P2PRISM - PEER TO PEER LEARNING WITH INDIVIDUAL PRISM FOR MODEL AGGREGATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Federated learning (FL) has made collaboration between nodes possible without explicit sharing of local data. However, it requires the participating nodes to trust the server and its model updates, the server itself being a critical node susceptible to failure and compromise. A loss of trust in the server and a demand to aggregate the model independently for oneself has led decentralized peer-to-peer learning (P2PL) to gain traction lately. In this paper, we highlight the never before exposed vulnerabilities of P2PL towards malicious attacks and show the differences between the behavior of P2PL and FL in a malicious environment. We then present a robust defense - P2PRISM as a secure aggregation protocol for P2PL.

1 INTRODUCTION

1.1 MOTIVATION FOR PEER-TO-PEER LEARNING

FL McMahan et al. (2017); Konečnỳ et al. (2016) has demonstrated how clients can benefit from collaboration by sharing their local gradient updates to the parameter server, which in turn aggregates Yin et al. (2018a); Blanchard et al. (2017); Guerraoui et al. (2018); Xia et al. (2019); Fung et al. (2020); Cao et al. (2020) all the received gradients to update the global model. For the sake of simplicity, we assume that the entire process is synchronous - the server waits to hear from all the clients before aggregation and all clients receive the same global model from the server after aggregation, that is, the clients always have to agree on the global model sent by the server and replace its local model with it before continuing with the local training. Although the aggregation technique being used may be known to all, but the actual aggregation is hidden from the clients for privacy concerns as it has been shown that access to a client's gradients can be used to recover its local data in an approximate or an exact way by optimization Geiping et al. (2020) or analytical Fowl et al. (2021) methods respectively. It is therefore not possible for clients to selectively choose other clients' gradients to aggregate even if it benefits them from any existing spatial locality among the clients. The clients have to trust the server to also aggregate the gradients in a byzantine-robust manner. Unless the server itself possesses a root dataset Cao et al. (2021) that correctly represents the entirety of data possessed by all clients as the ground truth, it is difficult for it to correctly identify malicious updates statistically without being extremely conservative and removing any suspected gradients leading to a significant loss of information. Whereas a node does have access to its own generated gradients as the benign ground truth and can make use of it, given the power to aggregate the model for itself. Due to the above mentioned reasons, and several others, a node is motivated to lose trust in a server and join a decentralized collaboration among the other nodes.

1.2 COMPARISON WITH FEDERATED LEARNING

In FL, the server aggregates the gradients from all the clients. However, in P2PL, a node may choose to communicate only with its neighbors in every round and locally aggregate the received models. If the graph formed by the nodes is not fully connected with equally weighed edges, the nodes are going to have models that differ from each other at every point in time even after their local aggregation. Consensus distance (δ) of the graph is defined as the average distance of each of the m local models (x_i) from the centroid ($\bar{x}$) of them all, known to an oracle.

$$\delta := \frac{1}{n} \sum_{i=0}^{i=n-1} \|x_i - \bar{x}\| \tag{1}$$

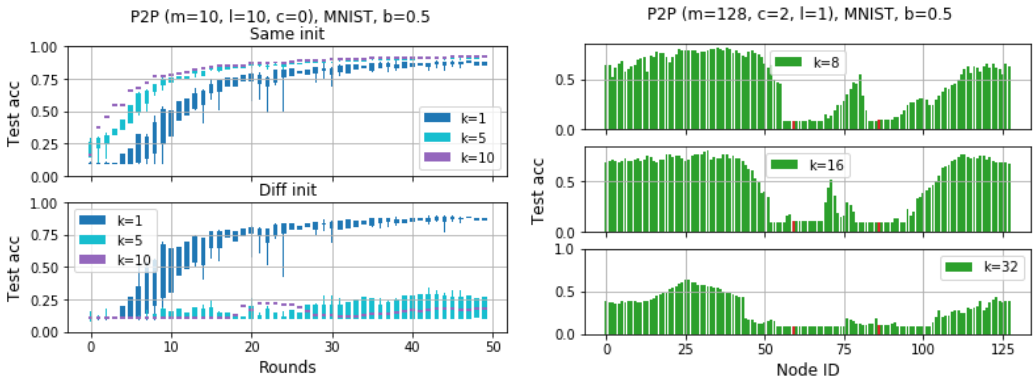

Figure 1: *Left*: Training a k-regular peer-to-peer graph with varying values of k with same and different model initialization. We see that it is necessary to initialize nodes with the same model to expect benefit from collaboration ($k > 1$). *Right*: Demonstrating the spread of the attack in a k-regular graph where only 2 out of 128 nodes (shown in red) are malicious. We can see that the larger the collaboration (k), the greater the spread of the attack if the aggregation used is insecure, such as gossip averaging in this case.

It is measured after every aggregation step. Needless to say this takes a zero value in FL and non-zero value in P2PL when the graph is not fully connected. A non-zero consensus value implies that sharing gradients as in FL is not appropriate in P2PL and the nodes need to share the actual model weights with their neighbors because the local gradients have been computed on already differing local models. The nodes, after performing local SGD on their own data, where the nodes' data can be assumed to be non-IID with bias $= b$ to incentivize collaboration, tend to diverge from each other. In FL, the server enforces full consensus among the clients, whereas in P2PL, the nodes perform one or more rounds of gossip averaging (GA) with their neighbors to keep the consensus distance in check. It is always wise to initialize all nodes in P2PL with the same model, that is, zero initial consensus distance to aid consensus control in the later stages of training, as also demonstrated in the toy experiment in Figure 1. We see that with same initialization, the nodes benefit from collaboration as k increases to 5 and 10 from 1 (individual training). This is also why every client in FL benefits from collaboration as their models are made to synchronize by the server. Whereas, when nodes are initialized differently, they hurt each other with collaboration, because they are all on different trajectories towards learning an optimal model, as the same ML problem may have multiple solutions depending on the initialization. In fact, we observe that the nodes are hurt even with $k = m = 10$ when after a single round of gossip, the consensus distance falls down to 0 and stays at 0. Hence, keeping the pre-gossip consensus distance low is as important as keeping the post-gossip consensus distance, which is why we recommend setting the initial pre-gossip consensus distance to 0 by initializing all the nodes with the same model.

Assuming a weighted mean aggregation, in FL, the server assigns weights to all the clients. In P2PL, every client is responsible for its own aggregation, and maintains a vector of weights (zero or non-zero) it assigns to every other node, including itself. These vectors when stacked on top of each other form the mixing matrix W of the graph and define its topology. If x is a matrix with x_i being its ith row, then the gossip averaging step is captured by the matrix operation Wx. It is to be noted that a node is only affected by its direct neighbors in one round of gossip averaging, but can be affected by an indirectly connected neighbor if multiple rounds are performed. For example two rounds of GA with matrix W results in the models $W \cdot Wx$ which is effectively a single round of GA with matrix W^2 where W^2 is less sparse than W for positive entries in W. It should also be obvious that more the number of gossip averaging steps, the tighter the consensus control Kong et al. (2021).

1.3 P2PL IN A BYZANTINE ENVIRONMENT

P2PL behaves differently from FL under a malicious setting. Assuming the same model poisoning attack in both the cases, with c out of m nodes being malicious or compromised, while the FL server has to deal with $f(= \frac{c}{m})$ fraction of malicious nodes, the nodes in P2PL have to deal with

a variable fraction of malicious neighbors which could be greater than f for many, depending on the distribution and connectivity of the nodes. Attackers can not only impact their direct neighbors, but the impact spreads like a disease with either multiple GA steps in the same round, or even with a single GA step when other nodes come in contact with the poisoned nodes in the next round before the infected could recover themselves during their local SGD. In this manner, the attacker also succeeds in magnifying the consensus distance among the nodes that makes training even more difficult. Depending on the graph topology and distribution of the attackers, although an attacker may not be able to impact every node as in FL, the impact on affected nodes can be higher because some nodes are bound to have $\geq f$ fraction of their neighbors malicious. Figure 1 demonstrates this phenomenon in P2PL in comparison with FL.

1.4 THREAT MODEL

We have extended the state-of-the-art model poisoning attack, SHEJWALKAR attack Shejwalkar & Houmansadr (2021) in FL to the P2PL case. The malicious nodes collaborate among themselves to access the complete knowledge of the current as well as the past state of the benign models. A perturbation unit vector is constructed along the direction in which the average benign model is moving, and would be scaled up and subtracted from the past averaged model to send the nodes along the direction of gradient ascent. The scaling factor is chosen by solving an optimizing function that balances between attack impact and stealth, given that we know the aggregation technique being used. The attack thus amplifies the distance between an infected and an uninfected node, leading them to disagreement and thereby breaking consensus among the nodes.

1.5 OUR DEFENSE - P2PRISM

Our defense is based on the key intuition that the direction of the received models should be of utmost importance during aggregation. This is difficult in FL as there is no way for the server to know which clients are benign to use as reference to compute the direction of the updates. Naturally, none of the aggregation techniques that are expected to be byzantine-robust in FL make use of this fact. In P2PL, a node can use its own local update as the reference to detect parameter-wise direction flip in the updates received from its neighbors. It can then filter out suspected updates based on a defence policy that we call as $PRISM$ before aggregating the received updates. We describe this policy in much detail in the next section, and we use it to completely undo the damage done by the attack, and reviving upto 100% of the benign nodes in a graybox attack and upto 88% of them in a k-regular graph in the presence of a whitebox attack.

In summary, our contributions are -

- We demonstrate the effect of byzantine nodes in P2PL and propose a defense against the state-of-the-art attack by leveraging the ground truth information available in P2PL as opposed to FL, reviving upto 100% of the nodes from the effect of the attack.
- We explore the behavior of P2PL across k-regular and power-law graph topologies and highlight the subtleties involved in training like model initialization that affect consensus among the nodes. We also empirically show how P2PRISM keeps the consensus distance under control.
- We evaluate P2PRISM on the image datasets of MNIST and FMNIST, and the NLP dataset of Shakespeare under graybox as well as fully whitebox adaptive attacker capabilities to prove effectiveness of our defense principle under the ultimate stress-test environment, where upto 88% of the nodes in a k-regular graph and 66% nodes in a power-law graph were completely revived even when 1/4-th of the nodes were under the control of a whitebox attacker.

2 DESIGN

P2PRISM is instantiated at every node to perform secure aggregation. It provides the node with a prism that blocks the passage of malicious updates into the aggregated model. P2PRISM uses a metric called flip-score Sharma et al. (2021) (FS) to detect malicious updates and helps a node keep a record of the reputation score of all its neighbors. The flip-score is expected to capture any large deflection as compared to a trusted benign update. A node i with model $x_{i,t}$ sets $\nabla x_{ii} = x_{i,t} - x_{i,t-1}$

as the trusted reference. All models that is receives from its neighbors j are adjusted according to this reference: $\nabla x_{ij} = x_{j,t} - x_{i,t-1}$. The flip-score of j as computed by i is

$$FS_{ij} = \sum_{k=0}^{|P|-1} (|\nabla x_{ij}[k]|)^2 \times (sign(\nabla x_{ij}[k]) \neq (sign(\nabla x_{ii}[k])). \tag{2}$$

where $|P|$ is the total number of parameters (weights and biases) in the model. Since a node can trust its own update ∇x_{ii} to go towards gradient descent, a high flip-score naturally captures an update that is likely going towards a gradient ascent - with either a large number of parameters going in the opposite direction with a small magnitude or a small number of them with a large magnitude. Cosine similarity is also based on a similar intuition but it does not capture a subtle attack that targets model parameters with zero gradients to prevent them from converging. With this, we can say that an abnormally large FS is indicative of an attack but it is not enough by itself without a defense policy.

In FL, byzantine-robust aggregation technqiues such as Trimmed Mean and Median Yin et al. (2018b), Krum Blanchard et al. (2017), Bulyan Mhamdi et al. (2018), FABA Xia et al. (2019), FLTrust Cao et al. (2021) and others assume a constant upper limit on the number of malicious nodes ($c_{max} > \frac{m}{2}$). That may work when the attackers are not in a majority, by trimming out a fixed number of updates in every iteration and penalizing the reputation of the respective clients. The attackers may, however form a local majority of variable size at different regions in P2PL, and complete security cannot be maintained with a constant c_{max}. We describe how P2PRISM handles this in the following steps.

1. **Finding the cutoff FS** - In P2PRISM, every node finds the minimum (FS_{min}) and median FS (FS_{med}) for all the updates it has received in a given iteration. It then sets a cutoff FS at $FS_{med} + \mu(FS_{med} - FS_{min})$, where $\mu \geq 0$ is the defense parameter. A lower μ sets a conservative policy whereas a large value sets a lenient one. For a robust computation, it is suggested to use a lower percentile instead of the absolute minimum, when the number of neighbors is sufficiently large (≥ 100 for example).
2. **Reward or penalize neighbors** - All neighbors with FS less than the cutoff are given a reward of 1 unit. All others are penalized by the amount $max(\frac{FS - FS_{med}}{FS_{med} - FS_{min}}, 5)$. The penalty is upper capped so that a node is allowed to redeem itself if it acts benign in the future.
3. **Update reputation** - The reward or penalty is added to the previous reputation score that was initialized to $\frac{1}{k}$ for a neighbor of a node that has $k-1$ neighbors. In this way, we make use of the past behavior of every neighbor. The reputation scores are normalized by dividing each value with the total positive reputation score of all neighboring nodes. All neighbors with a negative reputation are filtered out by this defense policy. .
4. **Aggregation** - The remaining updates, including a node's self update, undergo mean aggregation. A weighted mean weighed by the reputation is intentionally avoided to prevent the effect of any existing bias. For example. Nodes with similar updates could assign each other a high weight and isolate themselves from the rest of the graph, and not benefit from collaboration to learn a generalized model. This would also result in a higher consensus distance among the nodes.

In order to protect a node surrounded by malicious majority, we also cap its FS_{med} to be $\leq 10 * FS_{min}$. Without this, the median value could be poisoned leading to a victory of the attackers. The pseudocode of P2PRISM is described in Algorithm 1. P2PRISM is instantiated at every node in the peer-to-peer graph. It is to be noted that we use FS only to trim out the anomalies, and the rest of the the aggregation follows simple gossip averaging, hence the convergence analysis follows from already existing work Assran et al. (2019), assuming that malicious updates are correctly detected and penalized based on FS, which we empirically verify in our experiments. We do not assume a doubly stochastic mixing matrix as it requires every column sum to be equal to 1, but for ideal security want all entries in the column for a malicious node to be 0. Hence, we only assume it to be column stochastic. The rest of the assumptions are standard as follows -

Assumption #1: x_j are all L-smooth, that is, for all v and w, $x_j(v) \leq x_j(w) + (v-w)^T \nabla x_j(w) + \frac{L}{2}\|v-w\|_2^2$.

Algorithm 1 Peer-to-peer learning with P2PRISM

Input: Node i, current model $x_i(t, \cdot)$,
Neighbor updates $x_j(t+1, \cdot)$, neighbors' reputation $W_R(t, j)$
Output: Aggregated model $x_i(t+1, \cdot)$
Parameters: μ

0 : Compute reference vector $\nabla x_{ij} = x_i(t+1, \cdot) - x_i(t, \cdot)$
1 : for every neighbor j **compute flip-score**:
$FS_{ij} = \sum_{k=0}^{|P|-1} (|\nabla x_{ij}[k]|)^2 \times (sign(\nabla x_{ij}[k]) \neq (sign(\nabla x_{ii}[k]))$
2 : Compute cutoff flipscore:
$FS_{min,i} = min_j\{FS_{ij}\}$
$FS_{med,i} = min(median_j\{FS_{ij}\}, 10 \times FS_{min,i})$
$FS_{i,cut} = FS_{med,i} + \mu \times (FS_{med,i} - FS_{min,i})$
3 : Penalize neighbors with high FS:
$W_R(t+1, j) = W_R(t, j) - max\{\frac{FS_{ij} - FS_{med,i}}{FS_{med,i} - FS_{min,i}}, 5\}$
4 : Reward the rest of the neighbors:
$W_R(t+1, j) = W_R(t, j) + 1$
5 : Normalize reputation weights: $W_R(t+1, j) = \frac{W_R(t+1,j)}{\sum_{j:W_R(t+1,j)>0} W_R(t+1,j)}$
6 : Aggregate gradients: $x_i(t+1, \cdot) = x_i(t, \cdot) + avg_{j:W_R(t+1,j)>0}(x_j(t+1, \cdot) - x_i(t, \cdot))$

Assumption #2: x_j are all μ-strongly convex, that is, for all v and w, $x_j(v) \geq x_j(w) + (v - w)^T \nabla x_j(w) + \frac{\mu}{2}\|v - w\|_2^2$.

Assumption #3: Let ξ_t^j be sampled uniformly at random from the local data of the $j-th$ client, then the variance of stochastic gradients of each client is bounded, that is, $\mathbb{E}\|\nabla x_j(x_{j,t}, \xi_t^j) - \nabla x_{j,t}(x_{j,t})\|^2 \leq \sigma_j^2$ for $j = 1, 2, ...k_i$.

Assumption #4: The expected squared norm of stochastic gradients is uniformly bounded, that is, $\mathbb{E}\|\nabla x_j((x_{j,t}, \xi_t^j)\|^2 \leq G^2$ for $j = 1, ...k_i$, and $t = 0, ..T-1$.

Assumption #5: To each mixing matrix $\mathbf{W}^{(k)}$, we can associate a graph with vertex set $\{1, \ldots, m\}$ and edge set $E^{(k)} = \{(i, j) : w_{i.j}^{(k)} > 0\}$. We assume that there exists finite, positive integers, B and Δ, such that the graph with edge set $\bigcup_{k=lB}^{(l+1)B-1} E^{(k)}$ is strongly connected and has diameter at most Δ for every $l \geq 0$.

In our work, we use the flip score to estimate the mixing matrix, which remains column stochastic for every iteration. Following the Theorem from (Assran et al., 2019), as Assumption 1-5 hold, we have the following convergence bound:

$$\frac{\sum_{k=0}^{T-1} \mathbb{E}\|\nabla f(\bar{x}^k)\|^2}{T} \leq \frac{12(f(\bar{x}^0) - f^* + \frac{L\sigma^2}{2})}{\sqrt{nT}},$$

where T is the total number of iterations, which satisfies $T \geq \max\{\frac{nL^4C^4 60^2}{(1-q)^4}, \frac{L^4C^4P_1^2 n}{(1-q)^4(f(\bar{x}^0) - f^* + \frac{L\sigma^2}{2}))^2}, \frac{L^2C^2P_2 n}{(1-q)^2(f(\bar{x}^0) - f^* + \frac{L\sigma^2}{2}))}\}$. And $P_1 = 4(\sigma^2 + 3G^2)n + \frac{\sum_{i=1}^n \|x_i^{(0)}\|^2}{n}$ and $P_2 = \sigma^2 + 3G^2L^2C^2 + 2\frac{\sum_{i=1}^n \|x_i^{(0)}\|^2}{n}$.

3 IMPLEMENTATION

We have performed our experiments on two graph topologies - k-regular and power-law graphs. Every node in the k-regular graph was made to communicate with $k-1$ other neighbors in order to do a k-gossip averaging in every round. In order to make the *in_degree* (the number of nodes from which updates are received, including the self node) equal to the *out_degree* (the number of neighbors to which one's updates are sent), the nodes were simulated to be situated around a ring, and a symmetric communication was established between a node with $\lceil \frac{k-1}{2} \rceil$ nodes ahead of it and $\lfloor \frac{k-1}{2} \rfloor$ nodes behind it. In a power-law graph, some nodes are popular with many connections, while

Table 1: Training parameters

Dataset	MNIST	FMNIST	Shakespeare
Model	DNN	DNN	LSTM
batch size	32	32	32
learning rate	0.01	0.01	0.1
n_rounds	500	500	1000

a large number of nodes have only few connections. It is difficult to simulate a symmetric graph where both the *in_degree* and *out_degree* strictly follow the power law, hence an asymmetric graph was constructed with only the *in_degree* of nodes adhering strictly to the power-law in order to test the effectiveness of P2PRISM aggregation. To simulate this, the minimum degree was chosen as user input k while the maximum *in_degree* (deg_0) equals $m/2$, where m is the total number of nodes in the network. 1 random node would be simulated with *in_degree* = deg_0, 2 with deg_1, 4 with deg_2, and so on, where $k = deg_{log_2 m} < deg_{log_2 m-1} < ... < deg_0 = m/2$ are in a geometric progression. There would be 1 another node with *in_degree* = m to ensure that the *in_degree* distribution is in adherence with the power-law and all nodes have been assigned some *in_degree* value from the available degree values.

c out of m nodes were randomly chosen to be malicious with their models crafted with the SHEJWALKAR attack with *unit vector* perturbation in every iteration. The source code was obtained from `https://github.com/vrt1shjwlkr/NDSS21-Model-Poisoning` and the same attack parameter ($\lambda_0 = 10.0$) was used. The gradient descent direction was estimated by accessing the current and past average of the models from the benign nodes. The attack assumes a mean-like aggregation being used to optimize its λ value.

Every node was made to perform one iteration of local training ($l = 1$) on its minibatch, after which one round of gossip was performed. We use the image datasets of MNIST and Fashion-MNIST Xiao et al. (2017), and the NLP dataset of Shakespeare Caldas et al. (2019) for our experiments, trained with a DNN and an LSTM respectively. We report the test accuracy for the image dataset and the test perplexity ($= 2^{\text{test loss}}$) for the NLP dataset as used in the standard practice. The training hyperparameters are described in Table 1.

The DNN used to train MNIST and FMNIST consists of two CNN layers with 30 and 50 channels respectively, comprising of 3x3 kernels, the two layers being separated by a 2x2 maxpool layer, and followed by two fully connected layers of 200 and 10 neurons. The LSTM used to train Shakespeare consists of 1 hidden layer with 128 neurons between the input and the output layers. For the defense, a default value of $\mu = 0.75$ was used against the graybox attack and $\mu = 0$ against the whitebox attack. Due to the small value of k used in simulations, we use the 0-percentile, that is the minimum flip-score value to set the upper bound of permissible flip-score. For comparison, the benign baseline of gossip averaging was used. In the malicious case, the baseline used was Trimmed Mean (TM) aggregation with $c_{max,i} = \lceil \frac{c*k_i}{m} \rceil$ for every node i where k_i is its number of neighbors. This trims the $c_{max,i}$ number of received gradient values from both the higher and lower end out of the k_i values received based on their magnitude for every model weight and bias. The gradients here refer to the difference between the current received model from a neighbor and one's own model from the past iteration. By default, we choose $m = 128$ and the data was distributed in a non-IID manner with a bias b of 0.5 for the image datasets, both of which had 10 classes. The nodes were divided into 10 groups and a data sample with label l was assigned to group index l with a probability b and to any other group with a probability $\frac{1-b}{9}$. Within the group, the data samples were distributed randomly among the clients. For the NLP dataset, the data was divided sequentially (which is also non-IID) into $m + 1$ chunks, one for each client and the last chunk for testing.

4 EVALUATION

We first demonstrate how P2PRISM stops the spread of a graybox attack - one that only knows that a mean-like aggregation is being used but is oblivious to the details of the defense; and helps the benign nodes recover. Then we show how P2PRISM keeps the consensus distance of the graph under control which is most necessary for decentralized peer-to-peer learning. We then proceed to stress test P2PRISM under a whitebox attack where the attacker also has access to the defense

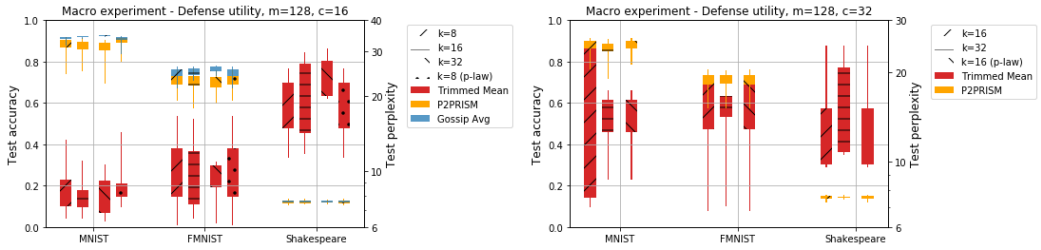

Figure 2: The figures show the final test accuracy (perplexity for Shakespeare) candles of the benign nodes comparing Trimmed Mean and P2PRISM under malicious conditions - $c/m = 1/8$ in the left and $c/m = 1/4$ in the right figure, along with the baseline benign results for Gossip Averaging in the left figure. For every dataset, the last bar corresponds to a power-law graph while the rest denote k-regular graphs.

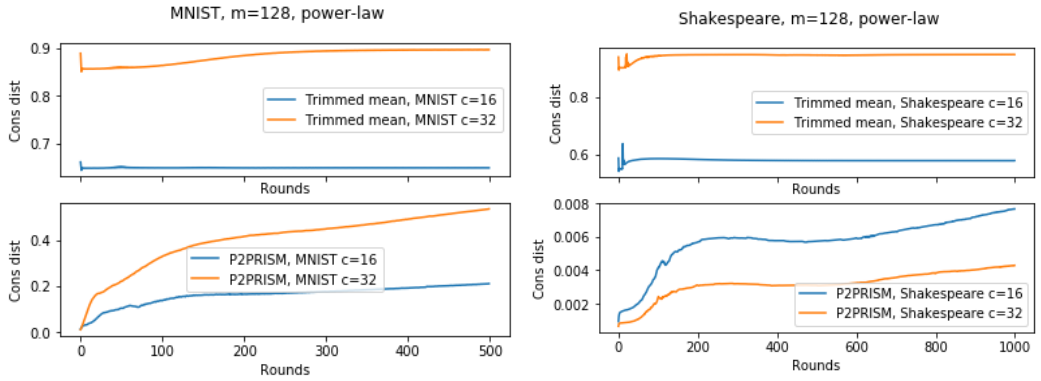

Figure 3: The figures show how consensus distance among the benign nodes changes during training. We show the results on one image dataset - MNIST on the left and on an NLP dataset - Shakespeare on the right. The graph topology chosen was power-law to mimic a realistic setting with $k = 8$ for $c = 16$ and $k = 16$ for $c = 32$. It is very clear how P2PRISM keeps the consensus control low as compared to Trimmed mean in a malicious environment.

algorithm, its parameters, and the cutoff flip-score in the past iteration beyond which all models were penalized. We find that the flip-score generated by malicious clients in the graybox attack are easily distinguishable from the benign ones due to the flip-score value being too large than the benign values and the detection rate is not affected by the defense parameter μ at all. Therefore, we use the whitebox attack to show the effect of varying μ where the malicious and benign flip-scores are indistinguishable.

Figure 2 shows the candle plot of test accuracy under different conditions. The lower and upper ends of the candle correspond to the 25-th and 75-th percentile final test accuracy (or test perplexity for Shakespeare dataset) respectively among the nodes, and the ends of the candle wicks correspond to the lowest and highest final test accuracy at the end of training. The attack completely damages P2PL without any secure aggregation to a random test accuracy of 10%. This is also evident from the level of impact that just 2 malicious nodes had in Figure 1. We do not show that in the plots here. We see that Trimmed Mean fails against the attack because of two reasons - 1) Magnitude based filtering is inefficient against a gradient ascent attack, and 2) a fixed value of c_{max} although suitable for FL does not work for P2PL where the fraction of malicious nodes differs in every neighborhood in any practical case. We also observe a clear separation between the performance of Trimmed Mean and P2PRISM where P2PRISM is very close to the benign standard of Gossip Averaging, and is in some cases even completely overlapping with it. A small dip in the test accuracy could be explained from the fact that the malicious nodes attacked in every iteration and hence nothing was learned from their local data which did contribute in the completely benign case. It is important to observe that P2PRISM excels for both the graph types - k-regular and power-law across all the 3 datasets used.

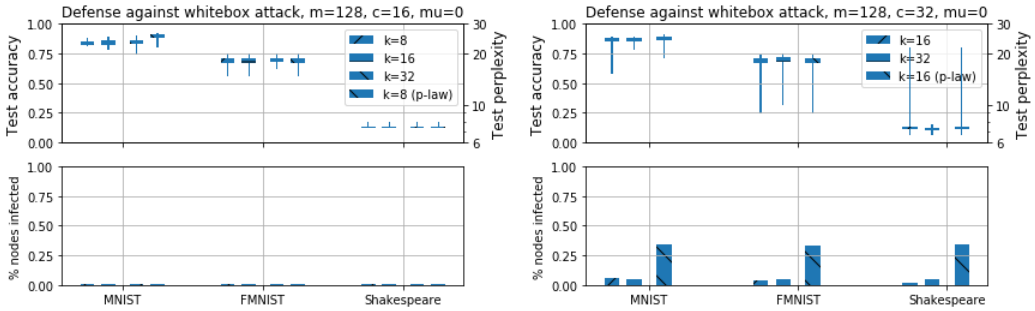

Figure 4: The figures show the performance of P2PRISM against a whitebox attacker that crafts a stealthy attack with full knowledge of the defense and its parameters to bypass the prism used in the defense. We consider the attack successful if it drops the test accuracy on MNIST below 0.8, and on FMNIST below 0.65, or if it raises the test perplexity on Shakespeare above 8. We observe that with $c = 16$, P2PRISM was successful in preventing all benign nodes from being infected. With $c = 32$, P2PRISM could restrict the number of infected nodes below 6% for a k-regular graph and under 34% for a power-law graph.

Next, we show how P2PRISM keeps the consensus among the benign nodes under control even in the presence of the attackers. In P2PL, every node can have a different model state, but as long as they are close to each other, they are all moving along the same trajectory towards the same global solution. We set the nodes to be at consensus initially. Naturally, as they train, they are going to drift from each other unless the graph is fully connected. The attackers try to spread dissensus among the benign nodes by infecting their neighbors. After a certain point, even if the attack is stopped, the existing dissensus may inhibit training as the nodes with disagreeing models collaborate. P2PRISM shows how, as in Figure 3, by identifying and filtering the malicious models based on one's flip-score, such a dissensus scenario is prevented. Having shown the success of P2PRISM against a graybox attack where the attacker was not aware of the details pf the defense being used by the nodes, we present a whitebox attacker at time t that not only has access to the defense parameter μ, but also the minimum, median, and cutoff flip-score of all the clients at $t-1$. The attacker uses this information to craft the attack at time t by optimizing the attack parameter λ so that the expected flip-score at t will be less than the cutoff flip-score at time $t-1$ for the node being targeted. This is a highly personalized attack and the information from the past, that is, from $t-1$ is used because the median, minimum, and cutoff values at t is determined by a node only after receiving all the updates, while the attacker is yet to craft its attack. Given that every node just runs one iteration of local training, and the training trajectory is smooth, the estimation done by the attacker based on one timestep in the past is extremely accurate. The malicious and benign flip-scores are now indistinguishable and the defense is bypassed. However, we observe that this stealthy attack loses its damaging impact as a trade-off most of the time. We see in Figure 4 that P2PRISM could save 100% of the nodes from being infected when $c/m = 16/128 = 1/8$. However, on increasing the number of attackers to 32, that position themselves at random locations in the graph, the attackers do form a majority at some neighborhoods and are able to attack the benign nodes at such vulnerable locations. This number is very small though, as can be seen in Figure 4, that the attacker could infect less than 6% of the benign nodes in a k-regular graph and less than 34% in a power-law graph. Although the power-law graph seems to be infected, it is only because the attack, being whitebox, stealthy, and controlling 1/4-th of the popular nodes becomes too powerful. This shows the limit of malice upto which P2PRISM can be expected to provide perfect security, after which, saving all nodes from being infected cannot be guaranteed. μ was set to zero to allow the defense to be conservative against an attacker which is too powerful.

We take this opportunity to also demonstrate the effect of sweeping μ. With higher μ, the flip-score cutoff increases, thereby giving the attackers an opportunity to be stealthy even with higher attack magnitude λ. Thus, the stealthy attack may still be damaging as the detection threshold is not as sensitive anymore. We see the same trend in Figure 5 where the whitebox attackers are able to impact more and more nodes as μ increases, however this increase is too gradual and it reflects the robustness of P2PRISM. This also empirically shows why we have chosen a default value of $\mu = 0.75$ for all our experiments with the graybox attacker. It is always recommended to use a low

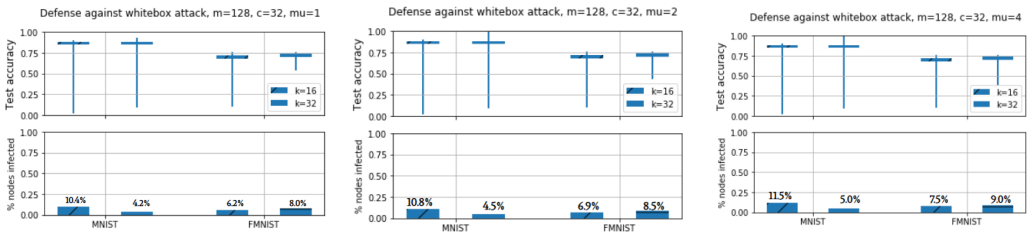

Figure 5: The figures show the effect of sweeping the defense parameter μ against a whitebox attack on MNIST and FMNIST in a k-regular graph with $m = 128, c = 32$. With increasing μ, the cutoff flip-score increases for every node and a relatively stronger attack can get the opportunity to be stealthy and bypass P2PRISM. Hence, we see a gradual increase in percentage of infected nodes. However, even with a high value of $mu = 4$, this increase is very slow the fraction of infected nodes stay below 12%, thus proving the robustness of P2PRISM, that is, low sensitivity with respect to tis parameter across a large sweep of its parameter.

μ for a conservative defense, however, this experiment proves that the performance of P2PRISM is not too sensitive with respect to its parameter μ.

5 Conclusion

In this paper, we have presented a secure aggregation for nodes participating in collaborative learning as each others' peers. We began by introducing the concept of peer-to-peer learning in comparison with federated learning. We saw how the consensus distance at model initialization can significantly affect the training process and stressed on initializing all the participating nodes with the same model. This is an extremely important aspect of the training process but has been ignored in the context of decentralized collaborative learning. We also discussed how P2PL is differently affected by byzantine nodes as compared to FL depending on the graph connectivity, the number of gossiping steps per round, and the location of the malicious nodes. Unlike FL, an attacker can have full access to its neighbors' models via gossiping and can launch a personalized attack. At the same time, every node has access to its own trusted benign gradient updates to compare all incoming updates with, and unlike FL, a node can perform a trusted secure aggregation for itself, given a robust defense policy. We leveraged this fact to construct our own defense policy, P2PRISM based on the concept of flip-score that can protect a node even if its surrounded by a malicious majority in its neighborhood. We evaluated P2PRISM against the state-of-the-art model poisoning attack and compared it against the secure aggregation protocol of Trimmed Mean by extending it to the decentralized learning domain. P2PRISM could successfully recover all benign models under a graybox attacker, while also maintaining a low consensus distance among the benign nodes across all the three datasets of MNIST, FMNIST, and Shakespeare on both the k-regular and power-law graph topologies. This was followed by a stress test of P2PRISM against a whitebox attacker that also has access to the defense parameter in addition to access to all the benign models, that can launch a highly personalized attack to each of its neighbors. We saw that P2PRISM could recover all the being node even in the presence of this extremely powerful attacker when only 1/8-th of the nodes were malicious. However, when the fraction of malicious nodes go up to 1/4, some benign nodes begin to be infected, but their fraction is contained to less than 6% for a k-regular graph. With a power-law graph in the presence of such a powerful adversary controlling 1/4-th of all the nodes, P2PRISM could only save 66% of the nodes from being infected. We also swept the defense parameter under the presence of this whitebox attacker to show the robustness of P2PRISM across a large usable range of the parameter. P2PRISM could successfully contain the infection within 12% of the benign nodes. With this, we conclude the evaluation of P2PRISM and propose its usage at every node in a decentralized learning architecture for the best possible security against model poisoning attacks.

REFERENCES

Mahmoud Assran, Nicolas Loizou, Nicolas Ballas, and Mike Rabbat. Stochastic gradient push for distributed deep learning. In *International Conference on Machine Learning*, pp. 344–353. PMLR, 2019.

Peva Blanchard, Rachid Guerraoui, and Julien Stainer. Machine learning with adversaries: Byzantine tolerant gradient descent. In *Advances in Neural Information Processing Systems*, pp. 119–129, 2017. Krum paper.

Sebastian Caldas, Sai Meher Karthik Duddu, Peter Wu, Tian Li, Jakub Konečný, H. Brendan McMahan, Virginia Smith, and Ameet Talwalkar. Leaf: A benchmark for federated settings. *arXiv preprint arXiv:1812.01097*, 2019.

Xiaoyu Cao, Minghong Fang, Jia Liu, and Neil Zhenqiang Gong. Fltrust: Byzantine-robust federated learning via trust bootstrapping. *arXiv preprint arXiv:2012.13995*, 2020.

Xiaoyu Cao, Minghong Fang, Jia Liu, and Neil Zhenqiang Gong. Fltrust: Byzantine-robust federated learning via trust bootstrapping. *Network and Distributed System Security Symposium*, pp. 1–18, 2021.

Liam Fowl, Jonas Geiping, Wojtek Czaja, Micah Goldblum, and Tom Goldstein. Robbing the fed: Directly obtaining private data in federated learning with modified models. *arXiv preprint arXiv:2110.13057*, 2021.

Clement Fung, Chris JM Yoon, and Ivan Beschastnikh. The limitations of federated learning in sybil settings. In *23rd International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2020)*, pp. 301–316, 2020.

Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. Inverting gradients-how easy is it to break privacy in federated learning? *Advances in Neural Information Processing Systems*, 33:16937–16947, 2020.

Rachid Guerraoui, Sébastien Rouault, et al. The hidden vulnerability of distributed learning in byzantium. In *International Conference on Machine Learning*, pp. 3521–3530. PMLR, 2018.

Jakub Konečnỳ, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*, 2016.

Lingjing Kong, Tao Lin, Anastasia Koloskova, Martin Jaggi, and Sebastian Stich. Consensus control for decentralized deep learning. In *International Conference on Machine Learning*, pp. 5686–5696. PMLR, 2021.

H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*, 2017.

El Mahdi El Mhamdi, Rachid Guerraoui, and Sébastien Rouault. The hidden vulnerability of distributed learning in byzantium. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pp. 3521–3530, 2018. Bulyan paper.

Atul Sharma, Wei Chen, Joshua Zhao, Qiang Qiu, Somali Chaterji, and Saurabh Bagchi. Tesseract: Gradient flip score to secure federated learning against model poisoning attacks. *arXiv preprint arXiv:2110.10108*, 2021.

Virat Shejwalkar and Amir Houmansadr. Manipulating the byzantine: Optimizing model poisoning attacks and defenses for federated learning. In *NDSS*, 2021.

Qi Xia, Zeyi Tao, Zijiang Hao, and Qun Li. Faba: An algorithm for fast aggregation against byzantine attacks in distributed neural networks. In *IJCAI*, pp. 4824–4830, 2019.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.

Dong Yin, Yudong Chen, Ramchandran Kannan, and Peter Bartlett. Byzantine-robust distributed learning: Towards optimal statistical rates. In *International Conference on Machine Learning*, pp. 5650–5659. PMLR, 2018a.

Dong Yin, Yudong Chen, Ramchandran Kannan, and Peter Bartlett. Byzantine-robust distributed learning: Towards optimal statistical rates. In *International Conference on Machine Learning*, pp. 5650–5659. PMLR, 2018b.