

BOSSA: A Decentralized System for Proofs of Data Retrievability and Replication

Dian Chen, Haobo Yuan, Shengshan Hu,
Qian Wang, *Senior Member, IEEE*, and Cong Wang, *Senior Member, IEEE*

Abstract—Proofs of retrievability and proofs of replication are two cryptographic tools that enable a remote server to prove that the users' data has been correctly stored. Nevertheless, the literature either requires the users themselves to perform expensive verification jobs, or relies on a “fully trustworthy” third party auditor (TPA) to execute the public verification. In addition, none of existing solutions consider the underlying incentive issues behind a rational server who is motivated to collect users' data but tries to evade the replication checking in order to save storage resources. In this work, we propose the first decentralized system for proofs of data retrievability and replication—BOSSA, which is incentive-compatible for each party and realizes automated auditing atop off-the-shelf blockchain platforms. We deal with issues such as proof enforcements to catch malicious behaviors, new metrics to measure the contributions, and reward distributions to create a fair reciprocal environment. BOSSA also incorporates privacy-enhancing techniques to prevent decentralized peers (including blockchain nodes) from inferring private information about the outsourced data. Security analysis is presented in the context of integrity, privacy, and reliability. We implement a prototype based on BOSSA leveraging the smart contracts of Ethereum blockchain. Our extensive experimental evaluations demonstrate the practicality of our proposal.

Index Terms—Proofs of Retrieval; Proofs of Replication; Decentralized system; Blockchain.

1 INTRODUCTION

As a new computing paradigm, cloud computing provides a convenient approach for ordinary users to enjoy powerful computing and storage resources. Users can outsource their data to a cloud service provider to get rid of complex local data management. Despite promising prospects, outsourcing data to a remote server \mathcal{S} also raises severe security concerns since it deprives the physical control of the data owner. One of the challenging problems is ensuring the correctness of data. A malicious server may discard the data that is rarely accessed for the purpose of saving resources, or cover up data loss accident for maintaining reputation. In addition, \mathcal{S} usually claims that the data will be stored together with several replicas for ensuring high reliability [1], it takes limited liability in their Service Level Agreements (SLAs) [2], especially for the data loss. And recent accidents happened in Tencent Cloud [3], other clouds [4] have already implied the fact that \mathcal{S} cannot be fully trusted.

Proofs of retrievability [5], [6], [7], [8], [9], [10], [11], [12] and proofs of replication [13], [14], [15], [16], [17], [18], [19] are two typical cryptographic methods allowing the server to prove that the original files as well as all the replicas are correctly

stored. In order to avoid keeping users on-line and free them from expensive auditing tasks, the state-of-the-art solutions for proofs of retrievability and replication heavily rely on a semi-trusted third party auditor (TPA) to execute public verifications. However, such methods suffer from the following two main drawbacks:

- **Collusion Attack.** The literature usually assumes that no collusion happens between TPA and the server \mathcal{S} . This strong assumption may be easily corrupted in practice driven by certain interests, and \mathcal{S} can easily create fake proofs that pass the verifications once it colludes with TPA. Furthermore, from the users' perspective, it is hard to tell whether or not collusion happened.
- **Corruption Attack.** Dependence on an “always-online” TPA is vulnerable to single point of failures caused by unpredictable accidents, *e.g.*, regional power outages, or malicious hacking attacks, like DDoS.

We owe these vulnerabilities to the great control power of TPA which plays a centralized role. Therefore, constructing a decentralized framework for auditing can fundamentally solve the above problems. The blockchain has been widely adopted for replacing the third party in the fields like e-voting [20], auction [21] and fair exchange [22] due to its non-repudiation and non-tampering properties. Replacing the TPA with the blockchain makes the whole auditing process trackable by users. As a result, the users are able to verify each step of proofs retrievability and replication. Moreover, the decentralized nature of the blockchain improves the tolerance of single point failures. In summary, the blockchain is a promising solution to defend against *collusion attack* and *corruption attack*.

In addition, for the data replica service that aims to guarantee the data reliability, the cloud may intentionally delete replicas for saving storage resources without concerns of being caught. Inspired by the idea of decentralized storage network (DSN) that

-
- D. Chen and Q. Wang are with the Key Laboratory of Aerospace Information Security and Trusted Computing, Ministry of Education, the School of Cyber Science and Engineering, Wuhan University, Wuhan, Hubei 430072, China, and also with the State Key Laboratory of Cryptology, P.O. Box 5159, Beijing 100878, China. (Corresponding author: Qian Wang.)
E-mail: {dianchen, qianwang}@whu.edu.cn
 - H. Yuan is with the School of Computer Science, Wuhan University, Wuhan, Hubei 430072, China.
E-mail: yuanhaobo@whu.edu.cn
 - S. Hu is with the School of Cyber Science and Engineering, Huazhong University of Science and Technology, Wuhan, Hubei 430074, China.
E-mail: hushengshan@hust.edu.cn
 - C. Wang is with the Department of Computer Science, City University of Hong Kong 999077, Hong Kong.
E-mail: congwang@cityu.edu.hk

tries to build a peer-to-peer network allowing peers buying and selling idle disk space, decentralizing data replication makes it much more difficult for the attackers to misbehave, as each step for the data transmission, storing and deleting is recorded and trackable.

1.1 Design Challenges and Our Solutions

In this work, we propose the Blockchain based OutSourcing Storage and Auditing (BOSSA) scheme, a brand-new general framework for proofs of retrievability and replication. BOSSA separately stores the original data on the server and the corresponding replicas on peers (also called farmers in BOSSA) of a decentralized network atop blockchain, and integrities of them will be checked respectively by carefully-designed smart contracts. BOSSA can naturally defend against the *collusion attack* and the *corruption attack*, while provides visibility-enabled data replication. BOSSA doesn't heavily rely on a specific consensus model like [23], which gives the opportunity that BOSSA can be built on off-the-shelf blockchain systems supporting Turing-complete smart contract [24], [25]. To design BOSSA, there are several challenges we have to cope with.

First of all, simply combining existing proof of retrievability schemes with the blockchain fails to achieve our desired goals since the blockchain only supports simple functionality. The most critical problem is that the blockchain cannot actively issue challenges and reacts only based on received transactions. This property allows the adversary to escape from auditing. To prevent such lazy behaviors, we make use of the feature that blocks are appended to the chain in an approximately-fixed rate, and propose a time-restricted proof forcing both the cloud and farmers to prove data availability.

Second, our scheme organizes peers, *i.e.*, farmers, to store the replicas of users' data. However, farmers are not fully trusted and may leave the network arbitrarily, which causes replicas loss, damaging the reliability of replicas. To address this problem, we design an incentive mechanism to motivate peers to store and provide replicas when needed. Specifically, we define a reward mechanism where farmers are periodically rewarded if they can provide valid proofs of replicas on a regular basis. Meanwhile, part of the reward is temporarily frozen until farmers' promised storage services expired, which discourages farmers from leaving the network. To further motivate farmers to share replicas, we define a metric called *contribution rate*, and connect farmers' frozen reward with the contribution rate, such that only farmers with 100% contribution rate can retrieve all the frozen rewards.

Third, most of blockchain platforms like Ethereum store transactions in plaintext, any participant is able to retrieve the information of transactions, including the inputs of smart contracts, the intermediate results, etc. For an auditing system naively built atop the blockchain, the open nature of the blockchain enables the adversary to record proofs and recover the original data. To prevent such privacy leakage, we incorporate privacy-enhancing techniques for the generation of proofs while ensuring their correctness.

1.2 Contribution

In summary, we make the following contributions:

- We propose a novel decentralized framework BOSSA for proofs of data retrievability and replication. Our design addresses security risks and caters to a more real-world

scenario, which has not been considered so far in existing works.

- To guarantee the reliability of BOSSA, we propose a time-restricted proof forcing \mathcal{S} and farmers to prove data availability, and a reward mechanism based on the new metric *contribution rate* to create a fair reciprocal environment. Privacy-enhancing techniques are also incorporated to protect users' private data.
- We provide the security analysis for BOSSA with respect to integrity, privacy, and reliability. A prototype of BOSSA is implemented on Ethereum, and evaluated. The experimental evaluations show that our proposal is feasible and incurs tolerable overheads for each party.

1.3 Related Work

Proofs of Retrievability/Data Possession. Both Proof of Data Possession (PDP) and Proofs of Retrievability (PoR) aim to ensure that the outsourced data is stored on \mathcal{S} correctly. Juels and Kaliski proposed the first PoR scheme [9] and Ateniese *et al.* [5] constructed the first PDP scheme. However, their schemes do not provide public verification, which requires users keeping on-line during the auditing process. Guan *et al.* [26] leveraged the indistinguishability obfuscation to realize public verification, but their scheme has poor performance due to the usage of the indistinguishability obfuscation. Shacham and Waters [10] then constructed a public verification scheme based on BLS signature [27], but their scheme may cause privacy-leakage during the auditing [12]. Armknecht [8], Xu [7] and Yang [6] proposed low-cost public PoR schemes by modifying the private verification scheme of [10], however, the TPA in their scheme holds the secret key for auditing, which is not suitable to public blockchain platforms.

Proofs of Replication. Different from PDP/PoR, Proofs of Replication focuses on ensuring that the replicas of original data are correctly stored on \mathcal{S} as it claimed. Most proofs of replication schemes are based on indistinguishable encryption/encoding [14], [16], [28] or noticeable replica generation time [13], [18], [19], [29]. MR-PDP proposed in [14] requires users to encrypt replicas and upload them to \mathcal{S} , and lets users audit the corresponding replicas. Hao *et al.* [28] leveraged similar methods, and their scheme supports public verification. Damgard *et al.* [16] generalized the indistinguishable encryption/encoding based model which utilizes a trapdoor function without which \mathcal{S} cannot generate correct replicas on-the-fly. However, these schemes bring heavy bandwidth and computation costs to users, as the user encrypts/encodes and uploads replicas. Other solutions let \mathcal{S} generate replicas, but the time cost of the replicas generation is noticeable, like RSA-based puzzle combining linear feedback shift registers [13], chained sequential encryption [17], and butterfly construction [15]. But these schemes rely on an assumption that the replicas generation is time-costly and cannot be parallelized. In BOSSA, we propose a new paradigm in which a decentralized storage network atop blockchain is used to store replicas. With the openness of the blockchain, \mathcal{S} 's action (*i.e.*, data replication) is publicly recorded, which is visible to the users. BOSSA neither brings heavy costs to the user, nor relies on the time-assumption. In addition, this paradigm also benefits \mathcal{S} as its storage is reduced (most of clouds' basic storage services promise to store several replicas in the same region).

Decentralized storage. The main purpose of decentralized storage schemes is to reuse idle storage resources of personal

computers (PC) by motivating PC users to share their unused disk space. Existing decentralized storage network schemes, like Storj [30], Sia [31], utilize Merkle tree-based data auditing scheme whose proof size is related to the size of data, which causes high communication cost for auditing large data. IPFS [32] plans to offer an open global p2p storage network, however, it lacks an incentive mechanism such that the storage resources may be provided by few nodes. In addition, it also fails to provide an auditing mechanism to check the integrity of data stored by nodes. Filecoin [23] is built atop IPFS and provides the aforementioned missing features. Except for incentive mechanism, Filecoin leverages a computationally expensive proofs of replica scheme [17] to ensure the correct data processing and builds a consensus mechanism. Similar idea is also presented in [33], [34], [35]. Compared to Filecoin, BOSSA is a more general framework which is compatible with other blockchain platforms like Ethereum.

1.4 Organization

The rest of the paper is organized as follows. Section 2 presents some backgrounds. In Section 3 we define the model of our scheme, threat model and security goals, before we introduce key techniques of our scheme in Section 4. And concrete construction of our scheme is described in Section 5. We analyze our scheme in Section 6, evaluate the performance of our scheme in Section 7. We conclude the paper in Section 8.

2 BACKGROUND

2.1 Blockchain and Smart Contract

Blockchain, a cryptographic primitive derived from Bitcoin proposed by Nakamoto Satoshi [36] in 2008, has been extensively studied and adopted in both academics [37], [38], [39] and industries [40]. Blockchain can be viewed as a distributed block-wise ledger allowing any participants to access and modify (appending only). Generally speaking, it can be generalized as the following formulation:

$$\begin{aligned} \text{Block}_n &= \text{Header}_n \parallel \text{Transactions}, \\ \text{Header}_n &= \text{Hash}(\text{Block}_{n-1}) \parallel \text{Nonce} \parallel \\ &\quad \text{Timestamp} \parallel \text{MerkleRoot}. \end{aligned}$$

A block consists of transactions' data and a header which generally contains information for consensus and security: a *Timestamp*, a hash of predecessor's block, a digest of transactions (*i.e.* *MerkleRoot*), and a *Nonce* used for solving consensus puzzles. Blocks are organized by the hash pointing to the predecessor block, forming a hierarchically expanding chain. Cryptographic mechanisms, like hashing, signature, guarantee that the blockchain is non-repudiation and non-tampering, and consensus models, *e.g.*, proof of work, proof of stake, motivate majority participants to maintain only one correct chain.

Bitcoin, the initial blockchain implementation, provides restricted customizable scripts, which limits it to a "cryptocurrency", instead of a general computational platform. The idea is later brought by the following implementations such as Ethereum [24], [41], HyperLedger [25], [42]. Ethereum, the first cryptocurrency supporting general computing, provides a virtual machine (EVM) and Turing-complete opcodes (EVM Bytes). The computation results of smart contracts inherit the nature of trust from the blockchain, hence Ethereum is widely applied in fields like e-voting [20], auction [21], fair exchange [22], etc.

2.2 Proofs of Retrievability

A Proofs of Retrievability (PoR) scheme typically contains three roles including a user, a remote server (also plays as a prover) and a verifier (*i.e.*, the TPA in public verification or the user itself in the private verification scenario). To prove that data is retrievable, a challenge-response protocol is executed between the prover and the verifier, and the data retrievability can be guaranteed with a high probability if the prover presents a valid response (proof). A PoR scheme [10] can be defined as follows:

Definition 1. A PoR scheme consists of the following four algorithms:

- $\text{Setup}(1^\lambda) \rightarrow (sk, pk)$: This probabilistic algorithm is run by the user. It takes the security parameter as input and outputs key pair (pk, sk) for the scheme setup.
- $\text{Store}(D, sk) \rightarrow (D', tag)$: The user runs this algorithm to encode data and generate metadata for auditing. This algorithm takes as inputs key sk and user's data D , and outputs tags tag and encoded data D' .
- $\text{Prove}(D', tag, pk, chal) \rightarrow \phi$: This algorithm is run by the prover to prove data integrity. It takes as input the verifier's challenge $chal$, outsourced data D' , the corresponding tags tag and public key pk , and returns a proof ϕ .
- $\text{Verify}(\phi, chal, pk, sk) \rightarrow (\top, \perp)$: This algorithm is run by the verifier. It takes as input the proof ϕ , public key pk (sk is needed in private verification), and determines whether the prover stores data honestly.

3 PROBLEM FORMULATION

3.1 System Model

The system model of BOSSA is depicted in Fig. 1. Our system involves the following entities: the cloud server \mathcal{S} , the cloud user \mathcal{U} , and a decentralized network including nodes with two kinds of roles: miners \mathcal{M} and farmers \mathcal{F} . \mathcal{M} represents the nodes who maintain the blockchain and execute smart contracts. \mathcal{F} represents the nodes joining in the decentralized storage network for renting out their idle storage resources.

In BOSSA, \mathcal{S} is willing to collect the original data from \mathcal{U} for accomplishing data analytics tasks such as prediction or recommendation [43], [44], [45]. Similar to existing works [7], [12], the data outsourced to \mathcal{S} is kept in plaintext in order to maximize their potential value (*e.g.*, data mining). To save resources and take full advantage of idle storage resources on \mathcal{F} , \mathcal{S} further outsources replicas to \mathcal{F} . At the same time, \mathcal{U} delegates \mathcal{M} to audit the original data on \mathcal{S} and replicas on \mathcal{F} , respectively.

The core functionalities of BOSSA are defined below.

Definition 2. The algorithms for proofs of retrievability and proofs of replication in BOSSA are defined as follows:

- $\text{Setup}(1^\lambda) \rightarrow (sk_{\mathcal{U}}, pk_{\mathcal{U}}, sk_{\mathcal{S}}, pk_{\mathcal{S}})$: Given the security parameter λ , this probabilistic algorithm outputs private and public key pairs for \mathcal{U} and \mathcal{S} .
- $\mathcal{U}.\text{Store}(D, sk_{\mathcal{U}}) \rightarrow (D^*, aux)$: This probabilistic algorithm takes the original data D and \mathcal{U} 's private key $sk_{\mathcal{U}}$ as input, and outputs D^* and the auxiliary information aux .
- $\mathcal{S}.\text{Seal}(D^*, ek) \rightarrow R$: This algorithm takes data D^* received from the user and an encryption key ek as input, and outputs a replica R .
- $\mathcal{S}.\text{Distribute}(R, sk_{\mathcal{S}}) \rightarrow \{R_i, aux'_i\}_{i \in [0, K]}$: This algorithm, taking as input the replica R and \mathcal{S} 's secret key $sk_{\mathcal{S}}$, outputs K replica blocks R_i with auxiliary information aux'_i .

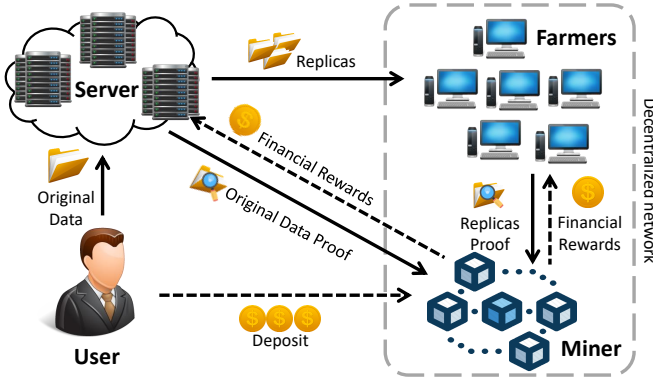


Fig. 1. System Model.

- $S.Prove(D^*, aux) \rightarrow \phi$: This algorithm takes as input the data D^* stored by S , the auxiliary information aux , then outputs the proof ϕ in a privacy-preserving manner.
- $F.Prove(R_i, aux'_i) \rightarrow \phi'_i$: F calculates and outputs the proof ϕ'_i based on replica R_i and aux'_i .
- $Verify_X(pk_X, \phi) \rightarrow \{\top, \perp\}$: This algorithm verifies S 's or F 's proof ϕ and outputs either \top (TRUE) or \perp (FALSE).

3.2 Threat Model

Following previous works [13], [14] we consider a rational S , who behaves correctly in most of the time unless misbehaving brings more benefits, e.g., deleting data blocks seldom accessed for resource-saving. The blockchain network can be regarded as a trusted and public entity. Concretely, most of the computational power of the blockchain is held by honest miners who will execute smart contracts faithfully [46]. Any node in the blockchain network can learn internal states and messages sent to the blockchain.

The farmers from the decentralized network, rent out their idle disk space to store U 's data replicas. We assume that there are massive farmers, the number of which is enough to store all the users' replicas. The farmers may leave the network irresponsibly if they lose the interest of storing replicas. Therefore, it is highly necessary to guarantee the availability of replicas, i.e., motivating the farmers to act honestly.

3.3 Security Goals

Storage Correctness. This property guarantees both outsourced original data and replica are retrievable. We formalize this property by adopting extraction algorithm in [10].

Definition 3. We say that the scheme guarantees storage correctness if the prover cannot forge a proof make the verifier accept, and there is an extraction algorithm such that for an ε -admissible prover (S or F) who can provide ε fraction of a valid proof, an efficient extraction algorithm will recover data with overwhelming probability.

Privacy Preservation. This property guarantees the nodes from the decentralized network (including M and F) cannot infer any private information about the data.

Definition 4. We say the scheme achieves privacy preservation if: 1) for a probabilistic polynomial time (PPT) adversary who acts as the verifier, namely M , it is computationally hard to extract private information from the prover's proof; 2) the replicas stored

in the decentralized network do not leak any private information about the original data.

Replica Reliability. This property guarantees that the original data can be recovered from replica blocks stored by F .

4 KEY TECHNIQUES

4.1 Building Blocks and Notations

4.1.1 Building Blocks

Definition 5. For a q -ary alphabet Σ_q , a (n, k, d) erasure code scheme consists of two algorithms $Encode: \Sigma^k \rightarrow \Sigma^n$ and $Decode: \Sigma^{n-d+1} \rightarrow \Sigma^k$. Typically, we say a code is a maximum distance separable (MDS) code, if $d = n - k + 1$.

Definition 6. Let $e: \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ be a bilinear map where \mathbb{G}_1 , \mathbb{G}_2 and \mathbb{G}_T are three multiplicative cyclic groups of prime order p . It has following properties:

- Given $u \in \mathbb{G}_1$, $v \in \mathbb{G}_2$ and $x, y \in \mathbb{Z}_p$, we have $e(u^x, v^y) = e(u, v)^{xy}$.
- e is non-degenerate, namely, $e(g_1, g_2) \neq 1$ where g_1, g_2 are generators of \mathbb{G}_1 and \mathbb{G}_2 respectively.
- There is an efficient algorithm to compute $e(u, v)$ for any $u \in \mathbb{G}_1$, $v \in \mathbb{G}_2$.

In addition, we define the following collision-resist hash functions:

- $h_1: \{0, 1\}^* \rightarrow \mathbb{G}_1$, it maps arbitrary string to an element of the group \mathbb{G}_1 .
- $h_2: \mathbb{G}_T \rightarrow \mathbb{Z}_p$, it maps elements of \mathbb{G}_T to finite set \mathbb{Z}_p .
- $h_3: \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$, it maps arbitrary string to constant-size bits, where λ is the security parameter.

Finally, we define $\{KGen, Enc, Dec\}$ a symmetric encryption algorithm with chose-plaintext-attack (CPA) security [47].

4.1.2 Other Notations

For ease of expression, we utilize L to represent the distributed ledger in the blockchain system. L is stored and maintained by M . We index an entity in the blockchain by $X.Addr$, e.g., $U.Addr$. In the following discussion, B_{now} represents the latest block number of the blockchain.

The dictionary $L.Roster$ stores $(D, B_{start}, B_{end}, services)$ and is indexed by farmer's address $F.Addr$. The attribute D is defined as the deposit paid by F . The list $services$ records replicas stored by F . B_{start} represents the time when F starts its storage services, and B_{end} represents the time when the services expire.

The dictionary $L.Audit$ is indexed by $name$, the unique identifier of U 's file (see Section 5.2). The attributes of $L.Audit$ consist of the tuple $(B_{last}, D, pk_U, ltr, p, r)$. For a file $name$, the attribute D in $L.Audit(name)$ denotes the deposit paid by U who outsources the file $name$. pk_U is the public key of U . The rest parameters will be explained later in our construction.

$L.Replica$ is a dictionary taking $name_{rep}$ as the key. It consists of the tuple $(D, pk_S, ltr, r, p, \mathcal{L})$, where D represents the deposit provided by U , and pk_S is the public key of S . The list \mathcal{L} consists of $(totalAsk, totalReply, D, sig_i, B_{last})$ and is indexed by $F.Addr$. The attribute D in \mathcal{L} represents the deposit paid by the farmer when it stores replicas. The other parameters will be explained in our construction.

TABLE 1
Notations

Notation	Description
\mathcal{L}	Distributed ledger of the blockchain
$X.Addr$	Address of party X in blockchain
$name$	An unique identifier for outsourced original file
$name_{rep}$	An unique identifier for outsourced replica
$services$	Services \mathcal{F} has accepted
B_{start}	The block number when \mathcal{F} starts its storage services
B_{end}	The block number when \mathcal{F} 's services are ended
B_{sel}	The number of block chosen by \mathcal{S} and \mathcal{F} to derive challenges
B_{now}	The number of latest block
B_{last}	The block of the previous block that chosen as challenge by \mathcal{S} and \mathcal{F}
\mathcal{D}	Deposit secured on the blockchain
\mathfrak{p}	The price that \mathcal{U} is going to give to \mathcal{F} or \mathcal{S}
r	A radius used for relaxing strictness of the proof
ltr	A time interval before both \mathcal{S} and \mathcal{F} present their next proof

4.2 Design Innovations

To prove data retrievability and possession of replicas, BOSSA incorporates several auditing techniques—compact proofs of retrievability of Shacham *et al.* [10] and privacy-preserving data auditing of Wang *et al.* [12]. Meanwhile, we have also made some changes that make our design more suitable for the blockchain environment. First, different from these schemes where the challenges are generated by the auditor, we require that the provers (*i.e.*, \mathcal{S} and \mathcal{F}) generate challenges by themselves since there are no active auditors in our scheme. To ensure an unbiased challenge, we leverage the latest block of the blockchain as a publicly traceable and uncontrollable randomness source. In order to prevent lazy behaviors of provers, we design a time-restricted proof mechanism to force them to prove. Second, directly offloading replicas to the decentralized network may cause privacy leakage. Moreover, if some nodes storing a portion of replicas leave the network, it could result in the uselessness of the whole replicas. Hence, we propose an algorithm $\mathcal{S.Seal}$ to encode replica to be outsourced and then encrypt the private data with the aid of the symmetric encryption algorithm. Last but not least, we extend the structure of the authenticator inside each replica block. Specifically, we add an extra identifier of the replica and the indices of the replica block into authenticators, to prevent \mathcal{F} from pretending to be storing all the data which is not at all. More details can be found in Section 5.2.

4.3 Time-Restricted Proof

We define the time-restricted proof as follows.

Definition 7. We use B_{last} to denote the number of the block from which the prover generates challenges for calculating the last valid proof, and B_{sel} denotes the block number of the block picked for the new proof. Let ltr be a time interval represented by the block numbers and a radius r represent the maximum gap between B_{sel} and B_{now} . Thus, during the verification procedure,




Fig. 2. Illustration of time-restricted proof. The prover must present the proof after $ltr = n$ blocks. As long as the prover's proof is accepted in yellow blocks (*i.e.*, b_{n+1}, b_{n+2} as $r = 2$), it can be regarded as providing a proof on time.

we say the prover provides proof on time only if the following equation holds:

$$B_{last} + ltr = B_{sel} \leq B_{now} \leq B_{last} + ltr + r.$$

We illustrate the time-restricted proof mechanism in Fig. 2. As \mathcal{M} cannot actively audit the prover (\mathcal{S} and \mathcal{F}), we demand that the prover always generates challenges from the latest block (and set $B_{sel} = B_{now}$) and provides the proof to \mathcal{M} periodically (after ltr blocks). The radius r is set for tolerating the processing delay of \mathcal{M} .

4.4 Contribution Rate

Considering the case that some farmers may be not willing to contribute replica blocks when the data is needed, we define a new metric, *contribution rate*, to measure \mathcal{F} 's contribution, which also directly affects \mathcal{F} 's final reward.

Definition 8. Let *totalAsk* be the amount of replica requests during \mathcal{F} 's services, and let *totalReply* be the amount of \mathcal{F} 's responses. Then, the contribution rate cr is calculated as follows:

$$cr = \begin{cases} \frac{totalReply}{totalAsk} & totalAsk > 0 \\ 1 & totalAsk = 0 \end{cases}, cr \in [0, 1].$$

We give an example for a better explanation. When \mathcal{S} gives replica blocks to \mathcal{F}_i , both *totalAsk* and *totalReply* are set to 0. In the case that \mathcal{S} needs replica blocks stored by \mathcal{F}_i , the counter *totalAsk* of \mathcal{F}_i is incremented by 1, and only when \mathcal{F}_i provides the stored data to \mathcal{S} , \mathcal{F}_i 's counter *totalReply* is incremented by 1. Later, when \mathcal{F}_i terminates the storage service, the monetary reward it gets is its deposit multiplied by the contribution rate, and the rest of the deposit is refunded to \mathcal{U} .

4.5 Reward Mechanism

\mathcal{S} will be rewarded for proving storage services. Suppose \mathcal{U} and \mathcal{S} have come to an agreement about a price \mathfrak{p} , each time, \mathcal{S} honestly proves the integrity of data, \mathcal{U} will pay \mathfrak{p} to \mathcal{S} .

Likewise, we adopt the same reward mechanism for \mathcal{F} with an extra modification. Considering the following scenario, after providing several proofs, \mathcal{F} 's reward has far exceeded its total deposit, then it can quit the network arbitrarily without worrying about the financial loss. Therefore, we set a factor $\xi \in (0, 1)$ to prevent this behavior. After a valid proof of \mathcal{F} , \mathcal{M} withholds $(1 - \xi) \cdot \mathfrak{p}$ of monetary reward instead of giving all of them to \mathcal{F} . The frozen part of monetary reward is refunded only when \mathcal{F} 's service expires. A proper selection of ξ (*e.g.*, $\xi = 0.5$) can make \mathcal{F} more sticky to the network without damaging its long-term benefits.

Outsourcing Original Data:

\mathcal{U} : 1: $(sk_{\mathcal{U}}, pk_{\mathcal{U}}, sk_{\mathcal{S}}, pk_{\mathcal{S}}) \leftarrow \mathcal{U}.\text{Setup}(1^\lambda)$
 2: $(D^*, aux) \leftarrow \mathcal{U}.\text{Store}(D, sk_{\mathcal{U}})$
 3: send D^*, aux to \mathcal{S}
 4: put deposit of D to \mathcal{M} and send ori-store, $name$, $pk_{\mathcal{U}}$, p_s , ltr_s and r_s to \mathcal{M}

\mathcal{M} : 1: upon receiving ori-store from \mathcal{U}
 2: initialize an empty item of $L.Audit$ indexed by $name$, set B_{last} as B_{now} , and store $(B_{last}, D, pk_{\mathcal{U}}, ltr_s, p_s, r_s)$ in $L.Audit(name)$
 3: broadcast ori-store-confirmed

Outsourcing Replicas:

\mathcal{U} : 1: send $sk_{\mathcal{S}}, pk_{\mathcal{S}}, name$ and $name_{rep}$ to \mathcal{S}
 2: put deposit of $D_{\mathcal{U}}$ in \mathcal{M} , and send rep-store, $name_{rep}$, $pk_{\mathcal{S}}$, ltr_f , r_f and p_f to \mathcal{M}

\mathcal{M} : 1: upon receiving rep-store from \mathcal{M}
 2: initialize an item of $L.Replica$ indexed by $name_{rep}$, initialize an empty dictionary \mathcal{L} and store the tuple $(D_{\mathcal{U}}, pk_{\mathcal{S}}, ltr_f, p_f, r_f, \mathcal{L})$ in $L.Replica(name_{rep})$
 3: broadcast rep-store-recruit

\mathcal{F} : 1: send deposit of $D_{\mathcal{F}}$ together with rep-store-enroll, $name_{rep}$ to \mathcal{M}

\mathcal{M} : 1: upon receiving rep-store-enroll from \mathcal{F}
 2: take the newest block as a PRF seed sd
 3: **assert** $h_3(\mathcal{F}.Addr \parallel sd) < th$
 4: **assert** $L.Roster(\mathcal{F}.Addr).B_{start} + \Delta < B_{now}$
 5: initialize a tuple $(0, 0, D_{\mathcal{F}}, \perp, \perp)$
 6: store the tuple in $L.Replica(name_{rep}).\mathcal{L}$ with the index $\mathcal{F}.Addr$
 7: broadcast rep-store-enrolled

\mathcal{S} : 1: generate an encryption key ek by KGen
 2: $R \leftarrow \mathcal{S}.\text{Seal}(D^*, ek)$
 3: $\{R_i, aux'_i\}_{1 \leq i \leq K} \leftarrow \mathcal{S}.\text{Distribute}(R, sk_{\mathcal{S}})$
 4: pick K farmers and for each farmer:
 5: $sig_i \leftarrow SSig_{ssk_{\mathcal{S}}}(h_3(R_i))$
 6: send (R_i, aux'_i, sig_i) to the farmer

\mathcal{F} : 1: store replica blocks locally
 2: send rep-store-finalize, sig_i to \mathcal{M}

\mathcal{M} : 1: upon receiving rep-store-finalize from \mathcal{F}
 2: update $L.Replica(name_{rep}).\mathcal{L}(\mathcal{F}.Addr)$ to $(0, 0, D_{\mathcal{F}}, sig_i, B_{now})$
 3: add $name_{rep}$ into $L.Roster(\mathcal{F}.Addr).services$
 4: broadcast rep-store-confirmed

Fig. 6. Data outsourcing.

Proving for reward. Both \mathcal{S} and \mathcal{F} prove data integrity for earning rewards through the time-restricted proof, and they are rewarded based on the reward mechanism proposed in Section 4.5. If the proof is valid, \mathcal{M} sets B_{last} to B_{sel} , then rewards the provers and updates \mathcal{U} 's deposit. Note that \mathcal{M} should check that \mathcal{F} 's service has not expired to prevent \mathcal{F} from cheating.

Fetching replica. When \mathcal{S} tries to reconstruct the original data from replicas, it sends message rep-fetch to \mathcal{M} together with $name_{rep}$. Upon receiving the message from \mathcal{S} , \mathcal{M} increases the $totalAsk$ of farmers by 1 who are in $L.Replica(name_{rep}).\mathcal{L}$, and broadcasts a message rep-fetch-request for acknowledging farmers. A rational farmer will share its local replica block R_i

Proving for Reward (\mathcal{S}):

\mathcal{S} : 1: set $B_{sel} = B_{now}$, $\phi \leftarrow \mathcal{S}.\text{Prove}(D^*, aux)$
 2: send ori-proof, ϕ and B_{sel} to \mathcal{M}

\mathcal{M} : 1: on receiving ori-proof, extract $name$ from ϕ
 2: load B_{last} , ltr , r from $L.Audit(name)$
 3: **assert** $B_{last} + ltr = B_{sel} \leq B_{now} \leq B_{last} + ltr + r$
 4: **if** $\text{Verify}_{\mathcal{S}}(\phi) = \top$ **then**
 5: $L.Audit(name).B_{last} = B_{sel}$
 6: sent monetary reward of $L.Audit(name).p$ to \mathcal{S} from \mathcal{U} 's deposit $L.Audit(name).D$
 7: **end if**
 8: broadcast ori-proof-verified

Proving for Reward (\mathcal{F}):

\mathcal{F} : 1: set $B_{sel} = B_{now}$, $\phi' \leftarrow \mathcal{F}.\text{Prove}(R'_i, aux')$
 2: send rep-proof, ϕ' and B_{sel} to \mathcal{M}

\mathcal{M} : 1: on receiving rep-proof, extract $name_{rep}$ from ϕ'
 2: get B_{last} from $L.Replica(name_{rep}).\mathcal{L}(\mathcal{F}.Addr)$ and get ltr , r from $L.Replica(name_{rep})$
 3: **assert** $B_{now} \leq L.Roster(\mathcal{F}.Addr).B_{end}$
 4: **assert** $B_{last} + ltr = B_{sel} \leq B_{now} \leq B_{last} + ltr + r$
 5: **if** $\text{Verify}_{\mathcal{F}}(\phi') = \top$ **then**
 6: $L.Replica(name_{rep}).\mathcal{L}(\mathcal{F}.Addr).B_{last} = B_{sel}$
 7: load p from $L.Replica(name_{rep}).p$
 8: deduct p from $L.Replica(name_{rep}).D$, send monetary reward of $\xi \cdot p$ to \mathcal{F} and store the rest into $L.Replica(name_{rep}).\mathcal{L}(\mathcal{F}.Addr).D$
 9: **end if**
 10: broadcast rep-proof-verified

Fig. 7. \mathcal{F} and \mathcal{S} proving for reward.

with \mathcal{S} for maximizing his profits (see Section 6.2.3). When \mathcal{S} receives enough replica blocks, it can reconstruct the original data and validate replicas through the comparison with its previously-stored hashes. For those farmers who contribute their local replica blocks, \mathcal{S} will notify \mathcal{M} to increase $totalReply$ with message rep-fetch-finalize. And \mathcal{M} broadcasts rep-fetch-confirmed for notifying farmers.

Farmer logout. When a farmer wants to leave the network and takes back all its deposit (including the frozen part of the reward), the following conditions must be met: 1) it is not storing replicas i.e., $L.Roster(\mathcal{F}.Addr).services = \emptyset$; 2) the storage service it promised should expire, i.e., $B_{end} \leq B_{now}$; 3) it proves correctly each time i.e., $0 \leq B_{end} - B_{last} < ltr$. The first two conditions allow the farmer to take back its deposit during the enrollment, and the last one allows it to obtain the deposit for storing replica blocks (which is refunded based on the contribution rate). Besides, a new farmer must be found for restoring its replicas. Only after transferring all the replicas to the new farmer, can it get back the frozen deposit. Details can be found in Fig. 8.

6 SYSTEM ANALYSIS

6.1 Choice of Challenge Size

The challenges against both \mathcal{S} and \mathcal{F} are based on the probabilistic framework [5], through which the workload of the provers can be noticeably reduced. This is because, during the auditing, the provers randomly sample a subset of stored data, as required by the verifier, instead of looking through all of the data. Meanwhile,

Replicas Transferring:

\mathcal{F}_s : 1: calculate $cm = h_3(r)$ based on a randomly picked seed r
2: send **rep-transfer**, $name_{rep}$ and cm to \mathcal{M}

\mathcal{M} : 1: upon receiving **rep-transfer**
2: load ltr from $L.Replica(name_{rep})$, load B_{last} from $L.Replica(name_{rep}).\mathcal{L}(\mathcal{F}_s.Addr)$, and load B_{end} from $L.Roster(\mathcal{F}_s.Addr)$
3: **assert** $B_{end} \leq B_{now}$
4: **assert** $0 \leq B_{end} - B_{last} < ltr$
5: store cm and publish $name_{rep}$
6: broadcast **rep-transfer-request**

\mathcal{F}_d : 1: upon receiving **rep-transfer-request** from \mathcal{M}
2: request r and replica block from \mathcal{F}_s
3: access $L.Replica(name_{rep})$ and load sig_i from $\mathcal{L}(\mathcal{F}_s.Addr)$
4: **assert** replica block is valid based on sig_i
5: put deposit of D to \mathcal{M} and send **rep-transfer-finalize**, r , $\mathcal{F}_s.Addr$ and $name_{rep}$ to \mathcal{M}

\mathcal{M} : 1: upon receiving **rep-transfer-finalize**
2: **assert** $cm = h_3(r)$
3: remove $name_{rep}$ from $L.Roster(\mathcal{F}_s.Addr).services$
4: remove \mathcal{F}_s from $L.Replica(name_{rep}).\mathcal{L}$
5: calculate contribution rate cr of \mathcal{F}_s
6: refund $L.Replica(name_{rep}).\mathcal{L}(\mathcal{F}_d.Addr).D$ to \mathcal{F}_s based on the contribution rate cr , and rest of deposit is refunded to \mathcal{U}
7: add a new tuple $(0, 0, D, sig_i, B_{now})$ into $L.Replica(name_{rep}).\mathcal{L}(\mathcal{F}_d.Addr)$
8: broadcast **rep-transfer-confirmed**

Farmer Logout:

\mathcal{F}_s : 1: give all of local replica blocks to new farmers
2: send **logout-request** to \mathcal{M} for logout

\mathcal{M} : 1: upon receiving **logout-request** from \mathcal{M}
2: load B_{end} from $L.Roster(\mathcal{F}_s.Addr)$
3: **assert** $B_{end} \leq B_{now}$
4: **assert** $L.Roster(\mathcal{F}_s.Addr).services$ is empty
5: give \mathcal{F}_s 's deposit back
6: broadcast **logout-confirmed**

Fig. 8. \mathcal{F} leaves the network.

with the aid of unbiased-random challenges (usually picked by the verifier), the cheater will be caught with an overwhelming probability, if the data loss is beyond a bound.

Supposing that the prover tampers with $n \cdot t$, $t \in (0, 1)$ blocks out of n data blocks in total, and he is asked to sample c blocks. We let X denote a discrete random variable representing the number of deleted blocks chosen during the proof, and we define P_X as the corresponding probability of X . Thus we have:

$$\begin{aligned}
 P_X &= P\{X \geq 1\} \\
 &= 1 - P\{X = 0\} \\
 &= 1 - \frac{n - n \cdot t}{n} \cdot \frac{n - 1 - n \cdot t}{n - 1} \cdots \frac{n - c + 1 - n \cdot t}{n - c + 1} \\
 &\geq 1 - (1 - t)^c.
 \end{aligned}$$

Clearly, if we fix the number of t , P_X is independent of the number of data blocks. For example, when $t = 1\%$, P_X can

reach 95% as long as the number of sampled blocks is not less than 300 (see Fig. 9).

The workloads of proving and verification are related to the number of challenges, namely the size of c . However, since both original data and replica are encoded with the erasure code, we can sacrifice storage space in exchange for efficient verification and proving. Recall that a (n_2, k_2) encoding is applied on each column of replica which is individually stored by farmers, and the redundancy rate is n_2/k_2 . If we raise the redundant rate, we could take a less rigorous challenge with a smaller challenge size. For example, if we take a $2\times$ redundancy-rate encoding, we can aggressively reduce challenge size to 7, as the data loss of 50% can be detected with a 99% probability (as depicted in Fig. 9) while the whole replica can still be recovered.

6.2 Security Analysis

6.2.1 Storage Correctness

Definition 9. The Computational Diffie-Hellman assumption is define as follows. Given a tuple $(g, g^x, g^y) \in \mathbb{G}$, where \mathbb{G} is a multiplicative cyclic group in order of p , g is a generator of \mathbb{G} , and $x, y \stackrel{R}{\leftarrow} \mathbb{Z}_p$, then the probability that a PPT adversary \mathcal{A} outputs g^{xy} is no more than a negligible probability ϵ :

$$Pr[\mathcal{A}(g, g^x, g^y) = g^{xy}] < \epsilon.$$

Theorem 1. In the random oracle model, if CDH assumption holds in bilinear groups, then BOSSA guarantees the storage correctness.

Proof (sketch). Firstly, we prove that an honest \mathcal{S} can always pass the verification. To show that, for the equation (1), we have:

$$\begin{aligned}
 LHE &= T \cdot e(\Sigma^\gamma, g_2) \\
 &= e(g_*, g_2)^{r_\delta} \cdot e(\bar{g}_U, v_U)^{r_m} \cdot \\
 &\quad e\left(\prod_{i \in I} (h_1(W_i)^{\nu_i} \cdot \bar{g}_U^{m_i \cdot \nu_i})^{x_U \cdot \gamma} \cdot g_*^{\rho \gamma}, g_2\right) \\
 &= e(g_*, g_2)^{r_\sigma + \rho \gamma} \cdot e\left(\prod_{i \in I} h_1(W_i)^{\nu_i \cdot \gamma} \cdot \bar{g}_U^{r_m + \gamma \mu'}, v_U\right) \\
 &= e(g_*, g_2)^\varsigma \cdot e\left(\left(\prod_{i \in I} h_1(W_i)^{\nu_i}\right)^\gamma \cdot \bar{g}_U^{\mu'}, v_U\right) = RHE.
 \end{aligned}$$

Thus, this property ensures the soundness of BOSSA, and guarantees the interests of \mathcal{S} . Then, we prove that for an ϵ -admissible \mathcal{S} who can correctly provide ϵ -fraction proofs, original data can be extracted from it.

We assume the verifier controls a random oracle h_2 and it answers hash queries from \mathcal{S} . Upon receiving the proof from \mathcal{S} , the verifier rewinds to the point when \mathcal{S} queries $h_2(T)$, and sends a different $\gamma^* = h_2(T)$ to \mathcal{S} . Having two different proofs (μ, Σ, ς) and $(\mu^*, \Sigma, \varsigma^*)$, the verifier extracts μ' and σ by calculating

$$\begin{aligned}
 \mu' &= \frac{\mu - \mu^*}{\gamma - \gamma^*} = \frac{(r_m + \gamma \mu') - (r_m + \gamma^* \mu')}{\gamma - \gamma^*}, \\
 \rho &= \frac{\varsigma - \varsigma^*}{\gamma - \gamma^*} = \frac{(r_m + \gamma \rho) - (r_m + \gamma^* \rho)}{\gamma - \gamma^*}, \\
 \sigma &= \frac{\Sigma}{g_*^\rho}.
 \end{aligned}$$

The pair (μ', σ) is the proof for verification in [10], and based on Theorem 4.2 in [10], if \mathcal{S} forges both μ' and σ , it cannot pass

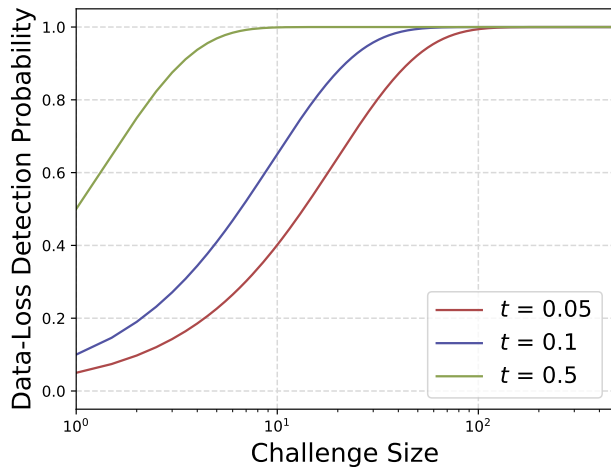


Fig. 9. Data-loss detection probability vs. challenge size. We give an illustration about challenge sizes for achieving various loss-detection probabilities under different degrees of t .

the verification, except with a negligible probability. Then, based on Theorem 4.3 in [10], if \mathcal{S} gives at least ε valid proofs to the verifier, the data can be recovered from the proofs through an extraction algorithm.

As for an ε -admissible \mathcal{F} , the storage correctness is achieved based on theorems in [10]. \square

6.2.2 Privacy Preservation

Theorem 2. *In the random oracle model, if Enc is CPA-secure and Diffie-Hellman problem is hard, BOSSA achieves privacy preservation defined in Definition 4.*

Proof (sketch). The replicas of original data are encrypted through Enc before being distributed, hence participants of the decentralized network cannot learn private information as long as they don't have the encryption key, and the auditing procedure also leaks no private information. So, the point is that there should be no privacy leakage during auditing \mathcal{S} .

To prove privacy preservation holds for \mathcal{S} , we show a simulator without knowing μ' and σ can provide a valid proof to \mathcal{M} . We assume it controls the random oracle h_2 and answers queries from \mathcal{M} . To provide a valid proof to \mathcal{M} , the simulator randomly picks γ and $\mu' \in \mathbb{Z}_p$ and $\sigma \in \mathbb{G}_1$, and calculates μ , ζ and Σ . Then, the simulator lets

$$T = e \left(\left(\prod_{i \in I} h_1(W_i)^{\nu_i} \right)^\gamma \cdot \bar{g}_U^\mu, v_U \right) \cdot \frac{e(g_*, g_2)^\zeta}{e(\Sigma^\gamma, g_2)}$$

and fills the random oracle $h_2(T)$ with γ . The simulator's proof is (T, ζ, μ, Σ) . It is clear that the simulator's proof is valid if the simulator answers γ when \mathcal{M} queries $h_2(T)$. The simulator knows nothing about μ' and σ , it means the proof leaks no information about \mathcal{S} . \square

6.2.3 Replica Reliability

According the analysis in Section 6.2.1, as long as \mathcal{F} provides a valid proof, it can convince the verifier that the data is stored properly. According to our reward distribution mechanism, rewards from the user will be given to \mathcal{F} . Since partial rewards

TABLE 2
Overheads of Proof Generation and Verification for 100M Data

	\mathcal{S}		\mathcal{F}
File size (MB)	100		100
Sample blocks	300	480	7
Proving (ms)	33.2	48.8	10.8
Verification (gas)	27,063,392	42,979,515	997,151
USD	12.18	19.34	0.45

are temporarily frozen in the deposit of \mathcal{F} , the more \mathcal{F} proves, the more deposit it will accumulate. When it wants to get the deposit back, the deposit is refunded based on the contribution rate cr , which is determined by whether \mathcal{F} contributes its local replicas when there is a request for replicas from \mathcal{S} . Hence, a rational \mathcal{F} who is eager to maximize its profits will actively contribute its local replica blocks.

Finally, we remark that as replicas are encoded by erasure codes, when \mathcal{F} goes offline accidentally during the replica fetching, BOSSA can still guarantee the reliability of the replicas, as long as the geographical locations of the farmers are sufficiently dispersed.

7 IMPLEMENTATION & EVALUATION

7.1 Implementation Setup

All the parties in our prototype are instantiated on a local Ubuntu 18.04 LTS with an Intel Core i5-8400 CPU, 12GB of RAM, and a 7200 RPM Seagate 1 TB hard disk. We adopt Ethereum as the underlying blockchain system in our implementation. The smart contracts are written using Solidity and deployed to Ganache-cli (a local Ethereum simulation) and Ropsten (an official testnet). Other parties are implemented with JavaScript combined with C/C++. We utilize Keccak256 as the hash function and 128-bit-of-security Barreto-Naehrig curve as the ellipse curve. We use ate-pairing [49] for bilinear pairing since it is based on the elliptic curve suitable for Ethereum. And we use Jerasure, an instance of Reed-Solomon coding, for implementing erasure codes. To evaluate economic costs, we take the exchange rate of $3 \cdot 10^{-9}$ Ether per gas (namely 3 Gwei per gas), and 150 USD per Ether (December, 2019).

7.2 Circumventing Gas Limitation of Ethereum

To prevent malicious users from launching DDoS attacks by sending infinite loops to miners, Ethereum uses *gas* to measure overheads for running smart contracts, and sets a limitation upon it. To circumvent the gas limitation, the implementation of our prototype is slightly different from our design. In BOSSA, the main gas costs are caused by the multiplications and hashes related to the size of challenges. In our prototype, we split those procedures into several sub-functions whose gas costs are below the gas bound. For example, if it is required to perform 300 times of \mathbb{G}_1 multiplications, we then run sub-functions 10 times, each of which runs \mathbb{G}_1 multiplication 30 times. This enables us to avoid exceeding the gas limitation. Note that this substitution increases total gas costs since we must store temporary variables and trigger multiple calls of the smart contracts.

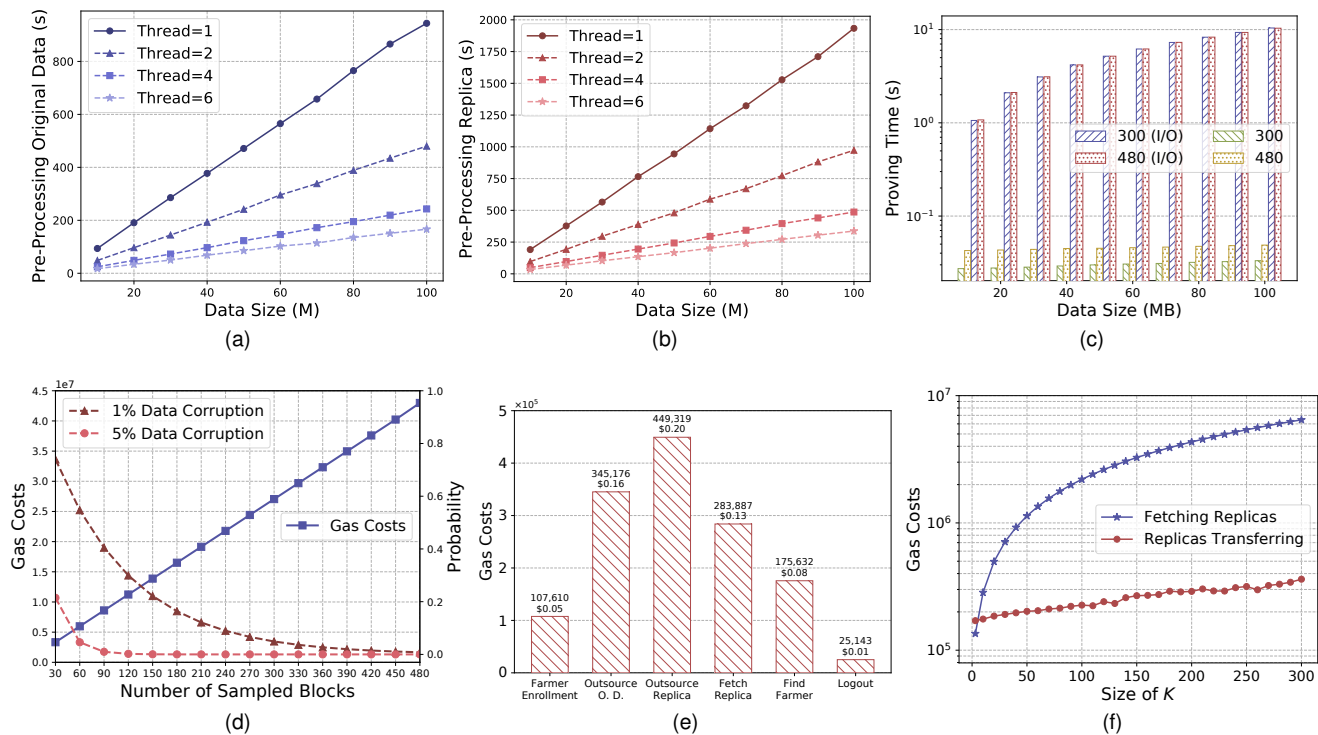


Fig. 10. Time and gas costs of BOSSA. (a) Time costs for pre-processing original data using multiple threads vs. data size. (b) Time costs for pre-processing one replica using multiple threads vs. data size. (c) The time costs of \mathcal{S} 's proof generation under different challenge sizes with and without I/O. (d) The probabilities of failing to detect data corruption and the proof verification gas costs vs. the number of sampled blocks. (e) Gas costs of \mathcal{M} ' actions when replica is stored by $K = 10$ farmers (O.D. stands for Original Data). (f) The gas cost of replica fetching and transferring vs. the number of farmers storing replica.

7.3 Experimental Results

Pre-processing. We firstly test time for generating key pairs (*i.e.*, the \mathcal{U} .Setup). The time cost is about 0.84ms, which is negligible. In Fig. 10(a), we measure the time costs of \mathcal{U} for pre-processing the original data before outsourcing it to \mathcal{S} (*i.e.*, the \mathcal{U} .Store algorithm) and Fig. 10(b) presents the time cost of \mathcal{S} to pre-process one replica before distributing it to farmers (*i.e.*, the algorithms of \mathcal{S} .Distribute and \mathcal{S} .Seal). We randomly generate testing data ranging from 10MB to 100MB. The result shows that the time cost increases linearly with the data size. The time costs are dominated by authenticators generation, while both replica encryption and replica coding incur negligible costs. The generation of authenticators of 100M size costs about 164s on \mathcal{U} , and 330s for a replica of 200M (as we encode original data into double size) on \mathcal{S} .

Proof generation and on-chain verification. The overheads of proof generation on \mathcal{S} and \mathcal{F} are only related to the size of the challenge. We test the overheads of sampling 7 sectors of a replica block, and sampling both 300 and 480 (for about 99% loss detection probability) blocks of original data as discussed in Section 6.1. The results, shown in Table 2, demonstrate our proving process is efficient, and it only costs tens of milliseconds. In Fig. 10(c), we evaluate the I/O effect on the proof generation on \mathcal{S} . We can see that the data size greatly affects the time cost. This is due to the fact that the time cost of proof generation is mainly dominated by the I/O cost (*e.g.*, reading data and authenticators from hard drivers).

For the on-chain verification, the gas costs are 27,063,392 and 42,979,515 when sampling 300 blocks and 480 blocks (Fig. 10(d)), respectively. On the other hand, the gas cost of verifying

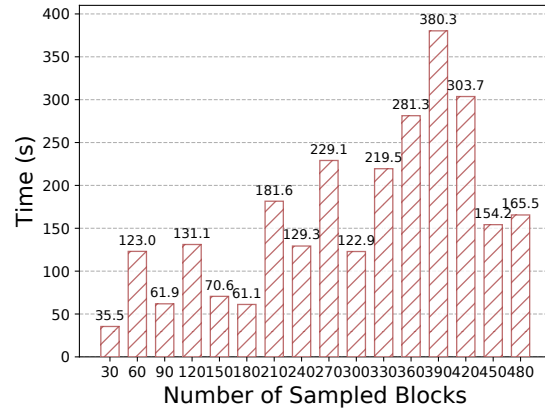


Fig. 11. Time cost of verification in Ropsten vs. the number of sampled blocks.

a farmer's proof is much lower, which is 997,151 gas (*i.e.*, about \$0.45). Note that this cost is for a single farmer, and the total costs are K times larger when there are K farmers to store replica. Since the underlying blockchain platform may affect the time cost of the verification process, we further deploy the smart contracts to Ropsten, an official public Ethereum testnet. We measure the time cost from sending the proof to recording verification result on the chain. The time cost for verifying \mathcal{S} 's proof is presented in Fig. 11. It can be seen that although heavily affected by the throughput of the testnet and the latency of the network, the time cost still roughly increases with the number of the sampled blocks. Note that techniques like sharding [50], [51] and state channel [52] can

further improve the efficiency of verification.

Interactions with the smart contract. In BOSSA, other processes, such as initialization, data fetching, also need to interact with \mathcal{M} (through the smart contracts). In Fig. 10(e), we evaluate the gas costs of those processes. We can see that those costs are much lower (about several cents) than that of proof verification. Meanwhile, among these processes, the costs of outsourcing original data and replicas are the highest. This is caused by the large call parameters (420B and 460B) for storing U 's and S 's public keys respectively.

Fig. 10(f) shows the impact of the number of farmers on the gas costs when fetching and transferring replicas. The gas costs increase linearly with the number of the farmers. This is due to the limitation of data structures supported in Ethereum. We add farmers' addresses into an array, such that it needs to traverse the whole array for searching a specific farmer. In addition, we have to simultaneously increase the counter *totalAsk* for each farmer by 1 when fetching replicas, and it charges an extra fee due to the changing of the blockchain states. To avoid exceeding the gas limitation, S can fetch replica blocks separately.

8 CONCLUSION

In this paper, we proposed BOSSA, a novel decentralized framework for proofs of data replication and retrievability. Different from existing schemes, BOSSA is built atop off-the-shelf blockchain platforms, where each participant is fairly treated and incentivized to faithfully follow the auditing protocol. We proposed a proof enforcement mechanism to catch lazy behaviors, and a new metric as well as a reward distribution mechanism to create a fair reciprocal environment. Our experiments show that BOSSA incurs tolerable costs and is feasible in practice.

ACKNOWLEDGMENTS

Qian Wang's research was supported in part by the NSFC under Grants 61822207 and U1636219, in part by the Outstanding Youth Foundation of Hubei Province under Grant 2017CFA047, and in part by the Fundamental Research Funds for the Central Universities under Grant 2042019kf0210. Cong Wang's research was supported in part by Research Grants Council of Hong Kong under Grants CityU 11212717 and CityU 11217819, and by the Innovation and Technology Commission of Hong Kong under ITF Project ITS/145/19. Shengshan Hu's research was supported by the Fundamental Research Funds for the Central Universities under Grant 2020kfyXJS075. Qian Wang is the corresponding author.

REFERENCES

- [1] "Data protection in amazon s3," <https://docs.aws.amazon.com/AmazonS3/latest/dev/DataDurability.html>.
- [2] H. Zhou, X. Ouyang, Z. Ren, J. Su, C. de Laat, and Z. Zhao, "A blockchain based witness model for trustworthy cloud service level agreement enforcement," in *Proc. of INFOCOM*. IEEE, 2019, pp. 1567–1575.
- [3] "Tencent cloud user claims \$1.6 million compensation for data loss," <https://technode.com/2018/08/06/tencent-cloud-user-claims-1-6-million-compensation-for-data-loss>.
- [4] "Google loses data as lightning strikes," <https://www.bbc.com/news/technology-33989384>.
- [5] G. Ateniese, R. Burns, R. Curtmola, J. Herring, L. Kissner, Z. Peterson, and D. Song, "Provable data possession at untrusted stores," in *Proc. of CCS*. ACM, 2007, pp. 598–609.

- [6] A. Yang, J. Xu, J. Weng, J. Zhou, and D. S. Wong, "Lightweight and privacy-preserving delegatable proofs of storage with data dynamics in cloud storage," *IEEE Transactions on Cloud Computing*, 2018.
- [7] J. Xu, A. Yang, J. Zhou, and D. S. Wong, "Lightweight delegatable proofs of storage," in *Proc. of ESORICS*. Springer, 2016, pp. 324–343.
- [8] F. Armknecht, J.-M. Bohli, G. O. Karame, Z. Liu, and C. A. Reuter, "Outsourced proofs of retrievability," in *Proc. of CCS*. ACM, 2014, pp. 831–843.
- [9] A. Juels and B. S. Kaliski Jr, "Pors: Proofs of retrievability for large files," in *Proc. of CCS*. ACM, 2007, pp. 584–597.
- [10] H. Shacham and B. Waters, "Compact proofs of retrievability," *Journal of Cryptology*, vol. 26, no. 3, pp. 442–483, 2013.
- [11] Q. Wang, C. Wang, K. Ren, W. Lou, and J. Li, "Enabling public auditability and data dynamics for storage security in cloud computing," *IEEE Transactions on Parallel and Distributed Systems*, vol. 22, no. 5, pp. 847–859, 2011.
- [12] C. Wang, S. S. Chow, Q. Wang, K. Ren, and W. Lou, "Privacy-preserving public auditing for secure cloud storage," *IEEE Transactions on Computers*, vol. 62, no. 2, pp. 362–375, 2013.
- [13] F. Armknecht, L. Barman, J.-M. Bohli, and G. O. Karame, "Mirror: Enabling proofs of data replication and retrievability in the cloud," in *Proc. of USENIX Security*, 2016, pp. 1051–1068.
- [14] R. Curtmola, O. Khan, R. Burns, and G. Ateniese, "Mr-pdp: Multiple-replica provable data possession," in *Proc. of ICDCS*. IEEE, 2008, pp. 411–420.
- [15] I. Leontiadis and R. Curtmola, "Secure storage with replication and transparent deduplication," in *Proc. of CODASPY*. ACM, 2018, pp. 13–23.
- [16] I. Damgard, C. Ganesh, and C. Orlandi, "Proofs of replicated storage without timing assumptions," in *Proc. of CRYPTO*. Springer, 2019, pp. 355–380.
- [17] J. Benet, D. Dalrymple, and N. Greco, "Proof of replication," Tech. Rep., 2017.
- [18] B. Fisch, "Poreps: Proofs of space on useful data," Tech. Rep., 2018.
- [19] B. Fisch, J. Bonneau, N. Greco, and J. Benet, "Scaling proof-of-replication for filecoin mining," Tech. Rep., 2018.
- [20] A. B. Aayed, "A conceptual secure blockchain-based electronic voting system," *International Journal of Network Security & Its Applications*, vol. 9, no. 3, pp. 1–9, 2017.
- [21] S. Wu, Y. Chen, Q. Wang, M. Li, C. Wang, and X. Luo, "Cream: a smart contract enabled collusion-resistant e-auction," *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 7, pp. 1687–1701, 2018.
- [22] S. Dziembowski, L. Eeckey, and S. Faust, "Fairswap: How to fairly exchange digital goods," in *Proc. of CCS*. ACM, 2018, pp. 967–984.
- [23] "Filecoin," <https://filecoin.io/filecoin.pdf>.
- [24] G. Wood, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum project yellow paper*, vol. 151, pp. 1–32, 2014.
- [25] C. Cachin, "Architecture of the hyperledger blockchain fabric," in *Workshop on distributed cryptocurrencies and consensus ledgers*, vol. 310, no. 4, 2016.
- [26] C. Guan, K. Ren, F. Zhang, F. Kerschbaum, and J. Yu, "Symmetric-key based proofs of retrievability supporting public verification," in *Proc. of ESORICS*. Springer, 2015, pp. 203–223.
- [27] D. Boneh, B. Lynn, and H. Shacham, "Short signatures from the weil pairing," in *Proc. of ASIACRYPT*. Springer, 2001, pp. 514–532.
- [28] Z. Hao and N. Yu, "A multiple-replica remote data possession checking protocol with public verifiability," in *Proc. of ISDPE*. IEEE, 2010, pp. 84–89.
- [29] D. Boneh, J. Bonneau, B. Bünz, and B. Fisch, "Verifiable delay functions," in *Proc. of CRYPTO*. Springer, 2018, pp. 757–788.
- [30] "Storj," <https://storj.io/white-paper>.
- [31] "Sia: Simple decentralized storage," <https://sia.tech/sia.pdf>.
- [32] J. Benet, "IpfS-content addressed, versioned, p2p file system," *arXiv preprint arXiv:1407.3561*, 2014.
- [33] A. Miller, A. Juels, E. Shi, B. Parno, and J. Katz, "Permacoin: Repurposing bitcoin work for data preservation," in *Proc. of S&P*. IEEE, 2014, pp. 475–490.
- [34] B. Sengupta, S. Bag, S. Ruj, and K. Sakurai, "Retricoin: Bitcoin based on compact proofs of retrievability," in *Proc. of ICDCN*. ACM, 2016, pp. 1–10.
- [35] S. Ruj, M. S. Rahman, A. Basu, and S. Kiyomoto, "Blockstore: A secure decentralized storage framework on blockchain," in *Proc. of AINA*. IEEE, 2018, pp. 1096–1103.
- [36] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008.

- [37] A. Kiayias, A. Russell, B. David, and R. Oliynykov, "Ouroboros: A provably secure proof-of-stake blockchain protocol," in *Proc. of CRYPTO*. Springer, 2017, pp. 357–388.
- [38] S. Hu, C. Cai, Q. Wang, C. Wang, X. Luo, and K. Ren, "Searching an encrypted cloud meets blockchain: A decentralized, reliable and fair realization," in *Proc. of INFOCOM*. IEEE, 2018, pp. 792–800.
- [39] S. Hu, C. Cai, Q. Wang, C. Wang, Z. Wang, and D. Ye, "Augmenting encrypted search: A decentralized service realization with enforced execution," *IEEE Transactions on Dependable and Secure Computing*, vol. PP, no. 99, pp. 1–1, DOI: 10.1109/TDSC.2019.2957091, 2019.
- [40] "Libra," <https://libra.org/en-US/white-paper/>.
- [41] V. Buterin *et al.*, "A next-generation smart contract and decentralized application platform," *White paper*, 2014.
- [42] E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. De Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich *et al.*, "Hyperledger fabric: a distributed operating system for permissioned blockchains," in *Proc. of EuroSys*. ACM, 2018, pp. 1–15.
- [43] S. Hu, L. Y. Zhang, Q. Wang, Z. Qin, and C. Wang, "Towards private and scalable cross-media retrieval," *IEEE Transactions on Dependable and Secure Computing*, vol. PP, no. 99, pp. 1–1, DOI: 10.1109/TDSC.2019.2926968, 2019.
- [44] Q. Wang, M. He, M. Du, S. S. M. Chow, R. W. F. Lai, and Q. Zou, "Searchable encryption over feature-rich data," *IEEE Transactions on Dependable and Secure Computing*, vol. 15, no. 3, pp. 496–510, 2018.
- [45] C. Wang, Q. Wang, K. Ren, N. Cao, and W. Lou, "Toward secure and dependable storage services in cloud computing," *IEEE Transactions on Services Computing*, vol. 5, no. 2, pp. 220–232, 2012.
- [46] M. Saad, J. Spaulding, L. Njilla, C. Kamhoua, S. Shetty, D. Nyang, and A. Mohaisen, "Exploring the attack surface of blockchain: A systematic overview," *arXiv preprint arXiv:1904.03487*, 2019.
- [47] J. Katz and Y. Lindell, *Introduction to modern cryptography*. Chapman and Hall/CRC, 2014.
- [48] I. S. Reed and G. Solomon, "Polynomial codes over certain finite fields," *Journal of the Society for Industrial and Applied Mathematics*, vol. 8, no. 2, pp. 300–304, 1960.
- [49] "ate-pairing," <https://github.com/herumi/ate-pairing>.
- [50] L. Luu, V. Narayanan, C. Zheng, K. Baweja, S. Gilbert, and P. Saxena, "A secure sharding protocol for open blockchains," in *Proc. of CCS*. ACM, 2016, pp. 17–30.
- [51] M. Zamani, M. Movahedi, and M. Raykova, "Rapidchain: scaling blockchain via full sharding," in *Proc. of CCS*. ACM, 2018, pp. 931–948.
- [52] A. Miller, I. Bentov, R. Kumaresan, and P. McCorry, "Sprites and state channels: Payment networks that go faster than lightning," in *Proc. of FC*. Springer, 2019, pp. 508–526.



Dian Chen received the B.E. degree from Wuhan University, Wuhan, China, in 2018. He is working toward Master degree in the School of Cyber Science and Engineering in Wuhan University. His research interests include information security and blockchain technology.



Haobo Yuan is an undergraduate student majoring in Computer Science and Technology at Wuhan University. His research interests include information security and blockchain technology.



Qian Wang received the Ph.D. degree from the Illinois Institute of Technology, Chicago, IL. He is currently a professor with the School of Cyber Science and Engineering, Wuhan University. His research interests include AI security, data storage, search and computation outsourcing security and privacy, wireless system security, and applied cryptography etc. He received National Science Fund for Excellent Young Scholars of China, in 2018. He is also an expert under National "1000 Young Talents Program" of China.

He is a recipient of the 2018 IEEE TCSC Award for Excellence in Scalable Computing (Early Career Researcher), and the 2016 IEEE Asia-Pacific Outstanding Young Researcher Award. He is also a co-recipient of several best paper awards from IEEE DSC'19, IEEE ICDCS'17, IEEE TrustCom'16, WAIM'14, and IEEE ICNP'11 etc. He serves as associate editors for the IEEE Transactions on Dependable and Secure Computing (TDSC), IEEE Transactions on Information Forensics and Security (TIFS), and IEEE Internet of Things Journal (IoT-J). He is a Senior Member of the IEEE and a Member of the ACM.



Shengshan Hu received the B.E. and Ph.D. degrees in computer science and technology from Wuhan University in 2014 and 2019, respectively. He is currently an Associate Professor in the School of Cyber Science and Engineering, Huazhong University of Science and Technology. His research interest focuses on privacy-enhancing technologies, AI security, and blockchain.



Cong Wang is an Associate Professor in the Department of Computer Science, City University of Hong Kong. His current research interests include data and network security, blockchain and decentralized applications, and privacy-enhancing technologies. He is one of the Founding Members of the Young Academy of Sciences of Hong Kong. He received the Outstanding Researcher Award (junior faculty) in 2019, the Outstanding Supervisor Award in 2017 and the President's Awards in 2019 and 2016, all from

City University of Hong Kong. He is a co-recipient of the IEEE INFOCOM Test of Time Paper Award 2020, Best Student Paper Award of IEEE ICDCS 2017, and the Best Paper Award of IEEE ICPADS 2018 and MSN 2015. His research has been supported by multiple government research fund agencies, including National Natural Science Foundation of China, Hong Kong Research Grants Council, and Hong Kong Innovation and Technology Commission. He serves/has served as associate editor for IEEE Transactions on Dependable and Secure Computing, IEEE Internet of Things Journal and IEEE Networking Letters, and TPC co-chairs for a number of IEEE conferences/workshops. He is a senior member of the IEEE, and member of the ACM.