

Accountable Storage

Giuseppe Ateniese¹, Michael T. Goodrich², Vassilios Lekakis³,
Charalampos Papamanthou³, Evripidis Paraskevas³, and Roberto Tamassia⁴

¹ Stevens Institute of Technology

² University of California, Irvine

³ University of Maryland

⁴ Brown University

Abstract. We introduce *Accountable Storage* (AS), a framework enabling a client to outsource n file blocks to a server while being able (any time after outsourcing) to provably compute how many bits were discarded or corrupted by the server. Existing techniques (e.g., proofs of data possession or storage) can address the accountable storage problem, with linear server computation and bandwidth. Instead, our optimized protocols achieve $O(\delta \log n)$ complexity (where δ is the maximum number of corrupted blocks that can be tolerated) through the novel use of invertible Bloom filters and a new primitive called *proofs of partial storage*. With accountable storage, a client can request a compensation of a dollar amount proportional to the number of corrupted bits (that he can now provably compute). We integrate our protocol with Bitcoin, supporting automatic such compensations. Our implementation is open-source and shows our protocols perform well in practice.

1 Introduction

Cloud computing is revolutionizing our digital world, posing new security and privacy challenges. E.g., businesses and individuals are reluctant to outsource their databases for fear of having their data lost or damaged. Thus, they would benefit from technologies that would allow them to manage their risk of data loss, just like insurance allows them to manage their risk of physical or financial losses, e.g., from fire or liability.

As a first step, a client needs a mechanism for verifying that a cloud provider is storing her entire database intact, and fortunately, *Provable Data Possession* (PDP) [3, 11, 13]) and *Proofs of Retrievability* (POR) [10, 19, 26, 27, 28], have been conceived as a solution to the integrity problem of remote databases. PDP and POR scheme can verify whether the server possesses the database originally uploaded by the client by having the server generate a proof in response to a challenge.

However, they leave unsettled several risk management issues. Arguably, an important question is

What happens if a PDP or POR scheme shows that a client's outsourced database has been damaged?

The objective of this work is to design new efficient protocols for *Accountable Storage* (AS) that enable the client to reliably and quickly assess the damage and at the same time automatically get compensated using the Bitcoin protocol.

To be precise, suppose Alice outsources her file blocks b_1, b_2, \dots, b_n to a potentially malicious cloud storage provider, Bob. Since Alice does not trust Bob, she wishes, at any point in time, to be able to compute the amount of damage, if any, that her file blocks have undergone, by engaging in a simple challenge-response protocol with Bob. For instance, she wishes to *provably* compute the value of a *damage metric*, such as

$$d = \sum_{i=1}^n w_i \cdot \|b_i \oplus b'_i\|, \quad (1)$$

where b'_i is the file currently stored by Bob at the time of the challenge, $\|\cdot\|$ denotes Hamming distance and w_i is a weight corresponding to file b_i . If $d = 0$, Alice is entitled to no dollar credit. Bob can easily prove to Alice that this is the case through existing protocols, as noted above. If $d > 0$, however, then Alice should receive a compensation proportional to the damage d , which should be provided automatically.⁵

Naive Approaches for AS. A PDP protocol [3, 4, 11, 13, 29] enables a server to prove to a client that all of the client’s data is stored intact. One could design an AS protocol by using a PDP protocol only for the portion of storage that the server possesses. This could determine the damage, d (e.g., when all weights w_i are equal to 1). However this approach requires using of PDP at the bit level, and in particular computing one 2048-bit tag for each bit of our file collection which is very storage-inefficient.

To overcome the above problem, one could use PDP at the block level, but at the same time keep some redundancy locally. Specifically, before outsourcing the n blocks at the server, the client could store δ extra check blocks locally (e.g., computed with an error-correcting code). The client could then verify through PDP that a set of at most δ blocks have gone missing and retrieve the lost blocks by executing the decoding algorithm on the remote intact $n - \delta$ data blocks and the δ local check blocks (then the recovered blocks can be used to compute d). This procedure has $O(n)$ communication, since the $n - \delta$ blocks at the server must be sent to the client. IRIS [28] is a system along these lines, requiring the whole file system be streamed over to the client for recovery.

Finally we note here that while PDP techniques combined with redundant blocks stored at the client can be used to solve the accountable storage problem (even inefficiently, as shown above), POR techniques cannot. This is because POR techniques (e.g., [26]) cannot provide proofs of retrievability for a certain portion of the file (as is the case with PDP), but only for the whole file—this is partly due to the fact that error-correcting codes are used on top of all the file blocks.

Our AS Protocol. Our protocol for assessing damage d from Relation 1 is based on recovering the actual blocks $b_1, b_2, \dots, b_\delta$ and XORing them with the corrupted blocks $b'_1, b'_2, \dots, b'_\delta$ returned by the server. For recovery, we use the invertible Bloom filter (IBF) data structure [12, 15]. An IBF is an array of t cells and can store $O(t)$ elements. Unlike a Bloom filter [7], its elements can be enumerated with high probability.

Let $B = \{b_1, b_2, \dots, b_n\}$ be the set of outsourced blocks and let δ be the maximum number of corrupted blocks that can be tolerated. In preprocessing, the client computes an IBF \mathbf{T}_B with $O(\delta)$ cells, on the blocks b_1, \dots, b_n . \mathbf{T}_B is stored locally. Computing

⁵ We highlight that such fine-grained compensation models, which work at the bit level as opposed to at the file block level, allow Alice to better manage her risk for damage even within the same file. For example, compensation for an unusable movie stored by Bob could be larger than that for a usable movie whose resolution has deteriorated by just 5%.

\mathbf{T}_B is similar to computing a Bloom filter: every cell of \mathbf{T}_B is mapped to a XOR over a set of at most n blocks, thus the local storage is $O(\delta)$. To outsource the blocks, the client computes homomorphic tags, \mathbb{T}_i (as in [3]), for each block b_i . The client then stores (b_i, \mathbb{T}_i) with the cloud and deletes b_1, b_2, \dots, b_n from local storage. In the challenge phase, the client asks the server to construct an IBF \mathbf{T}_K of $O(\delta)$ cells on the set of blocks K the server currently has—this is the “proof” the server sends to the client. Then the client takes the “difference” $\mathbf{T}_L = \text{subtract}(\mathbf{T}_B, \mathbf{T}_K)$ and recovers the elements of the difference $B - K$ (since $|B - K| \leq \delta$ and \mathbf{T}_L has $O(\delta)$ cells). Recovering blocks in $B - K$ enables the client to compute d using Relation 1. Clearly, the bandwidth of this protocol is proportional to δ (due to the size of the IBFs), and not to the total number of outsourced blocks n . Our optimized construction in Section 5 achieves sublinear server and client complexities as well.

Fairness through Integration with Bitcoin. The above protocol assures that Bob (the server) cannot succeed in persuading Alice that the damage of her file blocks is $d' < d$. After Alice is persuaded, compensation proportional to d must be sent to her. But Bob could try to cheat again. Specifically, Bob could try to give Alice a smaller compensation or even worse, disappear. To deal with this problem, we develop a modified version of the recently-introduced timed commitment in Bitcoin [2]. At the beginning of the AS protocol, Bob deposits a large amount, A , of bitcoins, where A is contractually agreed on and is typically higher than the maximum possible damage to Alice’s file blocks. The Bitcoin-integrated AS protocol of Section 6 ensures that unless Bob fully and timely compensates Alice for damage d , then A bitcoins are automatically and irrevocably transferred to Alice. At the same time, if Alice tries to cheat (e.g., by asking for compensation higher than the contracted amount), our protocol ensures that she gets no compensation at all while Bob gets back all A of his bitcoins.

Structure of the Paper. Section 2 presents background on IBFs and Bitcoin, Section 3 gives definitions, and Sections 4 and 5 present our constructions. We present our Bitcoin protocol in Section 6, our evaluation in Section 7 and conclude in Section 8.

2 Preliminaries

Let τ denote the security parameter, δ denote an upper bound on the number of corrupted blocks that can be tolerated, n denote the number of file blocks, and b_1, b_2, \dots, b_n denote the file blocks. Each block b_i has λ bits. The first $\log n$ bits of each block b_i are used for storing the index i of the block, which can be retrieved through function $\text{index}()$. Namely $i = \text{index}(b_i)$. Let also h_1, h_2, \dots, h_k be k hash functions chosen at random from a universal family of functions \mathcal{H} [9] such that $h_i : \{0, 1\}^\lambda \rightarrow \{1, 2, \dots, t\}$ for some parameter t .

Invertible Bloom Filters. An *Invertible Bloom Filter (IBF)* [12, 15] can be used to compactly store a set of blocks $\{b_1, b_2, \dots, b_n\}$: It uses a table (array) \mathbf{T} of $t = (k+1)\delta$ cells. Each cell of the IBF’s table \mathbf{T} contains the following two fields⁶: (1) dataSum: XOR of blocks b_i mapped to this cell; (2) hashSum: XOR of cryptographic tags (to

⁶ Note that we do not use the count field, as in [12, 15].

Algorithm $\mathbf{T} \leftarrow \text{update}(b_i, \mathbf{T})$ for each $j = 1, \dots, k$ do Set $\mathbf{T}[h_j(b_i)].\text{dataSum} \oplus = b_i$; Set $\mathbf{T}[h_j(b_i)].\text{hashSum} \oplus = \tau_i$; return \mathbf{T};	Algorithm $\mathbf{T}_D \leftarrow \text{subtract}(\mathbf{T}_A, \mathbf{T}_B)$ for each $i = 1, \dots, t$ do $\mathbf{T}_D[i].\text{dataSum} = \mathbf{T}_A[i].\text{dataSum} \oplus \mathbf{T}_B[i].\text{dataSum}$; $\mathbf{T}_D[i].\text{hashSum} = \mathbf{T}_A[i].\text{hashSum} \oplus \mathbf{T}_B[i].\text{hashSum}$; return \mathbf{T}_D;
--	--

Fig. 1: Update and subtraction algorithms in IBFs.

be defined later) τ_i for all blocks b_i mapped to this cell. As in Bloom filters, we use functions h_1, \dots, h_k to decide which blocks map to which cells.

An IBF supports simple algorithms for insertion and deletion via algorithm `update` in Figure 1. For $B \subseteq A$, one can also take the *difference* of IBFs \mathbf{T}_A and \mathbf{T}_B , to produce an IBF $\mathbf{T}_D \leftarrow \text{subtract}(\mathbf{T}_A, \mathbf{T}_B)$ representing the difference set $D = A - B$. Finally, given \mathbf{T}_D , we can enumerate its contents by using algorithm `listDiff` from [12]:

Lemma 1 (Adjusted from Eppstein et al. [12]). *Let $B \subseteq A$ be two sets having $\leq \delta$ blocks in their difference $A - B$, let \mathbf{T}_A and \mathbf{T}_B be the respective IBFs constructed using k hash functions and let $\mathbf{T}_D \leftarrow \text{subtract}(\mathbf{T}_A, \mathbf{T}_B)$. All IBFs have $t = (k + 1)\delta$ cells and their `hashSum` field is computed using a function mapping blocks to at least $k \log \delta$ bits. Then there is an algorithm `listDiff`(\mathbf{T}_D) that recovers $A - B$ with probability $1 - O(\delta^{-k})$.*

Bitcoin Time-Locked Transactions. Bitcoin [23] is a decentralized digital currency system where transactions are recorded on a public ledger (the blockchain) and are verified through the collective effort of *miners*. A bitcoin address is the hash of an ECDSA public key. Let A and B be two bitcoin addresses. A *standard* transaction contains a signature from A and mandates that a certain amount of bitcoins is to be transferred from A to B . If A 's signature is valid, the transaction is inserted into a block which is then stored in the blockchain.

Bitcoin allows for more complicated transactions, whose validation requires more than just a signature. In particular, each transaction can specify a *locktime* which contains a timestamp t at which the transaction is locked (before time t , even if a valid signature is provided, the transaction is not finalized). Slightly changing the notation from [2], a Bitcoin transaction \mathbf{T}_x can be easily represented as the table below, where `Prev` is the transaction (say \mathbf{T}_y) that \mathbf{T}_x is redeeming, `InputsToPrev` are inputs that \mathbf{T}_x is sending to \mathbf{T}_y so that \mathbf{T}_y 's redeeming can take place, `Conditions` is a program written in the Bitcoin scripting language (outputting a boolean) controlling whether \mathbf{T}_x can be redeemed or not (given inputs from another transaction), `Amount` is the value in bitcoins, and `Locktime` is the locktime. For standard transactions, `InputsToPrev` is just a signature with the sender's secret key, and `Conditions` implements a signature verification with the recipient's public key. Moreover, standard transactions have locktime set to 0, meaning they are locked and final. Andrychowicz *et al.* [2] uses time-locked Bitcoin transactions to describe a bitcoin contract mechanism for timed commitments [8] through the blockchain. With timed commitments in place, they show that it is possible to run fairly any secure multi-party protocols with the stipulation that parties pay a fine in bitcoins if they cheat, that

Prev :
InputsToPrev :
Conditions :
Amount :
Locktime :

is, if all parties follow the protocol specification then a function on private inputs is computed otherwise any party that deviates from (e.g., interrupts) the protocol will inescapably pay a sum of bitcoins to the others.

3 Accountable Storage Definitions

We now define an AS scheme. An AS scheme does not allow the client to compute damage d directly. Instead, it allows the client to use the server's proof to retrieve the blocks \mathcal{L} that are not stored by the server any more (or are stored corrupted). By having the server send the current blocks he stores in the position of blocks in \mathcal{L} (in addition to the proof), computing the damage d is straightforward.

Definition 1 (δ -AS scheme) A δ -AS scheme \mathcal{P} is the collection of four PPT algorithms:

1. $\{\text{pk}, \text{sk}, \text{state}, \mathbb{T}_1, \dots, \mathbb{T}_n\} \leftarrow \text{Setup}(b_1, \dots, b_n, \delta, 1^\tau)$ takes as inputs file blocks b_1, \dots, b_n , a parameter δ and the security parameter τ and returns a public key pk , a secret key sk , tags $\mathbb{T}_1, \dots, \mathbb{T}_n$ and a client state state .
2. $\text{chal} \leftarrow \text{GenChal}(1^\tau)$ generates a challenge for the server;
3. $\mathcal{V} \leftarrow \text{GenProof}(\text{pk}, \beta_{i_1}, \dots, \beta_{i_m}, \mathbb{T}_{i_1}, \dots, \mathbb{T}_{i_m}, \text{chal})$ takes as inputs a public key pk , a collection of $m \leq n$ blocks and their corresponding tags. It returns a proof of accountability \mathcal{V} ;
4. $\{\text{reject}, \mathcal{L}\} \leftarrow \text{CheckProof}(\text{pk}, \text{state}, \mathcal{V}, \text{chal})$ takes as inputs a public key pk and a proof of accountability \mathcal{V} . It returns a list of blocks \mathcal{L} or reject.

Relation to Proofs of Storage. A δ -AS scheme is a generalization of proof-of-storage (PoS) schemes, such as [3, 19]. In particular, a 0-AS scheme (i.e., where we set $\delta = 0$) is equivalent to PoS protocols, where there is no tolerance for corrupted/lost blocks.

Definition 2 (δ -AS scheme correctness) Let \mathcal{P} be a δ -AS scheme. Let $\{\text{pk}, \text{sk}, \text{state}, \mathbb{T}_1, \dots, \mathbb{T}_n\} \leftarrow \text{Setup}(b_1, \dots, b_n, \delta, 1^\tau)$ for some set of blocks $B = \{b_1, \dots, b_n\}$. Let now $\mathcal{L} \subseteq B$ such that $|\mathcal{L}| \leq \delta$, $\text{chal} \leftarrow \text{GenChal}(1^\tau)$ and $\mathcal{V} \leftarrow \text{GenProof}(\text{pk}, B - \mathcal{L}, \mathbb{T}(B - \mathcal{L}), \text{chal})$, where $\mathbb{T}(B - \mathcal{L})$ denotes the tags corresponding to the blocks in $B - \mathcal{L}$. A δ -AS scheme is correct if the probability that $\mathcal{L} \leftarrow \text{CheckProof}(\text{pk}, \text{state}, \mathcal{V}, \text{chal})$ is at least $1 - \text{neg}(\tau)$.⁷

To define the security of a δ -AS scheme, the adversary adaptively asks for tags on a set of blocks $B = \{b_1, b_2, \dots, b_n\}$ that he chooses. After the adversary gets access to the tags, his goal is to output a proof \mathcal{V} , so that if \mathcal{L} is output by algorithm CheckProof , where $|\mathcal{L}| \leq \delta$, then (a) either \mathcal{L} is *not* a subset of the original set of blocks B ; (b) or the adversary does *not* store all remaining blocks in $B - \mathcal{L}$ intact.

Such a proof is invalid since it would allow the verifier to either recover the *wrong* set of blocks (e.g., a set of blocks whose Hamming distance from the corrupted blocks is a lot smaller) or to accept a corruption of more than δ file blocks.

Definition 3 (δ -AS security) Let \mathcal{P} be a δ -AS scheme as in Definition 1 and \mathcal{A} be a PPT adversary. We define security using the following steps.

⁷ Function $\lambda : \mathbb{N} \rightarrow \mathbb{R}$ is $\text{neg}(\tau)$ iff \forall nonzero polynomials $p(\tau)$ there exists N so that $\forall \tau > N$ it is $\lambda(\tau) < 1/p(\tau)$.

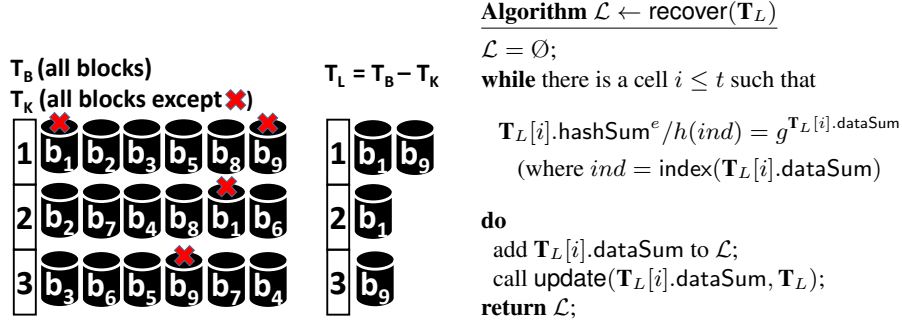


Fig. 2: **(Left)** On input b_1, b_2, \dots, b_9 , the client outputs an IBF \mathbf{T}_B of three cells using two hash functions. The server loses blocks b_1 and b_9 . \mathbf{T}_K is computed on blocks b_2, b_3, \dots, b_8 and \mathbf{T}_L contains the lost blocks b_1 and b_9 . **(Right)** The algorithm for recovering the lost blocks.

1. **Setup.** \mathcal{A} chooses $\delta \in [0, n)$, blocks $B = \{b_1, b_2, \dots, b_n\}$ and is given $\mathbb{T}_1, \dots, \mathbb{T}_n$ and pk output by $\{\text{pk}, \text{sk}, \text{state}, \mathbb{T}_1, \dots, \mathbb{T}_n\} \leftarrow \text{Setup}(b_1, \dots, b_n, \delta, 1^\tau)$.⁸
2. **Forge.** \mathcal{A} is given $\text{chal} \leftarrow \text{GenChal}(1^\tau)$ and outputs a proof of accountability \mathcal{V} . Suppose $\mathcal{L} \leftarrow \text{CheckProof}(\text{pk}, \text{state}, \mathcal{V}, \text{chal})$. We say that the δ -AS scheme \mathcal{P} is secure if, with probability at least $1 - \text{neg}(\tau)$: (i) $\mathcal{L} \subseteq B$; and (ii) there exists a PPT knowledge extractor \mathcal{E} that can extract all the remaining file blocks in $B - \mathcal{L}$.

Note here that if the set \mathcal{L} is empty, then the above definition is equivalent to the original PDP security definition [3]. Also note that the notion of a knowledge extractor is similar to the standard one, introduced in the context of proofs of knowledge [5]. If the adversary can output an accepting proof, then he can execute GenProof repeatedly until it extracts the selected blocks.

4 Our Basic Construction

We now give an overview of our basic construction: On input blocks $B = \{b_1, \dots, b_n\}$ in local storage, the client decides on a parameter δ (meaning that he can tolerate up to δ corrupted files) and computes the local state, tags, public and secret key by running $\{\text{pk}, \text{sk}, \text{state}, \mathbb{T}_1, \dots, \mathbb{T}_n\} \leftarrow \text{Setup}(b_1, \dots, b_n, \delta, 1^\tau)$. In our construction the tag \mathbb{T}_i is set to $(h(i)g^{b_i})^d \pmod N$, as in [3], where $h(\cdot)$ is a collision-resistant hash function, N is an RSA modulus and (e, d) denote an RSA public/private key pair. The client then sends blocks b_1, \dots, b_n and tags $\mathbb{T}_1, \dots, \mathbb{T}_n$ to the server and locally stores the state state, which is an IBF of the blocks b_1, b_2, \dots, b_n .

At challenge phase, the client runs $\text{chal} \leftarrow \text{GenChal}(1^\tau)$ that picks a random challenge s and sends it to the server. To generate a proof of accountability (see Figure 2-left) with GenProof, the server computes an IBF \mathbf{T}_K on the set of blocks that he (believes he) stores, along with a proof of data possession [3] on the same set of blocks.

⁸ \mathcal{A} could also choose blocks adaptively, after seeing tags for already requested blocks. Our proof of security handles that.

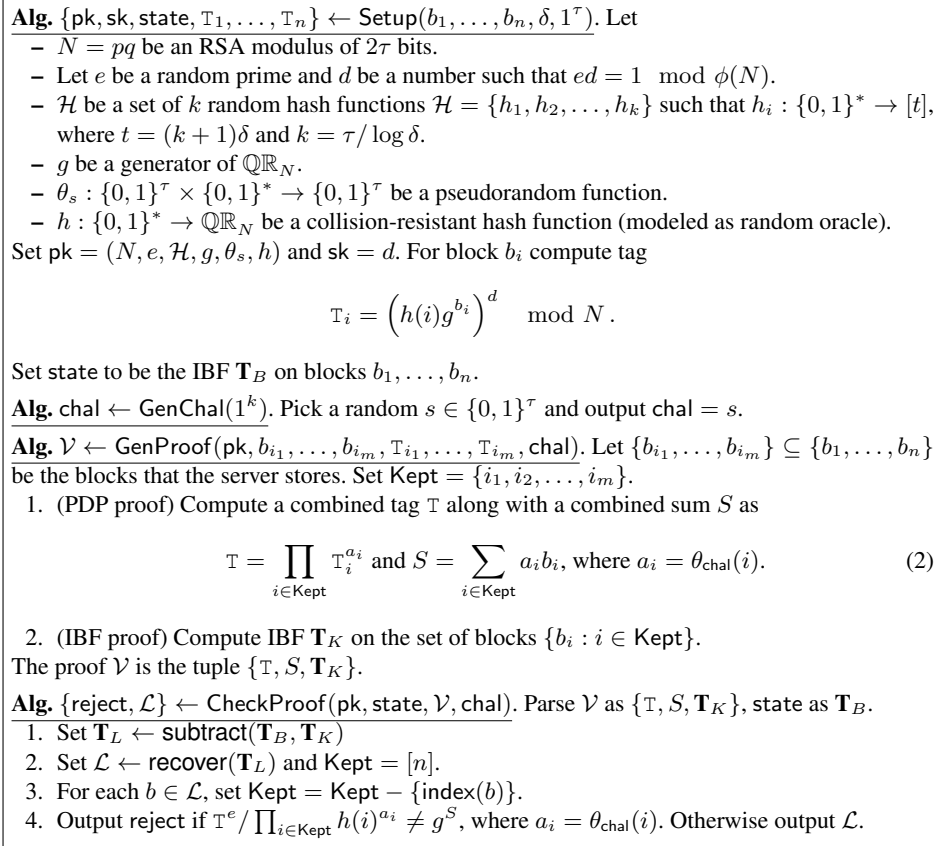


Fig. 3: Our δ -AS scheme construction.

The indices of these blocks are stored in a set Kept . For the computation of the PDP proof, the server uses randomness derived from the challenge s .

To verify the proof, the client takes the difference $\mathbf{T}_L = \text{subtract}(\mathbf{T}_B, \mathbf{T}_K)$ and executes algorithm `recover` from Figure 2-right, which is a modified version of `listDiff` from [12]. Algorithm `recover` adds blocks whose tags verify to the set of lost blocks \mathcal{L} . Then it checks the PDP proof for those block indices corresponding to blocks that were not output by `recover`. If this PDP proof does not reject, then the client is persuaded that the server stores everything except for blocks in \mathcal{L} . To make sure `recover` does not fail with a noticeable probability, our construction sets the parameters according to the following corollary. The detailed algorithms of our construction are in Figure 3.

Corollary 1 *Let τ be the security parameter and B and K be two sets such that $K \subseteq B$ and $|B - K| \leq \delta$. Let \mathbf{T}_B and \mathbf{T}_K be IBFs constructed by algorithm `update` of Figure 1 using $\tau / \log \delta$ hash functions. The IBFs \mathbf{T}_B and \mathbf{T}_K have $t = (\tau / \log \delta + 1)\delta$ cells and*

employ tags in the hashSum field that map blocks to τ bits. Then with probability at least $1 - 2^{-\tau}$, algorithm `recover(subtract($\mathbf{T}_B, \mathbf{T}_K$))` will output $\mathcal{L} = B - K$.

Our detailed proof of security is given in the Appendix. The local state that the client must keep is an IBF of $t = (k + 1)\delta$ cells, therefore the asymptotic size of the state is $O(\delta)$. For the size of the proof \mathcal{V} , the tag \mathbb{T} has size $O(1)$, the sum S has size $O(\log n + \lambda)$ and the IBF \mathbf{T}_K has size $O(\delta)$. Overall, the size of \mathcal{V} is $O(\delta + \log n)$. For the proof computation, note that algorithm `GenProof` must first access at least $n - \delta$ blocks in order to compute the PDP proof and then compute an IBF of δ cells over the same blocks, therefore the time is $O(n + \delta)$. Likewise, the verification algorithm needs to verify a PDP proof for a linear number of blocks and to process a proof of size $O(\delta + \log n)$, thus its computation time is again $O(n + \delta)$.

Theorem 1 (δ -AS scheme). *Let n be the number of blocks. For all $\delta \leq n$, there exists a δ -AS scheme such that: (1) It is correct according to Definition 2; (2) It is secure in the random oracle model based on the RSA assumption and according to Definition 3; (3) The proof has size $O(\delta + \log n)$ and its computation at the server takes $O(n + \delta)$ time; (4) Verification at the client takes $O(n + \delta)$ time and requires local state of size $O(\delta)$; (5) The space at the server is $O(n)$.*

We now make two observations related to our construction. First, note that the server could potentially launch a DoS attack, by *pretending it does not store some of the blocks* so that the client is forced to spend cycles retrieving these blocks. This is not an issue, since as we will see later, the server will be penalized for that, so it is not in its best interest. Second, note that the tags that the client initially uploads are publicly verifiable so anyone can check their validity—therefore the client cannot upload bogus tags and blame the server later for that.

Streaming and Appending Blocks. Our construction assumes the client has all blocks available in the beginning. This is not necessary. Blocks b_i could come one at a time, and the client could easily update its local state with algorithm `update($b_i, \mathbf{T}, 1$)`, compute the new tag \mathbb{T}_i and send the pair (b_i, \mathbb{T}_i) to the server for storage. This also means that our construction is partially-dynamic, supporting append-only updates. Modifying a block is not so straightforward due to replay attacks. However techniques from various fully-dynamic PDP schemes could be potentially used for this problem (e.g., [13]).

5 Sublinear Construction Using Proofs of Partial Storage

In the previous construction, the server and client run in $O(n + \delta)$ time. In this section we present optimizations that reduce the server and client performance to $O(\delta \log n)$. Recall that the proof generation in Figure 3 has two distinct, linear-time parts: First, proving that a subset of blocks is kept intact (in particular the blocks with indices in `Kept`), and second, computing an IBF on this set of blocks. We show here how to execute both these tasks in sublinear time using (i) partial proofs of storage; (ii) a data structure based on segment trees that the client must prepare during preprocessing.

Proofs of Partial Storage. In our original construction, we prove that a subset of blocks is kept intact (in particular the blocks with indices in `Kept`) using a PDP-style proof, as

originally introduced by Ateniese et al. [3]. In our new construction we will replace that part with a new primitive called *proofs of partial storage*. To motivate proofs of partial storage, let us recall how proofs of storage [26] work.

Proofs of storage provide the same guarantees with PDP-style proofs [3] but are much more practical in terms of proof construction time. In particular, one can construct a PoS proof in constant time as follows. Along with the original blocks b_1, b_2, \dots, b_n the client outsources an additional n redundant blocks $\beta_1, \beta_2, \dots, \beta_n$ computed with an error-correcting code such as Reed-Solomon, such that any n out of the $2n$ blocks $b_1, b_2, \dots, b_n, \beta_1, \beta_2, \dots, \beta_n$ can be used to retrieve the original blocks b_1, b_2, \dots, b_n . Also, the client outsources tags T_i (as computed in Algorithm Setup in Figure 3) for all $2n$ blocks. Now, during the challenge phase, the client picks a constant-sized subset of random blocks to challenge (out of the $2n$ blocks), say $\tau = 128$ blocks. Because the subset is chosen at random every time, the server, with probability at least $1 - 2^{-\tau}$, will pass the challenge (i.e., provide verifying tags for the challenged blocks) only if he stores at least half of the blocks $b_1, b_2, \dots, b_n, \beta_1, \beta_2, \dots, \beta_n$ —which means that the original blocks b_1, b_2, \dots, b_n are recoverable.

Unfortunately, we cannot use proofs of storage as described above directly, since we want to prove that a *subset of the blocks* is stored intact, and the above construction applies to the whole set of blocks. In the following we describe how to fix this problem using a segment-tree-like data structure.

Our New Construction: Using a Segment Tree. A segment tree T is a binary search tree that stores the set B of n key-value pairs (i, b_i) at the leaves of the tree (ordered by the key). Let v be an internal node of the tree T . Denote with $\text{cover}(v)$ the set of blocks that are included in the leaves of the subtree rooted on node v . Let also $|v| = |\text{cover}(v)|$. Every internal node v of T has a label $\text{label}(v)$ that stores:

1. All blocks $b_1, b_2, \dots, b_{|v|}$ contained in $\text{cover}(v)$ along with respective tags T_i . The tags are computed as in Algorithm Setup in Figure 3;
2. Another $|v|$ redundant blocks $\beta_1, \beta_2, \dots, \beta_{|v|}$ computed using Reed-Solomon codes such that *any* $|v|$ out of the $2|v|$ blocks $b_1, b_2, \dots, b_{|v|}, \beta_1, \beta_2, \dots, \beta_{|v|}$ are enough to retrieve the original blocks $b_1, b_2, \dots, b_{|v|}$. Along with every redundant block β_i , we also store its tag T_i .
3. An IBF T_v on the blocks contained in $\text{cover}(v)$;

By using the segment tree, one can compute functions on *any subset* of $n - \delta$ blocks in $O(\delta \log n)$ time (instead of taking $O(n - \delta)$ time): For example, if $i_1, i_2, \dots, i_\delta$ are the indices of the omitted δ blocks, the desired IBF T_K can be computed by combining (i.e., XORing the dataSum and hashSum fields and adding the count fields):

- The IBF T_1 corresponding to indices from 1 to $i_1 - 1$;
- The IBF T_2 corresponding to indices from $i_1 + 1$ to $i_2 - 1$;
- ...
- The IBF $T_{\delta+1}$ corresponding to indices from $i_\delta + 1$ to i_n .

Each one of the above IBFs can be computed in $O(\log n)$ time by combining a logarithmic number of IBFs stored at internal nodes of the segment tree and therefore the total complexity of computing the final IBF T_K is $O(\delta \log n)$. Similarly, a partial proof of storage for the lost blocks with indices $i_1, i_2, \dots, i_\delta$ can be computed by returning

- A proof of storage corresponding to indices from 1 to $i_1 - 1$;

- A proof of storage corresponding to indices from $i_1 + 1$ to $i_2 - 1$;
- ...
- A proof of storage corresponding to indices from $i_\delta + 1$ to i_n .

Again, each one of the above proofs of storage can be computed by returning $O(\log n)$ partial proofs of storage so in total, one needs to return $O(\delta \log n)$ proofs of storage. Note however that our segment tree increases our space to $O(n \log n)$ and also setting it up requires $O(n \log n)$ time. Therefore we have the following:

Theorem 2 (Sublinear δ -AS scheme). *Let n be the number of blocks. For all $\delta \leq n$, there exists a δ -AS scheme such that: (1) It is correct according to Definition 2; (2) It is secure in the random oracle model based on the RSA assumption and according to Definition 3; (3) The proof has size $O(\delta \log n)$ and its computation at the server takes $O(\delta \log n)$ time; (4) Verification at the client takes $O(\delta \log n)$ time and requires local state of size $O(\delta)$; (5) The space at the server is $O(n \log n)$.*

6 Bitcoin Integration

After the client computes the damage d using the AS protocol described in the previous section, we would like to enable automatic compensation by the server to the client in the amount of d bitcoins. The server initially makes a “security deposit” of A bitcoins by means of a special bitcoin transaction that automatically transfers A bitcoins to the client unless the server transfers d bitcoins to the client before a given deadline. Here, the amount A is a parameter that is contractually established by the client and server and is meant to be larger than the maximum damage that can be incurred by the server.⁹ We have designed a variation of the AS protocol integrated with bitcoin that, upon termination, achieves one of the following outcomes within an established deadline:

- If both the server and the client follow the protocol, the client gets exactly d bitcoins from the server and the server gets back his A bitcoins.
- If the server does not follow the protocol (e.g., he tries to give fewer than d bitcoins to the client, fails to respond in a timely manner, or tries to forge an AS proof), the client gets A bitcoins from the server automatically.
- If the client requests more than d bitcoins from the server by providing invalid evidence, the server receives all A deposited bitcoins back and the client receives nothing.

To guarantee the above outcomes, we implement the security deposit of A bitcoins by the server via a special Bitcoin transaction $\text{safeGuard}(x, t)$, and the related transactions retBtcs and fuse , and depicted with a diamond in Figure 4. Transaction $\text{safeGuard}(x, t)$ is based on the timed commitment over Bitcoin by Andrychowicz et al. [2], where x is the committed value and t is the bitcoin locktime. The functionality associated with $\text{safeGuard}(x, t)$ guarantees the following: (1) if x is known by the server then $\text{safeGuard}(x, t)$ can be opened and the server (and only the server) can get his A bitcoins back; (2) after time t all A bitcoins will go to the client. Namely, until either x is revealed or t has passed, the bitcoins of the cloud in the transaction $\text{safeGuard}(x, t)$

⁹ Of course, this is just a simple setting, a proof of concept. Clearly other technical and financial instruments can be used to improve this approach if committing such a large amount of A bitcoins is too demanding.

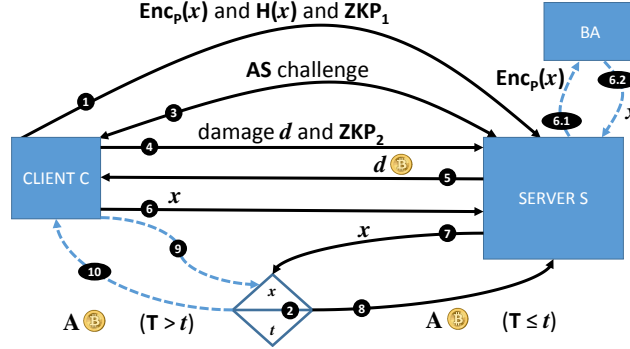


Fig. 4: Integration of the AS protocol with Bitcoin to enable an automatic compensation of the client in the case of data loss.

are effectively frozen. We note that the safeGuard transaction we are implementing has an important difference from the one introduced in [2]: the committed value x is chosen by the verifier (client) and not by the committer (server).

We now describe our protocol in detail, as depicted in Figure 4. Let S denote the server and C the client. Our protocol involves a trusted “Bitcoin Arbitrator” (BA). However, we note that the BA is only contacted by the server S and in case of disputes. For each step $i = 1, \dots, 10$, there is a deadline, t_i , to complete the step, where timelock $t \gg t_{10}$. We also assume neither the client nor the server can forge the timestamped transcript of the protocol, which can be verified by BA. This can be accomplished via standard techniques, e.g., through authenticated sessions that are posted on the public Bitcoin blockchain. (Note that, in this case, BA is effectively operating as a *smart contract* that receives and accesses messages stored on the blockchain.) In general, we assume a trusted and tamper-resilient channel between the client and the server that BA can efficiently validate.

- **Step 1:** C picks a random secret x and sends the following items to S : (i) an encryption $\text{Enc}_P(x)$ of x under BA’s public key, P ; (ii) a cryptographic hash of x , $H(x)$; and (iii) a zero-knowledge proof, ZKP_1 , that $H(x)$ and $\text{Enc}_P(x)$ encode the same secret x . If ZKP_1 does not verify, S aborts the protocol.
- **Step 2:** S posts bitcoin transaction $\text{safeGuard}(x, t)$ for A bitcoins with timelock t , as done in [2]. If this transaction is not posted within time t_2 (the server is not following the protocol), C aborts the protocol.
- **Step 3:** S and C run the AS protocol from the previous section, which results in C computing the damage, d . If the AS protocol rejects or S delays it past time t_3 , C jumps to Step 9.
- **Step 4:** C notifies S that the damage is d and sends a zero-knowledge proof, ZKP_2 , to S for that. If C fails to do so by time t_4 or ZKP_2 does not verify, S jumps to Step 6.
- **Step 5:** S sends d bitcoins to C . If S has not done so by time t_5 , C jumps to Step 9.
- **Step 6:** C sends secret x to S . If S has not received x by t_6 , S contacts BA and sends the timestamped transcript of the protocol up to that moment. BA checks the

transcript and if it is valid¹⁰, BA sends x to S . Note that the transcript must be tamper-proof and should contain the encryption of x and all messages exchanged up to that moment.

- **Step 7:** If S has secret x , S posts transaction `retBtcs` (i.e., S opens the timed commitment using x).
- **Step 8:** If transaction `retBtcs` is valid, S receives A bitcoins before timelock t .
- **Step 9:** C waits until time t and posts transaction `fuse`.
- **Step 10:** If transaction `fuse` is valid, C receives A bitcoins.

It is easy to see that when the above protocol terminates, one of the three outcomes described at the beginning of this section is achieved. Also, we note here that for the zero-knowledge proofs ZKP_1 and ZKP_2 , we can use a SNARK with zero-knowledge [25], that was recently implemented and shown to be practical.

The safeGuard Transaction. We now describe transaction `safeGuard`(x, t). The scheme is set up so that `safeGuard`(x, t) is created by S but where x is known only to C — S only knows $h = H(x)$. This is feasible to achieve through the first step of the protocol described in Figure 4. Since S knows h , he issues the following Bitcoin transaction:

Prev :	aTransaction
InputsToPrev :	$\text{sig}_S([\text{safeGuard}])$
Conditions :	$body, \sigma_1, \sigma_2, x :$ $H(x) = h \wedge \text{ver}_S(body, \sigma_1)$ \vee $\text{ver}_S(body, \sigma_1) \wedge \text{ver}_C(body, \sigma_2)$
Amount :	$A \mathbb{B}$
Locktime :	0

The above transaction redeems a transaction called `aTransaction` that has at least $A \mathbb{B}$ as value and can be redeemed by transactions whose `InputsToPrev` are of the type $body, \sigma_1, \sigma_2, x$ (i.e., two signatures on the same transaction and an integer x) and satisfy the conditions specified. Once the client reveals x to the server, the server posts the following transaction `retBtcs` to recover his A bitcoins.

Prev :	safeGuard
InputsToPrev :	$[\text{retBtcs}], \text{sig}_S([\text{retBtcs}]), \perp, x$
Conditions :	$body, \sigma :$ $\text{ver}_S(body, \sigma)$
Amount :	$A \mathbb{B}$
Locktime :	0

If the server does not cooperate (e.g., see Step 5 in the protocol), the client publishes the transaction `fuse` after time t and receives a compensation of A bitcoins from the server.

¹⁰ Note that the transcript contains the Bitcoin transaction that sends d bitcoins to the client, which the BA can locate on the public blockchain.

Prev :	safeGuard
InputsToPrev :	$[fuse], sig_S([fuse]), sig_C([fuse]), \perp$
Conditions :	$\frac{body, \sigma :}{ver_C(body, \sigma)}$
Amount :	$A \mathbb{B}$
Locktime :	t

It is very important to notice here that the locktime of the above transaction is t , meaning that if the transaction is posted earlier than t , it is not going to be accepted. This is what enables $A \mathbb{B}$ to be transferred to C if the server does not know x (which C controls!). Finally we note that the transaction Fuse is possible only because client C has already the server's signature $sig_S([fuse])$ on the body $[fuse]$, which includes the locktime t . We can assume that this signature is obtained in the beginning of the protocol.

Global safeGuard. The protocol above protects the client at each AS challenge. But the cloud provider could stop interacting, simply disappear, and never be reachable by the client. Our accountable framework thus establishes that there must be a *global* safeGuard transaction at the time the client and the server initiate their business relationship (i.e., when the client uploads the original file blocks and they both sign the SLA).

This global transaction is meant to protect the client if: (1) the server cannot be reached at all or (2) refuses to post the safeGuard transaction during any AS challenge, or (3) posts the safeGuard transaction but asks BA to recover the bitcoins in it without participating in the AS challenge or sending the d bitcoins to C . The global safeGuard transaction has the same format as the per-challenge safeGuard transaction and can be handled by the same BA. The only significant difference is that now the BA must mediate and interact with both the client and the server before returning any bitcoins to the server and, in case, reinitiate a new global safeGuard transaction between the client and the server.

A global safeGuard transaction prevents a malicious server from cheating but creates a scalability problem given that the server has to escrow a large amount of bitcoins for every client/customer. We do not address this problem technically but we expect it can be mitigated through financial mechanisms (securities, commodities, credit, etc.) typically deployed for traditional escrow accounts.

Removing the Bitcoin Arbitrator. Even though BA is only involved in case of disputes, it is preferable to remove it completely. Unfortunately, this seems impossible to achieve *efficiently* given the limitations of the Bitcoin scripting language. We sketch in this section two possible approaches to remove the BA. These will be further explored in future work.

The first approach relies on a secure two-party computation protocol. In a secure two-party computation protocol (2PC), party A inputs x and party B inputs y and they want to compute $f_A(x, y)$ and $f_B(x, y)$ respectively, without learning each other's input other than what can be inferred from the output of the two functions. Yao's seminal result [31] showed that oblivious transfer implies 2PC secure against honest-but-curious adversaries. This result can be extended to generically deal with malicious adversaries through zero-knowledge proofs or more efficiently via the *cut-and-choose* method [20] or LEGO and MiniLEGO [14, 24] (other efficient solutions were proposed in [18, 30]).

To remove the BA, it is enough to create a symmetric version of our original scheme where both parties create a safeGuard transaction and then exchange the secrets of both commitments through a fair exchange protocol embedded into a 2PC. The secrets must be verifiable in the sense that the fair exchange must ensure the secrets open the initial commitments or fail (as in “committed 2PC” by Jarecki and Shmatikov [18]). Unfortunately, generic techniques for 2PC results in quite impractical schemes and this is the reason why we prefer a practical solution with an arbiter. An efficient 2PC protocol with Bitcoin is proposed in [22] but it does not provide fairness since the 2PC protocol can be interrupted at any time by one of the parties. In the end, since this generic approach is too expensive in practice, we will not elaborate on it any further in this paper.

Another promising approach to remove the BA is to adopt smart contracts. Smart contracts are digital contracts that run on a blockchain. Ethereum [1] is a new cryptocurrency system that provides a Turing-complete language to write such contracts, which is expected to enable several decentralized applications without trusted entities. Smart contracts will enable our protocol to be fully automated without any arbitrators or trusted parties in between. To use a smart contract to run our protocol, the contract is expected to receive a deposit from the server, inputs from both parties, and then it will decide the money flow accordingly based on the CheckProof result. The contract in that case will maintain some properties to ensure fair execution, for example both parties should be incentivized to follow the protocol, and if any party does not follow the protocol (by aborting for example), there should be a mechanism to end the protocol properly for the honest party. To make our protocol fit in the smart contract model, we will need to address the fact that the CheckProof computation would be too expensive to be performed by the contract, due to its overhead. In Ethereum for example, the participants must pay for the cost of running the contract, which is run by the miners.

In order to address the points above, we can employ zero knowledge SNARKs [6] to help reduce miners’ overhead, and thus reduce the computational cost of the verification algorithm running on the network, while preserving the secrecy of the secret inputs. Therefore, instead of running CheckProof by the contract, the client Alice (who has the secret input) will run the CheckProof on the server input and her state, and then show the result along with a succinct proof to the contract. The verification of such proof takes few milliseconds to be executed, which would reduce the computation of the contract substantially in that case. This will require just one-time trusted setup to generate the cryptographic keys needed for the protocol, otherwise Alice could generate fake proofs. A contract realizing our functionality can be implemented on Ethereum except for the zk-SNARK verification. Although zk-SNARKs can conceptually be implemented there (as the language is Turing-complete), it would not be efficient without native support, due to the cryptography involved. Due to many applications zk-SNARKs have in cryptocurrencies, we expect to see its support in the near future.

7 Evaluation

We prototyped the proposed Accountable Storage (AS) scheme in Python 2.7.5. Our implementation is open-source¹¹ and consists of 4K lines of source code. We use the

¹¹ <https://github.com/evripidispar/delta-AccountableStorage>

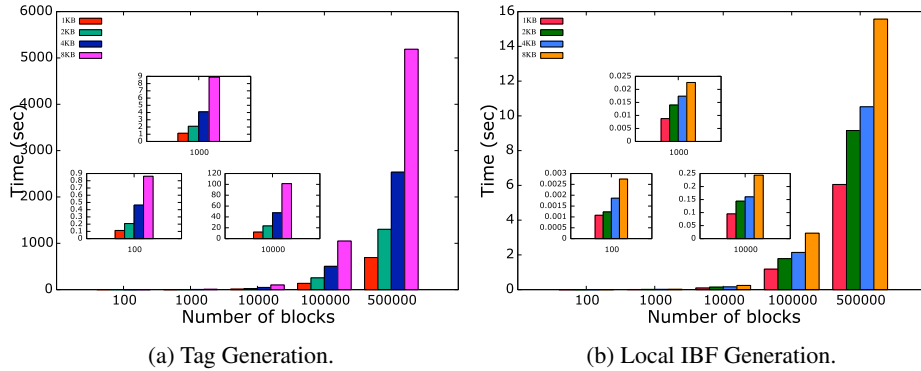


Fig. 5: Preprocessing overheads.




Fig. 6: Proof Generation and Recover time.

pycrypto library 2.6.1 [21] and an RSA modulus N of size 1024 bits. We serialize the protocol messages using Google Protocol Buffers [16] and perform all the modulo exponentiation operations using GMPY2 [17], which is a C-coded Python extension module that supports fast multiple precision arithmetic (the use of GMPY2 gave us 60% speedup in exponentiations in comparison with the regular python arithmetic library).

We divide the prototype in two major components. The first is responsible for data pre-processing, issuing proof challenges and verifying proofs. The second produces proof every time it receives a challenge. Both modules utilize the IBF data structure to produce and verify proofs. Our prototype uses parallel computing via the Python

multiprocessing module to carry out many of the heavy, but independent, cryptographic operations simultaneously. We used a *single-producer, many-consumers* approach to divide the available tasks in a pool of [8-12] processes-workers. The workers use message passing to coordinate and update the results of their computations. This approach significantly enhanced the performance of preprocessing as well as the proof generation and checking phase of the protocol. Our parallel implementation provides an approximate 5x speedup over a sequential implementation.

Finally note that since it can be easily estimated, we have not evaluated the Bitcoin part of our protocol which is dominated by the time it takes for transactions to be part of the blockchain. Nowadays this latency is approximately 10 minutes.

Experimental Setup: Our experimental setup involves two nodes, one implementing the server and another implementing the client functionality. The two nodes communicate through a Local Area Network (LAN). The two machines are equipped with an Intel 2.3 Ghz Core i7 processor and have 16 GB of RAM.

Our data are randomly generated filesystems. Every file-system includes different number of equally-sized blocks. The number of blocks ranges from 100 to 500000 and the different sizes of blocks used are 1KB, 2KB, 4KB and 8KB. The total filesystem size varies from 100 KByte to 4.1 GByte. Our experiments consist of 10 trials of challenge/proof exchanges between the client and the server for different filesystems. Throughout the evaluation we report the average values over these 10 trials.

In our experiments, we select the tolerance parameter δ , which indicates the maximum amount of data blocks that can be lost, to be equal to $\log_2(n)$. One other possible choice of δ is to set it equal to \sqrt{n} . We select the logarithm of the number of blocks as δ , because this provides a harder condition on how many blocks can be lost or corrupted from the cloud server.

For the IBF construction, we have used message authentication code scheme instead of XORing the tags of the data blocks for efficiency purposes. The message authentication code (MAC) used is an HMAC. In addition, the selected number of hash functions used for the IBF construction is $k = 6$. This choice of hash functions leads to a very low probability of failure of the recovery algorithm, which depends on the values of k and δ .

Preprocessing Overheads: We first examine the memory overhead of the preprocessing phase, which is shown in Table 1. The first column describes the available number of blocks in a filesystem and the second represents the total size of the tags needed. The preprocessing memory overhead is proportional to the number of blocks in a filesystem.

Figure 5 shows the CPU-time-related overheads of the preprocessing of the protocol. These overheads are divided to tag generation and the creation of the client state represented by the IBF T_B . The tag generation time (Figure 5a) increases linearly with both the available number of blocks and the size of each block. While this cost is significant for large file systems, it is an operation that client performs only once at the setup phase. On the other hand, the cost of construction of the IBF (Figure 5b) is negligible; the IBF construction of our biggest filesystem only takes around 16 seconds.

Challenge-Proof Overheads: We now examine memory and CPU-related overheads for the challenge-proof exchange and the recovery phase. The last four columns of

Table 1: Memory Footprint of the AS Scheme (KB)

n	Tag Size (KB)	Proof Size (KB)			
		1KB	2KB	4KB	8KB
10^2	32	341	680	1356	2710
10^3	236	510	1017	2029	4055
10^4	2644	735	1466	2927	5848
10^5	24895	904	1898	3526	7051
$5 * 10^5$	118326	1017	2028	4049	8034

Table 1 show the proof sizes (in KB) for $\delta = \log_2(n)$, which increase proportionally to the block size.

Every subgraph of Figure 6 shows how different block sizes affect the performance of the challenge-proof exchange for a given number of blocks. The left bar in the figure shows the proof generation time and the right bar the proof check along with the time to recover the lost blocks. Larger block sizes increase the time-overhead of challenge-proof exchange.

8 Conclusions

In this paper we put forth the notion of accountability in cloud storage. Unlike existing work such as proof-of-storage schemes and verifiable computation, we design protocols that respond to a verification failure, enabling the client to assess the damage that has occurred in a storage repository. We also present a protocol that enables automatic compensation of the client, based on the amount of damage, and is implemented over Bitcoin. Our implementation shows that our system can be used in practice.

Acknowledgments

Research supported in part by an NSF CAREER award CNS-1652259, NSF grants CNS-1525044, CNS-1526950, CNS-1228639 and CNS-1526631, a NIST award and by the Defense Advanced Research Projects Agency (DARPA) under agreement no. AFRL FA8750-15-2-0092. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

Bibliography

- [1] Ethereum: A platform for decentralized applications. www.ethereum.org/.
- [2] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. Secure multiparty computations on bitcoin. In *IEEE SSP*, 2014.
- [3] Giuseppe Ateniese, Randal Burns, Reza Curtmola, Joseph Herring, Lea Kissner, Zachary Peterson, and Dawn Song. Provable data possession at untrusted stores. In *ACM CCS*, 2007.
- [4] Giuseppe Ateniese, Roberto Di Pietro, Luigi V. Mancini, and Gene Tsudik. Scalable and efficient provable data possession. In *SecureComm*, 2008.
- [5] Mihir Bellare and Oded Goldreich. On defining proofs of knowledge. In *CRYPTO*, 1992.
- [6] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. Snarks for C: Verifying program executions succinctly and in zero knowledge. In *CRYPTO*, 2013.
- [7] Burton H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Comm. ACM*, 13:422–426, 1970.
- [8] Dan Boneh and Moni Naor. Timed commitments. In *CRYPTO*, 2000.
- [9] I. L. Carter and M. N. Wegman. Universal classes of hash functions. In *ACM STOC*, 1977.
- [10] David Cash, Alptekin Küpçü, and Daniel Wichs. Dynamic proofs of retrievability via oblivious ram. In *EUROCRYPT*, 2013.
- [11] Reza Curtmola, Osama Khan, Randal C. Burns, and Giuseppe Ateniese. MR-PDP: Multiple-replica provable data possession. In *ICDCS*, 2008.
- [12] David Eppstein, Michael T. Goodrich, Frank Uyeda, and George Varghese. What’s the difference? Efficient set reconciliation without prior context. In *SIGCOMM*, 2011.
- [13] Chris Erway, Alptekin Küpçü, Charalampos Papamanthou, and Roberto Tamassia. Dynamic provable data possession. In *ACM CCS*, 2009.
- [14] Tore Kasper Frederiksen, Thomas Pelle Jakobsen, Jesper Buus Nielsen, Peter Sebastian Nordholt, and Claudio Orlandi. Minilego: Efficient secure two-party computation from general assumptions. In *EUROCRYPT*, 2013.
- [15] Michael T. Goodrich and Michael Mitzenmacher. Invertible Bloom Lookup Tables. *ArXiv e-prints*, January 2011.
- [16] Google. Google protocol buffers. www.developers.google.com/protocol-buffers/.
- [17] Case Van Hosen. Gmpy2: Multiple-precision arithmetic for python. www.gmpy2.readthedocs.org/en/latest/intro.html/.
- [18] Stanisław Jarecki and Vitaly Shmatikov. Efficient two-party secure computation on committed inputs. In *EUROCRYPT*, 2007.
- [19] Ari Juels and Jr. Burton S. Kaliski. PORs: proofs of retrievability for large files. In *ACM CCS*, 2007.
- [20] Yehuda Lindell and Benny Pinkas. An efficient protocol for secure two-party computation in the presence of malicious adversaries. In *EUROCRYPT*, 2007.

- [21] Dwayne C. Litzenger. Pycrypto - the python cryptography toolkit. www.dlitz.net/software/pycrypto/.
- [22] Daniel Malinowski Marcin Andrychowicz, Stefan Dziembowski and Lukasz Mazurek. Fair two-party computations via the bitcoin deposits. In *FC*, 2014.
- [23] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. www.bitcoin.org/bitcoin.pdf.
- [24] Jesper Buus Nielsen and Claudio Orlandi. LEGO for two-party secure computation. In *TCC*, 2009.
- [25] Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In *IEEE SSP*, 2013.
- [26] Hovav Shacham and Brent Waters. Compact proofs of retrievability. In *ASIACRYPT*, 2008.
- [27] Elaine Shi, Emil Stefanov, and Charalampos Papamanthou. Practical dynamic proofs of retrievability. In *ACM CCS*, 2013.
- [28] Emil Stefanov, Marten van Dijk, Alina Oprea, and Ari Juels. Iris: A scalable cloud file system with efficient integrity checks. In *ACSAC*, 2012.
- [29] Qian Wang, Cong Wang, Jin Li, Kui Ren, and Wenjing Lou. Enabling public verifiability and data dynamics for storage security in cloud computing. In *ESORICS*, 2009.
- [30] David P Woodruff. Revisiting the efficiency of malicious two-party computation. In *EUROCRYPT*, 2007.
- [31] Andrew Chi-Chih Yao. How to generate and exchange secrets. In *SFCS*, 1986.

Appendix

RSA assumption

Definition 4 Let $N = pq$ be an RSA modulus, where p and q are τ -bit primes. Given N , e and g , where g is randomly chosen from \mathbb{Z}_N^* and e is a prime of $\Theta(\tau)$ bits, there is no PPT algorithm that can output $y^{1/e} \bmod N$, except with probability $\text{neg}(\tau)$.

Proof of security of construction in Figure 3.

Setup. \mathcal{A} chooses parameter $\delta \in [0, n)$, blocks $B = \{b_1, b_2, \dots, b_n\}$ and is given pk and T_1, \dots, T_n as output by $\{\text{pk}, \text{sk}, \text{state}, T_1, \dots, T_n\} \leftarrow \text{Setup}(b_1, \dots, b_n, \delta, 1^\tau)$. The random oracle is programmed so that it returns $r_i^e g^{-b_i}$ on input i for some random r_i , i.e., $h(i) = r_i^e g^{-b_i} \bmod N$.

Forge. \mathcal{A} is given $\text{chal} \leftarrow \text{GenChal}(1^\tau)$, computes proof of accountability \mathcal{V} and returns \mathcal{V} . Suppose $\mathcal{L} \leftarrow \text{CheckProof}(\text{pk}, \text{sk}, \text{state}, \mathcal{V}, \text{chal})$. We must show that with probability $\geq 1 - \text{neg}(\tau)$ it is (i) $\mathcal{L} \subseteq B$; and (ii) there exists a PPT knowledge extractor \mathcal{E} that can extract all the remaining file blocks in $B - \mathcal{L}$.

1. **Showing $\mathcal{L} \subseteq B$.** Note that all blocks in \mathcal{L} are output by Algorithm recover of Figure 2. In this algorithm a block b'_i can enter \mathcal{L} only if its tag verifies. Suppose now $b'_i \notin B$ (namely $b'_i \neq b_i$) and $\text{tag}^e / h(i) = g^{b'_i}$ for some arbitrary tag computed by the adversary. But since $h(i) = r_i^e g^{-b_i}$, this can be written as

$\text{tag}^e / r_i^e g^{-b_i} = g^{b'_i}$ which gives $g^{b'_i - b_i} = (\text{tag}/r_i)^e$. Since e is a prime and $b_i - b'_i \neq 0$, there exist α and β such that $(b'_i - b_i) \times \alpha + e \times \beta = 1$, giving $g^{1/e} = g^{-\beta} (\text{tag}/r_i)^{e \times \alpha}$, breaking the RSA assumption—see Definition 4.

2. **Showing there exists a PPT knowledge extractor \mathcal{E} that can extract all the remaining file blocks in $B - \mathcal{L}$.** We now show how to build an extractor that, after $\ell = |B - \mathcal{L}|$ interactions with the adversary, he can extract the blocks $\{b_i : i \in B - \mathcal{L}\}$. The extractor will challenge the adversary exactly ℓ times, each time with different randomness. Let $S_1, S_2, \dots, S_\ell, \mathbb{T}^{(1)}, \mathbb{T}^{(2)}, \dots, \mathbb{T}^{(\ell)}$ be the sums and tags he receives by \mathcal{A} during each challenge, as in Equation 2. We have two cases:

- (a) $S_j = \sum_{i \in B - \mathcal{L}} a_{ij} b_i$, for $j \in B - \mathcal{L}$ (a_{ij} denotes the randomness of the j -th challenge corresponding to the i -th block). In this case, the extractor can solve a system of ℓ linear equations and retrieve the original blocks $\{b_i : i \in B - \mathcal{L}\}$.
- (b) Suppose there exists $j \in B - \mathcal{L}$ such that $S_j \neq \sum_{i \in B - \mathcal{L}} a_{ij} b_i = S$. For simplicity of notation, let's set $\mathbb{T}^{(j)} = \mathbb{T}$ and $S_j = \bar{S}$. Then by the CheckProof algorithm we have

$$\frac{\mathbb{T}^e}{\prod_{i \in B - \mathcal{L}} h(i)^{a_i}} = g^{\bar{S}}.$$

But since $h(i) = r_i^e g^{-b_i}$ we have that

$$g^{\bar{S} - S} = \left(\frac{\mathbb{T}}{\prod_{i \in B - \mathcal{L}} r_i^{a_i}} \right)^e = \mathcal{Z}^e.$$

Therefore we have $\mathcal{Z}^e = g^{\bar{S} - S}$. Again, since e is prime and $S \neq \bar{S}$ we can use the same trick as before, and break the RSA assumption. \square