

PIRATE: A Blockchain-based Secure Framework of Distributed Machine Learning in 5G Networks

Sicong Zhou, Huawei Huang, Wuhui Chen, Zibin Zheng, and Song Guo

Abstract—In the fifth-generation (5G) networks and the beyond, communication latency and network bandwidth will be no more bottleneck to mobile users. Thus, almost every mobile device can participate in the distributed learning. That is, the availability issue of distributed learning can be eliminated. However, the model safety will become a challenge. This is because the distributed learning system is prone to suffering from byzantine attacks during the stages of updating model parameters and aggregating gradients amongst multiple learning participants. Therefore, to provide the byzantine-resilience for distributed learning in 5G era, this article proposes a secure computing framework based on the sharding-technique of blockchain, namely PIRATE. A case-study shows how the proposed PIRATE contributes to the distributed learning. Finally, we also envision some open issues and challenges based on the proposed byzantine-resilient learning framework.

I. INTRODUCTION

Machine learning has spawned a lot of useful applications, such as computer vision, and natural language processing, etc. However, the parties who benefit from the technology are mostly large organizations, e.g., commercial companies and research institutes. Individuals who only hold devices with limited computing-ability cannot take part in machine learning tasks.

In fact, the combined power of individuals by their holding devices has been much underestimated. For instance, the mobility and the large population of the mobile users can bring flexible computing possibilities. Taking these advantages into account, the distributed learning [1] enables individual devices to learn collaboratively. However, device heterogeneity and network conditions have been viewed as the challenges of distributed learning.

On one hand, a recent literature [2] shows that asynchronous distributed learning starts to have a sound theoretical basis. In the synchronous manner of distributed-learning, the unequal training times in each iteration induce the inevitable waiting dissipation on the fast training nodes. This problem brought by device heterogeneity will be eliminated with the growing maturity of the asynchronous learning.

On the other hand, communication latency and bandwidth are still viewed as the bottleneck resources of distributed machine learning [3]. Nowadays, the network bandwidth of most users is still insufficient to support the distributed-learning tasks. This situation makes distributed learning highly

unavailable for the majority. Fortunately, latency and bandwidth will be no more obstacles to distributed learning in the era of the fifth-generation (5G) networks, because the network conditions of individuals will be substantially improved.

Generally, the ideal factors that affect the distributed learning include the following aspects:

- High availability: any device can perform learning anytime.
- High scalability: the learning framework should support high concurrency, high communication efficiency and low storage complexity.
- Decentralization: the intervention of a centralized third-party is minimum.
- Byzantine-resilient model safety: the future distributed learning should ensure byzantine-resiliency [4], which indicates that the distributed learning can endure arbitrary attacks.

Since the availability of distributed-learning systems can be significantly improved in the 5G networks, we can reasonably envision a distributed learning system in which the heterogeneous devices are available anytime. This is because the 5G technologies enable various devices to participate in distributed-learning tasks with the sufficient network bandwidth. In such a sound distributed learning system, another challenge is that the learning system is prone to suffering from *byzantine attacks* [4]. Therefore, more sophisticated approaches that can ensure the byzantine-resilience are required.

Recently, byzantine-resilient machine learning under master/slave settings has gained much attention [5]–[7]. In a byzantine-resilient distributed-learning task, two things are risk-prone: 1) gradient aggregation, and 2) model parameters.

Conventionally, the byzantine-resilient distributed-learning tasks are conducted under centralized settings, in which the byzantine-tolerant components rely on a globally trusted third-party as the *parameter server*. The problem is that the workload-handling capacity of such the centralized parameter server is usually a bottleneck while performing the distributed learning. Moreover, to provide reliable services against the vulnerability of single-point-of-failure (SPOF), the redundant deployment of resources is entailed at the centralized third-party. Therefore, the centralized byzantine-resilient learning induces a high operational-expenditure (OPEX).

To achieve high availability while fulfilling the requirement of OPEX-efficiency and byzantine-resiliency, the decentralized learning system would ideally be decentralized. For example, a recent study [3] claims that the decentralized learning system is more efficient in terms of communication and storage than

S. Zhou, H. Huang, W. Chen and Z. Zheng are with the School of Data and Computer Science, Sun Yat-Sen University, China. Email: huanghw28@mail.sysu.edu.cn

Song Guo is with the Department of Computing, the Hong Kong Polytechnic University. Email: song.guo@polyu.edu.hk

the centralized learning systems. However, to the best of our knowledge, the byzantine-resilient learning in a decentralized manner has not been well studied.

Although the decentralized scheme shows the promising efficiency, a great challenge in terms of model-safety comes along with it. That is, the aggregation stage of distributed learning can be easily falsified, since the distributed-learning task becomes a collaboration among multiple peers. Once the model parameters are contaminated at one of the peers, this local-model contributor would bring limited or even negative contribution to the holistic learning task. Thus, to achieve the decentralized byzantine-resilience for distributed learning, a *state machine replication* (SMR) protocol [4] is required. As a representative paradigm of SMR protocols, the blockchain technology can provide such the critical byzantine-resilience to distributed learning. In detail, blockchains can provide protection to the model parameters, particularly at the stage of gradient aggregation, while ensuring a high level of availability, cost-efficiency and decentralization.

To this end, this article proposes a sharding-based blockchain framework, which is named PIRATE, for byzantine-resilient distributed-learning under the decentralized 5G computing environment.

II. PRELIMINARIES OF DISTRIBUTED MACHINE LEARNING IN 5G NETWORKS

A. Consensus Protocols for Decentralized Learning in 5G

To achieve the global agreement within a decentralized setting, a byzantine-tolerant consensus protocol for SMR is needed. Such the consensus protocol should ensure that the honest nodes can reach an agreement on the order of correct execution, i.e., model updates, even when a certain amount of byzantine effort plays a role [4]. In this section, we first briefly review the existing byzantine tolerant consensus protocols towards SMR, and then analyze what protocol is applicable for decentralized learning in the 5G era. For brevity, we call the byzantine-tolerant consensus protocol the consensus protocol in the remainder of this paper.

Consensus protocols can be categorized into 2 types: the competition-based and the communication-based. The leader of competition-based consensus, e.g., the Proof of Work (PoW) adopted by bitcoin (<https://bitcoin.org>), needs to earn his leadership through a “fair” competition. The communication-based consensus protocols (e.g., Hotstuff [8]) select leaders through a deterministic way, or based on an unbiased randomness generated collaboratively.

In the context of blockchains, for the competition-based consensus protocols, blocks are appended on the chain *before* consensus while the communication-based protocols append blocks thereafter. For the competition-based methods, higher scalability inevitably incurs higher chance of forking. Plagued by the byproduct, competition-based methods struggle to achieve a high scalability.

While communication-based consensus protocols have no concern of forks (instant finality), they require multiple rounds of communication to reach an agreement. The high communication overhead also hurts the scalability of communication-

based protocols. Therefore, the communication-based protocols are commonly used at a smaller scale.

Sharding-based consensus protocols [9], as a hybrid approach, can achieve scalable consensus in the permissionless blockchain. Such hybridization benefits from the instant finality of the communication-based methods, and permissionless resiliency of the competition-based methods. One representative work of the sharding-based approaches is *RapidChain* [9], which achieves a throughput of 7,300 tx/sec while maintaining a total resiliency of 1/3, i.e., the RapidChain system functions properly in face of 1/3 byzantine-malicious participants.

In 5G networks, we need a consensus protocol that is available for a large scale of participants when they are training large models. An ideal consensus protocol should be scalable and permissionless. Sharding-based protocols make full use of integrated resources by splitting a workload and amortizing all tasks among multiple committees. Such scalable strategy can fit perfectly in the environment of distributed machine learning. However, permissionless distributed-learning in a sharding-based community is challenging. Unlike a financial system where good behaviors can be stimulated by a certain incentive mechanism, when performing the distributed-learning tasks, only the honest behavior is not enough. Because participants cannot adjust the network conditions and computation abilities of their devices according to their behaviors. Thus, a consistent assessment of training-reliability throughout the training process is required for the distributed-learning task to function efficiently at each participant. Accordingly, we adopt a permissioned version of sharding-based consensus protocol in our proposed framework.

B. Configurations of Distributed Machine Learning

Machine Learning problems mostly rely on certain specific optimization problems. In the distributed machine learning, each participant can only yield a partial solution. In order to orchestrate a global optimization, as shown in Fig. 1, computing nodes need to carry out the following steps. First, each node computes local gradients based on their local data. Then, nodes communicate with each other to get a globally aggregated gradient for model update. Based on the style how they get a global aggregation, two typical styles of configuration have been proposed:

Master/slave style: A centralized parameter server aggregates gradients and synchronizes the aggregated result at each computing node.

Decentralized style: No centralized server is entailed. Each node can only communicate with its neighbors to exchange model parameters.

We then discuss the superiority of decentralization on distributed machine learning:

- **Communication efficiency.** A recent work [3] demonstrated that the decentralized settings can better exploit the bandwidth resource, avoid traffic jams and share workloads among peers than the centralized master/slave settings.
- **Cost efficiency.** When the scale of participation grows substantially, no single party should be responsible for

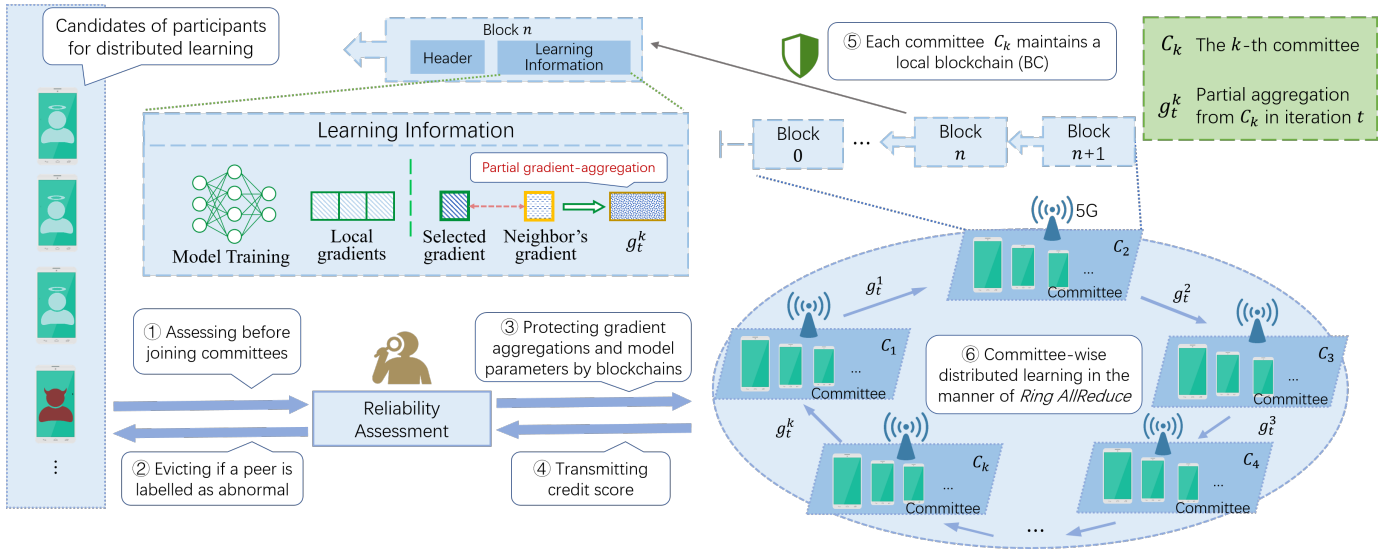


Fig. 1. The proposed PIRATE framework has two critical components: 1) reliability assessment, and 2) blockchain (BC) systems for distributed SGD (D-SGD). Gradient aggregations and model parameters are protected by blockchains. Meanwhile, reliability assessment determines the participants of distributed learning tasks.

maintaining the system. Analogous to the situation of cloud netdisk services nowadays, if a service provider plays such a dominant role, the OPEX cost of maintaining the centralized system eventually transfer to clients. As a feasible solution, service providers enforce clients to choose between low quality of service (e.g., very limited download speed) and expensive membership fee. This is an opposite of win-win for both providers and clients, provided that there are better solutions.

- **Reliability.** The centralized design of master/slave settings suffer from SPOF problem. Once the centralized server is overloaded or under attack, the whole system ceases to function. Thus, the communication burdens and the attack risks on one single server impair the reliability of the system.

III. STATE-OF-THE-ART STUDIES OF BYZANTINE-RESILIENT MACHINE LEARNING

A. Byzantine-Resilient Machine Learning

To train models in a large-scale multi-party network, serious security problems are needed to address if adopt the Sharing technique. As the scale of participants grows, the behaviors of computing nodes become more unpredictable. The distributed stochastic gradient descent (D-SGD) [10] framework should tolerate byzantine attacks, i.e., arbitrary malicious actions, such as sending harmful gradients and corrupting training models as shown in Fig. 2. Current studies on Byzantine-resilient machine learning mostly focus on protecting gradient aggregation. However, the parameters of training models owned by each data trainer are also vulnerable to be falsified. We elaborate the protection of gradient aggregation in the next section, and analyze existing frameworks that can protect both aspects of model training.

A blockchain-based learning method was proposed by Chen *et al.* [11], in which the proposed architecture *LearningChain*

is able to simultaneously protect gradient aggregation and model parameters, by storing them together on-chain. By exploiting the traceability of blockchain, erroneous global parameters produced by malicious nodes can be rolled back to its former unfalsified state. Historical parameter records cannot be falsified due to the tamper-proof character of blockchain. The proposed byzantine-tolerant gradient aggregation with a low computation-complexity is called “ l -nearest gradients aggregation”. This method ensures that if byzantine computing nodes yield local malicious gradients to prevent the convergence of learning algorithms, their effort would be mitigated. However, *LearningChain* still utilizes a master/slave setting for distributed gradient descent where the parameter server is randomly selected by PoW competition. In addition, the on-chain data could be potentially oversized, because any node in the system would have to store all the historical model parameters and gradients. Such architecture is prone to lead to a significant traffic congestion and a substantial storage overhead. When the scale of learning grows substantially, the storage and communication overheads could potentially weaken the advantage of distributed learning.

In terms of reliability, rollback can possibly fail if two consecutive byzantine leaders collude. That is the case when the previous byzantine leader contaminates the model updates and the subsequent leader properly updates but based on the contaminated model parameters. The honest leader following the two previous byzantine leaders is unable to detect the model contamination, since the model-examination process terminates when the honest leader approves the immediate update proposal, even though the update is based on contaminated model parameters. The essential problem is that, the model update is examined by only one leader when a proposal is submitted. In contrast, our proposed framework adopts a more decentralized setting, in which all nodes can naturally participate in validating the model updates, and every node maintains its own training model rather than accepting updates

from a third party.

B. Byzantine Protection on Gradients

Before the stage of updating training models, computing nodes need to aggregate their local gradients. The conventional aggregation solution of simple linear combinations (e.g., averaging [12]) cannot tolerate one byzantine worker [5]. Thus, byzantine protection on gradients aggregation has gained much growing attention. Basically, the existing byzantine-based approaches can be classified into 2 categories: the tolerance-based methods and the detection-based methods.

Blanchard *et al.* [5] proposed a byzantine tolerant method called *Krum*, and a variation called *Multi-Krum* which is an interpolation between *Krum* and averaging [12]. Instead of using a straight-forward linear combination of multiple local gradients, *Krum* precludes gradients too far away from the majority and chooses one local gradient based on a spatial score. To improve convergence speed, *Multi-Krum* chooses a number of local gradients based on the same score as *Krum*, and outputs their average. Experiments show that even with 33% of omniscient byzantine workers, the error rate is almost the same as that of 0%. However, *Krum's* computation complexity is $O(n^2)$. Compare with the averaging approach whose computation complexity is $O(n)$, *Krum* may impede the scalability of distributed learning.

The heuristic byzantine tolerant aggregation solution *l*-nearest gradients proposed by Chen *et al.* [11] cannot guarantee safety against omniscient attacks by byzantine workers. The aggregation solution is to aggregate *l* gradients closest, based on their cosine distances, to the sum of the received gradients. If an omniscient byzantine worker manages to acquire all local gradients for other workers in time, the byzantine worker can yield a local gradient that changes the global sum arbitrarily [5]. Sacrificing security, the algorithm gains a good time complexity $O(n)$.

One byzantine-detection method proposed by Li *et al.* [7] is designed for federated learning (FL). Existing byzantine-tolerant aggregation methods are mostly inefficient due to the non-identically and independently distributed training data. Experiments show that their detection-based method has a better performance than tolerance-based methods in federated learning. In their setting, a credit score was assigned by the pre-trained anomaly-detection model to each computing node. Since the weight of the local gradient was determined by the credit score, the weighted sum aggregation can filter out the byzantine local gradients.

Besides the tolerance and detection methods, another machine learning method was proposed by Ji *et al.* [13] to learn the gradient aggregation. Different from the deterministic aggregation solution, they model the aggregation as a learning problem. They utilize an Recurrent Neural Network (RNN) aggregator with all local gradients as input, and additional loss information, i.e., the loss based on a small sample of dataset, for each worker. The RNN aggregator optimizes an objective that depends on the trajectory of the original optimization problem.

The tolerance-based gradient aggregation methods are mostly designed under an independent identically distributed

(i.i.d) assumption. Therefore, in the settings of federated learning where data is non-i.i.d, most tolerance-based methods do not perform well. The relaxation of data distribution brings challenge to the tolerance-based methods. However, the tolerance-based methods have the benefit of simplicity that do not require additional training. Non-i.i.d tolerance-based methods are challenging and interesting at the same time.

As shown in Table. I, we compare the performance of different protection approaches for gradient-aggregation, and analyze their computation complexities, under two settings, i.e., the normal setting and the federated learning setting.

C. Protection over Decentralized Machine Learning

Recall that we have analyzed how a decentralized learning system is more promising than the centralized in terms of efficiency and reliability at large scale networks. However, the decentralized byzantine-resilient aggregation and the parameter protection for byzantine models have been ignore.

In a decentralized scheme, each node, including the malicious one, has a great impact on the global aggregation. As shown in Fig. 2, the existing protection methods are not applicable to the decentralized settings because of the following reasons.

- Attacks on partial aggregation are detrimental. Every node aggregates a partial aggregation provided by another node. For byzantine nodes, they have more attack patterns undermining learning tasks, such as sending falsified partial aggregation results or sending nothing to interrupt the aggregation process. Thus, without a quorum of validators in a multi-party network, partial aggregation process cannot be trusted.
- Without a trusted third party, anomaly-detection would be challenging. Credit scores of learning can be easily falsified without proper validation mechanisms.
- Moreover, synchronization of model parameters could be a problem. Every node needs to maintain a local training model. Once contaminated by attackers, computing nodes will have no actual contribution to the holistic learning task.

In summary, blockchain as a decentralized SMR system, can provide a quorums of validators and practical synchronization mechanism to achieve byzantine-resiliency in decentralized learning.

IV. OUR PROPOSAL - PIRATE: A MACHINE LEARNING FRAMEWORK BASED ON SHARDING TECHNIQUE

A. Overview of PIRATE

Generally, we propose a framework of blockchain-based protection framework for the distributed-machine learning that is named PIRATE and managed in a decentralized manner. To protect both the gradient aggregation and model parameters, the proposed approach utilizes a sharding-based blockchain protocol and gradient anomaly-detection [7]. Furthermore, with anomaly detection, the malicious nodes who yield harmful gradients can be replaced by eliminating their effort.

The proposed framework can be also applicable for federated learning under a decentralized configuration, where no a centralized parameter server is required.

TABLE I
GRADIENT PROTECTION METHODS

Method types	The tolerance-based		The detection-based	The learning-based
Representative studies	Krum [5]	l -nearest gradients [11]	Anomaly Detection [7]	Learning to learn [13]
Resiliency under 30% attack	High	Medium	Unknown	High
Resiliency under 30% attack (FL)	Low	Unknown	High	Unknown
Resiliency under majority attack	Low	Medium	Unknown	High
Resiliency under majority attack (FL)	Low	Unknown	High	Unknown
Computation complexity	$O(n^2)$	$O(n)$	$O(n)$	Model-related
Other functions except <i>aggregation</i>	None		Autoencoder training	Aggregator training

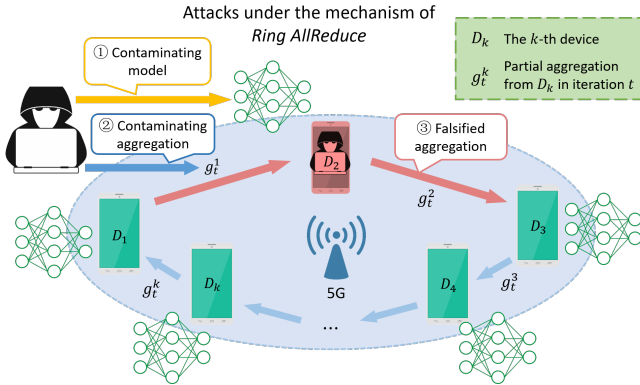


Fig. 2. While adopting the *Ring AllReduce*¹ mechanism, malicious attackers can perform attacking from both inside and outside. ① Attackers from the outside can contaminate training models in target nodes. ② The outside attackers can also attack partial gradient-aggregation. ③ Byzantine computing nodes can send harmful aggregations that damage the convergence of learning tasks.

B. Permission/Access Control

In a distributed learning task with high availability, reliability assessment is essential for all participants, especially the mobile devices, which have volatile states. For instance, devices could be busy running some CPU-consumed tasks or simply charging in a sleeping mode. In the perspective of network conditions, they can be well connected in a net cafe or without signal at all while travelling through a tunnel. Allowing devices that are in bad states to keep their learning tasks would severely slow down the entire learning process. Thus, the volatile states of devices demand real-time reliability assessment to conduct permission control.

We propose a centralized solution for permission control. Fig. 1 depicts the permission/access control of PIRATE. Similar to federated learning, before actually contributing to the global learning task, all nodes are assessed by a centralized component based on their computation ability, network condition, join/leave prospect and historical credit scores. Nodes once granted permission, can join a committee using the Bounded Cuckoo rule [9].

During training, validated credit scores generated by committees are transmitted to the permission-control center. Nodes with low accumulated credit scores would be evicted from the system.

C. Sharding-based Blockchain Protection towards Decentralized Distributed-Learning

We propose a sharding-based blockchain mechanism for the protection of decentralized distributed-learning. Existing byzantine-resilient aggregation solutions are not applicable under the decentralized settings. Having considered the increased freedom of byzantine nodes, our mechanism provides a consensus protocol on model parameters and gradient aggregation. Taking the advantages of 5G techniques, the sharding-based protection can work efficiently.

We randomly split the computing nodes into multiple committees, in which partial aggregations are agreed among committee members. Since aggregations are no longer centralized, the burden of aggregation workload is mitigated.

Let n denote the total number of computing nodes, c denote the size of a committee, we assume that through permission management, every nodes computation time for local gradients is roughly the same. All nodes have mini-batches assigned according to their computation resources. We then discuss the key steps of the proposed framework.

Random committee-construction. All nodes would have a random identity. According to the identities, committees of size c are formed. Every committee member knows the identity of all honest peers with high probability. After a certain number of rounds of training, a portion of nodes would be replaced by reliable nodes to prevent slowly-adaptive adversary [14].

Intra-committee consensus. In the first round, every member gossips its local gradient with the timestamp of the gossiping among its committee. In every step, a leader would aggregate two things: an agreed selection of c^2/n gradients and the neighbors aggregation from last step. Meanwhile, members of the committee are required to validate the result. After a global stabilization time (GST), an agreement is reached by a partially synchronous consensus protocol Hotstuff [8] in every committee. With a sufficient number of members having agreed and partially signed (threshold signature) on the data, the data is tamper-proof. Only the partial aggregation result is to be transmitted to the neighbor committee.

Global consensus. Similar to the *Ring Allreduce*¹ process, committees communicate locally agreed data with their neighbor committees. After $2(n/c - 1)$ steps, every node in the system would finally have a globally agreed aggregation.

¹<https://github.com/baidu-research/baidu-allreduce>

Reconfiguration. After some iterations of training, some malicious nodes in the committees are replaced. We use Cuckoo Rule [9] against join/leave attack while maintaining a same resiliency of 1/3 after reconfiguration.

D. Intra-Committee Consensus

As illustrated in figure 3, a learning task proceeds in *iterations*, wherein each iterations consist of *steps*. In an iteration, all nodes start by computing their local gradients and ends by updating the model parameters as in all distributed learning procedures. In each step, all nodes take part in a consensus process to ensure that byzantine efforts are limited.

When a local gradient is computed, members take part in repeated consensus steps. A consensus step (CS) includes:

- **Component 1:** local gradient selection,
- **Component 2:** neighbor committee aggregation,
- **Component 3:** aggregation result.

Fig. 3 depicts the intra-committee consensus process. We define the process of finishing a partial aggregation a *step*, the process of finishing a model update an *iteration*. We utilize the three-phase chained Hotstuff consensus protocol [8], wherein a consensus step require 4 phases of communication. Each phase is driven by a leader issuing a block containing verifiable quorum certificates. For component 1, committee members can either collaboratively select c^2/n local gradients, or coordinate in a round robin manner to choose c^2/n local gradients. Having a neighbor committee aggregation and a set of local gradients, both leaders and members can aggregate them using the detection-based BFT aggregation [7]. A pre-trained anomaly detection model would assign a weight to each gradient according to the anomaly score. If the anomaly score surpasses a threshold, zero weight would be assigned accordingly, thereby filtering harmful gradients that hinders convergence. Meanwhile, members are required to store historical credit scores of each other until a new committee is formed. Before the committee reconfiguration, credit scores are transmitted to the permission control center.

For component 2, members wait for the leader of the neighbor committee to broadcast the tamper-proof neighbor aggregation from last step. Since component 2 is an agreed aggregation result from the neighbors last step, it can be easily verified. Finally for component 3, the incumbent leader broadcast the partial aggregation and a hash index of training parameters for members to verify and agree on. Having sufficient quorum certificates, the leader would broadcast the decided aggregation in the committee and in the neighbor committee. If the leader chooses to withhold the result from the neighbor committee, the neighbor committee can ask a random committee member for the result.

As shown in Fig. 3, in order to reach a consensus decision (CS4), three phases of validation are required. This means that only 1/4 of blocks are generating aggregations. To address this issue, we can pipeline consensus steps as encapsulated commands to achieve a better performance. As Fig. 3 shows, utilizing the frequent view-change of Hotstuff, each leader would be responsible for driving four consensus steps in different phases, with each phase validating an aggregation

proposal. In an ideal scenario where no byzantine leader is elected, every block generates an agreed aggregation in average. Members are required to store four sets of gradients for validation, with each set composed of its own local gradient, corresponding aggregation of neighbor committee and an aggregation proposal from incumbent leader. Since PIRATE does not rely on rollback to recover from misdirection of the byzantine leadership, computing nodes are free from the storage burden of the entire history of all gradients.

V. CASE STUDY

A. Implementation and Methods

We implement a prototype of PIRATE with python 3.7 for performance evaluation. On one machine, we run 50 to 100 instances each to simulate a single committee. Assuming after permission control, devices spend a same amount of time for computing gradients. We simulate this process by having instances wait for a same amount of time to generate an equal-sized chunk of data (28MB). Then, instances transmit chunks of data to simulate a decentralized SGD process. The machine is a mini PC (model serial number: NUC8i5BEK), with a quad-core i5-8259U processor (up to 3.80 GHz). To simulate the network condition of 5G, we assume every message has a 10ms latency. And the uplink bandwidth is uniformly-distributed ranging from 80Mbps to 240Mbps, while the downlink bandwidth is set to 1Gbps for each node.

In order to minimize the communication overhead of gradient broadcast, we fix the ratio between n/c^2 to 4/1, to ensure exactly one local gradient and neighbors gradient is aggregated in each consensus step.

B. Performance Evaluation

We compare our prototype with another blockchain-based SGD framework LearningChain without the presence of malicious node. We first compare the gradient storage overhead of the two frameworks. We then compare the iteration time measured by the time used to broadcast a block that contains the necessary information for nodes to progress in D-SGD.

As Fig. 4 shows, gradient storage overhead for PIRATE nodes are constant as iteration progresses, while LearningChains storage has a linear growth. In each iteration, PIRATE stores only the leaders gradient, the neighbor committees gradient and the local gradient itself has computed. In LearningChain, nodes are required to store the history of all leader-announced gradients, and the local gradients broadcasted by all nodes.

Fig. 4 shows that PIRATE outperforms LearningChain on iteration time, i.e., the time used for each node to get a model update. The major cost in each iteration is the broadcast of gradients. Broadcasting to a larger number of nodes significantly increase the cost. PIRATE shows a superior performance in terms of iteration time for each committee. This is because nodes are required to broadcast to only c members, meanwhile, consensus decisions are reached in a parallel manner.

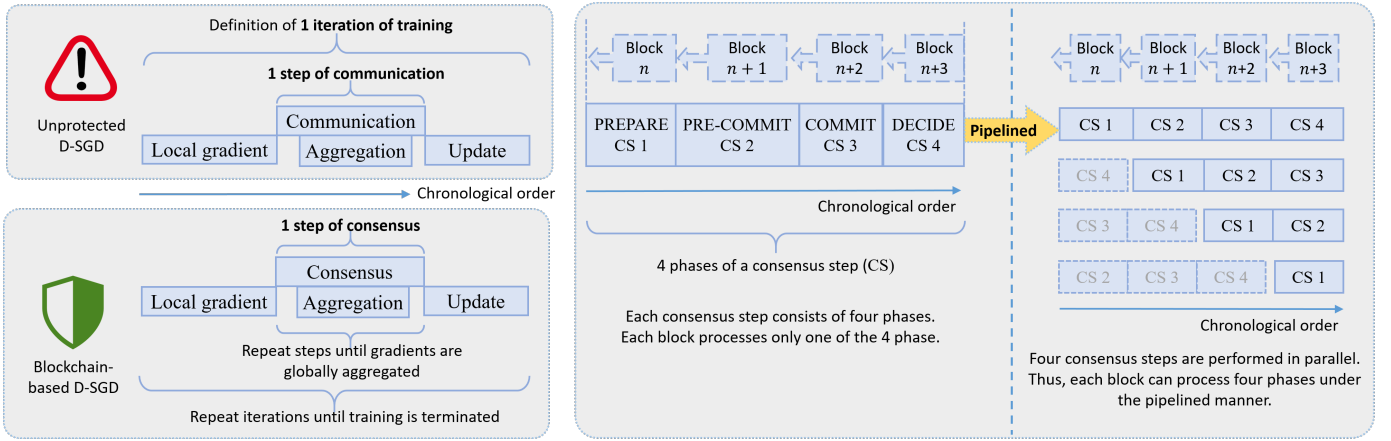


Fig. 3. We define the process of finishing a partial aggregation a *step*, the process of finishing a model update an *iteration*. A consensus step (CS) has 4 phases, each driven by leader. The protocol can be pipelined for performance enhancement, ideally executing 1 aggregation each block in average.

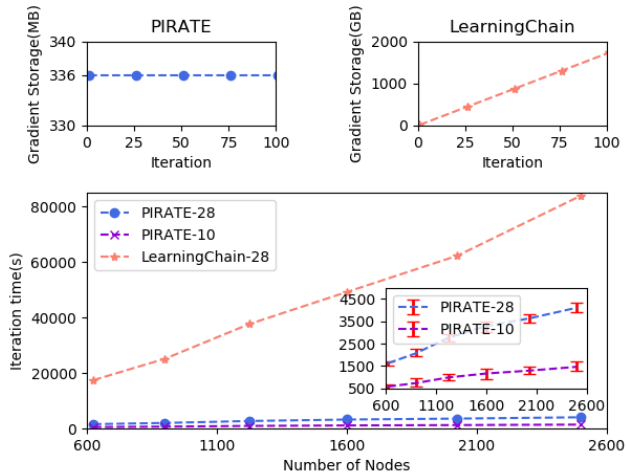


Fig. 4. The left-top figure: Gradient Storage of PIRATE vs. iteration with single gradient size of 28MB. The right-top figure: Gradient Storage of LearningChain vs. iteration with the single-gradient size of 28MB. The Bottom figure: Iteration time of PIRATE and LearningChain vs. number of nodes with the single-gradient sizes of 28MB and 10MB.

VI. OPEN ISSUES

The open issues are envisioned as follows.

- **Decentralized Permission Control.** Other than join/leave Sybil attacks, candidates having inferior computation ability, bad historical credit scores and unstable network conditions can undermine the efficiency of distributed learning. Without a centralized permission control, reliability assessment is challenging, especially for realtime states. However, for a decentralized system, latency induced by communication and verification inevitably affects timeliness. We can utilize an anchor chain to manage non-realtime states of devices to achieve further decentralization. The anchor chain is maintained by all candidates. Permissions to join shard chains, (i.e. learning committees) are jointly managed by the anchor Chain and a real-time assessment component.
- **Protection Against Model Poisoning Attack.** For fed-

erated learning, model-poisoning attack is a challenging issue in terms of safety. The attack can be successful even for a *highly constrained* byzantine node [15]. By exploiting the non-i.i.d property of data shards, byzantine attackers can send poisoned local updates that do not hurt convergence. Such harmful local updates can still affect the global model that triggers misclassifying. The attack method is also “sneaky” enough to bypass accuracy check by a central server. In a decentralized environment, where nodes are less constrained in terms of communicating, computing and validating, model poisoning attack can be even more threatening.

- **Incentive Mechanism.** In a decentralized setting, each node shares greater responsibility than the nodes in a centralized setting. This is because they are constantly required to conduct actions, like validations and broadcasting, which seem non-rewarding through the perspective of an individual. Since PIRATE is built based on a blockchain system, incentive mechanisms of participating can be naturally guaranteed.

VII. CONCLUSION AND FUTURE WORK

To guarantee the high availability of distributed learning in 5G era, a distributed-learning framework with high efficiency, decentralization and byzantine-resiliency is in urgent need. To fill this gap, we propose PIRATE, a byzantine-resilient D-SGD framework under the decentralized settings. Utilizing a sharding-based blockchain protocol, model parameters and gradient aggregations can be well protected. We then evaluate the feasibility of the PIRATE in a scenario where the large scale participants are with large-sized local models. Simulation results show that the proposed PIRATE is more efficient than an existing LearningChain solution in terms of both communication time and storage complexity.

We will further analyze the robustness and resiliency of PIRATE by conducting extensive experiments. Currently, we are developing a byzantine-resilient aggregation algorithm dedicated for PIRATE, by considering the efficiency of computation and communication.

VIII. ACKNOWLEDGEMENT

This work is partially supported by National Natural Science Foundation of China (61902445, 61872310), partially by Fundamental Research Funds for the Central Universities of China under grant No. 19lgy222, and partially by Natural Science Foundation of Guangdong Province of China under Grant 2019A1515011798.

REFERENCES

- [1] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and distributed computation: numerical methods*. Prentice hall Englewood Cliffs, NJ, 1989, vol. 23.
 - [2] X. Lian, W. Zhang, C. Zhang, and J. Liu, "Asynchronous decentralized parallel stochastic gradient descent," *arXiv preprint arXiv:1710.06952*, 2017.
 - [3] X. Lian, C. Zhang, H. Zhang, C.-J. Hsieh, W. Zhang, and J. Liu, "Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent," in *Advances in Neural Information Processing Systems*, 2017, pp. 5330–5340.
 - [4] L. Lamport, "Time, clocks, and the ordering of events in a distributed system," *Commun. ACM*, vol. 21, no. 7, pp. 558–565, Jul. 1978. [Online]. Available: <http://doi.acm.org/10.1145/359545.359563>
 - [5] P. Blanchard, E. M. El Mhamdi, R. Guerraoui, and J. Stainer, "Machine learning with adversaries: Byzantine tolerant gradient descent," in *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds. Curran Associates, Inc., 2017, pp. 119–129.
 - [6] Y. Chen, L. Su, and J. Xu, "Distributed statistical machine learning in adversarial settings: Byzantine gradient descent," *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 1, no. 2, pp. 44:1–44:25, Dec. 2017. [Online]. Available: <http://doi.acm.org/10.1145/3154503>
 - [7] S. Li, Y. Cheng, Y. Liu, W. Wang, and T. Chen, "Abnormal client behavior detection in federated learning," *arXiv preprint arXiv:1910.09933*, 2019.
 - [8] M. Yin, D. Malkhi, M. K. Reiter, G. G. Gueta, and I. Abraham, "Hotstuff: Bft consensus in the lens of blockchain," 2018.
 - [9] M. Zamani, M. Movahedi, and M. Raykova, "Rapidchain: Scaling blockchain via full sharding," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2018, pp. 931–948.
 - [10] M. Zinkevich, M. Weimer, L. Li, and A. J. Smola, "Parallelized stochastic gradient descent," in *Advances in neural information processing systems*, 2010, pp. 2595–2603.
 - [11] X. Chen, J. Ji, C. Luo, W. Liao, and P. Li, "When machine learning meets blockchain: A decentralized, privacy-preserving and secure design," 12 2018, pp. 1178–1187.
 - [12] B. T. Polyak and A. B. Juditsky, "Acceleration of stochastic approximation by averaging," *SIAM Journal on Control and Optimization*, vol. 30, no. 4, pp. 838–855, 1992.
 - [13] J. Ji, X. Chen, Q. Wang, L. Yu, and P. Li, "Learning to learn gradient aggregation by gradient descent," in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*. International Joint Conferences on Artificial Intelligence Organization, 7 2019, pp. 2614–2620. [Online]. Available: <https://doi.org/10.24963/ijcai.2019/363>
 - [14] E. Kokoris-Kogias, P. Jovanovic, L. Gasser, N. Gailly, E. Syta, and B. Ford, "Omniledger: A secure, scale-out, decentralized ledger via sharding," in *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2018, pp. 583–598.
 - [15] A. N. Bhagoji, S. Chakraborty, P. Mittal, and S. Calo, "Analyzing federated learning through an adversarial lens," *arXiv preprint arXiv:1811.12470*, 2018.
- Huawei Huang** (M'16) (corresponding author, huanghw28@mail.sysu.edu.cn) received his Ph.D in Computer Science and Engineering from the University of Aizu, Japan. He is currently an associate professor with the School of Data and Computer Science, Sun Yat-Sen University, China. His research interests mainly include distributed learning and blockchains optimization. He serves as a visiting scholar with the Hong Kong Polytechnic University (2017-2018); a post-doctoral research fellow of JSPS (2016-2018); an assistant professor with Kyoto University, Japan (2018-2019). He is a member of ACM.
- Wuhui Chen** (chenwuh@mail.sysu.edu.cn) is an associate professor in Sun Yat-sen University, China. He received his bachelor's degree from Northeast University, China, in 2008. He received his master's and Ph.D. degrees from University of Aizu, Japan, in 2011 and 2014, respectively. From 2014 to 2016, he was a JSPS research fellow in Japan. From 2016 to 2017, he was a researcher in University of Aizu, Japan. His research interests include Edge/Cloud Computing, Cloud Robotics, and Blockchain.
- Zibin Zheng** (SM16) (zhzibin@mail.sysu.edu.cn) received the Ph.D. degree from the Chinese University of Hong Kong, Hong Kong, in 2012. He is a Professor with the School of Data and Computer Science, Sun Yat-sen University, Guangzhou, China. His current research interests include service computing and cloud computing. Prof. Zheng was a recipient of the Outstanding Ph.D. Dissertation Award of the Chinese University of Hong Kong in 2012, the ACM SIGSOFT Distinguished Paper Award at ICSE in 2010, the Best Student Paper Award at ICWS2010, and the IBM Ph.D. Fellowship Award in 2010. He served as a PC member for IEEE CLOUD, ICWS, SCC, ICSOC, and SOSE.
- Song Guo** (M'02-SM'11-F'20) (song.guo@polyu.edu.hk) received his Ph.D. degree in computer science from the University of Ottawa, Canada. He is currently a full professor at Department of Computing, The Hong Kong Polytechnic University. His research interests mainly include cloud and green computing, big data, and cyber-physical systems. He serves as an Editor of several journals, including IEEE TPDS, TETC, TGCN, and IEEE Communications Magazine. He is a senior member of IEEE and ACM, and an IEEE Communications Society Distinguished Lecturer.

Sicong Zhou (chowsch2@mail2.sysu.edu.cn) is currently with the School of Data and Computer Science, Sun Yat-Sen University, China. His research interests mainly include distributed learning and blockchains optimization.