Leakage Resilient Proofs of Ownership in Cloud Storage, Revisited*

Jia Xu and Jianying Zhou

Infocomm Security Department, Institute for Infocomm Research, Singapore
{xuj,jyzhou}@i2r.a-star.edu.sg

Abstract. Client-side deduplication is a very effective mechanism to reduce both storage and communication cost in cloud storage service. Halevi *et al.* (CCS '11) discovered security vulnerability in existing implementation of client-side deduplication and proposed a cryptographic primitive called "proofs of ownership" (PoW) as a countermeasure. In a proof of ownership scheme, any owner of the same file can prove to the cloud storage server that he/she owns that file in an efficient and secure manner, even if a bounded amount of any efficiently extractable information of that file has been leaked. We revisit Halevi *et al.*'s formulation of PoW and significantly improve the understanding and construction of PoW. Our contribution is twofold: Firstly, we propose a generic and conceptually simple approach to construct *Privacy-Preserving* Proofs of Ownership scheme, by leveraging on well-known primitives (i.e. Randomness Extractor and Proofs of Retrievability) and technique (i.e. sample-then-extract). Our approach can be roughly described as Privacy-Preserving PoW = Randomness Extractor + Proofs of Retrievability. Secondly, in order to provide a better instantiation of Privacy-Preserving-PoW, we propose a novel design of randomness extractor with large output size, which improves the state of art by reducing both the random seed length and entropy loss (i.e. the difference between the entropy of input and output) simultaneously.

Keywords: Cloud Storage, Client-side Deduplication, Proofs of Ownership, Leakage Resilience, Privacy-Preserving, Proofs of Retrievability, Randomness Extractor, Sample-then-Extract.

1 Introduction

Cloud storage service (e.g. Dropbox, Skydrive, Google Drive, iCloud, Amazon S3) is becoming more and more popular in recent years [2]. The volume of personal or business data stored in cloud storage keeps increasing [3,4,5]. In face to the challenge of rapidly growing volume of data in cloud, deduplication technique is highly demanded to save disk space by removing duplicated copies of the same file (Single Instance Storage). SNIA white paper [6] reported that the deduplication technique can save up to 90% storage, dependent on applications.

* This work is supported by Singapore A*STAR project SecDC-112172014. The full version of this work is available at Cryptology ePrint Archive, Report 2013/514 [1].

I. Boureanu, P. Owesarski, and S. Vaudenay (Eds.): ACNS 2014, LNCS 8479, pp. 97–115, 2014.

Traditional deduplication technique (i.e. server side deduplication [7,8,9,10]) in centralized storage system removes duplicated copies residing in the same server. Unlike server-side deduplication, client-side deduplication in cloud storage system will identify duplicated copies such that one copy resides in the cloud storage server and the other resides remotely in the cloud client, and saves the uploading bandwidth (time, respectively) for the duplicated file. In both server and client side deduplication, all owners of the deduplicated file will be provided a soft link to the unique copy of that file stored in the centralized storage or cloud storage respectively. In contrast to server-side deduplication which saves only storage on server side, client-side deduplication saves not only server storage but also network bandwidth and transmission time, and benefits both cloud server and client.

However, how to implement client-side deduplication *securely* in an untrusted environment, is far more challenging than it first appears [11,12]. Arguably, the root cause of the difference between security requirements of server-side and client-side deduplication, is that server-side deduplication is executed in the trusted server, while client-side deduplication is distributively executed between the trusted[1] cloud server and potentially untrusted cloud client. Here the cloud user is considered as potentially untrusted, since anyone from the untrusted Internet could become a cloud user and the cloud server is unable to distinguish honest users from malicious users (i.e adversaries) in general.

Server side deduplication may simply apply a collision resistant hash function (say SHA256) to identity duplicated files in the storage server, and remove the extra copies to achieve "single instance storage". An existing implementation of client-side deduplication (called as "hash-as-a-proof" method) is as below: Cloud storage server keeps a lookup table, which records hash value of each file in its storage. Cloud user Alice, who tries to upload file F to the cloud storage, will firstly send hash value $\mathsf{hash}(F)$ to the cloud server. If $\mathsf{hash}(F)$ is not found in the lookup table, then Alice should upload file F to the cloud storage and cloud server will update the lookup table by adding entry $\mathsf{hash}(F)$. Otherwise, cloud server has a copy of F already, which could be uploaded by other users. Consequently Alice's uploading process will be saved, and Alice is allowed to download F from cloud server on demand. In the above method, the knowledge of hash value $\mathsf{hash}(F)$ is treated as a "proof" that Alice owns file F. Previously, Dropbox[2] applied the above "hash-as-a-proof" method on block-level cross-users deduplication [12][13].

Halevi *et al.* [12] targets the critical security vulnerability in the above "hash-as-a-proof" method, where the leakage of a short hash value $\mathsf{hash}(F)$ would lead (or amplify) to leakage of entire file F to outside adversary. Their work proposes a cryptographic primitive called "proofs of ownership" (PoW) to address

[1] The cloud server is trusted in data integrity and availability in this work.

[2] In Feb 2012, we noticed that Dropbox disabled the deduplication across different users, probably due to recent vulnerabilities discovered in their original cross-user client-side deduplication method. This also indicates the importance and urgency in the study of security in client-side deduplication.

such leakage amplification vulnerability. The distinguishable feature of Halevi *et al.* [12] from all of previous study in security of deduplication (e.g. convergent encryption [7,8,14]), is that Halevi *et al.* [12] adopts a *bounded leakage model* to characterize the untrusted environment in which the client-side deduplication runs. Their formulation requires that, after a setup between one owner of file F and the cloud storage server, any owner of F can efficiently *prove* (in the sense of "interactive proof system" [15]) to the cloud storage server that he/she indeed owns file F without really transmitting F, even if a bounded amount of any efficiently extractable information of F has been leaked via some owner (considered as the accomplice or colluder) of F intentionally or unintentionally.

In this work, we revisit Halevi *et al.* [12]'s formulation, and extend it in two aspects: (1) We shift a significant amount of workload (precisely, the setup procedure) from cloud server to a cloud user, which reflects our understanding of real world setting—the average computation power allocated to each online user by cloud server is typically smaller than the computation power of an average cloud user. (2) We protect data privacy against verifier (e.g. the cloud storage server), during the interactive proof protocol. Halevi *et al.* [12]'s formulation does not address privacy protection of user data against the cloud storage server. Prudent users may have reasons to not trust the cloud server. For example, the cloud server may be hacked (e.g.[16]), making it a single point of failure of user data privacy. In addition, the cloud server may make careless technical mistakes [17,18], which may expose user data to unauthorized persons. In this work, we will trust cloud storage server in data availability and integrity (which is the research topic of proofs of storage [19,20]), but not trust it in data privacy.

1.1 Overview of Our Result

Under the framework of Halevi *et al.* [12], in a secure PoW scheme, if the input file F has k bits min-entropy to the view of adversary at the very beginning and at most T ($< k - \lambda$) bits of message about F is leaked at adversary's (adaptive) choice, then the adversary should not be able to convince the cloud storage server that he/she owns file F with significant probability.

1.1.1 Generic Construction of Privacy-Preserving-PoW. Intuitively, our generic construction of Privacy-Preserving-PoW is as below: At first, apply a *proper*[3] randomness extractor over file F to output $T + 2\lambda$ ($< k$) bits almost-uniform random number Y_F. Next, apply a *proper* proofs of retrievability (POR [19]) scheme over Y_F. Since the output Y_F of the randomness extractor is statistically close to true uniform randomness, any adversary that learns at most T bits arbitrary information of F, cannot output the $T + 2\lambda$ bits long value Y_F entirely with significant probability, and thus cannot succeed in the verification of POR scheme. The difference $(k - T)$ is like the *entropy loss* in randomness extractor, thus the smaller the difference $(k - T)$ is , the better the PoW scheme is in aspect of leakage resilience.

[3] See Theorem 1 and Theorem 2 for the explanation of "proper" randomness extractor and "proper" POR.

Our result can be combined with convergent encryption or Message-Locked Encryption [7,8,21,10,22], in order to construct strong leakage-resilient client-side deduplication scheme for encrypted data in cloud storage and thus protect data privacy against both outside adversary and curious cloud server.

We remark that formulating and constructing privacy-preserving PoW scheme are very challenging. Previous work by Ng *et al.* [23] made the first attempt towards this goal, but gave an unsatisfactory solution: As pointed out by Xu *et al.* [21], Ng *et al.* [23] formulates the privacy property *locally* for each block and their scheme suffers from "divide and conquer" attack: If an input file with N blocks has 1 bit min-entropy in each block *independently*, then this file could be recovered by an outside adversary via brute force search in time $\mathcal{O}(N)$ instead of $\mathcal{O}(2^N)$.

1.1.2 Improved Randomness Extractor. Unfortunately, the state of art [24,25] (with restriction of small seed size and practical computation cost) of randomness extractor only gives us a PoW with $k-T = \Omega(|F|)$ and requires relatively large random seed. We propose a new randomness extractor with shorter random seed and results in a PoW with $k - T = \mathcal{O}(|F|^{1-c})$ for any constant $c \in (0, 1)$.

Table 1. Compare our PoW scheme with existing works. Unsatisfactory items are highlighted in italic font and red color.

Scheme	Distribution of input	Seed Size	Computation complexity	Privacy-Preserving	Security Model
PoW1 [12]	Any	$\mathcal{O}(\lambda)$	***Expensive [12]***	***No (Leaking whole file F)***	Stand. Model
PoW2 [12]	Any	***$\geq$ 6T*** †	***Prohibitively expensive [12]***	***No***	Stand. Model
PoW3 [12]	***Generalized block-fixing distribution***	$\mathcal{O}(\lambda)$	Practical	***Unclear***	***Rand. Oracle; Unjustified Assump.‡***
This work	Any	$\mathcal{O}(\lambda)$	Practical	Yes	Stand. Model

†T may take value 64MB.
‡ Theorem 3 in [12] relies on an unproven assumption that the code generated by the third construction PoW3 is "good" and authors of [12] admits that it is very hard to analyze this unproven assumption. See text surrounding Theorem 3 in [12].

Table 2. Compare randomness extractors with output size $\ell\rho$, where ℓ could take value as large as $2^{21} \approx 2$ millions. The input is file F. Unsatisfactory items are highlighted in italic font and red color.

Scheme	Distribution of input	Randomness complexity	Computation complexity	Entropy Loss	Security Model
$\mathsf{HMAC}(s_1,F)\|\|\cdots\|\|\mathsf{HMAC}(s_\ell,F)$	Any	$\ell\lambda$	$\ell\|F\|$	small	Random Oracle
Inner Product Universal Hash [26]	Any	$2\|F\|$	$\Omega(\|F\|\log(\ell\rho))$	$2\log(1/\epsilon)$	Stand. Model
[24]	Any	$\mathcal{O}(\ell\lambda)$	$2\|F\|\log\ell$	$\Omega(\|F\|)$	Stand. Model
This work	Any	$\mathcal{O}(\lambda)$	$2\|F\|\log\ell$	$\mathcal{O}(\|F\|^{1-c})$ †	Stand. Model

†$c \in (0, 1)$

1.2 Contributions

Our main contributions can be summarized as below:

1. We propose a generic and conceptually simple paradigm to construct proof of ownership scheme: PoW=Randomness Extractor + Proofs of retrievability. To the best of our knowledge, this is the first work that bridges the proof of ownership and randomness extractor. Our result improves previous works on PoW in the following aspects: (1) Complete proof of security in standard model for *any* distribution of input file, while still being practical. (2) The first generic framework to construct PoW and benefited from the future advance in randomness extractor or proofs of retrievability. (3) Privacy-Preserving against verifier (e.g. cloud storage server). A detailed comparison between our work and existing PoW schemes is given in Table 1 (on page 100).
2. We propose a novel construction of randomness extractor with large output size, which improves existing work [24] by reducing both the seed length and entropy loss (i.e. the difference between entropy of input and output) *simultaneously*. This new randomness extractor may have independent interest. A detailed comparison between our work and existing randomness extractors is given in Table 2 (on page 100).

1.3 Organizations

We introduce preliminaries and background in Section 2 and formulation in Section 3. We present our overall solution in a modular approach in Section 4 and Section 5: At first in Section 4, we propose the construction of Privacy-Preserving-PoW and analyze its security, by treating an important component (i.e randomness extractor) as black-box. Next, Section 5 constructs the required randomness extractor with rigorous analysis and completes the description of the proposed solution. Section 6 concludes this paper. Due to space constraint, experiment result and most detailed proofs will be available only in full paper [1].

2 Preliminaries and Background

2.1 Notations and Definitions

Key notations in this paper are defined in Table 3 (on page 102).

Definition 1 (Statistical Difference). *The* statistical difference *between two random variables* $\mathbf{X}$ *and* $\mathbf{Y}$ *on the same space* $\mathcal{U}$ *is defined as*

$$\mathsf{SD}(\mathbf{X}, \mathbf{Y}) = \frac{1}{2} \sum_{a \in \mathcal{U}} \left| \mathsf{Pr}[\mathbf{X} = a] - \mathsf{Pr}[\mathbf{Y} = a] \right| \tag{1}$$

Some useful background information about statistical difference is provided in full paper [1].

2.2 Proofs of Retrievability

We adopt the formulation of proofs of retrievability from existing works [27,28] and make some syntactical modifications according to our needs to construct proofs of ownership scheme.

Table 3. Key Notations

Notation	Semantics
λ	The security parameter.
PPT	Probabilistic polynomial time (w.r.t. security parameter λ, if not explicitly stated otherwise).
$[n]$	The set of integers $1, 2, 3, 4, \ldots, n$.
$\mathsf{h}(\cdot)$	Full domain collision resistant hash function (e.g. SHA256).
$F[i]$	The projection of bit-string F onto i-th coordinate (i.e. the i-th bit of F, $1 \leq i \leq \|F\|$).
$F[\{i_1, \ldots, i_n\}]$	The projection of bit-string F onto the subset of coordinates (i.e $F[i_1]\|F[i_2]\| \ldots \|F[i_n]$, where $1 \leq i_1 < i_2 < \ldots < i_n \leq \|F\|$).
$\mathbf{H}_\infty(X)$	min-entropy of random variable X.
$\mathsf{SD}(X, Y)$	Statistical difference between random variables X and Y.
$X \approx_\epsilon Y$	$\mathsf{SD}(X, Y) \leq \epsilon$; X is ϵ-close to Y.
$B\|_{A=a}$	The conditional distribution of B given that $A = a$ for jointly distributed random variables (A, B).
$x \sim \mathcal{D}$	Sample x according to distribution $\mathcal{D}$.
$U_{\|n\|}$	Independent uniform random variable over $\{0,1\}^n$.
$U_{\|n\|,1}, U_{\|n\|,2}, \ldots$	Independently and identically distributed uniform random variables over $\{0,1\}^n$.

Definition 2 (Proofs of Retrievability). *A proofs of retrievability (POR) scheme consists of PPT algorithms* $\mathsf{KeyGen}, \mathsf{Tag}, \mathsf{GenChal}, \mathsf{GenProof}$ *and* Verify*, which are described as below*

- $\mathsf{KeyGen}(1^\lambda) \rightarrow (pk, sk)$*. The key generation algorithm takes a security parameter* λ *as input and outputs a pair of public-private key* (pk, sk)*.*
- $\mathsf{Tag}(sk, \{F_i\}_{i=1}^n) \rightarrow \{\sigma_i\}_{i=1}^n$*. The tag generation algorithm computes an authentication tag* σ_i *for each file block* F_i*.*
- $\mathsf{GenChal}(pk, n, c) \rightarrow (C, \Psi_F, \Psi_\sigma)$*. The challenger generation algorithm takes as input the public key pk, erasure encoded file size n (in term of blocks), and the sample size c, and outputs a sample* $C \subset [n]$ *with* $|C| = c$ *and meta-data* (Ψ_F, Ψ_σ)*.*
- $\mathsf{GenProof}(pk, \{(F_i, \sigma_i)\}_{i=1}^n, C, \Psi_F, \Psi_\sigma) \rightarrow (\bar{F}, \bar{\sigma})$*, where* $\bar{F} := \mathsf{GenProof}_{\mathsf{data}}(pk, \{F_i\}_{i=1}^n, C, \Psi_F)$ *and* $\bar{\sigma} := \mathsf{GenProof}_{\mathsf{tag}}(pk, \{\sigma_i\}_{i=1}^n, C, \Psi_\sigma)$*. The algorithm* $\mathsf{GenProof}_{\mathsf{data}}$ *takes as input the public key pk, file blocks* F_i*'s, a sample set* $C \subset [n]$*, and meta-data* Ψ_F*, and outputs an aggregated file block denoted as* $\bar{F}$*. The algorithm* $\mathsf{GenProof}_{\mathsf{tag}}$ *takes as input the public key pk, authentication tags* σ_i*'s, a sample set* $C \subset [n]$*, and meta-data* Ψ_σ*, and outputs an aggregated authentication tag denoted as* $\bar{\sigma}$*.*
- $\mathsf{Verify}(K, \bar{F}, \bar{\sigma}, \Psi_F, \Psi_\sigma, C) \rightarrow$ `Accept` *or* `Reject`*. If K is private key sk, then the POR scheme supports private key verifiability; if K is public key pk, then the POR scheme supports public key verifiability.*

We remark that the above formulation is syntactically different from original [27,28] in the sense that we explicitly decompose the algorithm $\mathsf{GenProof}$ into two sub-routines: $\mathsf{GenProof}_{\mathsf{data}}$ and $\mathsf{GenProof}_{\mathsf{tag}}$, where $\mathsf{GenProof}_{\mathsf{data}}$ processes selected data blocks F_i $(i \in C)$ and $\mathsf{GenProof}_{\mathsf{tag}}$ processes corresponding authentication tags σ_i's. Many existing works (e.g. [27,28] and Merkle Hash Tree based POR) support such decomposition, but a few works (e.g. [19]) do not.

For some POR schemes [27,28], meta-data Ψ_F and Ψ_σ are two seeds from which a list of coefficients $\{\alpha_i\}_{i \in C}$, $\{\beta_i\}_{i \in C}$ can be generated, and the aggregated values are $\bar{F} = \sum_{i \in C} \alpha_i F_i$ and $\bar{\sigma} = \sum_{i \in C} \beta_i \sigma_i$.

Definition 3 (Soundness of POR [19,27,28]). *Let* $\epsilon \in (0, 1)$*. A POR scheme is* ϵ*-sound, if there exists a PPT extractor algorithm, such that for any prover*

which can convince the verifier to accept with probability $\geq \epsilon$*, then the extractor can output the original file with overwhelming high probability (1 - negl) by executing POR proof protocol with the prover.*

Readers may find more details about POR in [19,27,30,28].

2.3 Randomness Extractor

Definition 4 (Strong Extractor). *We say* $\mathsf{Ext} : \{0,1\}^{\ell_{\text{in}}} \times \{0,1\}^{\ell_s} \to \{0,1\}^{\ell_{\text{out}}}$ *is a strong* (k, ϵ)*-extractor, if for any distribution* X *over* $\{0,1\}^{\ell_{\text{in}}}$ *with at least* k *bits min-entropy, the following inequality holds*

$$\mathsf{SD}\Big((\mathsf{Ext}(X;s), s),\ (U_{\ell_{\text{out}}}, s)\Big) \leq \epsilon \tag{2}$$

where the seed s *is uniformly randomly chosen from* $\{0,1\}^{\ell_s}$ *and* $U_{\ell_{\text{out}}}$ *is a uniform random variable over* $\{0,1\}^{\ell_{\text{out}}}$*.*

It is well known that the output size $\ell_{\texttt{out}}$ of any randomness extractor can not exceed the min-entropy k of the input (i.e. $\ell_{\texttt{out}} < k$), and the difference $(k - \ell_{\texttt{out}})$ is called the "entropy loss" of the randomness extractor.

3 Formulation: Proofs of Ownership, Revisited

Halevi *et al.* [12] proposed the formulation of proofs of ownership. In this section, we revisit and revise their formulation and propose our definition for privacy-preserving proofs of ownership.

Definition 5 (Proofs of Ownership [12]). *A proof of ownership scheme (*PoW*) consists of a probabilistic algorithm* S *and a pair of probabilistic interactive algorithm* $\langle \mathsf{P}, \mathsf{V} \rangle$*, which are described as below:*

- $\mathsf{S}(F, 1^\lambda) \to \psi$*: The randomized summary function* S *takes a file* F *and the security parameter* λ *as input, and outputs a short summary value* ψ*, where the bit-length of* ψ *is short and independent on file size* $|F|$*.*
- $\langle \mathsf{P}(F), \mathsf{V}(\psi) \rangle \to$ `Accept` or `Reject`*: The prover algorithm* P *which takes as input a file* F*, interacts with the verifier algorithm* V *which takes as input a short summary value* ψ*, and outputs either* `Accept` *or* `Reject`*.*

We are only interested in efficient PoW *scheme, such that* V *is polynomial time algorithm w.r.t. security parameter* λ *and both* S *and* P *are polynomial algorithms in* $|F|$ *and* λ*.*

Definition 6 (Completeness of PoW [12]). *A PoW scheme* $(\mathsf{S}, \langle \mathsf{P}, \mathsf{V} \rangle)$ *is* complete*, if for all positive integer* λ *and for any file* $F \in \{0,1\}^{poly(\lambda)}$*, it holds that*

$$\langle \mathsf{P}(F), \mathsf{V}(\mathsf{S}(F, 1^\lambda)) \rangle \text{ always outputs } \texttt{Accept}.$$

3.1 Two Players Setting and Three Players Setting of PoW

In the original framework [12], PoW runs by two players: verifier and prover. In this paper, we will redefine this system model by introducing a third player, called summarizer, who is responsible to preprocess the data file F during the setup. The PoW scheme in three players setting executes in this way: Summarizer

(e.g. data owner of F) runs summary function to obtain $\psi := \mathsf{S}(F, 1^\lambda)$ and sends ψ to verifier (e.g. the cloud storage server). Then prover (e.g. some cloud user claiming to own file F), who runs algorithm $\mathsf{P}(F)$, interacts with the verifier, who runs algorithm $\mathsf{V}(\psi)$. A dishonest prover (e.g. dishonest cloud user) may replace the prover algorithm P with any other PPT program of his/her choice.

Definition 7 (Two/Three Players setting of PoW). *For any PoW scheme* $(\mathsf{S}, \langle \mathsf{P}, \mathsf{V} \rangle)$*, the two players setting and three players setting are described as below:*

- *in a* **two players setting***, the summary algorithm* S *and verifier algorithm* V *are executed by the first player*—verifier *(e.g. cloud storage server), and the prover algorithm* P *is executed by the second player*—prover *(e.g. cloud user);*
- *in a* **three players setting***, the summary algorithm* S *is executed by the first player*—summarizer *(e.g. cloud user owning file F), the verifier algorithm* V *is executed by the second player*—verifier *(e.g. cloud storage server), and the prover algorithm* P *is executed by the third player*—prover *(e.g. another cloud user claiming to own F).*

Our three players setting will further relieve the computation burden of the cloud storage server, and might make our scheme easier to be adopted by cloud storage servers in real applications—This is exactly our initial motivation to introduce the new three players setting of PoW. We believe that, the average computation resource that a cloud storage server allocates to each online user, is typically less than the computation resource of an average cloud user. Additionally, the fact that many cloud storage servers (e.g. Dropbox, Skydrive, and Google Drive) provide free service to public users, further justifies our attempt to shift some computation burden from cloud server to cloud user.

The change from two players setting to three players setting also leads to the change of trust model and thus impact the security formulation. In the original two players setting of PoW [12], preserving privacy of input file F during the interactive proof $\langle \mathsf{P}, \mathsf{V} \rangle$ (like in zero-knowledge proof) is meaningless, since the verifier, who runs V, also runs the summary function $\mathsf{S}(F, 1^\lambda)$ and has direct access to file F. Therefore, the verifier has to be trusted in data confidentiality of input file F in this two players setting. In contrast, in our three players setting, preserving privacy of F during the interactive proof $\langle \mathsf{P}, \mathsf{V} \rangle$ (like in zero-knowledge proof) is an interesting problem, if the verifier (e.g. cloud storage server) is not trusted in data confidentiality.

3.2 Soundness of PoW

Intuitively, PoW aims to prevent leakage amplification in client-side deduplication: If an outside adversary *somehow* obtain a bounded amount ($\leq T$ bits) of messages about the target user file F via out-of-band leakage, then the adversary cannot obtain the whole file F by participating in the client-side deduplication with the cloud storage server.

The security game $\mathsf{G}_{\mathcal{A}}^{\mathsf{PoW}}(k, T)$ between a PPT adversary $\mathcal{A}$ and a challenger w.r.t. PoW scheme $(\mathsf{S}, \langle \mathsf{P}, \mathsf{V} \rangle)$ is defined as below. Here k is the lower bound

of min-entropy of the distribution of the challenged file F at the beginning of the game, and the adversary is allowed to learn at most T bits message related to file F (possibly including random coins chosen when processing F) from the challenger via the leakage query.

Setup. The description of $(\mathsf{S}, \langle \mathsf{P}, \mathsf{V} \rangle)$ is made public. Let $\mathcal{D}$ be a distribution over $\{0,1\}^M$ with min-entropy $\geq k$, where $\mathcal{D}$ is chosen by the adversary $\mathcal{A}$ and M is any public positive integer constant. The challenger samples file F according to distribution $\mathcal{D}$ and runs the summary algorithm to obtain $\psi := \mathsf{S}(F, 1^\lambda)$.

Learning. The adversary $\mathcal{A}$ can adaptively make polynomially many queries to the challenger, where each query is in one of the following types and concurrent queries of different types are not allowed[4]. Furthermore, the total amount of messages output by all leakage queries should not be greater than the threshold T, i.e. $\mathcal{Y}_{\mathrm{I}} + \mathcal{Y}_{\mathrm{II}} \leq T$, where $\mathcal{Y}_{\mathrm{I}}$ and $\mathcal{Y}_{\mathrm{II}}$ will be defined below.

- Prove-Query: The challenger, running the verifier algorithm V with input ψ, interacts with the adversary $\mathcal{A}$ which replaces the prover algorithm P, to obtain $b := \langle \mathcal{A}, \mathsf{V}(\psi) \rangle$. The adversary $\mathcal{A}$ is given the value of b.
- Leak-Query-I($\mathcal{P}$): This query consists of a description of a PPT algorithm $\mathcal{P}$ (a variant version of prover algorithm). The challenger responses this query by computing the output y of $\mathcal{P}(F)$ after interacting with $\mathsf{V}(\psi)$ (i.e. $y := \mathcal{P}(F)^{\mathsf{V}(\psi)}$) and sending y to the adversary $\mathcal{A}$. Denote with $\mathcal{Y}_{\mathrm{I}}$ the sum of bit-lengths of all responses y's for this type of queries.
- Leak-Query-II($\mathcal{L}$): This query consists of a description of a PPT algorithm $\mathcal{L}$. Let $\mathsf{transcript}_{\mathsf{S}}$ denote the transcript of all steps of operations in the execution of algorithm "$\psi := \mathsf{S}(F, 1^\lambda)$" in the above **Setup** phase. The challenger responses this query by computing the output $y := \mathcal{L}(\mathsf{transcript}_{\mathsf{S}})$ and sending y to the adversary $\mathcal{A}$. Denote with $\mathcal{Y}_{\mathrm{II}}$ the sum of bit-lengths of all responses y's for this type of queries.

Challenge. The adversary $\mathcal{A}$ which replaces the prover algorithm P, interacts with the challenger, which runs the verifier algorithm V with input ψ, to obtain $b := \langle \mathcal{A}, \mathsf{V}(\psi) \rangle$. The adversary $\mathcal{A}$ wins the game, if $b = \mathtt{Accept}$.

Definition 8 (Soundness of PoW (Refining [12])). *A* PoW *scheme is* (k, T, ϵ)*-sound in* three players setting*, if for any PPT adversary* $\mathcal{A}$*,* $\mathcal{A}$ *wins the security game* $\mathsf{G}_{\mathcal{A}}^{\mathsf{PoW}}(k, T)$ *with probability not greater than* $\epsilon + negl(\lambda)$*.*

$$\mathsf{Pr}[\mathcal{A} \text{ wins the security game } \mathsf{G}_{\mathcal{A}}^{\mathsf{PoW}}(k, T)] \leq \epsilon + negl(\lambda). \tag{3}$$

The (k, T, ϵ)*-soundness definition in* two players setting *is the same as the above, except that the adversary* $\mathcal{A}$ *is not allowed to make* Leak-Query-II *in the security game* $\mathsf{G}_{\mathcal{A}}^{\mathsf{PoW}}(k, T)$ *(i.e.* $\mathcal{Y}_{\mathrm{II}} = 0$*).*

[4] Concurrent Prove-Query and Leak-Query would allow the adversary to replay messages back and forth between these two queries, and eliminate the possibility of any secure and efficient solution to PoW. Therefore, the framework of Halevi *et al.* [12] do not allow concurrent queries of different types in the security formulation. We clarify that, concurrent queries of the same type can be supported. Thus, in the real application, the cloud storage server (verifier) can safely interact with multiple cloud users (prover) w.r.t. the same file concurrently.

We remark that (1) the (k, T, ϵ)-soundness definition in two players setting is essentially the same as the original formulation [12], and (2) soundness in three players setting implies soundness in two players setting, but not vice versa.

3.3 Privacy-Preserving PoW

Intuitively, we say a PoW scheme is privacy-preserving against the verifier, if everything about file F that the verifier can learn after participating the PoW scheme w.r.t. F, can be computed from the short summary value of F and some almost-perfect uniform random number.

Definition 9 (Privacy-Preserving). *A PoW scheme* $(\mathsf{S}, \langle \mathsf{P}, \mathsf{V} \rangle)$ *is* (k, T, ϵ)*-privacy-preserving against the verifier (in the three players setting), if for any distribution* $\mathcal{D}$ *over* $\{0,1\}^M$ *with at least* k *bits min-entropy, for every PPT interactive algorithm* V^**, there exists a PPT algorithm* Sim *and a random variable* Z *over domain* $\{0,1\}^{T+\lambda+\Omega(\lambda)}$*, such that*

- $\mathsf{SD}(Z, U_{|Z|}) \leq \epsilon$*, where* $U_{|Z|}$ *is the uniform random variable over* $\{0,1\}^{|Z|}$*;*
- *for any function* $f : \{0,1\}^M \to \{0,1\}$*, and any (leakage) function* $\mathcal{L} : \{0,1\}^M \to \{0,1\}^{\leq T}$*, the following two probabilities (taken over file* $F \sim \mathcal{D}$ *and the random coins of related algorithms) are equal*

$$\mathsf{Pr}\Big[\mathsf{V}^*\big(\psi \| \mathcal{L}(F)\big)^{\mathsf{P}(F)} = f(F)\Big] = \mathsf{Pr}\Big[\mathsf{Sim}\big(\psi \| \mathcal{L}(F), Z\big) = f(F)\Big],$$

where $\psi := \mathsf{S}(F, 1^\lambda)$ *and* $\mathsf{V}^*(\mathsf{S}(F, 1^\lambda) \| \mathcal{L}(F))^{\mathsf{P}(F)}$ *denotes the output of (dishonest) verifier* V^* *taking the summary value* $\mathsf{S}(F, 1^\lambda)$ *and leakage information* $\mathcal{L}(F)$ *as input and having interaction with interactive prover algorithm* $\mathsf{P}(F)$*.*

As we discussed before, preserving privacy against the verifier for any PoW scheme in the two players setting, is impossible.

3.4 Clarification on Leakage of User ID and Password

We admit that, as the same as Halevi *et al.* [12], this work will consider leakage of user account (i.e. id and password) as out of scope. We assume the user account is associated to user's real identity (e.g. mobile phone number) and sibyl account is hard to create. Thus, leakage of user file stored in cloud storage by disclosure of user account could be traced back to the source and the corresponding account could be disabled without affecting honest users.

4 Generic Construction of Proofs of Ownership

4.1 Some Unsatisfactory Approaches

At first, putting privacy-preserving property aside, we review some straightforward approaches and existing works for PoW as below.

4.1.1 Compute fresh MACs online on Both Sides. To prove his/her ownership of a file F, the prover can compute a MAC (i.e. Message Authentication Code) value over F with a random nonce as key, where the random nonce is chosen by the verifier. To verify the correctness of this MAC value, the verifier need to re-compute the MAC value of F under the same key. This approach is secure, but rejected for two reasons: (1) in some applications of PoW, the verifier

does not have access to the file F; (2) the stringent requirement on efficiency (including disk IO efficiency) given by Halevi *et al.* [12] does not allow verifier to access entire file F during the interactive proof.

4.1.2 Pre-compute MACs offline. In the summary phase, t number of keys $s_1, \ldots, s_t$ are randomly chosen and t number of MAC values $\mathsf{MAC}_{s_i}(F)$'s are computed correspondingly. The summary value of file F is $\{(i, s_i, \mathsf{MAC}_{s_i}(F)) : i \in [t]\}$. In the i-th proof session, the verifier sends the MAC key s_i to the prover and expects $\mathsf{MAC}_{s_i}(F)$ as response.

This approach is not secure in the setting of PoW [12], since a single malicious adversary could consume up all of t pre-computed MACs easily by impersonating or colluding with t distinct cloud users.

4.1.3 Proofs of Retrievability. Some instance of POR (e.g. [27,32,30]) can serve as PoW. The first construction (i.e. PoW1 as in Table 1) of Halevi *et al.* [12] is just the Merkle Hash Tree based POR scheme (MHT-POR), which combines error erasure code and Merkle Hash Tree proof method[5]. The drawback of this approach is that, the relatively expensive error erasure code[6] is applied over the whole input file, while in our approach, error erasure code is applied over the output of the randomness extractor, which is much shorter than the whole input file.

We notice that recent work by Zheng and Xu [33] attempts to equip proofs of storage (POR or PDP) with deduplication capability. However, their work is not in the leakage setting of Halevi *et al.* [12].

4.1.4 Pairwise-Independent Hash with Large Output Size. The second construction of PoW in Halevi *et al.* [12] is based on pairwise independent hash family (a.k.a 2-independent or 2-universal hash family). A large input file is hashed into a constant size (say about $3T = 3 \times 64\texttt{MB}$) hash value and then apply the merkle hash tree proof method over the hash value. This construction is secure, but very in-efficient in both computation and randomness complexity. Furthermore, large random seed also implies large communication cost required to share this seed among all owners of the same file. It is worth pointing out that Halevi *et al.* [12] overlooked the disadvantage in large randomness complexity (i.e. at least twice of hash output size, say about $2 \times 3T = 6 \times \texttt{64MB}$), although they admitted that this construction is *prohibitively* expensive in computation for practical data size.

A quick thought to reduce the seed length is to apply pseudorandomness generated from a short true random seed. However, in the leakage setting of PoW, any short seed could be leaked to the adversary by some colluded owner of target file. Consequently, the standard computational indistinguishability argument of pseudorandom number generator (or pseudorandom functions) is not applicable.

[5] Merkle Hash Tree proof method proves the correctness of a leaf value by presenting as a proof all sibling values along the path from the questioned leaf to the root of Merkle Hash Tree, and verification requires only the root value.

[6] In typical usage of error erasure code, block length is some small constant (say 223 bytes for (255, 223)-reed-solomon code). However, in the usage of POR, the block length has to be as large as the input file, which makes the coding much slower than typical case.

It is unclear whether this pseudorandomness approach works or not without new sophisticated proof (or disproof). Similar issue is discussed in the study of proofs of retrievability by Dodis *et al.* [30], which adopts sampling technique with public coin as seed to replace pseudorandomness.

4.1.5 PoW with respect to Particular Distribution. The third construction of PoW in Halevi *et al.* [12] is the most efficient one among all of three constructions proposed by Halevi *et al.* [12]. In the third construction, the size of random seed is dramatically reduced by treating hash function `SHA256` as a random oracle. However, their proof (in random oracle model) of this construction is incomplete: firstly, the distribution of input file is restricted as "generalized bit/block-fixing distribution"[7]; secondly, their proof assumes their algorithm will generate a "good linear code" and the authors admit that it is "very hard to analyze" this unproven assumption (See texts around Theorem 3 in [12]).

We emphasize that, information leakage of file F may have different forms. For example, some plain bits $F[i]$'s are leaked, or some aggregated information of file F (e.g. a hash value) is leaked. In the latter case, file F is hardly considered as fitting in (generalized) fixed-bit/block distribution.

Gabizon *et al.* [35] proposed a randomness extractor for input under bit-fixing distribution. Such extractor can be combined with our generic construction to obtain a secure PoW scheme for bit-fixing input file and with complete security proof in standard model.

Other works on deduplication/PoW include Pietro and Sorniotti [36], which treats a projection $(F[i_1], \ldots, F[i_\lambda])$ of file F onto λ randomly chosen bit-positions $(i_1, \ldots, i_\lambda)$ as the "proof" of ownership of file F. Similar to the "hash-as-a-proof" method, this work is extremely efficient but insecure in the bounded leakage setting [12]. Readers may find more related works in Xu *et al.* [21].

4.2 Our Approach: PoW = Randomness Extractor + POR

Intuitively, our generic construction extracts $(T + 2\lambda)$ bits message Y from the input file F and then apply a proofs of retrievability scheme over Y. It is worth noting that in our usage of proofs of retrievability scheme, algorithm $\mathsf{POR.GenProof}_{\mathsf{data}}$ runs by prover and algorithm $\mathsf{POR.GenProof}_{\mathsf{tag}}$ runs by verifier[8], while in the literature [19,27,28], both of these two algorithms run by prover. It is easy to see that, such modification will preserve the soundness of POR scheme.

The detailed construction is given in Figure 1 (on page 110). Before presenting a formal statement in Theorem 2 for the PoW scheme in Figure 1 which

[7] A M bits long file F with k bit entropy under "generalized bit-fixing distribution" is generated in this way: (1) Independently choosing k uniform random bits; (2) deriving all other $(M - k)$ bits from these k random bits (Halevi *et al.* [12] applies linear transformation); (3) the file F is a random permutation of these k random bits and $(M - k)$ derived bits. If in the above step (2), all $(M - k)$ bits are constant, then the resulting distribution is called "bit-fixing distribution" with entropy k.

[8] All tag values are stored with the verifier instead of the provers, in order to prevent any potential leakage of partial information of Y from its tag values to the (dishonest) provers.

constructed from a generic randomness extractor algorithm and a generic POR scheme, we will prove a stronger result in Theorem 1 for the special case that the POR scheme is instantiated with MHT-POR[9] scheme in the construction of PoW. The reason that MHT-POR can achieve a stronger result is that, the security of MHT-POR relies on the cryptographic one-way function without trapdoor (precisely the collision resistance hash function). In contrast, most other POR schemes rely on cryptographic trapdoor one-way function (e.g. factorization), and such *short* trapdoor (or private key) might be leaked via some colluded file owner in our stringent security model in three player setting. Once the short trapdoor is leaked to the adversary, the POR scheme can be easily broken.

Theorem 1. *Suppose* Extractor : $\{0,1\}^M \times \{0,1\}^{\ell_s} \rightarrow \{0,1\}^{T+2\lambda}$ *is a strong* (k,ϵ)*-extractor, and the POR scheme is the Merkle Hash Tree based scheme* MHT-POR *(as described in Sec 2.2.1 in the full paper [1]), which is* ϵ*-sound. Then the PoW scheme constructed in Figure 1 is* (k,T,ϵ)*-sound and* (k,T,ϵ)*-privacy-preserving in the three players setting. (Proof is in full paper [1])*

Most POR schemes [27,28] require a short private key (e.g. the factorization of a RSA modulus, the secret key of some pseudorandom function) to work and thus cannot resist Type-II leak query Leak-Query-II, from which the adversary could learn the short private key and break the POR scheme. Therefore, for such POR schemes with private key, we have to disable Type-II leak query by switching to the two players setting as below.

Theorem 2. *Suppose* Extractor : $\{0,1\}^M \times \{0,1\}^{\ell_s} \rightarrow \{0,1\}^{T+2\lambda}$ *is a strong* (k,ϵ)*-extractor and* POR *is an* ϵ*-sound POR scheme. Then the PoW scheme constructed in Figure 1 is* (k,T,ϵ)*-sound in the two players setting.*

We compare two instantiations of our generic approaches in Table 4 (on page 109).

Table 4. Two instantiations of PoW=RE+POR

Choice of POR	Setting	Summary Value Size (bits)	Communication cost (bits)
MHT-POR	2P,3P	λ	$\lambda \cdot \log_{1-\alpha} \epsilon \cdot \log(T/\alpha)$
Brent-Waters-POR [27]	2P	$T/(\alpha s)$ †	$(s+3)\lambda + 440$

† : s is a system parameter of POR [27] and can take any positive integer value.

5 Randomness Extractor with Large Output Size

In this section, we propose in Figure 2 (on page 111) a novel randomness extractor with large output size using the well-known "sample-then-extract" approach: Repeatedly sample a subset of bits from a weak random source and then apply an existing extractor with small output size over the sample.

Intuitively, the sampling lemma [24,25] states that "if one samples a random subset of bits from a weak random source, the min-entropy rate (i.e. ratio of min-entropy to bit-length) of the source is nearly preserved". Precisely if $X \in \{0,1\}^n$

[9] Detailed description of Merkle Hash Tree based POR (MHT-POR) is given in Sec 2.2.1 of the full paper [1]

$\mathsf{S}(F, 1^\lambda)$ Summary function.

Input: An M-bit file $F \in \{0,1\}^M$ and security parameter λ in unary form.

Extract: Choose random seed s from domain $\{0,1\}^{\ell_s}$ and compute $Y := \mathsf{Extractor}(F; s)$.

Expand: Apply Erasure-Correcting-Code on Y to obtain $\hat{Y} = (\hat{Y}_1, \hat{Y}_2, \ldots, \hat{Y}_n)$ such that Y can be completely recovered from any αn blocks among $\{\hat{Y}_1, \hat{Y}_2, \ldots, \hat{Y}_n\}$, where constant $\alpha \in (0,1)$ is some system parameter. Generate POR-key pair $(pk, sk) := \mathsf{POR.KeyGen}(1^\lambda)$, and authentication tags $\{\sigma_i\}_{i=1}^n := \mathsf{POR.Tag}(sk, \{\hat{Y}_i\}_{i=1}^n)$. Let $\pi_F = (pk, sk, \{\sigma_i\}_{i=1}^n)$.
Note: As mentioned in [12], in the construction of PoW, the decoding algorithm of the above Erasure-Correcting-Code is not required to be practical, since the decoding algorithm will not be invoked in the legitimate application of PoW.

Output: The summary value of file F is $\psi = (s, \alpha, \pi_F)$. Output ψ.

$\langle \mathsf{P}(F), \mathsf{V}(\psi) \rangle$ Interactive proof system between verifier (cloud storage server) and prover (cloud storage client).

Input: The prover has file F as input and the verifier has a summary value $\psi = (s, \alpha, \pi_F)$ as input, where $\pi_F = (pk, sk, \{\sigma_i\}_{i=1}^n)$.

V1: Verifier finds $c = \lceil \log_{1-\alpha} \epsilon \rceil$ (i.e. c is the smallest integer such that $(1-\alpha)^c \leq \epsilon$) and computes $(C, \Psi_F, \Psi_\sigma) := \mathsf{POR.GenChal}(pk, n, c)$. Verifier sends $(C, s, \alpha, pk, \Psi_F)$ to the prover.

P1: Prover runs the extractor algorithm to obtain $Y := \mathsf{Extractor}(F; s)$, and re-generate the erasure code $\hat{Y}$ from Y using the same Erasure-Correcting-Code with the same parameter α. Prover divides $\hat{Y}$ into n blocks $\hat{Y}_1, \ldots, \hat{Y}_n$ and computes $\bar{F} := \mathsf{POR.GenProof}_{\mathsf{data}}(pk, \{\hat{Y}_i\}_{i=1}^n, C, \Psi_F)$. Prover sends $\bar{F}$ to verifier.

V2: Verifier computes $\bar{\sigma} := \mathsf{POR.GenProof}_{\mathsf{tag}}(pk, \{\sigma_i\}_{i=1}^n, C, \Psi_\sigma)$ and $b := \mathsf{POR.Verify}(K, \bar{F}, \bar{\sigma}, \Psi_F, \Psi_\sigma) \in \{\mathtt{Accept}, \mathtt{Reject}\}$, where K is pk if the POR scheme supports public key verification; otherwise K is sk.

Output: Output $b \in \{\mathtt{Accept}, \mathtt{Reject}\}$.
*Note: The subset C requires $|C| \log n$ bits communication cost. We can reduce this communication cost by using Goldreich [29]'s (δ, γ)-*hitter sampler[a] *to represent C compactly with only $\log n + 3\log(1/\gamma)$ bits of public random coins.*

[a] Goldreich [29]'s (δ, γ)-hitter guarantees that, for any subset $W \subset [1, n]$ with size $|W| \geq (1-\delta)n$, $\mathsf{Pr}[C \cap W \neq \emptyset] \geq 1 - \gamma$. Readers may refer to [29,30] for more details.

Fig. 1. PoW = RE + POR: A Generic Construction of PoW using Randomness Extractor $\mathsf{Extractor}(\cdot\ ;\ \cdot)$ and POR scheme ($\mathsf{KeyGen}, \mathsf{Tag}, \mathsf{GenChal}, \mathsf{GenProof}_{\mathsf{data}}, \mathsf{GenProof}_{\mathsf{tag}}, \mathsf{Verify}$). The completeness of the constructed PoW scheme is straightforward.

has δn min-entropy and $X[S] \in \{0,1\}^t$ is the projection of X onto a random set $S \subset [n]$ of t positions, then with high probability, $X[S]$ is statistically close to a random variable with $\delta' t$ min-entropy. We consider the difference $(\delta t - \delta' t)$ as the entropy loss in sampling t bits. Nisan and Zuckerman(Lemma 11 in [24]) gave

a sampling algorithm where $\delta' = c\delta/\log(1/\delta)$ for some small positive constant c. Vadhan (Lemma 6.2 in [25]) improved their result and allows $\delta' = (\delta - 3\tau)$ for sufficiently small positive constant τ.

We brief the existing approach [24,38] as below: (1) Independently and randomly choose l number of seeds, in order to get l samples $X_1, \ldots, X_l$ from the input weak source F, which has min-entropy rate δ. (2) Show that $(X_1, \ldots, X_l)$ is a δ'-block-wise source with δ' close to δ, i.e. for each $i \in [l]$, conditional on $(X_1, \ldots, X_i)$, the random variable X_{i+1} has min-entropy rate at least δ'. (3) Apply existing randomness extractor on the *structured* weak random source $(X_1, \ldots, X_l)$ to generate almost-uniform random output $(y_1, \ldots, y_l)$.

Roughly speaking, in the analysis of the above approach in [24,38], to extract each block y_i, the remaining min-entropy of the input F reduces by $|X_i|$ bits—the bit-length of X_i. Unlike previous works [24,25,38], we do not generate block-wise source as intermediate product, and manage to show that the remaining min-entropy of the input F, after extracting each block y_i, reduces by $|y_i|$ bits—the bit-length of y_i which is much smaller than $|X_i|$. Readers may find definition and calculation of remaining (or conditional) min-entropy $\tilde{\mathbf{H}}_\infty(A|B)$ of variable A given variable B in the full paper [1]. In this jargon, we manage to switch the conditional variable B from X_i (as previous works) to y_i in the analysis of our new design.

Extractor($F; s, s'$) This extractor algorithm will serve as a subroutine to construct PoW scheme.

Input: An M-bit file $F \in \{0,1\}^M$; $s \in \{0,1\}^{r_0}$ and $s' \in \{0,1\}^{r_1}$ are true random seeds, where $r_0 + r_1 = \rho$.

Sample-then-Extract-Loop:

Let $s_1 := s$ and $s'_1 := s'$. Let $h_F := \mathtt{SHA256}(F)$ with $|h_F| \leq \rho$.

For each i from 1 to ℓ:

Sample: Independently and randomly sample t *distinct* indices from the set $[M]$, using random seed s_i, to obtain $S_i := \mathsf{Samp}([M], t;\ s_i) \subset [M]$.

Extract: Compute $y_i := \mathsf{Ext}(h_F \| \ F[S_i];\ s'_i) \in \{0,1\}^\rho$. Let s_{i+1} be the prefix of bit-length r_0 of bit-string y_i, and s'_{i+1} be the suffix of bit-length r_1 of bit-string y_i.

Note: The hash value h_F is added into the input of Ext*, in order to ensure that any change in file F will lead to significant change in the output of randomness extractor.*

Output: Let $Y := y_1 \| y_2 \| \ldots \| y_\ell \in \{0,1\}^{\rho\ell}$. The output is Y.

Fig. 2. A Novel Randomness Extractor with Large Output Size and Short Seed. Ext is some existing strong randomness extractor and Samp is some existing sampling algorithm.

Theorem 3. *Let $t = M^c$ and $\tau = M^{-c}$ for constant $c \in (0,1)$. Let* Ext $: \{0,1\}^{t+256} \times \{0,1\}^{r_1} \to \{0,1\}^\rho$ *be a strong (k_0, ϵ_0)-extractor. Let* Samp *be an (μ, θ, γ)-averaging sampler [25,38]. Then the algorithm* Extractor $: \{0,1\}^M \times$

$\{0,1\}^\rho \to \{0,1\}^{\rho\ell}$ *constructed in Figure 2 is a* (k_1, ϵ_1)*-extractor, where* $\rho = \lambda + \log(M/t) + \log(1/\gamma) \cdot poly(1/\theta)$*,* $\rho \cdot \ell = k_1 - (k_0 + 3)M^{1-c}$*, and* $\epsilon_1 = 5\ell(\epsilon_0 + \gamma + 2^{-\lambda} + 2^{-\Omega(\tau M)})$.

We make the following remarks: (1) Our algorithm in Figure 2 requires about $1/\ell$ fraction of the amount of random bits required by [24], since [24] requires that all of sampling seeds $s_1, s_2, \ldots, s_\ell$ should be independent randomness. (2) The choice of value $t = M^c$ ensure that there will be sufficient remaining min-entropy in the last sample (worst case), and this value of sample size t would be much larger than required for the first few samples (good cases). One may use different sample size t_i for the i-th sample ($t_1 < t_2 < t_3 \ldots < t_\ell = M^c$), in order to reduce the IO reading. (3) Alternatively, we may choose hitter-sampler [29] as in [24] instead of averaging sampler, in order to reduce the seed length ρ (only $\mathcal{O}(\lambda + \log M)$ bits) at the cost of larger value of t. (4) In practice, one may use Tabulation Hashing [39] or CBC-MAC or HMAC as the underlying extractor algorithm Ext (possibly in the companion with hitter sampler which allows small ρ), as analyzed by Dodis *et al.* [40].

To prove Theorem 3, we introduce Lemma 4 and Lemma 5.

Lemma 4 (Amplification). *Suppose the algorithm* $\overline{\mathsf{Ext}} : \{0,1\}^M \times \{0,1\}^\rho \to \{0,1\}^\rho$ *defined as*

$$\overline{\mathsf{Ext}}(X; (s, s')) \stackrel{\text{def}}{=} \mathsf{Ext}\Big(\mathtt{SHA256}(X) \parallel X[\mathsf{Samp}(s)];\ s'\Big) \tag{4}$$

is a strong (k_2, ϵ_2)*-extractor. Then* $\mathsf{Extractor} : \{0,1\}^M \times \{0,1\}^\rho \to \{0,1\}^{\rho\ell}$ *constructed in Figure 2 is a* (k_1, ϵ_1)*-extractor, where* $k_1 \geq k_2 + \rho(\ell - 1) + \lambda$ *and* $\epsilon_1 = 5\ell(\epsilon_2 + 2^{-\lambda})$.

Our proof for Lemma 4 in full paper [1] is an analog of *hybrid proof technique* for (computational) indistinguishability [41].

Lemma 5 (Theorem 6.3 [25], sample-then-extract). *Let* $1 \geq \overline{\delta} \geq 3\tau > 0$*. Suppose that* $\mathsf{Samp} : \{0,1\}^{r_0} \to [M]^t$ *is an* (μ, θ, γ) *averaging sampler with distinct samples for* $\mu = (\overline{\delta} - 2\tau)/\log(1/\tau)$ *and* $\theta = \tau/\log(1/\tau)$ *and that* $\mathsf{Ext} : \{0,1\}^{t+256} \times \{0,1\}^{r_1} \to \{0,1\}^\rho$ *is a strong* $(k_0 = (\overline{\delta} - 3\tau)t,\ \epsilon_0)$*-extractor. Let* $\rho = r_0 + r_1$ *and define* $\overline{\mathsf{Ext}} : \{0,1\}^M \times \{0,1\}^\rho \to \{0,1\}^\rho$ *by*

$$\overline{\mathsf{Ext}}(X; (s, s')) \stackrel{\text{def}}{=} \mathsf{Ext}\Big(\mathtt{SHA256}(X) \parallel X[\mathsf{Samp}(s)];\ s'\Big) \tag{5}$$

Then $\overline{\mathsf{Ext}}$ *is a strong* (k_2, ϵ_2)*-extractor with* $k_2 = \overline{\delta}M$ *and* $\epsilon_2 = \epsilon_0 + \gamma + 2^{-\Omega(\tau M)}$*. Note: As mentioned in [25],* τ *could be arbitrarily small and approaches 0. In this paper, we set* $\tau = M^{-c}$ *for some constant* $c \in (0,1)$.

Computational Complexity. Recall that, in order to reduce computation cost, we could choose different sample size t_j for iteration j, where $t_1 < t_2 < \ldots < t_\ell = t = M^c$. The computational complexity of our proposed randomness extractor can be measured by the total number of bits read (or sampled) from the file (double counting repeated bits), i.e. the sum of t_j for $j \in [\ell]$. We will give an upper bound on the sum of t_j.

Lemma 6 (Complexity). *Suppose* $M^{1-c} \geq 2$. *The total number of bits (i.e.* $\sum_{j=1}^{\ell} t_j$*) of input file* F *accessed by the randomness extractor in Figure 2 is in* $\mathcal{O}(M \log \ell)$.

Note: (1) If the underlying extractor Ext *is Tabulation Hashing, then the constant behind the big-O notation is very small—around 2. (2) Multiple access to the same bit will be counted with its frequency. (3) The proof of this lemma is in full paper [1].*

We remark that the extractor algorithm in Figure 2 can be modified into m concurrent threads/processes, while increasing the seed size by m times.

6 Conclusion and Open Problems

We were the first one to bridge construction of PoW with randomness extractor and proofs of retrievability. We also proposed a novel randomness extractor with large output size, which improves existing works in both seed length and entropy loss (i.e. the difference between entropy of input and output). Our proofs of ownership scheme can be applied in client-side deduplication of encrypted (unencrypted, too) data in cloud storage service, and the new randomness extractor may have independent interest.

Whether "partition-then-extract" approach works for *any* distribution of input file and how to apply pseudo-entropy extractor (e.g Yao-Entropy extractor) to construct proofs of ownership scheme, remain two open problems.

References

1. Xu, J., Zhou, J.: Leakage Resilient Proofs of Ownership in Cloud Storage, Revisited. Cryptology ePrint Archive, Report 2013/514 (2013), http://eprint.iacr.org/2013/514
2. iHS iSuppli: Cloud Storage Services Now Have Over 375M Users, Could Reach 500M By Year-End, http://goo.gl/BO6zWy
3. Blog, A.: Amazon S3 goes exponential, now stores 2 trillion objects, http://goo.gl/NUIEny, http://gigaom.com/2013/04/18/amazon-s3-goes-exponential-now-stores-2-trillion-objects/
4. Blog, W.A.S.T.: Windows Azure Storage – 4 Trillion Objects and Counting, http://blogs.msdn.com/b/windowsazurestorage/archive/2012/07/20/windows-azure-storage-4-trillion-objects-and-counting.aspx
5. Blog, D.: Over 175 million people using Dropbox and more than a billion files synced each day, https://blog.dropbox.com/2013/07/dbx/
6. SNIA: Understanding Data De-duplication Ratios. white paper, http://www.snia.org/sites/default/files/Understanding_Data_Deduplication_Ratios-20080718.pdf
7. Douceur, J., Adya, A., Bolosky, W., Simon, D., Theimer, M.: Reclaiming space from duplicate files in a serverless distributed file system. In: ICDCS 2002: International Conference on Distributed Computing Systems (2002)
8. Douceur, J., Bolosky, W., Theimer, M.: US Patent 7266689: Encryption systems and methods for identifying and coalescing identical objects encrypted with different keys (2007)

9. Storer, M., Greenan, K., Long, D., Miller, E.: Secure Data Deduplication. In: StorageSS 2008: ACM International Workshop on Storage Security and Survivability, pp. 1–10 (2008)
10. Bellare, M., Keelveedhi, S., Ristenpart, T.: Message-Locked Encryption and Secure Deduplication. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 296–312. Springer, Heidelberg (2013), http://eprint.iacr.org/2012/631
11. Harnik, D., Pinkas, B., Shulman-Peleg, A.: Side Channels in Cloud Services: Deduplication in Cloud Storage. IEEE Security and Privacy Magazine, Special Issue of Cloud Security 8(6) (2010)
12. Halevi, S., Harnik, D., Pinkas, B., Shulman-Peleg, A.: Proofs of ownership in remote storage systems. In: CCS 2011: ACM Conference on Computer and Communications Security, pp. 491–500 (2011), http://eprint.iacr.org/2011/207
13. Dropship: Dropbox api utilities (April 2011), https://github.com/driverdan/dropship
14. Storer, M., Greenan, K., Long, D., Miller, E.: Secure data deduplication. In: Proceedings of the 4th ACM International Workshop on Storage Security and Survivability, StorageSS 2008, pp. 1–10 (2008)
15. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof systems. SIAM Journal on Computing 18(1), 186–208 (1989)
16. Wikipedia: PlayStation Network outage, http://en.wikipedia.org/wiki/PlayStation_Network_outage
17. wired.com: Dropbox Left User Accounts Unlocked for 4 Hours Sunday, http://www.wired.com/threatlevel/2011/06/dropbox/, http://blog.dropbox.com/?p=821
18. Twitter: Tweetdeck, http://money.cnn.com/2012/03/30/technology/tweetdeck-bug-twitter/
19. Juels, A., Kaliski, Jr., B.: Pors: proofs of retrievability for large files. In: CCS 2007: ACM Conference on Computer and Communications Security, pp. 584–597 (2007)
20. Ateniese, G., Burns, R., Curtmola, R., Herring, J., Kissner, L., Peterson, Z., Song, D.: Provable data possession at untrusted stores. In: CCS 2007: ACM Conference on Computer and Communications Security, pp. 598–609 (2007)
21. Xu, J., Chang, E.C., Zhou, J.: Weak Leakage-Resilient Client side Deduplication of Encrypted Data in Cloud Storage. In: ASIACCS 2013: Proceedings of the 8th ACM Symposium on Information, Computer and Communications Security (Full Paper), pp. 195–206 (2013), http://eprint.iacr.org/2011/538
22. Bellare, M., Keelveedhi, S., Ristenpart, T.: DupLESS: Server-Aided Encryption for Deduplicated Storage (will appear in Usenix Security Symposium 2013). Cryptology ePrint Archive, Report 2013/429 (2013), http://eprint.iacr.org/2013/429
23. Ng, W.K., Wen, Y., Zhu, H.: Private data deduplication protocols in cloud storage. In: SAC 2012: Proceedings of the 27th Annual ACM Symposium on Applied Computing, pp. 441–446 (2012)
24. Nisan, N., Zuckerman, D.: Randomness is linear in space. Journal of Computer and System Sciences 52(Special issue on STOC 1993) , 43–52 (1996)
25. Vadhan, S.: Constructing Locally Computable Extractors and Cryptosystems in the Bounded-Storage Model. J. Cryptol. 17(1), 43–77 (2004)
26. Stinson, D.R.: Universal hash families and the leftover hash lemma, and applications to cryptography and computing. Journal of Combinatorial Mathematics and Combinatorial Computing 42, 3–31 (2002)
27. Shacham, H., Waters, B.: Compact Proofs of Retrievability. In: Pieprzyk, J. (ed.) ASIACRYPT 2008. LNCS, vol. 5350, pp. 90–107. Springer, Heidelberg (2008)

28. Xu, J., Chang, E.C.: Towards efficient proof of retrievability. In: ASIACCS 2012: Proceedings of the 7th ACM Symposium on Information, Computer and Communications Security (Full Paper) (2012), http://eprint.iacr.org/2011/362
29. Goldreich, O.: A Sample of Samplers - A Computational Perspective on Sampling (survey). Electronic Colloquium on Computational Complexity (ECCC) 4(20) (1997)
30. Dodis, Y., Vadhan, S., Wichs, D.: Proofs of Retrievability via Hardness Amplification. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 109–127. Springer, Heidelberg (2009)
31. Xu, J., Chang, E.C., Zhou, J.: Leakage-Resilient Client-side Deduplication of Encrypted Data in Cloud Storage. Cryptology ePrint Archive, Report 2011/538 (2011), http://eprint.iacr.org/2011/538
32. Chang, E.C., Xu, J.: Remote Integrity Check with Dishonest Storage Server. In: Jajodia, S., Lopez, J. (eds.) ESORICS 2008. LNCS, vol. 5283, pp. 223–237. Springer, Heidelberg (2008), http://eprint.iacr.org/2008/346
33. Zheng, Q., Xu, S.: Secure and efficient proof of storage with deduplication. In: CODASPY 2012: ACM conference on Data and Application Security and Privacy, pp. 1–12 (2012)
34. Barak, B., Dodis, Y., Krawczyk, H., Pereira, O., Pietrzak, K., Standaert, F.-X., Yu, Y.: Leftover Hash Lemma, Revisited. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 1–20. Springer, Heidelberg (2011)
35. Gabizon, A., Raz, R., Shaltiel, R.: Deterministic Extractors for Bit-Fixing Sources by Obtaining an Independent Seed. SIAM Journal on Computing 36(4), 1072–1094 (2006)
36. Pietro, R.D., Sorniotti, A.: Boosting Efficiency and Security in Proof of Ownership for Deduplication. In: ASIACCS 2012: ACM Symposium on Information, Computer and Communications Security (Full Paper) (2012)
37. Ateniese, G., Burns, R., Curtmola, R., Herring, J., Khan, O., Kissner, L., Peterson, Z., Song, D.: Remote data checking using provable data possession. ACM Transactions on Information and System Security 14, 12:1–12:34 (2011)
38. Vadhan, S.: Pseudorandomness. Foundations and Trends in Theoretical Computer Science 7(1-3), 1–336 (2012)
39. Patrascu, M., Thorup, M.: The power of simple tabulation hashing. In: STOC 2011: ACM Symposium on Theory of Computing, pp. 1–10 (2011)
40. Dodis, Y., Gennaro, R., Håstad, J., Krawczyk, H., Rabin, T.: Randomness Extraction and Key Derivation Using the CBC, Cascade and HMAC Modes. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 494–510. Springer, Heidelberg (2004)
41. Goldreich, O.: Foundations of Cryptography. Basic Applications, vol. 2. Cambridge University Press (2004)