

Boosting Efficiency and Security in Proof of Ownership for Deduplication

Roberto Di Pietro
Università di Roma Tre, Italy
dipietro@mat.uniroma3.it

Alessandro Sorniotti
IBM Research – Zurich, Switzerland
aso@zurich.ibm.com

ABSTRACT

Deduplication is a technique used to reduce the amount of storage needed by service providers. It is based on the intuition that several users may want (for different reasons) to store the same content. Hence, storing a single copy of these files is sufficient. Albeit simple in theory, the implementation of this concept introduces many security risks. In this paper we address the most severe one: an adversary (who possesses only a fraction of the original file, or even just partially colluding with a rightful owner) claiming to possess such a file. The paper's contributions are manifold: first, we introduce a novel Proof of Ownership (POW) scheme that has all features of the state-of-the-art solution while incurring only a fraction of the overhead experienced by the competitor; second, the security of the proposed mechanisms relies on information theoretical (combinatoric) rather than computational assumptions; we also propose viable optimization techniques that further improve the scheme's performance. Finally, the quality of our proposal is supported by extensive benchmarking.

Categories and Subject Descriptors

H.3.5 [Information Systems]: Information storage and retrieval—*Online information services*

General Terms

Security

Keywords

Cloud Security, Deduplication, Proof of Ownership

1. INTRODUCTION

The rapid surge in cloud service offerings has resulted in a sharp drop in prices of storage services, and in an increase in the number of customers. Through popular providers, like Amazon s3 and Microsoft Azure, and backup services,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ASIACCS '12, May 2–4, 2012, Seoul, Korea.

Copyright 2012 ACM 978-1-4503-1303-2/12/05 ...\$10.00.

like Dropbox and Memopal, storage has indeed become a commodity. Among the reasons for the low prices, we find a strong use of multitenancy, the reliance on distributed algorithms run on top of simple hardware, and an efficient use of the storage backend thanks to compression and deduplication.

Deduplication is the process of avoiding having to store the same data multiple times. It leverages the fact that large data sets often exhibit high redundancy. Examples include, for instance, common email attachments, financial records, with common headers and semi-identical fields, and popular media content—such as music, video—likely to be owned (and stored) by many users.

There are four different deduplication strategies, depending on whether deduplication happens at the client side (i.e. before the upload) or at the server side, and whether deduplication happens at a block level or at a file level. Deduplication is most rewarding when it is triggered at the client side, as it also saves the bandwidth of the upload. For these reasons, deduplication is a critical enabler for a number of popular and successful storage services (e.g. Dropbox, Memopal) that offer cheap remote storage to the broad public by performing client-side deduplication, thus saving both the network bandwidth and the storage costs associated to processing the same content multiple times.

Security Threats to Deduplication.

Harnik *et al.* [10] have identified a number of threats that can affect a storage system performing client-side deduplication. These threats, briefly reviewed in the following, can be turned into practical attacks by any user of the system.

A first set of attacks targets the privacy and confidentiality of users of the storage system. For instance, a user can check whether another user has already uploaded a file by trying to upload it as well and by checking whether the upload actually takes place. This attack is particularly relevant for rare files that may reveal the identity of the user who performed the upload.

A different type of attack attempts to subvert the intended use of a given storage system. Two users who have no direct connectivity may, for instance, try to use the storage system as a covert channel. For instance, to exchange a bit of information, the two users would pre-agree on two files; then the transmitting user would upload one of the two files, and the receiving user would detect which one gets deduplicated and would output either 0 (for the first file) or 1 (for the second).

Finally, users may abuse a storage system by turning it

into a content distribution network: users who wish to exchange large files leveraging the large bandwidth available to the servers of the storage systems may upload only a single copy of such a file and share the short token that triggers deduplication (in most cases, the hash digest of the file) among all users who wish to download the content. Real-world examples of such an approach include Dropship [6].

Proof of Ownership (POW).

To remedy the security threats mentioned above, the concept of Proof of Ownership (POW) has been introduced [9]. POW schemes essentially address the root-cause of the aforementioned attacks to deduplication, namely, that the proof that the client owns a given file (or block of data) is solely based on a static, short value (in most cases the hash digest of the file), whose knowledge automatically grants access to the file.

POW schemes are security protocols designed to allow a server to verify (with a given degree of assurance) whether a client owns a file. The probability that a malicious client engages in a successful POW run must be negligible in the security parameter, even if the malicious client knows a (relevant) portion of the target file. A POW scheme should be efficient in terms of CPU, bandwidth and I/O for the server and all legitimate clients: in particular, POW schemes should not require the server to load the file (or large portions of it) from its back-end storage at each execution of POW.

Halevi *et al.* [9] have introduced the first practical cryptographic protocol that implements POW. This seminal work, however, suffers from a number of shortcomings that might hinder its adoption. The first is that the scheme has extremely high I/O requirements at the client-side: it either requires clients to load the entire file in memory or to perform random block accesses with an aggregate total I/O higher than the size of the file itself. Secondly, the scheme takes a heavy computational toll on the client. Thirdly, its security is admittedly based on assumptions that are hard to verify.

Contributions.

In this paper, we present a novel scheme for secure Proof of Ownership. Our scheme attains several ambitious goals: i) its I/O and computational costs do *not* depend on the input file size; ii) it is *very efficient* for a wide range of systems parameters; iii) it is *information-theoretically secure*; and, iv) it requires the server to keep a per-file state that is a *negligible fraction* of the input file size.

Roadmap.

The remainder of this paper is organised as follows: Section 2 reviews the state of the art; Section 3 defines system and security models; Section 4 presents our basic scheme and its two improvements; Section 5 describes the implementation and benchmarks; Section 6 contains a discussion on the performance and an optimization that captures all the preceding ones, while Section 7 presents our conclusions.

2. RELATED WORK

Douceur *et al.* [5] study the problem of deduplication in a multitenant system in which deduplication has to be reconciled with confidentiality. The authors propose the use of

convergent encryption, i.e., deriving keys from the hash of the plaintext, so that two users will produce the same ciphertext from the same plaintext block, and the ciphertext can then be deduped. Storer *et al.* [16] point out some security problems of convergent encryption, proposing a security model and two protocols for secure data deduplication.

The seminal work of Harnik *et al.* [10] first discusses the shortcomings of client-side deduplication, and later presents some basic solutions to the problem. In particular, attacks to privacy and confidentiality can be addressed without a full-fledged POW scheme by triggering deduplication only after a small, but random, number of uploads.

Whereas POW deals with the assurance that a client indeed possesses a given file, Provable Data Possession (PDP) and Proof of Retrievability (PoR) deal with the dual problem of ensuring—at the client-side—that a server still stores the files it ought to. PDP is formally introduced by Ateniese and colleagues [3, 2]. A number of earlier works already address remote integrity checking, see the Related Work Section of [2] for more details. Ateniese *et al.* [4] present a dynamic PDP scheme based on symmetric cryptography, and show how relaxing the requirement of public verifiability allows a much more lightweight scheme. The scheme is dynamic in that it allows data blocks to be appended, modified and deleted. Erway *et al.* [7] present formal definitions of Dynamic PDP together with two protocols allowing also block insertion. PoR schemes, introduced by Juels and Kaliski [11] combine message authentication code-based data verification with error-correcting codes to allow a client to download pre-determined subsets of blocks and check whether their MAC matches the pre-computed one: the use of ECC ensures that small changes in the data are detected with high probability.

2.1 The State-of-the-Art Solution

Next we shall describe in detail the POW scheme presented by Halevi *et al.* [9], as it represents the state-of-the-art solution our solution will be compared against.

The authors present three schemes that differ in terms of security and performance. All three involve the server challenging the client to present valid sibling paths for a subset of leaves of a Merkle trees [14]. Both the client and the server build the Merkle tree; the server only keeps the root and challenges clients that claim to possess the file. The Merkle tree is built on a buffer, whose content is derived from the file, and pre-processed in three different ways for the three different schemes.

The first scheme applies erasure coding on the content of the original file; the erasure-coded version of the file is the input for construction of the Merkle tree.

The second scheme pre-processes the file with a universal hash function instead of erasure coding, to the same end: the file is hashed to an intermediate reduction buffer whose size is sufficiently large to discourage its sharing among colluding users, but not too big to be impractical. The authors settle for a size of 64 MiB. This buffer is then used as an input for the construction of the Merkle tree.

The third scheme, which is the one we will compare our solution with, follows the same approach, but substitutes universal hashing with a mixing and a reduction phase that “hash” the original file into the reduction buffer mentioned above. In the remainder of this paper, we shall refer to this scheme as b-POW.

This scheme has two phases (see Figure 1 of [9]): the first phase populates the reduction buffer by xoring each block of the original file in four random positions in the buffer (after performing a bit shift). The mixing phase amplifies the confusion and diffusion in the reduction buffer by xoring together random positions of the reduction buffer.

3. SYSTEM MODEL

The system is composed of two main principals, the client \mathcal{C} and the server \mathcal{S} . Both \mathcal{C} and \mathcal{S} are computing nodes with network connectivity. \mathcal{S} has a large back-end storage facility and offers its storage capacity to \mathcal{C} ; \mathcal{C} uploads its files and can later download them. During the upload process, \mathcal{S} attempts to minimize the bandwidth and to optimize the use of its storage facility by determining whether the file the client is about to upload has already been uploaded by another user. If so, the file does not need to be uploaded and we say it undergoes deduplication.¹ Note that a trivial solution would be to transfer the file from the client to the server, and later have the checks performed on the server-side. However, this solution is highly bandwidth demanding, and also sacrifices another benefit of deduplication: the reduction of the completion time on both the client and server-side.

A further requirement on the server-side is to minimize accesses to its back-end storage system: for example, a protocol that requires to access the file content at each interaction with a client to evaluate the potential for deduplication would not meet this requirement. We assume, however, that \mathcal{S} has a front-end storage facility, whose capacity is a small fraction of the capacity of the back-end one, and that can be used to store per-file information. We finally assume that server-side computational power is abundant and cheap, especially if the required computation does not need to be executed immediately but can be deferred to moments of low system load.

\mathcal{C} is assumed to have limited resources in terms of computational power and I/O capability, and therefore one of the design guidelines of the scheme is to minimize the scheme’s client-side computational and I/O footprint. \mathcal{C} and \mathcal{S} engage in an interactive protocol. As previously mentioned, minimization of the latency of this protocol is another important objective.

3.1 Adversarial Model

In the context of POW protocols, \mathcal{S} is considered to be a trusted entity that abides by the rules of the protocol as its correct execution is in \mathcal{S} ’s best interest. \mathcal{C} , in contrast, is considered to be a malicious principal and consequently it cannot be assumed that it is bound by the rules of the protocol.

Given a target file f_* , the objective of a malicious client $\tilde{\mathcal{C}}$ is to convince the server that $\tilde{\mathcal{C}}$ owns f_* , despite this not being the case. It is assumed that $\tilde{\mathcal{C}}$ does not know f_* in its entirety; however we assume that $\tilde{\mathcal{C}}$ knows an arbitrarily large fraction of it. The estimated upper bound on the fraction of f_* known to $\tilde{\mathcal{C}}$ will be one of the inputs of the system, playing a role in the scheme’s security analysis. Several malicious clients can collude and share information, for instance, about past protocol rounds. $\tilde{\mathcal{C}}$ may even collude

¹The privacy issues raised by this solution are out of the scope of this paper, and some preliminary solutions have been proposed in [10].

with other clients that indeed possess f_* , and may receive arbitrary information about the file from them, including its content—but never in its entirety. However we assume that $\tilde{\mathcal{C}}$ cannot interact with clients that possess f_* during the challenge between \mathcal{S} and $\tilde{\mathcal{C}}$ over f_* , as otherwise such clients could easily circumvent the security of the protocol by answering instead of $\tilde{\mathcal{C}}$.

4. OUR SCHEME

In this section, we shall describe our POW solution. Our scheme consists two separate phases: in the first phase, the server receives a file for the first time and it pre-computes the responses for a number of POW challenges related to the file. Computation of POW challenges for a given file is carried out both upon receiving an upload request for a file that is not yet present at the server, and when the stock of the previously computed challenge/responses has been depleted. The number of challenges to be precomputed is a tunable system parameter.

The second phase is triggered by the client when it sends to the server a unique identifier for a file it wishes to prove possession of. The server chooses an unused challenge from the pre-computed ones for that file and sends it to the client; the client derives the response based on its knowledge of the file, and sends the response to the server. The server then checks whether the client’s response matches the pre-computed one.

In the following sections, we will detail our scheme. We will do so incrementally, starting with an initial scheme s-POW, and later presenting two improved variants, s-POW1 and s-POW2. After the outcome of our benchmarking in Section 5 and the discussion in Section 6 we will show in Section 6.1 how to build a final scheme, s-POW3, that combines the best features of the other solutions.

4.1 s-POW

The basic idea behind s-POW is that the probability that a malicious user is able to output the correct value of K bits of the file, with each bit selected at a random position in the file, is negligible in the security parameter k – assuming an upperbound on the subset of bits of the file known to the attacker. Therefore, a response to a challenge will be a bit strings of size K , constructed as the concatenation of K bits of the original file, picked at K random positions.

Let us describe s-POW in more detail. The server keeps a hash-map data structure \mathfrak{F} that maps strings of finite size to 4-tuples; these tuples contain a file pointer ptr , an array of responses $res[]$ and two indexes, id_c and id_u . The first index keeps track of the highest challenge computes so far, while the second counts the number of challenges used. By default, both indexes are initialized to zero; $res[]$ is initialized with an array of empty strings, and ptr is associated to an unassigned pointer. The search key into the hash-map is the hash digest of the file; given a digest d , $\mathfrak{F}[d]$ represents the tuple mapped to d : $\mathfrak{F}[d] = \perp$ if d has not yet been associated with any tuple.

Also, let H be a cryptographic hash function and F_s a pseudo-random number generator taking s as seed. For the sake of simplicity, we assume that F_s generates integers ranging from zero to the size of the file in bits minus one. `get_bit` is a macro taking as input a file pointer and a bit position

²Protection from this attack is outside the scope of this paper.

ALGORITHM 1: Server-side algorithm: the server precomputes the challenges for a file.

Input: A hash digest d ; a number n of responses that need to be pre-computed and a response bit length K .

Output: An updated response vector.

```

begin
   $f_d \leftarrow \mathfrak{F}[d]$ ;
  for  $i \in [0, n - 1]$  do
     $ctr \leftarrow f_d.id_c + i$ ;
     $s \leftarrow F_{S_{MK}}(d || ctr)$ ;
    for  $j \in [0, K - 1]$  do
       $pos \leftarrow F_s(j)$ ;
       $res[i] = res[i] || get\_bit(f_d.ptr, pos)$ ;
    end
  end
   $f_d.id_c = ctr + 1$ ;
  return  $\perp$ ;
end

```

and producing as output the corresponding bit value. Finally, let S_{MK} be the server master secret.

Algorithm 1 describes the operations that occur at the server-side when either a new file has been uploaded, or the precomputed responses of an old file have been exhausted and new ones need to be generated. The server computes n challenges at a time: this allows optimization of the I/O operations. For each challenge, a fresh new random seed is computed using index id_c , the digest d of the file, and the server master key S_{MK} . Then, K random positions are generated using F , and the bits in the corresponding positions are concatenated to form the response to the id_c -th challenge.

Algorithm 2 describes how a client replies when challenged by the server; the client essentially uses the challenge seed s received from the server that is needed to generate the K random position over the file, and collects the bit-value of the file at these K random positions to form the response $resp$.

ALGORITHM 2: Client-side algorithm: the client computes the response to a challenge of the server.

Input: A file f and a challenge seed s .

Output: A response bit string.

```

begin
  let  $res$  be an empty string;
  for  $j \in [0, K - 1]$  do
     $pos \leftarrow F_s(j)$ ;
     $res = res || get\_bit(f, pos)$ ;
  end
  return  $res$ ;
end

```

Algorithm 3 shows the overall protocol executed between client and server. The protocol starts with the client computing the hash of the file and sending it to the server with a request to store the associated file. The server checks whether the file already exists in the hash map. If not, the file needs to be uploaded and no challenge takes place. If a challenge is required, the server picks the first unused challenge for the given file, computes the associated seed and sends it to the client. The client is then able to invoke Algorithm 2 and compute the response, which is sent back to the server. The server checks the response for equality with

the precomputed one and outputs success or failure based on this check. At this stage, the server will assign that file to the set of files belonging to the client, so that later on the client can access it. Finally, if all precomputed challenges have been used up, the server invokes Algorithm 1 to repopulate the response vector.

ALGORITHM 3: The protocol of s-POW, expressed as a distributed algorithm run between \mathcal{C} and \mathcal{S} .

```

 $\mathcal{C}$  : upon upload of file  $f$  do
   $d \leftarrow H(file)$ ;
  send to  $SRV$  a store file request with  $d$ ;
end
 $\mathcal{S}$  : upon receipt of a store file request do
  if  $\mathfrak{F}[d] \neq \perp$  then
     $s \leftarrow F_{S_{MK}}(d || \mathfrak{F}[d].id_u)$ ;
    send to  $\mathcal{CLI}$  a challenge request with  $s$ ;
  else
    initialize  $\mathfrak{F}[d]$ ;
    receive  $f$  from  $\mathcal{CLI}$ ;
     $\mathfrak{F}[d].ptr \leftarrow f$ ;
  end
end
 $\mathcal{C}$  : upon receipt of a challenge request do
  invoke Algorithm 2 on input  $file$  and  $s$  to get  $resp$ ;
  send to  $SRV$  a challenge response with  $resp$ ;
end
 $\mathcal{S}$  : upon receipt of a challenge response do
  if  $resp = \mathfrak{F}[d].res[\mathfrak{F}[d].id_u \bmod n]$  then
     $\mathcal{CLI}$  succeeds;
  else
     $\mathcal{CLI}$  fails;
  end
   $\mathfrak{F}[d].id_u = \mathfrak{F}[d].id_u + 1$ ;
  if  $\mathfrak{F}[d].id_u \equiv 0 \bmod 0$  then
    invoke Algorithm 1;
  end
end
end

```

4.1.1 Security Analysis of s-POW

As introduced in Section 3, the goal of the adversary $\tilde{\mathcal{C}}$ is to pass the check performed by \mathcal{S} during the file uploading phase, while not owning the file in its entirety. In this way, $\tilde{\mathcal{C}}$ could later gain access to the file actually stored on the server. In the following, we analyse the security of our solution that is based on challenging the client on the value of K bits randomly chosen over the file that $\tilde{\mathcal{C}}$ claims to possess. Before exploring the security of s-POW, we remind the reader that the cryptographic digest d of the file f does not play any role in the security of the scheme, as we assume that this short value can be obtained by $\tilde{\mathcal{C}}$.

In accordance with the working hypothesis given in Section 3, we can assume that $\tilde{\mathcal{C}}$ owns (or has access to) a fraction $p = (1 - \epsilon)$ of the file. When confronted with a single-bit challenge posed by the server, two cases can occur: first, the requested bit belongs to the portion of the file in $\tilde{\mathcal{C}}$'s availability – let us indicate this event with w . This can happen with probability: $P(w) = (1 - \epsilon)$. Otherwise, we can assume $\tilde{\mathcal{C}}$ performs a (possibly educated) guess that results in a success probability g . Therefore, $\tilde{\mathcal{C}}$ can succeed on a single-bit challenge ($P(succ_1)$), under the assumption

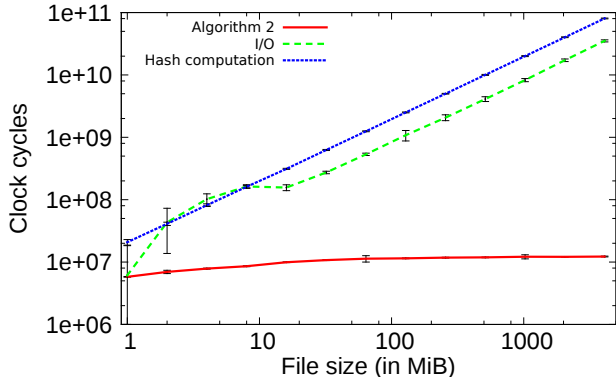


Figure 1: Comparison of the running time of each of the three main operations executed in s-POW as the input file size grows.

that $\epsilon > 0$, with probability

$$\begin{aligned}
 P(\text{succ}_1) &= P(\text{succ}_1 \wedge (w \vee \bar{w})) \\
 &= P(\text{succ}_1|w)P(w) + P(\text{succ}_1|\bar{w})P(\bar{w}) \\
 &= P(w) + gP(\bar{w}) \\
 &= (1 - \epsilon) + g(1 - (1 - \epsilon)) \\
 &= 1 - \epsilon(1 - g)
 \end{aligned}$$

However, \tilde{C} is confronted with K challenges, each being i.i.d. from the others. Therefore, the probability that \tilde{C} can successfully pass the check ($P(\text{succ})$) is

$$P(\text{succ}) = (1 - \epsilon(1 - g))^K \quad (1)$$

Equation 1 completely characterizes our security model. Indeed, once reasonable values are set for ϵ and g , and given a security parameter k , an appropriate value of the parameter K of s-POW – assuring that $P(\text{succ}) \leq 2^{-k}$ – can be simply derived as:

$$K = \left\lceil \frac{k \ln 2}{\epsilon(1 - g)} \right\rceil \quad (2)$$

Note that – in its present formulation (cf. Algorithm 1) – s-POW trades information-theoretical security for improved space efficiency, by deriving challenge seeds from a master secret. A simple way of achieving information-theoretical security would be in Algorithm 1 to generate a fresh random seed s for each new challenge, and to save it together with its pre-computed response.

Finally, note that equations 1 and 2 also highlight that K is not affected by the length of the file, because the only parameters involved are ϵ , the fraction of the file *unknown* to \tilde{C} and the value g .

4.2 s-POW1

In this section we propose a first improvement to s-POW. Given the extreme simplicity of s-POW, it may seem that there is very little room for optimization. Indeed, as we have seen in the previous section, the number K of bits the client is challenged on is a function of the security parameter of the scheme. Therefore, any reduction in K will inevitably alter the overall security of the scheme.

Section 5 describes the results of the benchmarking on our scheme. We will, however, already mention one of the re-

sults here to give the reader an idea of our improved scheme: Figure 1 shows the evolution of the clock cycles spent by the client in the execution of the three main components of s-POW as the size of the file grows. The three components are the I/O time, i.e., the time spent to access the file on disk and to load it in main memory, the hash time, i.e., the time spent in the computation of the hash digest of the file, and finally the time spent in the execution of Algorithm 2. Notice that the I/O and the hash time are by far more expensive than Algorithm 2. It is natural therefore to try to reduce the cost of these two components.

Let us recall that the computation of the hash is required because the server needs to be able to uniquely identify the file among those already stored (if this is the case) to compute the appropriate challenge seed and to compare the response of the client with the pre-computed one. However, the cryptographic properties of standard hash functions (one-wayness, preimage resistance and second preimage resistance) are not strictly needed in this setting. Indeed, the properties we are looking for are such that the hash function may be replaced with another function that: i) has a small probability of producing the same output given different files as input; ii) is computationally less expensive than a hash; and, iii) minimizes the required I/O.

Algorithm 2 is an excellent candidate for such a function, as it has a very small computational footprint and requires only a minimum number of I/O operations to retrieve the bits that constitute the output of a challenge. Consequently, we modify the overall protocol as shown in Algorithm 4.

ALGORITHM 4: Changes in the protocol of s-POW to achieve s-POW1.

Input: A file f .

\mathcal{C} : **upon** upload of file f **do**

 | invoke Algorithm 2 on input f and s_{pub} to get d ;

 | send to \mathcal{SRV} a store file request with d ;

end

The main difference with respect to the original version of the protocol for s-POW as shown in Algorithm 3 is that – at the client-side – Algorithm 2 (on input of a public seed s_{pub} , randomly generated and published as a parameter of the system) is invoked instead of the hash function H to generate the file digest d . As we shall see in Section 5, this change achieves a significant improvement in the performance of the scheme, especially at the client-side. Indeed, it is no longer necessary to scan the entire file to compute its hash, but only to perform a relatively small number of random accesses to its content.

To tolerate the scenario in which multiple files have the same digest d (produced by the invocation of s-POW on input of the public seed), the server has to keep a one-to-many map (instead of the previously used one-to-one map) between the output of the indexing function and the tuple containing indexes, file pointer, and array of pre-computed responses. In this scenario, the server would receive a single *resp* and compare it with all the pre-computed responses for all files indexed by the same value of d . If none of them matches, the client is uploading a new file. If one matches, the server concludes that the client owns the file associated with the matching precomputed response. However, this approach comes at the expense of a slightly higher usage rate of the precomputed challenges. Indeed, imagine that

there are two files $f1$ and $f2$ with the same digest d : a client owning $f1$ engages in the POW protocol with the server and receives a challenge seed $F_{S_{MK}}(d||i)$ for some value i of the current counter; that seed can no longer be used for a client owning $f2$ because if the challenge is leaked, a user colluding with \tilde{C} could precompute the correct response and send it to \tilde{C} . Not reusing challenges that have been disclosed implies that the usage rate of challenges for files indexed by the same key d is equal to the sum of the rates of requests for each of these files. However this does not constitute a problem, because the server has abundant computational power and can regularly schedule the pre-computation of challenges in periods of low system load.

4.2.1 Security Analysis of s-POW1

The security of the scheme is unchanged: indeed, even if an attacker were able to produce the correct value d for a given file (for instance, by receiving it from an accomplice), it would still need to generate the correct response to the challenge of the server. However we have shown in Section 4.1.1 that the probability of this happening is negligible in the security parameter.

The cryptographic hash function H previously used for indexing was collision resistant by definition. As explained above, H has been replaced with an s-POW invocation on input of a public seed. We therefore need to quantify the collision probability of such an indexing function, i.e. the probability that the digests $d1$ and $d2$ of two different files $f1$ and $f2$ are equal. We can derive this probability by assuming that two files are similar with a given probability z (z expressing the percentage of the bits of the two files in the same position that show the same value). Hence, for M files we have the probability of collision $P(coll)$:

$$\begin{aligned} P(coll) &\leq \binom{M}{2} P(d1 = d2) \\ &\leq \frac{M^2}{2} P(d1 = d2) = \frac{M^2}{2} z^K \end{aligned} \quad (3)$$

where K is the parameter of s-POW mentioned in the previous section. The above probability can still be considered negligible for practical instantiations of the scheme. For instance, for $M = 10^9$, $z = .95$ and³ $K = 1830$, $P(coll) \leq 2^{-75}$.

However, let us assume that collisions *do* happen. Then, as explained above, an invocation of the mapping \mathfrak{F} on d would return a set of m files. We then need to quantify the additional advantage that an adversary might have in passing the proof of ownership given that the server has to compare the client's response with m pre-computed ones instead of a single one. Let r_i be the event that $resp_*$, the response received by \tilde{C} equals $resp_i$, the i -th precomputed response of one of the m files in the set. Then it follows that the probability $P(succ)$ of \tilde{C} of succeeding the test over at least one out of the m files is

$$\begin{aligned} P(succ) &= P(r_1 \vee \dots \vee r_m) \leq mP(succ_{r_i}) \\ &= m(1 - \epsilon(1 - g))^K \end{aligned} \quad (4)$$

where the term $(1 - \epsilon(1 - g))^K$ comes from Equation 1. From 4 we conclude that m can therefore become an additional parameter of the system and can contribute to the

³See Section 5 on the sizing of the parameter K .

determination of the parameter K , even though its effect over K is only logarithmic: major changes in m will have very little effect over the value of K .

Another aspect we need to consider for the case in which an invocation of the mapping \mathfrak{F} on d returns a set of m files, is the probability that there are collisions among the m pre-computed responses for a given value of the index $\mathfrak{F}[d].id_u$; that is, the $\mathfrak{F}[d].id_u$ -th pre-computed response for file f_i is equal to the $\mathfrak{F}[d].id_u$ -th pre-computed response of file f_j , $i \neq j$, and the digest d_i of f_i is equal to the digest d_j of f_j . However, this happens with a negligible probability as shown in Equation 3, by substituting M with m ; moreover, $m \ll M$.

4.3 Distribution of File Sizes and s-POW2

Further improvements may be achieved if another, less expensive candidate for the indexing function of the file can be found.

Here we consider using the size $f.size$ of a file f as a candidate for the indexing function. This approach clearly meets the last two requirements outlined in Section 4.2, because it optimizes both I/O and computation.

ALGORITHM 5: Changes in the protocol of s-POW to obtain s-POW2.

Input: A file f .
 \mathcal{C} : **upon** upload of file f **do**
 | $d \leftarrow f.size$;
 | send to SRV a store file request with d ;
end

Algorithm 5 shows the changes to the client-side introduced in this version of the protocol. As we can see, no computation – beside determining the size of the file – is required of the client.

We have explained in Section 4.2 how to cope with collisions in the indexing. However, we still need to verify whether the file size constitutes a good indexing function, i.e. whether in practice, the likelihood that two different files have the same size is tolerably small.

To this end, we have studied the distribution of file sizes over two independent datasets, the Evans and Kuenning dataset [8], containing information on approximately 3 million files and the Agrawal *et al.* dataset [1], for which we only focused on a subset of approximately 200 million files. The two datasets capture (among other information) the sizes of the files observed in the computers of an academic campus and of a large corporation, respectively. Both datasets contain snapshots of the entire content of filesystems in an academic and an industrial environment, respectively. Both datasets are extremely relevant for our scheme, because they would correspond to users backing up (possibly in a storage cloud) their entire hard-disk drives.

The objective of our analysis is to verify the intuition that – especially for large files, i.e. those for which computing another indexing function is more expensive – the size of a file can become a very effective file indexing function. To this end we have extracted from both datasets a unique file identifier (e.g. the hash of the filename) and the file sizes. After purging doubles, we have counted the number of files with equal size.

Figure 2 shows the results of the analysis. For each dataset we show two curves: the first one plots the distribution (with

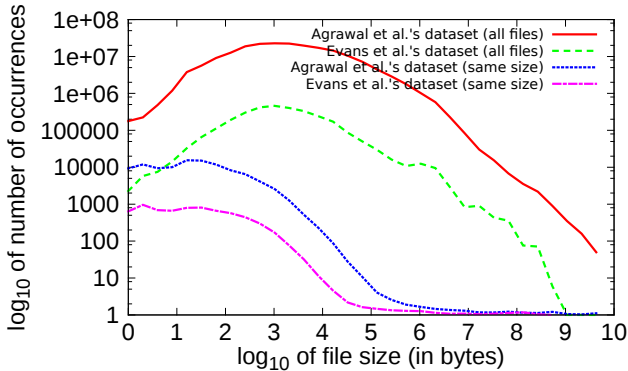


Figure 2: Plot of the distribution of file sizes for files in the entire dataset (with power-of-two bins) and of the average number of files per bin (with power-of-two bins) with the same size for both datasets.

power-of-two bins) of file sizes for files in the entire dataset; the second plots the average number of files with the same size per bin. Notice that we have tried to only count *different* files with the same size: when available, the hash of the file has been used to establish whether two files were the same; otherwise, we have used a combination of file name and creation timestamp.

First of all, we can observe that for both datasets, the distribution of file sizes has a similar shape, albeit that of the first dataset is shifted higher by two orders of magnitude because of the larger file population. Both distributions of files with the same size have similar shape: we can see that in both cases there is a relatively constant number of files with the same size up to around 10 KiB, after which both curves plunge. In both scenarios, the size virtually starts to behave as a unique identifier for files larger than approximately 1 MiB, even though the number of files in the considered bin is still relatively high: for example, in the 1 MiB to 2 MiB range, we have 1,784,136 and 10,819 files in the first and second dataset respectively and the average number of files with the same size is 1.9 and 1.2 respectively. Clearly, the number of files with the same size decreases also because each bin is wider and less and less populated: however these distributions portray very likely the ones a file-storage service may receive as input from its customers—and it is therefore of high significance for this paper. Far from claiming to be exhaustive, our study nonetheless strongly supports the use of file-size for indexing purposes in our scenario.

4.3.1 Security Analysis of s-POW2

Similar considerations to those made in Section 4.2.1 apply to s-POW2: the influence of the number of files with the same size on the choice of the system’s parameters has been captured in Equation 4. Intuitively, the approach suggested in s-POW2 is particularly effective for files with large size, because: i) as shown, the probability of collision is low; and, ii) avoiding the computation of another indexing function on a very large file is particularly cost-effective.

In Section 6.1 we leverage this consideration and those made in the previous sections to obtain a particularly effective scheme.

5. RUNNING POW

To evaluate the effectiveness of our scheme, we have implemented both b-POW and s-POW and its two variants, s-POW1 and s-POW2.

The code has been developed in C++ using the OpenSSL crypto library for all cryptographic operations and using the Saito-Matsumoto implementation [15] of the Mersenne twister [13] for the pseudo random number generator. The code implements both the client-side and the server-side of all schemes. The interactions between client and server as well as the data exchange have been virtualised so as to not consider networking-related delays and to focus only on local (client and server) I/O and computation.

5.1 Experimental Settings

We have run our implementation of both schemes on a 64-bit RedHat box with an Intel Xeon 2.27GHz CPU, 18 GiB of RAM and an IBM 42D0747 7200 RPM SATA hard disk drive. All schemes operate on input files that have been generated at random; the input file size ranges from 1 MiB to 4 GiB, doubling the size at each step. The files are reasonably well defragmented, with a maximum of 34 different extents on the 4 GiB file.

The parameters for b-POW have been chosen in strict adherence with the choices made in [9]. We have also used the same security parameter $k = 66$. Our scheme has two parameters, $\epsilon = (1 - p)$ and g . The values of these parameters are needed to derive a value for K in Equation 2. We have chosen p , the upperbound on the fraction of the file known to the adversary, as $p \in \{0.5, 0.75, 0.9, 0.95\}$. As for g , this parameter measures the probability that the adversary successfully guesses the value of a bit without knowing it. To assign a reasonable value to g , we have analysed what the probability of guessing a bit in an ASCII file with lowercase letters written in the English language is, arguably a relatively conservative case with low entropy in the input distribution. Given the letter frequency analysis in [12], the probability that a given bit equals one is 0.52731. In addition, Equation 2 shows that slight changes in the value of g do not sensibly affect the value K . We have therefore chosen $g = 0.5$.

Each configuration has been run for at least 200 times; before each repetition, cached data, dentries and inodes have been flushed (at both the client- and at the server-side) to ensure accurate measurements. To perform the comparison of the different schemes, the code has been instrumented by surrounding relevant code blocks and functions with calls to extract the Intel Time Stamp Counter through the RDTSC assembly instruction. The figures below have been generated by reporting the mean value and the standard deviation (using a box plot) of the extracted clock cycle count.

5.2 Client-Side

Here we compare the client-side performances of b-POW with those of s-POW, s-POW1 and s-POW2. The implementation of b-POW first loads the file into main memory, where the various phases of the scheme are performed: the reduction phase (which also results in the computation of the SHA256 hash digest of the file), the mixing phase and the calculation of a binary Merkle tree on the resulting reduction buffer.

On the client-side of s-POW, the input file is loaded into memory, the hash digest is computed and then Algorithm 2




Figure 3: Comparison of the running time of b-POW with that of s-POW for different values of K as the input file size grows.

is executed. Figure 3 shows the results of the experiments assessing the performances of the two competing solutions: s-POW is faster than b-POW – from ten times to twice as fast. The complexity of both schemes grows at an approximately equal rate as the input file size grows. The reason for this is that – as previously mentioned – reading the file and computing the hash are by far the predominant operations for both schemes. The discontinuity in the curve of b-POW – noticeable around 64 MiB – is due to the fact that 64 MiB is the maximum size for the reduction buffer. Therefore the computational cost of the reduction phase reaches its maximum at 64 MiB and remains constant afterwards.

For s-POW1, the computation of the hash is replaced by an initial invocation of Algorithm 2. Note that, as access to the entire content of the file is no longer needed, the file is no longer loaded into memory. Indeed, only random disk accesses are needed to fetch the required bits. Figure 4 shows that this second version improves the scheme’s performance with respect to that of b-POW. We can see how the computational cost of our scheme reaches a plateau for sufficiently large files, because – regardless of the input file size – the computation required is essentially constant. The growth rates of the two schemes are now markedly different: b-POW grows linearly with the input file size whereas s-POW1 is asymptotically constant. In addition, the influence of the parameter K starts to become visible.

5.3 Server-Side

At the server-side, we have identified two main phases: the initialization phase and the regular execution phase. The initialization phase corresponds to the first upload of the file; in both schemes, this phase starts with the computation of a hash digest of the file. Then comes the reduction and mixing phases for b-POW, or Algorithm 1 (with n set to 10000) for our scheme. The implementation of Algorithm 1 has been optimized by pre-computing all bit position indexes at once (for all n pre-computed challenges) and by sorting them before performing the file access operations to fetch the corresponding bits. This optimization allows us to only scan the file at most once, thus avoiding the performance penalty associated with random, non-sequential file accesses.

The regular execution phase includes the operations that have to be executed by the server upon each interaction with the client. In b-POW, this phase requires verification of the

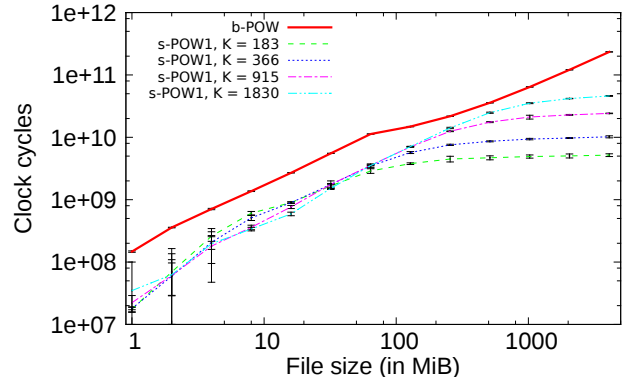


Figure 4: Comparison of the running time of b-POW with that of s-POW1 for different values of K as the input file size grows.

correctness of the sibling path in the computed Merkle tree for a super-logarithmic number of leaves (we have picked this number to be 20 as in [9]). In contrast, in our scheme, if we factor out the table lookup required to retrieve (from the received file index) the correct data structure holding the state for the given file, our protocol only needs to verify the equality of two short bit strings. The related overhead is therefore negligible. However, our scheme also requires regular re-executions of Algorithm 1 to pre-compute new challenges: we will therefore include this in the regular execution phase.

Figure 5 shows the performance of the initialization phase: the cost of b-POW grows – as explained in the previous section – with the same rate as the cost of reading the entire file. Our scheme exhibits an essentially constant computational cost up to a certain point, and a cost similar to that of b-POW (linear with the cost of reading the entire file) from that point on. The reason for this is that the overhead of generating the $n \cdot K$ challenges, sorting them and maintaining the data structure with all bit vectors is constant: for small files, this overhead is higher than the cost of reading the entire file and thus prevails. After the input file has reached a critical size however, the cost of reading the file becomes dominant. The reason for this asymptotic behaviour is that with high probability, reading $n \cdot K$ bit positions in the file requires fetching most data blocks of the file, which is roughly equivalent to reading the entire file.

Figure 6 compares the performance of 10000 repetitions of the regular execution phase for both schemes. b-POW exhibits an essentially constant computational cost as the number of leaves of the Merkle tree is relatively low and does not grow past 64 MiB. The computational costs for this phase of our scheme are the same as those shown in Figure 5 minus the hash computation which is no longer required. We would like to emphasize however that Figure 6 shows a comparison between the *on-line* computation required by the verification phase of b-POW and the *off-line* computation required to generate the challenges: the former requires readily available computation power regardless of the load of the system (since delaying client requests is not acceptable); the latter is a computation that can be carried out when the system load is low.




Figure 5: Comparison of the running time of the server initialization phase of b-POW with that of s-POW for different values of K as the input file size grows.

6. COMPARISON AND DISCUSSION

In light of the analysis performed in the previous section we are now ready to compare the state of the art solution and our proposals. Table 1 compares b-POW, s-POW and s-POW1 in terms of computational cost, I/O, storage and bandwidth requirements; we omitted s-POW2 from the comparison as it has the same asymptotic costs as s-POW1.

On the client-side s-POW and s-POW1 are far less demanding than b-POW from both the computational and the I/O perspective. This is a highly desirable characteristic, as the end user will receive a better service. Such lower demands on the client-side are compensated by an increase of the footprint on the server-side. However the design of the scheme allows the server to distribute computation and I/O over time, and to carry them out in moments of low system load.

From a computational perspective on the client-side, both b-POW and s-POW are dominated by the cost of calculating the hash of the file whereas s-POW1 has a constant cost that only depends on the security parameter K , and is independent of the input file size. Similar considerations can be made when investigating the client-side I/O requirements.

On the server-side we have made separate considerations for the initialization and for the regular execution phase. In the initialization phase, b-POW and s-POW are once more dominated by the computational and I/O cost of the hash calculation, whereas s-POW1 only requires the pre-computation of n challenges. The regular execution phase is particularly cheap for b-POW as no I/O and only constant computation are required. s-POW and s-POW1 require regularly replenishments of the stock of precomputed challenges. However, this operation can be performed offline in moments of low system load. Furthermore, files are often regularly read at the server-side as part of standard management tasks anyway (e.g. periodic integrity checks, backup, replication); in this case, the I/O cost of the response pre-computation phase, which is by far the predominant cost, can be factored out.

As for server-side storage, b-POW requires only the root of the Merkle tree to be stored, whereas s-POW and s-POW1 require storing the pre-computed challenges. We emphasize however that the number of responses to be pre-computed is




Figure 6: Comparison of the running time of the server regular execution phase of b-POW with that of s-POW for different values of K as the input file size grows.

a tunable parameter. Furthermore, the size of the responses is independent of the input file size – usually only a negligible fraction of the latter. For instance, storing 1000 responses, each 1830 bits long, would require less than 230 KiB.

Finally, the b-POW scheme requires the exchange of a super-logarithmic number of sibling paths of the Merkle tree between client and server, whereas s-POW and s-POW1 only require the exchange of a k -bit string.

ALGORITHM 6: Changes to s-POW to obtain s-POW3.

```

C : upon upload of file  $f$  do
  if  $h \leftarrow H(\text{file})$  available for  $f$  then
    |  $d \leftarrow h$ ;
  else
    if  $f.\text{size} < t_1$  then
      |  $d \leftarrow H(\text{file})$ ;
    else if  $t_1 \leq f.\text{size} < t_2$  then
      | invoke Algorithm 2 on input  $\text{file}$  and  $s_{\text{pub}}$  to get  $d$ ;
    else
      |  $d \leftarrow f.\text{size}$ 
    end
  end
  send to  $SRV$  a store file request with  $d$ , including the
  indexing mechanism used;
end

```

6.1 Putting it All Together: s-POW3

Based on the analysis of each version of our scheme, we now suggest a further variant combining the different optimizations presented to obtain the best performance. The rationale behind this last optimization is based on the following considerations: i) the cost of computing a standard cryptographic hash (e.g. SHA-1) for small files is negligible; ii) based on the file size distribution in common datasets, and on the distribution of files with the same size, the file size is a good indexing function for file sizes of 1 MiB and more; iii) often the hash digest of a file is already available at the client-side; for instance, several peer-to-peer file sharing clients compute and/or store the hash of downloaded files.

Algorithm 6 shows the changes required in Algorithm 3 to obtain s-POW3. In particular, the server always accepts a digest from the client if available, and uses it as the lookup key (this requires the server to compute the digest of new

	b-POW	s-POW	s-POW1
Client-side computation	$O(m)$ hash	$O(m)$ hash	$O(k)$ PRNG
Client-side I/O	$O(m)$	$O(m)$	$O(k)$
Server-side computation (initialization phase)	$O(m)$ hash	$O(m)$ hash	$O(nk)$ PRNG ^a
Server-side computation (regular execution phase)	$O(1)$	$O(nk)$ PRNG ^a	$O(nk)$ PRNG ^a
Server-side I/O (initialization phase)	$O(m)$	$O(m)$	$O(nk)$ ^b
Server-side I/O (regular execution phase)	0	$O(nk)$ ^b	$O(nk)$ ^b
Server-side storage	$O(1)$	$O(nk)$ ^a	$O(nk)$ ^a
Bandwidth	$O(k \log k)$	$O(k)$	$O(k)$

^a For the precomputation of n challenges.

^b For the generation of n challenges; if n is sufficiently large, it is bounded by $O(m)$.

Table 1: Performance analysis of the schemes; m represents the input file size, k is the security parameter. As explained in Section 4.1.1, the more precise formulation for $O(k)$ in our scheme is shown in Equation 2.

files, but this is very likely performed for integrity protection anyway). If not, two thresholds are set: for files smaller than t_1 , the client has to compute the hash digest; for files in the t_1 to t_2 range, the approach of s-POW1 is used; whereas for very large files (size greater than t_2), the file size (and thus the approach of s-POW2) is used.

7. CONCLUSIONS

We have presented a suite of novel security protocols to implement proof of ownership in a deduplication scenario. Our scheme is provably secure and achieves better performance than the state-of-the-art solution in the most sensitive areas of client-side I/O and computation. On the server-side, I/O and computation can be conveniently deferred to moments of low system load. Note that the proposed solutions are fully customizable in the system parameters. Finally, extensive simulation results support our findings.

Acknowledgments

The authors would like to thank Geoff Kuenning and John Douceur for their help in accessing statistics on file sizes; Alexandra Shulman-Peleg and Christian Cachin for their insightful feedback; Charlotte Bolliger for her precious corrections.

Roberto Di Pietro has been partially supported by a Chair of Excellence granted by University Carlos III, Madrid.

8. REFERENCES

- [1] N. Agrawal, W. J. Bolosky, J. R. Douceur, and J. R. Lorch. A five-year study of file-system metadata. In *FAST*, pages 31–45, 2007.
- [2] G. Ateniese, R. C. Burns, R. Curtmola, J. Herring, O. Khan, L. Kissner, Z. N. J. Peterson, and D. Song. Remote data checking using provable data possession. *ACM Trans. Inf. Syst. Secur.*, 14(1), 2011.
- [3] G. Ateniese, R. C. Burns, R. Curtmola, J. Herring, L. Kissner, Z. N. J. Peterson, and D. X. Song. Provable data possession at untrusted stores. In *ACM Conference on Computer and Communications Security*.
- [4] G. Ateniese, R. Di Pietro, L. V. Mancini, and G. Tsudik. Scalable and efficient provable data possession. In *Proceedings of the 4th international conference on Security and privacy in communication networks*, SecureComm '08, 2008.
- [5] J. R. Douceur, A. Adya, W. J. Bolosky, D. Simon, and M. Theimer. Reclaiming space from duplicate files in a serverless distributed file system. In *ICDCS*, pages 617–624, 2002.
- [6] driverdan. dropship - dropbox api utilities. <https://github.com/driverdan/dropship>, 2011.
- [7] C. Erway, A. Küpçü, C. Papamanthou, and R. Tamassia. Dynamic provable data possession. In *Proceedings of the 16th ACM conference on Computer and communications security*, CCS '09, 2009.
- [8] K. M. Evans and G. H. Kuenning. A study of irregularities in file-size distributions. In *International Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS '02)*, 2002.
- [9] S. Halevi, D. Harnik, B. Pinkas, and A. Shulman-Peleg. Proofs of ownership in remote storage systems. In *ACM Conference on Computer and Communications Security*, pages 491–500, 2011.
- [10] D. Harnik, B. Pinkas, and A. Shulman-Peleg. Side channels in cloud services: Deduplication in cloud storage. *IEEE Security & Privacy*, 8(6), 2010.
- [11] A. Juels and B. S. K. Jr. Pors: proofs of retrievability for large files. In *ACM Conference on Computer and Communications Security*, pages 584–597, 2007.
- [12] R. Lewand. *Cryptological mathematics*. Classroom resource materials. Mathematical Association of America, 2000.
- [13] M. Matsumoto and T. Nishimura. Mersenne twister: A 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Trans. Model. Comput. Simul.*, 1998.
- [14] R. C. Merkle. A digital signature based on a conventional encryption function. In *CRYPTO*, pages 369–378, 1987.
- [15] M. Saito and M. Matsumoto. Simd-oriented fast mersenne twister: A 128-bit pseudorandom number generator. In *In Monte Carlo and Quasi-Monte Carlo Methods 2006*. Springer-Verlag, 2007.
- [16] M. W. Storer, K. M. Greenan, D. D. E. Long, and E. L. Miller. Secure data deduplication. In *StorageSS*, pages 1–10, 2008.