakdown points* can reach $\frac{1}{2}$ as shown in Section 6. Furthermore, none of these synchronous GARs can tolerate crash-stop failure. Except one of the GARs, named Hogwild! [26], employs a coordinator to monitor the number of workers which is time-consuming.

²<https://en.wikipedia.org/wiki/Barycenter>

B. Problem settings

We consider the problem of minimizing an loss function that has the form of a sum, $Q(w) = \frac{1}{n} \sum_{i=1}^n Q_i(w)$. Each of function Q_i is associated with the i -th training set. The sum-minimization problem also arises for empirical risk minimization (ERM) [38]. In this case, $Q_i(w)$ is the value of the loss function, and $Q(w)$ is the empirical risk. When used to minimize the above function, a standard gradient descent method would perform the following iterations:

$$\mathbf{w}^{t+1} := \mathbf{w}^t - \gamma \nabla Q(\mathbf{w}^t) = \mathbf{w}^t - \frac{\gamma}{n} \sum_{i=1}^n \nabla Q_i(\mathbf{w}^t) \quad (1)$$

where t is the index of epoch, γ is learning rate, and the aggregation rule $F = \frac{1}{n} \sum_{i=1}^n \nabla Q_i(\mathbf{w}^t)$.

We assume that α fraction of the n workers are Byzantine and $1 - \alpha$ fraction are correct workers. The Byzantine workers will not obey the rule of GARs by sending arbitrary messages or pretending they are valid among n correct workers. In addition, correct workers can be crashed down due to software, hardware, or communication issues. In our assumption, we assume that a system can only tolerate at most 50% of Byzantine workers. If f workers are Byzantine, then $n - f$ correct workers must be larger than f , we then get $n > 2f$. When $n \leq 2f$, the GARs still can choose the minority group of gradients and aggregate them for an update, but the result is not under Byzantine resilience assumptions because it breaks the principle of minority versus majority. Under this premise, median is the most appealing unbiased estimator, however, it can be maliciously altered to a new one and no longer be accurate.

C. Relevant Attacks

1) *Crash-stop failure*: A crash failure [30] occurs when a node suffers from an omission failure once, and then continues to not respond.

2) *Bit-flipping Attack*: Bit-flipping attack [29] is an attack which the attacker can change the ciphertext to result in a predictable change of the plaintext.

3) *Random Gaussian Attack*: Random Gaussian Attack also known as Gaussian Noise [1], it's a statistical noise having a probability density function equal to normal distribution.

D. Relevant Definitions

We introduce the coordinate-wise median and some other relevant definitions as follows, which serve as cornerstones for our algorithms.

Definition 1 (*Coordinate-wise median*): For vectors V_i in \mathbb{R}^d , $i \in [1, n]$, the coordinate-wise median $\mathbf{g} = \text{med}\{V_i : i \in [1, n]\}$ is a vector with its k -th coordinate being $g_k = \text{med}\{V_i^k : i \in [1, n]\}$ for each $k \in [d]$, where med is a one-dimensional median and \mathbf{g}_k is one column of k medians.

Definition 2 (*Bounded variance*) $\forall i \in \{1, \dots, n\}$, $\text{diag}(\mathbb{E}[(\mathbb{E}V_i - \boldsymbol{\mu})(\mathbb{E}V_i - \boldsymbol{\mu})^T]) \leq \sigma^2$.

Definition 3 (*Absolute skewness*): For one-dimensional

random variable X , define the absolute skewness as $\gamma(X) = \mathbb{E}[(\frac{X-\mu}{\sigma})^3] = \frac{\mathbb{E}[(X-\mu)^3]}{\sigma^3} = \frac{\mathbb{E}[(X-\mu)^3]}{(\mathbb{E}[(X-\mu)^2])^{3/2}} = \frac{K_3}{K_2^{3/2}}$, and let S_k be the skewness of random distribution, say $S_k = \frac{m_3}{m_2^{3/2}} = \frac{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^3}{[\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2]^{3/2}}$, where \bar{x} defined as sample mean, m_2 as the second central moment (variance), and m_3 defined as third central moment.

Definition 4 (*Lipschitz continuity*): $\forall (w, w')$, h is L -Lipschitz if $\|h(w) - h(w')\| \leq L\|w - w'\|_2$

Definition 5 (*Smoothness*): h is L' -smooth if and only if $\forall (w, w')$, $\|h'(w) - h'(w')\|_2 \leq L'\|w - w'\|_2$

Definition 6 (*Strong convexity*): h is μ -strongly convex if and only if $\forall (w, w')$, $h(w') \geq h(w) + h'(w)^\top (w' - w) + \frac{\mu}{2}\|w' - w\|_2^2$.

IV. PARTIAL SYNCHRONOUS BYZANTINE-RESILIENT SGD

We now introduce our partial gradient aggregation rule, ParSGD, which satisfies both weak and strong Byzantine resilience. The detailed distributed synchronous SGD is shown in *Algorithm 1*.

In a normal distribution, the median equals to mean. Namely, if $|g_k - \mu_k|$ is relatively small, then we can achieve good training accuracy. *Theorem 1* proved that the *breakdown point* of median can reach 50%, but it doesn't mean it will always be 50% because it can easily be altered if Byzantine workers are pretending. In this section, we proposed two corollaries, *Corollary 1&2*, which help explain why to take the mean of closest gradients near median is an excellent solution.

Theorem 1. *With the fraction of Byzantine workers as α and the sample median as the estimator, the asymptotic breakdown point [17] for α is 1/2.*

Proof. If we have m data points and we let a minority of $\lfloor \frac{m-1}{2} \rfloor$ points become outliers leaving the rest of the fixed points, and the median stays with the majority. Usually median may change, but it does not become arbitrarily bad, then the sample breakdown point is

$$\begin{aligned} \frac{\lfloor \frac{m-1}{2} \rfloor}{m} &= \lfloor \frac{m-1}{2m} \rfloor \\ &\Rightarrow \frac{1 - \frac{1}{m}}{2}, \end{aligned}$$

the more m we have, the breakdown point will be closer to one-half. The asymptotic breakdown point is one-half. \square

A. Important Notations

The notations used in this paper is summarized in Table 2.

B. Partial Synchrony

We first consider timing assumption in partial synchronous SGD. In a synchronous system, there is a known fixed upper bound δt^3 on the time required for a message to be sent from one worker to another. In an asynchronous system, no

³We define t represents the index of an epoch, and δt represents a bounded time.

Notation	Meaning
i	index of worker
n	number of total workers
T	number of epochs
t	index of epochs
δt	unknown bounded time
c	number of crash workers
f	number of Byzantine workers
γ	learning rate
\mathbf{V}	all vector of d-dimensional gradients
\mathbf{U}_{f+1}	f vectors near median, including the median
$\bar{\mathbf{U}}$	take the mean of \mathbf{U}_{f+1}
g_k	the value of median in k -th dimension
\mathbf{g}	a vector of selected median
μ_k	average value in k -th dimension
F	gradient aggregation function

TABLE II: Important Notations

Algorithm 1: Partial Synchronous Byzantine SGD

Workers: $i = 1, \dots, n$

for $t = 1, \dots, T$ **do**

Waiting and receive \mathbf{w}^t from the server;
 Training, compute, and send the new gradients
 $\mathbf{V}_i^{t+1} = \nabla Q_i^t(\mathbf{w}^t)$ to the server

end

Server:

input: δt (Empirically and manually initialize a enough time δt);

for $t = 1, \dots, T$ **do**

Broadcast \mathbf{w}^t to all the workers;
 Wait only $2 * \delta t$ bounded time for $n - c$ gradients
 $\{\mathbf{V}_i^t : i \in [n - c]\}$ arrive;
 Recalculate and update the δt if all the gradients
 are being collected;
 Calculate f as defined in Definition 7, $f = \frac{n-c-1}{2}$;
 Compute $\bar{\mathbf{U}}$ as defined in Corollary 1 and 2;
 Update the parameter $\mathbf{w}^{t+1} = \mathbf{w}^t - \gamma^t \bar{\mathbf{U}}$;

end

fixed upper bound δt exists. And partial synchrony [14], fixed bound δt exists, but it's not known a priori⁴.

Proposition 1. *If a trusted GAR executes aggregation $F(\cdot)$ during a bounded and not-a-priori time δt , then $n - c$ active vectors are eventually collected, where n refers to the number of workers, and c refers to the crash nodes.*

In this proposition, the server only needs to initialize an estimated time δt which is enough for collecting all the gradients in one round. In the rest of the epochs, the GAR will recalculate the aggregation time δt , and only wait for $2 * \delta t$ bounded time which is not known a priori.

Proof. We initially set an unknown bounded time $2 * \delta t$ to

⁴https://en.wikipedia.org/wiki/A_priori_and_a_posteriori

collect $n - c$ active workers, the proof is simple and straightforward. We assume c number of workers may be crashed or compromised by malicious failures (crash-stop failure is also a kind of Byzantine failure). In this case, Byzantine workers may spend more training time than normal workers because of the extra time for Byzantine infection. Simply, we only wait for $2 * \delta t$ time in aggregation for the crash nodes, eventually, $n - c$ vectors will be collected. We exclude c crash workers which may be caused by network delay, malicious attack, or arbitrary system failures. \square

C. Byzantine Resilience

We introduced strong and weak Byzantine resilience in Section 1. In order to explain how our ParSGD work both in strong and weak Byzantine resilience, we formally define our strong and weak Byzantine resilience GARs in partial synchronous SGD and two Corollaries as follows.

Definition 7 (Weak Byzantine Resilience): *Given a system of n components, b of which are Byzantine workers, including c random crash-stop vectors as $\mathbf{C}_1, \dots, \mathbf{C}_c$ and f malicious workers pretending as corrects ones as $\mathbf{B}_1, \dots, \mathbf{B}_f$, where $0 \leq b \leq n, b = f + c$, and $b, c, f \in \mathbb{N}$. Let $\mathbf{V}_1, \dots, \mathbf{V}_{n-b}$ be the independent identically distributed(i.i.d.) random in \mathbb{R}^d , $\mathbf{V}_i \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, and with $\mathbb{E}V_i^k = \mu_k$, k refers to dimensional index. F is said to be weak Byzantine resilience if it satisfies*

$$F = F(\underbrace{\mathbf{V}_1, \dots, \mathbf{B}_1, \dots, \mathbf{B}_f, \dots, \mathbf{V}_{n-f-c}, \mathbf{C}_1, \dots, \mathbf{C}_c}_{n-c})$$

$\underbrace{\hspace{10em}}_n$

- (i) $n > 2f + c$, and
 - $f = 0, c = 0$ if Byzantine free
 - $0 \leq c < n, 0 \leq f \leq \lceil \frac{n-c-1}{2} \rceil$ if $f + c > 0$
- (ii) $|g_k - \mu_k| \leq \epsilon$, where ϵ is negligible, g_k refers to the median.

Definition 8 (Strong Byzantine Resilience): *F is said to be strong Byzantine resilience if it satisfies condition (i) and (ii) in Definition 7, and (iii) $|g_k - V_i^{far}| < |g_k - B_j^{close}|$, where g_k is the median, V_i^{far} refers to the farthest correct gradient, and B_j^{close} refers to the closest Byzantine gradient.*

Theorem 1 proved Median is an unbiased estimator which can tolerate at most $\frac{1}{2}$ Byzantine workers. Even there are $f \simeq \frac{n}{2} - 1$ Byzantine workers, Median can still find a new correct gradient as the new median to tolerate the arbitrary gradients attack. Based on Theorem 1 and Definition 7 and 8, we get two corollaries below.

Corollary 1: *Under the assumption of Definition 7, 8 and Theorem 1, suppose we have random vectors $\tilde{\mathbf{V}}_i \in \mathbb{R}^d, i \in [1, n - c]$ containing normal gradients and malicious ones. The set of f -neighbor vectors of \mathbf{g} is defined as $\mathbf{U}_{f+1} := \{\{\tilde{\mathbf{V}}_i\}_{i=1}^f, \mathbf{g}\}$, where $\{\tilde{\mathbf{V}}_i\}_{i=1}^f = \arg \min \sum_{i=1}^f \sum_{k=1}^d |V_{(i)}^k - g_k|$, for different $V_{(i)}^k \in V_j^k, i \in [1, f], j \in [1, n - f - c]$. Here \mathbf{U}_{f+1} is*

obtained by comparing the distance between each vector and median per dimension. We can conclude that none of the vectors in \mathbf{U}_{f+1} contains Byzantine values. The estimator to update the gradient is $\bar{\mathbf{U}} := \frac{1}{f+1}(\sum_{i=1}^{f+1} \hat{\mathbf{V}}_i + \mathbf{g})$.

Proof. Suppose we have a set $S := \{\mathbf{V}_1, \dots, \mathbf{V}_{n-f-c}, \mathbf{B}_1, \dots, \mathbf{B}_f\}$ taking value in \mathbb{R}^d . In the k -th dimension, let $V_1^k, \dots, V_{n-f-c}^k$ and B_1^k, \dots, B_f^k be the reordered gradients such that $|V_1^k - g_k| \leq \dots \leq |V_{n-c}^k - g_k|$ and $|B_1^k - g_k| \leq \dots \leq |B_f^k - g_k|$, where g_k refers to the sample median. Under the assumption of Definition 7 and 8, $|g_k - V_i^{f \text{ ar}}| < |g_k - B_j^{\text{close}}|$, meaning that

$$|V_1^k - g_k| \leq \dots \leq |V_{n-f-c}^k - g_k| \leq |B_1^k - g_k| \leq \dots \leq |B_f^k - g_k| \quad (2)$$

Definition 7(i) limits the f must smaller than $n - f - c$. We defined

$$n > 2f + c \Rightarrow n \geq 2f + c + 1$$

$$\text{then, } f < \lceil \frac{n - c - 1}{2} \rceil$$

We take the largest f in n workers,

$$n = 2f + c + 1$$

$$n - f - c = f + 1$$

Here, $f + 1$ equals to $n - f - c$, we get the first $n - f - c$ correct gradients and append them to the \mathbf{U}_{f+1} . \square

Corollary 2: Under the assumption of Definition 7 and Theorem 1, suppose we have random vectors $\tilde{\mathbf{V}}_i \in \mathbb{R}^d, i \in [1, n - c]$ containing normal gradients and malicious ones. The set of f -neighbor vectors of \mathbf{g} is defined as $\mathbf{U}_{f+1} := \{\{\hat{\mathbf{V}}_i\}_{i=1}^f, \mathbf{g}\}$, where $\{\hat{\mathbf{V}}_i\}_{i=1}^f = \arg \min \sum_{i=1}^f \sum_{k=1}^d |V_{(i)}^k - g_k|$, for different $V_{(i)}^k \in V_j^k, i \in [1, f], j \in [1, n - f - c]$. We can conclude that $V_{(i)}^k$ may contain Byzantine vector elements, but all the vectors in \mathbf{U}_{f+1} still move in the bounded deviation direction to make the training converge. The estimator to update the gradient is $\bar{\mathbf{U}} := \frac{1}{f+1}(\sum_{i=1}^{f+1} \hat{\mathbf{V}}_i + \mathbf{g})$

Proof. Since some Byzantine workers pretend to be correct ones, the median vector \mathbf{g} is important for the gradient direction. As we know the breakdown point of α is $\frac{1}{2}$, we will show the $\mathbb{E}[\mathbf{g}] = \boldsymbol{\mu}$. Because we perform the median selection per coordinate, the proof only needs to show $\mathbb{E}[g_k] = \mu_k$, on the k -th dimension, with $k \in [1, d], k \in \mathbb{N}$.

Considering a set of i.i.d normal variable $\{X_j\}_{j=1}^m$ and $X_j \in \mathbb{R}, \mu$ and σ^2 are the mean and variance respectively. Let $X_{j:n}$ be the j -th order in the a sorted sequence $\{X_{1:m}, \dots, X_{m:m}\} = \{X_j\}_{j=1}^m$. Suppose m is odd, the median of the set is $M := X_{\frac{m-1}{2}+1:m}$, which has the probability density of this order statistics as:

$$f_{med}^{odd}(x) = \frac{m+1}{2} \binom{m}{\frac{m-1}{2}} f(x) (F(x)(1-F(x)))^{\frac{m-1}{2}} \quad (3)$$

where $f(x)$ and $F(x)$ are the density and cumulative distribution functions respectively. Since the normal distribution is symmetric, $F(x) = 1 - F(2\mu - x)$, we can easily get $f_{med}^{odd}(x) = f_{med}^{odd}(2\mu - x)$, namely

$$\mathbb{E}[X_{\frac{m-1}{2}+1:m}] = \mu \quad (4)$$

When m is a even number, the median is $M := \frac{1}{2}(X_{\frac{m-1}{2}+1:m} + X_{\frac{m-1}{2}:m})$. So the joint probably density function is:

$$f_{med}^{even}(x_1, x_2) = \left(\frac{m}{2}\right)^2 \binom{m}{\frac{m}{2}} f(x_1) f(x_2) (F(x_1)(1-F(x_2)))^{m-1} \quad (5)$$

Again, using the symmetric property, we have

$$\mathbb{E}[M] = \mathbb{E}[2\mu - M] \quad (6)$$

which means $\mathbb{E}[M] = \mu$, still. From eq.(4) (6), we know that the sample median of a normal distribution is an unbiased estimator. When Byzantine samples existing, provided that $\alpha < \frac{1}{2}$, we can easily see the sample median as the unbiased estimator still holds.

We then prove the weak Byzantine resilience. Under the assumption of Definition 7(ii) $|g_k - \mu_k| \leq \epsilon$, where ϵ is negligible. The \mathbf{V}_i in a probability distribution is symmetric about the mean, presuming that data near the mean are more frequent in occurrence than data far from the mean. Even some Byzantine workers pretend they are correct among V_i^k , they still have high likelihood values with a given normal distribution, and we say the sample mean of the set \mathbf{U}_{f+1} still in the right direction. \square

D. Complexity Analysis

Computing the median of a list of n of unordered elements takes $\mathcal{O}(n)$ time with standard selection algorithms for a 1-dimensional vector, and the time complexity to get the median which is only $\mathcal{O}(1)$. ParSGD has one more step than Median, which is to find f closest gradients, and the time complexity of this step is $\mathcal{O}(n)$. Usually, there are d -dimensional vectors, so the overall time complexity is $\mathcal{O}(d \cdot n)$.

V. CONVERGENCE ANALYSIS

In this section, we analyze the convergence of the SGD using ParSGD algorithm defined in Section 4. All vectors in \mathbf{U}_{f+1} are in the normal distribution, and all the elements of the vectors are from correct values in strong Byzantine assumption. The absolute value between k -th dimension of $\bar{\mathbf{U}}$ and μ_k is bounded in a small value (η) which will help lead a right direction of gradient descent. Same as the weak Byzantine assumption, some workers may pretend they are correct gradients in \mathbf{U}_{f+1} , but the $f + 1$ gradients in $\tilde{\mathbf{V}}_i$ are pre-selected through f closest neighbors near the median, and we proved median is an unbiased estimator which can tolerate at most $\frac{1}{2}$ Byzantine workers. In Proposition 2, we bounded η and found it is only related to the maximum of covariance σ and f .

Proposition 2. Suppose we have random vectors $\tilde{V}_i \in \mathbb{R}^d, i \in [1, n - c]$ containing normal gradients and malicious ones. We assume that V_i be any independent and i.i.d d -dimensional vector, $V_i \sim \mathcal{N}(\mu, \Sigma)$, and g_k refers to the median of \tilde{V}_i^k . σ is the upper bound of the variance under Definition 2. If $n - c > 2f$ and $|g_k - \mu_k| \leq \epsilon$, let $u_k = \bar{U}^k$ be the k -th dimensional sample mean of U_{f+1} . We defined η by

$$\mathbb{E}(u_k - \mu_k)^2 = \frac{f\sigma^2}{(f+1)^2}, \quad \eta = \frac{\sqrt{f}\sigma}{f+1} \leq \frac{\sigma}{2}$$

Proof. Under strong Byzantine resilience, \tilde{V}_i only composes of V_i . All the computation is under a normal distribution condition. The bounded variance can be written as

$$\mathbb{E}(V_i^k - \mu_k)^2 \leq \sigma^2 \quad (7)$$

$$\mathbb{E}((V_i^k)^2 + \mu_k^2 - 2V_i^k\mu_k) \leq \sigma^2 \Leftrightarrow \mathbb{E}((V_i^k)^2) \leq \sigma^2 + \mu_k^2 \quad (8)$$

Also, $u_k = \bar{U}^k$ and $\mathbb{E}(u_k - \mu_k)^2$ can be calculated as

$$\begin{aligned} u_k &= \frac{1}{f+1} \left(\sum_{i=1}^{f+1} V_i^k + g_k \right) \\ &= \frac{1}{f+1} (\Delta + g_k), \\ \text{where, } \Delta &= \sum_{i=1}^{f+1} V_i^k \end{aligned} \quad (9)$$

$$\begin{aligned} \text{then, } \mu_k^2 + u_k^2 - 2u_k\mu_k &= \frac{1}{(f+1)^2} (\Delta^2 + g_k^2 \\ &+ 2\Delta g_k) + \mu_k^2 + \frac{2}{f+1} (\Delta + g_k)\mu_k \end{aligned} \quad (10)$$

$$\begin{aligned} \mathbb{E}(u_k - \mu_k)^2 &\leq \mathbb{E} \left(\frac{1}{(f+1)^2} (\Delta^2 + g_k^2 + 2\Delta g_k) \right. \\ &\quad \left. + \mu_k^2 + \frac{2}{f+1} (\Delta + g_k)\mu_k \right) \end{aligned} \quad (11)$$

As defined condition (ii) in Definition 7, $|g_k - \mu_k| < \epsilon$, then

$$\mathbb{E}(g_k - \mu_k)^2 = \mathbb{E}[g_k^2] - \mu_k^2 \leq \epsilon^2 \Leftrightarrow \mathbb{E}[g_k^2] \leq \epsilon^2 + \mu_k^2 \quad (12)$$

$$\mathbb{E}(\Delta) = f\mu_k, \quad \mathbb{E}(\Delta)^2 = f \cdot (\sigma^2 + \mu_k^2) + \frac{f \cdot (f-1)}{2} \mu_k^2 \quad (13)$$

We finally bounded the u_k to mean combined with formula 8, 11, 12 and 13,

$$\begin{aligned} \mathbb{E}(u_k - \mu_k)^2 &\leq \frac{1}{(f+1)^2} (f \cdot (\sigma^2 + \mu_k^2) + f \cdot (f-1) \mu_k^2 \\ &\quad + \mu_k^2 + \epsilon^2 + 2f\mu_k^2) - \mu_k^2 \\ &\leq \frac{f\sigma^2 + \epsilon^2}{(f+1)^2} \Leftrightarrow \frac{f\sigma^2}{(f+1)^2} \end{aligned}$$

then standard deviation of $\mathbb{E}(u_k - \mu_k)^2$ equals to $\frac{\sqrt{f}\sigma}{f+1}$ and η is bounded by σ and f in no more than $1/2\sigma$ when $f = 1$, and the more workers we have the smaller η is, then we say $\eta \leq 1/2\sigma$.

For the Weak Byzantine resilience, \tilde{V}_i^k may contain V_i^k and B_i^k . However, B_i^k pretends to be V_i^k , which means $\mathbb{E}(\tilde{V}_i^k -$

$\mu_k)^2 \leq \sigma^2$. Then the u_k is updated as $u_k = \frac{1}{f+1} (\sum_{i=1}^{f+1} \tilde{V}_i^k + g_k) = \frac{1}{f+1} (\Delta + g_k)$, where $\Delta = \sum_{i=1}^{f+1} \tilde{V}_i^k$. Since B_i^k pretends to be correct, we assume $\mathbb{E}B_i^k = \mu_k$. $\mathbb{E}(\Delta)$ and $\mathbb{E}(\Delta)^2$ remain the same, which means the overall result remains the same. Now that the proposed gradient estimator \bar{U} is bounded in a small deviation from the ground truth, the proposed algorithm will converge as the regular gradient descent with mean as the estimator. So we have proved ParSGD algorithm is Byzantine resilience and the training will converge even it is a weak Byzantine resilience. \square

VI. EXPERIMENTAL EVALUATION

We implemented and evaluated ParSGD in a simulated mode on a Tesla P100 Nvidia setting. Our algorithm is evaluated by 3 common attacks, Crash-stop [30], Bit-flipping [29] and Random Gaussian attack. We summarize the results of our experiment as follows:

- ParSGD does not wait for all workers' updates because it collects gradients in partial synchronous mode. It reaches the best training accuracy under the Crash-stop failure compared with Mean, Median, and Krum.
- Under the attack of Bit-flipping, ParSGD gets the best Top-1 accuracy among Mean, Median, and Krum in MNIST. When the Byzantine workers f reach 22 (44% Byzantine workers), some fluctuations occur in Top-1 accuracy affected by the median, but the Loss still converges best. Krum and Mean cannot converge both in CIFAR10 and MNIST when Byzantine workers f reach 20, as shown in Figure 2 and Figure 4.
- With Random Gaussian attack, compared with Krum which has the inherent advantage of resistance to random Gaussian attack when Byzantine workers $f < \frac{n}{2}$, ParSGD still achieves the best training accuracy and Loss benchmarking with both MNIST and CIFAR10.

A. Overviews

1) *Datasets:* We conduct experiments on benchmark MNIST [21], handwritten digits for image processing, which contains 60,000 training images and 10,000 testing images. We also conduct experiments on benchmark CIFAR10 [20], which consists of 60,000 32x32 colour images in 10 classes, with 6,000 images per class.

2) *Evaluation Settings:* During the evaluation, we execute 200 epochs with 50 workers. We set the learning rate to 0.05 and the batch size to 100. In each epoch, we use cross-entropy loss function with Top-1 and Top-5 accuracy evaluation metrics. For the distributed ML model, we run a multi-layer convolutional network, which has four 3x3 convolution layers (the first two layers with 64 channels, the last two layers with 128, each followed with a 2x2 max pooling), and a fully connected layer with 128 units and ReLU activation, with a last output layer.

We normalize all data and divide the data into two classes: training data (eighty of all data), and testing data (twenty of all data).

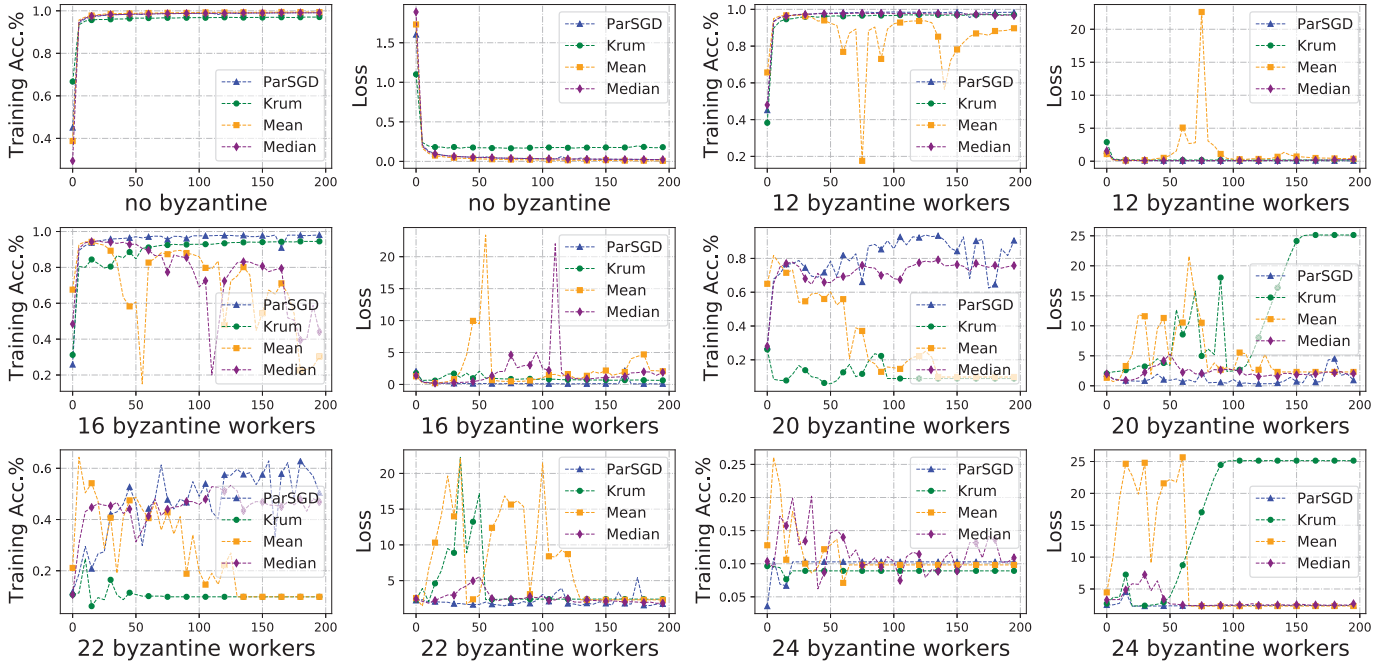


Fig. 2: MNIST: Top-1 Accuracy and Loss under Bit-Flip Attack




Fig. 3: MNIST: Top-1 Accuracy and Loss under Random Gaussian Attack

B. Crash-stop failure

Crash-stop failure usually happens at a sudden stop in an emergent situation, and it is a type of failure that causes the component of a system to stop operating. We tested ParSGD under 5 crashed workers and 22 Byzantine workers. The result shows that only ParSGD can tolerate Crash-stop failure among Median, Mean and Krum, because ParSGD is in partial synchronous mode which only collects $n - c$ vectors from active workers.

C. Bit-flipping Attack

The formal definition of a Bit-flipping attack is an attack in which the attacker can change the ciphertext to result in a predictable change of the plaintext. In our experiment, we simulate the Bit-flipping attack by changing f vectors to opposite descent directions. More specifically, the adversary p_i first calculates each true gradient vector Δw_j , and then updates $-c_i \Delta w_j$ to the server, where c_i is a random constant (or one for simple). The result is shown in Figure 2.

We evaluate the number of Byzantine workers f from 12 to 24. We found Mean cannot tolerate even one Byzantine worker, while Median and Krum cannot converge at nearly

18 Byzantine workers. Our ParSGD can still converge when Byzantine number f reaches up to 22, see the training accuracy of 22 Byzantines in Figure 2. The accuracy of ParSGD is influenced by the median (median may pretend correctly), but the Loss still converges best even Byzantine number f reaches up to 24, see the Loss of 24 Byzantines in Figure 2.

D. Random Gaussian Attack

We simulate using random Gaussian attack to compromise the above 4 aggregation rules. We take the standard deviation to 0.1, 1, and 200, and take the mean to $-1e8$ or 0. Our experimental results prove that as long as *the majority of the workers are correct*, the accuracy of ParSGD can always be stable. We found Krum is the best solution here because of its inherent advantage of finding the minimum vector among $n - f$ vectors in strong Byzantine resilience (Definition 8), as shown in Figure 3. However, in weak Byzantine resilience (Definition 7), if all the Byzantine workers pretend they are correct, the training accuracy of Krum will be increasingly threatened. ParSGD still works well when there are 24 Byzantine workers and both of their Top-5 accuracies can reach up to 99%. Due

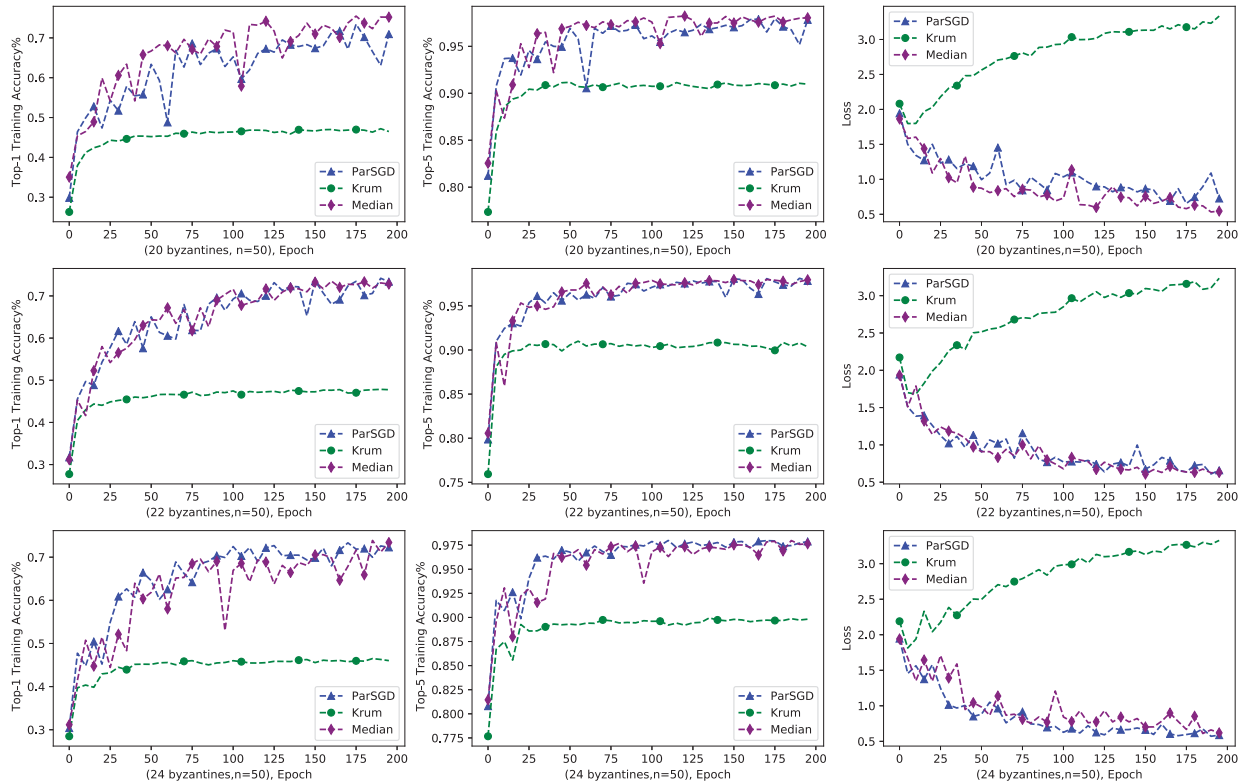


Fig. 4: CIFAR10: Top-1, Top-5 Accuracy and Loss under Random Gaussian Attack.

to space limitations, we only show the Top-1 results here since the Top-5 results have roughly the same conclusion.

We also benchmark ParSGD with CIFAR10 under the random Gaussian attack, as shown in Figure 4. We only list ParSGD, Krum, and Median here, as we have proved that Mean can not tolerate even one Byzantine worker. In the meantime, all the algorithms can tolerate less 30% Byzantine workers except Mean, so we only show our results for 20, 22, and 24 Byzantine workers here ($n = 50$ and $c = 0$). We evaluate both the Top-1 and Top-5 accuracy in this CIFAR10 evaluation.

We compromised f d -dimensional vectors by replacing f wrong vectors. These vectors are in normal distribution by setting mean to 0 and variance to 200, pretending all the generated gradients are correct. The results show that the accuracy of Median and ParSGD are almost the same in Top-1 and Top-5, but ParSGD has relatively fewer fluctuations when the Byzantine number reaches up to 24. In addition, ParSGD converges best as shown in Figure 4, see the Loss of 24 Byzantines.

In the nutshell, consider all evaluations, we found only Median and ParSGD get converged correctly. As we said, Krum can hardly converge even Byzantine number f less than correct workers in weak Byzantine resilience.

E. Discussion and Limitation

During the evaluation, we find that the training accuracy of ParSGD is easily affected by the median because our algorithm

is finding f nearest neighbors close to the median. When there are more than 45% Byzantine workers under Bit-flipping attack and random Gaussian, the median can easily be altered to a wrong one. Namely, if the median is pretty close to the mean in an unknown distribution, we can achieve good training accuracy. But in real-world training settings, the vectors are not completely in a normal distribution, that's why the accuracy starts to fluctuate when Byzantine workers near 50%. For example, under the Gaussian attack when the variance is relatively normal, f Byzantine workers will pretend they are correct, so a wrong gradient of a Byzantine node may be selected to be a new median. The result of collecting f gradients near median may count a lot of Byzantine workers, but because they pretend to be normal, the training finally gets converge, however, it just has some fluctuations during training. Compared with Krum and Median, our algorithm ParSGD performs a good result when there are more than $\frac{1}{3}$ Byzantine workers. We demonstrate our algorithm, ParSGD, is practical and meaningful in an untrusted distributed environment.

VII. CONCLUSION

We propose a novel gradient aggregation rule, ParSGD, in partial synchronous mode, which can tolerate Crash-stop failures and nearly 50% of Byzantine workers in distributed and collaborative training. Compared with Mean, Median, and Krum, we get the best accuracy under three common attacks. The algorithm has a provable convergence analysis in both

strong and weak assumptions. We will try to apply this solution to federated learning settings in the future.

ACKNOWLEDGEMENTS

We gratefully acknowledge the support of the NSF through grant IIP-1919159. We also acknowledge the support of the IBM research team.

REFERENCES

- [1] https://en.wikipedia.org/wiki/gaussian_noise.
- [2] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16)*, pages 265–283, 2016.
- [3] D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic. Qsgd: Communication-efficient sgd via gradient quantization and encoding. *Advances in Neural Information Processing Systems*, 30:1709–1720, 2017.
- [4] Baidu Research. Tensorflow allreduce and Baidu allreduce. <https://github.com/baidu-research>, 2018. Accessed: 2017-10-23.
- [5] A. Bessani, J. Sousa, and E. E. Alchieri. State machine replication for the masses with bft-smart. In *2014 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, pages 355–362. IEEE, 2014.
- [6] P. Blanchard, E. M. El Mhamdi, R. Guerraoui, and J. Stainer. Machine learning with adversaries: Byzantine tolerant gradient descent. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 118–128, 2017.
- [7] L. Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, pages 177–186. Springer, 2010.
- [8] M. Castro and B. Liskov. Practical byzantine fault tolerance and proactive recovery. *ACM Transactions on Computer Systems (TOCS)*, 20(4):398–461, 2002.
- [9] L. Chen, H. Wang, Z. Charles, and D. Papailiopoulos. Draco: Byzantine-resilient distributed training via redundant gradients. In *International Conference on Machine Learning*, pages 903–912. PMLR, 2018.
- [10] G. Damaskinos, E. M. El Mhamdi, R. Guerraoui, A. H. A. Guirguis, and S. L. A. Rouault. Aggregathor: Byzantine machine learning via robust gradient aggregation. In *The Conference on Systems and Machine Learning (SysML), 2019*, number CONF, 2019.
- [11] G. Damaskinos, R. Guerraoui, R. Patra, M. Taziki, et al. Asynchronous byzantine machine learning (the case of sgd). In *International Conference on Machine Learning*, pages 1145–1154. PMLR, 2018.
- [12] J. Dean and S. Ghemawat. Mapreduce: simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.
- [13] S. Duan, C. Liu, X. Wang, Y. Wu, S. Xu, Y. Yesha, and H. Zhang. Intrusion-tolerant and confidentiality-preserving publish/subscribe messaging. In *2020 International Symposium on Reliable Distributed Systems (SRDS)*, pages 319–328. IEEE, 2020.
- [14] C. Dwork, N. Lynch, and L. Stockmeyer. Consensus in the presence of partial synchrony. *Journal of the ACM (JACM)*, 35(2):288–323, 1988.
- [15] E.-M. El-Mhamdi, R. Guerraoui, A. Guirguis, L. N. Hoang, and S. Rouault. Genuinely distributed byzantine machine learning. In *Proceedings of the 39th Symposium on Principles of Distributed Computing*, pages 355–364, 2020.
- [16] M. Fang, X. Cao, J. Jia, and N. Gong. Local model poisoning attacks to byzantine-robust federated learning. In *29th {USENIX} Security Symposium ({USENIX} Security 20)*, pages 1605–1622, 2020.
- [17] C. J. Geyer. Breakdown point theory notes. *Class Notes on Nonparametric Statistics*, 2006.
- [18] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, et al. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*, 2019.
- [19] V. Kalavri, V. Vlassov, and S. Haridi. High-level programming abstractions for distributed graph processing. *IEEE Transactions on Knowledge and Data Engineering*, 30(2):305–324, 2017.
- [20] A. Krizhevsky, G. Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [21] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [22] M. Li, L. Zhou, Z. Yang, A. Li, F. Xia, D. G. Andersen, and A. Smola. Parameter server for distributed machine learning. In *Big Learning NIPS Workshop*, volume 6, page 2, 2013.
- [23] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017.
- [24] E. M. E. Mhamdi, R. Guerraoui, and S. Rouault. The hidden vulnerability of distributed learning in byzantium. *arXiv preprint arXiv:1802.07927*, 2018.
- [25] A. Miller, Y. Xia, K. Croman, E. Shi, and D. Song. The honey badger of bft protocols. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 31–42, 2016.
- [26] F. Niu, B. Recht, C. Ré, and S. J. Wright. Hogwild!: A lock-free approach to parallelizing stochastic gradient descent. *arXiv preprint arXiv:1106.5730*, 2011.
- [27] D. Ongaro and J. Ousterhout. In search of an understandable consensus algorithm. In *2014 {USENIX} Annual Technical Conference ({USENIX}{ATC} 14)*, pages 305–319, 2014.
- [28] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. 2017.
- [29] A. S. Rakin, Z. He, and D. Fan. Bit-flip attack: Crushing neural network with progressive bit search. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1211–1220, 2019.
- [30] F. B. Schneider. Byzantine generals in action: Implementing fail-stop processors. *ACM Transactions on Computer Systems (TOCS)*, 2(2):145–154, 1984.
- [31] A. Sergeev and M. Del Balso. Horovod: fast and easy distributed deep learning in tensorflow. *arXiv preprint arXiv:1802.05799*, 2018.
- [32] J.-y. Sohn, D.-J. Han, B. Choi, and J. Moon. Election coding for distributed learning: Protecting signsgd against byzantine attacks. *Advances in Neural Information Processing Systems*, 33, 2020.
- [33] C. Xie, O. Koyejo, and I. Gupta. Generalized byzantine-tolerant sgd. *arXiv preprint arXiv:1802.10116*, 2018.
- [34] C. Xie, S. Koyejo, and I. Gupta. Zeno: Distributed stochastic gradient descent with suspicion-based fault-tolerance. In *International Conference on Machine Learning*, pages 6893–6901. PMLR, 2019.
- [35] C. Xie, S. Koyejo, and I. Gupta. Zeno++: Robust fully asynchronous sgd. In *International Conference on Machine Learning*, pages 10495–10503. PMLR, 2020.
- [36] D. Yin, Y. Chen, R. Kannan, and P. Bartlett. Byzantine-robust distributed learning: Towards optimal statistical rates. In *International Conference on Machine Learning*, pages 5650–5659. PMLR, 2018.
- [37] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, I. Stoica, et al. Spark: Cluster computing with working sets. *HotCloud*, 10(10-10):95, 2010.
- [38] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.