

Fantastyc: Blockchain-based Federated Learning Made Secure and Practical

William Boitier, Antonella Del Pozzo, Álvaro García-Pérez, Stephane Gazut, Pierre Jobic, Alexis Lemaire, Erwan Mahe, Aurelien Mayoue, Maxence Perion, Tuanir Franca Rezende, Deepika Singh, Sara Tucci-Piergiovanni
Université Paris-Saclay, CEA, List F-91120, Palaiseau, France
 {name.surname@cea.fr}

Abstract—Federated Learning is a decentralized framework that enables multiple clients to collaboratively train a machine learning model under the orchestration of a central server without sharing their local data. The centrality of this framework represents a point of failure which is addressed in literature by blockchain-based federated learning approaches. While ensuring a fully-decentralized solution with traceability, such approaches still face several challenges about integrity, confidentiality and scalability to be practically deployed. In this paper we propose Fantastyc, a solution designed to address these challenges that have been never met together in the state of the art.

Index Terms—Distributed Systems, Blockchain, Federated learning, Fault tolerance

I. INTRODUCTION

Federated Learning (FL) has emerged as an important paradigm in modern large-scale machine learning [1]. Unlike in traditional centralized learning where models are trained using large datasets stored in a central server, in federated setting, the training data remains distributed over a large number of clients, which may be phones, network sensors, hospitals or companies. A federated model is then trained under the supervision of a central server by aggregating only ephemeral locally-computed updates, thereby ensuring a basic level of privacy by design. However, FL still faces key challenges [2] to be fully accepted in both institutional and industrial sectors. First, the training process should provide guarantees to ensure its robustness to malicious actors who try to degrade the accuracy of the model (integrity) or retrieve sensitive information about isolated training data (privacy). Second, the federated process should not collapse even if there is a large number of participants to the training step (scalability). Third, the server which orchestrates the federated process is a central player which potentially represents a single point of failure. While large companies or organizations can play this role in some application scenarios, a reliable and powerful central server may not always be available or desirable in more collaborative learning scenarios.

In the literature, a growing research effort is dedicated to propose fully decentralized learning through blockchain-based Federated Learning (BC-based FL) [3]–[10]. Blockchains offer a secure, shared ledger among users, where information is stored and verified by a peer-to-peer network of nodes. The

peer-to-peer network adds new information to the ledger thanks to a consensus mechanism tolerant to malicious participants. As a result, blockchains provide a means to eliminate the need for trust in a central server. Indeed, the blockchain could be used to validate and store updates and models to secure the FL process. Additionally, blockchains incorporate native incentive mechanisms. These mechanisms encourage users to participate honestly in the system, leveraging the audit trail of users participation registered in the blockchain.

To the best of our knowledge, however, current BC-based FL solutions do not address all the challenges as a whole, which makes them not fully secure or practical. A fundamental issue lies in the scalability of blockchains, particularly concerning the latency required to register transactions, a factor influenced by the size of the registered data and the number of clients. In this paper, our objective is to develop a BC-based FL system that ensures both model integrity with proven guarantees, addresses privacy concerns, and maintains practicality even with a large number of clients. Our goal is to implement realistic model training scenarios, akin to those used in classical benchmarking for FL image classification tasks (such as Federated Extended MNIST), involving thousands of clients and models of several megabytes in size.

To this end, in our design, we leverage blockchain technology with a specific emphasis on reducing interactions with it, all the while capitalizing on its inherent advantages. To enhance efficiency, especially given the resource-intensive nature of validating updates and models, we offload the validation process – meant to be handled by the blockchain – to a dedicated set of servers. These servers handle aggregation and generate validity proofs, ensuring the integrity of the FL process. Clients can efficiently verify model integrity by examining these validity proofs. Moreover, we tackle the challenge of managing large-sized data that cannot be directly stored on-chain. Our solution involves employing a fault-tolerant key-value store for registering client updates and models as follows. In each round of the FL process, clients updates are gathered by servers, stored in the storage and cryptographically anchored in the blockchain alongside validity proofs, to record client participation. By reading the updates anchored in the blockchain, each server retrieves an identical set of client updates before aggregating them and produce the next model. Fig. 1 illustrates these ideas, highlighting the

This work is funded by The Fantastyc Carnot project.

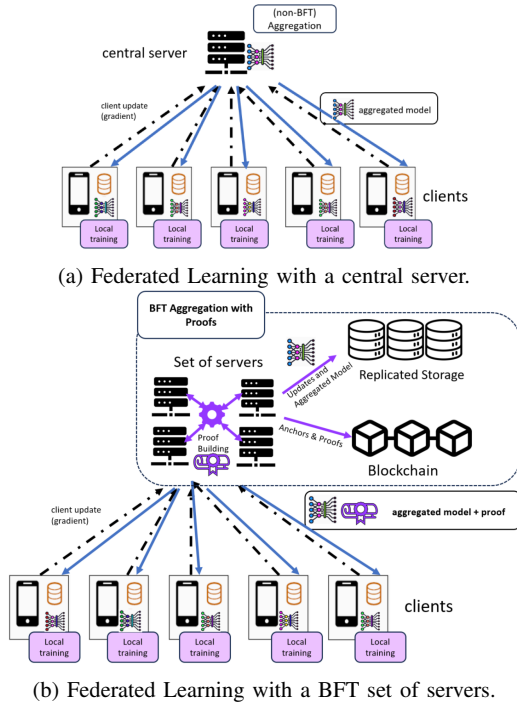


Fig. 1: Difference between Federated Learning (a) and our solution (b).

additional components we use in our architecture with respect to a centralized, non-Byzantine tolerant solution. The figure focuses on the exchange of data in a single FL iteration. In the centralized version, the clients performing the local training send their updates (the gradients) to the server for that iteration. The server performs aggregation and sends back the aggregated model to clients. In our solution, all participants register in the system via the blockchain before engaging in the learning process. The aggregation is then handled by a set of participants who register to play the role of servers (where a minority can be Byzantine). Servers perform robust aggregation [11] to filter out Byzantine client contributions while engaging in a protocol to certify the validity of the obtained aggregated model to send back to clients. At the conclusion of this protocol, a validity proof is generated. Both clients' updates and aggregated models are stored in a replicated storage, while their anchors (i.e., cryptographic hashes) are stored in the blockchain, along with the proofs. Interestingly, servers also write via the proof the contributions from the clients they retained for the aggregation, establishing a trail of client contributions to be exploited for rewarding.

We also employ a light-weight confidentiality technique to face the risk of data leakage. The method makes it difficult for attackers to recover the personal data from users' dataset. Notably, this module does not add any communication overhead while only suffering small accuracy loss.

Overall, Fantastyc achieves the following properties:

- **Robustness:** Our solution exhibits robustness against

Byzantine attacks [12] from both clients and servers, safeguarding the accuracy of the federated model against adversarial attempts to compromise its integrity.

- **Enhanced Byzantine Tolerance:** Our solution guarantees model integrity and availability by assigning the blockchain the sole responsibility of managing clients' participation and contributions. This approach allows the set of servers to maintain an optimized fault-tolerant threshold, requiring a majority of honest servers rather than the conventional two-thirds [12], making it a more secure option.
- **Efficient Resource Use and Proofs:** Our solution optimizes blockchain usage by storing updates and models within an external distributed storage connected to the blockchain. In this setup, the blockchain exclusively holds a cryptographic fingerprint of the data (the anchor), accompanied by a validity proof. This proof certifies the availability and integrity of the data within the storage. This approach ensures that clients can always retrieve computed models (availability) providing them with a guarantee of retrieving a trustworthy model (integrity). It is important to note that prevalent architectures in the blockchain space, such as those utilizing external storage systems like IPFS, lack computational integrity even when they offer mechanisms for making redundant copies (proof-of-replication [13]). Indeed, the primary goal of these architectures is to offer tamper-proof guarantees: clients can only ascertain whether someone has tampered with their own data, without ensuring that the data are the output of correct computations.
- **Confidentiality:** Our solution limits the risks of information leakage about the client's data with little loss of accuracy and without adding any communication overhead. Each client has to complete a simple data processing routine before participating in the training process that consists in mixing his own data samples and applying a random sign-flipping pattern.
- **Scalability:** Our solution confines the model aggregation to a small number of servers – outside the blockchain. This ensures a blockchain latency independent of the model size. Additionally, a verifiable random selection process allows us to efficiently poll a subset of clients in different rounds, enabling scalability to accommodate the participation of a large number of clients. Notably, we demonstrate the capability to handle realistic FL processes in a geo-replicated deployment, for the first time for a blockchain-based FL solution.

The rest of the paper is organized as follows. Section II introduces FL and Blockchain. Section III discusses the current state of the art of BC-based FL solutions and their limitations. Section IV delves into design choices of the proposed solution that enhance robustness, Byzantine tolerance, efficient resource utilization and confidentiality. Section V presents algorithms and correctness arguments. Section VI discusses the implementation and evaluation of the solution in a geo-

replicated distributed system. Section VII concludes the work.

II. BACKGROUND AND PRELIMINARIES

A. Federated Learning

In Federated learning (FL), we consider a set of C clients with n_c data that participate in the training of a collaborative model $w \in \mathbb{R}^d$ under the orchestration of a central server but without sharing their local data due to privacy issues. More formally, if we define $f_c : \mathbb{R}^d \rightarrow \mathbb{R}$ as the local objective function for a client c , the goal is to minimize the overall objective function f considering the isolated datasets:

$$\min_w f(w) = \min_w \sum_{c=1}^C \frac{n_c}{N} f_c(w) \quad \text{where } N = \sum_{c=1}^C n_c \quad (1)$$

The FL process consists of a succession of rounds. At each round t , the server selects a subset C_t of $K \leq C$ clients that take part in training and sends them the current global model w_t . Then, each selected client k uses its dataset to minimize his proper objective function f_k with stochastic gradient descent (SGD), a number of epochs E and a mini-batch size B to update parameters w_t^k . Each client k then sends updates $g_t^k = w_t - w_t^k$ to the server who averages the received parameters to obtain a new global model w_{t+1} . The aggregation rule depicted by Eq.2 is known as FedAvg.

$$w_{t+1} = w_t - \sum_{k \in C_t} \frac{n_k}{n} g_t^k \quad \text{where } n = \sum_{k \in C_t} n_k \quad (2)$$

B. Blockchain

Blockchain is a decentralized immutable ledger, which consists of a sequence of linked blocks, each containing a set of transactions. Transactions allow blockchain users to immutably record information, which constitutes a common knowledge shared among blockchain users. In particular, this information is transparent, tamper-resistant, and totally ordered. At the core of each blockchain, the consensus mechanism plays a pivotal role in appending new blocks without breaking the chain structure. In our context, we focus on blockchains based on BFT-Consensus, such as Cosmos [14] and Tezos [15], which provide immediate finality, ensuring that once a block is added, it cannot be revoked.

Most blockchains, like Tezos and Ethereum [16], feature smart contracts [17] for programmability. Smart contracts are self-executing scripts tied to specific blockchains, deployed and executed through transactions, placing both the generic applications and the blockchain logic on the same layer. In contrast, Cosmos introduces the ABCI (Application BlockChain Interface), separating the blockchain from application logic. With Cosmos ABCI, the blockchain only records transactions, while applications independently process and update their state, enhancing modularity and deployment flexibility.

While blockchain offers security and decentralization, it also faces scalability limitations, particularly in transactions per second and communication latency. For instance, in Tendermint [18] (now CometBFT [19]), the underlying consensus mechanism for Cosmos, the average throughput is

approximately 430 transactions per second with a latency of 3.45 seconds in the presence of 128 correct nodes [20]. These constraints stem from the consensus mechanism and the distributed nature of blockchain, resulting in substantial memory costs as each node maintains an updated replica. Therefore, blockchain usage must be approached with frugality, emphasizing the efficient and careful use of resources to ensure sustainability and optimize performance.

III. STATE OF THE ART

BlockFL [5] is a pioneering work that introduces a conceptual architecture where clients record their local updates in a public (proof-of-work) blockchain and perform local aggregation as soon as a sufficient number of local updates are written on the blockchain. This approach might be suited for small-sized models that can fit within blocks, even if the paper does not present a concrete solution and evaluation. Furthermore, the architecture relies on each client being associated with a single miner for sending and downloading blocks. However, a malicious miner could potentially pose a threat by refusing to include a client's local update in the blockchain or by remaining silent, causing availability issues. Finally, in the proposed architecture, miners are expected to validate the correctness of client updates by confirming that each update has been executed across the entire client dataset, because a client might send fake updates without incurring the cost of training. This verification relies on both miners having knowledge of each client's data size and the presence of a trusted execution environment at each client, ensuring the integrity of the clients' contributions.

FLChain [4] shares similarities with BlockFL in its use of blockchains to record local updates. However, it diverges from BlockFL by introducing an additional cryptographic layer to encrypt the local updates stored in the blockchain. Aggregation is then executed by (a subset of) blockchain nodes on these encrypted updates, and the resulting aggregated model is stored in the blockchain. Subsequently, aggregated models are collectively decrypted by clients, with each client holding the ability to decrypt only a share of each aggregated model. The approach also includes incentive mechanisms designed to motivate participants to provide uncorrupted shares. Much like BlockFL, this approach does not explicitly address concerns related to Byzantine clients conducting poisoning attacks. Also, due to the substantial utilization of the blockchain and the collaborative decrypting algorithm, it is better suited for smaller-sized models and a few number of clients.

VFChain [10] shares architectural concepts similar to FLChain and BlockFL. However, it distinguishes itself by providing finely-tuned mechanisms for the dynamic election of an aggregation committee. Nonetheless, the storage of aggregated models in the blockchain is maintained, and experimental evaluations have been conducted with a participant count not exceeding 10. It is important to note that in this system, clients are regarded as honest but curious, while Byzantine behavior is attributed solely to aggregators, with a tolerance for up to one-third of them. In BFLC [21], a similar architecture with a

committee is used for efficiency, whereas a robust aggregation function is implemented by committee members to defend the system from Byzantine clients.

BFEL [6], in comparison to the aforementioned approaches, introduces a more sophisticated architectural design. It combines a public blockchain with private blockchains, with each private blockchain dedicated to a specific FL task. This approach optimizes storage efficiency in multi-task scenarios and enhances privacy. Additionally, to alleviate communication overhead resulting from the exchange of large models, BFEL applies a gradient compression method. In this system, a committee of verifiers plays a crucial role in filtering out potentially poisoned models submitted by clients. However, it is worth noting that the verification method relies on each verifier using a testing dataset to evaluate the updates, instead of employing more robust aggregation techniques that might offer greater reliability. Experimental evaluations have been conducted with two parallel images classification tasks, in a local machine, without exceeding 22 clients.

BlockFLA [7] also introduces a hybrid architecture, where a private blockchain is integrated with a public blockchain. In this setup, model updates are entirely stored within the private blockchain and are subsequently linked to the public blockchain for accountability via a smart contract. While the paper mentions the existence of a verification function within the smart contract, it does not provide specific details regarding its functioning.

Finally, Biscotti [8] refines and elaborates on principles stipulated by FLChain [4] by introducing a three-role structure: noisers create unbiased noise for clients to mask local updates, verifiers endorse masked updates through robust filtering, and aggregators confidentially aggregate the endorsed updates. Aggregators work with encrypted data to ensure privacy. The encrypted aggregated model is stored on the blockchain within commitments to endorsed updates. Unlike our approach, Biscotti, by storing the model on the blockchain, can only handle small models (models with thousands of parameters as opposed to millions in our approach). Additionally, there's uncertainty about the purpose of the noise particularly considering the authors' statement that they do not utilize it at the end due to significant utility loss. Without noise, verifiers access unmasked client updates, posing a privacy risk if even a single verifier is compromised. Therefore, finding the right trade-offs between noise and utility is essential to ensure privacy while using this architecture.

It is also worth to mention some other research efforts, such as those in [3] and [9], that have delved into solutions within the Ethereum blockchain framework to replace the model-aggregating server with smart contracts. However, deploying these smart contracts on public blockchain systems can incur significant costs, mainly due to miner fees. These costs are associated with the computational workload and the substantial storage requirements resulting from full replication on every blockchain node. As a result, these approaches are typically more suitable for handling simpler models and tasks.

Fantastyc distinguishes itself from existing state-of-the-art

solutions by minimizing blockchain utilization through a novel off-chain Byzantine-tolerant protocol leveraging distributed storage servers, ensuring the integrity and availability of off-chain data and computations without depleting system resources. Moreover, our construction seamlessly adapts its protocol to limit the risk of information leakage without introducing any additional overhead while integrating secure guarantees to remain robust in the presence of both Byzantine clients and servers. Notably, it offers an enhanced level of Byzantine tolerance compared to existing solutions, requiring a majority of honest servers rather than the conventional two-thirds, making it a more secure option. Interesting, because of its modular architecture, our solution offers as well the flexibility to be deployed as a private, public, or hybrid blockchain while allowing users to initialize learning tasks with any type of aggregation, selection, and rewarding function.

IV. DESIGN CHOICES AND KEY IDEAS

A. Robustness and Enhanced Byzantine Tolerance

Every distributed process is vulnerable to misbehaving nodes that can disrupt the learning procedure. Indeed, arbitrary failures due to hardware defaults may occur or some actors may also try to deliberately corrupt the model. Following the terminology of distributed computing, these nodes are called Byzantine [12]. We refer to the other nodes as correct or honest. The robustness of a BC-based FL process is thus ensured when model integrity is preserved despite the presence of both Byzantine clients and servers.

1) *Byzantine Clients*: In a BC-based FL process, the servers cannot check whether the updates they receive from the clients are coherent with licit training data, which makes the learning procedure vulnerable to update manipulation. Considering the standard aggregation rule FedAvg (eq.2), it is very easy for Byzantine clients to hinder the model's convergence by sending erroneous updates such as $-\lambda_{boost}g_t^k$ instead of g_t^k (where $\lambda_{boost} \in [1, +\infty[$ is used to boost the strength of the attack). To mitigate the negative impact of Byzantine nodes, it is common to use a non-linear robust aggregator such as the coordinate-wise median [11]. Considering f_c Byzantine clients, this Byzantine-tolerant aggregation rule requires at least $2f_c + 1$ participants to each round of training.

2) *Byzantine Servers*: In the context of ensuring robustness against Byzantine attacks, a critical challenge lies in guaranteeing that clients do not retrieve poisoned models from potentially malicious servers. This challenge can be faced harnessing the power of blockchain technology to ensure the integrity of the models collected by clients. To illustrate this approach, let us envision the blockchain as a ledger – an immutable, ordered list of records accessible for both clients and servers to write and read. We present two distinct strategies:

(i) In the first strategy, each server collects a (possibly different) set of client updates, performs robust aggregation, and subsequently records the resulting model on the blockchain. It is worth noting that due to network delays and message losses, each server may receive a different subset of updates.

In this scenario, clients retrieve from the ledger, in each round, the aggregated models from a count of servers not exceeding a specified fault threshold, denoted as $n_s - f_s$, where f_s accounts for potentially silent or slow servers. Clients then engage in robust aggregation to derive the model for the next round.

(ii) In the second strategy, clients take on the responsibility of directly writing their updates to the blockchain. For each round, servers await the writing in the ledger of updates from a predefined number of clients. In this scenario read updates are the same set at each server. Subsequently, servers engage in robust aggregation and synchronize to cast their votes in favor of models that fall within an epsilon range¹. When a model accumulates at least $f_s + 1$ votes, it signifies that at least one honest server has endorsed it, rendering it safe for use by any client. A client retrieves then a model with $f_s + 1$ server votes.

Assuming the client count (n_c) surpasses that of servers (n_s), the first strategy is better in terms of ledger storage space used. However, it necessitates a higher fault tolerance threshold, requiring $n_s = 3f_s + 1$ (to account for slow servers in the round and malicious speaking servers). Additionally, it mandates clients to perform robust aggregation on their end. The second approach requires a fault tolerance threshold of only $n_s = 2f_s + 1$ (to account for slow servers in the round and to guarantee that at least $f_s + 1$ votes are cast for a model) and more importantly it spares clients from the computational overhead of robust aggregation. We opt then for the second approach to avoid overload clients with robust aggregation computation and because the better tolerance threshold. To maintain practicality, we also prevent clients from directly writing to the blockchain. Instead, we propose a byzantine tolerant protocol where servers gather updates and anchor the collected set in the blockchain with the help of an external fault-tolerant distributed storage, as detailed next.

B. Efficient Resource Use by Proof of Availability & Integrity and Optimally Replicated Fault-tolerant Key-value Store

To efficiently utilize the blockchain, we propose a Byzantine-tolerant protocol that allows for the reliable storage of both client updates and aggregated models in a fault-tolerant key-value store. Because we need to achieve consensus only on the client updates set, the only information recorded in the blockchain is the cryptographic hook of client updates contained in the store, referred to as Proof of Availability & Integrity (PoA&I). The PoA&I is generated each time a value (update or model) is consistently and reliably replicated at enough key-value store replicas. This proof guarantees clients and servers to always retrieve the value (*availability*) and ensures that the value is the result of a correct computation (*computational integrity*) – for instance the correct model of a robust aggregation. In the blockchain, consensus is reached

¹ It is noteworthy that, owing to the computational intricacies of floating-point calculations, models generated from identical inputs may exhibit slight differences. Empirical observations, using models with 500,000 parameters, have shown disparities up to three decimal places, underscoring the level of precision involved.

on the set of proofs uniquely associated with correct client updates retrievable by servers in the store.

At high level, the protocol to generate the PoA&I works as follows: as soon as a server gathers a value v for which a PoA&I is required, the server will verify if the value is valid. If valid, it stores the pair $(hash(v), v)$ in the key-value store where the hash of v is used as key, and then broadcasts a signed message with the hash of the valid value, called the *hash-key*, to the other servers. As soon as a server receives a $f_s + 1$ quorum of signed valid hash-keys from other servers, it produces a PoA&I². In the first step of the protocol, the verification performed to validate values depends on the type of value: in the case v is a client update, the server verifies if the update comes from the set of clients selected for the current round; if v is a model aggregated by other servers, then the value is considered valid if the received model and the locally computed model are not significantly different, typically within an epsilon distance.

To ensure persistence and counteract potential attackers attempting to erase or modify off-chain data from the storage, the solution mandates the hash-key always to be the cryptographic hash of the corresponding value, if it is not the case, the store reject the value as it is not possible to associate to a key a value whose hash differs from the key itself.

1) *One blockchain transaction per round*: To minimize the number of blockchain transactions, instead of recording each client update's PoA&I on the blockchain immediately, servers wait to generate a sufficient number of client updates' PoA&I. When this threshold is reached, they write a set of updates' PoA&I as a candidate input for the next robust aggregation on the blockchain by sending a transaction. Since transactions on the blockchain are totally ordered, the first set of updates, each with a valid PoA&I, recorded on-chain is the one utilized for robust aggregation by all correct servers. It is important to note that after agreeing on a set of PoA&Is, servers perform robust aggregation and generate a PoA&I for valid models computed and stored. The PoA&I of the aggregated model is then sent to clients, who can retrieve models directly from the store without interacting with the blockchain.

2) *Optimal replication*: It is worth noting that our approach decouples the ordering (blockchain side), the integrity (server side) and availability (storage side) of the information. This decoupling allows to implement those properties employing the optimal number of nodes. As discussed in Section VI-C, the number of servers is at least $2f_s + 1$. Concerning the key-value store, $f_r + 1$ replicas suffices, where f_r is the upper bound on the number of Byzantine replicas and each replica is replicating the content of the key-value store. Thanks to the PoA&I associated to the hash-key, a client querying the store can assess the validity of a returned value by checking if the retrieved value, once hashed, is equal to the required hash-key. As a consequence, to get the value, it is sufficient to

² PoA&I is similar in spirit to the Proof of Availability and Retrieval [22], which, however, does not target computational integrity. Indeed, [22] describes a mechanism to produce a proof of availability to implement a secure pull communication mechanism for blocks inside blockchains.

get a correct answer from at most one correct storage replica, which exists, given the existence of the PoA&I.

Let us stress that, only updates associated with a PoA&I are considered part of a valid updates set. This ensures that any set recorded in the blockchain, if valid, contains updates available in the storage and produced by selected clients. This mechanism provides certainty to a server; if they miss some updates necessary to aggregate the next model, they are assured they can retrieve them in the storage. A similar approach applies when servers interact with clients. Indeed, once servers compute the next aggregated model, they store it in the distributed storage, and create a proof before sending it directly to clients. This approach provides clients with verifiable evidence of integrity and availability, mitigating concerns related to computational manipulation. Detailed workflows and algorithms can be found in Section V.

C. Confidentiality

By design, FL intrinsically protects the data stored on each device by sharing model updates (*i.e.* gradients) instead of the original data. While sharing gradients was thought to leak significantly less information about the clients compared to the raw data, recent papers developed a “gradient inversion attack” by which honest-but-curious servers are able to reconstruct parts of the client’s private data [23], [24]. Homomorphic encryption is the gold standard for a privacy-preserving solution but its large computational demand still makes it not compatible with scalability despite recent improvements [25]. Based on literature, we propose countermeasures to limit the risks of information leakage. We advise clients who want to participate in the training process to apply InstaHide [26] data augmentation to their dataset. In a nutshell, InstaHide has two key components. The first step consists in building virtual training samples in the following way: to encode an example $x \in \mathbb{R}^d$, a client picks $s - 1$ other samples from his private dataset and s random nonnegative coefficients $\{\lambda_j\}_{j=1}^s$ that sum to 1. A composite (one-hot encoding) label is also created using the same set of mixing coefficients:

$$x = \lambda_1 x + \sum_{j=2}^s \lambda_j x_j \quad \text{and} \quad y = \lambda_1 y + \sum_{j=2}^s \lambda_j y_j \quad (3)$$

As this step permits to train a neural network on linear combinations of examples instead of raw data, an attacker could only retrieve the virtual examples. The capacity of inferring the raw data from the virtual ones requires the knowledge of the mixing coefficients but this information seems hard to obtain in a realistic scenario. However, InstaHide adds an another layer of security that consists in picking a random sign-flipping pattern $\sigma \in \{-1, 1\}^d$ and outputting the encryption $\tilde{x} = \sigma \circ (\lambda_1 x + \sum_{j=2}^s \lambda_j x_j)$ where \circ is coordinate-wise multiplication of vectors.

V. SYSTEM DETAILED DESIGN

In this section we first introduce the system model and a formalisation of the PoA&I, before presenting the workflow of the FL process, which includes algorithms to build the PoA&I

(Section V-A) and methods to manage the non-determinism of the computations (Section V-A2).

1) *System model*: Our system is composed of a finite set of processes³: clients, servers, storage replicas, Model Owner (MO), all having access to a blockchain. The blockchain, an immutable, ordered list of records, allows processes to communicate and share a common knowledge upon the information securely recorded on it. Clients, servers and storage replicas can also directly communicate. More formally, we assume that the system comprises $2f_s + 1$ servers S and $f_r + 1$ storage replicas R , where f_s and f_r represent the upper bounds on Byzantine faulty processes (processes that exhibit an arbitrary behaviour). We assume model owners to be correct. Additionally, we rely on the reliable communication assumption that if a correct process p sends a message m to another correct process q , then eventually q receives m . Moreover, if a correct process sends a transaction t to the blockchain, then t is eventually totally ordered in the blockchain. All messages are signed by the sender. We also consider that servers use the same pseudo-random function `select_clients`, that takes as input the blockchain (which contains the list of clients) to verifiably select the client set for the next round.

2) *PoA&I*: We refer to the PoA&I as $P(\text{TAG}, v)$, a set of cryptographic hashes of v signed by $f_s + 1$ servers. Let `locally_valid`(TAG, v) be a function that returns true if v passes the computational integrity check relative to TAG. A correct server signs $h(v)$ if `locally_valid`(TAG, v) holds and after storing v on the key-value store. If a process p receives $P(\text{TAG}, v)$, then `locally_valid`(TAG, v) holds for at least one correct server⁴, and eventually, p gets v from the key-value store.

A. Solution workflow and algorithms

We outline the workflow for an FL task, denoted fl_i , from the moment the model owner inputs the model for training to the moment the model owner retrieves the trained model. Algorithms 1-2 detail the core part of the process (pseudo-code does not explicit any particular fl_i , as it can be easily generalized to handle multiple FL tasks in parallel). Algorithm 1 details the server pseudo-code, as described from step 4 to step 7 in the following workflow, and how the PoA&I works. Algorithm 2 specifies how processes (clients, servers or model owner) interact with the secure distributed key-value store to get a value once they collect its PoA&I. Due to its simplicity, we omit the pseudo-code for the replica storage process. Instead, we describe its behavior in comments within Algorithms 1 and 2 when needed. *Workflow steps*:

- 1) The process starts with the model owner that sends a transaction to the blockchain to run a new FL task fl_i . This transaction carries all the information (*e.g.* the hash-key of the model to train in the store) and $param_{fl_i}$, the

³ The term process is used here to denote a computer thread, distinct from the usage of FL process referring to the Federated Learning process. Both terms are employed within their respective contexts to avoid any confusion.

⁴ When `locally_valid` is deterministic, it holds for all correct servers.

Algorithm 1 Server i

```
1:  $bc$ ; ▷ Blockchain object
2:  $t \leftarrow 0$ ; ▷ Current round
3:  $m$ ; ▷ Minimum number of expected clients' updates to aggregate them
4:  $\mathcal{P}(G_t)_{cand} \leftarrow \emptyset$ ; ▷ Set of updates' PoA&I  $P(h(g_t^k))$  candidate for round  $t$ 
5:  $\mathcal{P}(G_t)_{commit} \leftarrow \emptyset$ ; ▷ Set of updates' PoA&I  $P(h(g_t^k))$  commit for round  $t$ 
6:  $C_t \leftarrow \emptyset$ ; ▷ Set of  $C_t$  elements, the selected clients for the round  $t$ 
7:  $LP \leftarrow \emptyset$ ; ▷ Set of local proofs for values

8: upon RECEIVE(UPDATE,  $g_t^k$ ) ▷ Update  $g$  by client  $k$  for round  $t$ 
9:   if  $k \in C_t$  then
10:     buildPoA&I(UPDATE,  $g_t^k$ );
11:      $C_t \leftarrow C_t \cup k$ ;

12: upon EVENT( $|\mathcal{P}(G_t)_{cand}| \geq m$  for the first time)
13:    $C_{t+1} \leftarrow \text{select\_clients}(bc)$ ; ▷ Based on the client's JOIN transactions
14:   buildPoA&I(CLIENTS,  $C_{t+1}$ );

15: upon DELIVERTX( $\langle NR, t, \mathcal{P}(G_t)_{cand}, C_{t+1}, P(C_{t+1}) \rangle$ )
16:   if verify_proof( $\mathcal{P}(G_t)_{cand} \wedge C_{t+1} \wedge \mathcal{P}(G_t)_{commit} = \emptyset$ ) then
17:      $\mathcal{P}(G_t)_{commit} \leftarrow \mathcal{P}(G_t)_{cand}$ 

18: upon EVENT( $\mathcal{P}(G_t)_{commit} \neq \emptyset$  for the first time)
19:    $G_t \leftarrow \text{getAll}(\mathcal{P}(G_t)_{commit} \setminus \mathcal{P}(G_t)_{cand})$ ;
20:    $w_{t+1}^i \leftarrow \text{robust\_aggregation}(G_t)$ ;
21:   buildPoA&I(MODEL,  $w_{t+1}^i$ );

22: upon BUILDPOA&I(TAG,  $v$ ):
23:   if locally_valid(TAG,  $v$ ) then
24:     if TAG  $\neq$  CLIENTS then ▷ Only models and updates are sent to the storage
25:       send  $\langle h(v), v \rangle$  to all  $r \in R$ ; ▷ Each storage replica, upon reception,
checks if  $h(v) = v$  before storing  $v$  with hash-key  $h(v)$ 
26:       if TAG = MODEL then
27:         send (MODEL,  $v$ ) to all  $j \in S$ ;
28:         send  $\langle LP, TAG, h(v)_i \rangle$  to all  $j \in S$ ; ▷  $h(v)_i$  is signed by server  $i$ 

29: upon RECEIVE(MODEL,  $v$ ) for the first time
30:   if locally_valid(TAG,  $v$ ) then
31:     send  $\langle h(v), v \rangle$  to all  $r \in R$ ;
32:     send  $\langle LP, TAG, h(v)_i \rangle$  to all  $j \in S$ ;

33: upon RECEIVE( $\langle LP, TAG, h(v)_j \rangle$ )
34:    $LP \leftarrow \{ \langle LP, TAG, h(v)_j \rangle \}$ ;
35:   if LP contains  $f + 1$   $\langle LP, TAG, h(v)_* \rangle$  then
36:      $P(h(v)) \leftarrow f + 1$   $\langle LP, TAG, h(v)_* \rangle$ ;
37:     if TAG = UPDATE then
38:        $\mathcal{P}(G_t)_{cand} \leftarrow \mathcal{P}(G_t)_{cand} \cup P(h(v))$ 
39:     else if TAG = CLIENTS then
40:       sendTx(NR,  $t, \mathcal{P}(G_t)_{cand}, C_{t+1} = v, P(h(v))$ ) to bc;
41:       ▷ If  $t = t_{fin}$ , the last round, then rather than  $C_t$ , this transaction
carries the PoA&I of the rewarding information.
42:     else if TAG = MODEL then
43:       send(MODEL,  $t, P(h(v))$ ) to clients;
44:        $t \leftarrow t + 1$ ;
```

Algorithm 2 Read from the key-value store at process i

```
1: upon GET( $P(\ell)$ )
2:   if verify_proof( $P(\ell)$ ) then ▷ verify if  $h(\ell)$  is signed by  $f_s + 1$  servers
3:     return read( $P(\ell)$ );

4: upon GETALL( $\mathcal{P}(Set)$ ) and (verify_proof( $P(\ell)$ ),  $\forall \ell \in Set$ ):
5:    $buffer \leftarrow \emptyset$ ;
6:   for all  $P(\ell) \in \mathcal{P}(Set)$  do
7:      $buffer \leftarrow buffer \cup \text{read}(P(\ell))$ ;
8:   return  $buffer$ ;

9: upon READ( $P(\ell)$ ):
10:  send(QUERY,  $h(\ell)$ ) to all  $r \in R$ ;
11:  wait until received(REPLY,  $\ell'$ ) and  $h(\ell) == h(\ell')$  return  $\ell'$ ;
```

parameters necessary for the task execution. The set-up round $t = 0$ starts;

- 2) Once the fl_i task is registered on the blockchain, all clients and servers are aware of it. Each client sends a transaction to the blockchain to state its interest to join (if any). Such a transaction stakes a certain amount of coins⁵ for fl_i , witnessing the client commitment to participate. When the client's transaction is registered on the blockchain we say that the client is registered for fl_i . Servers wait for sufficiently many clients to be registered (the number of awaited clients is set in $param_{fl_i}$).
- 3) Once there are sufficiently many clients registered for fl_i , each server selects deterministically out of them a subset C_1 to perform the learning for the first round and builds $P(C_1)$, the PoA&I for it (cf. Algorithm 1). Finally, it sends a transaction START_TASK to announce to the clients in C_1 that they are selected for the first round. Once this transaction is registered on the Blockchain the set-up round is over and the computation moves to the first learning round $t = 1$.
- 4) Each client k in C_t , retrieves a previously aggregated model associated to a PoA&I, (cf. Algorithm 2), performs the learning on its local data to produce the update g_t^k , signs it, and sends it directly to all servers;
- 5) Each server upon reception of g_t^k for the first time, builds a PoA&I for it (cf. Algorithm 1).
- 6) When a server collects sufficiently many updates's PoA&I for round t (as determined in $param_{fl_i}$), it has a candidate update set, $\mathcal{P}(G_t)_{cand}$. It computes C_{t+1} , the clients subset to perform the learning in the next round and build $P(C_{t+1})$. Finally, it sends a transaction NR (New Round) with $\mathcal{P}(G_t)_{cand}$, C_{t+1} and $P(C_{t+1})$;
- 7) When a server i sees a NR transaction on-chain for the first time for a round t , it retrieves the missing updates and the previous model (if necessary) from the storage;
- 8) When a server p has all the information, it computes the aggregated model w_t^p (by performing robust aggregation to Byzantine clients as defined in Section IV-A1), sends it to the other servers⁶ to build a PoA&I for it. Finally, it sends the PoA&I for w_t^p to clients (which is used in steps 4 to retrieve the model). Round t ends and we move to $t + 1$.
- 9) The workflow loops between steps 4 and 8 (a round) as long as the learning does not halt (the halt condition is set in $param_{fl_i}$). At the end of the learning, each server produces information to assess quantitatively clients' contributions to the fl_i task, builds a PoA&I and sends a transaction with those information. This transaction allows clients to be automatically rewarded and the model owner to retrieve the model from the storage.

Interestingly, as we detailed, for each round $t \geq 1$ the

⁵ "Staking" in the context of cryptocurrency refers to the process of actively participating in the operations of a blockchain network. This involves locking up a certain amount of cryptocurrency as collateral to perform these tasks.

⁶ It suffices having only $f_s + 1$ selected servers that diffuse their aggregated model to have at least one ϵ -valid model that obtains a PoA&I. We omit that to keep the code simple.

blockchain registers a single transaction per round.⁷ An interested reader can find a sequence diagram of the described workflow in [27].

1) *BuildPoA&I - Computational Integrity check*: The PoA&I production applies to the client set C_t , client k 's update g_t^k , aggregated model w_t , and rewarding information r . Each of these elements undergoes a `locally_valid` check (cf. Algorithm 1). The specific validation process depends on their type. When handling the update $\langle \text{UPDATE}, g_t^k \rangle$, the check involves verifying the correct formatting of g_t^k and ensuring that $k \in C_t$. In client set case, as C_t results from the same deterministic computation executed by all correct servers, there is no need for additional verification. The same applies for rewarding information r . However, the aggregated model, resulting from a non-deterministic computation, may slightly differ across correct servers. Each server p potentially produces a different w_t^p . As formalized in the next section, we allow for an ϵ gap between aggregated models. Thus, p considers another model w_t^q valid if ϵ close to w_t^q .

2) *Floating points and non-determinism handling*: While servers perform computations using the same update set, slight variations arise due to the non-determinism of floating-point representation. This discrepancy poses a challenge when a server i has to validate the model w_j^t produced by server j having its own model w_i^t as reference.

On a positive note, the disparity between computed models is small enough (cf. note 1) to introduce an acceptance vector gap that we denote as $\epsilon \in \mathbb{R}^d$. We assume that ϵ is static and provided by the model owner when specifying fl_i . We assume that ϵ is sufficiently large to prevent a deadlock: a round in which there are no models issued by servers that obtain a PoA&I (if their difference is bigger than ϵ).

Assumption 1: Let $\mathcal{W}_{fl_i,t}^C$ be the set of all aggregated models produced by correct servers in a round t for the task fl_i via a robust aggregation function as defined in Section IV-A1. Let $\epsilon \in \mathbb{R}^d$ be the acceptance vector gap for the fl_i task. We assume that each pair of models in $\mathcal{W}_{fl_i,t}^C$ is ϵ -close. We say that an aggregated model w^t is ϵ -valid if it is ϵ -close to at least one model in $\mathcal{W}_{fl_i,t}^C$.

Due to the lack of space, we defer the correctness arguments to [27].

VI. EXPERIMENTAL EVALUATION

This section evaluates the practicality of Fantastyc in a realistic federated learning scenario. Initially, we characterize the required number of clients to be selected in each round of the FL process to achieve good accuracy (Section VI-B). Subsequently, we assess the accuracy loss when implementing confidentiality and robust aggregation methods (Section VI-C). Both evaluations are conducted in a scenario with only one server and co-located clients. Following that, we delve into the effects of server decentralization in a geo-replicated deployment (Section VI-D). We evaluate communication overheads introduced by our decentralization approach, specifically

⁷ In the practice there is one transaction per server, but processes wait only for the first valid one on-chain to proceed with their computation.

analyzing the contributions of the blockchain and the PoA&I to round latency. As a baseline, we measure the round latency with one server and clients spread across regions on the Internet. This baseline provides the performance benchmark for a classical FL process with multiple clients and a single server – a non Byzantine fault-tolerant configuration with $f_s = 0$. We will study the latency overhead due to different fault tolerance thresholds.

It is worth noting that our evaluation does not deal with the demonstration of high throughput because we focus on only one FL task. Let us also stress that FL operates in rounds that can be lengthy, not imposing strict latency constraints. Furthermore, we do not compare our approach with other blockchain-based methods due significant differences in architectures, use and choice of blockchains, along with the impracticality of storing large models in the blockchain (see Section III for a discussion on related work) – an issue that limits realistic implementations.

A. System implementation and deployment

We implemented our blockchain-based solution using a BFT protocol that ensures immediate finality and can be flexibly deployed as a private, public, or private side-chain like in Cosmos⁸. For this reason we chose Tendermint⁹ (now CometBFT) [18], which is the consensus protocol of Cosmos. For the implementation of our servers we use the ABCI for decoupling the application logic from the blockchain (see Section II-B). All the algorithms presented in Section V are implemented as part of the application logic. The FL clients and the model owner communicate directly with the blockchain as Tendermint clients. We use the pub/sub/query protocol Zenoh¹⁰ to implement the fault-tolerant key-value store. The application logic at the servers is implemented in Python¹¹ and the FL training at the clients is based on the TensorFlow library¹².

B. Experimental setup - model and dataset

We use the Federated Extended MNIST dataset¹³ dedicated to the handwritten digit classification task as experimental setup. This MNIST version contains 10 classes and comes with author id in such a way that its federated version was built by partitioning the data based on the writer. The dataset contains 3,579 writers (each with train/test data split for a total of 240,000 train and 40,000 test images) that permits to test the robustness and scalability of our solution. Evaluations were done with a standard CNN (the same as the one used in [1]) composed of two $5 * 5$ convolution layers (the first with 32 channels, the second with 64 channels, each followed with $2 * 2$ max pooling), a fully connected layer with 512 units and ReLu activation, and a final softmax output layer (1,663,370 total parameters). For local training, the number of epochs and the minibatch size were set to $E = 10$ and $B = 5$ respectively.

⁸ <https://cosmos.network/>

¹⁰ <https://zenoh.io/>

¹² <https://www.tensorflow.org/>

⁹ <https://docs.tendermint.com/v0.34/>

¹¹ <https://www.python.org/>

¹³ <https://www.nist.gov/itl/products-and-services/emnist-dataset>

As benchmark, we report the accuracy of the federated model on the test set throughout the training process without malicious clients or countermeasures to prevent information leakage. We experiment with the number K of participants per round, which controls the amount of multi-client parallelism (Fig.2.a). We show the impact of varying K about the speedup of the learning process and it appears that there is no more advantage in increasing the client fraction above $K = 50$ (Fig.2.b). Thus, for further experiments we fix $K = 50$, which strikes a good balance between computational efficiency and convergence rate.




Fig. 2: Study about the impact of multi-client parallelism varying K (while setting $B = 5$; $E = 10$). (a) Test set accuracy vs. communication rounds varying the number of participants (b) Number of communication rounds to reach a target accuracy of 95% depending on the number of participants per round.

C. Privacy and robustness trade-offs evaluation

We have conducted experiments to evaluate accuracy loss due to the addition of our modules dedicated to privacy and integrity issues (Table I). Without malicious clients or countermeasures to prevent information leakage, the use of FedAvg as aggregation rule permits to reach the best accuracy of 98.6% after 500 rounds of training. We illustrate next the necessary trade-off between privacy, integrity and accuracy.

The risk of privacy leakage decreases as the number of samples s used to build InstaHide composite images increases but in the same time, the training process suffers small accuracy loss of 1.1% to 3.7% when s varies from 2 to 5.

To evaluate the robustness of our solution to Byzantine attacks, we consider that there are $f_c = \{10, 20, 30\}$ malicious clients among 50 participants at each round of training. By only sending $-\lambda_{boost}g_t^k$ instead of g_t^k (where $\lambda_{boost} = 5$), it is very easy for Byzantine clients to impeded a training process based on FedAvg as the model comes back to random with 10% accuracy. The use of a Byzantine-tolerant aggregation rule such as the median is thus essential to withstand Byzantine attacks and assure the converge of the model – as long as the number of honest participants remains a majority.

D. Experimental setup - distributed system

We run our experiments in the cloud computing service Akamai-Linode¹⁴, which offers dedicated CPU and storage resources that can be allocated in geographically distant regions. For the experiments, we evaluate a deployment where within

¹⁴ <https://www.linode.com/>

f_c	0					10	20	30	10	20	30
Mix	no	2	3	4	5	no			4		
FedAvg	98.6	97.5	96.5	95.6	94.9	10.0	10.0	10.0	10.0	10.0	10.0
median	96.7	95.2	93.8	92.0	90.5	96.3	94.7	10.0	90.9	89.7	10.0

TABLE I: Model accuracy after 500 rounds of training with $K = 50$ clients and considering FedAvg;median as aggregation rule. We deploy InstaHide with $s = \{2; 3; 4; 5\}$ mix samples and simulate $f_c = \{0, 10, 20, 30\}$ byzantine clients with $\lambda_{boost} = 5$.

each node we collocate a server, a blockchain validator, and a storage replica (a peer, in Zenoh’s parlance). We consider several configurations where the nodes and the clients are spread in up to four different regions (EU, US-west, US-east, Asia). At each region we collocate one or more instances (dedicated Linode machines with 32GB of memory) which host the nodes. For the FL clients, we let each Tendermint client (called a proxy client) to generate the workload of a set of realistic participants (recall a participant is an FL client that takes part in the training). At each round, each proxy client sends the payloads from the set of participants it contains in sequential fashion. We deploy up to two instances per region, each of them with a number of nodes that varies between one and 16. We evenly distribute the clients among instances deployed in all regions, where each instance may contain up to 8 proxy clients. Each proxy client is capable of generating the workload for between 2 to 5 participants. Typically, we ensure a balanced distribution by limiting the number of instances to a maximum of four and adjusting the number of participants per proxy client to fall within the range of 10, 50, or 100 participants in total. The communication bandwidth between the regions is in the range of 3Mbps–10Mbps.

E. Scalability evaluation

We first study the latency associated to a single round and then we show the communication overhead within the overall FL process, with varying number of regions, nodes, and clients.

Round analysis. Round latency includes the block time (time to add in the blockchain a block that includes the transaction sent in the round), the PoA&I time (to generate a valid PoA&I), as well as the contribution to complete steps 4–8 in the workflow of Section V. Importantly, at each round there is only one write access to the blockchain (step 6).

Blocktime contribution. Fig.3 plots latency results with $K = 50$ participants per round and a varying number of nodes, training the CCN model considered in Section VI-B, whose payload corresponds to 12MB. Preliminary experiments yield a practical limit of 16 nodes, above of said limit the latency degrades very quickly due to depletion of resources at the Linode instances. The top two curves correspond to round latency with payload 12MB, with either four regions and one instance per region (blue) or two regions and two instances per region (orange). The best round latency with 16 nodes is

96s with two regions. The next two curves show the round time respectively with a mocked model of size 6KB. The best latency with 16 nodes is 20s with four regions. Finally, the base line plots the block times associated to the configurations of the four curves above (green and red for 12MB payload and four or two regions respectively, purple and brown for 6KB payload and four or two regions respectively). These experiments confirm that the block time is independent of the payload size, which is to be expected given Section IV-B and smoothly increases with the number of nodes in nominal network conditions.




Fig. 3: Evaluation of the impact of the number of nodes with a constant number of 50 clients.

PoA&I contribution. Fig.4 presents the latency breakdown across different intervals of the learning round, maintaining the same client configuration as described earlier for 4, 10, 16 nodes, respectively, evenly deployed over 2 regions, each with 2 instances. In the plot, the time for the client to gather the model from the storage (computed in the previous round) and verify model’s PoA&I (step 4) is represented in blue, the time for servers to collect client updates sent by clients and generate the PoA&I (step 5) is in orange, the block time (step 6) is in green, the time for servers to verify updates’ PoA&I and get the last model from the storage (step 7) is in red, and the time to aggregate updates into the new model and generate the PoA&I of the aggregated model (step 8) is in purple.




Fig. 4: Breakdown of latency within a round with 50 clients.

As for the PoA&I contribution to latency for client updates (included in the orange bar), we can see that the PoA&I generation for client updates is quick; witnessed by the comparable latency between the baseline (without PoA&I) and the 4-node configuration. This is due to a voting process involving small-sized votes (less than 1KB) and the fact that each PoA&I production is run asynchronously as soon as a client update

is received by the server, and this for each update received in the round.

The purple bar shows the time spent to generate the PoA&I for the model. For this PoA&I, each node publishes the model for the other server nodes. The only linear increase in time with the number of nodes is because the latency is dominated by the publishing of the model, while the latency to get $f_s + 1$ votes is negligible, due to the extremely small size of quickly published vote messages. As for the blue bar, the client only needs to retrieve one model with a valid PoA&I. This explains the constant contribution of the step 4 no matter of the number of nodes. Without a PoA&I, each client would be required to wait for at least $f_s + 1$ computed models, which would have a prohibitive cost with a large number of client participants.

Communication overheads. For a comprehensive evaluation of the results, Table II illustrates the communication overhead for various fault-tolerance thresholds f_s , for a number of participants per round $K = 10$, $K = 50$, $K = 100$ (which implies $f_c=4$, $f_c=20$ and $f_c=40$). We can consider the threshold for servers, such that $n_s = 2f_s + 1$, but since in the considered deployment the server is collocated with a blockchain validator in the same node, we can also fairly consider $n_s = 3f_s + 1$ – this depends on granularity of the failure model, at node level or at software component level. We consider both, distinguishing them within the interval $[f_s^{min}, f_s^{max}]$. To provide a clearer understanding of the overhead, we present the total time \mathbf{T} spent for training calculated as the latency per round multiplied by 500, and time to reach good accuracy T_a ($> 95\%$), which is 37 rounds for $K \geq 50$ and 43 rounds for $K = 10$ (Section VI-B). From the results it can be noted that on average for the smaller model the overhead increases linearly with the number of nodes. On the other hand, for the large model, the overhead increases sub-linearly with the number of nodes. In general, our approach is able to optimise the case of the larger models. As observed in the latency breakdown (Fig.4), this is due to the time for a client to retrieve and verify the new model from its PoA&I, as it waits for the first good reply (blue bar) making this time constant with respect to the number of nodes, the asynchronous process to generate a PoA&I (orange bar) leveraging parallelism when the number of servers increases, and an acceptable contribution of the block time experienced in nominal network conditions.

n_s	1 (baseline)	4	16
$[f_s^{min}, f_s^{max}]$	[0,0]	[1,1]	[5,7]
$[T, T_a]$ 6KB, K=10	[7, 0.6]min	[12, 1.1]min	[135, 11.6]min
$[T, T_a]$ 6KB, K=50	[12, 0.9]min	[45, 3.3]min	[197, 15]min
$[T, T_a]$ 6KB, K=100	[17, 1.3] min	[70, 5.2] min	[247, 18]min
$[T, T_a]$ 12MB, K=10	[1, 0.09] hr	[2.4, 0.2] hr	[8.5, 0.73] hr
$[T, T_a]$ 12MB, K=50	[3.4, 0.25] hr	[5, 0.4] hr	[13, 1] hr
$[T, T_a]$ 12MB, K=100	[6.7, 0.5] hr	[9.4, 0.7] hr	[15.3, 1.1] hr

TABLE II: Communication overhead with different fault thresholds.

VII. CONCLUSION AND PERSPECTIVES

This paper introduced a blockchain-based federated learning system that is provably secure, improves Byzantine tolerance, limits the risks of sensitive information leakage and showcases practical applicability by amortizing the cost of growing workloads with the number of server replicas in nominal network conditions. Future work will primarily concentrate on two concurrent axes: the first involves further improving the underlying communication overhead by implementing gradient compression techniques, and evaluation in more dynamically changing environments. studying privacy/utility trade-offs. The second axis revolves around incentive mechanisms that can leverage the participation count of FL contributors in each round to address the issue of free-riding¹⁵. Indeed, even if completely eliminating free-riding without access to each client's specific dataset remains an unsolved challenge [28], robust incentive mechanisms encouraging agents to exhibit honest behavior can be designed, assuming the presence of an auditor [29]. Blockchain-based auditor mechanisms are then a promising avenue of research. As part of our future work, we also explore a mechanism for dynamically adjusting ϵ . That would allow to lift Assumption 1, enabling to start the FL process with a small ϵ , without risking a deadlock. Finally, we could also make our solution robust to backdoor attacks [30]. As our architecture is modular, it seems straightforward to integrate a post-training defense such as the one presented in [31].

REFERENCES

- [1] H. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 2017, pp. 1273–1282.
- [2] P. Kairouz et al., "Advances and open problems in federated learning," 2019, arXiv preprint 1912.04977.
- [3] P. Ramanan, K. Nakayama, and R. K. Sharma, "Baffle : Blockchain based aggregator free federated learning," *IEEE International Conference on Blockchain (Blockchain)*, 2019.
- [4] X. Bao et al., "Flchain: A blockchain for auditable federated learning with trust and incentive," in *2019 5th International Conference on Big Data Computing and Communications (BIGCOM)*, 2019.
- [5] H. Kim, J. Park, M. Bennis, and S.-L. Kim, "Blockchained on-device federated learning," *IEEE Communications Letters*, 2020.
- [6] J. Kang et al., "Scalable and communication-efficient decentralized federated edge learning with multi-blockchain framework," in *Blockchain and Trustworthy Systems*, 2020.
- [7] H. B. Desai, M. S. Ozdayi, and M. Kantarcioglu, "Blockfla: Accountable federated learning via hybrid blockchain architecture," in *Proceedings of the Eleventh ACM Conference on Data and Application Security and Privacy*, ser. CODASPY '21, 2021.
- [8] M. Shayan, C. Fung, C. J. M. Yoon, and I. Beschastnikh, "Biscotti: A blockchain system for private and secure federated learning," *IEEE Transactions on Parallel and Distributed Systems*, vol. 32, no. 7, 2021.
- [9] C. Korkmaz et al., "Chain fl: Decentralized federated machine learning via blockchain," *Second International Conference on Blockchain Computing and Applications (BCCA)*, 2020.
- [10] Z. Peng et al., "Vfchain: Enabling verifiable and auditable federated learning via blockchain systems," *IEEE Transactions on Network Science and Engineering*, 2022.
- [11] D. Yin, Y. Chen, R. Kannan, and P. Bartlett, "Byzantine-Robust Distributed Learning: Towards Optimal Statistical Rates," in *International Conference on Machine Learning*. PMLR, 2018, pp. 5650–5659.
- [12] L. Lamport, R. Shostak, and M. Pease, "The Byzantine Generals Problem," *ACM Transactions on Programming Languages and Systems*, vol. 4, no. 3, pp. 382–401, 1982.
- [13] P. Labs, "Proof of replication," <https://filecoin.io/proof-of-replication.pdf>, 2017.
- [14] "Cosmos: The internet of blockchains." [Online]. Available: <https://github.com/cosmos/cosmos>
- [15] M. Bourgoin, "An overview of the tezos blockchain." [Online]. Available: <https://tezos.com/>
- [16] V. Buterin, "Ethereum white paper: A next generation smart contract & decentralized application platform," 2013. [Online]. Available: <https://github.com/ethereum/wiki/wiki/White-Paper>
- [17] N. Szabo, "Smart contracts," 1994. [Online]. Available: <https://www.fon.hum.uva.nl/rob/Courses/InformationInSpeech/CDROM/Literature/LOTwinterschool2006/szabo.best.vwh.net/smart.contracts.html>
- [18] E. Buchman, J. Kwon, and Z. Milosevic, "The latest gossip on bft consensus," arXiv preprint arXiv:1807.04938, 2018.
- [19] "Cometbft. abci+," accessed on: November 2023. [Online]. Available: <https://docs.cometbft.com/v0.37/spec/abci/>
- [20] D. Cason et al., "The design, architecture and performance of the tendermint blockchain network," in *International Symposium on Reliable Distributed Systems (SRDS)*, 2021.
- [21] Y. Li, C. Chen, N. Liu, H. Huang, Z. Zheng, and Q. Yan, "A blockchain-based decentralized federated learning framework with committee consensus," *IEEE Network*, vol. 35, no. 1, 2021.
- [22] S. Cohen, G. Goren, L. Kokoris-Kogias, A. Sonnino, and A. Spiegelman, "Proof of availability and retrieval in a modular blockchain architecture," in *International Conference on Financial Cryptography and Data Security*. Springer, 2023, pp. 36–53.
- [23] L. Zhu, Z. Liu, and S. Han, "Deep leakage from gradients," in *Advances in Neural Information Processing Systems*, 2019.
- [24] J. Geiping, H. Bauermeister, H. Dröge, and M. Moeller, "Inverting gradients - how easy is it to break privacy in federated learning?" in *Advances in Neural Information Processing Systems*, 2020.
- [25] A. Choffrut, R. Guerraoui, R. Pinot, R. Sirdey, J. Stephan, and M. Zuber, "Practical homomorphic aggregation for byzantine ml," 2023.
- [26] Y. Huang, Z. Song, K. Li, and S. Arora, "InstaHide: Instance-hiding schemes for private distributed learning," in *Proceedings of the 37th International Conference on Machine Learning*, 2020.
- [27] W. Boitier, A. D. Pozzo, Álvaro García-Pérez, S. Gazut, P. Jobic, A. Lemaire, E. Mahe, A. Mayoue, M. Perion, D. Singh, T. F. Rezende, and S. Tucci-Piergiovanni, "Fantastyc: Blockchain-based federated learning made secure and practical," 2024. [Online]. Available: <https://arxiv.org/abs/2406.03608>
- [28] Y. Fraboni, R. Vidal, and M. Lorenzi, "Free-rider attacks on model aggregation in federated learning," *International Conference on Artificial Intelligence and Statistics*, 2021.
- [29] X. A. Gao, J. R. Wright, and K. Leyton-Brown, "Incentivizing evaluation with peer prediction and limited access to ground truth," *Artificial Intelligence*, vol. 275.
- [30] E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin, and V. Shmatikov, "How to backdoor federated learning," in *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics (AISTATS)*, vol. 108, Palermo, Italy, 2020, pp. 2938–2948.
- [31] F. Castellon, D. Singh, A. Mayoue, and C. Gouy-Pailler, "Fuba: Federated uncovering of backdoor attacks for heterogeneous data," in *2023 5th IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA)*, 2023, pp. 55–63.

¹⁵ Free-riding in this context is related to the problem of an agent that will prefer not endeavoring the cost of learning while presenting fake gradients.